

OPNFV Lab-as-a-Service (LaaS)

<https://wiki.opnfv.org/display/INF/Lab+as+a+Service>

- Objective
 - Offer openly accessible server resources to developers – to test-drive, demo, test, extend OPNFV and related solutions
 - Offer readily installed scenarios, per the choice of the developer
 - Support a wide variety of use-cases (developer (virtual/bare-metal deployments), demo, CI, XCI, x-community (e.g. ONAP):
<https://wiki.opnfv.org/display/INF/LaaS+scenarios>)
- Approach
 - LaaS is a cloud/bare-metal environment, providing automatic (“one click”) provisioning, deployment/delivery, decommissioning of resources

LaaS: Initial use-cases and required number of servers

OPNFV LaaS initial use-cases

Use Case	Description
Snapshot deploy	Spawn a deployment from snapshot
Initial master deploy	Deploy certain OPNFV scenario from scratch from master using an existing artifact
Initial stable deploy	Deploy certain OPNFV scenario from scratch using a released OPNFV version
Build and deploy	Build the artifact and deploy certain OPNFV scenario for the given patch
Deploy with addons	Deploy certain OPNFV scenario and provide additional scenarios to deploy/develop other components of the stack such as ONAP
OPNFV for ONAP developer	Deploy and integrate OPNFV scenario and ONAP instance for developer use
OPNFV for ONAP X-CI	Deploy and integrate OPNFV scenario and ONAP instance for X-CI
OPNFV+ONAP CI/CD	Deploy and integrate OPNFV scenario and ONAP instance for full CI/CD/testing
Deploy OS	Provide a machine with OS installation only

See [LaaS scenarios](#) for a detailed description of LaaS use cases and associated requirements.



LaaS - Number of Servers

The following are typical numbers to be expected, following requests seen from the community that have been observed over the past few months (see e.g. <https://jira.opnfv.org/browse/INFRA-157?jql=project%203D%20INFRA%20AND%20text%20~%20vpod%20ORDER%20BY%20created%20DESC>) as well as what is expected for the future (especially for XCI and ONAP)

- OPNFV Developer (software) - vPOD: There are quite a few community requests for vPODs (see)
 - 12 x vPOD = 12 x Server (Intel) (note: ARM does not support nested virtualization yet)
- OPNFV Developer and XCI (hardware) - physical POD:
 - 3 x Pharos POD (x86)
 - 1 x Pharos POD (ARM)
- ONAP + OPNFV Developer - physical POD:
 - 1 x Pharos POD (x86)
 - 1 x Pharos POD (ARM)

Overall requirement:

- Based on the above
 - 36 Servers (x86) (12 x vPOD + 4 * 6 Pharos POD)
 - 12 Servers (ARM)
- Sparring / specific requests
 - 2 Servers (x86)
 - 2 Servers (ARM)
- **TOTAL:**
 - 38 Servers (x86)
 - 14 Servers (ARM)

Reference: <https://wiki.opnfv.org/display/INF/Lab+as+a+Service>

LaaS hardware requirements: Details

Hardware requirements blessed by Infra-WG meeting on Aug 21, 2017

LaaS Hardware

A key objective of LaaS is to create a development environment on the fly, whether it is vPOD, Pharos baremetal POD or and OPNFV-ONAP POD. LaaS will avoid dedicating a particular server to a particular role, but rather allow for flexible allocation following demand. As a consequence, LaaS assumes a single type of server which can be used for a vPOD, a Pharos POD, or an ONAP POD.

Server Hardware (x86 based)

- Intel XEON E5-2699 v4 2.20 GHz, 145W, 22C, 55MB (high number of cores (22), well suited to run VMs) - https://ark.intel.com/products/91317/Intel-Xeon-Processor-E5-2699-v4-55M-Cache-2_20-GHz
- 512 GB of DDR4 ECC RAM
- 2 x 10G DPDK-compliant NICs (e.g. Intel X710: 4 port 10GE) to allow for dedicated ports for admin / public / tenant / test networks
- 0.9 TB of disk (2 x 480G SSD)
- Lights-out-management (must supply IPMI compatible solution)

Server Hardware (ARM based)

- Cavium ThunderX ARM processor (64bit ARMv8 architecture, 48 cores per processor, 2.0GHz BGA 2601), e.g. ASA ASA9103-48C-TX
- 256 GB of DDR4 ECC RAM
- 2 x 10G and DPDK capable NICs (4 port 10GE each); In case of on-board NICs (like e.g. with ThunderX), a minimum of 4 ports of 10G should be supplied
- 0.9 TB of disk (2 x 480G SSD)
- Lights-out-management (must supply IPMI compatible solution)

Network Switches

Switches should support 10/40/100Gbps ports to future proof the setup (switch hardware PODs from 10 Gbps to either 40 Gbps or 100 Gbps to also support performance testing), e.g. Cisco Nexus 92160YC-X.

Assuming the above number of servers: 38 x86 based servers, and 14 ARM based servers and further assuming that a maximum of 6 ports per server would be wired up to the switch, a total of 312 switch ports would be needed.

This means that just for connecting the server-ports (not counting ports on uplink/spine switches), a total of 7 48-port switches would be needed.

Status

- Quotation for the equipment
 - Requests for quotes based on the list of hardware requirements sent for
 - “Buy the equipment”
 - “Host the equipment”
 - “Metal-as-a-Service”
 - Some responses received, several still pending
- LaaS Development
 - LaaS driven by Infra WG
 - UNH is currently working on developing the basic functionality – see the [link](#)
 - Onboarding of installers to LaaS community discussion

Ask to the TSC

- Does the TSC endorse the LaaS program for OPNFV as described on <https://wiki.opnfv.org/display/INF/Lab+as+a+Service>