

# NFV Interoperability Evaluation Results

**Phase 1: Multivendor NFV infrastructure-to-virtual network function (VNF) interoperability**

# contents

Introduction.....	03
The Participants.....	04
Test Setup and Coverage .....	04
The Results Matrix .....	07
Interoperability Findings and Challenges .....	08
Further Interoperability Findings and Challenges .....	09
NFV interoperability evaluation report cards.....	12
Alcatel-Lucent Cloudband .....	12
Cisco Systems NFVI .....	20
Huawei FusionSphere .....	27
Juniper Contrail .....	30

## INTRODUCTION

Of all the industry test results out there, this is the one that people have been waiting for.

In a historic first for the communications networking industry, Light Reading, in partnership with EANTC on behalf of the New IP Agency (NIA), is happy and proud to announce the publication of the world's first independent interoperability evaluation of NFV infrastructure, focused (in Phase 1) on multivendor NFV infrastructure-to-virtual network function (VNF) interoperability.



These results are just a first step on a very important journey for an industry that is in dramatic upheaval. Never before have network operators, faced with the unknown and daunting challenges of introducing virtualized functions and their supporting infrastructure platforms, needed independent test results as much as they need them now: Such independent evaluations can cut down development times, save operators resources, and help speed up further tests, field trials and live deployments.

That's why we have been moving fast. The Phase 1 evaluation, which was publicized, set up and completed in just a few months (an extraordinarily short time for such a procedure), included 12 vendors, which between them submitted four NFV infrastructure (NFVi) platforms and 17 VNFs.

The NFVi vendors, which all offered platforms based on OpenStack, were:

- Alcatel-Lucent
- Cisco Systems
- Huawei Technologies
- Juniper Networks

The VNF vendors were:

- Alcatel-Lucent
- Cisco Systems
- Cobham Wireless
- Hitachi Communication Technologies America
- Huawei Technologies
- IneoQuest Technologies
- Juniper Networks
- Metaswitch Networks
- NetNumber
- Netrounds
- Procera Networks
- Sonus Networks

The VNFs came in many shapes and forms, from virtual routers, EPC (evolved packet core) and IMS components, firewalls, test probes, deep packet inspection (DPI) modules and session border controllers.

The testing focused on the functional interoperability of NFVi platforms and various VNFs, but did not include the sort of tests more commonly associated with communications infrastructure, such as performance, scalability and resilience testing: In addition, these tests did not include any management and orchestration (MANO) evaluations (they will be the focus of a later evaluation phase).

The tests carried out by the EANTC team, which were all conducted remotely (another new feature!), were of NFVi-VNF interoperability combinations that had never been tried before in such open conditions. As a result, not all were successful, but as this is still a very nascent market, it was expected that some of the evaluations were not be completed successfully.

In fact, the success rate — 25 of the 39 combinations tested, or 64%, passed — was "a great result," noted EANTC Managing Director Carsten Rossenhövel.

That higher-than-expected pass rate was due, in part, to the high degree of industry participation: During the seven weeks of testing, 55 engineers from various vendors (including test system partner Ixia) were involved, with support from 14 marketing and communications

executives, seven EANTC staff and three from Light Reading/NIA: In total, 79 people were directly involved in the process.

So does that mean that there was a 36% fail rate? It could be viewed that way, but that would be shortsighted. The point of these evaluations is not to identify technology that 'does not work' — it is to see what can be done better, what needs to be fixed. The NIA exists to advance the industry and speed up the process of getting useful technology to market: When the evaluations highlight an interoperability combination that doesn't work, that means the vendors can fix the problems! This can only be good for the industry.

The evaluation also highlighted many other interesting and useful outcomes, not least of which was insight into the performance of OpenStack. For some OpenStack supporters, it is a viable alternative to standardization: "If only everybody would base their implementations on OpenStack, all interop issues would be solved," could be a phrase you might hear in virtualization circles.

Well, according to EANTC's results, OpenStack is far from ready to play an interop role. "There were tons of interop issues despite the fact that all NFVis were based on OpenStack," noted EANTC's Rossenhövel. That's because there are far too many options (including how to manage a VNF), too many versions (with backwards compatibility an oversight), and items left open for vendor consideration (including license management).

So there are many positives to take away from the Phase 1 test results, but there is much, much more to be done: A test combination that passed confirms that the initial lifecycle management tasks work between the parties on a functional basis, but that still leaves many aspects of application-layer functionality, performance, resiliency and orchestration to be evaluated. Those aspects will be the focus of future evaluations.

In the meantime, the following pages will provide you will a full matrix of the test results and detailed insights into the evaluations, the outcomes and the conclusions. So dive in!

— Ray Le Maistre, Editor-in-Chief, Light Reading and Carsten Rossenhövel, Managing Director, European Advanced Networking Test Center AG (EANTC) (<http://www.eantc.de/>), an independent test lab in Berlin. EANTC offers vendor-neutral network test facilities for manufacturers, service providers, and enterprises.

## THE PARTICIPANTS

The invitation to participate in the NIA's Phase 1 test was issued in August. Once we had published the invitation, NFV infrastructure (NFVi) and virtual network function (VNF) vendors responded very quickly and in a very positive manner. In total, four NFVis participated in the first test campaign:

- Alcatel-Lucent CloudBand
- Cisco NFVi
- Huawei FusionSphere
- Juniper Contrail

All NFVis were configured as data center solutions, ready for deployment at centralized service provider infrastructure locations. In addition, 17 VNFs from 12 vendors lined up to be tested with these NFVis. They were:

- Alcatel-Lucent VSR (virtualized service router)
- Alcatel-Lucent VMG (virtualized mobile gateway, part of the mobile Evolved Packet Core [EPC])
- Alcatel-Lucent VMM (virtualized mobility manager, part of the mobile EPC)
- Cisco ASA (virtual firewall)
- Cisco CSR1000v (virtual router)
- Cobham Wireless TeraVM (Virtual Application Emulation and Security Validation)
- Hitachi vMC (virtual mobile core — consisting of EPC components MME, SGSN, SCGW, uEPC)
- Huawei VNE (virtual router)
- Ineoquest IQDialogue ASM (virtual video test probe)
- Ineoquest DVA (virtual video quality probe)
- Juniper vMX (virtual router)
- Juniper vSRX (virtual firewall)
- Metaswitch Perimeta vSBC (virtual session border controller [SBC])
- NetNumber TITAN (virtual centralized signaling and routing control [CSRC])
- Netrounds Test Agent (virtual IP application probe)
- Procera PacketLogic/V (virtual deep-packet inspection [DPI] filter)
- Sonus SBC SWe (virtual SBC)

The wide variety of participating VNFs shows that the industry is progressing well in terms of virtualizing all sorts of network functions. Some VNF types, such as virtual routers, are typically quite straightforward regarding their internal structure and the migration from physical to virtual. Initially we had expected that most of the VNFs put forward for this evaluation would be of this “low-hanging fruit” type. As it turned out, quite a few of the more complex virtual mobile core implementations (EPC, IMS, SBC) were submitted for testing, indicating that these types of applications are more ready for NFV than we had thought.

This initial round of testing, being focused on basic lifecycle management, did not do justice to complex VNFs: As one of the vendors rightfully pointed out, the “cloudification” of network functions involves substantial modification of the physical network function code base. An interesting future test area will focus on how VNF vendors enable elastic, interoperable cloud-based services.

## TEST SETUP AND COVERAGE

This section describes the technical coverage, logical topologies and physical test bed, and explains what the ‘Pass’ entries in the test results matrix (see next section) actually mean.

(If you’re desperate to see the results, please feel free to skip forward and come back later.)

When we started the NIA test program, a couple of factors quickly became obvious:

1. This is a big industry with a lot of new VNFs from many established vendors and startups. Asking all vendors to travel somewhere for some joint testing slots would not scale. We had to come up with a distributed, remote support scenario with scheduled test slots.
2. “Virtual” does not mean “interchangeable.” Initially we thought about running all the tests on a common infrastructure. An expectation of NFV is the removal of hard-coded linkages between software and hardware and, consequently, the NFVi software should support a range of third-party hardware. Some vendors informed us this would be possible but it would be a separate testing topic that was outside the scope of this first testing phase. It should be noted that this open interoperability does add time and cost to testing and therefore, in practice, may limit the degree of openness deployed in networks. Using public cloud providers was not an option for our evaluation setup, not only because of hardware interoperability but also because none of them provided the networking infrastructure (NICs, switches, internal IP addressing) and flexible/secure connectivity we needed.
3. Reporting results transparently has always been a big asset of EANTC tests. To help the industry to accelerate and to save service providers lab testing cycles, our results need to be documented and detailed enough that they can be referenced clearly and reproduced by anybody anywhere.
4. At this early stage of the industry, more time is actually required to ‘bring up’ VNFs — that is, integrate them with individual NFVis — than for the actual testing. The integration time ranged between a couple of hours to four weeks per VNF. This is a major factor; even exchanging the required information between the vendors and the test lab is not easy, as none of these descriptors are sufficiently standardized.
5. The human factor still governs all activities. Tests can be as virtual as possible, but the integration and troubleshooting support is



still provided by humans. (Well, so we assume... we did not see all faces on Webex...) Support times need to be coordinated, cultural differences need to be taken into account, and so on. The evaluation involved people located in 11 countries and multiple time zones, adding delays and reducing available test time, which ultimately impacted the number of VNF combinations tested.

## PHASE 1 EVALUATION SCOPE

This report covers the first campaign of the New IP Agency's NFVi-VNF interoperability program. We chose this topic because there is the highest likelihood in the NFV environment that virtual network functions will be provided by multiple entities and they will all need to work with the deployed NFVi. In addition, there is already a sizeable set of options and combinations — there are already more than 300 VNFs and 20 NFVis commercially available already.

We focused the tests on VNF lifecycle management, as set out in the ETSI GS NFV-MAN 001 Appendix B, the well-known 'ETSI Phase 1 MANO' document. Our intent was not to test the individual applications for functionality but to evaluate the interoperability of the VNFs' management aspects with different NFVis.

Future NIA testing phases will focus on high availability, management and orchestration, fault management and other aspects of the NFV framework as soon as the industry is ready for multivendor interoperability testing in these areas.

## TEST BED TOPOLOGY

The test topology was fairly straightforward. The four NFVi vendors, our testing partner Ixia and EANTC set up the hardware in our lab in Berlin. Each NFVi vendor provided its own hardware (for the full details, see the appendix):

- Alcatel-Lucent Cloudband came with a full OpenStack deployment (three control and eight compute nodes including local storage, plus a router for internal connectivity).
- Cisco NFVi provided a full OpenStack deployment as well, including three redundant control nodes, three compute nodes, a triple-redundant storage system, two build nodes and a router plus a switch connecting the infrastructure internally.
- Huawei provided FusionSphere with three hot-standby control nodes, one compute node with a lot of local storage, and a physical router for internal connection between servers.

- Juniper provided one Contrail server which implemented the OpenStack environment using the virtual deployment option: one virtual control node, one virtual compute node, and a virtual router. This satisfied the minimum testing requirements in this phase.



*The NIA Phase 1 NFV interoperability test bed at EANTC's headquarters in Berlin.*

## TEST METHODOLOGY

With each NFVi-VNF combination, we covered the following mandatory test cases:

- 1. VNF Package on-boarding**  
Load the VNF package image into the NFVi using either the virtual infrastructure manager (VIM) or the VNF manager; verify that the loaded package is in an integral shape and contains all the required configurations, scripts and binary sources.
- 2. Standalone Instantiation**  
Instantiate a single VNF copy including resource allocation by the VIM, initiating the boot sequence and other configurations.
- 3. Graceful Termination**  
Validate that the NFVi releases resources upon graceful termination and that VNFs can be restarted without any further manual intervention required.
- 4. Forceful Termination**  
"Power-off" the VNF forcefully through the VIM; verify that all resources are released and related runtime entries cleared and that the VNF can be restarted subsequently.

## 5. VNF Package Deletion

Verify that the deletion of the VNF package removes all related files and configurations without affecting other packages.

In addition, there were the following ancillary test cases:

- 1. VNF Instance Modification**  
Verify that certain modifications of a running VNF instance trigger a notification event by the NFVi and that the VNF continues to function properly: Add one more virtual network interface; configure VLAN and IP for the newly added interface; add one more CPU core; retrieve current negotiation.
- 2. Restart of One Virtual Instance**  
Verify the effects of restarting one VNF instance during live operation of a second VNF of the same type; validate that the VNF instance gets restarted correctly while traffic through the second VNF is unaffected.
- 3. Persistence and Stability**  
Test how the VNF survives forced removal of a virtual port in use and subsequent reconnection of that virtual port. The test used two variants — interface detach and port status update.

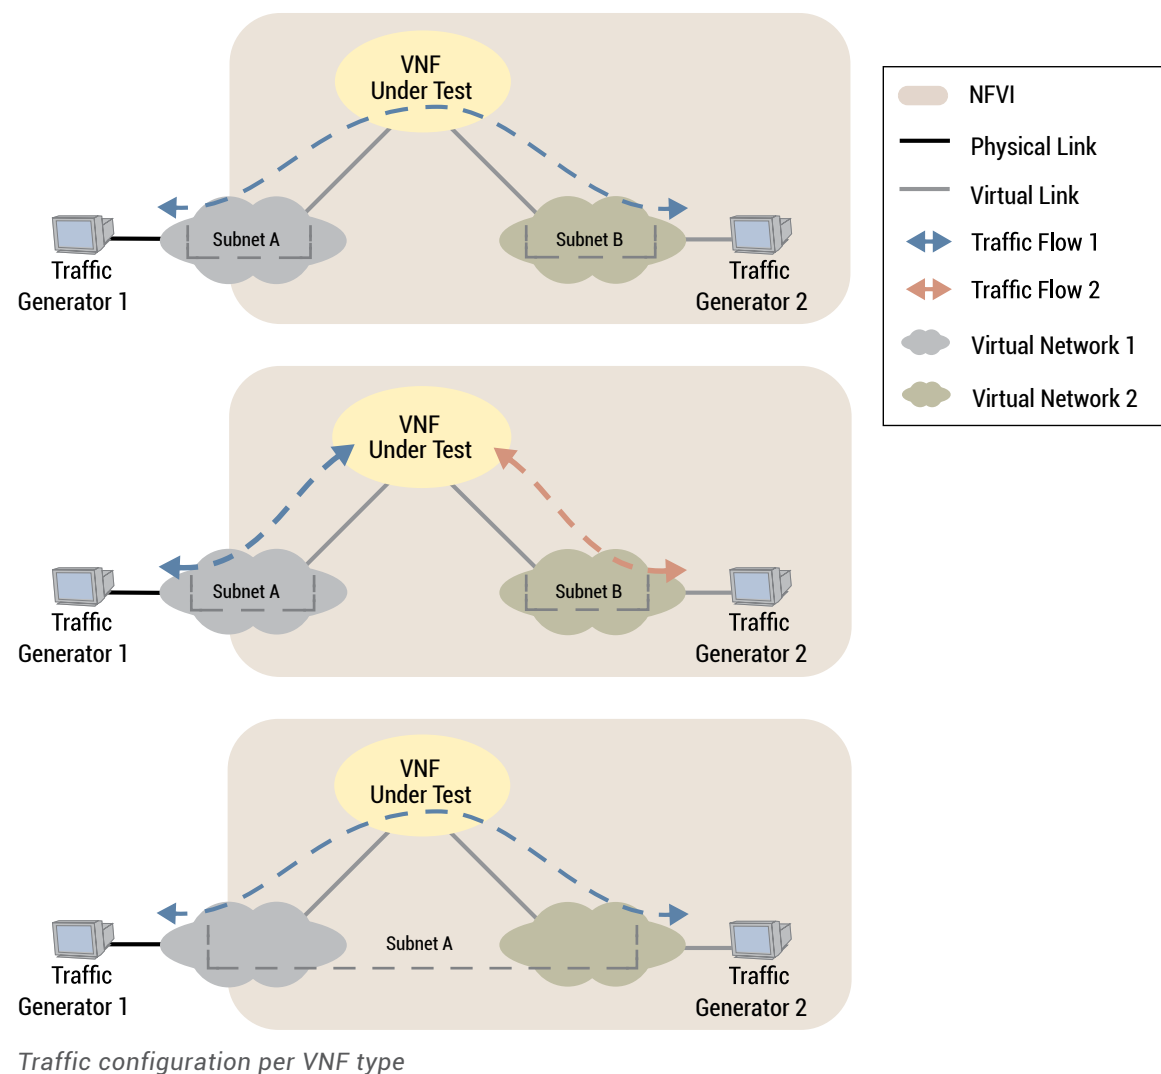
## DATA PLANE VERIFICATION (TRAFFIC CONFIGURATION)

From the start, we wanted to verify the virtual switch data plane connectivity as part of the interoperability test campaign. The virtual switch is an important part of the NFVi; the correct connectivity of virtual ports with a VNF is fundamental to any network-related VNF's functions. When estimating potential areas of potential interoperability issues with OpenStack, we guessed that network functions would yield the highest rate of problems: Traditional enterprise cloud applications have been running on OpenStack for years. They are compute- and storage-heavy but have never cared so much for versatile networking options (and performance, of course, but that is out of scope here).

We wanted to evaluate the vSwitch/VNF interoperability in as complete a way as possible, so we defined a network configuration by attaching the VNF to a physical tester port on one side (the “left network”) and to a virtual tester port on the other side (the “right network”).

With help from Ixia, we configured and pre-staged the test setup in our lab. The physical port was quickly connected with a legacy Ixia Ethernet load generator port – we did not require anything special, since this test focuses only functional data plane aspects. In addition, the Ixia team deployed a centralized Chassis VM and Client user interface VM (both currently running on Microsoft Windows), plus a virtual load module per NFVi (running Linux as the guest OS). The client user interface VM managed both the virtual and real chassis and sent bidirectional traffic when needed. We used two Ixia test tools – IxNetwork for VMs that were based on IP traffic and IxLOAD for VMs that required application-layer emulation.

All virtual routers, virtual firewalls and virtual DPI systems were tested with IP traffic. Most of the virtual EPC, IMS and SIP solutions were tested with Voice over IP (SIP) traffic, except for a few that were configured to just pass native IPv4 traffic. Due to the issues that some NFVis and VNFs have with IPv6 at this point, we deviated from our generic rule to always use dual-stack IPv4 and IPv6 data plane configurations. Instead, the whole test bed (shamefully!) used only IPv4 for VNFs and for management.



## THE RESULTS MATRIX

The results matrix is the heart of this report. It confirms successful test combinations of NFVis and VNFs. For a successful test, all mandatory (critical) test cases had to be executed and passed:

- VNF package on-boarding
- VNF instantiation
- VNF forceful and graceful termination
- VNF package deletion

In addition, there were a number of ancillary test cases – multiple instance co-existence, modification of vSwitch parameters during VNF operation – whose results are not reported in this matrix but which are discussed in the next section (Interoperability Findings and Challenges).

Given the large number of potential test combinations ( $4 \times 17 = 68$ ) and the limited testing time of only two months, EANTC did not aim for a full-mesh test campaign as we knew it would not be possible to complete. We aimed to get two test combinations per VNF completed, which worked out in some cases but not in others, often due to integration, configuration or troubleshooting delays (or simply lack of EANTC test slots). Any cell that is left empty in the matrix below relates in most all cases to a combination that was not tested. Fundamental interoperability issues are reported in the next section. In total we planned for 39 test combinations, of which 25 passed, six finally failed, seven were postponed due to configuration or support issues or troubleshooting, and one was still in progress at the time of the report deadline.



Of the 39 tests, 25 passed (64%)

The usual question at this point is: Why does EANTC not report failed test combinations individually? Well, that wouldn't serve our target audience for two key reasons: First, vendors that took the effort to participate should not be punished – and they will all go back to their labs and immediately improve the implementations with the goal of coming back to pass the tests next time. The vendors that should be called out are those that do not invest in multivendor interoperability. As a reader, feel free to inquire with the companies that opted not to participate in this (or the next) round of testing.

The second reason for not publishing failures is that we want to conduct tests using the most advanced implementations. Experience has shown that vendors bring only their most proven, legacy solutions if there is a plan to publicly identify any failures. Obviously that would be counter-productive and not make any sense for the industry, specifically not in an NFV context. Our tests can only help the industry to progress if we are testing with the very latest implementations: That's why we need to safeguard the testing environment with non-disclosure agreements, making sure failures are permissible and encouraged, because failures help vendors and the industry to progress.

Participant vendors are permitted to review the test report to improve its quality and accuracy, but they are not permitted to request selective publication of specific results only. Of course vendors can ask to opt out of the results publication completely, in which case their name would be mentioned without products or results associated.

All VNF vendors were asked to select NFVis to test with at the beginning of the test campaign, based on a first-come, first-serve approach. Vendors that signed their contracts and assembled their support teams early had an advantage. Some combinations had been assigned by EANTC, disregarding vendor preferences, to balance the test bed.

We decided not to test interoperability of a VNF with an NFVi from the same vendor: While there is some merit to completing the table, we would expect that there will be a much lower likelihood of interoperability issues when testing solutions from a single vendor. We declined any such test requests; the associating cells are marked as N/A.

Due to shipping logistics, Huawei's FusionSphere was available for only 40% of the testing time, which explains the small number of combinations tested.

Test Results Matrix	Alcatel-Lucent Cloudband	Cisco NFVI	Huawei FusionSphere	Juniper Contrail
Alcatel-Lucent VSR	N/A	PASS		PASS
Alcatel-Lucent VMG	N/A			PASS
Alcatel-Lucent VMM	N/A	PASS	PASS	
Cisco ASAv	PASS	N/A		
Cisco CSR1000v	PASS	N/A	PASS	
Cobham Wireless TeraVM				PASS
Hitachi vMC	PASS			PASS
Huawei VNE	PASS	PASS	N/A	
Ineoquest IQDialogue ASM		PASS		
Ineoquest DVA		PASS		
Juniper vMX	PASS			N/A
Juniper vSRX	PASS			N/A
Metaswitch Perimeta vSBC		PASS		PASS
NetNumber TITAN	PASS			PASS
Netrounds Test Agent	PASS	PASS		
Procera PacketLogic/V			PASS	
Sonus SBC SWe				PASS

## INTEROPERABILITY FINDINGS AND CHALLENGES

We unveiled a number of important findings and encountered a number of challenges during the Phase 1 evaluation.

### VNF CONFIGURATION OPTIONS WITH AND WITHOUT HEAT

There are many tools that OpenStack provides to instantiate a VNF. They range from a simple instantiation with nova commands and a manual configuration, or pushing a configuration file with the boot command itself, to using a Heat template or even a separate VNF Manager. Each of these methods has its merits and challenges.

Once the VNF has booted, it needs to find its network-layer and optionally application-layer configuration file in a second step. This configuration file will define network interface configurations, passwords, licenses, etc. Again, there are different ways to provide this functionality:

### NOVA

The basic way to instantiate a VNF is using OpenStack's nova commands (or the GUI wrapper). The operator has to create a flavor, defining resource allocations for virtual CPU, storage and RAM. Subsequently the VNF can be started with a nova boot command. Next, the network layer configuration needs to be supplied to the VNF – this is done by pushing a management network port to the VNF as the first network parameter. The VNF usually requests an IP address through directed DHCP, which can then be used to contact it and start application-layer configuration.

### NOVA WITH CONFIG DRIVE

It is possible to simplify the second step of the VNF instantiation process by supplying the VNF with a configuration drive access. The operator will upload the VNF config file in advance and will provide the file name as part of the nova boot command. This avoids the need to bootstrap through a management IP address. Specifically, if the VNF does not want to use DHCP to identify its IP configuration, this option is helpful. With this method, the VNF first has to complete its bootup process before it can access the configuration file.

### CLOUDINIT SCRIPT

The CloudInit script comes in when the Config Drive cannot be used – in case the VNF needs to be configured at boot time, not after boot time. In this case, the CloudInit script(s) can be pushed by nova and can be accessed via the VNF at boot time.

The three options above are supported by any OpenStack implementation.

But there are options that may, or may not, be supported by any OpenStack implementation.

### HEAT

Heat is an OpenStack template that is compiled and then converted into a series of commands that can instantiate a service along with all the dependent resources (networks, policies and VNFs). Heat requires support from the NFVi, which the following implementations provided at the time of our test:

- Alcatel-Lucent Cloudband
- Juniper Contrail

HEAT was the most straightforward and cleanest configuration option for instantiating and tearing down VNFs. However, we observed some challenges defining the templates: Sometimes version numbers did not match (backwards compatibility), and some versions of HEAT are limited in translation – they take only unique network IDs (UUIDs) instead of network names.

VNF tested with NFVI using...	Alcatel-Lucent Cloudband	Cisco NFVI	Huawei FusionSphere	Juniper Contrail
Alcatel-Lucent VSR		Nova		Nova
Alcatel-Lucent VMG				Nova + User Data
Alcatel-Lucent VMM		Nova	Nova	
Cisco ASAv	Nova + Config File			
Cisco CSR1000v	Nova + Config File		Nova + Config File	
Cobham Wireless TeraVM				Nova
Hitachi vMC	Nova + User Data			Nova + User Data
Huawei VNE	Nova	Nova		
Ineoquest IQDialogue ASM		Nova + Cloud-init		
Ineoquest DVA		Nova + Cloud-init		
Juniper vMX	Nova + Heat			
Juniper vSRX	Nova + Heat			
Metaswitch Perimeta vSBC		Nova		Nova
NetNumber TITAN	Nova + Heat			Nova + Heat
Netrounds Test Agent	Nova	Nova		
Procera PacketLogic/V			Nova	
Sonus SBC SWe				VNF Manager + YAML VNF Templates

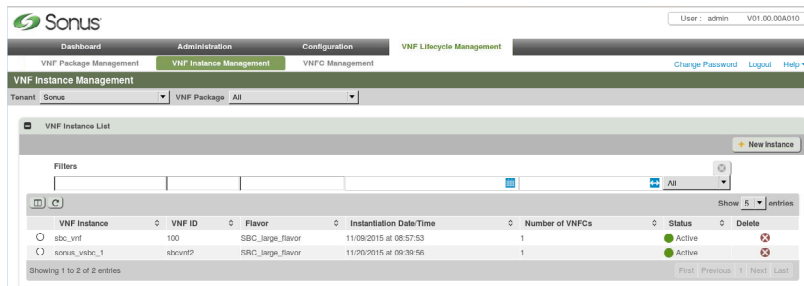
Source: EANTC



## VNF MANAGER

A VNF could also provide its own management entity — a VNF Manager, or VNFM — that could shield all the details of its bootup configuration process from the operator. The VNFM would manage the complete instantiation and configuration process opaquely. There has been an extensive debate in the ETSI NFV group whether VNFs should bring their own VNFM or use a generic VNF manager.

Some vendors elected to go with OpenStack Heat templates for the sake of open interoperability. Only one vendor, Sonus, brought a VNF manager. In this instance, Sonus provided a YAML-formatted VNFD (VNF Descriptor) file describing the VNF parameters, which were then used by the VNF manager to configure the VNF.



The Sonus VNF Manager

## FURTHER INTEROPERABILITY FINDINGS AND CHALLENGES

There was much to learn from the Phase 1 process.

### DIFFERENCES IN OPENSTACK VENDOR IMPLEMENTATIONS

One of the big fun areas with open source projects is their wealth of deployment options and general lack of backwards compatibility awareness. This is one of the reasons why RedHat has been successful with RHEL (RedHat Enterprise Linux): It validates and ensures consistency and backwards compatibility of Linux packages in its distribution. Unfortunately, such a thing is still lacking in OpenStack.

Three of the four NFVi vendors implemented OpenStack version Juno (released Q1, 2015), while one provided Icehouse (released Q2, 2014), which had an incompatible Heat version. Some of the more complex VNFs did not manage to boot up on Icehouse, and it would have been too cumbersome to configure them manually without Heat templates.

We intentionally did not request specific OpenStack versions. In the near future, service providers will be faced with data centers all based on different hardware and software combinations. Backwards compatibility is an important aspect to look at and clearly the industry needs to focus on it more extensively.

When the Icehouse NFVi vendor wanted to upgrade its servers, we noticed the next issue: In-service OpenStack upgrades will be a major challenge. It will be an interesting test area for the future to test in-service NFVi upgrades: One NFVi vendor was already willing to demonstrate this feature in the next round of testing.

### MULTIVENDOR TROUBLESHOOTING OF VNF STARTUP ISSUES

When there are issues in bringing up VNFs, it is obviously important to identify the root cause efficiently. In our test bed, all vendors were fully motivated to identify and resolve issues swiftly, as each issue resolved prior to customer deployment saves a lot of budget and stress. Nobody was interested in fingerpointing, which may be different in a service provider environment, where a lot of business may be at stake if a NFVi/VNF combination does not work.

The multivendor interoperability troubleshooting techniques and methods are different than those required for operational problems. While it is important to get a very quick overview of the battlefield in

case of operational issues, functional interoperability problems are best solved if full and detailed information can be made available from either side for analysis by the involved vendors.

To this end, OpenStack offers many logs. It turned out that one of the most insightful troubleshooting “tools” was simply the console log of the VNF, including its bootup messages. We had multiple cases where a VNF would crash during the boot process: Accessing its kernel panic or crash message was instrumental for the VNF vendor(s) to identify the root cause.

Simply getting out a text file with the console log required different methods in different NFVis: Graphical user interfaces shielding all the dirty CLI details from the operator do not help in this case. For most NFVi vendors it was possible to access the console logs in one way or another, but one of the participating NFVis unfortunately did not save the console log in instances where a VNF crashed early during the boot process. The simple live view of the console was insufficient, as text scrolled too fast — even recording the console output with a camera did not help. In the end the VNF vendor had to give up and the combination was declared failed due to the inability to identify the root cause. NFVi vendors should include feature-rich and verified troubleshooting solutions for VNF interoperability testing.

### OPENSTACK IP ADDRESSING AND VNF IP ADDRESSING

Traditionally, OpenStack assigns a single IPv4 address (using Directed DHCP) to each virtual network port of the VM during bootup. That's probably reasonable for enterprise cloud environments, but for virtual network functions, this mechanism is less useful as it may limit the VNF too much, or cause confusion to the operator if different VNFs use different methods.

The four OpenStack NFVis in our test supported a range of IP address assignment options as follows:

1. Directed DHCP for all ports, pushing an IP address by OpenStack directly — can be used for the management port only.
2. Directed DHCP — for all ports.
3. Plain DHCP: OpenStack maps complete subnets to virtual network ports, spins up a DHCP server and will subsequently always respond to any DHCP requests on that port.
4. Boot config: Bypassing DHCP, the boot configuration may include static IP addresses as mentioned in the startup option subsection

above. While the VNF may ignore the OpenStack IP addressing, OpenStack management tools will still assume the supposedly assigned IP addresses during the virtual port creation. This will create operator confusion.

Initially, we did not focus on IP addressing, so we went along with any suggestions NFVi and VNF vendors made. As we moved ahead with the VNF configuration testing, we noticed that a stricter rule — preferably using Directed DHCP — would ease the IP address management in a multivendor compatible way.

### LICENSE MANAGEMENT CONSIDERATIONS

All of the VNFs under test were commercial solutions. As such, they require some sort of licensing. As it turned out during the test campaign, licensing was one of the least mature areas. By their nature, none of the open source projects addresses commercial licensing in a standardized way.

Traditional methods of software licensing do not work well or do not seem to fit in a virtualized environment. VNFs are instantiated, moved around the data center, terminated and re-instantiated as needed. Each time, they are given a new unique identification number so it is impossible to tie their existence to any static license. Thus far, most of the participating VNFs lacked a proper licensing mechanism that would fulfill all our requirements. Many have expressed their awareness that licensing needs to be improved and made an integral part of the VNF management in the future.

In the test bed, we observed the following ways to implement licenses:

#### 1. Local license manager

The VNF requires a local license manager on instantiation. This solution required some planning to install a license manager in our lab environment and to route requests to the VNF over the internal management network. Once that was done, the local license manager solution worked fine. From the vendor perspective, this solution may still not be secure (see discussion below this list).

#### 2. Public license server

A license server in the cloud sounds easy: No setup is required and it appears to be the most manageable and secure option for VNF vendors. Operationally, it may or may not be applicable to production environments. A service provider would need to open their management network to specific Internet locations and services, at least for outbound HTTPS connections. Even in our

lab environment, we did not permit direct communication between VNFs and the Internet, so this option was out of question. For one vendor that needed it (because a local license manager was offered for production deployments only), we helped ourselves with an application-layer gateway (HTTP forwarding).

#### 3. Pre-licensed image for a time limit

Any of the following three options apply traditional licenses (i.e. encrypted strings) to the VNF. Since it is difficult to identify any unique aspect of a VNF, it is probably most realistic to just not tie the license to anything but the system clock. In this case, however, there is no full guarantee that the customer might not reuse a license multiple times. Typically, the license is applied during instantiation time. The extra (manual) steps required make this option useful only in case of small-scale deployments that do not require automatic scale-out operations.

#### 4. Pre-licensed image tied to MAC address

Some vendors, looking for unique identification of the VNF instance, resorted to the MAC address. This is obviously useless as the MAC address can be set with nova commands. It just ensures that there won't be two VNF instances using the same license on the same Ethernet/vswitch management network segment.

#### 5. Pre-licensed image tied to UUID

Some vendors found a unique identification of the VNF instance — the UUID. This is actually a correct and suitable unique ID assigned by OpenStack. It cannot be manipulated by the operator. Unfortunately, the UUID changes on each and every administrative action. A license tied to the UUID does not allow relocation of the VNF, or even termination/re-instantiation.

None of the license management options seemed very suitable. We feel that the industry has to put more thought into standardization in this area.

The unfortunate reality is that it is not possible to come up with a secure licensing solution in an elastic environment without the vendor being able to establish a trusted, uncloneable “anchor” point for VM licensing. The most favorable-looking option from the list above — option 1, the local license manager — provides this “uncloneability” characteristic only if used in conjunction with a physical license token (such as a USB key) that ties the server to a fixed set of host machines, or a Web-based license soft token.

With any license server (whether local or in the Internet), resilience is a concern. What happens if the license server becomes unavailable? One

vendor informed us that their clients raise an alarm if the connection to the license server is lost; the customer has a whole month to react to the issue before the system enters its locked-down unlicensed state. (We have not tested this aspect yet.)

In the lab, there was an additional solution that worked most easily out of the box: Some vendors implemented a license-less mode where the VNF would just support a tiny bit of throughput, for example 0.2 Mbit/s. Cisco's and Huawei's virtual routers, for example, supported such a mode. Since we wanted to conduct functional tests only, these modes were sufficient.

What works great in the lab may be an operational concern: If a license server might become unavailable so that the VNF falls back to tiny-throughput mode, the issue might go undetected for a while if the VNF is idling. Orchestrators should probably include license management monitoring methods for each VNF instance and the license manager itself — especially if the license manager might become a single point of failure. (This is just some food for thought, but worth noting we think.)

### VSWITCH CONFIGURATION OPTIONS

Some VNFs — implementations that consisted of multiple VMs — required jumbo frame support (beyond 1518 byte-sized Ethernet packets) for internal communication. Some of the participating NFVis supported jumbo frames by proprietary OpenStack code modifications as part of their software release; one implemented a temporary workaround that required manual intervention each time a VNF was instantiated. The lesson learned is that jumbo frame support cannot be taken for granted: One vendor told us that OpenStack Kilo will support jumbo frames out of the box, but another vendor contested any dependency on OpenStack versions.

### EARLY THOUGHTS ON VNF/NFVI SECURITY

All NFVis, with the exception of one, fully trusted any management requests received via their REST APIs. Any OpenStack commands received by the NFVi were executed without requiring transport-level security (HTTPS) or validating a certificate. With this, it would be relatively easy to manipulate an NFVi once transport-level access has been achieved by any tenant. This is unfortunate, as OpenStack supports HTTPS and certificates. The one vendor that used and required this transport-level security (HTTPS), however, used self-signed certificates, which made it difficult for some VNFs

to interoperate. Certificates had to be exchanged between NFVi and VNF externally, and some rejected the self-signed certificates. This highlighted an issue well known from the encryption industry: Certification management is not straightforward and provides a challenge comparable to the license management described above.

In OpenStack, tenants get granular operational rights assigned to instantiate VNFs. Sometimes, simple configuration on OpenStack was not sufficient: We needed to upgrade some of the tenants to gain administration rights in order to spin up their VNFs successfully. The standard OpenStack configuration does not allow sufficient customization of access rights to tenants, as would be required in a service provider environment with different administrative groups operating the NFVi and a range of VNFs.

For example, a `heat_stack_owner` permission should allow a tenant to use Heat for instantiating a VNF. However, sometimes this was not sufficient because tenants needed to create internal networks for communication between multiple VMs and part of the VNF. This was only possible with a more privileged account.

## CPU FLAGS AND INSTRUCTION SETS

Intel continues to evolve x86 instruction sets. NFV requires much more network throughput (i.e. copying to/from network interface cards) than other x86 use cases have needed in the past. Intel developed the Data Plane Development Kit (DPDK) to help developers access the network data plane more efficiently. Some of these modifications use Intel's x86 instruction set extension SSSE3 (Supplemental Streaming SIMD Extensions 3). Although we did not test performance in this program, we still came across instruction set issues: Some VNFs waited for certain CPU compatibility flags to be sent by the hypervisor (KVM). SSSE3 was supported by all CPUs in the test bed, but one of the VNFs required the appropriate flag to be transferred from KVM, as it wanted to check whether Intel DPDK 2.1 requirements would be satisfied. The test combination got stuck because the CPU flag could not be properly conveyed. DPDK was officially out of scope in this functional interoperability evaluation, but nevertheless is seemingly needed.

In another case, a VNF required Intel's Integrated Performance Primitives (IPP) library, which, in turn, is based on the SSSE4, SSSE3, SSE, AVX, AVX2 and AVX512 instruction set extensions. These are not supported by all CPU types, obviously – specifically only by newer Xeon processors, not by the Atom family. In fact, hardware

independence and compute/network performance are conflicting goals; VNFs with high performance requirements may well be compatible only with a subset of NFVi hardware platforms.

---

*That concludes the written report. On the following pages you will find the NFV interoperability evaluation report cards for each of the 25 interoperability evaluations that provided a 'Pass' result.*

## VNF-NFVI Interoperability Test Report Card: Alcatel-Lucent Cloudband

Test Campaign	Phase 1 / Campaign 1
Date	2015-11

### VNF

Vendor	Hitachi
Product	Mobility Management Element (MME)
Type	Virtual Mobile Core
Version	8.2.36

### Test Results

Critical Test Cases	Pass / Fail / Inconclusive / Not Applicable / Not Executed
Package Onboarding	PASS
VNF Instantiation	PASS
Forceful VNF Teardown	PASS
Graceful VNF Teardown	PASS
Package Deletion	PASS

### VNF Configuration

Number of Modules	5
Data plane virtual NICs	2
Instantiation Method	Nova + User Data
VNF Management	CLI
License Management	Manual

### Test Configuration

Test Equipment	Ixia IxNetwork, Ixia X16 w/ STXS-4, Ixia Virtual Chassis w/ Virtual Load Module
Connectivity	Left Network: Physical; Right Network: Virtual
Test Bed Topology	Figure 1
Service under Test	Diameter Interface Availability

Test Plan Version	1.0
Lab	EANTC, Germany

### NFVI

Vendor	Alcatel-Lucent
Product	Cloudband
Type	Centralized Data Center
Version	3.0

### NFVI Configuration

CPU	Intel(R) Xeon(R) CPU E5-2660 v2 @ 2.20GHz
Chipset	C600 family
Number of Cores	20
Cores per CPU	10
LL CACHE	25MB
QPI/Speed	2 x 8 GT/s
Memory	128GB
Network Interface Cards	Broadcom Corporation NetXtreme II BCM57810 10 Gigabit Ethernet
BIOS	HP-P75
Host OS	RHEL 7.0
Kernel version	3.10.0-123.20.1.el7.x86_64
Other	N/A
Hypervisor	KVM - QEMU 2.1.2
Virtual Infrastruct Mgmt	OpenStack Juno



## VNF-NFVI Interoperability Test Report Card: Alcatel-Lucent Cloudband

Test Campaign	Phase 1 / Campaign 1
Date	2015-11

### VNF

Vendor	Juniper Networks
Product	vMX
Type	Router
Version	JUNOS 14.1R5.4

### Test Results

Critical Test Cases	Pass / Fail / Inconclusive / Not Applicable / Not Executed
Package Onboarding	PASS
VNF Instantiation	PASS
Forceful VNF Teardown	PASS
Graceful VNF Teardown	PASS
Package Deletion	PASS

### VNF Configuration

Number of Modules	2
Data plane virtual NICs	2
Instantiation Method	HEAT
VNF Management	CLI
License Management	Manual

### Test Configuration

Test Equipment	Ixia IxLoad, Ixia X16 w/ STXS-4, Ixia Virtual Chassis w/ Virtual Load Module
Connectivity	Left Network: Physical; Right Network: Virtual
Test Bed Topology	Figure 1
Service under Test	Layer 3 Routing

Test Plan Version	1.0
Lab	EANTC, Germany

### NFVI

Vendor	Alcatel-Lucent
Product	Cloudband
Type	Centralized Data Center
Version	3.0

### NFVI Configuration

CPU	Intel(R) Xeon(R) CPU E5-2660 v2 @ 2.20GHz
Chipset	C600 family
Number of Cores	20
Cores per CPU	10
LL CACHE	25MB
QPI/Speed	2 x 8 GT/s
Memory	128GB
Network Interface Cards	Broadcom Corporation NetXtreme II BCM57810 10 Gigabit Ethernet
BIOS	HP-P75
Host OS	RHEL 7.0
Kernel version	3.10.0-123.20.1.el7.x86_64
Other	N/A
Hypervisor	KVM - QEMU 2.1.2
Virtual Infrastruct Mgmt	OpenStack Juno

## VNF-NFVI Interoperability Test Report Card: Alcatel-Lucent Cloudband

Test Campaign	Phase 1 / Campaign 1
Date	2015-11

### VNF

Vendor	Cisco Systems
Product	ASAv
Type	Firewall
Version	9.5.1.200

### Test Results

Critical Test Cases	Pass / Fail / Inconclusive / Not Applicable / Not Executed
Package Onboarding	PASS
VNF Instantiation	PASS
Forceful VNF Teardown	PASS
Graceful VNF Teardown	PASS
Package Deletion	PASS

### VNF Configuration

Number of Modules	1
Data plane virtual NICs	2
Instantiation Method	Nova + Config File
VNF Management	CLI
License Management	Manual

### Test Configuration

Test Equipment	Ixia IxNetwork, Ixia X16 w/ STXS-4, Ixia Virtual Chassis w/ Virtual Load Module
Connectivity	Left Network: Physical; Right Network: Virtual
Test Bed Topology	Figure 1
Service under Test	Layer 3 Routing

Test Plan Version	1.0
Lab	EANTC, Germany

### NFVI

Vendor	Alcatel-Lucent
Product	Cloudband
Type	Centralized Data Center
Version	3.0

### NFVI Configuration

CPU	Intel(R) Xeon(R) CPU E5-2660 v2 @ 2.20GHz
Chipset	C600 family
Number of Cores	20
Cores per CPU	10
LL CACHE	25MB
QPI/Speed	2 x 8 GT/s
Memory	128GB
Network Interface Cards	Broadcom Corporation NetXtreme II BCM57810 10 Gigabit Ethernet
BIOS	HP-P75
Host OS	RHEL 7.0
Kernel version	3.10.0-123.20.1.el7.x86_64
Other	N/A
Hypervisor	KVM - QEMU 2.1.2
Virtual Infrastruct Mgmt	OpenStack Juno

## VNF-NFVI Interoperability Test Report Card: Alcatel-Lucent Cloudband

Test Campaign	Phase 1 / Campaign 1
Date	2015-11

### VNF

Vendor	Cisco Systems
Product	CSR1000v
Type	Router
Version	IOS XE 03.16.00.S, CSR1000V 15.5(3)S Release (fc6)

### Test Results

Critical Test Cases	Pass / Fail / Inconclusive / Not Applicable / Not Executed
Package Onboarding	PASS
VNF Instantiation	PASS
Forceful VNF Teardown	PASS
Graceful VNF Teardown	PASS
Package Deletion	PASS

### VNF Configuration

Number of Modules	1
Data plane virtual NICs	2
Instantiation Method	Nova + Config File
VNF Management	CLI
License Management	Manual

### Test Configuration

Test Equipment	Ixia IxNetwork, Ixia X16 w/ STXS-4, Ixia Virtual Chassis w/ Virtual Load Module
Connectivity	Left Network: Physical; Right Network: Virtual
Test Bed Topology	Figure 1
Service under Test	Layer 3 Routing

Test Plan Version	1.0
Lab	EANTC, Germany

### NFVI

Vendor	Alcatel-Lucent
Product	Cloudband
Type	Centralized Data Center
Version	3.0

### NFVI Configuration

CPU	Intel(R) Xeon(R) CPU E5-2660 v2 @ 2.20GHz
Chipset	C600 family
Number of Cores	20
Cores per CPU	10
LL CACHE	25MB
QPI/Speed	2 x 8 GT/s
Memory	128GB
Network Interface Cards	Broadcom Corporation NetXtreme II BCM57810 10 Gigabit Ethernet
BIOS	HP-P75
Host OS	RHEL 7.0
Kernel version	3.10.0-123.20.1.el7.x86_64
Other	N/A
Hypervisor	KVM - QEMU 2.1.2
Virtual Infrastruct Mgmt	OpenStack Juno

## VNF-NFVI Interoperability Test Report Card: Alcatel-Lucent Cloudband

Test Campaign	Phase 1 / Campaign 1
Date	2015-11

### VNF

Vendor	Huawei
Product	VNE
Type	Router
Version	V100R002C00SPC013T

### Test Results

Critical Test Cases	Pass / Fail / Inconclusive / Not Applicable / Not Executed
Package Onboarding	PASS
VNF Instantiation	PASS
Forceful VNF Teardown	PASS
Graceful VNF Teardown	PASS
Package Deletion	PASS

### VNF Configuration

Number of Modules	1
Data plane virtual NICs	2
Instantiation Method	Nova
VNF Management	CLI
License Management	Manual

### Test Configuration

Test Equipment	Ixia IxLoad, Ixia X16 w/ STXS-4, Ixia Virtual Chassis w/ Virtual Load Module
Connectivity	Left Network: Physical; Right Network: Virtual
Test Bed Topology	Figure 1
Service under Test	Layer 3 Routing

Test Plan Version	1.0
Lab	EANTC, Germany

### NFVI

Vendor	Alcatel-Lucent
Product	Cloudband
Type	Centralized Data Center
Version	3.0

### NFVI Configuration

CPU	Intel(R) Xeon(R) CPU E5-2660 v2 @ 2.20GHz
Chipset	C600 family
Number of Cores	20
Cores per CPU	10
LL CACHE	25MB
QPI/Speed	2 x 8 GT/s
Memory	128GB
Network Interface Cards	Broadcom Corporation NetXtreme II BCM57810 10 Gigabit Ethernet
BIOS	HP-P75
Host OS	RHEL 7.0
Kernel version	3.10.0-123.20.1.el7.x86_64
Other	N/A
Hypervisor	KVM - QEMU 2.1.2
Virtual Infrastruct Mgmt	OpenStack Juno



## VNF-NFVI Interoperability Test Report Card: Alcatel-Lucent Cloudband

Test Campaign	Phase 1 / Campaign 1
Date	2015-11

### VNF

Vendor	Netnumber
Product	TITAN EDGE
Type	Centralized Signaling and Routing Control (CSRC)
Version	TITAN-7.6.2_v3.2.1

### Test Results

Critical Test Cases	Pass / Fail / Inconclusive / Not Applicable / Not Executed
Package Onboarding	PASS
VNF Instantiation	PASS
Forceful VNF Teardown	PASS
Graceful VNF Teardown	PASS
Package Deletion	PASS

### VNF Configuration

Number of Modules	1
Data plane virtual NICs	2
Instantiation Method	HEAT
VNF Management	CLI
License Management	Manual

### Test Configuration

Test Equipment	Ixia IxLoad, Ixia X16 w/ STXS-4, Ixia Virtual Chassis w/ Virtual Load Module
Connectivity	Left Network: Physical; Right Network: Virtual
Test Bed Topology	Figure 2
Service under Test	DNS Server

Test Plan Version	1.0
Lab	EANTC, Germany

### NFVI

Vendor	Alcatel-Lucent
Product	Cloudband
Type	Centralized Data Center
Version	3.0

### NFVI Configuration

CPU	Intel(R) Xeon(R) CPU E5-2660 v2 @ 2.20GHz
Chipset	C600 family
Number of Cores	20
Cores per CPU	10
LL CACHE	25MB
QPI/Speed	2 x 8 GT/s
Memory	128GB
Network Interface Cards	Broadcom Corporation NetXtreme II BCM57810 10 Gigabit Ethernet
BIOS	HP-P75
Host OS	RHEL 7.0
Kernel version	3.10.0-123.20.1.el7.x86_64
Other	N/A
Hypervisor	KVM - QEMU 2.1.2
Virtual Infrastruct Mgmt	OpenStack Juno

## VNF-NFVI Interoperability Test Report Card: Alcatel-Lucent Cloudband

Test Campaign	Phase 1 / Campaign 1
Date	2015-11

### VNF

Vendor	Netrounds
Product	Virtual Test Agent (VTA)
Type	End-to-end Service Verification
Version	2.11.3

### Test Results

Critical Test Cases	Pass / Fail / Inconclusive / Not Applicable / Not Executed
Package Onboarding	PASS
VNF Instantiation	PASS
Forceful VNF Teardown	PASS
Graceful VNF Teardown	PASS
Package Deletion	PASS

### VNF Configuration

Number of Modules	1
Data plane virtual NICs	2
Instantiation Method	Nova
VNF Management	Console, Web (Netrounds Control Center)
License Management	Web-Hosted (Netrounds Control Center)

### Test Configuration

Test Equipment	Ixia IxLoad, Ixia X16 w/ STXS-4, Ixia Virtual Chassis w/ Virtual Load Module
Connectivity	Left Network: Physical; Right Network: Virtual
Test Bed Topology	Figure 2
Service under Test	HTTP Client Probe

Test Plan Version	1.0
Lab	EANTC, Germany

### NFVI

Vendor	Alcatel-Lucent
Product	Cloudband
Type	Centralized Data Center
Version	3.0

### NFVI Configuration

CPU	Intel(R) Xeon(R) CPU E5-2660 v2 @ 2.20GHz
Chipset	C600 family
Number of Cores	20
Cores per CPU	10
LL CACHE	25MB
QPI/Speed	2 x 8 GT/s
Memory	128GB
Network Interface Cards	Broadcom Corporation NetXtreme II BCM57810 10 Gigabit Ethernet
BIOS	HP-P75
Host OS	RHEL 7.0
Kernel version	3.10.0-123.20.1.el7.x86_64
Other	N/A
Hypervisor	KVM - QEMU 2.1.2
Virtual Infrastruct Mgmt	OpenStack Juno

## VNF-NFVI Interoperability Test Report Card: Alcatel-Lucent Cloudband

Test Campaign	Phase 1 / Campaign 1
Date	2015-11

### VNF

Vendor	Juniper Networks
Product	vSRX
Type	Firewall
Version	JUNOS 12.1X47-D20.7

### Test Results

Critical Test Cases	Pass / Fail / Inconclusive / Not Applicable / Not Executed
Package Onboarding	PASS
VNF Instantiation	PASS
Forceful VNF Teardown	PASS
Graceful VNF Teardown	PASS
Package Deletion	PASS

### VNF Configuration

Number of Modules	1
Data plane virtual NICs	2
Instantiation Method	HEAT
VNF Management	CLI
License Management	Manual

### Test Configuration

Test Equipment	Ixia IxLoad, Ixia X16 w/ STXS-4, Ixia Virtual Chassis w/ Virtual Load Module
Connectivity	Left Network: Physical; Right Network: Virtual
Test Bed Topology	Figure 1
Service under Test	Layer 3 Routing

Test Plan Version	1.0
Lab	EANTC, Germany

### NFVI

Vendor	Alcatel-Lucent
Product	Cloudband
Type	Centralized Data Center
Version	3.0

### NFVI Configuration

CPU	Intel(R) Xeon(R) CPU E5-2660 v2 @ 2.20GHz
Chipset	C600 family
Number of Cores	20
Cores per CPU	10
LL CACHE	25MB
QPI/Speed	2 x 8 GT/s
Memory	128GB
Network Interface Cards	Broadcom Corporation NetXtreme II BCM57810 10 Gigabit Ethernet
BIOS	HP-P75
Host OS	RHEL 7.0
Kernel version	3.10.0-123.20.1.el7.x86_64
Other	N/A
Hypervisor	KVM - QEMU 2.1.2
Virtual Infrastruct Mgmt	OpenStack Juno

## VNF-NFVI Interoperability Test Report Card: Cisco NFVI

Test Campaign	Phase 1 / Campaign 1
Date	2015-11

### VNF

Vendor	IneoQuest
Product	IQDialogue ASM
Type	Streaming and CDN Monitoring Tool
Version	3.1.216.5

### Test Results

Critical Test Cases	Pass / Fail / Inconclusive / Not Applicable / Not Executed
Package Onboarding	PASS
VNF Instantiation	PASS
Forceful VNF Teardown	PASS
Graceful VNF Teardown	PASS
Package Deletion	PASS

### VNF Configuration

Number of Modules	1
Data plane virtual NICs	2
Instantiation Method	Nova + User Data (cloud-init)
VNF Management	CLI, Java Web GUI
License Management	Manual

### Test Configuration

Test Equipment	Ixia IxLoad, Ixia X16 w/ STXS-4, Ixia Virtual Chassis w/ Virtual Load Module
Connectivity	Left Network: Physical; Right Network: Virtual
Test Bed Topology	Figure 2
Service under Test	HTTP Streaming Client

Test Plan Version	1.0
Lab	EANTC, Germany

### NFVI

Vendor	Cisco Systems
Product	NFVI
Type	Centralized Data Center
Version	To be filled by Cisco

### NFVI Configuration

CPU	TBD
Chipset	TBD
Number of Cores	TBD
Cores per CPU	TBD
LL CACHE	TBD
QPI/Speed	8 GT/s
Memory	128 GB (Control) 256 GB (Compute)
Network Interface Cards	Cisco UCS VIC 1227
BIOS	C220M4.2.0.6a.0.051220151501
Host OS	Redhat RHEL 7.1
Kernel version	3.10.0-229.7.2.el7.x86_64
Other	TBD
Hypervisor	TBD
Virtual Infrastruct Mgmt	Openstack (Juno)



## VNF-NFVI Interoperability Test Report Card: Cisco NFVI

Test Campaign	Phase 1 / Campaign 1
Date	2015-11

### VNF

Vendor	Huawei
Product	VNE
Type	Router
Version	V100R002C00SPC013T

### Test Results

Critical Test Cases	Pass / Fail / Inconclusive / Not Applicable / Not Executed
Package Onboarding	PASS
VNF Instantiation	PASS
Forceful VNF Teardown	PASS
Graceful VNF Teardown	PASS
Package Deletion	PASS

### VNF Configuration

Number of Modules	1
Data plane virtual NICs	2
Instantiation Method	Nova
VNF Management	CLI
License Management	Manual

### Test Configuration

Test Equipment	Ixia IxLoad, Ixia X16 w/ STXS-4, Ixia Virtual Chassis w/ Virtual Load Module
Connectivity	Left Network: Physical; Right Network: Virtual
Test Bed Topology	Figure 1
Service under Test	Layer 3 Routing

Test Plan Version	1.0
Lab	EANTC, Germany

### NFVI

Vendor	Cisco Systems
Product	NFVI
Type	Centralized Data Center
Version	To be filled by Cisco

### NFVI Configuration

CPU	TBD
Chipset	TBD
Number of Cores	TBD
Cores per CPU	TBD
LL CACHE	TBD
QPI/Speed	8 GT/s
Memory	128 GB (Control) 256 GB (Compute)
Network Interface Cards	Cisco UCS VIC 1227
BIOS	C220M4.2.0.6a.0.051220151501
Host OS	Redhat RHEL 7.1
Kernel version	3.10.0-229.7.2.el7.x86_64
Other	TBD
Hypervisor	TBD
Virtual Infrastruct Mgmt	Openstack (Juno)

## VNF-NFVI Interoperability Test Report Card: Cisco NFVI

Test Campaign	Phase 1 / Campaign 1
Date	2015-11

### VNF

Vendor	Netrounds
Product	Virtual Test Agent (VTA)
Type	End-to-end Service Verification Probe
Version	2.11.3

### Test Results

Critical Test Cases	Pass / Fail / Inconclusive / Not Applicable / Not Executed
Package Onboarding	PASS
VNF Instantiation	PASS
Forceful VNF Teardown	PASS
Graceful VNF Teardown	PASS
Package Deletion	PASS

### VNF Configuration

Number of Modules	1
Data plane virtual NICs	2
Instantiation Method	Nova
VNF Management	Console, Web (Netrounds Control Center)
License Management	Web-Hosted (Netrounds Control Center)

### Test Configuration

Test Equipment	Ixia IxLoad, Ixia X16 w/ STXS-4, Ixia Virtual Chassis w/ Virtual Load Module
Connectivity	Left Network: Physical; Right Network: Virtual
Test Bed Topology	Figure 2
Service under Test	HTTP Client Probe

Test Plan Version	1.0
Lab	EANTC, Germany

### NFVI

Vendor	Cisco Systems
Product	NFVI
Type	Centralized Data Center
Version	To be filled by Cisco

### NFVI Configuration

CPU	TBD
Chipset	TBD
Number of Cores	TBD
Cores per CPU	TBD
LL CACHE	TBD
QPI/Speed	8 GT/s
Memory	128 GB (Control) 256 GB (Compute)
Network Interface Cards	Cisco UCS VIC 1227
BIOS	C220M4.2.0.6a.0.051220151501
Host OS	Redhat RHEL 7.1
Kernel version	3.10.0-229.7.2.el7.x86_64
Other	TBD
Hypervisor	TBD
Virtual Infrastruct Mgmt	Openstack (Juno)

## VNF-NFVI Interoperability Test Report Card: Cisco NFVI

Test Campaign	Phase 1 / Campaign 1
Date	2015-11

### VNF

Vendor	Metaswitch
Product	Perimeta
Type	Session Border Controller (SBC)
Version	Core Software V3.8.40_SU8_P01.00

### Test Results

Critical Test Cases	Pass / Fail / Inconclusive / Not Applicable / Not Executed
Package Onboarding	PASS
VNF Instantiation	PASS
Forceful VNF Teardown	PASS
Graceful VNF Teardown	PASS
Package Deletion	PASS

### VNF Configuration

Number of Modules	2
Data plane virtual NICs	2
Instantiation Method	Nova
VNF Management	CLI, SFTP
License Management	Manual Configuration + License Server

### Test Configuration

Test Equipment	Ixia IxLoad, Ixia X16 w/ STXS-4, Ixia Virtual Chassis w/ Virtual Load Module
Connectivity	Left Network: Physical; Right Network: Virtual
Test Bed Topology	Figure 1
Service under Test	SIP Signalling + Media

Test Plan Version	1.0
Lab	EANTC, Germany

### NFVI

Vendor	Cisco Systems
Product	NFVI
Type	Centralized Data Center
Version	To be filled by Cisco

### NFVI Configuration

CPU	TBD
Chipset	TBD
Number of Cores	TBD
Cores per CPU	TBD
LL CACHE	TBD
QPI/Speed	8 GT/s
Memory	128 GB (Control) 256 GB (Compute)
Network Interface Cards	Cisco UCS VIC 1227
BIOS	C220M4.2.0.6a.0.051220151501
Host OS	Redhat RHEL 7.1
Kernel version	3.10.0-229.7.2.el7.x86_64
Other	TBD
Hypervisor	TBD
Virtual Infrastruct Mgmt	Openstack (Juno)

## VNF-NFVI Interoperability Test Report Card: Cisco NFVI

Test Campaign	Phase 1 / Campaign 1
Date	2015-11

### VNF

Vendor	Alcatel-Lucent
Product	Virtualized Mobility Manager (VMM)
Type	Mobility Management Entity (Part of EPC)
Version	R33.62.01

### Test Results

Critical Test Cases	Pass / Fail / Inconclusive / Not Applicable / Not Executed
Package Onboarding	PASS
VNF Instantiation	PASS
Forceful VNF Teardown	PASS
Graceful VNF Teardown	PASS
Package Deletion	PASS

### VNF Configuration

Number of Modules	3
Data plane virtual NICs	2
Instantiation Method	Nova
VNF Management	CLI
License Management	Manual

### Test Configuration

Test Equipment	Ixia IxNetwork, Ixia X16 w/ STXS-4, Ixia Virtual Chassis w/ Virtual Load Module
Connectivity	Left Network: Physical; Right Network: Virtual
Test Bed Topology	Figure 1
Service under Test	Layer 3 Routing

Test Plan Version	1.0
Lab	EANTC, Germany

### NFVI

Vendor	Cisco Systems
Product	NFVI
Type	Centralized Data Center
Version	To be filled by Cisco

### NFVI Configuration

CPU	TBD
Chipset	TBD
Number of Cores	TBD
Cores per CPU	TBD
LL CACHE	TBD
QPI/Speed	8 GT/s
Memory	128 GB (Control) 256 GB (Compute)
Network Interface Cards	Cisco UCS VIC 1227
BIOS	C220M4.2.0.6a.0.051220151501
Host OS	Redhat RHEL 7.1
Kernel version	3.10.0-229.7.2.el7.x86_64
Other	TBD
Hypervisor	TBD
Virtual Infrastruct Mgmt	Openstack (Juno)



## VNF-NFVI Interoperability Test Report Card: Cisco NFVI

Test Campaign	Phase 1 / Campaign 1
Date	2015-11

### VNF

Vendor	Alcatel-Lucent
Product	Virtualized Service Router (VSR)
Type	Router
Version	13.0.R5

### Test Results

Critical Test Cases	Pass / Fail / Inconclusive / Not Applicable / Not Executed
Package Onboarding	PASS
VNF Instantiation	PASS
Forceful VNF Teardown	PASS
Graceful VNF Teardown	PASS
Package Deletion	PASS

### VNF Configuration

Number of Modules	1
Data plane virtual NICs	2
Instantiation Method	Nova
VNF Management	CLI
License Management	Manual

### Test Configuration

Test Equipment	Ixia IxNetwork, Ixia X16 w/ STXS-4, Ixia Virtual Chassis w/ Virtual Load Module
Connectivity	Left Network: Physical; Right Network: Virtual
Test Bed Topology	Figure 1
Service under Test	Layer 3 Routing

Test Plan Version	1.0
Lab	EANTC, Germany

### NFVI

Vendor	Cisco Systems
Product	NFVI
Type	Centralized Data Center
Version	To be filled by Cisco

### NFVI Configuration

CPU	TBD
Chipset	TBD
Number of Cores	TBD
Cores per CPU	TBD
LL CACHE	TBD
QPI/Speed	8 GT/s
Memory	128 GB (Control) 256 GB (Compute)
Network Interface Cards	Cisco UCS VIC 1227
BIOS	C220M4.2.0.6a.0.051220151501
Host OS	Redhat RHEL 7.1
Kernel version	3.10.0-229.7.2.el7.x86_64
Other	TBD
Hypervisor	TBD
Virtual Infrastruct Mgmt	Openstack (Juno)

## VNF-NFVI Interoperability Test Report Card: Cisco NFVI

Test Campaign	Phase 1 / Campaign 1
Date	2015-11

### VNF

Vendor	IneoQuest
Product	Expedus DVA IP
Type	Video Monitoring Tool
Version	4.95.00.010

### Test Results

Critical Test Cases	Pass / Fail / Inconclusive / Not Applicable / Not Executed
Package Onboarding	PASS
VNF Instantiation	PASS
Forceful VNF Teardown	PASS
Graceful VNF Teardown	PASS
Package Deletion	PASS

### VNF Configuration

Number of Modules	1
Data plane virtual NICs	2
Instantiation Method	Nova + User Data (cloud-init)
VNF Management	CLI, Web GUI
License Management	Manual

### Test Configuration

Test Equipment	Ixia IxLoad, Ixia X16 w/ STXS-4, Ixia Virtual Chassis w/ Virtual Load Module
Connectivity	Left Network: Physical; Right Network: Virtual
Test Bed Topology	Figure 2
Service under Test	Multicast Video Probe

Test Plan Version	1.0
Lab	EANTC, Germany

### NFVI

Vendor	Cisco Systems
Product	NFVI
Type	Centralized Data Center
Version	To be filled by Cisco

### NFVI Configuration

CPU	TBD
Chipset	TBD
Number of Cores	TBD
Cores per CPU	TBD
LL CACHE	TBD
QPI/Speed	8 GT/s
Memory	128 GB (Control) 256 GB (Compute)
Network Interface Cards	Cisco UCS VIC 1227
BIOS	C220M4.2.0.6a.0.051220151501
Host OS	Redhat RHEL 7.1
Kernel version	3.10.0-229.7.2.el7.x86_64
Other	TBD
Hypervisor	TBD
Virtual Infrastruct Mgmt	Openstack (Juno)

## VNF-NFVI Interoperability Test Report Card: Huawei FusionSphere

Test Campaign	Phase 1 / Campaign 1
Date	2015-11

### VNF

Vendor	Alcatel-Lucent
Product	Virtualized Mobility Manager (VMM)
Type	Mobility Management Entity (Part of EPC)
Version	R33.62.01

### Test Results

Critical Test Cases	Pass / Fail / Inconclusive / Not Applicable / Not Executed
Package Onboarding	PASS
VNF Instantiation	PASS
Forceful VNF Teardown	PASS
Graceful VNF Teardown	PASS
Package Deletion	PASS

### VNF Configuration

Number of Modules	3
Data plane virtual NICs	2
Instantiation Method	Nova
VNF Management	CLI
License Management	Manual

### Test Configuration

Test Equipment	Ixia IxNetwork, Ixia X16 w/ STXS-4, Ixia Virtual Chassis w/ Virtual Load Module
Connectivity	Left Network: Physical; Right Network: Virtual
Test Bed Topology	Figure 1
Service under Test	Layer 3 Routing

Test Plan Version	1.0
Lab	EANTC, Germany

### NFVI

Vendor	Huawei
Product	FusionSphere
Type	Centralized Data Center
Version	5.1

### NFVI Configuration

CPU	Intel® Xeon Haswell EP E5-2690v3
Chipset	Intel® C600
Number of Cores	24
Cores per CPU	12
LL CACHE	30M
QPI/Speed	9.6 GT/s
Memory	DDR4 RDIMM 16G*24
Network Interface Cards	Broadcom 5719
BIOS	V1.38
Host OS	SUSE Linux
Kernel version	3.9.9-0.8-default
Other	N/A
Hypervisor	KVM 3.6
Virtual Infrastruct Mgmt	OpenStack Juno

## VNF-NFVI Interoperability Test Report Card: Huawei FusionSphere

Test Campaign	Phase 1 / Campaign 1
Date	2015-11

### VNF

Vendor	Cisco Systems
Product	CSR1000v
Type	Router
Version	IOS XE 03.16.00.S, CSR1000V 15.5(3)S Release (fc6)

### Test Results

Critical Test Cases	Pass / Fail / Inconclusive / Not Applicable / Not Executed
Package Onboarding	PASS
VNF Instantiation	PASS
Forceful VNF Teardown	PASS
Graceful VNF Teardown	PASS
Package Deletion	PASS

### VNF Configuration

Number of Modules	1
Data plane virtual NICs	2
Instantiation Method	Nova + Config File
VNF Management	CLI
License Management	Manual

### Test Configuration

Test Equipment	Ixia IxNetwork, Ixia X16 w/ STXS-4, Ixia Virtual Chassis w/ Virtual Load Module
Connectivity	Left Network: Physical; Right Network: Virtual
Test Bed Topology	Figure 1
Service under Test	Layer 3 Routing

Test Plan Version	1.0
Lab	EANTC, Germany

### NFVI

Vendor	Huawei
Product	FusionSphere
Type	Centralized Data Center
Version	5.1

### NFVI Configuration

CPU	Intel® Xeon Haswell EP E5-2690v3
Chipset	Intel® C600
Number of Cores	24
Cores per CPU	12
LL CACHE	30M
QPI/Speed	9.6 GT/s
Memory	DDR4 RDIMM 16G*24
Network Interface Cards	Broadcom 5719
BIOS	V1.38
Host OS	SUSE Linux
Kernel version	3.9.9-0.8-default
Other	N/A
Hypervisor	KVM 3.6
Virtual Infrastruct Mgmt	OpenStack Juno

## VNF-NFVI Interoperability Test Report Card: Huawei FusionSphere

Test Campaign	Phase 1 / Campaign 1
Date	2015-11

### VNF

Vendor	Procera Networks
Product	Packetlogic Real-time Enforcement
Type	Traffic Inspection and Policy Enforcement
Version	15.1.5.12-1508130111

### Test Results

Critical Test Cases	Pass / Fail / Inconclusive / Not Applicable / Not Executed
Package Onboarding	PASS
VNF Instantiation	PASS
Forceful VNF Teardown	PASS
Graceful VNF Teardown	PASS
Package Deletion	PASS

### VNF Configuration

Number of Modules	1
Data plane virtual NICs	2
Instantiation Method	Nova
VNF Management	CLI, Client GUI
License Management	Manual

### Test Configuration

Test Equipment	Ixia IxLoad, Ixia X16 w/ STXS-4, Ixia Virtual Chassis w/ Virtual Load Module
Connectivity	Left Network: Virtual; Right Network: Virtual
Test Bed Topology	Figure 3
Service under Test	Inline HTTP Traffic Inspection

Test Plan Version	1.0
Lab	EANTC, Germany

### NFVI

Vendor	Huawei
Product	FusionSphere
Type	Centralized Data Center
Version	5.1

### NFVI Configuration

CPU	Intel® Xeon Haswell EP E5-2690v3
Chipset	Intel® C600
Number of Cores	24
Cores per CPU	12
LL CACHE	30M
QPI/Speed	9.6 GT/s
Memory	DDR4 RDIMM 16G*24
Network Interface Cards	Broadcom 5719
BIOS	V1.38
Host OS	SUSE Linux
Kernel version	3.9.9-0.8-default
Other	N/A
Hypervisor	KVM 3.6
Virtual Infrastruct Mgmt	OpenStack Juno

## VNF-NFVI Interoperability Test Report Card: Juniper Contrail

Test Campaign	Phase 1 / Campaign 1
Date	2015-11

### VNF

Vendor	Sonus
Product	Virtualized SBC Swe
Type	Session Border Controller (SBC)
Version	SBC Swe V04.02.03A627, Orchestrator V01.00.00A010

### Test Results

Critical Test Cases	Pass / Fail / Inconclusive / Not Applicable / Not Executed
Package Onboarding	PASS
VNF Instantiation	PASS
Forceful VNF Teardown	PASS
Graceful VNF Teardown	PASS
Package Deletion	PASS

### VNF Configuration

Number of Modules	2
Data plane virtual NICs	2
Instantiation Method	VNF Manager (SonusVNF Manager) + yaml VNFD templates
VNF Management	VNF Manager (SonusVNF Manager)
License Management	Manual

### Test Configuration

Test Equipment	Ixia IxLoad, Ixia X16 w/ STXS-4, Ixia Virtual Chassis w/ Virtual Load Module
Connectivity	Left Network: Physical; Right Network: Virtual
Test Bed Topology	Figure 1
Service under Test	SIP Signalling + Media

Test Plan Version	1.0
Lab	EANTC, Germany

### NFVI

Vendor	Juniper
Product	Contrail
Type	Centralized Data Center
Version	2.2

### NFVI Configuration

CPU	Intel® Xeon CPU E5-2670 v3, 2.3 GHz
Chipset	TBD
Number of Cores	48
Cores per CPU	24
LL CACHE	30720 KB
QPI/Speed	TBD
Memory	128GB
Network Interface Cards	Intel 82599ES, Intel X540-AT2
BIOS	TBD
Host OS	Ubuntu 14.04.2 LTS
Kernel version	3.16.0-30-generic
Other	N/A
Hypervisor	KVM/Qemu 2.0
Virtual Infrastruct Mgmt	OpenStack Icehouse



## VNF-NFVI Interoperability Test Report Card: Juniper Contrail

Test Campaign	Phase 1 / Campaign 1
Date	2015-11

### VNF

Vendor	Alcatel-Lucent
Product	Virtualized Mobile Gateway (VMG)
Type	EPC Mobile GW
Version	TIMOS-MG-C-7.0.R7

### Test Results

Critical Test Cases	Pass / Fail / Inconclusive / Not Applicable / Not Executed
Package Onboarding	PASS
VNF Instantiation	PASS
Forceful VNF Teardown	PASS
Graceful VNF Teardown	PASS
Package Deletion	PASS

### VNF Configuration

Number of Modules	3
Data plane virtual NICs	2
Instantiation Method	Nova + User Data
VNF Management	CLI
License Management	Manual

### Test Configuration

Test Equipment	Ixia IxNetwork, Ixia X16 w/ STXS-4, Ixia Virtual Chassis w/ Virtual Load Module
Connectivity	Left Network: Physical; Right Network: Virtual
Test Bed Topology	Figure 1
Service under Test	Layer 3 Routing

Test Plan Version	1.0
Lab	EANTC, Germany

### NFVI

Vendor	Juniper
Product	Contrail
Type	Centralized Data Center
Version	2.2

### NFVI Configuration

CPU	Intel® Xeon CPU E5-2670 v3, 2.3 GHz
Chipset	TBD
Number of Cores	48
Cores per CPU	24
LL CACHE	30720 KB
QPI/Speed	TBD
Memory	128GB
Network Interface Cards	Intel 82599ES, Intel X540-AT2
BIOS	TBD
Host OS	Ubuntu 14.04.2 LTS
Kernel version	3.16.0-30-generic
Other	N/A
Hypervisor	KVM/Qemu 2.0
Virtual Infrastruct Mgmt	OpenStack Icehouse

## VNF-NFVI Interoperability Test Report Card: Juniper Contrail

Test Campaign	Phase 1 / Campaign 1
Date	2015-11

### VNF

Vendor	Alcatel-Lucent
Product	Virtualized Service Router (VSR)
Type	Router
Version	13.0.R5

### Test Results

Critical Test Cases	Pass / Fail / Inconclusive / Not Applicable / Not Executed
Package Onboarding	PASS
VNF Instantiation	PASS
Forceful VNF Teardown	PASS
Graceful VNF Teardown	PASS
Package Deletion	PASS

### VNF Configuration

Number of Modules	1
Data plane virtual NICs	2
Instantiation Method	Nova
VNF Management	CLI
License Management	Manual

### Test Configuration

Test Equipment	Ixia IxNetwork, Ixia X16 w/ STXS-4, Ixia Virtual Chassis w/ Virtual Load Module
Connectivity	Left Network: Virtual; Right Network: Virtual
Test Bed Topology	Figure 1
Service under Test	Layer 3 Routing

Test Plan Version	1.0
Lab	EANTC, Germany

### NFVI

Vendor	Juniper
Product	Contrail
Type	Centralized Data Center
Version	2.2

### NFVI Configuration

CPU	Intel® Xeon CPU E5-2670 v3, 2.3 GHz
Chipset	TBD
Number of Cores	48
Cores per CPU	24
LL CACHE	30720 KB
QPI/Speed	TBD
Memory	128GB
Network Interface Cards	Intel 82599ES, Intel X540-AT2
BIOS	TBD
Host OS	Ubuntu 14.04.2 LTS
Kernel version	3.16.0-30-generic
Other	N/A
Hypervisor	KVM/Qemu 2.0
Virtual Infrastruct Mgmt	OpenStack Icehouse

## VNF-NFVI Interoperability Test Report Card: Juniper Contrail

Test Campaign	Phase 1 / Campaign 1
Date	2015-11

### VNF

Vendor	Hitachi
Product	Mobility Management Element (MME)
Type	Virtual Mobile Core
Version	8.2.36

### Test Results

Critical Test Cases	Pass / Fail / Inconclusive / Not Applicable / Not Executed
Package Onboarding	PASS
VNF Instantiation	PASS
Forceful VNF Teardown	PASS
Graceful VNF Teardown	PASS
Package Deletion	PASS

### VNF Configuration

Number of Modules	5
Data plane virtual NICs	2
Instantiation Method	Nova + User Data
VNF Management	CLI
License Management	Manual

### Test Configuration

Test Equipment	Ixia IxNetwork, Ixia X16 w/ STXS-4, Ixia Virtual Chassis w/ Virtual Load Module
Connectivity	Left Network: Physical; Right Network: Virtual
Test Bed Topology	Figure 1
Service under Test	Diameter Interface Availability

Test Plan Version	1.0
Lab	EANTC, Germany

### NFVI

Vendor	Juniper
Product	Contrail
Type	Centralized Data Center
Version	2.2

### NFVI Configuration

CPU	Intel® Xeon CPU E5-2670 v3, 2.3 GHz
Chipset	TBD
Number of Cores	48
Cores per CPU	24
LL CACHE	30720 KB
QPI/Speed	TBD
Memory	128GB
Network Interface Cards	Intel 82599ES, Intel X540-AT2
BIOS	TBD
Host OS	Ubuntu 14.04.2 LTS
Kernel version	3.16.0-30-generic
Other	N/A
Hypervisor	KVM/Qemu 2.0
Virtual Infrastruct Mgmt	OpenStack Icehouse

## VNF-NFVI Interoperability Test Report Card: Juniper Contrail

Test Campaign	Phase 1 / Campaign 1
Date	2015-11

### VNF

Vendor	Cobham Wireless
Product	TeraVM
Type	Virtual Test Agent
Version	TVM v3.0.8-169, TVM-C v11.4-613

### Test Results

Critical Test Cases	Pass / Fail / Inconclusive / Not Applicable / Not Executed
Package Onboarding	PASS
VNF Instantiation	PASS
Forceful VNF Teardown	PASS
Graceful VNF Teardown	PASS
Package Deletion	PASS

### VNF Configuration

Number of Modules	2
Data plane virtual NICs	2
Instantiation Method	Nova
VNF Management	CLI, Java Web GUI
License Management	Manual

### Test Configuration

Test Equipment	Ixia IxNetwork, Ixia X16 w/ STXS-4, Ixia Virtual Chassis w/ Virtual Load Module
Connectivity	Left Network: Virtual; Right Network: Virtual
Test Bed Topology	Figure 2
Service under Test	HTTP Client Test

Test Plan Version	1.0
Lab	EANTC, Germany

### NFVI

Vendor	Juniper
Product	Contrail
Type	Centralized Data Center
Version	2.2

### NFVI Configuration

CPU	Intel® Xeon CPU E5-2670 v3, 2.3 GHz
Chipset	TBD
Number of Cores	48
Cores per CPU	24
LL CACHE	30720 KB
QPI/Speed	TBD
Memory	128GB
Network Interface Cards	Intel 82599ES, Intel X540-AT2
BIOS	TBD
Host OS	Ubuntu 14.04.2 LTS
Kernel version	3.16.0-30-generic
Other	N/A
Hypervisor	KVM/Qemu 2.0
Virtual Infrastruct Mgmt	OpenStack Icehouse

## VNF-NFVI Interoperability Test Report Card: Juniper Contrail

Test Campaign	Phase 1 / Campaign 1
Date	2015-11

### VNF

Vendor	Metaswitch
Product	Perimeta
Type	Session Border Controller (SBC)
Version	Core Software V3.8.40_SU8_P01.00

### Test Results

Critical Test Cases	Pass / Fail / Inconclusive / Not Applicable / Not Executed
Package Onboarding	PASS
VNF Instantiation	PASS
Forceful VNF Teardown	PASS
Graceful VNF Teardown	PASS
Package Deletion	PASS

### VNF Configuration

Number of Modules	2
Data plane virtual NICs	2
Instantiation Method	Nova
VNF Management	CLI, SFTP
License Management	Manual Configuration + License Server

### Test Configuration

Test Equipment	Ixia IxNetwork, Ixia X16 w/ STXS-4, Ixia Virtual Chassis w/ Virtual Load Module
Connectivity	Left Network: Physical; Right Network: Virtual
Test Bed Topology	Figure 1
Service under Test	SIP Signalling + Media

Test Plan Version	1.0
Lab	EANTC, Germany

### NFVI

Vendor	Juniper
Product	Contrail
Type	Centralized Data Center
Version	2.2

### NFVI Configuration

CPU	Intel® Xeon CPU E5-2670 v3, 2.3 GHz
Chipset	TBD
Number of Cores	48
Cores per CPU	24
LL CACHE	30720 KB
QPI/Speed	TBD
Memory	128GB
Network Interface Cards	Intel 82599ES, Intel X540-AT2
BIOS	TBD
Host OS	Ubuntu 14.04.2 LTS
Kernel version	3.16.0-30-generic
Other	N/A
Hypervisor	KVM/Qemu 2.0
Virtual Infrastruct Mgmt	OpenStack Icehouse

## VNF-NFVI Interoperability Test Report Card: Juniper Contrail

Test Campaign	Phase 1 / Campaign 1
Date	2015-11

### VNF

Vendor	Netnumber
Product	TITAN EDGE
Type	Centralized Signaling and Routing Control (CSRC)
Version	TITAN-7.6.2_v3.2.1

### Test Results

Critical Test Cases	Pass / Fail / Inconclusive / Not Applicable / Not Executed
Package Onboarding	PASS
VNF Instantiation	PASS
Forceful VNF Teardown	PASS
Graceful VNF Teardown	PASS
Package Deletion	PASS

### VNF Configuration

Number of Modules	1
Data plane virtual NICs	2
Instantiation Method	HEAT
VNF Management	CLI
License Management	Manual

### Test Configuration

Test Equipment	Ixia IxLoad, Ixia X16 w/ STXS-4, Ixia Virtual Chassis w/ Virtual Load Module
Connectivity	Left Network: Physical; Right Network: Virtual
Test Bed Topology	Figure 2
Service under Test	DNS Server

Test Plan Version	1.0
Lab	EANTC, Germany

### NFVI

Vendor	Juniper
Product	Contrail
Type	Centralized Data Center
Version	2.2

### NFVI Configuration

CPU	Intel® Xeon CPU E5-2670 v3, 2.3 GHz
Chipset	TBD
Number of Cores	48
Cores per CPU	24
LL CACHE	30720 KB
QPI/Speed	TBD
Memory	128GB
Network Interface Cards	Intel 82599ES, Intel X540-AT2
BIOS	TBD
Host OS	Ubuntu 14.04.2 LTS
Kernel version	3.16.0-30-generic
Other	N/A
Hypervisor	KVM/Qemu 2.0
Virtual Infrastruct Mgmt	OpenStack Icehouse