

Configuring your environment so that deployed Virtual Servers are VEPA (802.1Qbg) enabled

Use the steps in this quick start guide to set up your environment so that the deployed virtual servers are enabled to support Virtual Ethernet Port Aggregator (VEPA). You can also use this procedure to set up different tenants, or groups of users, with separate network connections using Network Control and VMControl from the IBM Flex System Manager management node.

Prerequisites:

Make sure that you have met the following prerequisites before you complete the steps in this quick start guide:

- If there are Redhat 6.4 KVM hypervisors running in the environment, verify that:
 - The kernel version is 2.6.32-345 or higher.
Note: If the KVM Hypervisor is Redhat 6.3, and it has a lower kernel version, the virtual servers will not be able to communicate to the outside due to a Linux kernel defect. As a work-around, start a packet capture on the interface to which the VM is connected. For example, if you create a virtual server using vlan 100 on interface eth0, you can start packet capture using the command `tcpdump -i eth0.100`

Lldpad package with version 0.9.45 or higher is installed and running. You can determine the Lldpad version by running the command:

```
lldpad -v
```

- If there are Power Systems hypervisors running in the environment, verify that
 - The Power Systems Firmware is at level AF770 or higher.
 - The Power Systems compute nodes support Virtual Server Network phase 2:

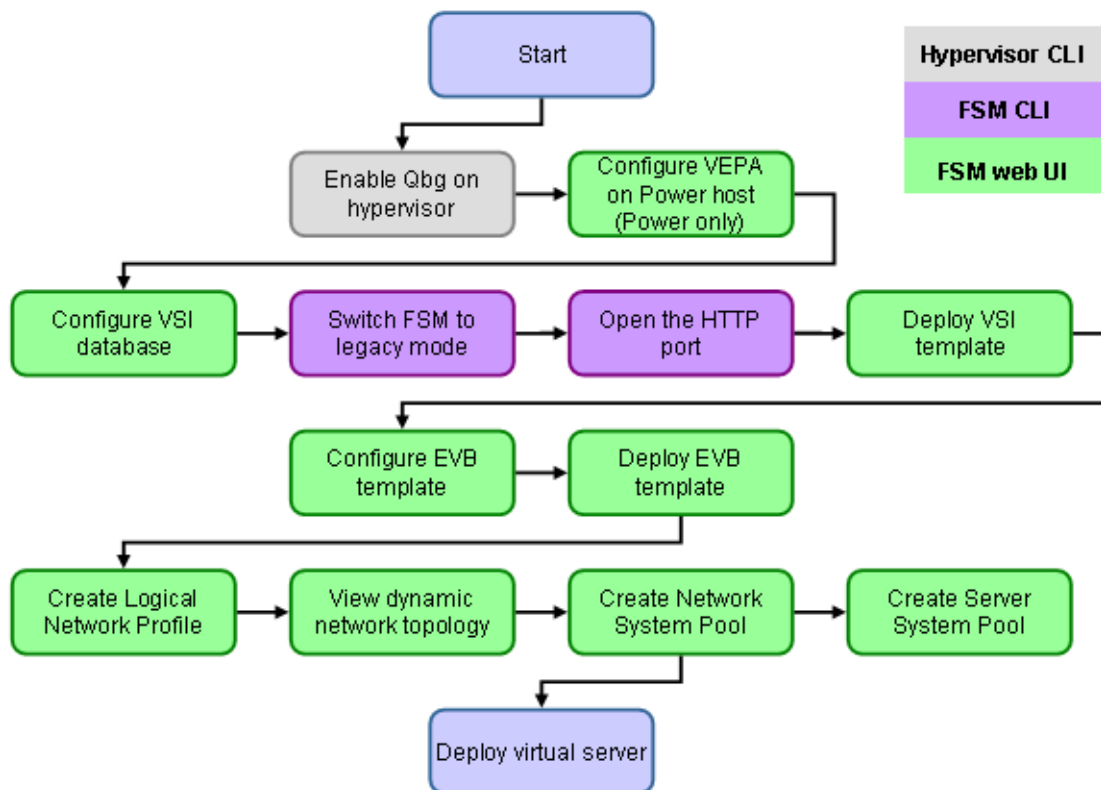
To determine the level of support, run the following commands:

- `lssysconn -r all`
Record the model and serial number of specific compute node
 - `lssyscfg -r sys -m <system> -F`
where `-m` is the model and serial number of the compute node
 - The Lldp daemon is running on the Virtual I/O server.
 - The Virtual I/O Server is configured with Shared Ethernet Adapter(s) backed by physical adapters.
- Verify that the physical adapters of the hypervisor are connected to an IBM Networking Switch that has a firmware supporting 802.1 qbg. IBM System Networking Switches started supporting 802.1 Qbg beginning with firmware version 7.2.

In addition to these prerequisites, the following terminology may be helpful as you perform the steps in this quick start guide:

- Server system pool – logically groups similar hosts and facilitates the relocation of virtual servers from one host to other hosts in the system pool.
- Network system pool – defines the set of switches and their attached servers that Automated Logical Network Provisioning will configure.
- Logical network profile – a template that defines a VLAN ID and VSI parameters required by a virtual machine or set of virtual machines.
- Uplink port – a port on a switch in a network system pool that connects the network system pool to a switch outside of the network system pool.
- VEB (Virtual Ethernet Bridging) - software that acts similar to a Layer 2 hardware switch providing inbound/outbound and inter-VM communication. A VEB works well to aggregate multiple VMs traffic across a set of links and to provide frame delivery between VMs based on MAC address. Where a VEB is lacking is network management, monitoring and security.
- VEPA (Virtual Ethernet Port Aggregator) - VEPA is standard for providing consistent network control and monitoring for Virtual Machines (of any type.) VEPA has been used by the IEEE as the basis for 802.1Qbg, Edge Virtual Bridging. The hypervisor forces each VM frame out to the external switch regardless of destination. This causes no change for destination MAC addresses external to the host, but for destinations within the host (another VM in the same VLAN) it forces that traffic to the upstream switch which forwards it back instead of handling it internally, called a *hairpin turn*.) This hairpin turn causes the requirement for the upstream switch to have updated firmware; typical STP behavior prevents a switch from forwarding a frame back down the port it was received on. The firmware update allows the negotiation between the physical host and the upstream switch of a VEPA port, which then allows this hairpin turn.
- VSI (Virtual Station Interface) - An internal point-to-point Ethernet connection that connects an Ethernet port of a VEB or VEPA to a vNIC. Each VSI carries a single MAC service instance. Sometimes the term VSI is used to refer to the reference point where the internal LAN attaches to the vNIC. Each VSI can be represented by a collection of a VM, vNIC, vSwitch and the connection between a vNIC and vSwitch port.
- EVB (Edge Virtual Bridge) - Edge Virtual Bridging (EVB) is an environment where physical end stations containing multiple Virtual Station Interfaces (VSIs) that require the services of adjacent bridges forming a LAN. EVB environments are unique in that virtual Network Interface Controller (NIC) configuration information is available to the EVB device that is not normally available to an 802.1Q bridge.
- Edge Switch – This document uses the term edge switch as the physical switch to which the Server NICs of the Hypervisor host is connected.
- VM – Virtual Machine. Also known by terms such as *virtual server*, *logical partition*, and *LPAR*.

The following chart summarized the steps required to to set up your environment so that the deployed virtual servers are enabled to support Virtual Ethernet Port Aggregator (VEPA):



Steps:

Complete the following steps to set up your environment so that the deployed Virtual Servers are VEPA capable.

1. Enable 802.1 qbg on Hosts (Hypervisors)

- a) If you are using KVM hypervisors, configure VEPA on the KVM host.
Note: Run these commands as root.

- 1) Find the network interface connected to your edge switch supporting 802.1 qbg. Suppose it is eth0, then run the following commands to enable reflective relay on this interface.

```
lldptool -L -i eth0 adminstatus=rxtx
```

```
lldptool -i eth0 -g ncb -L adminstatus=rxtx
```

```
lldptool -T -i eth0 -g ncb -V evbCfg -c enableTx=yes
```

```
lldptool -T -i eth0 -g ncb -V evbCfg -c fmode=reflectiverelay
```

```
lldptool -T -i eth0 -g ncb -V evbCfg -c  
capabilities=rte,ecp,vdp
```

```
lldptool -i eth0 -T -g ncb -V evbCfg enableTx=yes
```

```
lldptool -i eth0 -T -g ncb -V evbCfg fmode=reflectiverelay
```

```
lldptool -i eth0 -T -g ncb -V evbCfg capabilities=rte,vdp,ecp
```

```
lldptool -T -i eth0 -V vdp enableTx=yes
```

```
lldptool -i eth0 -T -g ncb -V vdp enableTx=yes
```

```
service lldpad restart
```

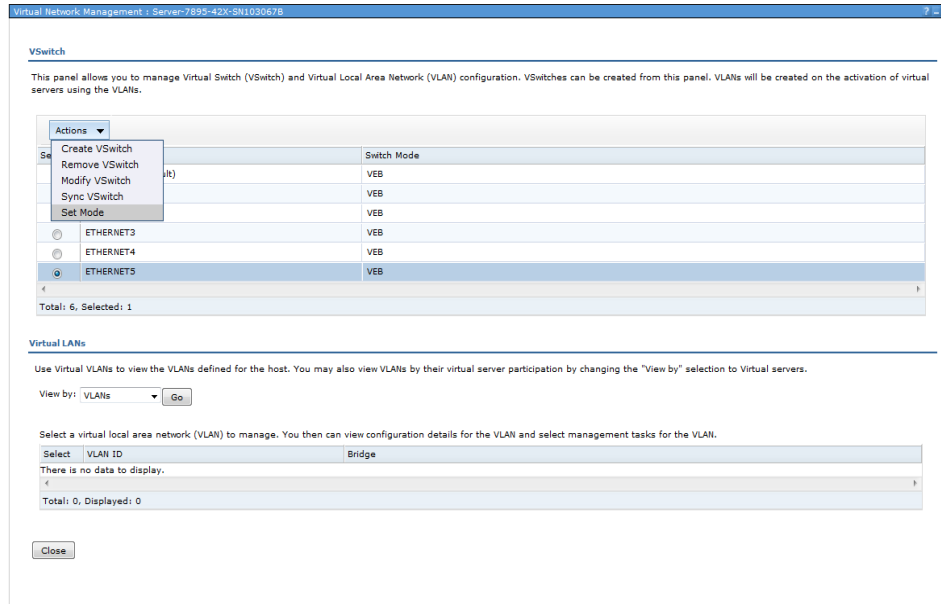
- 2) Verify that the configuration is done correctly by running this command:
`/opt/ibm/icc/bin/cimcli -n root/virt ei KVM_SwitchService`

Verify that the output contains:

IsVSISupported = TRUE;

VSIInterface = "eth0";

- b) If you are using Power Systems hypervisors, configure VEPA on Power host:
 - 1) Select the power System from the Resource explorer, right - click **System Configuration** → **Virtual Resources** → **Virtual Network Management**.
 - 2) From the list of virtual switches, select the vSwitch that is connected to the physical adapter connected to the edge switch. Select **Actions** and then set the mode as VEPA.



3) On the Virtual I/O Server managing the Power Server, complete the following steps to ensure that the IBM Flex System Manager detects the correct network topology.

a. Login to the Virtual I/O server and run the following command to obtain root access:
`oem_setup_env`

b. Verify that lldp is running using the following command:
`ps -ef | grep -i lldpd`

If the service is not running, start it with the following command:
`startsrc -s lldpd`

c. Find the Shared Ethernet Adapters by running the command
`lsdev | grep Shared`

d. For each Shared Ethernet Adapter, run the following command
`chdev -l <SEA name> -a lldpsvc=yes`

e. Verify the neighbor information of this SEA by running command:
`lldpctl show neighbor <SEA name>`

Make sure that Reflective Relay is shown under both the capabilities and current configuration section.

2. Create a VSI Database configuration Template:

- Log in to the IBM Flex System Manager user interface with a user account that has sufficient privileges to configure devices managed by the IBM Flex System Manager management node.
- On the Home page, select the Plug-ins tab.

- c) Select **Configuration Manager**.
- d) Under Configuration tasks, select **Create configuration template**.
- e) On the Configuration Templates page, click **Create**. Complete the wizard using the following values:
 - Template type – **Operating System**.
 - Configuration to create a template – **802.1Qbg VSI Database configuration**.
 - Name – The name to be used for the template. For example: VSIDBTemplate.
- f) Complete the wizard using the following values:
 - (Optional) On the ACL Filters page, create some ACL filters depending on the ACLs to be used by your Virtual Machines.
 - On the VSI Types page, click on create, specify the VSI Type Id (corresponding to Index Number), version, manager Id, Name, VLANs to be used by your VM and add optional ACL filters created in previous page. Add more VSI Types if required.
 - Note down the vlan Id, VSI Type Id (corresponding to Index Number), version and manager Id so that these needs to be provided in a later step when you create a Logical Network Profile.
 - On Summary page, verify the data and click on **Finish**.

More information can be obtained from

http://pic.dhe.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.8731.doc/com.ibm.sdnm.adv.helps.doc/fnc0_r_learnmore_VSI_config.html

3. Configure VSI Database on the Flex System Manager:

- a) Put the Flex System Manager into legacy mode and enable http access by running the following commands


```
smcli setFsmSecPolicy -l legacy
chnetcfg -s add -app-name http
```

More information can be obtained from:

http://pic.dhe.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.8731.doc/com.ibm.sdnm.adv.helps.doc/fnc0_r_learnmore_about_hosting_VSI_db_within_FSM.html

- b) Deploy the VSI DB template on to the FSM Operating System endpoint:
 - 1) On the Plug-ins tab on the Home page, select **Configuration Manager**.
 - 2) Under Configuration tasks, select View configuration templates.
 - 3) Select the VSI Database Configuration template created in the previous step, and click on deploy button on top.
 - 4) Select the Flex System Manager Operating System end point as the target, and deploy the template.
 - 5) Run the task and verify that the deployment has completed successfully.

4. Create an EVB configuration Template for your edge switch:

- a) On the Plug-ins tab on the Home page, select **Configuration Manager**.
- b) Under Configuration tasks, select **Create** configuration template.

- c) On the Configuration Templates page, click Create. Complete the wizard using the following values:

Template type – **Ethernet Switch**

Configuration to create a template – EVB Configuration for <your edge switch>

Name – The name to be used for the template. For example: EVBConfigtemplate

- d) Complete the wizard by following steps:
- 1) On EVB profile Index page, fill the EVB profile index for ports connected to your hypervisors. The profile index can be same for all the ports.
 - 2) In EVB Profile configuration page, Reflective Relay and VSI Discovery will be enabled for the profile indexes selected in previous step.
 - 3) Keep the default values in VSI Database Configuration page.
 - 4) Click **Finish** on the summary page, and save the template.

More information can be found at

http://pic.dhe.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.8731.doc/com.ibm.sd.nm.adv.helps.doc/fnc0_r_learnmore_EVB_config.html

5. Deploy the EVB Configuration on the edge switch:

Deploying the EVB configuration template on the edge switch will configure the switch in a way that it will access the FSM to get the VSI information, and the switch ports will have reflective relay and VSI discovery enabled.

- a) Deploy the EVB Configuration template on the edge switch:
On the Plug-ins tab on the Home page, select **Configuration Manager**.
- 1) Under Configuration tasks, select **View** configuration templates.
 - 2) Select the EVB Configuration template created in the previous step, and click on deploy button on top.
 - 3) Select your edge switch connected to your host hypervisor the target, and deploy the template.
 - 4) Run the task and verify that the deployment has completed.

6. Create the Network System Pool, Server System Pool and deploy VMs.

Refer the Quick Start Guide *Configuring your environment so VLANs are dynamically created*, for creating Logical Network Profile, Network System Pool, Server System Pool and deploying Virtual Appliances . The Quick Start Guide is available in Table 1 at the following location:

http://pic.dhe.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.commontasks.doc/commontasks_virtualization.html

Note: The step about specifying the VSI parameters in the Logical Network Profile is marked optional in the quick start guide. For VEPA environments, that step is required, and the VSI type, Version and manager ID, and the VLAN used in the Logical Network Profile has to match the values used when

creating the VSI Database Configuration Template.

To verify that the VEPA works correctly, start a ping test between two VEPA virtual servers, and verify that the edge switch port connected to the physical adapter of the hypervisor is showing the traffic. This can also be verified by logging on to the physical switch using Telnet or SSH and issuing the command `show virt evb vdp vm` in ISCLI to verify the switch is aware of the VMs provisioned.