

CLOUDBAND WITH OPENSTACK AS NFV PLATFORM

STRATEGIC WHITE PAPER | NFV INSIGHTS SERIES

OpenStack® is an open-source technology with many capabilities that are needed in any Network Functions Virtualization (NFV) environment, and many telco service providers are interested in using OpenStack in their NFV infrastructures. However, to realize the benefits of NFV, service providers need NFV platforms that provide additional capabilities to support distributed clouds, enhanced network control, lifecycle management and high performance data planes. This paper examines these gaps and explains how Red Hat® and the Alcatel-Lucent CloudBand™ team work together to build a solution that is optimized for telco NFV environments.

About the NFV Insights Series

NFV represents a major shift in the telecommunications and networking industry. NFV applies virtualization and cloud principles to the telecommunications domain, something that appeared to be impossible until recently due to the stringent performance, availability, reliability, and security requirements in communication networks. Many service providers are now keen to implement NFV to help them gain an advantage through automation and responsiveness to deliver an enhanced customer experience and reduce operational costs. This series of whitepapers addresses some of the key technical and business challenges on the road to NFV.

TABLE OF CONTENTS

OpenStack, NFV and CloudBand / 1

 The CloudBand and Red Hat approach / 2

Distribution / 4

Networking / 5

Automated lifecycle management / 6

NFV infrastructure operations / 7

High performance data plane / 8

Conclusion / 8

Acronyms / 9

OPENSTACK, NFV AND CLOUDBAND

In 2010, RackSpace® and NASA jointly launched OpenStack®, an open-source cloud computing platform. Since then the OpenStack community has gained tremendous momentum with over 200 member companies.

Originally, OpenStack was not designed with carrier requirements in mind. Some early technical guidelines and constraints for building a carrier-grade cloud were given in a white paper by SCOPE Alliance, published in May of 2011 [1]. SCOPE Alliance is an industry exercise driven by network equipment providers to encourage and promote carrier-grade platforms based on openness while maintaining technical carrier-grade requirements.

In 2012, a group of major telecommunication service providers came together and founded an initiative to apply virtualization and cloud principles to the telecommunication domain. The term Network Functions Virtualization (NFV) was coined as a label for this initiative. Service providers called for vendors to build virtualized network functions and NFV platforms to help them become more agile in delivering services, and to reduce equipment and operational cost.

To realize the promise of NFV, service providers need a horizontal NFV platform that provides:

- Compute, storage and networking resources for the virtual network functions (vNF)
- A centralized platform for NFV Orchestration (NFVO)
- A centralized Virtual Network Function Manager (VNFM) that manages the lifecycles of the vNFs

Adopting this approach will make it possible to break open the many costly application silos that exist today.

The NFV community does not often speak of the concept of cloud although everybody knows that effectively an NFV platform is a cloud, albeit a specific cloud, optimized for virtual network functions. Alcatel-Lucent is working with many large service providers around the globe. Most of them are looking for an open and multi-vendor NFV platform and they are convinced that their NFV platform should be based on OpenStack. OpenStack is an open, non-proprietary technology with many cloud core capabilities that are needed in any NFV environment. As OpenStack is being driven by a highly active open source community (openstack.org) it would not make sense to replace it with a proprietary solution dedicated to NFV.

However, OpenStack and its supporting open source community are not focused on NFV, hence OpenStack is lacking in some capabilities that are critical requirements for an NFV platform. The NFV framework defined by the ETSI NFV Industry Specification Group can be referenced in order to better understand these critical NFV requirements. A group of communications service providers initiated this group to define the requirements for NFV as well as a set of frameworks and architectures to address these requirements.

Some of the key requirements for communications service providers are:

- The ability to deal efficiently with a distributed NFV infrastructure, which is necessary for assuring low signal latencies and disaster resiliency; this infrastructure must, however, also be manageable as a single distributed cloud with global views, statistics and policies
- The ability to set up and manage the local and WAN structures required for carrier applications in a programmable manner
- The ability to automate the lifecycle of complex carrier applications with many component functions that have different deployment, scalability and healing requirements
- The ability to manage high performance data plane network functions on commercial off-the-shelf servers


It is likely that OpenStack’s support of at least some of these requirements will improve over time. In some areas, however, it is unclear if that will ever happen. The OpenStack community is driven by IT companies, such as Red Hat®, Mirantis, RackSpace, SUSE®, HP™, IBM®. There is little representation from the telco industry on the OpenStack board of directors. The only service provider member is AT&T®. Requirements that are specific to NFV have not been high on the agenda. However, recently, an NFV subgroup was formed to focus on these requirements.

To address identified gaps, major industry players are planning to establish “Open Platform for NFV” as a Linux™ Foundation Collaborative Project. The intention is to create a carrier-grade open source reference platform for NFV, based on OpenStack. Industry peers will build this platform together to advance the evolution of NFV and to ensure consistency, performance and interoperability among multiple open source components.

The CloudBand and Red Hat approach

The Alcatel-Lucent CloudBand and Red Hat teams promote the inclusion of more NFV requirements in the OpenStack upstream. In addition, CloudBand adds to OpenStack in the critical under-addressed areas identified in Figure 1.

Figure 1. OpenStack and CloudBand focus

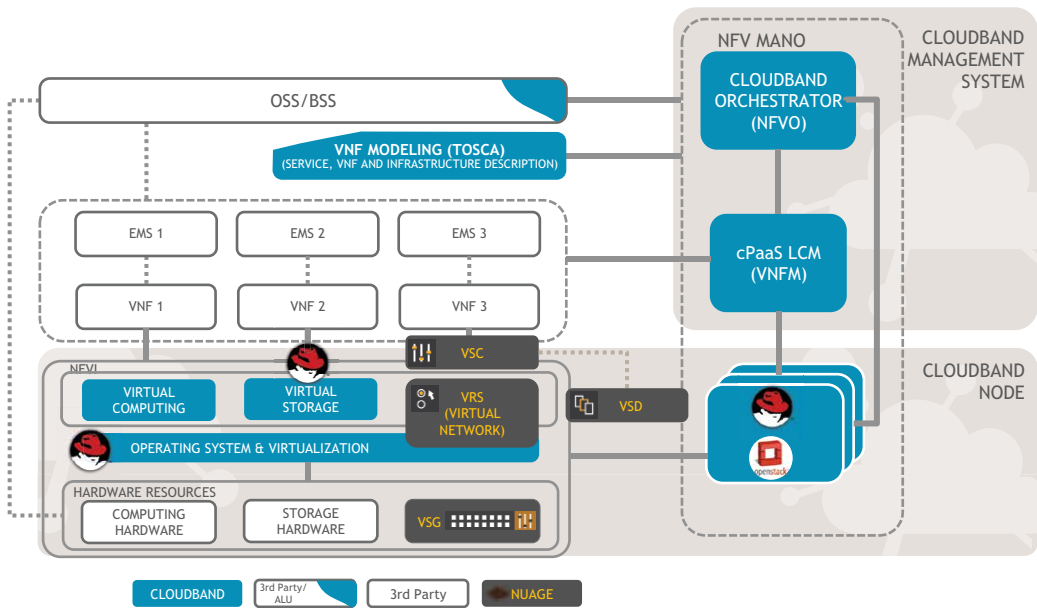


CloudBand comprises two parts: the CloudBand Management System and the CloudBand Node (Figure 2).

The **CloudBand Management System** incorporates the CloudBand NFVO and the CloudBand cPaaS (Carrier PaaS) as the VNFM. The CloudBand Management System interfaces with the CloudBand Node using standard OpenStack APIs as well as additional function APIs (Software-Defined Network (SDN) controller API for example).

The **CloudBand Node** uses Red Hat Enterprise Linux OpenStack Platform as the virtual infrastructure manager (VIM) without modifying it. Due to the momentum of OpenStack, incorporating the ongoing OpenStack software enhancements in a modified NFV OpenStack version is not tenable and would lead to a constant lag between OpenStack and this modified version. Instead, the CloudBand Node adds complementary functionality to OpenStack using the open APIs and connectors of OpenStack. In addition, the CloudBand Node software stack includes Inktank™ Ceph (now part of Red Hat) as the virtual storage solution, and KVM and Libvirt on top of Red Hat Enterprise Linux at the virtualization layer, both as part of the NFV infrastructure (NFVI). The importance of Red Hat’s “upstream first” open source product development strategy cannot be underestimated when considering the CloudBand solution for NFV deployments. Also, the CloudBand Node software stack is hardware agnostic as are the Red Hat Enterprise Linux OpenStack Platform and Ceph open source technologies.

Figure 2. CloudBand and Red Hat architecture mapped to ETSI-NFV




The rest of this paper focuses on five areas where CloudBand builds on OpenStack to create a solution that is optimized for telco NFV environments: distribution, networking, automated lifecycle management, NFV infrastructure operations and high performance data plane.

DISTRIBUTION

In the IT world, enterprises are striving to consolidate their datacenters to reduce cost. This is not always the optimal choice for NFV. Many NFV applications require a real-time response with low latency. NFV applications also need to be highly available and survive disasters. Many of the additional requirements are described in the SCOPE Alliance White Paper [1]. Service providers need the flexibility to deploy network functions in a distributed infrastructure — at the network core, metro area, access, and possibly even a customer’s premises.

Figure 3. Distributed NFV infrastructure



OpenStack support for distribution is quite limited, mainly in terms of scalability, but it can still be achieved through a complex configuration of cells, regions and availability zones. Each OpenStack region provides separate API endpoints with no coordination between regions. Typically, one or more regions are located in one datacenter. The Cells component provides a single API endpoint aggregating multiple regions. With Cells, placement (“scheduling”) of workloads across cells is by explicit specification or by random selection. The Cells component does not contain a placement algorithm able to choose the best geographic location based on the needs of the application. The Horizon GUI is restricted to a single region at any one time. There is no GUI that would show an aggregated view of the NFV cloud infrastructure. The OpenStack Glance virtual machine image manager is also limited to a single region. This means that the NFV operator would have to deploy images manually to the regions where it is needed.

CloudBand supports distributed NFV cloud infrastructures in a variety of ways. CloudBand provides aggregated northbound APIs that allow NFV applications and BSS/OSS to deal with the different locations as a single cloud. CloudBand provides a policy-based placement algorithm that computes a quasi-optimal location based on server utilization at the different locations, affinity and anti-affinity rules, and other parameters. In addition, CloudBand supports Placement Zones, which span geo-distributed datacenters by “aggregating” CloudBand Node OpenStack availability zones. Virtual machines belonging to vNFs can then be instantiated across these Placement Zones according to pre-defined business policies. CloudBand provides built-in load balancing services for scalable distributed network functions. The CloudBand graphical user portal gives users an aggregated view of the infrastructure according to their different roles and responsibilities. The image management component of CloudBand manages a single catalog of images automatically and assures that images are made available where they are needed. User accounts and key pairs are also managed at the global level.

NETWORKING

Virtual network functions vary widely in their network demands. Due to their distribution throughout an NFV infrastructure, the baseline requirement for an NFV network is connectivity — not only within the datacenter but also across the wide area. For security reasons, different network functions should only be connected to each other if they need to exchange data, and the NFV control, data and management traffic should be separated. As network functions are decomposed, for example into data plane components and a centralized control plane component, the wide area connectivity between these components needs to remain as highly reliable as with traditional integrated architectures. Enough network resources should be available to ensure that surging traffic from other applications cannot adversely affect the NFV applications. The network should be resilient against equipment failures and force majeure disasters. Latency and jitter requirements vary from hundreds of milliseconds for some control and management systems to single digit milliseconds for mobile gateways and cloud radio access networks.


NFV networks will typically consist of a semi-static physical infrastructure along with a much more dynamic overlay network layer to address the needs of virtual network functions. The overlay layer needs to respond quickly to changing service demand, new service deployments and so on.

OpenStack Neutron is the networking component offering abstractions, such as Layer 2 and Layer 3 networks, subnets, IP addresses, and virtual middleboxes. Neutron has a plugin-based architecture. Networking requests to Neutron are forwarded to the Neutron plugin installed to handle the specifics of the present network. Neutron is limited to a single space of network resources typically associated with an OpenStack region. There is no capability to federate multiple network domains and manage WAN capabilities.

CloudBand delivers NFV networking abstractions that extend from the datacenter across the WAN toward multiple locations (Figure 4). CloudBand Route Domains and Network Templates support flexible aggregations of networks. CloudBand is open to interface with any networking framework using standard OpenStack Neutron APIs and plugins. CloudBand can interface with existing MPLS networks and other legacy networks using a process called VPN stitching.

For example, CloudBand supports Nuage Networks Virtualized Services Platform (VSP), a second generation SDN solution. Nuage Networks VSP supports the federation of multiple SDN network domains each with their own SDN controller. Networks are synchronized in both directions between CloudBand and Nuage Networks VSP. The SDN controller, in this case Nuage Networks VSP, has been recently identified by ETSI NFV Industry Specification Group and is now referenced as the WAN Infrastructure Manager (WIM).


Figure 4. CloudBand network support, leveraging Nuage Networks VSP as the SDN controller (WIM)



AUTOMATED LIFECYCLE MANAGEMENT

One of the greatest advantages of NFV as a software-based solution is the ability to automate operational processes including the application lifecycle from deployment to monitoring, scaling, healing and upgrading, all the way to phase out. Studies have shown that through this automation, service providers can reduce OPEX in some cases by more than 50 percent.

Figure 5. Application lifecycle



OpenStack Heat allows users to write templates to describe virtual applications (“stacks”) in terms of their component resources, such as virtual machines including nested stacks. Currently, Heat is based on AWS™ CloudFormation™. For the future, the OASIS Topology and Orchestration Specification for Cloud Applications (TOSCA), or other description formats are under discussion. The focus of Heat is on the definition and deployment of application stacks. OpenStack Solum is a new project designed to make cloud services easier to consume and integrate into the development process. It will also provide some

of the missing lifecycle automation functions. There is some initial work on auto-scaling by combining the measurement capabilities of OpenStack Ceilometer with Heat. Heat is currently limited to a single OpenStack region.

CloudBand provides automated lifecycle management for NFV applications (vNFs) that is both distribution-aware and network-aware. As part of the application onboarding process, a vNF Descriptor, also known as a recipe, is created. The recipe describes the components of the application, the dependencies between its components, application monitoring parameters and the actions taken when certain lifecycle events occur. Policies provide the user with control over the network functions, such as the ability to establish priorities among applications and resources.

CloudBand Lifecycle Management operates at the top cloud level and can leverage the CloudBand smart placement algorithms. Recipes can specify network templates to be instantiated when an application is deployed. Virtually all lifecycle phases involve the monitoring of the application components (virtual network function components (vNFC)). CloudBand can accomplish this with agents installed inside a component virtual machine or it can work in a remote-agent mode. The remote-agent mode is necessary for vNF components that do not easily allow installation of an agent inside the virtual machine. CloudBand can use aggregate algorithms that detect possible problems and trends in different places to trigger changes in the application and keep it running in an optimal way. The OASIS TOSCA is an advanced domain-specific language for modeling cloud applications. ETSI, vendors and service providers have embraced TOSCA as an emerging standard language to describe a topology of cloud-based web services, their components, relationships, and the procedures that manage them. TOSCA is still evolving and Alcatel-Lucent is contributing significantly to its evolution, for example, in the area of network lifecycle management.

NFV INFRASTRUCTURE OPERATIONS

The distribution of NFV infrastructures across many locations in the service provider network will pose specific challenges and impact the operational processes and support systems. A distributed infrastructure means that cloud nodes at different locations are added, upgraded and/or removed more frequently than in a centralized cloud. These processes should be done remotely whenever possible to avoid truck rolls and sending engineers across the coverage area.

OpenStack TripleO (OpenStack on OpenStack) is an experimental addition to the OpenStack family. The project aims at automating the installation, upgrade and operation of OpenStack clouds using OpenStack's own cloud facilities. TripleO uses Heat to deploy an OpenStack instance on top of a bare-metal infrastructure.

CloudBand Nodes are specifically designed for a distributed NFV infrastructure. A CloudBand Node comprises a pre-built cloud software stack-in-a-box based on OpenStack. The Node contains extensive software for commissioning and monitoring itself. A CloudBand Node can be shipped to the customer location as a fully mounted and cabled physical rack. It can be up and running within as little as four hours compared to days or weeks with other systems. This is achieved by the CloudBand Automatic Infrastructure Deployment (ACID) tool, which automates the complex software stack deployment and upgrade procedures. The Node can be installed and run in a lights-out mode with no need to send engineers on location.

The CloudBand solution contains many other capabilities to help optimize NFV operations. The CloudBand Management System provides support for different roles and responsibilities with different access rights. The system is multi-tenant: different organizations can use the system simultaneously, without seeing each other's resources and activities. CloudBand also implements a growing set of functions to enable easy root cause analysis for system and application failures.

HIGH PERFORMANCE DATA PLANE

Many carrier network functions, such as deep packet inspection, media gateways, session border controllers, mobile core serving gateways and packet data network gateways are implemented today on special-purpose hardware to achieve high packet processing and input/output throughput. Running those functions on current off-the-shelf servers with current hypervisors can lead to a performance degradation by a factor of ten. The industry is currently working on new technologies that have the potential to improve data plane performance on commercial off-the-shelf servers in some cases to nearly the levels of special-purpose hardware.

Data plane performance is a fringe activity in the **OpenStack** community. There is some preliminary work on PCI passthrough to support Intel® Single Root I/O Virtualization technology.

CloudBand — as an NFV platform — is committed to offering a variety of technologies to optimize data plane performance. Alcatel-Lucent has established a partnership with Intel, the leader in data plane optimization solutions. Objectives of this partnership are the exchange of requirements and collaboration on bringing these technologies into the CloudBand Node and into the OpenStack community.

CONCLUSION

Service providers are looking for an open, multi-vendor NFV platform that allows them to realize the benefits of NFV. CloudBand uses OpenStack as the VIM and complements OpenStack in three areas to deliver a unique platform optimized for NFV.

- The OpenStack community is not strongly focused on some of the NFV requirements, such as distribution, networking, operational optimization, and data plane optimization. An NFV platform, such as CloudBand, goes beyond the scope of OpenStack in these areas to help customers realize the promise of NFV in the areas of CAPEX and OPEX reduction and improved service agility.
- OpenStack is still under heavy development in many areas. As it matures, OpenStack will become more stable and richer in functionality, allowing it to better meet NFV requirements in certain areas. However, it is not expected to meet all requirements. Alcatel-Lucent works with partner Red Hat and others to fill the gaps through contributions to the OpenStack community.
- OpenStack is free software but not ready to use out of the box. OpenStack is a collection of a growing number of software modules. Any installation of OpenStack needs to be integrated and many configuration choices need to be made. In any production environment, bug fixes and security patches need to be ported to all active versions and the system needs to be tested rigorously. CloudBand uses Red Hat Enterprise Linux OpenStack Platform, a familiar and widely supported environment, in the CloudBand Node along with additional software elements to build a plug-and-play NFV infrastructure.

Service providers will find that CloudBand enables them to fully leverage the advantages of NFV.

For further reading:

- [1] Telecom grade cloud computing, May 2011, SCOPE Alliance.
(http://scope-alliance.org/sites/default/files/documents/CloudComputing_Scope_1.0.pdf)
- [2] OpenStack and Red Hat, white paper sponsored by Red Hat, IDC.
(<https://engage.redhat.com/content/outlook-for-openstack-201309131116>)
- [3] Why Service Providers Need an NFV Platform, Alcatel-Lucent.
(<http://resources.alcatel-lucent.com/?cid=170811>)

ACRONYMS

| | |
|---------|--|
| API | Application programming interface |
| BSS | Business support system |
| CAPEX | Capital expenditures |
| EMS | Element management system |
| GUI | Graphical user interface |
| LCM | Lifecycle management |
| MP-BGP | Multiprotocol Border Gateway Protocol |
| NFV | Network Functions Virtualization |
| NFVI | Network Functions Virtualization Infrastructure |
| NFVO | Network Functions Virtualization Orchestrator |
| OASIS | Organization for the Advancement of Structured Information Standards |
| OPEX | Operating expenditures |
| OSS | Operations support system |
| PCI | Peripheral Component Interconnect |
| SDN | Software-Defined Networking |
| TOSCA | Topology and Orchestration Specification for Cloud Applications |
| TripleO | OpenStack on OpenStack |
| vBNG | Virtualized Broadband Network Gateway |
| vCDN | Virtualized Content Delivery Network |
| vCPE | Virtualized customer premises equipment |
| VIM | Virtualized infrastructure manager |
| vIMS | Virtualized IP Multimedia Subsystem |
| vNF | Virtualized network function |
| vNFC | Virtualized network function component |
| VNFM | Virtualized network function manager |
| VPN | Virtual private network |
| VRS | Nuage Networks Virtual Routing and Switching |
| vSBC | Virtualized Session Border Controller |
| VSC | Nuage Networks Virtualized Services Controller |
| VSD | Nuage Networks Virtualized Services Directory |
| VSG | Nuage Networks Virtualized Services Gateway |
| VSP | Nuage Networks Virtualized Services Platform |
| WIM | WAN Infrastructure Manager |