ONAP open lab – A fully virtual flavor

Ver 0.3

Helen Chen, William Chen, Lingli Deng, Chris Donley, Ranny Haiby, Jason Hunt, Catherine Lefevre, Tomer Oster, Oliver Spatscheck, Chengli Wang, Gary Wu, Kang Xi, Yang Xu

May, 2017

Purpose of building an open lab

- The primary focus of the ONAP lab is to support the CI/CD of the board approved release use cases
  - Run the release gating test cases
- By having additional lab(s) we may :
  - Provide easy access for all developers in the community
  - Support ONAP demos
  - Support Interoperability testing with multi-vendor's hardware and software in "real" environments
  - Showcase use cases outside the board approved ones

We assume there would be multiple labs (physical and virtual) and multiple usages (CI/CD, Demos, Integration, etc.). The assignment of usages to labs is out of the scope of this document, but the goal is not to limit any possible assignment

Implementation options

- Physical (Will of course use virtualization on top of the lab hardware)
  - Requires dedicated hardware and real estate
    - May come from community member contributions – no CapEx upfront
  - Requires support personnel – Already committed by some community members
  - Has fixed capacity with typically long lead times for expansion
  - Limited SLA – Downtime may occur

- Virtual (cloud hosted)
  - No platform maintenance required (handled by the cloud provider)
  - Elastic capacity
  - Globally accessible
  - “Practice what we preach” – full virtualization
  - Incurred OpEx
  - Well defined SLA

Challenges of a virtual lab

- Cost
  - Annual budget for resource consumption
  - Resource capping (per vendor? Only paying members? Etc.)
- Running vendor VNFs on a public cloud
  - Integration with the cloud platform (CPU, storage and networking)
  - Security and confidentiality – Who is liable in case of data breech that will compromise vendor assets?
  - Integration with cloud platform management (may be different than plain OpenStack)
  - Performance – how to match on premise capacity to cloud offering
- Connectivity to physical hardware

Addressing the challenges

- Cost and capping
  - Virtual lab footprint can be kept minimal during normal times and scaled out when there is higher demand
    - When there is a need for more testing (e.g. prior to a release)
    - When there is an outage in the physical labs
    - To support an event (trade show, developer conference, etc.)
  - Agree on capping policy in the community
    - May require assigning community lab managers that will be responsible for lab scaling
    - Will require a usage tracking tool
- Repeatability
  - Create lab blueprints, or “packages” that will be reusable and allow scaling out of a virtual lab or creating a new one efficiently

Addressing the challenges

- Running vendor VNF on public cloud
  - Running OpenStack (and later other VIMs) on the public cloud could provide a standard platform.
  - OpenStack may be run on bare metal servers or using nested virtualization (former has better performance, latter higher portability). VMWare may be a challenge.
  - Minimum requirement for the cloud is to support X86
  - ONAP may run directly on public cloud or on hosted OpenStack

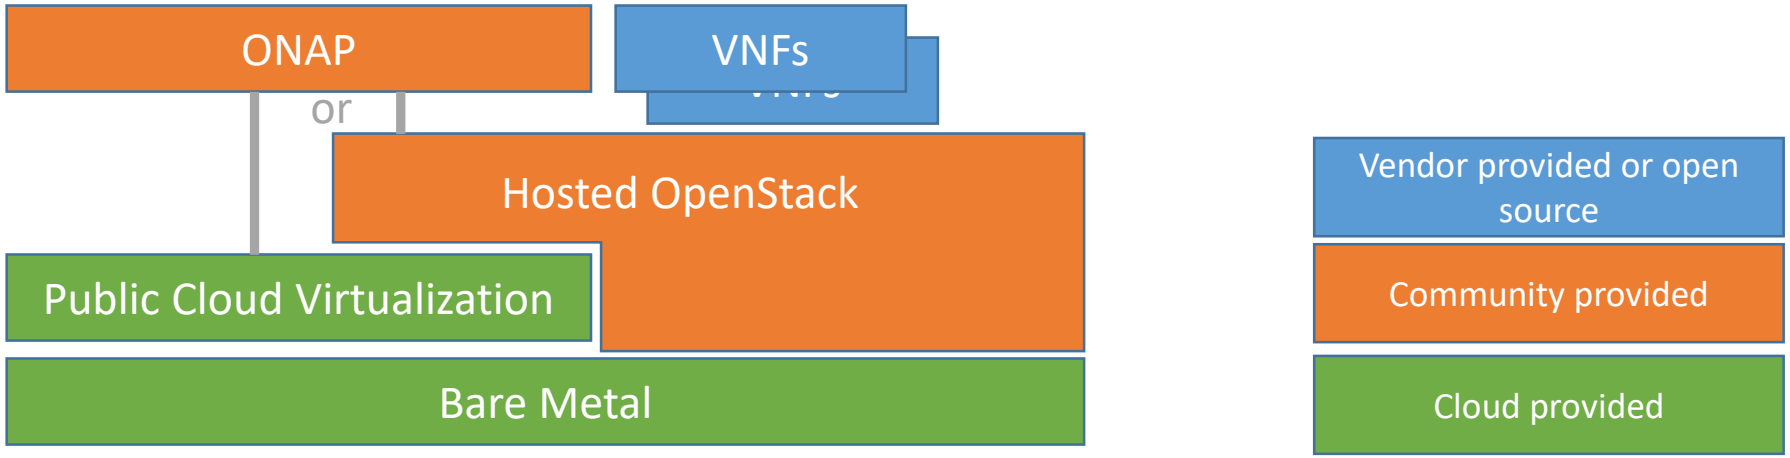

Addressing the challenges

- Connectivity to physical hardware
  - May be solved by using security gateways as part of the lab that will maintain secure tunnels with physical labs of the operators and vendors
  - The tunnels will be used for both VNF-PNF traffic and ONAP-PNF control traffic.
  - Some VNFs may run on the Vendor/Operator lab

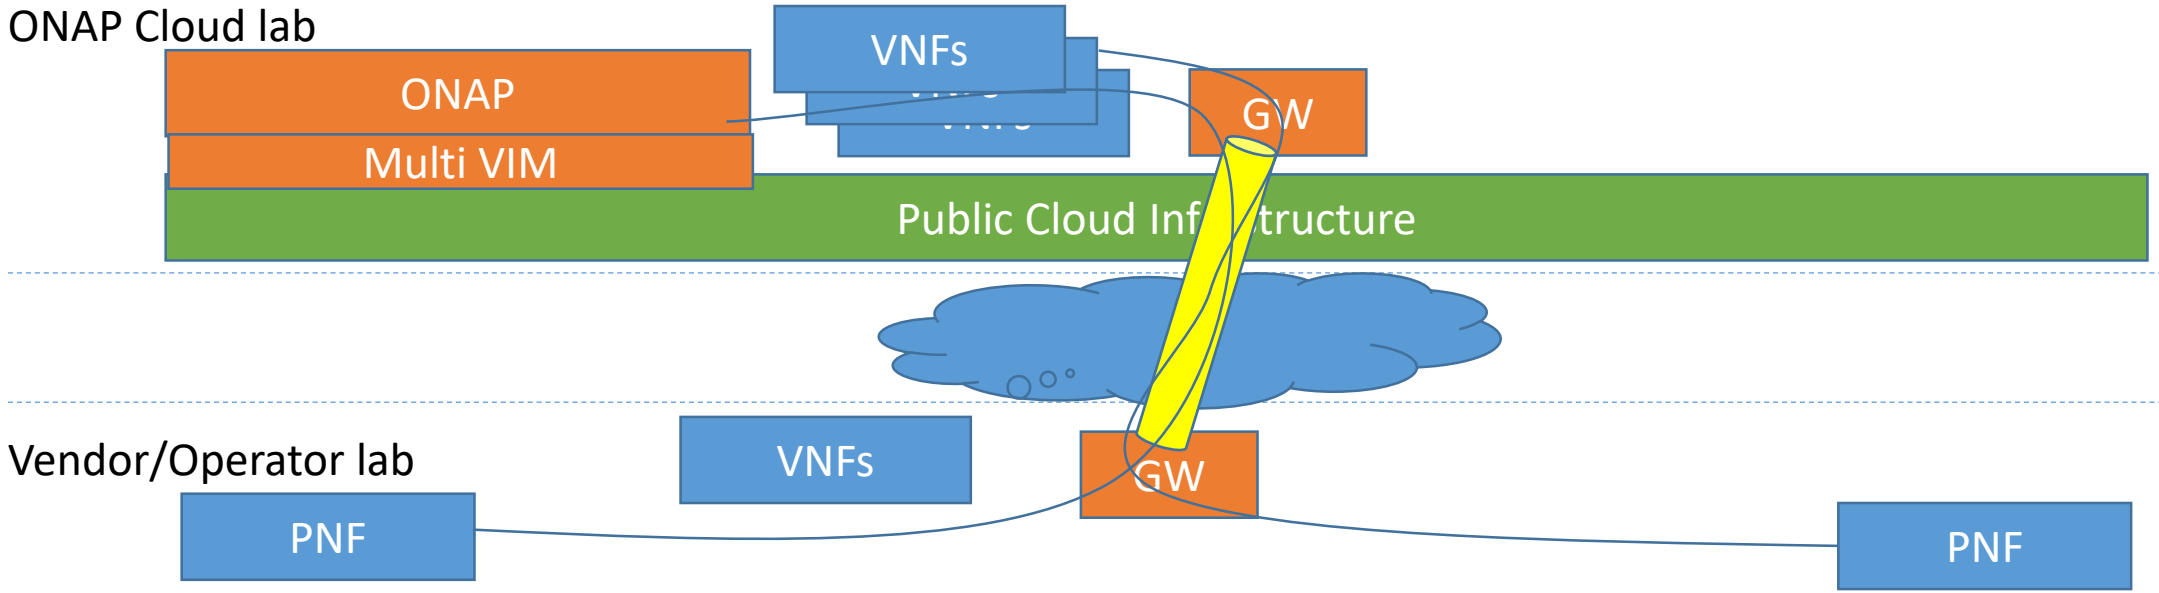

Addressing the challenges

- Security and liability
  - There needs to be a clearly assigned responsible entity that will be in charge of running the lab, much like any other, physical lab.

Possible usage scenarios

Scenario 1
ONAP in vLAB, VNFs in pLAB
(This may be easier to accomplish for release 1)

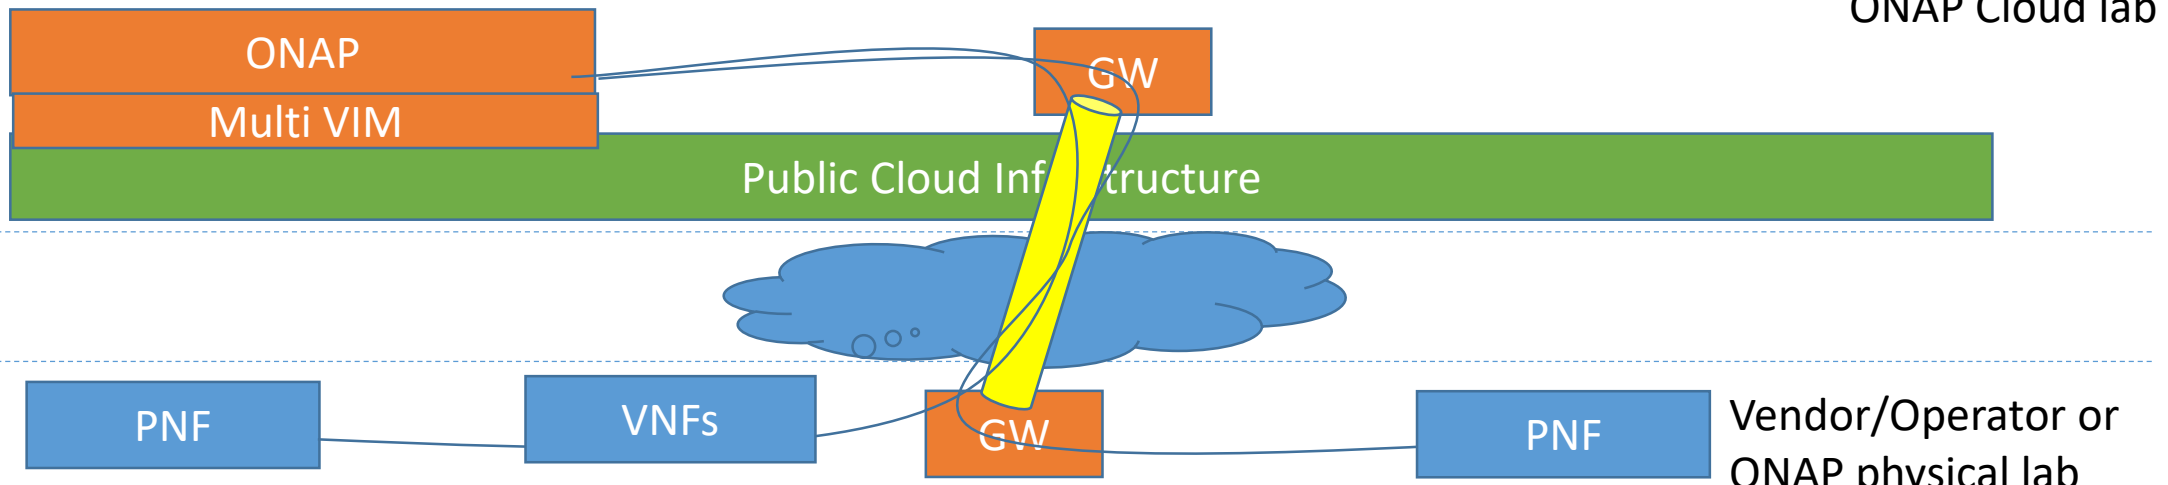

Scenario 2
ONAP and VNFs in vLAB

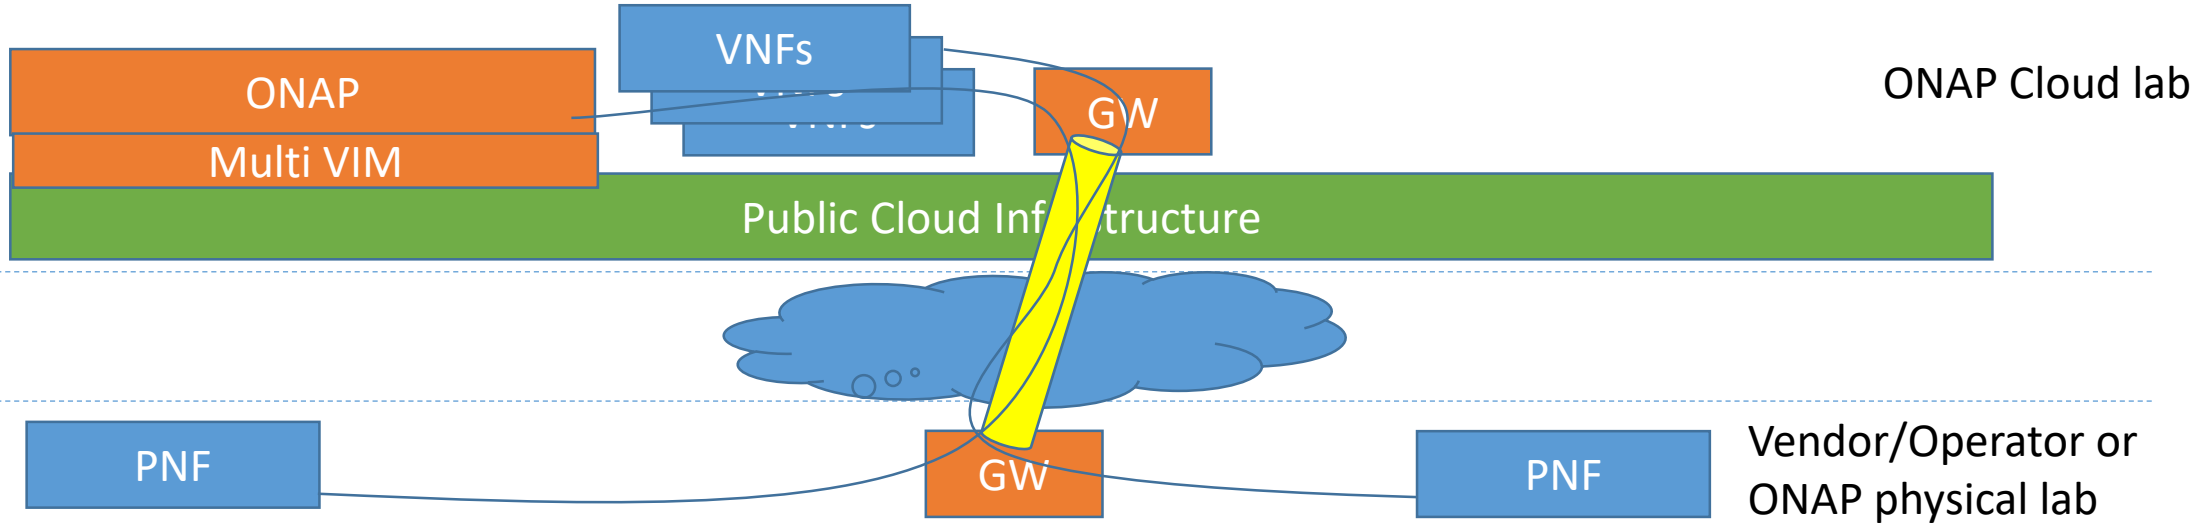

A mesh of labs

- Why stop at two?

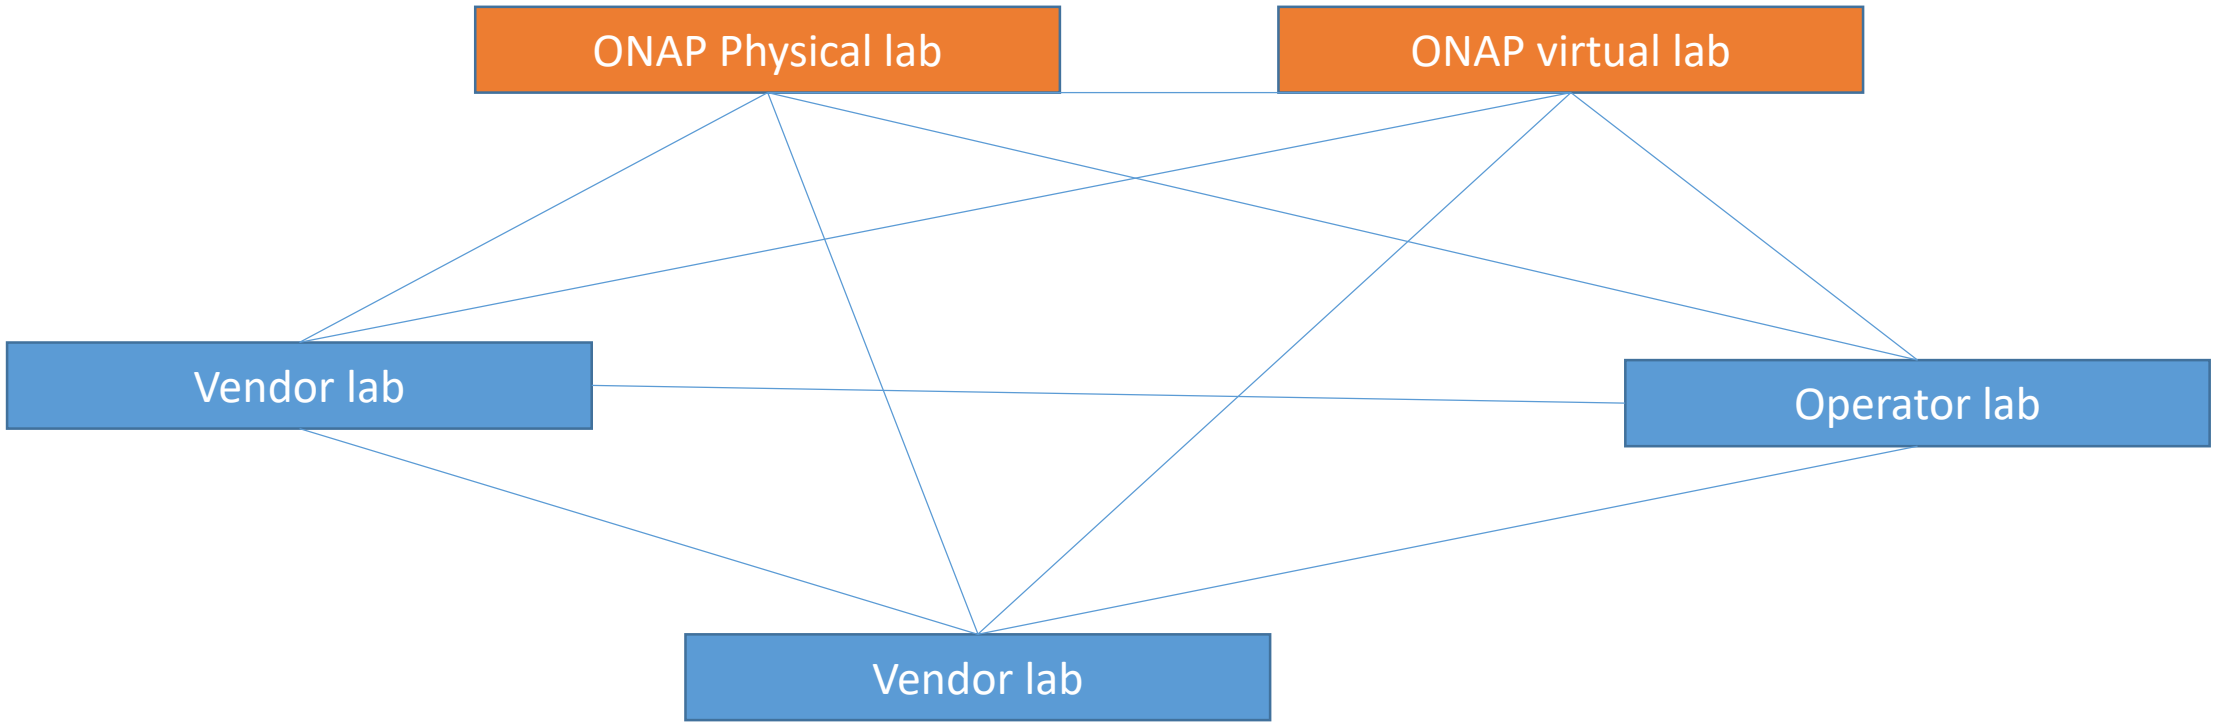

Resource tally sheet

- Based on https://wiki.onap.org/display/DW/Lab+Resource

		Memory	Storage	Network	Qty
Lab					
Jump server	Xeon-E5-2658A 2.2GHz equivalent	256G	2T	2X10G	1
Security GW	Xeon-E5-2658A 2.2GHz equivalent	256G	100G	2X10G	2
ONAP					
ONAP	Xeon-E5-2658A 2.2GHz equivalent	256G	2T	2X10G	3
OpenStack					
Controllers	Xeon-E5-2658A 2.2GHz equivalent	256G	2T	2X10G	2
Compute	Xeon-E5-2658A 2.2GHz equivalent	256G	2T	2X10G	3

Public cloud requirements

- X86 based platform (e.g. not ARM, etc.) – The majority of VNFs are only certified to run on X86
- Support for running KVM virtualization and OpenStack VIM
  - Advantage for bare-metal option as it will avoid the need for nested virtualization
- Support for overlay secure tunnels between VMs in the cloud and physical labs
  - E.g. IPSec or SSL
- Compute, storage and networking resources as specified in the previous slide
- SLA – A committed reliability/resiliency of the provided cloud resources
- <More requirements to follow based on use case analysis>

Summary and next steps

- We identified several usages (CI/CD, integration, development, etc.) for a virtual lab
- We identified several scenarios (local VNFs, remote VNFs, etc.)
- Release 1 scope may be limited and with probably be focused on the scenario where VNFs are in a remote lab and only ONAP is running in a virtual lab.
- Next steps would involve:
  - finalizing the exact requirements – This can only happen once the detailed use cases are agreed upon
  - Approaching cloud providers with a request for proposal that will include the requiremtns
  - Approving the

ONAP
OPEN NETWORK AUTOMATION PLATFORM

Backup (for future use)

Cost assessment

- Annual cost estimate

	Minimum (steady state)	Max (burst)
Compute		
Network		
Storage		
Personnel		
???		
Total		