

OpenStack: How can VXLAN traffic be offloaded to InfiniBand ?

How can VXLAN traffic be offloaded to InfiniBand ?

If the hardware on which a Bright OpenStack cluster is running has InfiniBand support, then the InfiniBand network can be used to offload the VXLAN traffic. In other words, the tenant (project) networks can run over the InfiniBand fabric instead of over ethernet.

This has been tested on Bright Openstack 7.3

Requirements:

1. Bright OpenStack 7.3
2. InfiniBand HCA installed on the hardware of the OpenStack cluster
3. IPoIB configured.

Steps :

1. Remove the current VXLAN network

```
#csh
```

```
%network
```

```
%use vxlanhostnet
```

```
%clear openstacknetworktype
```

```
%clear baseaddress
```

```
%clear broadcastaddress
```

```
%clear openstackphysicalnetworkname
```

```
%commit
```

OpenStack: How can VXLAN traffic be offloaded to InfiniBand ?

```
%network
```

```
%use ibnet <----assuming ibnet is the name of the current infiniband network
```

```
%set openstacknetworktype vxlan\ host
```

```
%set openstackvlanrange 1:500000 <---- can be anything, from 1 to maximum  
allowed VXLANs
```

```
%set openstackphysicalnetworkname phyvxlanhostnet
```

```
%set baseaddress <required base address>
```

```
%set broadcastaddress <required broadcast address>
```

```
%clear gateway
```

```
%commit
```

A proper IP address must be configured for all the InfiniBand interfaces of the OpenStack cluster. The addresses must be assigned by the Bright Cluster Manager to the InfiniBand network, by setting the network to `ibnet`. For example:

```
%device use node001
```

```
%interfaces
```

```
%user ib0
```

```
%set ip <ip address>
```

```
%set network ibnet
```

```
%commit
```

The OpenStack cluster, including the head node, should then be rebooted:

OpenStack: How can VXLAN traffic be offloaded to InfiniBand ?

```
%device
```

```
%reboot -n <all the nodes>
```

After the cluster is up and running again, it is then possible to create project networks. For example (using cmsh rather than neturon commands):

```
#cmsh
```

```
%openstack
```

```
%networks
```

```
%add vxlan-network1
```

```
%set networktype vxlan
```

```
%commit
```

```
%subnets
```

```
%add subnets-vxlan1
```

```
%set network vxlan-network1
```

```
%set cidr 10.100.0.0/16
```

```
%allocationpools
```

```
%add 10.100.1.2
```

```
%set end 10.100.255.254
```

```
%commit
```

Unique solution ID: #1326

Author: ahmed

Last update: 2016-10-14 16:06