oVirt – Neutron Integration

July 2013
Mike Kolesnik

Agenda

oVirt

- oVirt network configuration
- Neutron overview
  - Integration benefits
  - External providers
  - Neutron as an external provider
- Under the hook
- Future work

oVirt

oVirt Network Configuration

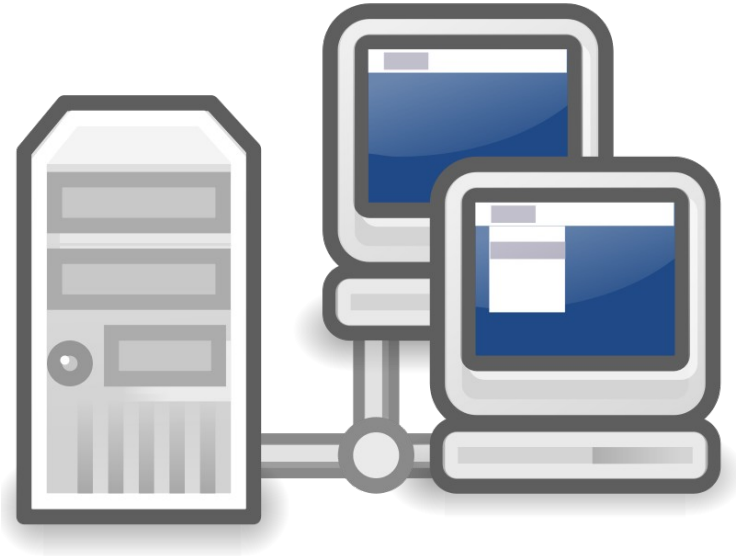

Network View

- Network - a logical entity that represents a layer 2 broadcast domain
- Defined within the scope of a data center

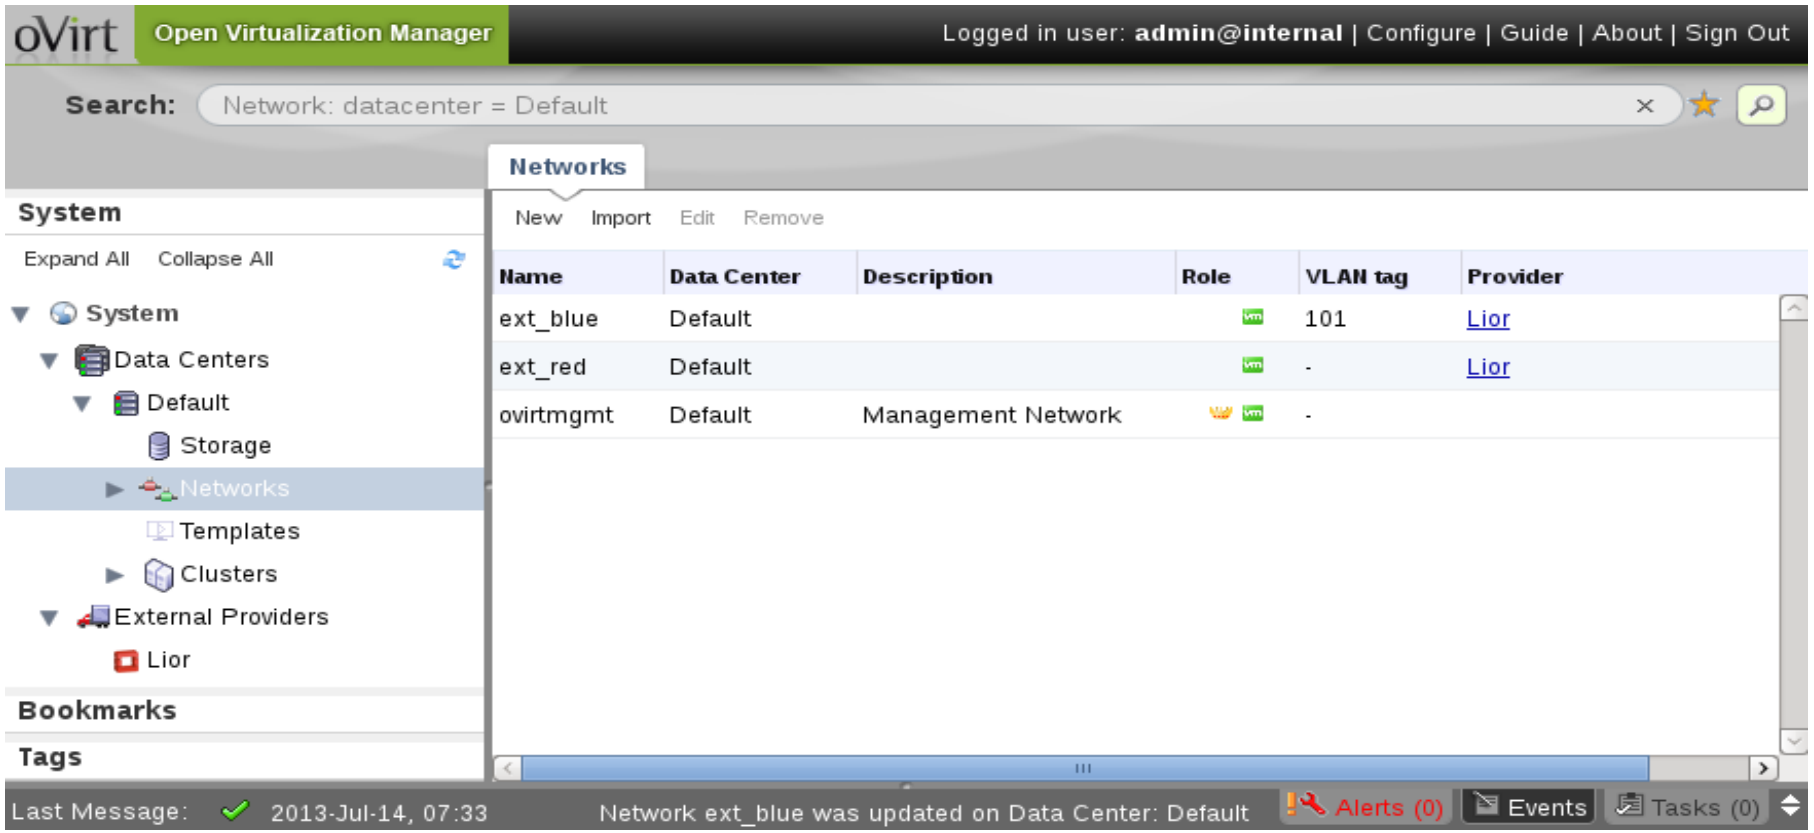

Adding a New Network

- Give it a name
- Define the network properties (VLAN, MTU, VM/Non-VM)

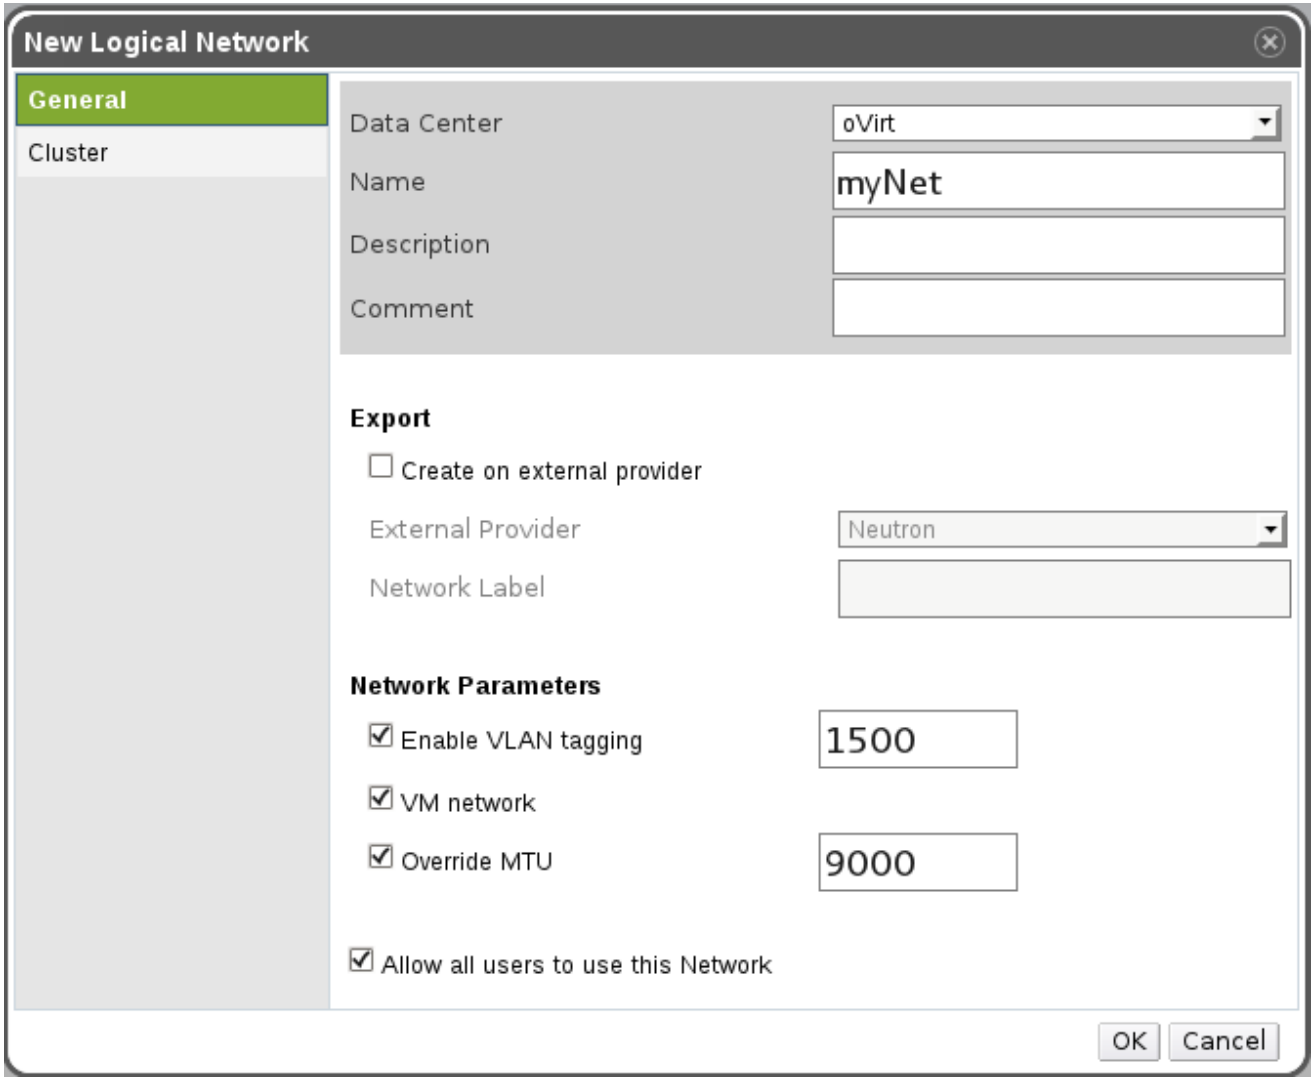

Adding a New Network

- Make the network available in the selected clusters.

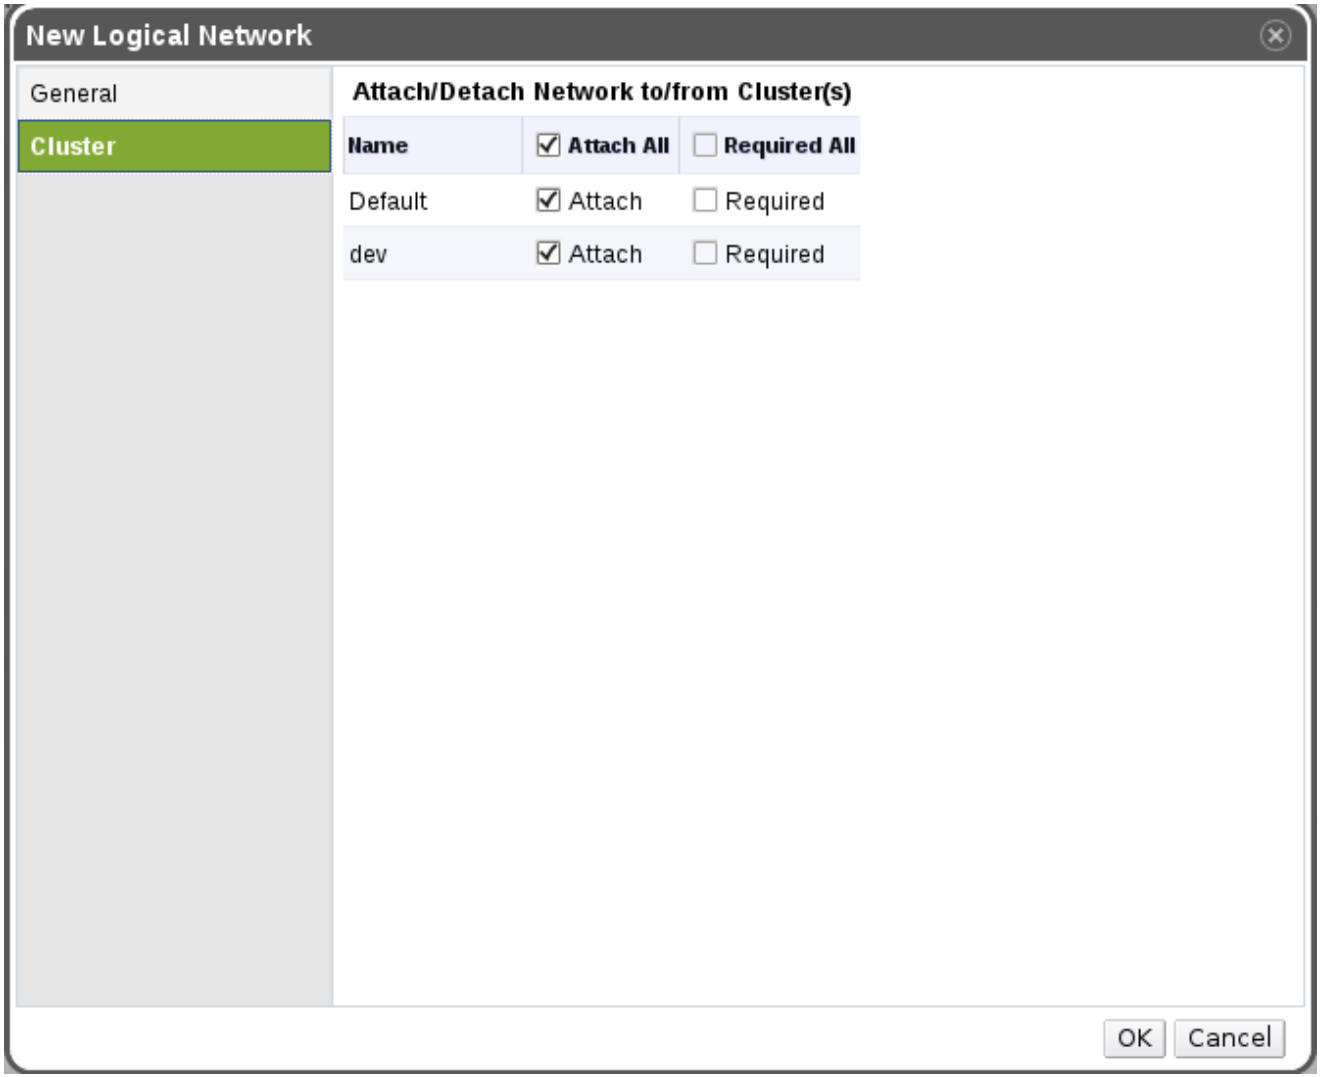

Host Level Configuration

- Optional vs. required networks
- Host level configuration:

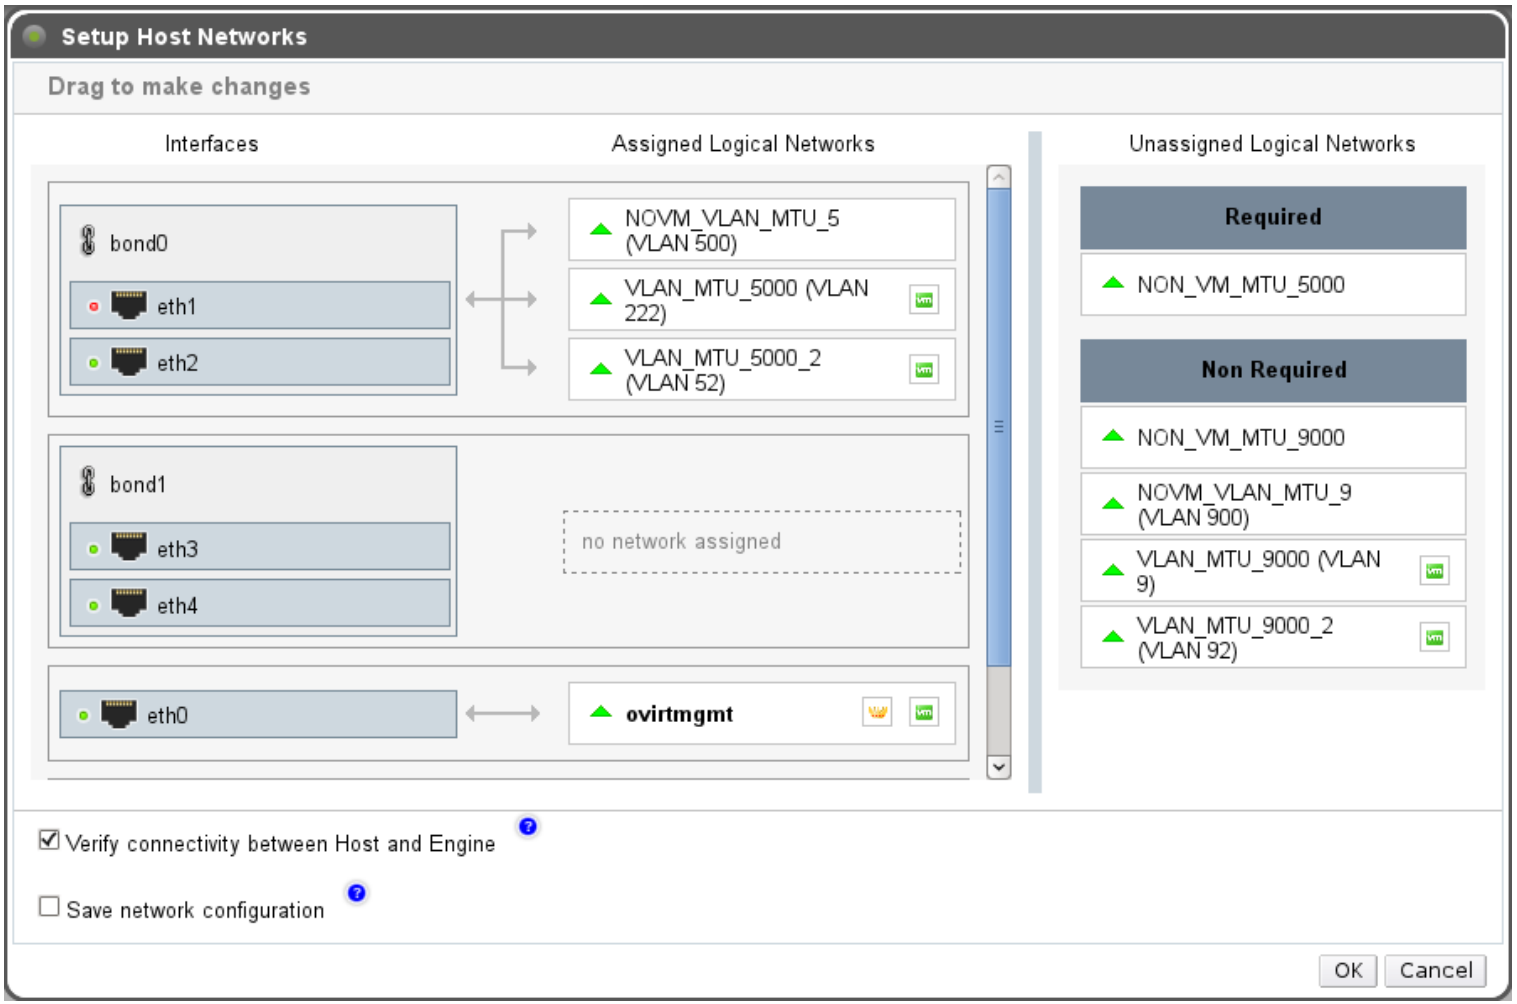

Adding Network to a vNIC

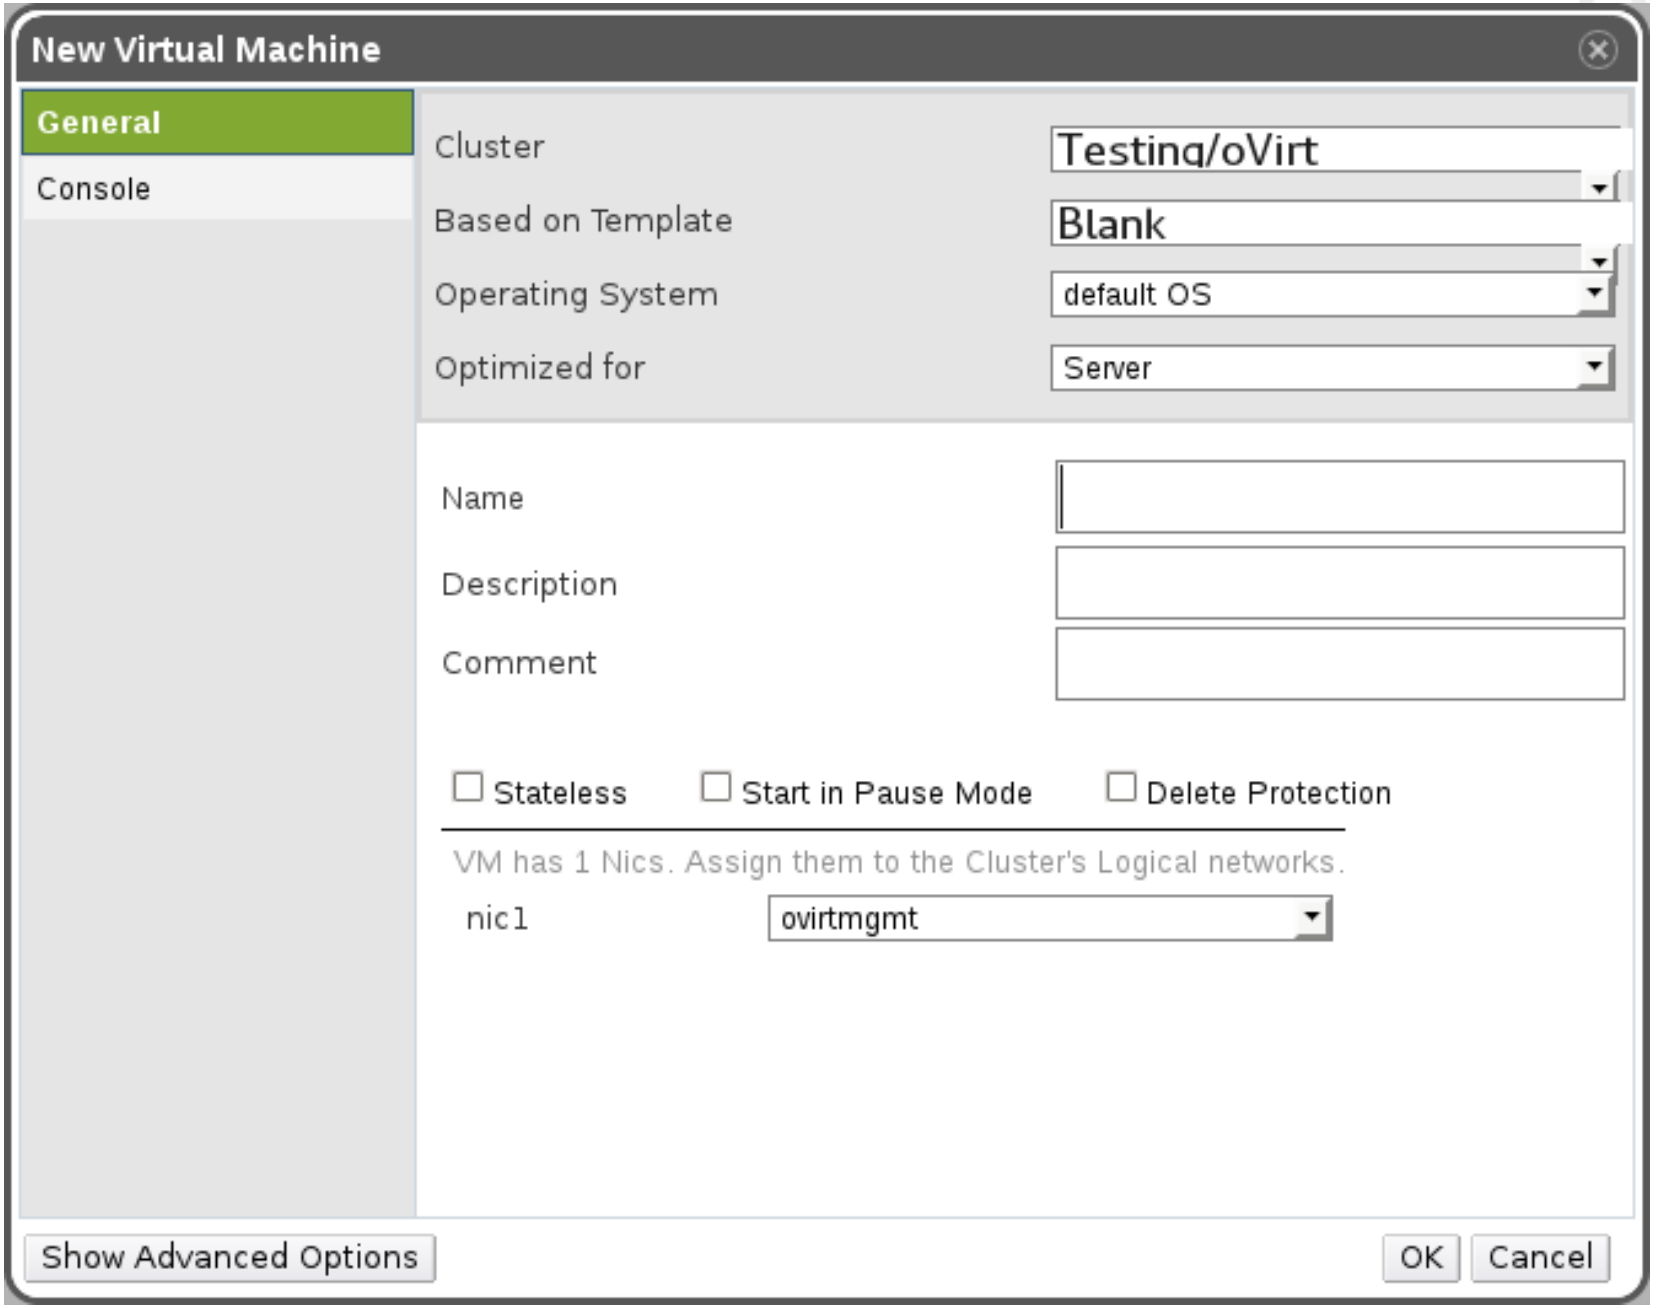

Supported Configuration - Linux Bridge

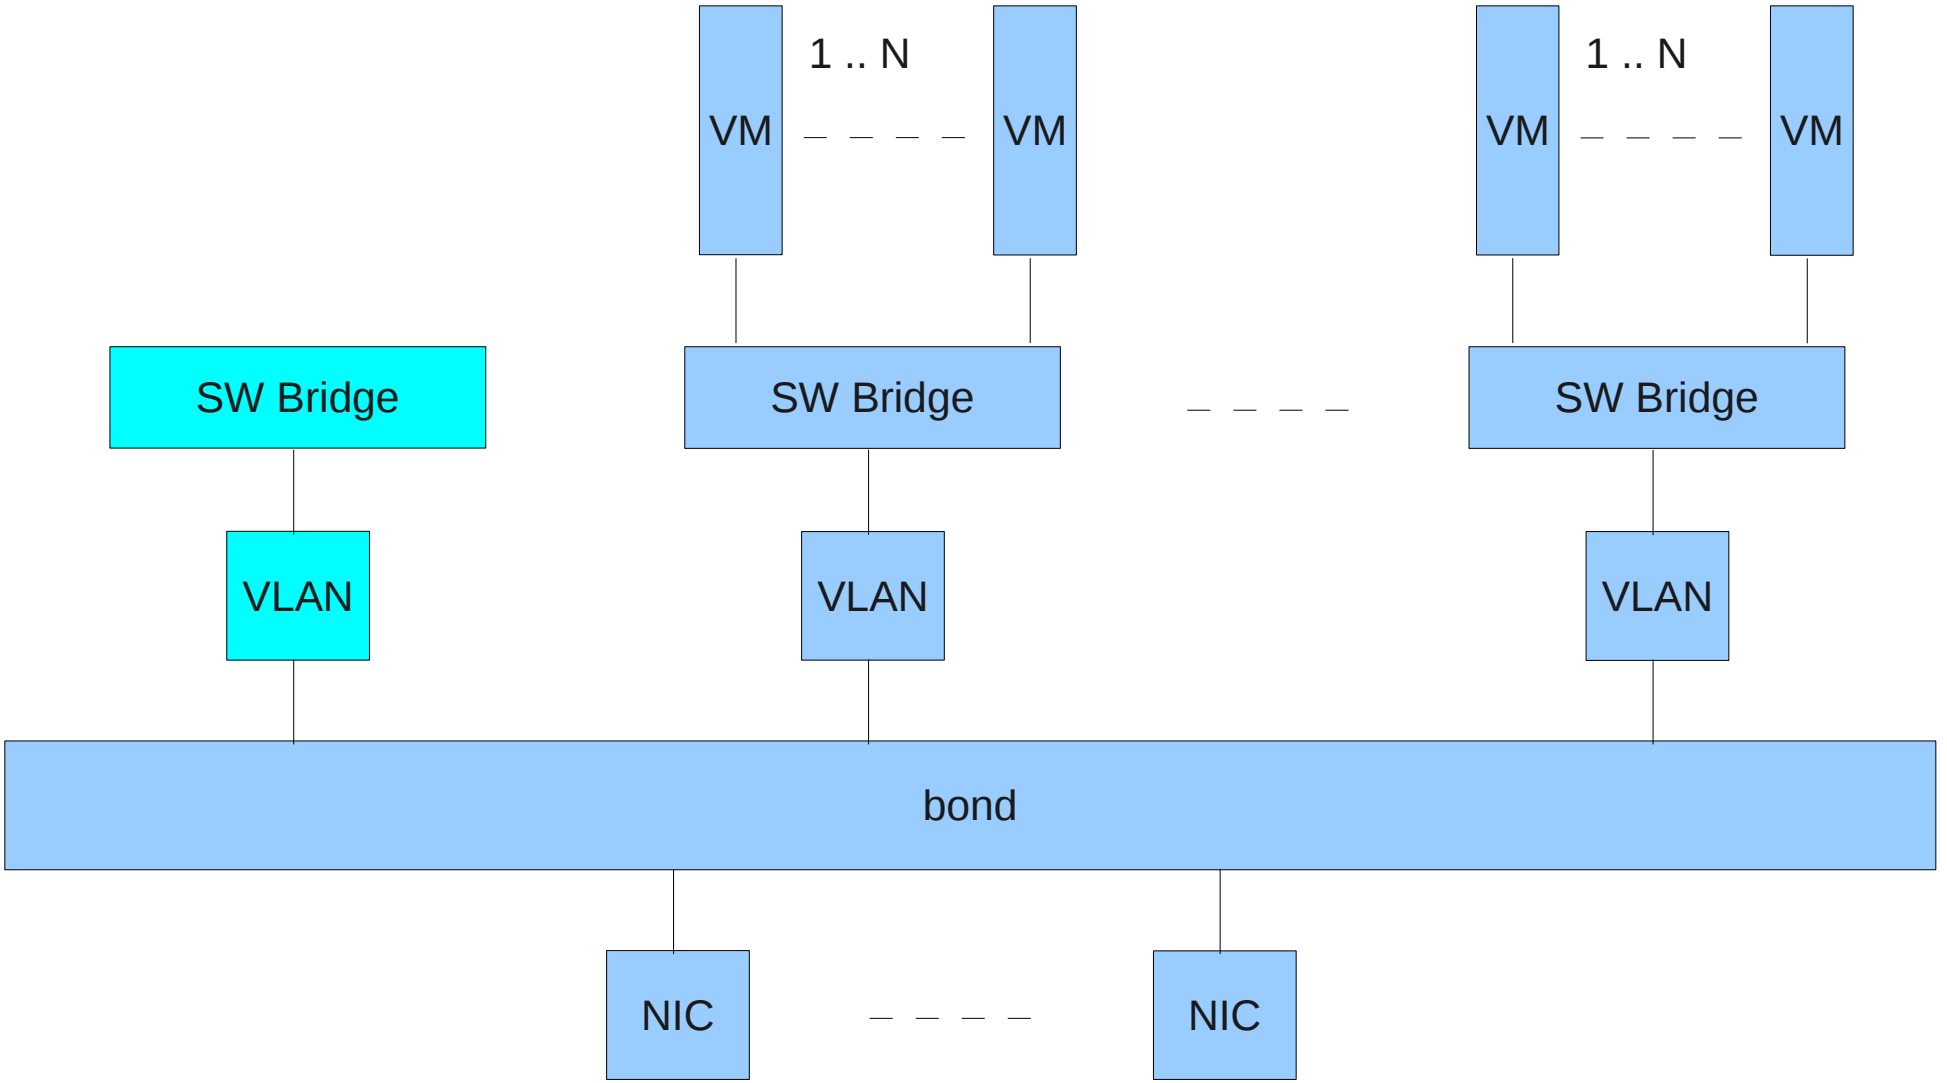

oVirt

Neutron Overview

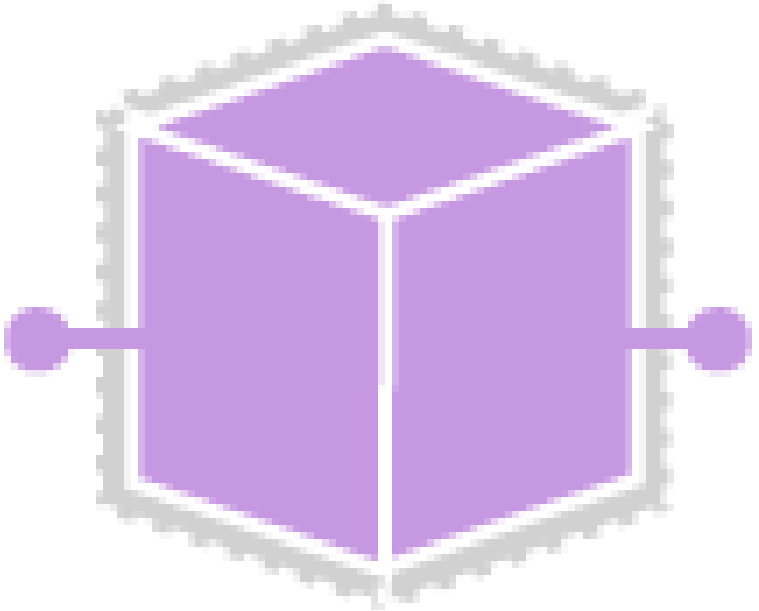

OpenStack Networking - Neutron

oVirt

- Neutron provides network connectivity-as-a-service
- It offers a plug-in architecture designed to support various network technologies through vendor specific plug-ins and API extensions
- Exposes REST API for accessing the service
- Available plugins: Linux Bridge, OVS, Cisco Nexus, NVP, Ryu, NEC, etc...

Neutron high level architecture

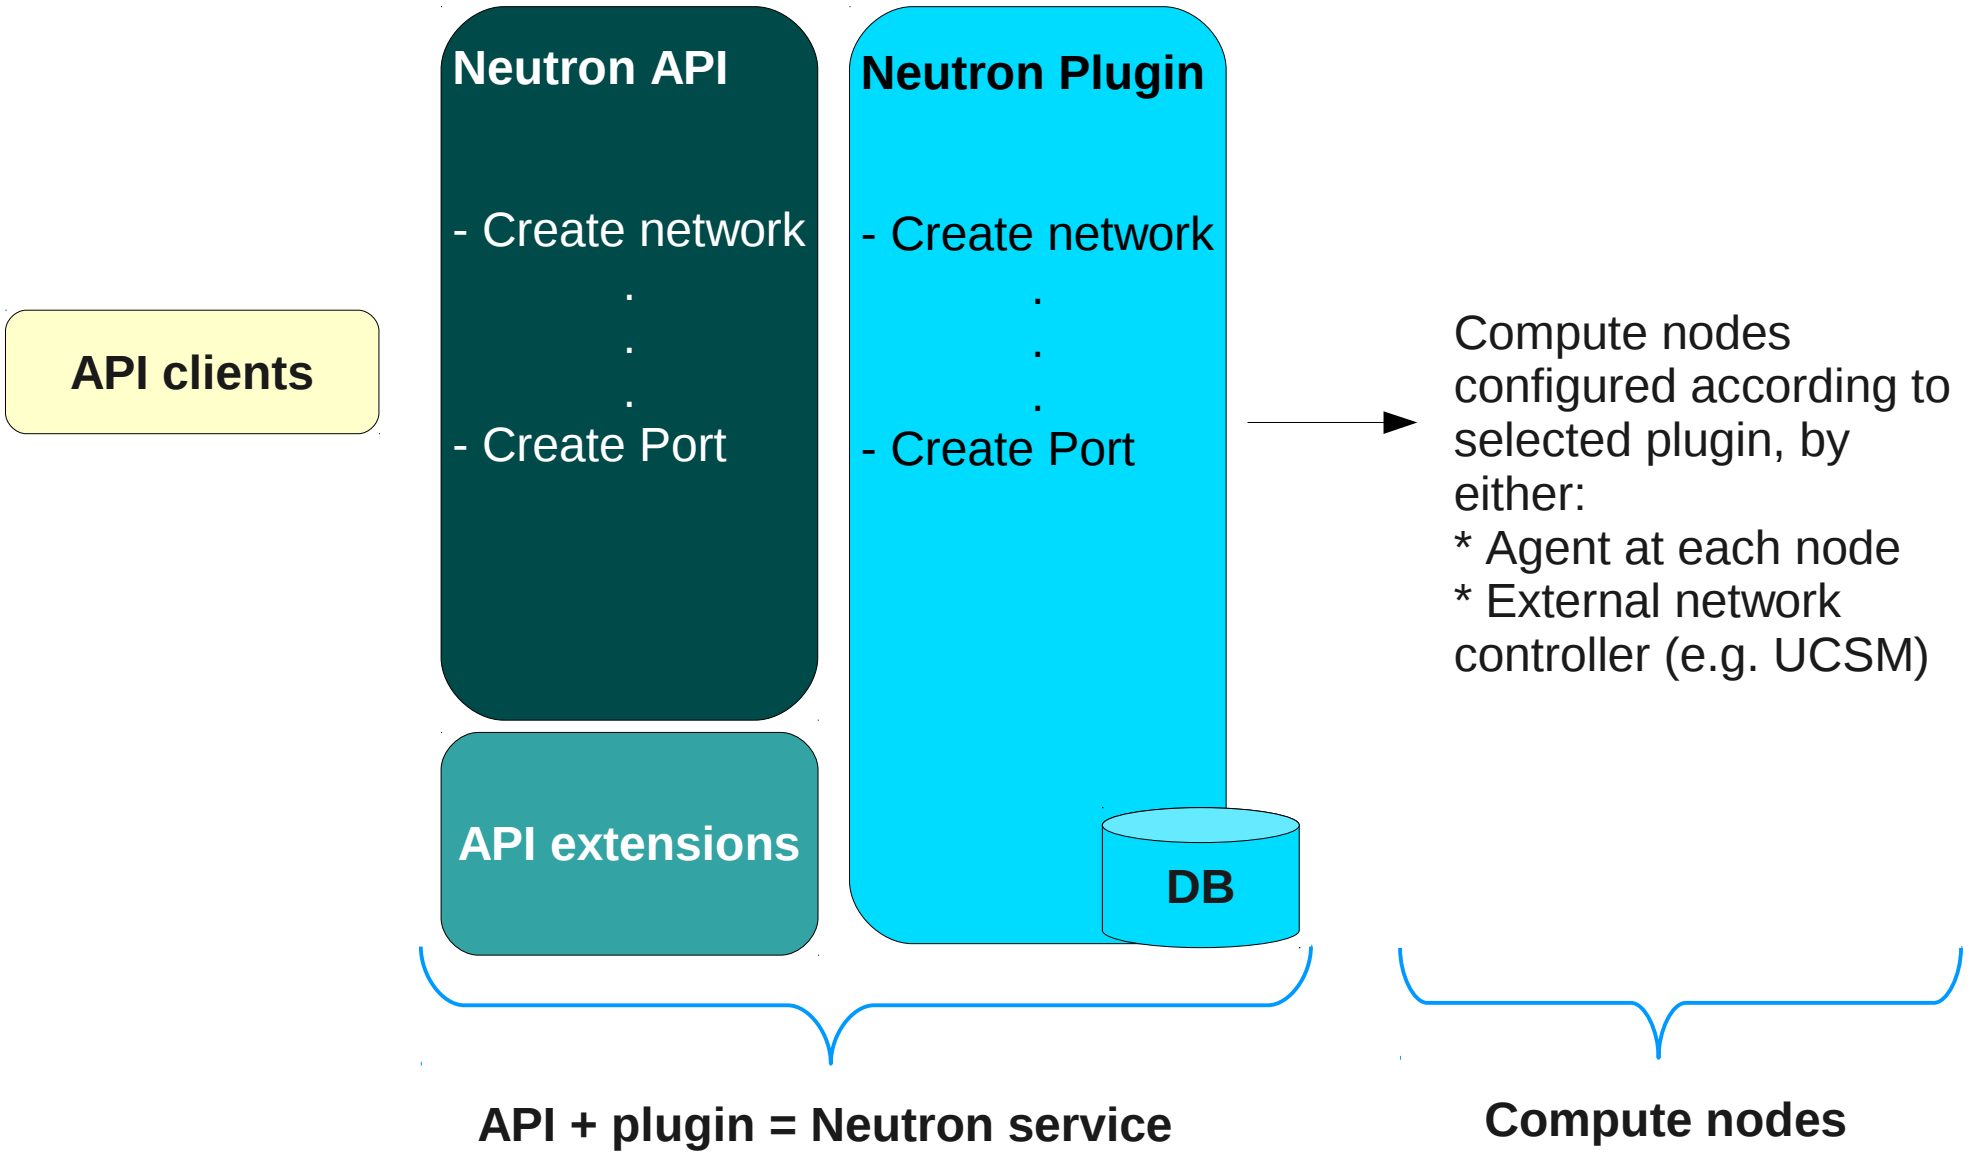

Integration Benefits

oVirt

- Add support in oVirt for the various network technologies provided via Neutron plugins
- Leverage L3 services modeled in Neutron
- Enjoy both worlds:
  - Neutron for managing VM networks
  - oVirt for managing infrastructure networks (Migration network, storage network etc.)
- Neutron networks are exposed side by side with oVirt networks which allows the user to use oVirt mature implementation for network configuration

oVirt

Integration bits

External Providers

- An external product that can be used to provide resources for oVirt
- Resources that can be provided: hosts, networks, etc..
- Configure once, use everywhere

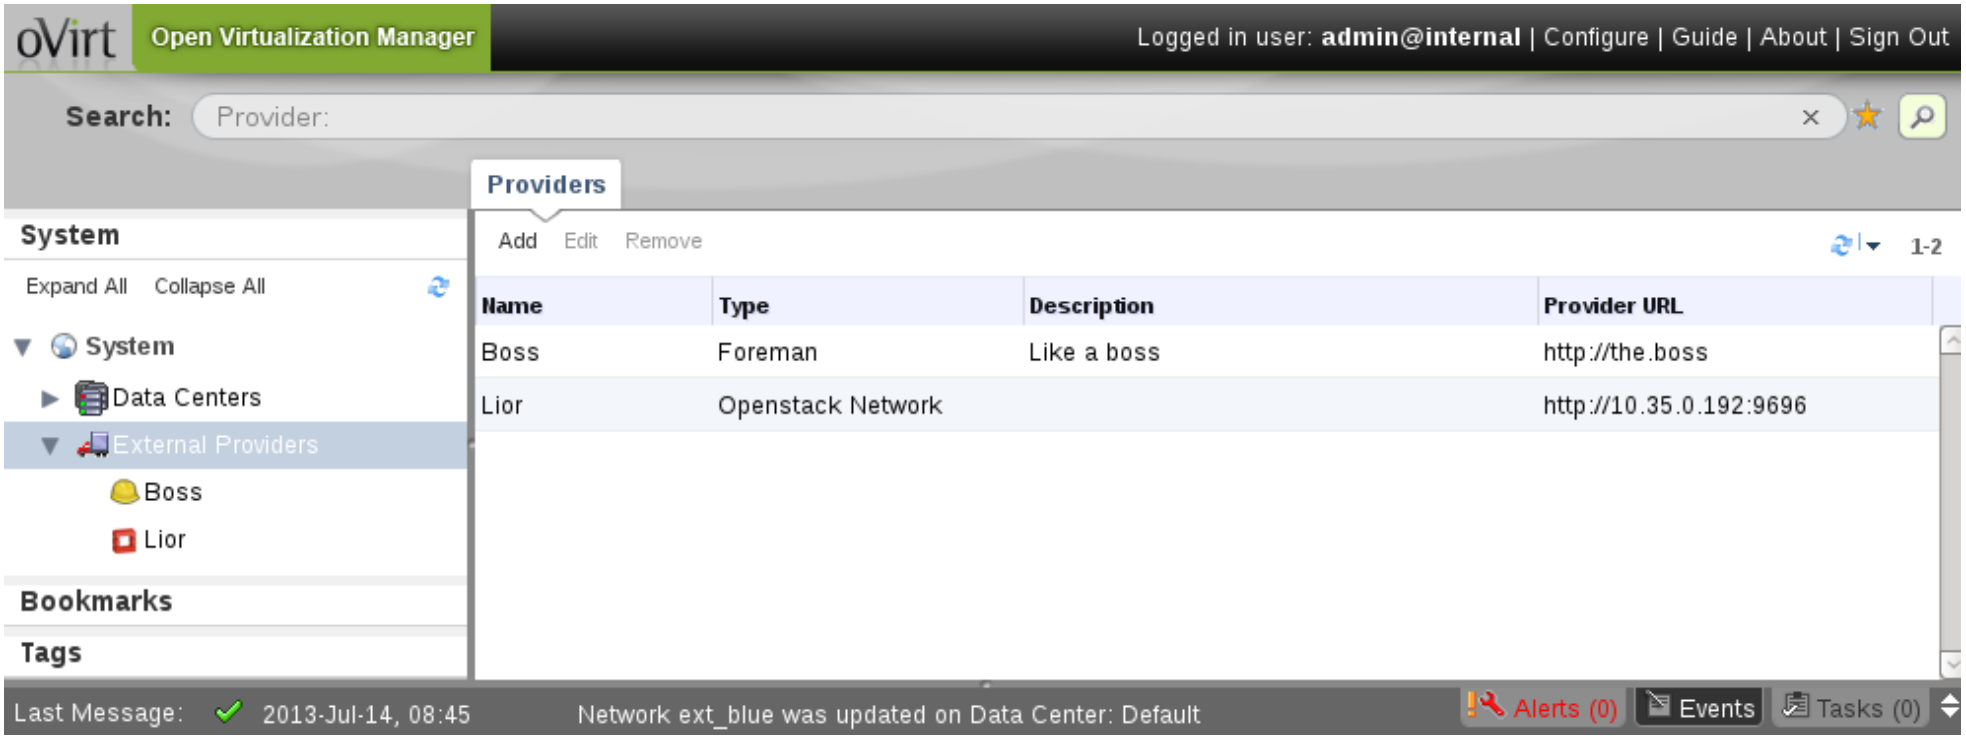

External Network Provider

- ***External network provider*** - an independent network manager which collaborates with oVirt by implementing a predefined API
- ***Internal network –*** provisioned in oVirt
- ***External network*** – provisioned by an external network provider and consumed within oVirt
  - External networks can be discovered in oVirt and then can be used within oVirt (currently in VMs)
  - User can configure permissions on external networks once they are imported, like they do for internal networks.

The Neutron External Provider

- Created as an external network provider
- Can be deployed with the user choice of plug-in
- Can be used in either of the flavors:
  - Neutron centric – Existing Neutron installation, oVirt is jusr a “user” of some of the networks.
  - oVirt centric – Neutron is an implementation detail, the networks are actually managed in oVirt.

How to Use Neutron in oVirt?

- It's simple! Just follow these few steps:
  1. Install a Neutron instance
  2. Add the instance as an external provider
  3. Add networks on the provider:
     - 3.1. Import networks
     - 3.2. Add a new network on the provider
  4. Install host with the provider's agent
  5. Use the network in a VM's NIC
  6. Run the VM

Step 1: Installing Neutron

oVirt

- Install Neutron service and configure your choice of plugin
- Install Keystone
  - Configure Keystone for the Neutron service

Step 1: Installing Neutron

oVirt

- Install Neutron service
- Install Keystone (configure for the Neutron instance)
- Configure Neutron:
  - Configure authentication for keystone
  - Configure qpid messaging bus
- Install one of the supported plug-ins:
  - Linux Bridge
  - Open vSwitch
- Configure the plugin

Step 2: Adding a Neutron Provider

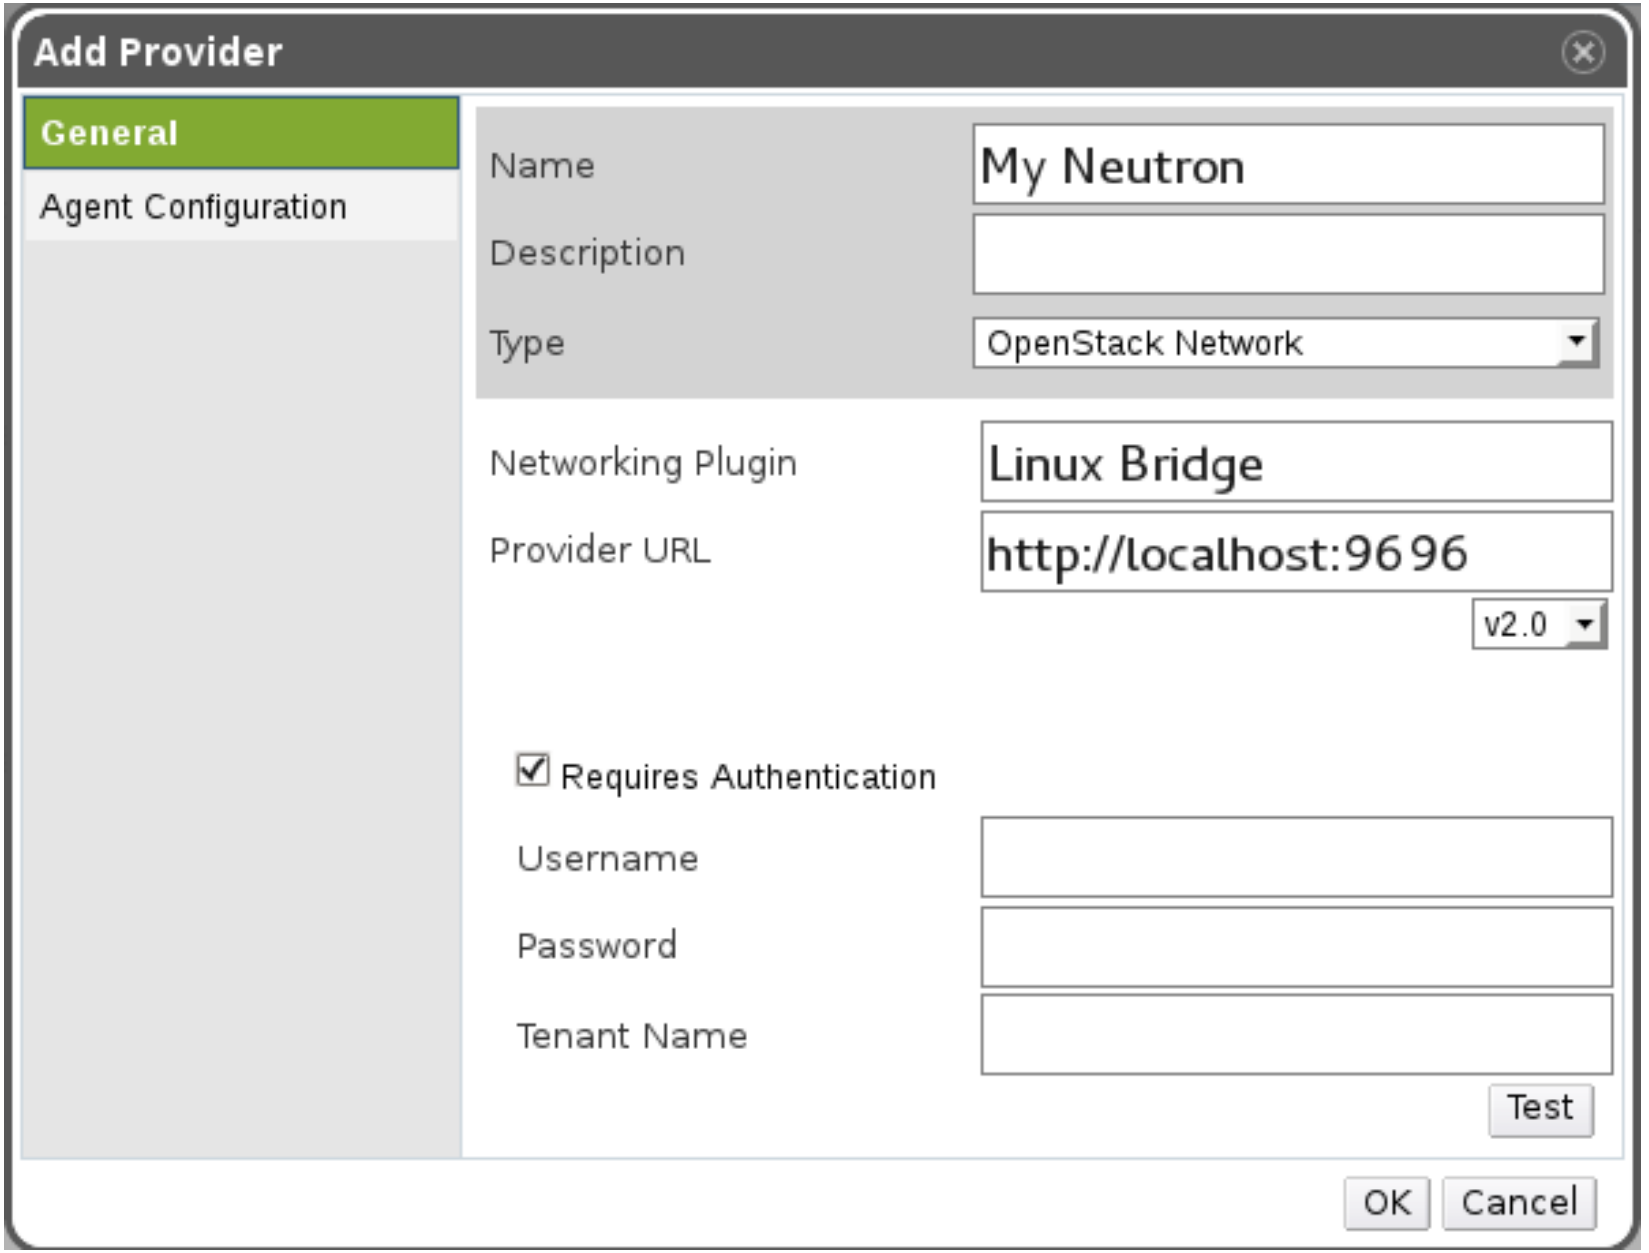

Step 2: Adding a Neutron Provider

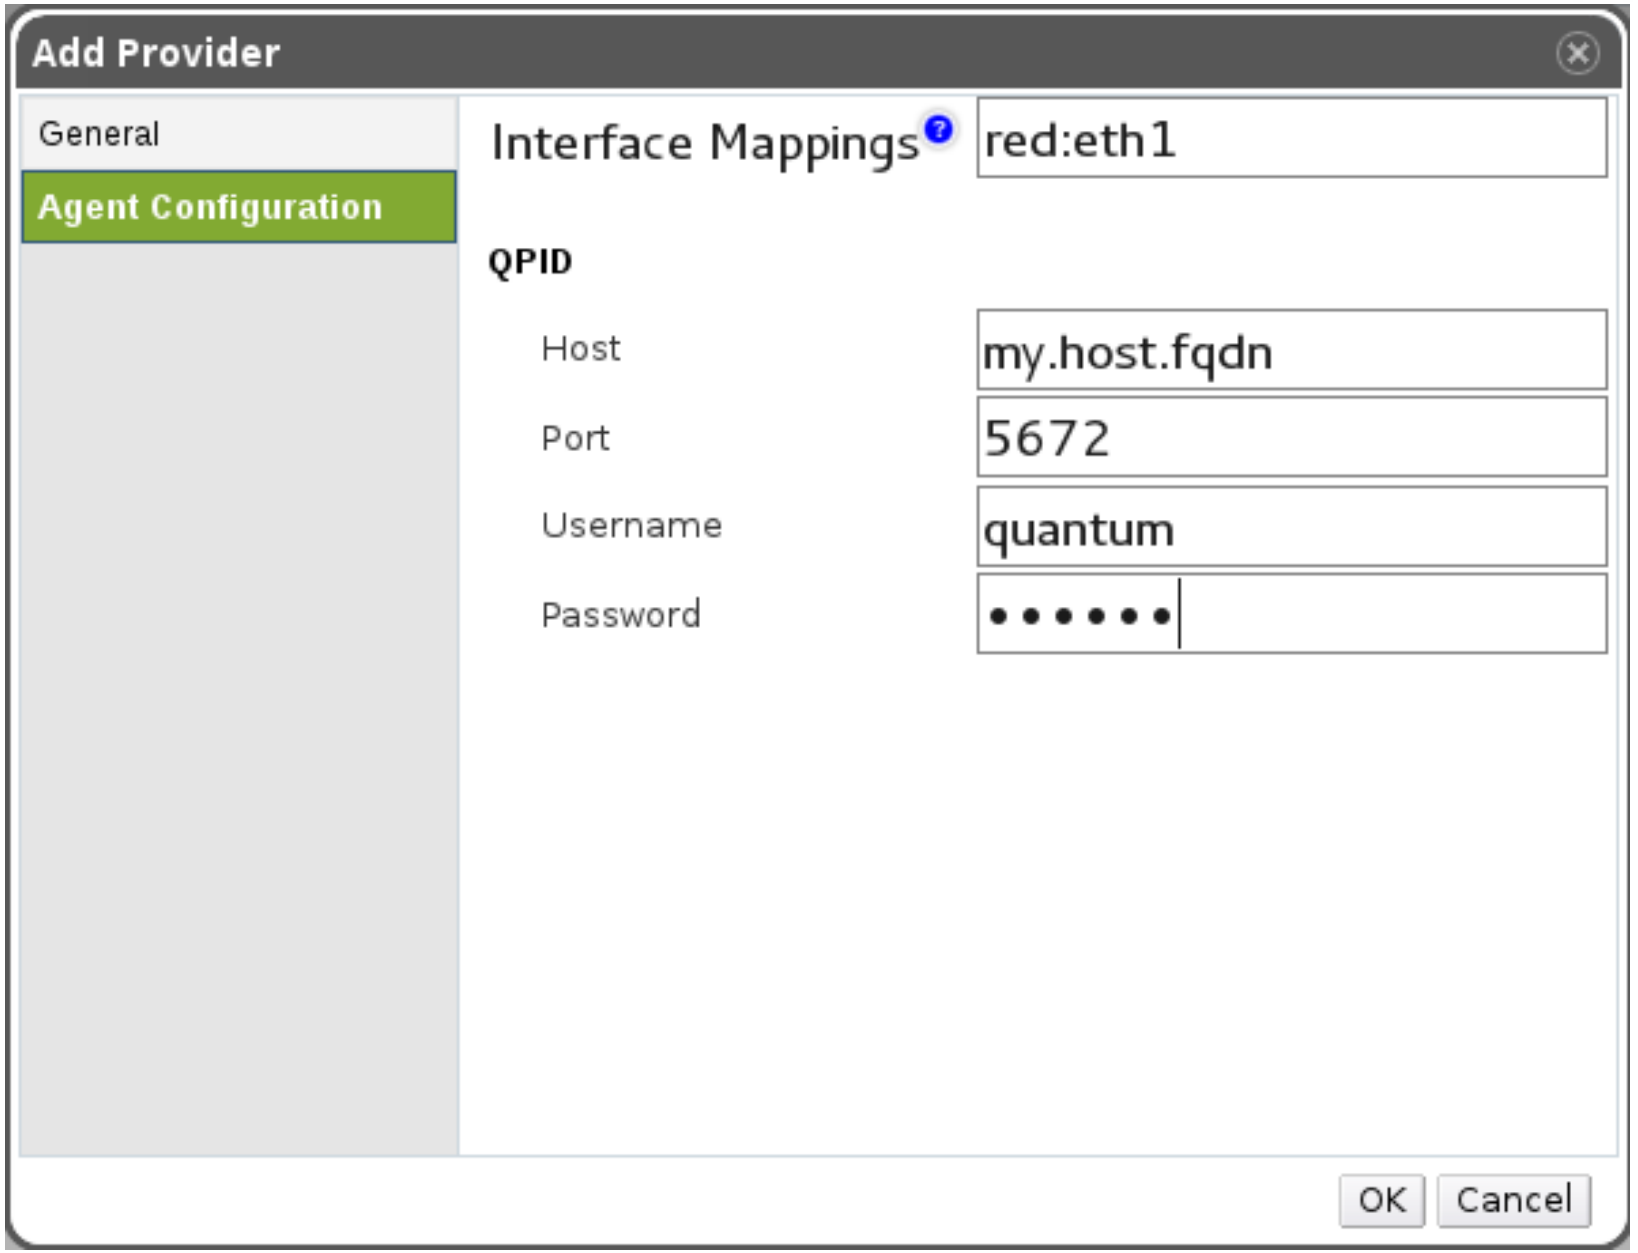

Step 3.1: Importing Networks

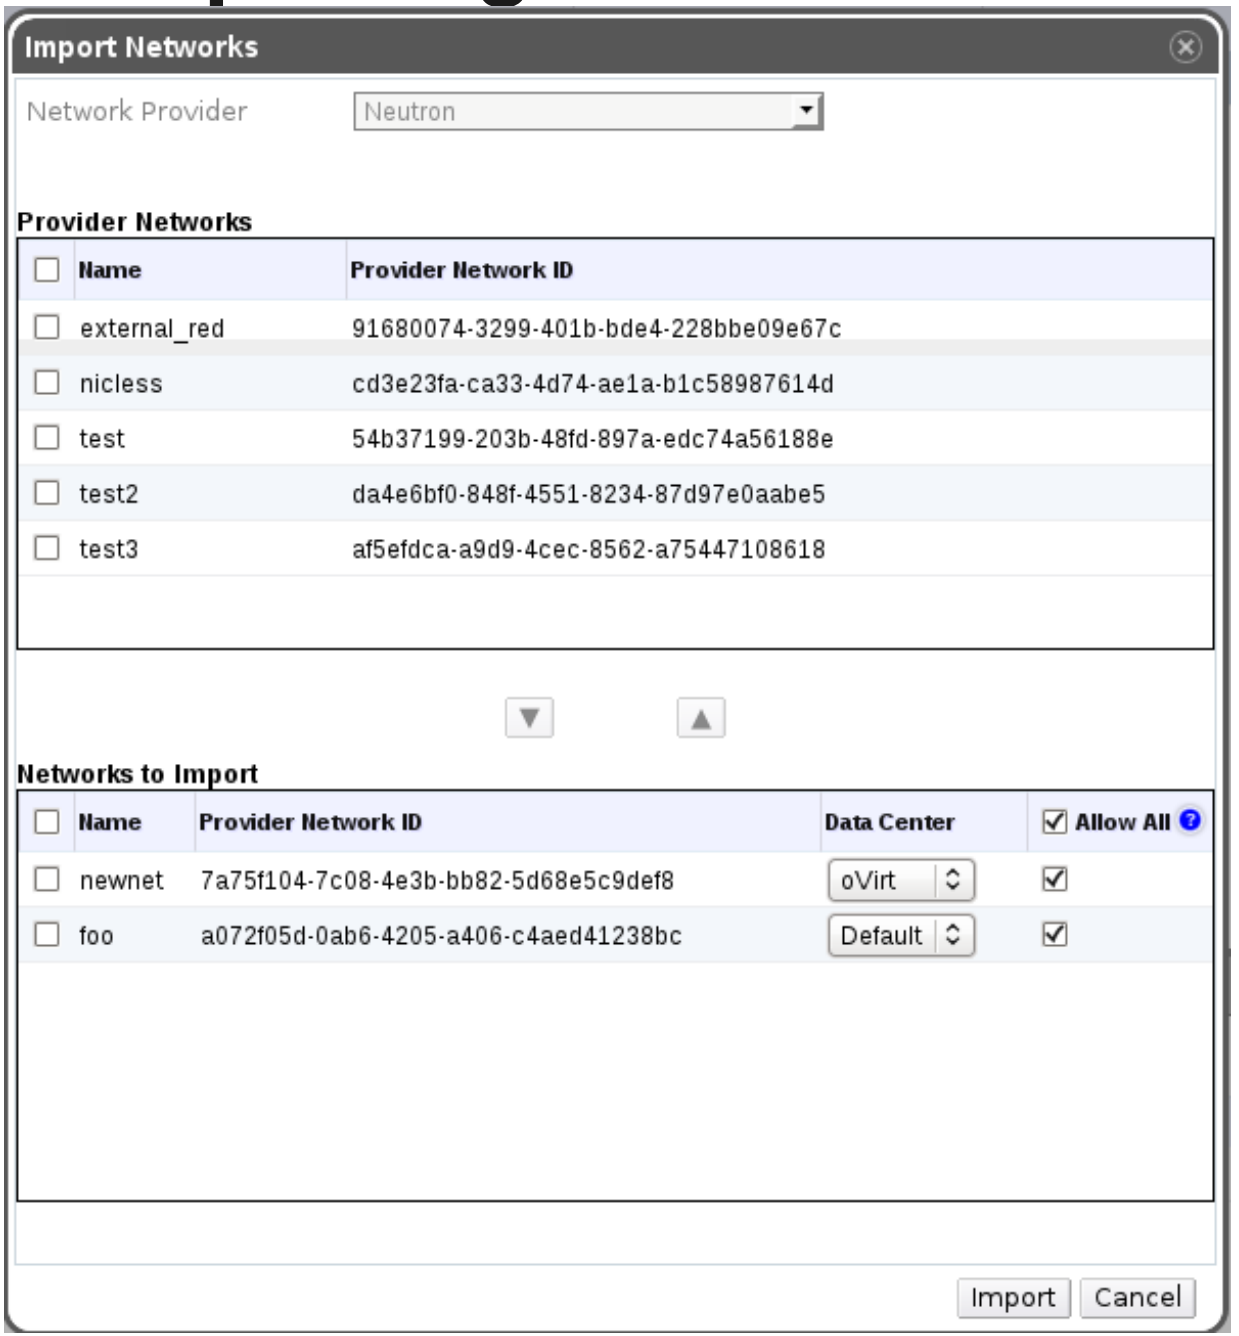

Step 3.2: Adding a New Network

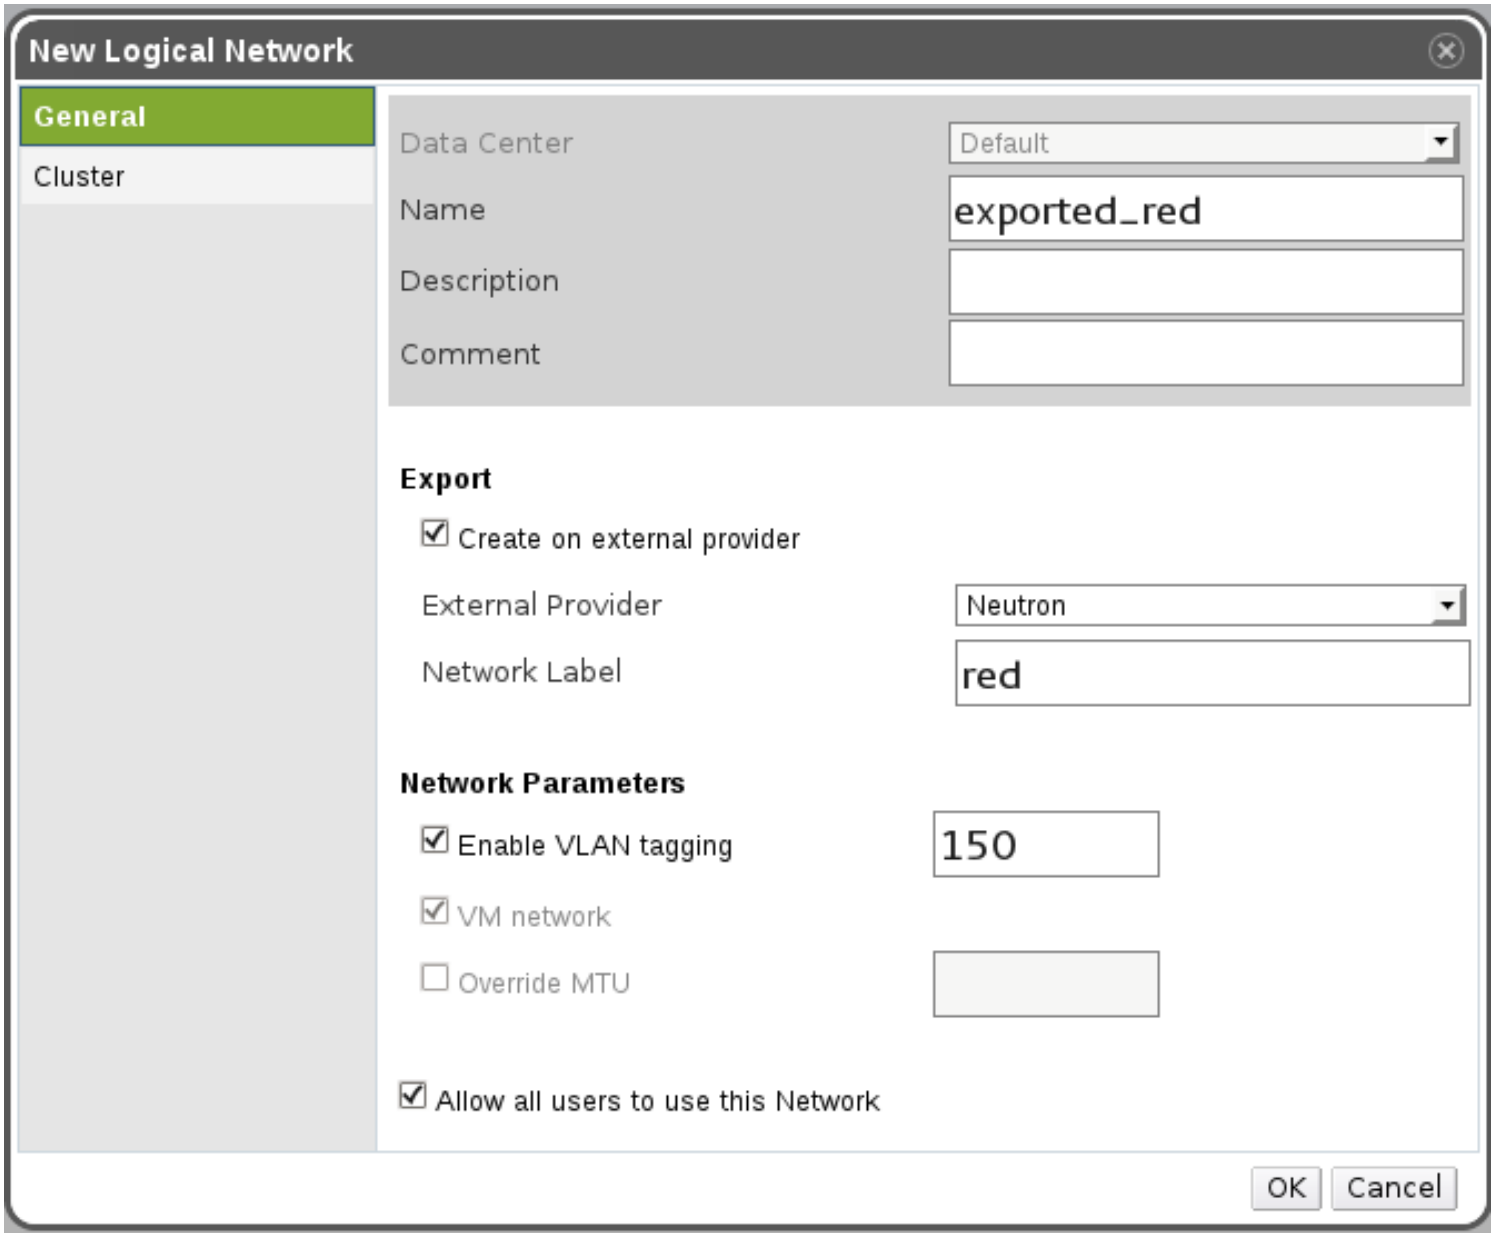

Step 4: Installing Host With Agent

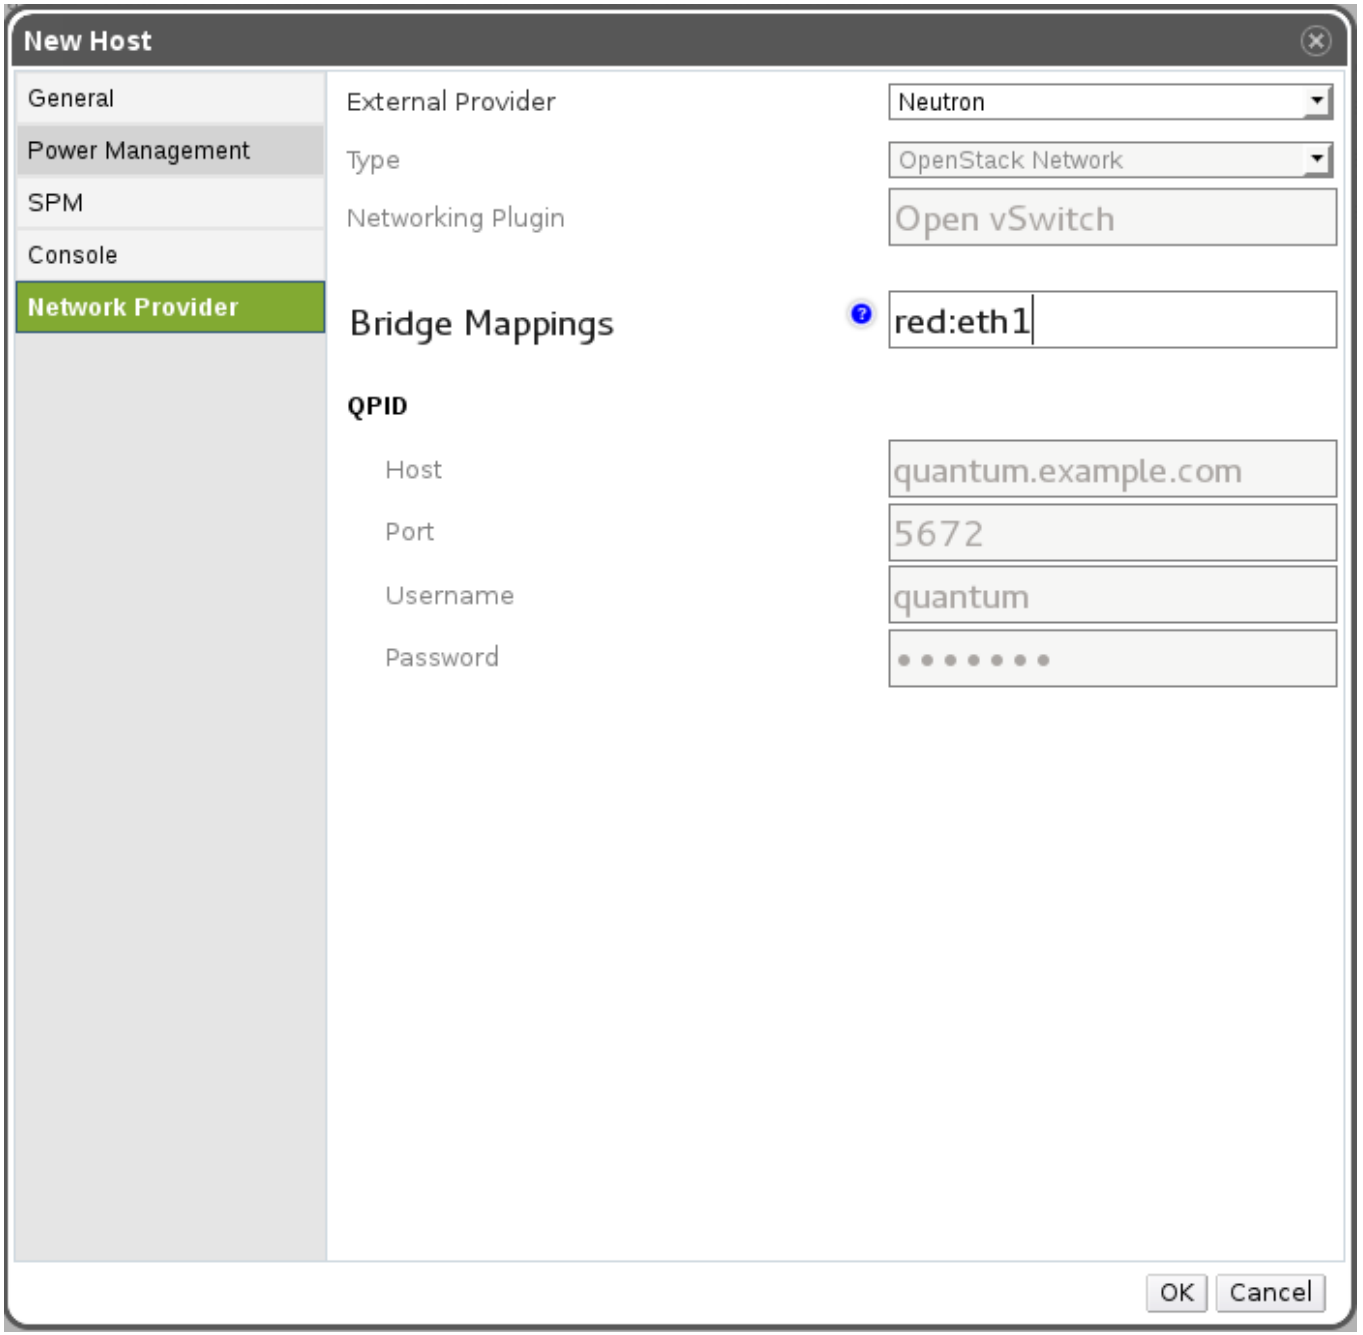

Step 5: Adding Network to a vNIC

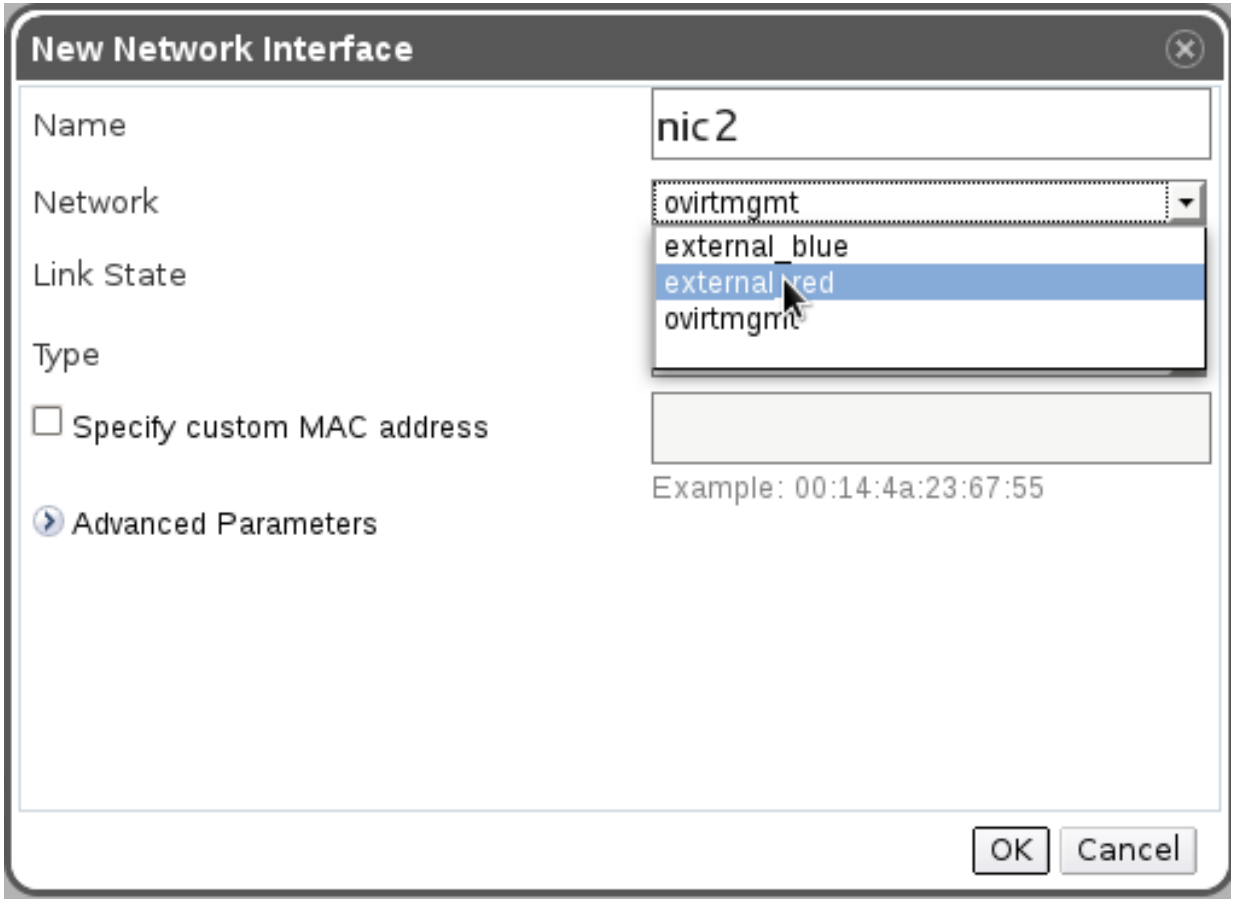

Step 6: Running the VM

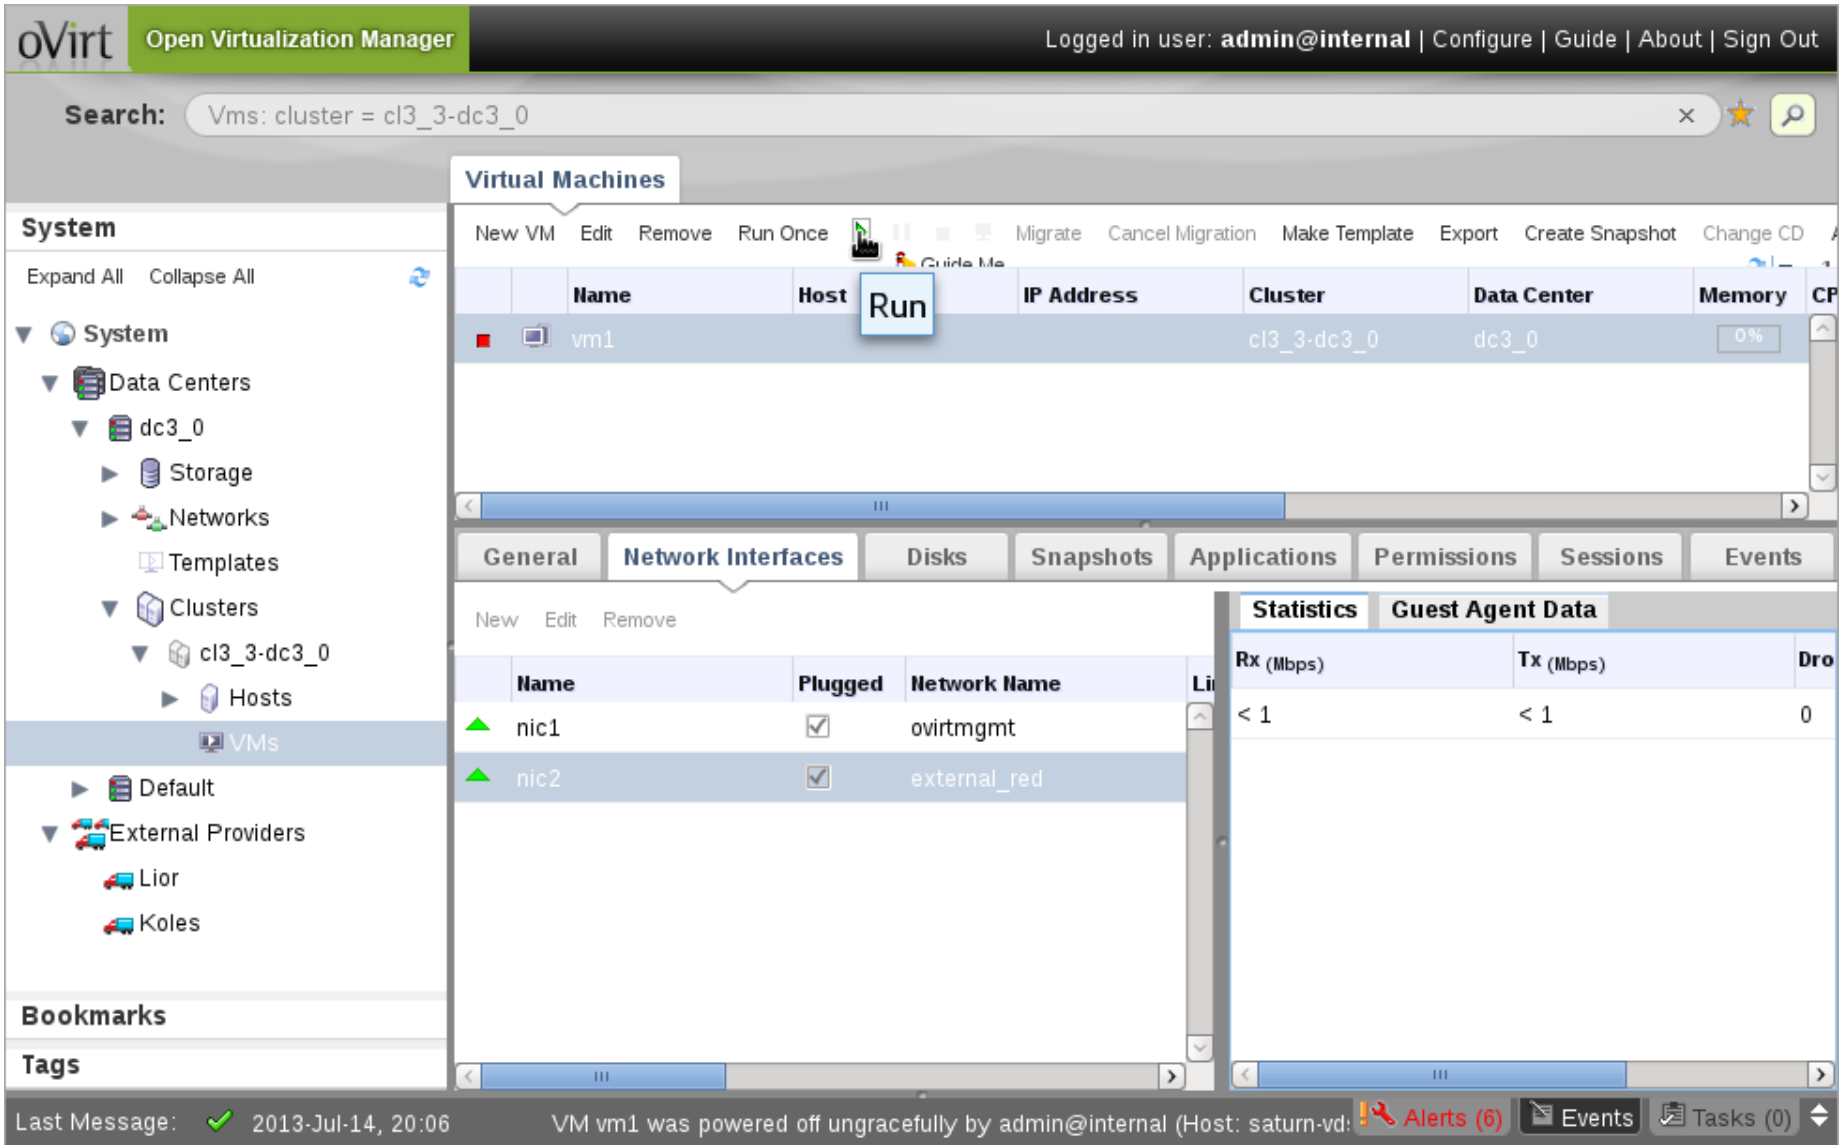

oVirt

Under the Hooks:
A deep dive to the hook internals

Hooks Overview

- VDSM is the oVirt “Compute agent”
  - Responsible for VM life cycle
  - Connects networking & storage resources
- Hooking mechanism in VDSM allows the administrator to define scripts that extend VM operation

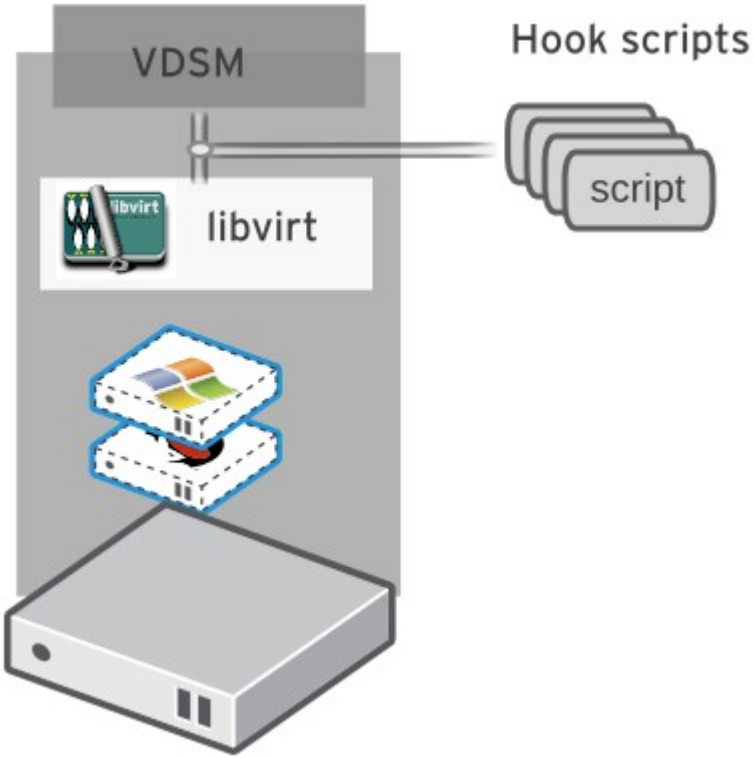

Under the Hook: Hooks Overview

oVirt

- VDSM is the oVirt “Compute agent”
  - Responsible for VM life cycle
  - Connects networking & storage resources
- Hooking mechanism in VDSM allows flexible extension
- Hook needs to be installed separately
- Hook failure can cause the VM to fail to start
- Post-hooks XML is copied on VM migration

Neutron Usage

oVirt

- We're utilizing hooks to connect the Neutron vNIC correctly
  - The hook takes care of the vNIC XML sent to libvirt
  - Local agent takes care of the connectivity on the host
- Current hook has support for:
  - Open vSwitch
  - Linux Bridge
  - Anything else can be added by a 3rd party vendor

Under the Hook:
Open vSwitch Agent

oVirt

Open vSwitch: How Does it Work?

oVirt

- Required:
  - Integration bridge & physical connections are preset on the host
- Once a vNIC exists on the integration bridge:
  - The agent sets the flows of information for the vNIC
- Henceforth, the agent monitors the port status:
  - If port admin state is changed it will be updated accordingly
  - If port is deleted, it's flows will be removed

Open vSwitch in Action

oVirt

Step 1:
A VM with a vNIC is started by libvirt

The tap device is connected to br-int
The port id is sent in the metadata of OVS

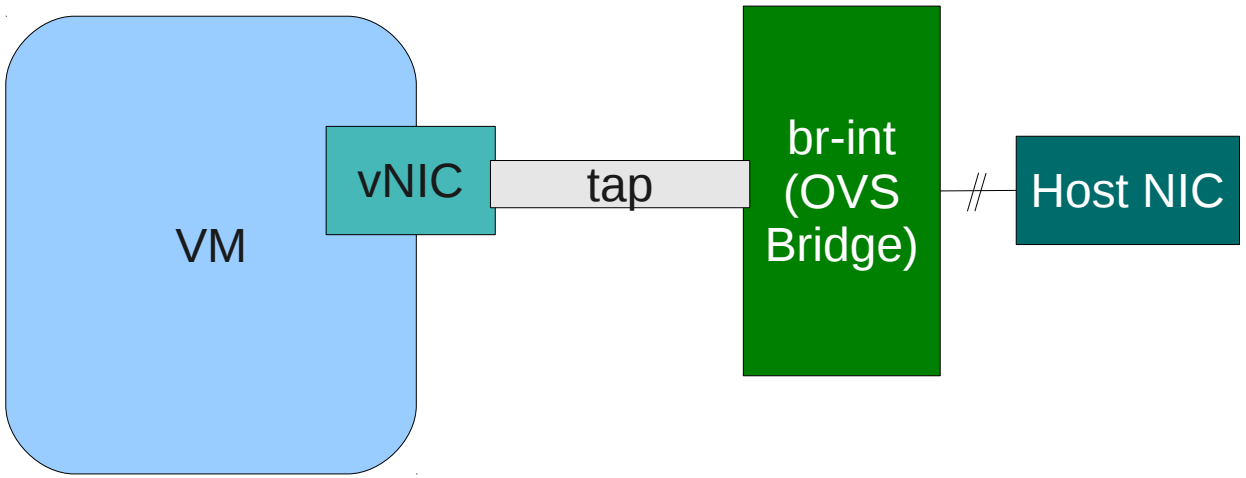

Open vSwitch in Action

Step 2:
Open vSwitch agent detects the tap

The agent queries the plugin for the port + network info

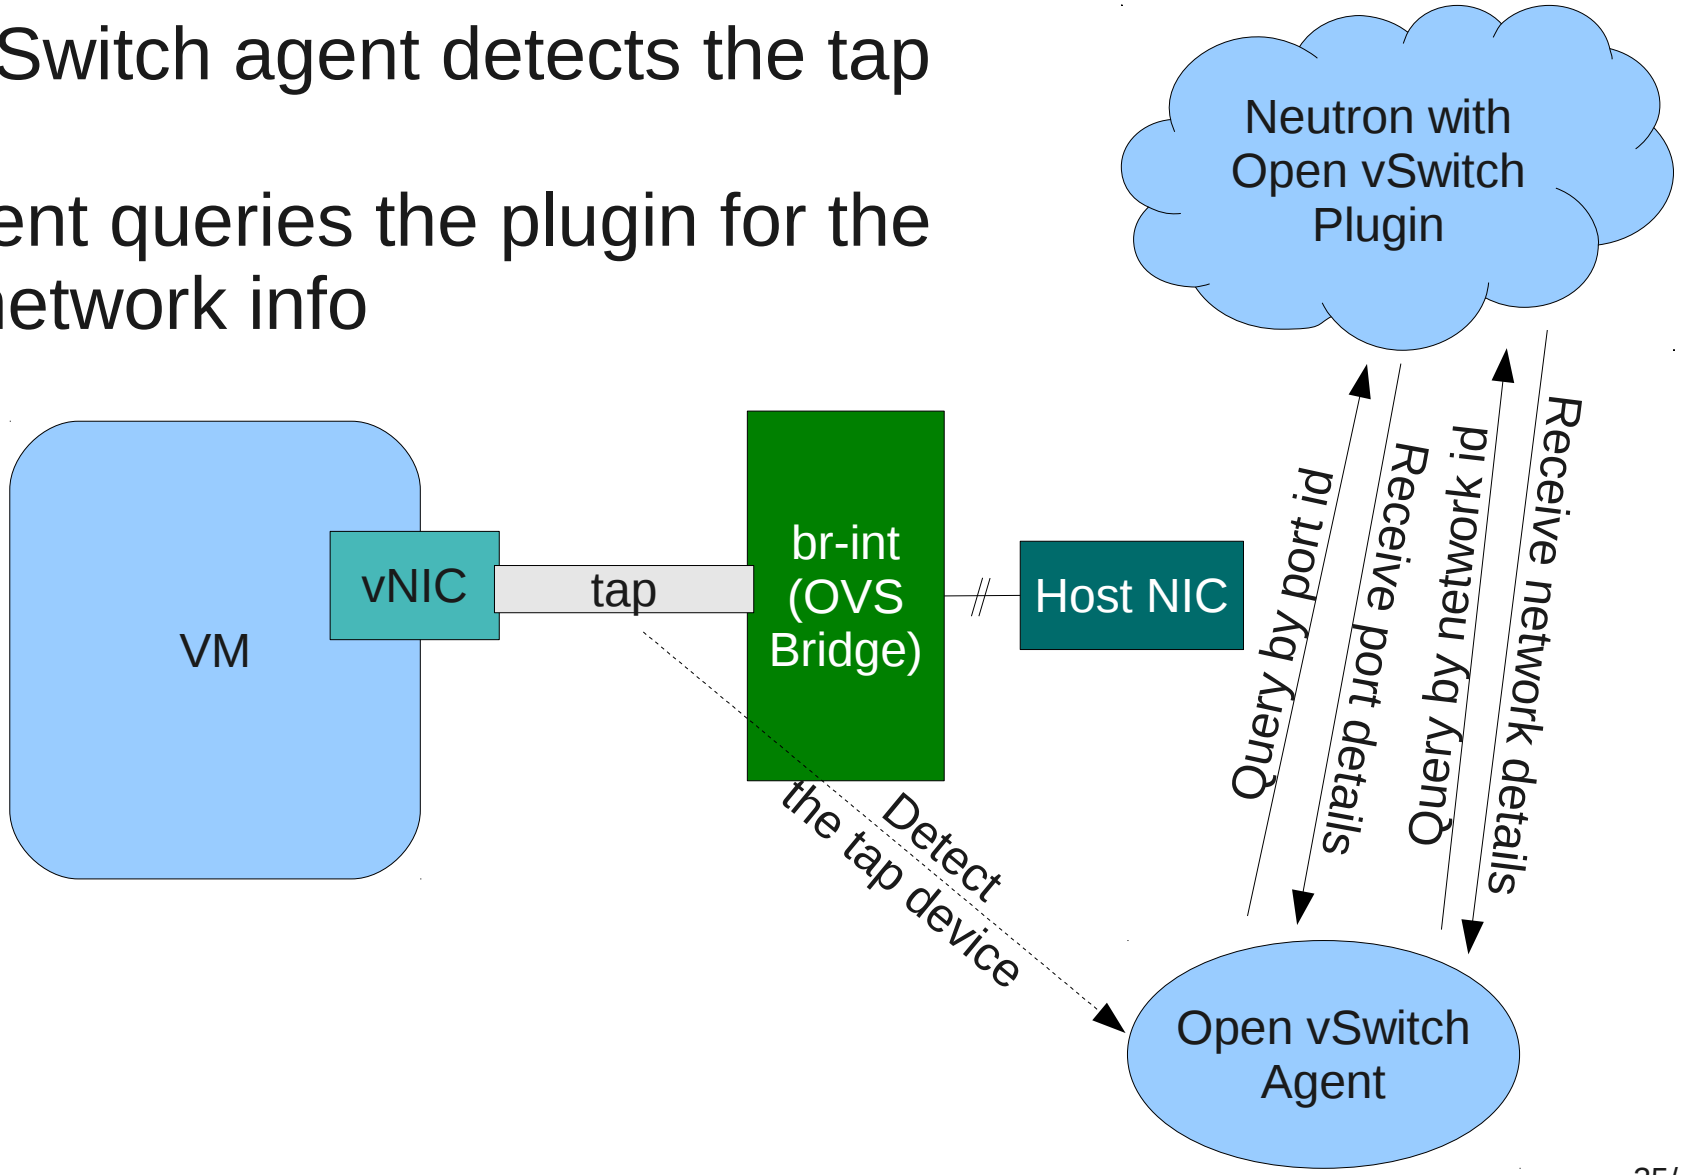

Open vSwitch in Action

Step 3:
Open vSwitch agent sets the necessary flows for the tap

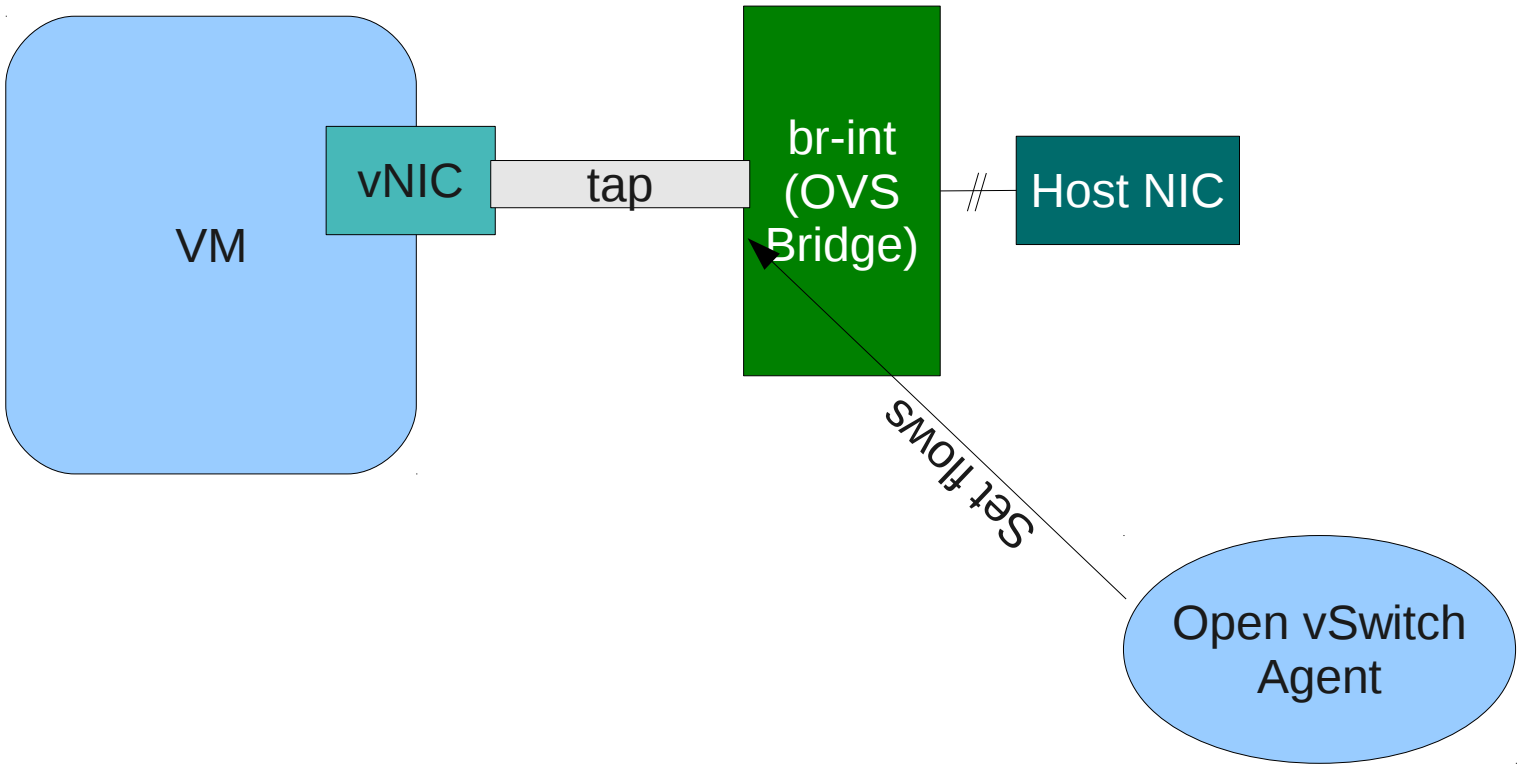

Under the Hook: Open vSwitch

oVirt

- On engine, a port is created on the Neutron network
- Pre-start (VM/hot plug):
  - Bridge is changed to br-int
  - A new element “<virtualport type='openvswitch'>” is added
  - Port id is set as a child of the new element
- Agent takes care of the host connectivity

Under the Hook:
Linux Bridge Agent

oVirt

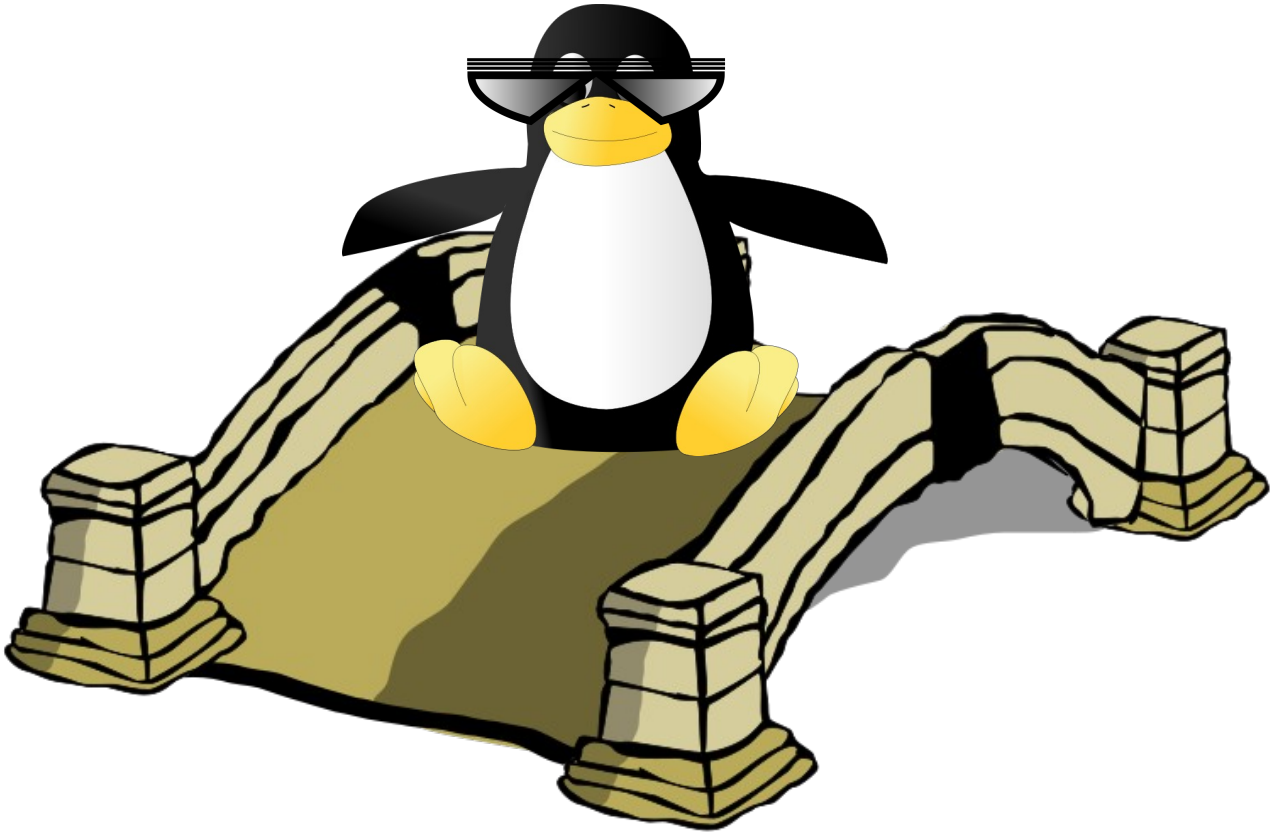

Linux Bridge: How Does it Work?

oVirt

- Once a vNIC exists on the host:
  - The agent ensures the brigde/VLAN exist
  - The agent connects the tap device to the bridge
- Henceforth, the agent monitors the port status:
  - If port admin state is changed it will be connected or disconnected accordingly
  - If port is deleted, it will be disconnected

Linux Bridge in Action

Step 1:
A VM with a vNIC is started by libvirt

The tap device is not connected to a bridge
The name of the device is “tap” + port id[:11]

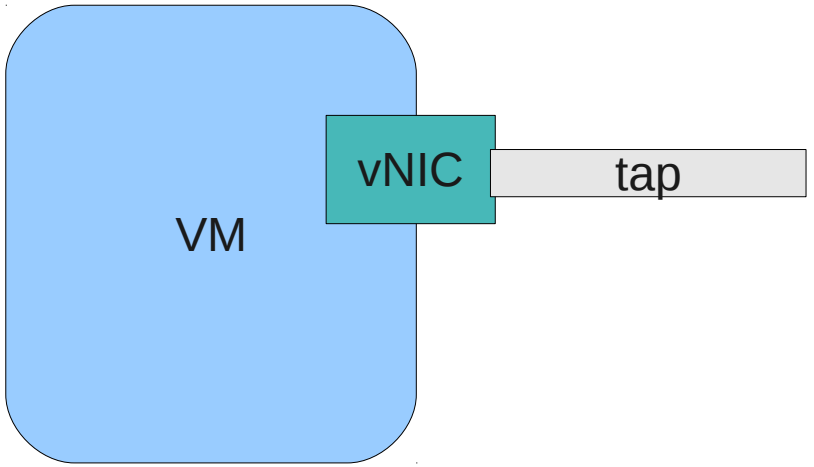

Linux Bridge in Action

Step 2:
Linux bridge agent detects the tap

The agent queries the plugin for the port + network info

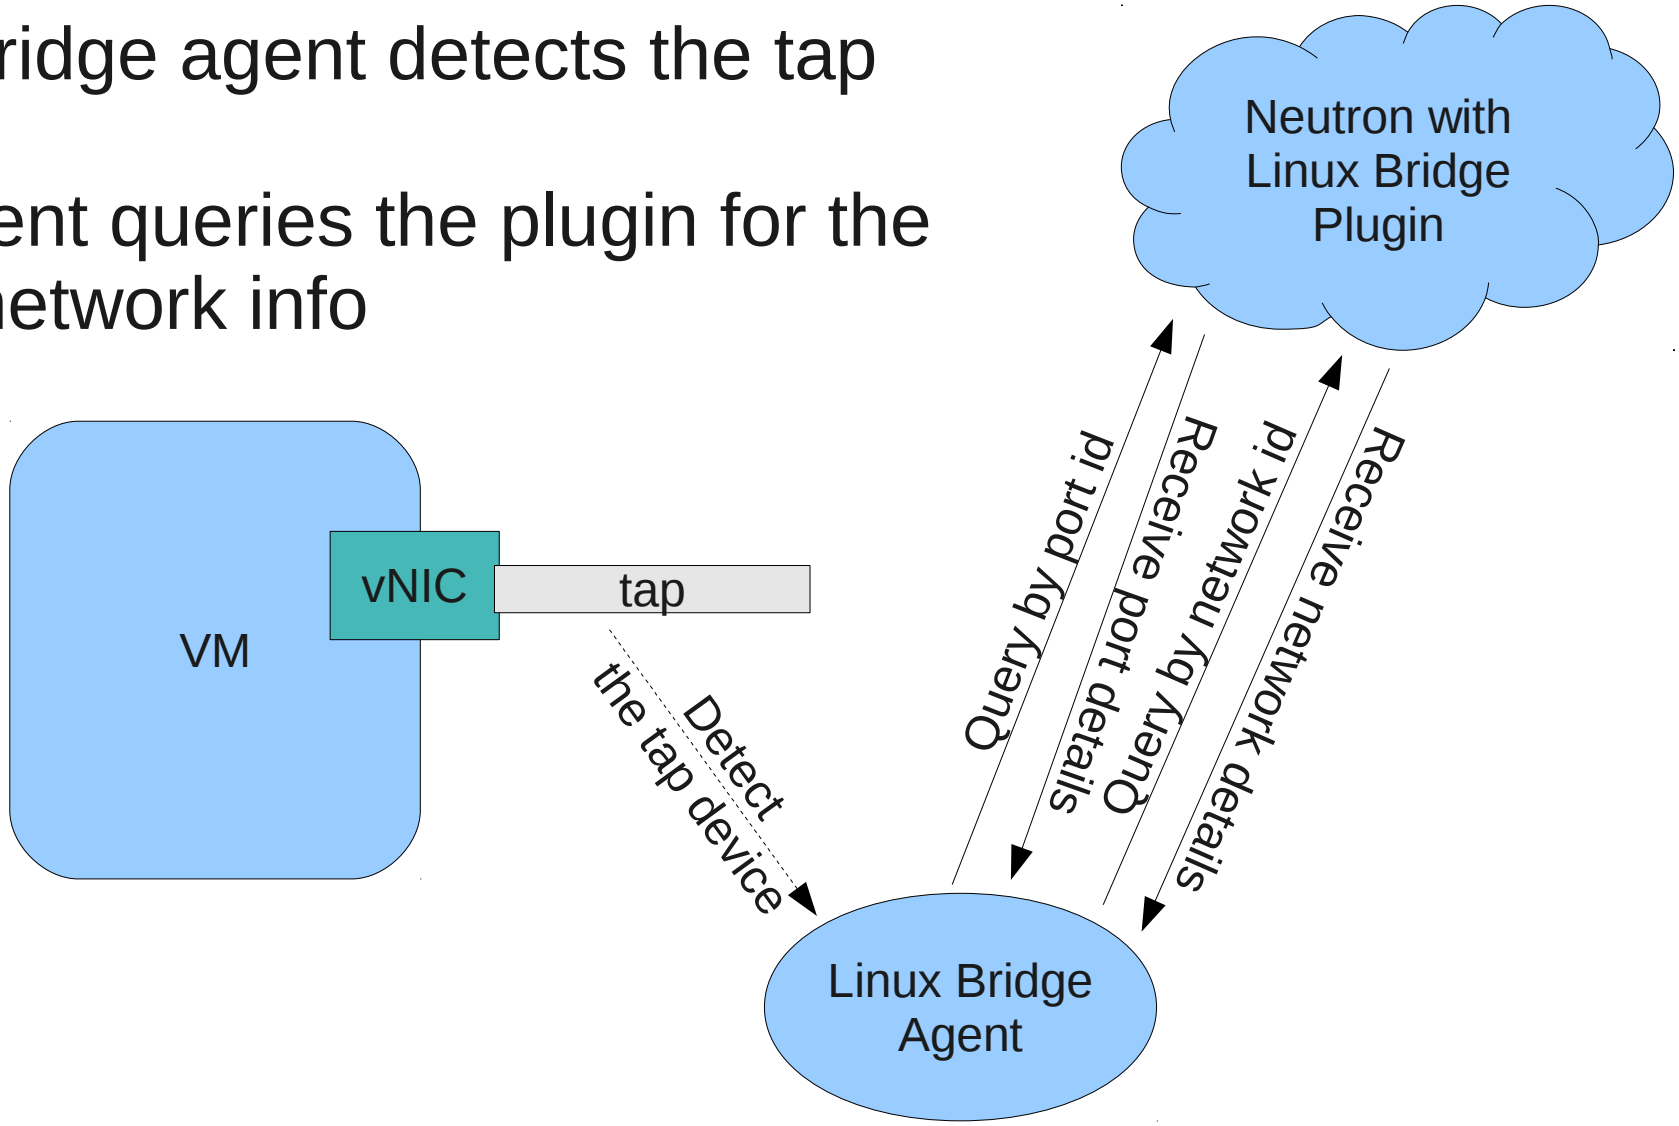

Linux Bridge in Action

oVirt

Step 3:
Linux bridge agent ensures VLAN + bridge on the host

Once the bridge exists, connects the tap device to the bridge

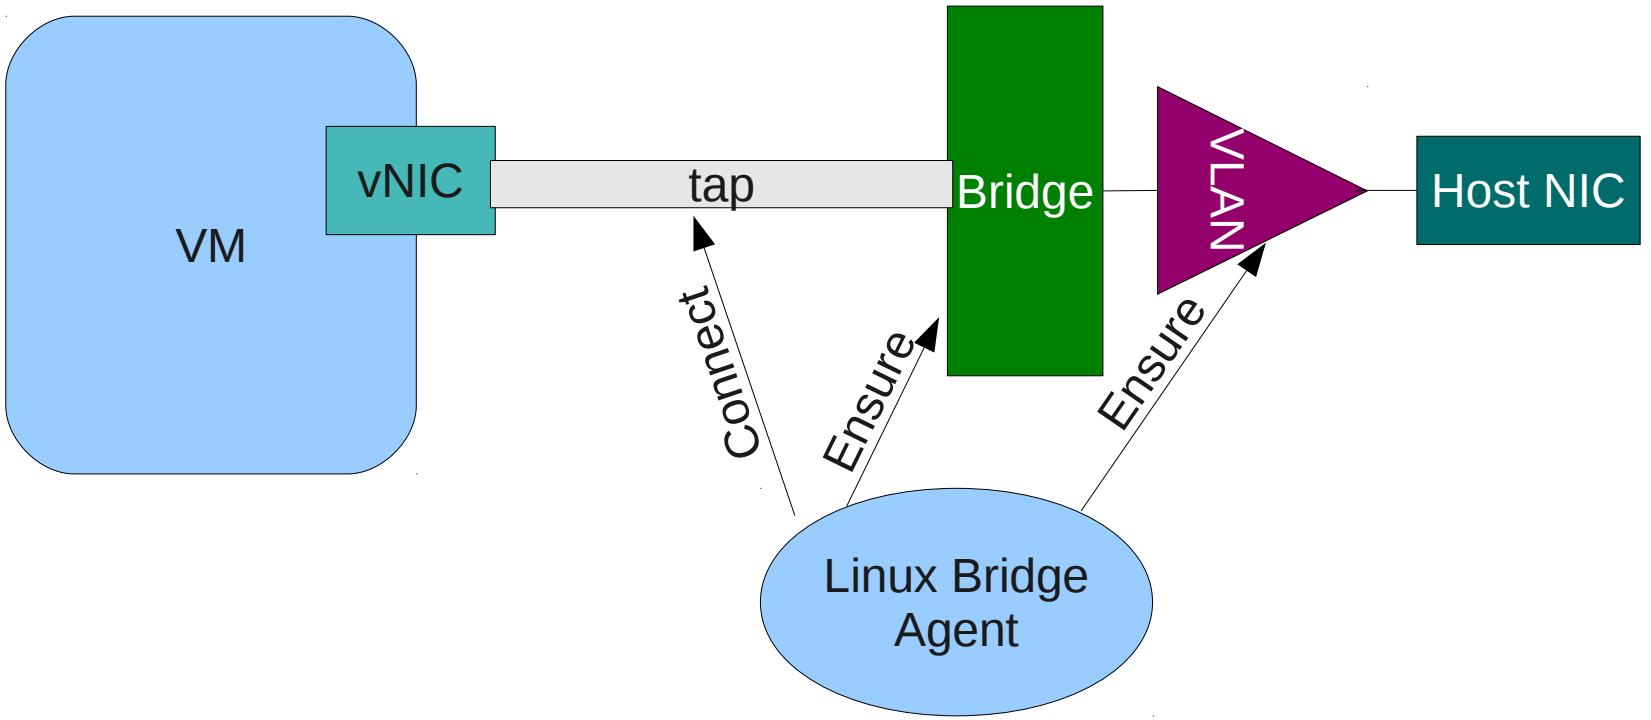

Under the Hook: Linux Bridge

oVirt

- On engine, a port is created on the Neutron network
- Pre-start (VM/hot plug):
  - Bridge is changed to dummy bridge
  - Tap name is set to “tap” + port id[:11]
- Post-start (VM/hot-plug/migrate):
  - Disconnect the tap from the dummy bridge
- Agent takes care of the host connectivity

oVirt

Future Work

Future Work

oVirt

- Improve VM scheduling, taking into account the networks availability on the host
  - Which host has access to which network
- Monitor vNIC connectivity after VM/vNIC started
- Integrate the tenant concept into oVirt
- Manipulate external networks from within oVirt
- Integrate L3 functionality into oVirt
- Support more Plugin types

Future Work and Open Questions

oVirt

- Auto-discovery mechanism
  - Open issues - Which Data Center? Which permissions?
- Import a network multiple times? With different properties, different SLA, etc.
- Multiple providers associated with a single oVirt network, enables support of multiple technologies for the same network on different physical segments

In Conclusion

oVirt

- oVirt network configuration
- Neutron overview
  - Integration benefits
  - External providers
  - Neutron as an external provider
- Under the hook
- Future Work

More info

oVirt

- Neutron
  - https://wiki.openstack.org/Neutron
- oVirt
  - http://www.ovirt.org/Network_Provider
- Mailing lists
  - users@ovirt.org
  - arch@ovirt.org
  - engine-devel@ovirt.org
  - vdsm-devel@lists.fedorahosted.org
- IRC Channel
  - #ovirt channel on irc.OFTC.net

oVirt

Thank You for Listening!

Questions ?