

Lab Specification and Setup

Lab and BM RA alignment towards CNTT RA2:

Cloud Infrastructure Software Profile Capabilities

Reference Model Section	Reference	Description	Requirement for Basic Profile	Requirement for Network Intensive Profile	BM RA Notes	OPNFV Intel Lab
4.2.5	e. cap .001	Max number of vCPU that can be assigned to a single Pod by the Cloud Infrastructure	At least 16 (1)	At least 16 (1)		[OK] POD 19: 22c/44t [OK] POD 21: 20c/44t
4.2.5	e. cap .002	Max memory in MB that can be assigned to a single Pod by the Cloud Infrastructure	at least 32 GB(1)	at least 32 GB(1)		[OK] POD 19: 64G [OK] POD 21: 192G
4.2.5	e. cap .003	Max storage in GB that can be assigned to a single Pod by the Cloud Infrastructure	at least 320 GB(1)	at least 320 GB(1)		[OK] POD 19: 3TB (HDD) [OK] POD 21: 800G (SSD)
4.2.5	e. cap .004	Max number of connection points that can be assigned to a single Pod by the Cloud Infrastructure	6	6	Possible through Multus + CNI (vEth, SR-IOV)	
4.2.5	e. cap .005	Max storage in GB that can be attached / mounted to Pod by the Cloud Infrastructure	Up to 16TB(2)	Up to 16TB(2)		
4.2.5	e. cap .006	CPU pinning support	Not required	Must support	Possible through CMK (or Privileged flag)	
4.2.5	e. cap .007	NUMA support	Not required	Must support	(Maybe) Possible through CMK (or other BM RA tool)	
4.2.5	e. cap .008	IPSec Acceleration using the virtio-ipsec interface	Not required	Optional	(??) Might be covered by QAT device plugin	QAT support depends on chipset
4.2.5	e. cap .009	Crypto Acceleration using the virtio-crypto interface	Not required	Optional	(??) Might be covered by QAT device plugin	QAT support depends on chipset
4.2.5	e. cap .010	Transcoding Acceleration	Not required	Not required	(??)	
4.2.5	e. cap .011	Programmable Acceleration	Not required	Not required	(??)	
4.2.5	e. cap .012	Enhanced Cache Management: L=Lean; E=Equal; X=eXpanded	E	E	(??)	
4.2.5	e. cap .013	SR-IOV over PCI-PT	Must support	Not required	Not sure if this is specifically supported	
4.2.5	e. cap .014	Hardware coprocessor support (GPU/NPU)	Not required	Not required	GPU Device Plugin	Depends on Lab HW

4.2.5	e. cap .015	SmartNICs	Optional	Not required	(Maybe) Not sure if supported.	Depends on Lab HW
4.2.5	e. cap .016	FPGA/other Acceleration H/W	Optional	Not required	FPGA Device Plugin	Depends on Lab HW
4.2.5	e. cap .017	Ability to monitor L2-L7 data from workload	n/a(3)	n/a(3)		
4.2.5	i. cap .014	Indicates the number of CPU cores consumed by the Cloud Infrastructure on the worker nodes	2	2	IIRC this is possible with BM RA. Expecting it to be "core 0" on each socket	
4.2.5	i. cap .015	Indicates the memory consumed by Cloud Infrastructure on the worker nodes	16 GB	16GB	TBD	
4.2.5	i. cap .016	Number of virtual cores per physical core; also known as CPU overbooking ratio that is required	1:1	1:1	Should be default with K8s request /limit	
4.2.5	i. cap .017	QoS enablement of the connection point (vNIC or interface)	Not required	Must support	TBD	
4.2.5	i. cap .018	Support for huge pages	Not required	Must support	Part of K8s	
4.2.5	i. pm. 001	Monitor worker node CPU usage, per nanosecond	Must support	Must support	TBD	
4.2.5	i. pm. 002	Monitor pod CPU usage, per nanosecond	Must support	Must support	TBD	
4.2.5	i. pm. 003	Monitor worker node CPU utilisation (%)	Must support	Must support	TBD	
4.2.5	i. pm. 004	Monitor pod CPU utilisation	Must support	Must support	TBD	
4.2.5	i. pm. 005	Measure external storage IOPs	Must support	Must support	TBD	
4.2.5	i. pm. 006	Measure external storage throughput	Must support	Must support	TBD	
4.2.5	i. pm. 007	Measure external storage capacity	Must support	Must support	TBD	

Cloud Infrastructure Software Profile Requirements

Reference Model Section	Reference	Description	Requirement for Basic Profile	Requirement for Network Intensive Profile	BM RA Notes	OPNFV Intel Lab
5.2.1	infra.com.cfg.001	CPU allocation ratio	1:1	1:1	See above	
5.2.1	infra.com.cfg.002	NUMA awareness	Must support	Must support	See above	

5.2.1	infra.com.cfg.003	CPU pinning capability	Must support	Must support	See above	
5.2.1	infra.com.cfg.004	Huge Pages	Must support	Must support	See above	
5.2.2	infra.stg.cfg.002	Storage Block	Must support	Must support	TBD	
5.2.2	infra.stg.cfg.003	Storage with replication	Not required	Must support	TBD	
5.2.2	infra.stg.cfg.004	Storage with encryption	Must support	Must support	TBD	
5.2.2	infra.stg.acc.cfg.001	Storage IOPS oriented	Not required	Must support	TBD	
5.2.2	infra.stg.acc.cfg.002	Storage capacity oriented	Not required	Not required	TBD	
5.2.3	infra.net.cfg.001	IO virtualisation using virtio1.1	Must support(1)	Must support(1)	Considering (1), this should be fine with SR-IOV Device Plugin and CNI	
5.2.3	infra.net.cfg.002	The overlay network encapsulation protocol needs to enable ECMP in the underlay to take advantage of the scale-out features of the network fabric.(2)	Must support VXLAN, MPLSoUDP, GENEVE, other	No requirement specified	TBD	TBD
5.2.3	infra.net.cfg.003	Network Address Translation	Must support	Must support	(??) Is this possible in K8s by default	
5.2.3	infra.net.cfg.004	Security Groups	Must support	Must support	TBD	
5.2.3	infra.net.cfg.005	SFC support	Not required	Must support	TBD	
5.2.3	infra.net.cfg.006	Traffic patterns symmetry	Must support	Must support	TBD	
5.2.3	infra.net.acc.cfg.001	vSwitch optimisation	Not required	Must support DPDK(3)	(??) Part of BM RA (with additional configuration needed?)	
5.2.3	infra.net.acc.cfg.002	Support of HW offload	Not required	Must support SmartNic	See above	
5.2.3	infra.net.acc.cfg.003	Crypto acceleration	Not required	Must support	See above	
5.2.3	infra.net.acc.cfg.004	Crypto Acceleration Interface	Not required	Must support	See above	

(1) Workload Transition Guidelines. might have other interfaces (such as SR-IOV VFs to be directly passed to a VM or a Pod) or NIC-specific drivers on guest machines transiently allowed until mature enough solutions are available with a similar efficiency level (for example regarding CPU and energy consumption).

(2) In Kubernetes based infrastructures network separation is possible without an overlay (e.g.: with IPVLAN)

(3) This feature is not applicable for Kubernetes based infrastructures due to lack of vSwitch however workloads need access to user space networking solutions.

Cloud Infrastructure Hardware Profile Requirements

Reference Model Section	Reference	Description	Requirement for Basic Profile	Requirement for Network Intensive Profile	BM RA Notes	OPNFV Intel Lab
-------------------------	-----------	-------------	-------------------------------	---	-------------	-----------------

5.4.1	infra.hw.cpu.cfg.001	Number of CPU (Sockets)	2	2	[OK] POD 19: x2 [OK] POD 21: x2
5.4.1	infra.hw.cpu.cfg.002	Number of Cores per CPU	20	20	[OK] POD 19: 22 [OK] POD 21: 20
5.4.1	infra.hw.cpu.cfg.003	NUMA	Not required	Must support	Supported by default (?)
5.4.1	infra.hw.cpu.cfg.004	Simultaneous Multithreading /Hyperthreading (SMT/HT)	Must support	Must support	Depends on BIOS configuration (can be enabled)
5.4.1	infra.hw.cac.cfg.001	GPU	Not required	Not required	Depends on HW configuration
5.4.2	infra.hw.stg.hdd.cfg.001	Local Storage HDD	No requirement specified	No requirement specified	[OK] POD 19: 3TB HDD [!] POD 21: Only SSD
5.4.2	infra.hw.stg.ssd.cfg.002	Local Storage SSD	Should support	Should support	[OK] POD 19: 190GB [OK] POD 21: 800GB
5.4.3	infra.hw.nic.cfg.001	Total Number of NIC Ports available in the host	4	4	OK, assuming no requirement on port speed (1/10Gbps)
5.4.3	infra.hw.nic.cfg.002	Port speed specified in Gbps (minimum values)	10	25	[OK] Basic Profile [!] Network Intensive Profile
5.4.3	infra.hw.pci.cfg.001	Number of PCIe slots available in the host	8	8	TBD
5.4.3	infra.hw.pci.cfg.002	PCIe speed	Gen 3	Gen 3	Should be OK
5.4.3	infra.hw.pci.cfg.003	PCIe Lanes	8	8	Should be OK
5.4.3	infra.hw.nac.cfg.001	Cryptographic Acceleration	Not required	Optional	Depends on HW configuration
5.4.3	infra.hw.nac.cfg.002	A SmartNIC that is used to offload vSwitch functionality to hardware	Not required	Optional(1)	Not likely, but depends on HW configuration
5.4.3	infra.hw.nac.cfg.003	Compression	No requirement specified	No requirement specified	TBD

(1) There is no vSwitch in case of containers, but a SmartNIC can be used to offload any other network processing.