



AN INTEL COMPANY

Building the Carrier Grade NFV Infrastructure

Wind River Titanium Cloud

WHEN IT MATTERS, IT RUNS ON WIND RIVER

EXECUTIVE SUMMARY

Network Functions Virtualization (NFV) is the telecommunications company’s version of IT virtualization, achieved by augmenting the latter with the carrier grade capabilities required for high availability, security, and performance, as well as for more efficient network management. In NFV, software-based virtualized network functions (VNFs) run on one or more virtual machines (VMs) and are chained together to create communications services. NFV solutions consist of three layers: a VNF layer running on an NFV infrastructure (NFVI) layer, and an NFV management and orchestration (NFV-MANO) layer which manages the VNF and NFVI layers. The NFV server, the basic building block for NFV carrier grade solutions, consists of NFV software running on industry-standard hardware.

This paper discusses the benefits and challenges of NFV, and examines how Wind River® Titanium Cloud, the industry’s first feature-complete NFV server, achieves an extremely high level of reliability, providing a solid foundation for building carrier grade NNF network equipment, networks, and services.

TABLE OF CONTENTS

Executive Summary 2
NFV Benefits 3
Service Outage Challenge 3
NFV Carrier Grade Server 3
Wind River Titanium Cloud 3
Reliability/Availability/Maintainability (RAM) Modeling 3
Sample Titanium Cloud Configuration 4
Titanium Cloud Results 4
Conclusion 5



NFV BENEFITS

NFV promises to revolutionize carrier networks by providing communications service providers (CSPs) with unprecedented flexibility in how they deploy and manage network equipment and services. Virtualization decouples network functions from hardware, allowing VNFs to be located anywhere in the network where they can be most cost-efficient.

The anticipated cost benefits are expected to be immediate: VNFs can be physically located in one or more data centers, reducing operational costs. Further cost reductions are achieved by consolidating multiple network functions and taking advantage of NFV's efficient equipment utilization, allowing reductions in the amount of equipment, spares inventory, equipment real estate, cabling, and power consumption. And the NFV server's building block approach results in shorter development cycles and improved operational efficiencies because of the commonality of tools and increased network management automation.

SERVICE OUTAGE CHALLENGE

The inherent benefits of NFV also enable flexibility in the level of service availability achieved. Protection groups can be set up in a wide range of sizes, service mixes, and configurations, and can be protected using an equally wide range of redundancy strategies. Virtualization enables efficient use of redundant resources, minimizing the cost of redundancy. But the anticipated cost savings of NFV can be eroded by inadequate fault management design, resulting in increased maintenance and outage costs due to poor fault coverage and diagnostics.

Outage costs include both CSP operational costs associated with diagnosing and repairing outage incidents, and the costs incurred by CSP customers as a result of service interruptions. A recent study of IT data center outages by Ponemon Institute Research reported an average outage cost of \$690,240 per incident, or \$6,828 per minute ("2013 Cost of Data Center Outages," Ponemon Institute, sponsored by Emerson Network Power, December 2013). These are total costs—that is, due to IT operations, reduced productivity, lost revenue, and business disruption. But these IT costs are relatively small when compared to public carrier network outage costs. For example, a France Telecom outage in July, 2011, left 28 million customers without phone and text messaging service for over 12 hours, costing France Telecom between \$12M and \$25M in repair costs and customer refunds (Leila Abboud,

"Analysis: France seeks influence on Telcos after outage," Reuter online U.S. edition, July 2012). Further, service outage costs vary depending on the type of service. A 2000 survey by Contingency Planning Research listed service outage costs as high as \$6.5M per hour for brokerage operations and \$2.6M per hour for credit card authorization ("Outage Cost Survey," Contingency Planning Research, 2000).

NFV CARRIER GRADE SERVER

The NFV server is the critical component for achieving carrier grade NFV solutions, since it is used to build both the NFVI and NFV-MANO layers. Its carrier grade features are needed to achieve high-availability and low-maintenance NFV solutions. Its fault tolerance and fault management features must ensure that failures are properly managed across the complete lifecycle of a failure event. Failures must be quickly detected and contained, recovered to adequately provisioned alternate resources, alarmed, repaired, tested, and returned to service. Key aspects that characterize carrier grade implementations include high fault detection coverage, fast recovery times, accurate fault isolation, and low hardware and software failure rates.

WIND RIVER TITANIUM CLOUD

Wind River Titanium Cloud is the industry's first fully integrated, open source-based, feature-complete NFV server. It has a full complement of carrier grade fault management features and employs advanced robustness techniques to form a solid foundation for building carrier grade NFV network equipment, networks, and services. The Titanium Cloud software stack runs on most standard multi-core platforms to configure compute servers for the NFVI layer and control servers for the NFV-MANO layer.

RELIABILITY/AVAILABILITY/MAINTAINABILITY (RAM) MODELING

To determine whether Titanium Cloud's fault-tolerant design can meet the "six-nines" (99.9999%) objective, custom Markov Models were built (Rob Paterson, "Wind River Titanium Cloud Reliability/Availability/Maintainability (RAM) Modeling Analysis," KerrNet Consulting Inc., January 2015). These models capture Titanium Cloud's response to hardware and software failures, repair actions, and software upgrades to define system states and inter-state transitions that capture the fault-tolerant behavior of Titanium Cloud. The resultant RAM model calculates Titanium Cloud RAM

metrics and can be used to evaluate different design options by varying design and operational parameters such as fault detection coverage, software failure rate, and repair time.

The RAM model uses steady-state useful-life hardware and software failure rates, and assumes an annual in-service upgrade to calculate the following metrics:

- **Compute server downtime:** The average minutes per year that a compute server is unavailable to support a set of VNFs
- **Control server downtime:** The average minutes per year that an NFVI solution is without management and orchestration
- **Unplanned maintenance actions ratio (UMAR):** The ratio of unplanned maintenance action rates for the product that employs fault tolerance to one that employs no fault tolerance

Since adding redundancy increases maintenance costs, this last metric is useful for comparing the amount of relative effort required to minimize downtime for different fault-tolerant design options.

SAMPLE TITANIUM CLOUD CONFIGURATION

Titanium Cloud consists of compute servers, which run VNFs, and control servers, which manage the NFVI and the VNFs. To calculate the solution’s RAM metrics, the server configuration in Figure 1 was modeled (Rob Paterson, “Wind River Titanium Cloud Reliability/Availability/Maintainability (RAM) Modeling Analysis,” KerrNet Consulting Inc., January 2015). The dashed line defines the demarcation scope of the analysis. The model excludes the redundant Gigabit (GE) switches and the CSP’s VNF software.

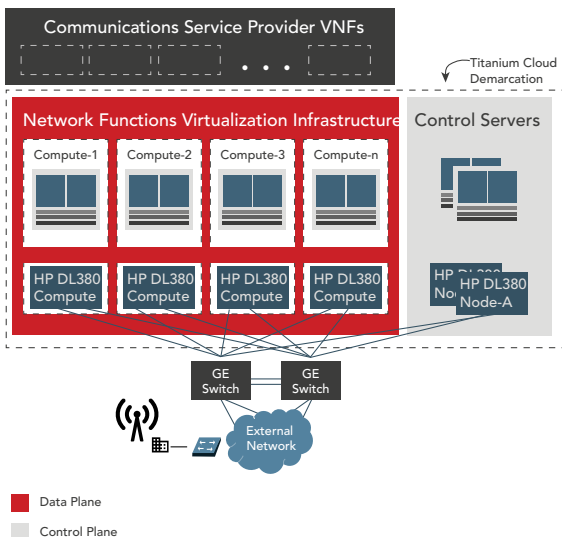


Figure 1. Wind River Titanium Cloud configuration

The sample configuration consists of eleven Titanium Cloud compute servers (ten active and one active-standby), which are responsible for processing all communications data between the network and the CSP’s VNFs. A compute server failure results in the automatic failover to the active-standby. This process is managed by the Titanium Cloud control server, which consists of dual servers (one active and one active-standby). Communication between control and compute servers is via dual Gigabit GE switches using fully redundant inter-server and inter-switch links. The configuration is a distributed virtualized one, such that there are no failure modes within the demarcation zone that would result in all VNFs served by the compute servers being out of service simultaneously (i.e., no “total system downtime” failure mode).

The 10:1 compute server configuration reserves one active-standby compute server that is processing software. All compute servers implement a software agent that is tested by the control server using a “heartbeat” function with a frequency in the hundreds of milliseconds. The control server that employs one active-standby hardware platform is also tested by the active control server at the same intervals.

Failure of compute server hardware (HP DL380 is used in the model) results in all affected VNFs being automatically reassigned to the hot standby compute server in less than 500 milliseconds with no impact on VNF communications. Titanium Cloud software failures are automatically recovered via sub-five second software process restarts.

Titanium Cloud implements a “hold-off” mechanism to handle transient failures that prevents recovery instability and therefore unnecessary outages triggered by certain external network failure modes such as slow GE switch failovers.

TITANIUM CLOUD RESULTS

The predicted results for an attended 10:1 Titanium Cloud configuration are summarized in Table 1. It compares Titanium Cloud with an IT-grade equivalent solution. It assumes that the IT-grade solution is maintained using carrier maintenance procedures; it thus compares the design capabilities. Similarly to Titanium Cloud, the IT-grade solution supports compute server hardware redundancy, but lacks the fault detection coverage and recovery speed of Titanium Cloud. Therefore the IT-grade solution is not capable of hitless software upgrades.

Table 1. Titanium Cloud Results

Metric	Wind River Titanium Cloud	IT-Grade Equivalent
Total compute system downtime (min/yr)	0	0
Individual compute server downtime (min/yr)	0.19	6.4
Control system downtime (min/yr)	0.23	7.4
Unplanned maintenance actions ratio	1.2	1.1

Because Titanium Cloud employs a distributed architecture, there is no common failure mode that would result in a total system outage. The predominant failure mode is single compute server outages, which would affect a group of VNFs. The predicted compute server system downtime is 0.19 minutes per year; in contrast, the IT-grade solution’s compute server downtime is 6.4 minutes per year. For carrier deployments this could result in tens of millions of dollars difference in outage costs.

The loss of management control (control system downtime) is 0.23 minutes per year for Titanium Cloud—significantly lower than the IT-grade solution, in which the predicted downtime for loss of control is 7.4 minutes per year. This is time spent in a vulnerable state, as the VNFs and the compute servers are operating without fault management, communication of alarms, and maintenance.

The system’s predicted annual UMAR is 1.2, which is relatively low because of the high ratio of active to inactive compute servers. This efficient use of redundancy resources translates to lower maintenance costs. The IT-grade solution has a marginally lower maintenance rate because of an unduplicated control server, but the outage cost avoidance as a result of lower downtime far exceeds the marginal cost increase in maintenance actions.

Titanium Cloud’s redundancy ratio can be increased beyond 10:1 configured. Using the RAM model, the redundancy ratio was varied up to 100:1 to determine the impact on compute server downtime. The results (see Figure 2) indicate that the redundancy ratio can be increased to 100:1 for attended locations and 30:1 for unattended offices without adverse effects on downtime. This results in a significant reduction in UMAR—1.05 and 1.08, respectively.

Impact of Redundancy Ratio on Compute Server Downtime (Spare Used for Upgrades)

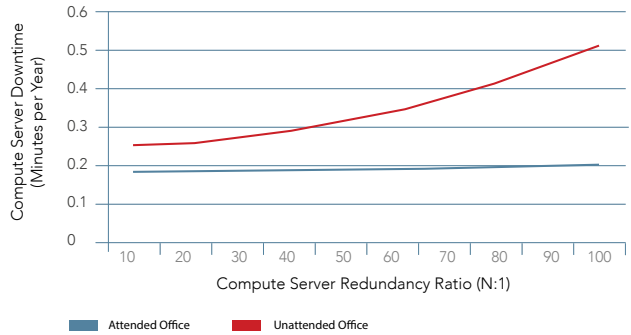


Figure 2. Impact of redundancy ratio on downtime

CONCLUSION

NFV marks a significant change in how services and equipment are deployed and managed in the public carrier network. NFV enables CSPs to consolidate many types of network functions into data centers, resulting in significant reductions in capital and operational expenditures, as well as in service introduction times. Wind River Titanium Cloud is the industry’s first fully integrated, open source-based, feature-complete NFV server. It supports a wide range of redundancy types and fault management features to achieve better than six-nines availability at minimal redundancy costs.

