LaaS – TSC Update 2017-11-21

LINUX FOUNDATION
COLLABORATIVE PROJECTS

OPNFV Lab-as-a-Service

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

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

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.

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

LaaS - Number of Servers

The following are typical numbers to be expected, following requests seem from the community that have been observed over the past few months (see e.g. https://jira.opnfv.org/browse/INFRA-157?jql=project%20%3D%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)
- Sparing / specific requests
  - 2 Servers (x86)
  - 2 Servers (ARM)
- **TOTAL:**
  - **38 Servers (x86)**
  - **14 Servers (ARM)**

LaaS – Priorities

- Make available as soon as possible → launch an “MVP”
- Establish provisional Access and Usage Policy during “MVP phase”
- Agree on a minimum set of fundamental principles
- Establish a task force to drive implementation and work out further details/proposals

MVP Phase

- Important to bring up LaaS asap in order to
  - Gain practical experience for tooling and appropriate usage policies in view of more sophisticated usage models to be launched later
  - Provide more development resources quickly where needed

- Recommended to launch an “MVP” consisting of
  - Booking requests only via dashboard/Web UI
  - Provision machines with either one of Ubuntu 16.04, CentOS7, openSUSE Leap 42.3
  - Only short booking periods, limited number of extensions → avoid locking down usage patterns during MVP phase
  - Installer onboarding in a second step after MVP phase

Access and Usage Policy – MVP phase

- Access will be limited to SSH only
- Allow only time-limited bookings (1 week, max. 4 extensions = 4 weeks total)
- Upon expiration of booking, terminate access, wipe out machine, and put back into the pool
- May have to restrict number of user accounts per project/machine due to capacity reasons (tbd)

Fundamental principles

- Guiding principle for eventual usage policy: maximize overall community benefit
  - Avoid hard limits for usage duration
  - Enable repurposing of resources for increased community benefit

- Installer onboarding requires supporting the relevant Infra WG initiatives
  - Full PDF/IDF support in order to enable full automation while keeping lab usage flexible

Implementation Task Force

- It is necessary to have a small team that oversees the setup, development and roll out of LaaS
- Work closely with the development team and set priorities
- Incorporating needs of sister communities (ODL, fd.io, ONAP, etc.)
- Work out requirements and solutions for tooling & post-MVP policies