

# LaaS scenarios

Use Case Name	Description	Expected User Types	Hardware requirements (Hardware examples)	Provided by	Notes / Discussion
developer (software)	<p>Access to a vPOD for the purpose of working on various OPNFV projects.</p> <p>This specific use case would support development of projects that don't require higher performance or specific hardware requirements.</p> <p>The provisioned environment can vary from just having base OS to all in one OpenStack to HA OPNFV Scenario.</p>	<ul style="list-style-type: none"> <li>developers from different type of projects (installer, feature, test, etc.) implementing new features, integrating components, writing test cases.</li> <li>students/interns.</li> </ul>	<p>single server with virtual pods</p> <p>Hardware example:</p> <ul style="list-style-type: none"> <li>Intel XEON E5-2699 v4 2.20 GHz, 145W, 22C, 55MB (high number of cores (22), well suited to run VMs) - <a href="https://ark.intel.com/products/91317/Intel-Xeon-Processor-E5-2699-v4-55M-Cache-2_20-GHz">https://ark.intel.com/products/91317/Intel-Xeon-Processor-E5-2699-v4-55M-Cache-2_20-GHz</a></li> <li>512 GB of DDR4 ECC RAM</li> <li>2 x 10 Gbps NICs</li> <li>0.9 TB of SSD (2 x 480G SSD)</li> </ul>	community	
developer (hardware)	<p>Access to a baremetal POD for the purpose of working on various OPNFV projects. This specific use case would support development of projects that require higher performance or specific hardware requirements.</p> <p>The provisioned environment can vary from just having jumpshost provisioned and the rest of the nodes with no OS at all to full blown HA OPNFV Scenario.</p>	<ul style="list-style-type: none"> <li>developers from different type of projects (installer, feature, test, etc.) implementing new features, integrating components, writing test cases.</li> </ul>	<p>Pharos pod hardware:</p> <p>Hardware example:</p> <ul style="list-style-type: none"> <li>6 Servers Intel <ul style="list-style-type: none"> <li>Intel XEON E5-2667 v4 3.20 GHz, 135W, 8C, 25MB <a href="https://ark.intel.com/products/92979/Intel-Xeon-Processor-E5-2667-v4-25M-Cache-3_20-GHz">https://ark.intel.com/products/92979/Intel-Xeon-Processor-E5-2667-v4-25M-Cache-3_20-GHz</a></li> <li>128 GB of DDR4 ECC RAM</li> <li>2 x 10 Gbps NICs or 100 Gbps NICs</li> <li>.9 TB of SSD (2 x 480G SSD)</li> </ul> </li> <li>6 Servers ARM <ul style="list-style-type: none"> <li>(details to be added)</li> </ul> </li> <li>Switch with &gt; 12 x 10Gbps ports</li> </ul>	community or LF	
demo	<p>Access to a pre-deployed pod, for the purpose of trying or demonstration a working deployment of OPNFV / OpenStack.</p>	end user(maybe not member of OPNFV)	<p>single server with virtual pods</p> <p>Hardware example:</p> <ul style="list-style-type: none"> <li>Intel XEON E5-2699 v4 2.20 GHz, 145W, 22C, 55MB (high number of cores (22), well suited to run VMs) - <a href="https://ark.intel.com/products/91317/Intel-Xeon-Processor-E5-2699-v4-55M-Cache-2_20-GHz">https://ark.intel.com/products/91317/Intel-Xeon-Processor-E5-2699-v4-55M-Cache-2_20-GHz</a></li> <li>512 GB of DDR4 ECC RAM</li> <li>2 x 10 Gbps NICs</li> <li>0.9 TB of SSD (2 x 480G SSD)</li> </ul>	community or LF	
xci		<p>upstream community:</p> <ul style="list-style-type: none"> <li>openstack</li> <li>opendaylight</li> <li><a href="#">FD.io</a></li> <li>ovs</li> </ul>	virtual pods or physical pods		<p><del>I'm not sure we can support this in the LaaS, as the pods are not expected to be controlled by Jenkins, or to be "long-lived", etc.</del></p> <p>XCI has no dependency to Jenkins, etc and the lifetime of the deployment is not relevant. The point with XCI is to have ability to deploy latest verified versions of different components or even patches.</p>

ONAP developer	OPNFV scenarios are deployed on-demand for ONAP developers to test against their ONAP deployments (in another lab, connected to OPNFV LaaS over VPN), as needed for specific NFVI features or to work with specific OPNFV distros (the rest of ONAP devs are fine with generic/public clouds for testing VNFs).	ONAP developer	<p>10 virtual servers (5-10 physical servers) on avg consumed for ONAP devs, per the assumptions:</p> <ul style="list-style-type: none"> <li>• A partially or fully allocated Pharos compliant POD server, to run a virtual OPNFV scenario deploy and deploy VNFs for testing <ul style="list-style-type: none"> <li>◦ typical small/medium VNF (e.g. vRouter, vFW, vLB, vDNS, vPE, vCE): 50% of a server</li> <li>◦ typical large VNF (e.g. vEPC, vIMS): 1 full server</li> </ul> </li> <li>• 25% of ONAP developers spending avg 4 hours per day testing with the OPNFV LaaS resources <ul style="list-style-type: none"> <li>◦ 5% of ONAP developers active on OPNFV LaaS resources at any one time (assuming even timezone-distribution)</li> <li>◦ thus for ~150 ONAP devs, &lt;10 active devs at any one time</li> </ul> </li> </ul> <p>See also: <a href="https://wiki.onap.org/display/DW/ONAP+Lab+Specification">https://wiki.onap.org/display/DW/ONAP+Lab+Specification</a></p>		
ONAP X-CI	A new or existing OPNFV deployment is used for X-CI runs to verify tests against ONAP code changes. An existing OPNFV deployment (refreshed as needed) can be used to test against multiple ONAP X-CI jobs to reduce X-CI overhead (resources, test duration). A new deploy can be used if the type of X-CI job requires it (why is TBD).	OPNFV X-CI team	At least one LaaS virtual server (existing deployment), or multiple (TBD) per X-CI job forecast		
OPNFV+ ONAP developer	New or existing OPNFV deployments are used for OPNFV devs to test OPNFV+ONAP integration. Several OPNFV scenarios can be used in parallel to verify multi-VIM or scenario-focused tests (e.g. NetReady).	OPNFV developer	Probably up to 5 physical servers (5-10 virtual servers for distinct scenarios) for ongoing ONAP integration testing.		