

NFVRG IETF98 Chicago

Wednesday April 5, 2017 Morning Session I 9:00-11:30am – Zurich D

Thursday April 6, 2017 Afternoon Session II 3:20-5:20pm – Vevey I

Chairs:

Diego Lopez
Ramki Krishnan

Secretary:

Sarah Banks

Note Takers: Sarah Banks, Jon Hudson (thank you!)

NFVRG session 1 Wednesday 9:00-11:30am

Welcome and administrative matters

Presenter: Chairs

Slides: <https://www.ietf.org/proceedings/98/slides/slides-98-nfvrg-00-welcome-and-admin-matters-02.pdf>

- Mailing list: <https://www.ietf.org/mailman/listinfo/nfvrg>
- Web site: <http://trac.tools.ietf.org/group/irtf/trac/wiki/nfvrg>
- Proceedings: <https://www.ietf.org/proceedings/95/nfvrg.html>
- Sarah is note taker. See recording for jabber scribe (thank you, jabber scribe!)
- Use the NFVRG list for announcements, limited exclusively to research-related events

Microservices on the Edge: The Infrastructure Impact

Presenter: Ramki Krishnan

Slides: <https://www.ietf.org/proceedings/98/slides/slides-98-nfvrg-sessa-01-microservices-on-the-edge-the-infrastructure-impact-04.pdf>

- Background on microservices
- Focused on infrastructure impact, start with enterprise microservices, move on to the edge
- Walked through an example
 - o One key takeaway: HW acceleration key for deterministic performance, particularly for latency sensitive applications
 - o Used encryption/decryption as an example
- Walked through another example – the travel booking example as a set of microservices
- Q: Georgios K (Huawei): Is the context of the microservices in the application area, or in the VNA?
- A: Ramki: Good question, we'll get to that further in the presentation
- Q: Georgios K (Huawei): regarding the API; you'd like to have the microservices that'll be developed by vendors?
- A: The microservices API is a flexible loose term. It's a broad loose flexible definition, and context driven.
- Use Case Summary
 - o Digging in on the IoT use case
 - o Getting data is great, but if you're unable to analyze in real time, it's useless
 - o Your implementation or architecture should be flexible to allow the addition of new microservices on the fly
 - o Q from Nokia: Where's the notion of microservices from your diagram?

- A: The separation here is about internal jobs, versus input from the customer. Alerting can be viewed as a mechanism; something that is satisfying a business need is a microservice. The reason for the split is that when you want to add a new capability, you don't have to rearchitect, you can add on the fly
- Q from Nokia: Maybe I'm missing something, I expected to hear about VMs, micro kernels, etc, but instead I'm hearing about business requirements
- A: There are different business needs, there's no one size fits all solution; your implementation could be very different based on the problem you're trying to solve.
- Q from Alan?: Are there lessons we should learn about the capabilities we roll out? What the interfaces should be, what the APIs should be?
- A: This is not describing the full solution, it's a new area of focus, and the intent is to trigger discussions around this.
- Q Kyle Lerosse, Sandvine: network functions use a lot of bandwidth. By decoupling them, would we just bottleneck at a different spot?
- A: Good question, the answer is based on the implementation; it depends on the use case, understand the use case and deployment, and develop a strategy based on that
- Containers – FCAPS framework, how they can be tied to this
- Performance management, security, and hardware security discussed
- Practical Deployment Discussion
 - NFV starting out as a SaaS
 - Sometimes have to run 3rd party apps
 - Viable for a predominantly containerized deployment (as long as there are no performance issues)
 - VNF/vendors, and operators, need to step up and help move the solution forward
 - Asking for participation within NFVRG
 - Q from Kalyani, Verizon: when services are put together by SPs, we want to know what APIs are needed to put them together; the second part of your presentation is where we need more work to be done
 - A: Yes; the first part of the presentation was to give some background, I agree

Distributed NFV in Scattered Premises

Presenter: Liang Geng, China Mobile

draft-geng-nfvrg-distributed-nfv

Slides: <https://www.ietf.org/proceedings/98/slides/slides-98-nfvrg-sessa-03-draft-geng-nfvrg-distributed-nfv-00.pdf>

- Background
 - High demand on flexibility, security, and reliable service quality guarantee, for the future of new services (industrial Internet, IoT, etc)
 - As a SP, to meet these demands, we have to provide dedicated end to end resources
 - How can NFV help?
 - Decoupling helps (HW and SW can evolve independently)
 - Isolated instances provide better security
 - Dedicated resources provide the performance guarantee
 - Centralized or Distributed?
 - Opinion: NFV not built for Centralized use cases
 - NFV tech is most likely to be tested and commercialized in centralized use cases
 - We've been doing NFV trials for core network, for 2-3 years
 - In distributed scenario, using the home gateway scenario, it used to be a thin box, now it's an intelligent device, like OSGI, install plugins on the device.. but that's not NFV, it's a plugin to add more services
 - Q from Diego: What do you mean, ARM?

- A: For the distributed devices, they're all mostly based on ARM. We're doing a test in our lab where we found, there should not be a big compatibility program, but if you're using an n86 control node, but the compute node runs with ARM, there are issues
 - We have this concept – centralize what you can, distribute what you must
 - We're starting to see a lot of musts coming
- Examples of Distributed NFV PoPs
 - Customer premise devices
 - Scattered transport network elements
 - Needs discussion
 - Thinking ahead though, if you're a tenant of network slicing (5G) then you might want a say in this scheduling and resources
 - For our traditional network items, there is a need to provide multiple domain provisioning
- Use Cases of Distributed NFV
- Other Issues
- Q Linda Dunbar, Huawei: what do you want those distributed functions, what kind of issues do you see today? What action(s) do you need?
- A: We have seen several deployments of this type of scenarios in enterprise (1 accelerator implemented as a VNF in CPE). The residential CPE, I would like to see them as an edge computing node. Those resources should be able to be offered to different tenants, different content service providers to provide service, not just China Mobile to offer triple play services. What about a third-party company to provide a brilliant home security service. PaaS is the model they have in mind
- Q from Cisco: How you manage to have a centralized state, coordination across the system when you have distributed VNFs; you need to ensure that the centralized controller that you have is able to have a federated state across the system. How do you plan to architect this in a way that you are able to scale this out?
- A: There is a project in OPNFV, Edge NFV, where the goal is to put requirements in this area; it's not proceeding very fast, and I've just become involved in this project. One of the items we're addressing is how to centrally manage all of these distributed instances. The gap has been identified.
- Q from Diego: I don't see OpenStack working here at all. This is an interesting idea, not all the items are related to the group, but I encourage you to discuss and bring some proposals.
- Q from unknown person: Do you think there are new standards, what kind of solution are you considering now?
- A: We have just started this, we haven't gotten there yet. In our lab, we've developed a proxy, where we saw OpenStack didn't work
- Comment from Huawei: BBF has work on this, TR-317

Control and User Plane Separation – Architecture of Cloud-based BNG

Presenter: Gu Rong

draft-gu-nfvrg-cloud-bng-architecture

Slides: <https://www.ietf.org/proceedings/98/slides/slides-98-nfvrg-sessa-04-draft-gu-nfvrg-cloud-bng-architecture-00.pdf>

- Walked through draft, discussion of main points, and presented lab tests, where the HA and QoS tests were an issue, as well as scale.
- Q from Nokia: How does load balancing work?
- A: We have centralized control plane. When the user comes, connects to one of the UP, and the UP device connects to CP, receive response, and do the forwarding

- Comment from Sarah Banks, NETSCOUT: I agree with Diego, this might not be the correct home for this, however, the BBF is interested in this work, and we can connect you with a few folks from the BBF

Multi-Domain network virtualization

Presenter: Luis Contreras

draft-bernardos-nfvrg-multidomain

Slides: <https://www.ietf.org/proceedings/98/slides/slides-98-nfvrg-sessa-06-draft-bernardos-nfvrg-multidomain-00.pdf>

- Problem Statement
 - o Availability of different infrastructure environments pertaining to distinct administrative domains
 - o There are no established mechanisms for providing access to multi-domain environments in a standardized way
 - o A solution is needed to deal with both multi-operator and single operator multi-domain problem
- Architecture Proposition (see presentation)
- Next Steps
 - o Asking for feedback from NFVRG
- Q Jeremy Fuller: I'm chair of the group working on all of these issues. One of the real challenges we have is that people forget about control – the actual running of the VNF, which requires a VNF manager – and where is that VNF manager? Is it in the home network, visitor network, how does it get IP addresses.. I'm encouraged to see this work, but if you could bring in the VNF manager, and how it knows where things are and where it's located, that'd be great, thank you
- A: OK thank you. We foresee different use case families, VNFaaS (the remote provider is the one who provides the VNF manager), SliceaaS (here the VNF Manager is in the original network provider)
- Q from Jeremy: where is the Element Manager?
- A: In the main provider/primary provider, we are working on this
- Comment from Ramki: Bring in how it works with different hardware vendors, that'd be helpful
- Q from Cox: Your IF3, is it connecting to NFV0, or across domain service orchestrator?
- A: (see Slide 8 for answer)
- Q: Joe from Huawei: If you want to create and end to end type of slice, which orchestrator has control?
- A: The customer will have a relationship with a provider, that provider would be the one giving the orchestration
- Q from Kalyani, Verizon: Discussion around element management, if there are multiple domains in the SP, you can have control in one place or another, but if it's between SPs, there'll be different kinds of control between different SP domains
- A: Some kind of interaction between EMS or OSSs is something we'd have to look at.

Verification of NFV services: Problem statement and challenges

Presenter: Myung-Ki Shin

draft-irtf-nfvrg-service-verification

Slides: <https://www.ietf.org/proceedings/98/slides/slides-98-nfvrg-sessa-07-draft-irtf-nfvrg-service-verification-00.pdf>

- Overview
 - o NFV relocates network functions from dedicated hardware appliances to generic servers, so they can run in software. However, incomplete and/or inconsistency configuration of VNF and FGs (like service chains) may leads to verification issues.
- Draft was adopted Nov 2015, the draft is stable, but asking for input specifically on security aspects of the draft.
- Q from Ramki: Separation of processing and data planes – how could we align this work with verification?

- A: OK
- Q Diego: do you plan to put any references to strategies on conflict resolution?
- A: I didn't have a clear image of this, but we can talk more about this

Open Source Mano: An update on OSM to the NFVRG

Presenter: Diego Lopez

Slides: <https://www.ietf.org/proceedings/98/slides/slides-98-nfvrg-sessa-02-update-on-osm-00.pdf>

- Update
- No questions came in (time was short)

NFVRG session 2 Thursday 3:20-5:20pm

Network Virtualization Research Challenges

Presenter: Carlos Bernardos

draft-irtf-nfvrg-gaps-network-virtualization

Slides: <https://www.ietf.org/proceedings/98/slides/slides-98-nfvrg-sessb-09-draft-irtf-nfvrg-gaps-network-virtualization-01.pdf>

- Document/draft Overview and review
- Q from Alex: There's additional comments to be added, on programmability, these things impact the other sections
- Q from Georgious: Support for draft being in a good shape, sending it to LC you'll get more comments, it's a good start
- Q from Ramki: We're discovering more and more new items, like micro services.. will you update the draft with these or?
- A: This is where we are now.. we'll stop here
- Q from Diego: Alex's comments make sense.. we should be careful with the short term aspects, something that concerns me a little bit will be to make the content disconnected from the IETF as possible – when talking about mapping to Working Groups, instead, map them to areas..
- Q: Laurent Ciavaglia: I didn't catch this
- Q from Ramki: I didn't capture this
- A from Laurent:
- Q Kyle LeRose, Sandvine: We're trying to say this is what the RG should be working on, isn't that the role of the charter?
- Q from Sarah Banks, NETSCOUT: I'm not sure this is the appropriate document to publish. What value does this document, as an RFC, provide? It's work we should be working on as a group, certainly, but I don't see the value of this document as an RFC.
- Q Al Morton: I think it's valuable to show consensus, I think it's valuable to show that we agree on this set of gaps at this point of time. Are we expecting folks in the RG to do the individually to the different IETF WGs to get this work done?
- Comment from Diego: When we started I was a bit concerned that we couldn't publish these snapshots, then I talked with the RFC Editors, and asked them if they could publish a consensus and they said no. You can obsolete the RFC with a new version, there's a mechanism for that. My thought is that since there was

only one voice against this, there is general consensus, we could go for LC, and that's my recommendation.

- Diego/Chair did a hum for consensus, there was consensus in the room, he will send it to the list, and see what we get there.

General Discussion

- General comments from chairs
 - o Didn't get the feedback and interaction we wanted
- Bring ideas, we'll match it against what the group wants to work on
 - o We'll be more selective moving forward
- Fate of policy based resource management
 - o Work done, but not complete
- There are a couple of options, Diego wants to ask the editors and original authors
- Q from Bert, Huawei: If this isn't relevant work then why ask them to bring it back?
- Q from Alex: This work has been around since 1992.. it's a gap that has grown. As soon as you have thousands or millions of policies, there's a problem in maintaining the system, more complex than the system you want to change. Maybe redefine the problem
- Q from Alex: resource management – this is an evolving concept. This seems to be a research topic, in particular, when it comes to controlling conflicts or better usage of the resources (one of the 2).
- Comment from Ramki: Would you think a survey, looking at the trade offs, positives/negatives, would that be of some value to the community?
- Comment from Diego: Take it to the list

Network Coding in the SHINE ESA Project

Presenter: Simon Pietro Romano

Slides: <https://www.ietf.org/proceedings/98/slides/slides-98-nfvrg-sessb-11-network-coding-in-the-shine-esa-project-00.pdf>

- Presentation on SHINE - Secure Hybrid in Network caching Environment
- Basic concepts – use coded multicast across the satellite-enabled trunks of the overall platform, and use either MPEG-DASH or WebRTC within the edge access networks.
- You can take content at the source, properly encode the content, and then simply send multiplexed frames while crossing satellite channels
- Will work with Angeles on her draft and collaborate
- Use VNCF-based SHINE as a scenario
- Q from Ramki: Is the use case here for remote sites?
- A: The assumption here is that you find yourself in need of using satellite. One of the cases is the underdeveloped or under structured regions, or you want to provide a streaming service to maritime fleets
- Q from Ramki: Network coding side, all software implementation?
- A: There is already software, we need some changes, and tailor implementation to our specific use case
- Q from Brandon Williams, Akamai: What makes this a good candidate for VNF deployment? Is it that the sat providers happen to have the hosting platform that you need already? With the caching part, it tends to require a significant amount of resource that can be specialized. How did you come to the conclusion this is a good candidate?
- A: These points you raise are of interest to us, we're working with a Satellite provider, and they're exactly interested in this kind of approach; making the Sat Provider behave like other providers when it comes to the types of services they offer. Virtualization is seen as necessary in order to survive in the market
- Q from Ramki: Is the goal to provide real time services and cached on demand?
- A: Focused on real time services, with streaming as real time as possible, but you start with stored video.

Multi-access Edge Computing (MEC) Applications

Presenter: Hannu Flinck

Slides: <https://www.ietf.org/proceedings/98/slides/slides-98-nfvrg-sessb-12-multi-access-edge-computing-mec-applications-00.pdf>

- Overview
- Review of ETSI MEC Release 1 work
- Review of several MEC Applications (see slides)
- Q from Dan, AT&T: Is this to demonstrate what ETSI is doing?
- A: The work stems from ETSI
- Q from Dan: There are other initiatives that are trying to tackle this problem differently, with a different architecture and different objectives. It's worth noting that edge technologies are going to become increasingly important
- A: I fully agree with this
- Q Guy, Cox: Does this augment, do you add additional interfaces to ETSI MANO?
- A: Shows a slide from backup slides
- Q from Georgios, Huawei: Are you planning to bring results from the experiments that you have done?
- A: That remains to be discussed.
- Q from Ramki: What do you think are the big research challenges you foresee?
- A: The basics have been covered, but things like handover support, the notion of stateless and stateful services, how do we manage those?

Data-Intensive Function Acceleration

Presenter: Ning Zong

Slides: <https://www.ietf.org/proceedings/98/slides/slides-98-nfvrg-sessb-13-data-intensive-function-acceleration-00.pdf>

- Motivation: data plane performance (including throughput, latency, jitter) is one of the key challenges
 - o It's challenging to have low-latency and high-performance VNFs in an NFV environment
- Accelerator discussion
 - o IFA018 specifies the interface between data-intensive VNF and accelerator
- No questions

Towards Integration of Slice Networking in NFV

Presenter: Alex Galis

Slides: <https://www.ietf.org/proceedings/98/slides/slides-98-nfvrg-sessb-14-towards-integration-of-slice-networking-in-nfv-00.pdf>

- Purpose of presentation: No gap in the beginning between slicing and virtualization; somehow they evolved individually, and are just now being discussed to come back together.
- Q from Ramki: How is network slicing different from the other types of (didn't capture it)
- Proposal of a unified slice definition
- Q from Ramki: several slices can coexist on the same physical node?
- A: Yes, and it SHOULD coexist on the same physical node. How many is an engineering question.

- Q from Kyle Lerosé: Is slicing also guaranteeing resources? Does it make sure cache is shared properly, for example? Is this something that needs focus?
- A: Good question. Answer is yes/no. It's a management and control plane level system. The data plane is there, but does not automatically dictate the rest. In theory, the slice should have its own management system.
- Q from Guy Meador, Cox: It seems like this is a network service
- A: Yes
- Q from Guy Meador, Cox: #4 (from slide) would also apply to network services
- A: Agree
- Q from Dave Dolson, Sandvine: Import that slices be recursive?
- A: Fundamentally yes

EOM