

Considerations for Benchmarking VNFs and their Infrastructure

Al Morton

Nov, 2015

[draft-ietf-bmwg-virtual-net](#) (01)

Current WG 01 Additions

COMMENTS on the bmwg-list:

- Addressed Ramki Krishnan's comments on section 4.5, power, see that section (7/27 message to the list).
- Addressed Saurabh Chattopadhyay's 7/24 comments on VNF resources and other resource conditions and their effect on benchmarking, see section 3.4.
- Addressed Marius Georgescu's 7/17 comments on the list (sections 4.3 and 4.4).

AND, comments from the extended discussion during IETF-93 BMWG session:

- Section 4.2: VNF footprint and auxilliary metrics (Maryam Tahhan),
- Section 4.3: Verification affect metrics (Ramki Krishnan);
- Section 4.4: Auxilliary metrics in the Matrix (Maryam Tahhan, Scott Bradner, others)

Charter and Next step

- **VNF and Related Infrastructure Benchmarking:** Benchmarking Methodologies have reliably characterized many physical devices. This work item extends and enhances the methods to virtual network functions (VNF) and their unique supporting infrastructure.
- **A first deliverable** from this activity will be a document that considers the new benchmarking space to ensure that common issues are recognized from the start, using background materials from industry and SDOs (e.g., IETF, ETSI NFV). ... (Milestone in Aug 2015)
- This draft is **Referenced** in ETSI NFV GS, OPNFV specs.
- **WGLC?**

Backup

Current WG 00 Additions

- Using “General Purpose” computing
- Motivated complete desc. of test context: User Story (what’s in the Black Box?)
- Sec3: SUT description = platform and VNFs and...
- Barry Constantine’s comments on the list:
 - “concurrent” VNFs, how many can platform support?
- Sec3.4 Consider interactions/dependencies (placement, HA, VM or Bare Metal)
- Sec 4.1 Scale and capacity benchmarks still needed.
- Sec 4.4 Resolved the question of Scale and the 3x3 Matrix
- new 4.5, Power consumption (need metric(s))

Assess Benchmark Coverage 3x4

SPEED

ACCURACY

RELIABILITY

SCALE

Activation/
Creation/Setu
p


Operation

De-Activation/
Deletion/Take
-Down

Report Results (Capacity = N units)

	SPEED	ACCURACY	RELIABILITY
Activation/ Creation/Setup			
Operation			
De-Activation/ Deletion/Take- Down			

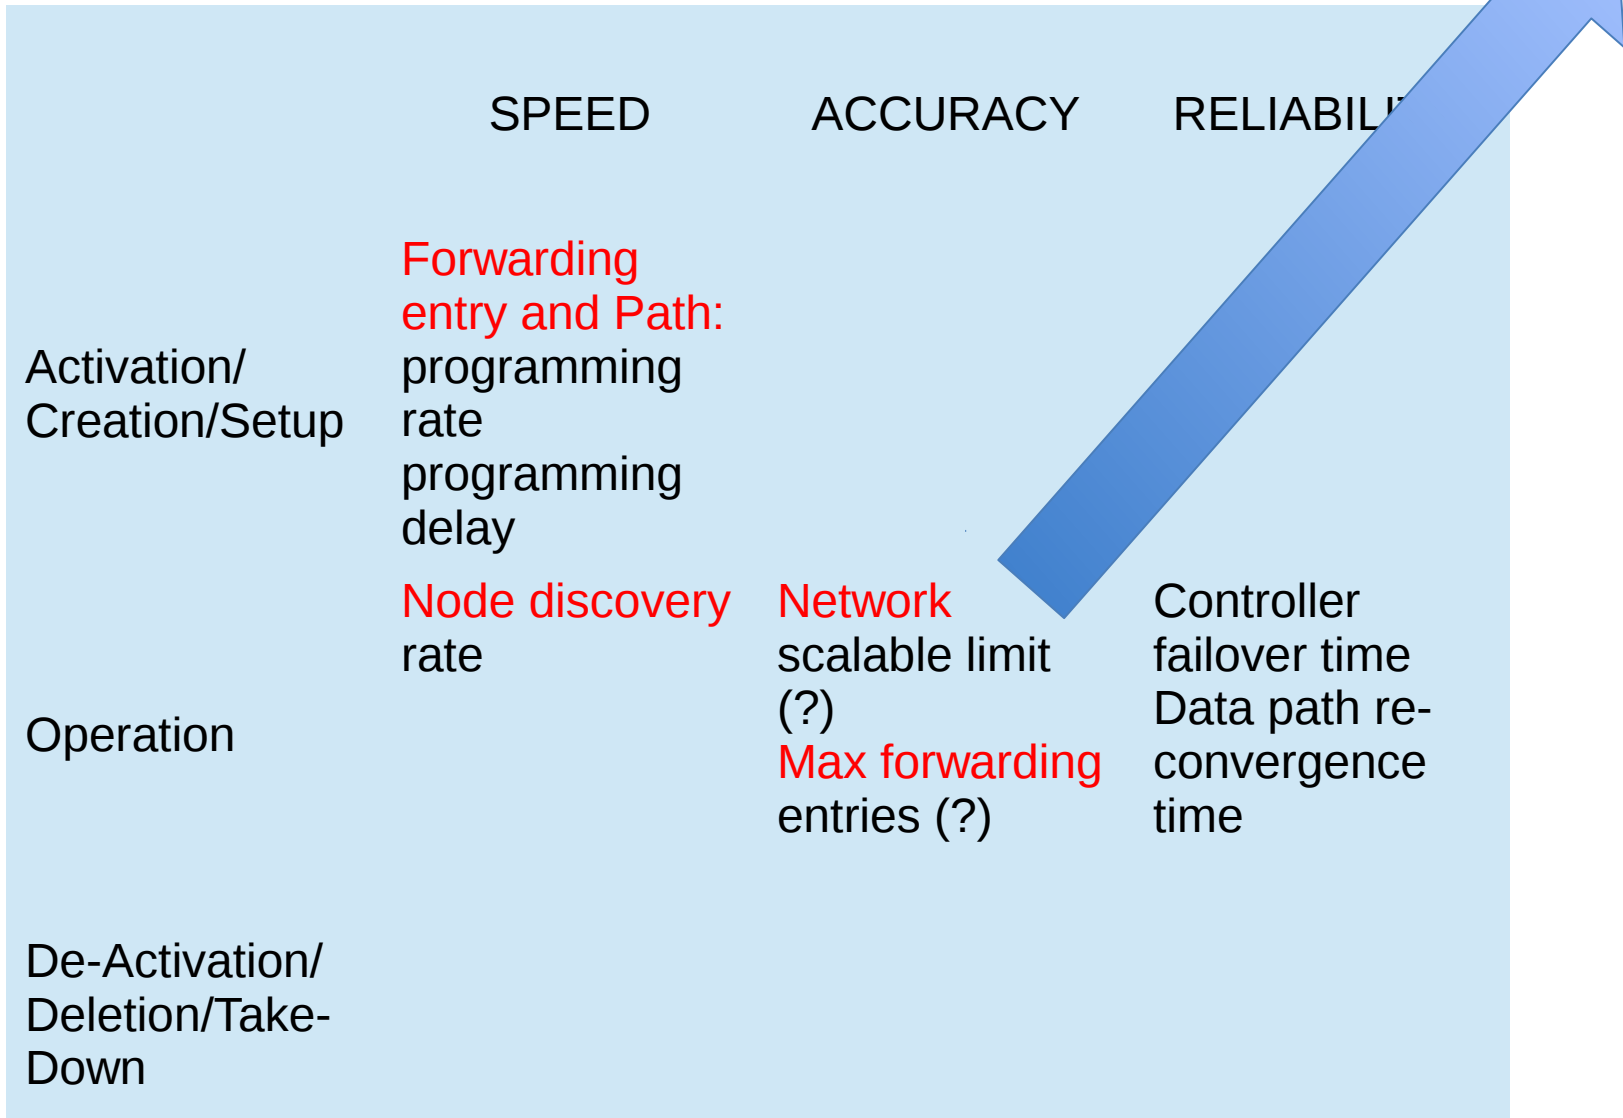
Report Results (Concurrent VNFs)



Some Related Work

- [ETSI NFV](#):
 - vSwitch Benchmarking Req (Acceleration-related)
 - Pre-deployment Testing of VNFs and Infrastructure
 - Interoperability Testing
- [OPNFV \(Open Platform for NFV\)](#):
 - Characterize vSwitch Performance for Telco NFV
 - Many other testing projects
- [OPEN DAYLIGHT](#):
 - Wrapped Cbench – [WCBENCH](#) – [Daniel Farrell](#)

SDN Controller Coverage:



Example: Quality Metric Coverage for Virtual Machines

	SPEED	ACCURACY	RELIABILITY
Activation/ Creation/Setup	<u>Successful Activation Time</u>	Incorrect Activations per total attempts	Failed/DOA Activations per total attempts
Operation	I/O Capacity Benchmarks on CPU, Memory, Storage	Incorrect outcomes per Operation attempts	Error/Stall outcomes per Operation attempts
De-Activation/ Deletion/Take- Down	Successful De- Activation Time	Incorrect De-Activations per total att.	Failed/no-resp. De-Activations per total att.

Version 01, Benchmarking Considerations

- Comparison with Physical Network Functions
 - Re-use of existing benchmarks, with review
- Continued Emphasis on Black-Box Benchmarks
 - Internal Metrics from Open Source are tempting
 - Supply both, may provide useful OPS insight
- New Benchmarks for a Dynamic World
 - Time to deploy VNFs, Time to Migrate,
- Assessment of Benchmark Coverage

Ver 02, HW & Test Considerations

Section 4.4

- How do we reflect Scale/Capacity Benchmarks in the 3x3 Matrix? Alternatives:
 - Add a new column
 - Include Scaleability under Reliability
 - Keep Size, Capacity, and Scale separate from the matrix and present results (using the matrix) with titles that give details of configuration and scale.
- Yes, results could be organized by Matrix, too.

Test Configuration (ver 00)

- o number of server blades (shelf occupation)
- o CPUs
- o caches
- o storage system
- o I/O

configurations that support the VNF:

- Hypervisor
 - o Virtual Machine
 - o Infrastructure Virtual Network

the VNF itself:

- specific function being implemented in VNF
 - o number of VNF components in the service function chain
 - o number of physical interfaces and links transited in the service function chain

characterizing perf at capacity limits may change? (ver 00)

- Charac. Infrastructure support of #? VMs:
 - N when all VM at 100% Util
 - $2*N$ when all VM at 50% Util ??
- #? VNF profile A, VNF profile B
 - Profiles may include I/O, storage, CPU demands
- Partition VNF performance
 - from single VNF in infinite I/O loop
- System errors occur as transients (longer dur.)
- VM and VNF flux: constant change in population while characterizing performance

Third level

Fourth level

Fifth level

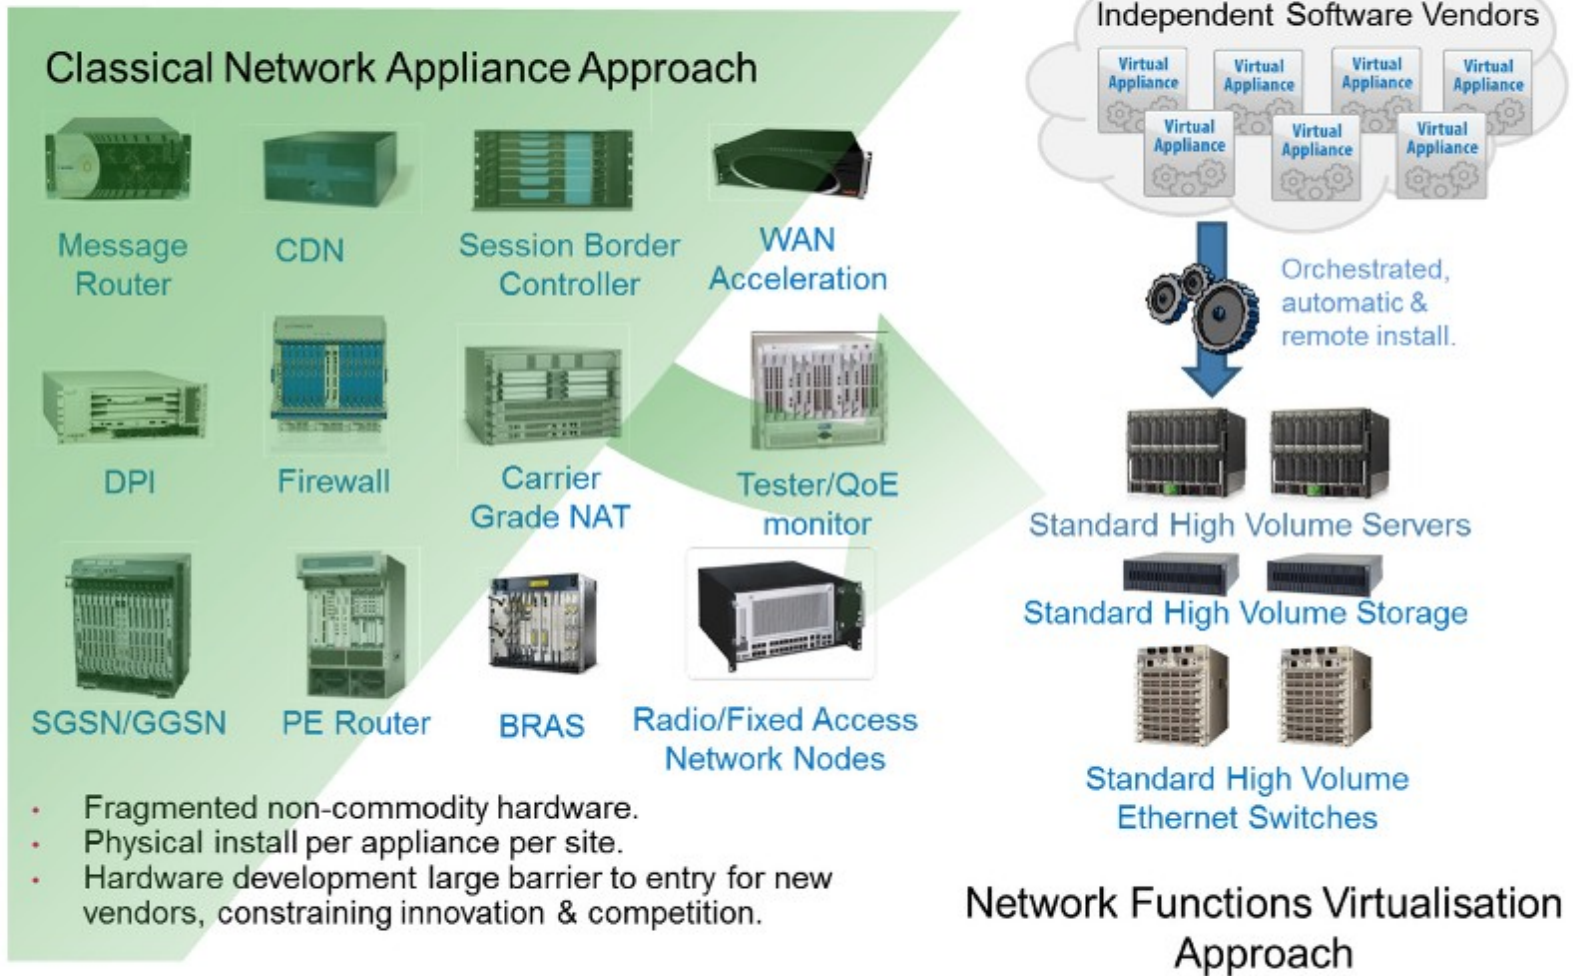


Figure 1: Vision for Network Functions Virtualisation

<http://www.etsi.org/technologies-clusters/technologies/nfv>