

- PCI Passthrough / SR-IOV** 3
- PCI Passthrough** 3
- Openstack SR-IOV** 3
- 3
- Mellanox CONNECTX-4/CONNECTX-5/CONNECTX-6 SR-IOV 3
- Intel X550 SR-IOV 4
- SR-IOV** 4

PCI Passthrough / SR-IOV

PCI Passthrough

nova.conf 가 Passthrough []

```
[filter_scheduler]
enabled_filters = AvailabilityZoneFilter, ComputeFilter,
ComputeCapabilitiesFilter, ImagePropertiesFilter,
ServerGroupAntiAffinityFilter, ServerGroupAffinityFilter,
PciPassthroughFilter
available_filters = nova.scheduler.filters.all_filters

[pci]
passthrough_whitelist = [{ "devname": "ens6", "physical_network": "physnet2"
}, { "vendor_id": "10de", "product_id": "13c2" }]
alias = { "vendor_id":"10de", "product_id":"13c2", "device_type":"type-PF",
"name":"GTX970" }
```

Openstack SR-IOV

- <https://docs.openstack.org/neutron/latest/admin/config-sriov.html>

QoS(Quality of Service) max_burst_kbps(max_kbps)

max_kbps Mbps

SR-IOV

neutron.conf

```
[securitygroup]
firewall_driver = noop
```

SR-IOV OpenStack Dashboard(horizon) CLI API

SR-IOV

Mellanox CONNECTX-4/CONNECTX-5/CONNECTX-6 SR-IOV

- <https://enterprise-support.nvidia.com/s/article/HowTo-Configure-SR-IOV-for-ConnectX-4-Connect>

X-5-ConnectX-6-with-KVM-Ethernet

Intel X550 SR-IOV

- <https://forum.proxmox.com/threads/enabling-sr-iov-for-intel-nic-x550-t2-on-proxmox-6.56677/>

SR-IOV

1. IOMMU

```
grubby --update-kernel=ALL --args="intel_iommu=on iommu=pt"
grubby --update-kernel=ALL --args="amd_iommu=on iommu=pt"

dracut -f -v
```

2. SR-IOV device

```
echo "echo '8' > /sys/class/net/ens6/device/sriov_numvfs" >> /etc/rc.local
```

/etc/kolla/config/nova.conf 가

- GPU Passthrough

```
[pci]
passthrough_whitelist = { "vendor_id": "10de", "product_id":"20f1" }
```

- SR-IOV

```
[pci]
passthrough_whitelist = { "devname": "ens6", "physical_network": "physnet2"
}
```

| | Compute | ens6 | VF가 |
|----------|---------|------|------|
| physnet2 | | . | |
| | id | , | PF |
| | | | . VF |

```
passthrough_whitelist = { "vendor_id": "8086", "product_id":"1563",
"physical_network": "physnet2" }
```

Passthrough

```
[pci]
passthrough_whitelist = [{ "devname": "ens6", "physical_network": "physnet2"
}, { "vendor_id": "10de", "product_id": "13c2" }]
alias = { "vendor_id":"10de", "product_id":"13c2", "device_type":"type-PF",
"name":"GTX970" }
```

3. Launching instances with SR-IOV ports

```
openstack network create --provider-physical-network sriovnet1 \
  --provider-network-type vlan --provider-segment 319 \
  sriov-net

openstack subnet pool create --share --pool-prefix 10.33.19.0/24 \
  --default-prefix-length 24 shared-default-subnetpool-v4

openstack subnet create --network sriov-net \
  --subnet-pool shared-default-subnetpool-v4 \
  sriov-subnet

net_id=$(openstack network show sriov-net -c id -f value)

openstack port create --network $net_id --vnic-type direct sriov-port

port_id=$(openstack port show sriov-port -c id -f value)

openstack server create --availability-zone sriov --flavor c1_r2_d30 --image
rhel8 \
  --nic port-id=$port_id test-sriov
```

From:

<https://atl.kr/dokuwiki/> - AllThatLinux!

Permanent link:

https://atl.kr/dokuwiki/doku.php/pci_passthrough_sr-iov?rev=1701673461

Last update: **2023/12/04 07:04**

