

Participating in CloudNFV: Let Me Count the Ways!



CloudNFV is an open project, and while we have resource constraints that will prevent us from answering every question and supporting every attempt to work with us, we do intend to advance our platform by integrating with others.

From the time our work became known, CloudNFV has generated a lot of press and a lot of interest, and many people have asked to contribute to the project in some way. We want that too. CloudNFV is an open activity, meaning that it's our goal to allow for open, multi-vendor, integration with our framework at every reasonable point of interfacing. Obviously we intend to support/expose the interface points mandated by the NFV ISG, but since that body has not approved a final specification and may not do so for some time, we are working in "pre-standards" territory here! We're also just ironing out our own procedures to work with others, and this document is our first step.

To start with, it's important that before you make any requests for information from us, you explore the website and read all of the public material available. There is a white paper that describes CloudNFV, and that document is particularly important in establishing the technical framework of CloudNFV, and framing the various activities that could be cooperative. We are also committed to producing webinars, Google+ Hangouts on Air, and other public tutorial material just as fast as you can. All of this should help you come up to speed and decide whether, and how, you'd like to work with CloudNFV.

When you request to join the project, you're automatically enrolled as a **Supporter**, which means that you are free to access the main area of the LinkedIn group and also to receive open project documents. Please realize that some of the contributed elements in CloudNFV are proprietary, so don't expect our members to reveal

- We're asking that everyone who wants to integrate with us read all of our public documentation and then work with us according to some specific guidelines.
- We're defining multiple models of participation to try to match our resources to partner capabilities, and by so doing advance both CloudNFV's functionality and our contributions back into the NFV ISG and the TMF
- Any significant work with us will require you to submit an integration proposal, and we'll prioritize our support of these proposals based on how much we think they advance the cause of NFV and how readily we can resource them.
- Anyone is free to enter into a bilateral commercial arrangement with any company involved in CloudNFV to advance their project faster

trade secrets with you! You're free to participate in any open tutorial activity we offer and **from time to time we'll poll Supporters to draw resources for further work**, but unless you're selected for that we can't provide you support or education beyond our documentation; we have only limited resources and for now at least those have to be directed to the next level of participation.

...which is **Integration Partner**. In order to be an Integration Partner you'll have to plan to integrate or develop an element of NFV. Integration Partners will have to be companies who are members of the NFV ISG. This category is for those who plan to develop or integrate with CloudNFV at some defined interface point. We recognize three broad classes of integration projects, described later in this document.

While our Integration Partners are normally going to be vendors, we also accept more broadly defined projects from standards groups, and in fact we would accept a project proposal from an SDO or sanctioned forum in any area where we recognize a relationship could be valuable. In these cases the purpose of the early discussions would be to define a range of activities that would create specific integration activity, or simply to formalize liaisons. Our only requirement here is that if none of our CloudNFV founding members are members of the body, we'll need to have the body admit someone from our group as a member without paying any special fees (remember those resource constraints!) We'll post notices of these organizational liaisons as they're formalized and define the way contributions under the liaisons would work, so members can take advantage of the process.

Because of those pesky resource limitations, we must require that all who want to join as an Integration Partner provide us with a plan, which must specify what you intend to contribute and how you plan to work with us. We describe the general integration requirements below, and you should reference these and classify your intentions into the appropriate functional areas. We will need to know when you expect to be ready to test integration, how you plan to link with our test framework (in person, online) and what resources you will need from us. Remember, the more you need us to do the more difficult it will be to get your plan approved! Once we've reviewed your plan, we'll indicate to you when we'll be ready

to support your activity and what you'll need to do ahead of time to be sure you're ready.

We recognize the desire of individuals to contribute to CloudNFV, but we simply can't address all of the opportunities that this might present unless we lay out a basic program. What we can support are individuals who plan to be **Open-Source Virtual Function Developers**, meaning that you want to develop open-source components of CloudNFV or open-source network functions for cloud deployment. We can only accept code contributed under an open-source license that permits free integration with commercial components. You must specify your license and provide a copy in order to join, and you'll only be able to test your software in "Virtual Function Plugfests" that we'll provide periodically. In order to offer software in this program, you'll need to comply with the integration rules for virtual functions described in the next section.

Integrating with CloudNFV: Technical Guidelines

CloudNFV has three basic layers—virtual functions, management and orchestration, and infrastructure/resources. We establish a general interface framework based on OpenStack for deployment, but to integrate with our order/management processes you'll need to be able to integrate with our Active Resource and Active Contract data models. This is done using what we call "Visualizers" which are small software objects that provide information derivation and presentation rules. You can populate our data models, view them, or both using the right Visualizers. As the ETSI ISG develops we'll define Visualizers for any standard interfaces they specify and you can adopt these as available. We also will support Visualizer-based interfaces to other SDOs that we've established liaison processes with. If there are none yet defined, or if you need something different, you'll have to develop your own Visualizers to make a connection to our data model. You should assume that **any integration project except (possibly) virtual function development will require that you develop (or fund) at least one Visualizer.**

- There are three layers to CloudNFV—virtual function logic, management and orchestration logic, and infrastructure. Each of these will have specific rules for integration.
- The easiest thing for a partner to do is to develop a virtual function for us to deploy and manage. If you can write software for OpenStack, you can probably write for us.
- Almost all other integration projects will involve our partner creating one or more software objects we call "Virtualizers" that will provide the connection between CloudNFV and their own contribution.

How much more there might be depends on your integration goals. The ISG framework for NFV (as of now) defines three functional layers; services/applications (the “network functions”), orchestration and management (the “virtualization”), and infrastructure (what things are virtualized onto). Our open integration plans reflect this division, and we recommend that interested parties frame their participation into this framework as well.

At a high level, here’s how participation works for each layer:

- If you are contributing network functions to be deployed by us, you will need to test these functions first in an OpenStack cloud environment to be sure they run, and to establish the Nova and Neutron (formerly Quantum) framework needed to deploy and connect them. You then describe them in our data model based on our templates/examples and you’re ready to test with us. Integration partners with pre-approved test plans can contribute proprietary logic to be integrated this way, but individuals must supply open-source applications only so we avoid intellectual-property or copyright issues with participants who are employed by a company but representing themselves in CloudNFV.
- If you are contributing a component of infrastructure, like a server, network, or software platform you’ll again need first to integrate this with OpenStack and test OpenStack on your platform. Server/software providers will have to comply with Nova APIs and for network equipment you’ll need to conform to a Neutron (formerly Neutron) plugin that can also control Open vSwitch. When you can run under OpenStack, you are ready to go. To integrate your resource telemetry with Active Resource, you’ll need to develop a Visualizer element to our specifications that will report your status items correctly.
- The NFV ISG defines orchestration as a monolithic function but we have some internal points where you’ll connect to data and resource information. This level of integration is the most complicated so be prepared to expend some effort here!

Need a little more guidance on one or more of these options? Read on!

Page 4

Preparing Network Functions for Virtualization

Services/applications integration is a matter of preparing service logic or cloud applications to be deployed and operationalized by CloudNFV. To do that, partners would have to create machine images compatible with our test bed: Linux, KVM, and OpenStack, and if you plan to pass high-volume traffic you'll need to include 6WIND's 6WINDGate optimization software. For each image we would then need a load description (VNFD) that described the hosting requirements and the interfaces to be connected. We will supply the template and examples for this. If any special parameterization of components is needed upon instantiation, the partner would have to provide a Visualizer element for Active Resource/Active Contract. We will provide specifications for that as well. Finally, any special network management interface will also require a Visualizer element, and specifications for that will also be made available.

If you are deploying multiple cooperative components, you'll need a VNFD for each component and also a **Package** descriptor that tells how to link your components using Neutron Models. We have a template and examples for this process as well.

Integrating Your Own Orchestration/Management

Participation in the orchestration and management layer requires that the partner support our O&M Visualizer interface, for which specifications and examples will be provided, to obtain the service and resource information needed from Active Contract and Active Resource. A partner who wants to do a complete orchestration/management integration would have to accept our Active Contract data for a deployment order, optimize and deploy the virtual functions, connect them, and then provide the management data connections to both Active Contract and Active Resource. We want to be honest and say that this is a seriously

Page 5

non-trivial task for both you and for us, and so we are not going to approve this kind of integration request lightly. You'll have to offer some dramatic advantage and be fully resourced. We recommend that if your goal is to implement the ISG NFV specifications in this area, you consider **federating** with us instead; see below for more details.

Offering Infrastructure Elements for Integration

To participate in the infrastructure layer, requirements are divided into compute platform and network platform.

Participating in the compute platform requires that the partner be able to run the load image associated with CloudNFV, which is based on Linux, KVM, 6WIND data path optimization, and Qosmos traffic telemetry. This platform will properly integrate with Active Resource and with our orchestration. To participate in the network platform requires that the partner support Neutron and integrate with Overture's Neutron framework.

There's no real functional difference between one version of the compute side of NFV infrastructure and another, so we're interested in this form of integration primarily to extend the footprint of CloudNFV to more labs. Be sure your proposal contains a commitment to do this or you'll have a lower priority.

Network integration with CloudNFV would be accomplished through a variety of means:

- An OpenStack plugin (Neutron) can be developed that will create the necessary Models. Because there can be only one plugin, you'll have to either support our baseline server/network configuration with yours, or propose a deployment in a separate lab and federate with us.
- An application that controls your network equipment through the one of the standard Neutron plugins (same restrictions as above).

- An application that links us to Open Daylight's northbound APIs and can control your equipment.
- An "Infrastructure Service" that defines its own service model structures and the logic needed to instantiate them on systems, networks, and even host applications on blades in network devices.

Overture has provided a Orchestration framework that will control their own switches outside of Neutron, so as long as your plugins can support your own switch and OVS they should be compatible with our framework.

Infrastructure Services are the most flexible way of integrating infrastructure elements, especially network and SDN technology. We have a special slide deck on this option, and if you want to review it for integration you should contact us for the details.

In both the network and server partnerships, you'll need to develop a Visualizer to populate Active Resource with your status and also expose any special variables you'll develop and use. If your devices need parameterization you will need a management Visualizer to provide for that and manage devices during state changes in the service/product lifecycle.

Special Consideration for Open-Source Contributions to CloudNFV

Most open-source contributions to CloudNFV would likely fit into one of the following:

- A contribution of network functionality, something that is deployed by CloudNFV and that implements some network service feature.

- We can't get involved in intellectual property issues, so we require that anyone contributing virtual function logic do so in open-source form.
- If you want to contribute proprietary virtual functions you'll have to do so as a company, provide a license to us for deployment, and comply with our Integration Partner program rules

- A contribution of management functionality, usually an interface between CloudNFV and a management system or application.

An ideal network function is one that is developed as a RESTful application and that can be replicated as needed by CloudNFV to improve availability or to improve performance. This means that you'll need to manage state carefully in your development, and that you'll need to decide how your network function divides work among multiple instances.

Federating with CloudNFV

“Federation” is the term we use to describe CloudNFV relationships that cross (simulated) provider boundary points. In effect, each CloudNFV installation is a CloudNFV Federation Domain, and can be connected to other domains in various ways.

The simplest federation model is the model of **NFVI-level federation**, corresponding roughly to a federation between IaaS cloud providers. NFVI federation involves three steps:

- Provide a network-to-network interface (NNI) to the partner network to provide addressability of the assets to be deployed there. This can be done using conventional network inter-carrier practices.
- Provide a **Federation Gateway** between Active Resource in the “home network” (CloudNFV) and the resource pool and telemetry in the partner network. This is a form of Visualizer that will include a set of policies to govern what can be collected from the partner (or populated back there) and how it will be integrated with Active Resource (which is our specification).
- Define the partner resources as objects in Active Resource, which allows us to collect data through the Federation Gateway and integrate it with the rest of CloudNFV.

- Do you have your own NFV implementation, even a partial one, that can accept orders for virtual functions and deploy them? We really want to work with you as a “federated partner” in CloudNFV.
- There are a variety of federation models we support, and we're sure you can find at least one that you can work with.
- Our federation approach offers controlled cross-visibility in management so assets in partner networks can be protected even while shared.

In this level of federation, CloudNFV will treat the partner as simply part of the resource pool, and will deploy and manage things there via the Federation Gateway. No NFV logic need be deployed in the partner network; all VNFs are from the home network and orchestration and management takes place only in the home network.

The second level of federation is where the partner network provides VNFs for deployment in their network, by the home network's orchestration and management. Instead of home-network VNFs, the VNFs must come from the partner network and so they must be available in OpenStack Glance there. This is a form of "PaaS" cloud federation, and all of the above federation requirements must be met, and in addition:

- The home CloudNFV must have a VNFD for each of the partner virtual functions in Glance there. These can be composed into normal service hierarchies, but the VNFD will be in the namespace of the partner.
- The Federation Gateway must include an interface between Active Contract and the partner Glance VNF repository. This is used to pull the VNF image back into local orchestration; it's essentially a framework for selective "Glance Federation".

In both the above forms of federation, the only copy of CloudNFV orchestration/management running is in the home network. Partner resources are gateway-integrated as provided by the federation agreement. The final form of federation is one where either CloudNFV or another suitable framework for deployment/management is running in the partner system as well. The following models are recognized:

- CloudNFV is running in one place but federation partners are running OpenStack in each of their infrastructures.
- CloudNFV is running both in the home network and the partner network.
- CloudNFV is running in the home network and another orchestration framework that's NFV ISG compatible is running in the partner network.

In all of these cases, the difference with this most complex model of federation is that some of the deployment functionality is resident in the partner network—anywhere from OpenStack to a complete implementation of NFV orchestration. The support for all these models requires an enhanced version of the Federation Connector to support the exchange between the partners at an appropriate functional level. For example, if federation involves the partner contributing a complete customer-facing service component to be composed into a retail service, then the service order for that must pass through the connector and the connections created by order fulfillment have to be visible through the NNI between the parties. The number of possible situations here is large enough that providing guidance in advance for all of them isn't practical. Please make a request for this level of federation in an application for Integration Partner status and we'll work through things when we process the request.

Closing Comments

If you expected us to close without another comment on those issues of resources, think again! CloudNFV has generated a lot of interest, but remember that we are not a startup, we're not funded by some giant forum process. We all have day jobs that have to pay the bills, and that means that we can't possibly answer all the questions and support all the activities that have been suggested to us already—even though we've not been open for outside participation. We will be managing resources carefully in order to advance the work of the ETSI ISG, improve the functionality and scalability of CloudNFV, and demonstrate our commitment to an open community of participants.

Integration partnerships will take time and resources, and so we will have to schedule these to avoid being overwhelmed. We ask that anyone with a specific idea present their plan to us as soon as possible after we've opened up, and that the plan be articulated to match the technical description of CloudNFV contained in our white paper and the integration opportunities that have been outlined earlier in this paper. You will also need to indicate when you're prepared to start, what you are willing to commit in terms of resource, and what you think you'll need from us. If you have specific carrier sponsorship, please list the carriers/contacts!

Based on the value of your integration, the timing of the receipt of your plan, the sponsorship, and the resourcing, we'll either propose an approximate timetable for work or tell you we don't think we can support you. In the latter case we'll explain our reasoning, and you're free to fix things up and reapply. You will go back into the queue at that point, so try to get things as right as possible the first time you submit a plan.

Please note that you can enter into a bilateral agreement to work with/integrate with any of the CloudNFV vendors, and in such an agreement you can negotiate terms for education, support, etc. This is what you should consider if you are unable to integrate with our framework based on our reference material. Be sure that such agreements cover the scope of issues you expect, and review the roles of each CloudNFV vendor to be sure you're working with the most appropriate one for the type of contribution you intend to make.

Our goal is to create a global test bed for the ETSI NFV work, an implementation that can if operators choose be run in any and every lab, everywhere on the planet, and by doing that demonstrate the value and range of NFV. We also want to show how virtualization can be integrated into, and even **become** the operations support of the future. We want to do all of this with as open a community as we can create and support, and we thank you for considering participation.