

Support of OVS-DPDK in Airship

Project Technical Lead: Georg Kunz. Elected 1/17/19.

Project Committer Details:

Initial Committers for a project will be specified at project creation. Committers have the right to commit code to the source code management system for that project.

A Contributor may be promoted to a Committer by the project's Committers after demonstrating a history of contributions to that project.

Candidates for the project's Project Technical Leader will be derived from the Committers of the Project. Candidates must self nominate by marking "Y" in the Self Nominate column below by Jan. 16th. Voting will take place January 17th.

Only Committers for a project are eligible to vote for a project's Project Technical Lead.

Please see [Akraio Technical Community Document](#) section 3.1.3 for more detailed information.

Committer	Committer Company	Committer Contact Info	Committer Bio	Committer Picture	Self Nominate for PTL (Y/N)
Andrew Wilkinson	Ericsson	andrew.wilkinson@ericsson.com			N
Georg Kunz	Ericsson	georg.kunz@ericsson.com			Y
Cheng Li	Intel	cheng1.li@intel.com			
Rihab Banday	Ericsson	rihab.banday@ericsson.com			

Use Case Details:

Feature project proposers: Ericsson, Nokia and Radysis (confirm) and (others - confirm).

The following BP proposals require support of OVS-DPDK in Airship:

Network Cloud : OVS-DPDK Unicycle Dell Blueprint Proposal

Radio Edge Cloud

Edge Video Processing

(any other dependent BPs?)

Initial list of high-level working items

The table below provides a list of high-level work items required to enable support for OVS-DPDK in Airship. This list is not considered to be complete but a starting point for design discussions in the feature project. Please feel free to add / modify / extend work items.

Task name	Description	Airship component	Implementation	Upstream reference
create openvswitch agent chart	create helm chart for openvswitch agent	openstack-helm	in place and being deployed when OVS is enabled	
create ovs-dpdk chart	create helm chart for openvswitch dpdk container	openstack-helm	Extend existing openvswitch chart with config parameters for DPDK.	https://review.openstack.org/#/c/626894/
deploy neutron openvswitch agent	ensure chart of openvswitch agent is deployed	treasuremap	In place and being deployed when OVS is enabled	

DPDK host config: enable 1G hugepages	modify kernel cmdline to enable 1G hugepages (hugepagez=1G hugepages=8 iommu=pt intel_iommu=on)	drydock	<p>Already available:</p> <p>Define the number of available hugepages in the node's HardwareProfile: https://github.com/openstack/airship-drydock/blob/master/python/tests/yaml_samples/deckhand_fullsite.yaml#L427</p> <p>Define the kernel parameters for enabling hugepages in the node's BaremetalNode configuration (kernel_param section): https://github.com/openstack/airship-drydock/blob/master/python/tests/yaml_samples/deckhand_fullsite.yaml#L366</p>	https://storyboard.openstack.org/#!/story/2004790
DPDK host config: mount hugepages	mount hugepages into local file system on system boot (hardcoded mount point)	Divingbell	<p>Two alternatives available:</p> <ul style="list-style-type: none"> utilize auto-mount capabilities of Ubuntu https://www.apt-browse.org/browse/ubuntu/xenial/main/i386/systemd/229-4ubuntu4/file/lib/systemd/system/dev-hugepages.mount <p>Issue: does not allow fine-grained control of mount-options (e.g. specify the size of the hugepages if there are multiple available)</p> <ul style="list-style-type: none"> Deploy Divingbell daemonset on compute nodes with a given chart configuration (values.yaml). <p>https://airship-divingbell.readthedocs.io/en/latest/#mounts</p>	https://storyboard.openstack.org/#!/story/2004790
DPDK host config: make hugepage mount point a config option	make mount point config option for e.g. use by helm charts	openstack-helm	config option of helm chart (see patchset). Need to figure out how to pass overrides to chart.	https://review.openstack.org/#/c/626894/
specify PCI IDs of NICs for use by DPDK	specify in site config which PCI IDs (NICs) should be decided to DPDK		Add parameters to OVS helm chart	https://review.openstack.org/#/c/626894/
DPDK host config: install DPDK kernel modules and tools on host OS	either install host OS dpdk package or build from source	drydock or divingbell	depends on DPDK driver we want to use: igb_uio, uio_generic_pci or vfio-pci. Only for igb_uio a custom built kernel module is needed.	https://review.openstack.org/#/c/626894/
DPDK host config: load DPDK kernel module in host OS	load dpdk kernel modules uio and igb_uio during host boot-up	drydock or divingbell	extend existing openvswitch-vswitchd init container	https://review.openstack.org/#/c/626894/
DPDK host config: bind NICs to DPDK	use dpdk-devbind to bind specified NICs to DPDK	drydock or divingbell	<p>look into re-using and/or adapting a tool used by kolla-ansible</p> <p>https://github.com/openstack/kolla-ansible/blob/master/tools/ovs-dpdkctl.sh</p>	https://review.openstack.org/#/c/626894/
DPDK host config: enable hugepage support for kubelet	enable hugepage support for k8s kubelet via feature gate option	promenade	hugepages is a beta feature since K8s 1.10 and enabled by default.	https://storyboard.openstack.org/#!/story/2004790
ensure communication between OVS agent and OVS	ensure common socket configuration	openstack-helm	already in place	
adapt OVS bridge configuration for OVS setup	<p>Work items:</p> <ul style="list-style-type: none"> create bridges with datapath type netdev add physical interface to physical bridge (br-phy) 	openstack-helm	<p>look into re-using and/or adapting a tool used by kolla-ansible</p> <p>https://github.com/openstack/kolla-ansible/blob/master/tools/ovs-dpdkctl.sh</p>	https://review.openstack.org/#/c/626894/
adapt Neutron (ovs-agent) configuration	Adapt neutron.conf and ml2 plugin config	openstack-helm	extend ml2 plugin configuration of neutron in openstack-helm	https://review.openstack.org/#/c/643284/
adapt IP address assignment to OVS DPDK bridge	<p>make sure that correct IPs get assigned to the OVS bridges running DPDK.</p> <p>Every dpdk bridge needs a separate IP address for the tunnel endpoints.</p>	openstack-helm		

create docker image with dpdk-enabled OVS	update the image build scripts of openstack-helm to include DPDK in the OVS image	openstack-helm-images	<p>openstack-helm-image repo:</p> <p>https://github.com/openstack/openstack-helm-images/tree/master/openvswitch</p> <p>Alternatives (kept just for reference)</p> <p>Kolla images: https://hub.docker.com/r/kolla/ubuntu-binary-ovsdpdk/ or https://hub.docker.com/r/kolla/ubuntu-source-ovsdpdk/</p>	https://review.openstack.org/#/c/650152/
update nova and neutron images with newer version of OVS	OVS in current images is outdated	openstack-helm-images		https://review.opendev.org/#/c/651060/
update site configuration to deploy ovs-dpdk	create a site configuration which actually deploys ovs-dpdk as data plane	treasuremap	enable openvswitch chart group	https://review.openstack.org/#/c/634207/
add LAG support to DPDK configuration	configure LAG support on DPDK NICs in OVS	openstack-helm		

Legend:

done / available
ongoing
blocking