

DefCon 23 Video Demo: MAC Flooding on Gentoo/Xen with Open vSwitch 2.0.0

DefCon 23 Video Demo: MAC Flooding on Gentoo/Xen with 802.1d Bridging

Open vSwitch on Gentoo – Setting Up Your First vSwitch

In my [last post](#) I outlined how to get Open vSwitch installed on Gentoo from source for version 1.11.0 as well as from portage using version 2.0.0. I also described how to associate Open vSwitch with with Xen based virtual machines. This guide will detail how to build your first virtual switch from scratch and create virtual switch ports associated with the virtual switch that persist upon a reboot of the host machine. The virtual ports can then be used with VirtualBox VMs and other generic applications that can make use of them.

If you followed my [last post](#) you should have Open vSwitch installed and the respective services running. You should also have created an initial bridge interface, this was called `xenbr0` and was created with the following commands:

where `enp4s0` is your physical interface that is bound to the bridge. This is necessary since this will be the “switch” port that “uplinks” the vSwitch to the rest of the physical network.

You should now be able to view the current bridge setup by using the following command:

Which should display output similar to this:

This shows that the bridge is created and has a single port attached to it labeled `enp4s0` that is bound to the physical interface `enp4s0`.

At this point you should be fine if you were using this for a Xen based setup since the hotplug scripts provided with `xen-tools` will take care of the virtual interface setup for each VM. However if you want to use Open vSwitch with VirtualBox based VMs on a Gentoo host read on ...

In order to create virtual switch ports that you can bind to your virtual machines a `tun/tap` interface must be created for each port that will be needed. So if you wish to create a 16 port switch 16 `tap` interfaces will be necessary. These interfaces can be added to the vSwitch as follows:

First verify that the `tun` module is loaded

If this does not run verify that you have `tun` compiled as a module in your kernel.

And that the tun module is set to load on boot up in /etc/conf.d/modules:

Next create the tap interfaces for each virtual switch port, in this example I will create an 8 port switch.

Now check that they were added:

Which should display:

This indicates that the virtual ports have been associated with the vSwitch xenbr0, however the interfaces still need to be created. In order to do this entries need to be made in /etc/conf.d/net that describes each tap interface and how they will be configured at bootup. In the case of these tap interfaces we want them to be brought up but not configured. So in /etc/conf.d/net add the following entries:

Notice how they are all set to "null" this indicates that on bootup they will not be assigned an IP address nor will they poll a DHCP server for an address.

Next create the device symlinks in /etc/init.d/ as you would for any other network interface so that they can be started on bootup:

Then add them to the default runlevel:

Now reboot the server and everything should start up correctly.

After the server comes back online verify that the new ports are all associated with the vSwitch.

Which should return output similar to this:

Finally fire up VirtualBox and create a new VM or go to the settings of an old VM and set the network interface to “bridged mode” and choose one of the new switch ports in the list. Fire up the VM and you should have a connection!

References:

[Xen Networking – Xen Openvswitch with Virtualbox](#)

Installing Open vSwitch on Gentoo (Xen Hypervisor)

The Gentoo ebuild for Open vSwitch does not seem to work with the latest available kernel as of this writing (*3.10.7-gentoo-r1*). This post is documentation of the process that I performed in order to successfully install Open vSwitch on a Gentoo server running the Xen hypervisor. This guide assumes that you already have a Gentoo environment configured and running with the Xen hypervisor available in portage.

Note: See the update in the comment section below for how to install openvswitch-2.0.0 from portage!

First make sure the following kernel settings are enabled for full Open vSwitch compatibility:

After you rebuild the kernel and reboot the machine you can load the openvswitch module by typing:

Next add an entry for the openvswitch module to `/etc/conf.d/modules` so it loads on each reboot:

In order to successfully install Open vSwitch it must be downloaded and installed from source. The latest source code can be downloaded [here](#).

In this guide the `openvswitch-1.11.0.tar.gz` file was downloaded and extracted to `/usr/src/openvswitch`. Perform the following commands to build and install Open vSwitch from the downloaded source code.

Open vSwitch should now have files installed in `/usr` and `/var`

The `ovs-*` commands should also now be available in your path

Next it is necessary to create the openvswitch DB

Startup the Open vSwitch database server

Initialize the database

Then start up openvswitch

In order to have Xen use Open vSwitch as its default virtual interface add the following entry to `/etc/xen/xl.conf`

The physical Ethernet interface that will be used with Open vSwitch has to be set to null in `/etc/conf.d/net`

Finally create the first Open vSwitch bridge called `xenbr0`

Note: See the update in the comment section below for how to install openvswitch-2.0.0 from portage!

References:

[How to Install Open vSwitch on Linux, FreeBSD and NetBSD](#)

[Xen Networking – Xen \(Setting up Open vSwitch networking\)](#)

[QEMU with Open vSwitch network](#)

Xen Cloud Platform (XCP) – Cloning Hard Drive Woes

The main hard drive seems to be flaky in one of my XCP servers. I decided to use [Clonezilla](#) to clone sda to another drive to see if it is in fact the hard drive. After cloning over the drive I found that my LVM storage group VG_XenStorage-xxx was not mounting, and XenCenter was giving off the following error when trying to connect to the server: **“This server cannot see any storage”**

Turns out the LVM volume group was inconsistent after the clone, my guess is because the hard drives were of the same capacity but different brands so there may have been some differences. Using the lvs and vgs commands did not show the LVM volume information, but instead displayed a kernel dump with a plethora of information. The main error being about the inconsistency of the volume group. In order to solve the problem I had to perform the following command:

Then restart the xapi service:

Time to see if we can make it crash again with the new drive!

Xen Cloud Platform (XCP) – XenCenter Snapshot Issue

When using XenCenter in conjunction with XCP you will find that you are unable to take snapshots of your virtual machines in XenCenter. XenCenter will spew up an error saying **“Snapshots require XenServer 5.5 or later”**.

This seems very strange since XCP is based on XenServer 5.6.0-fp1. It turns out that the XCP xapi needs to be patched to fool XenCenter into thinking that XCP is 5.6.0. Here is how you do it.

First stop the xapi service:

Then backup the xapi binary:

Now patch the binary with the new version information:

Start xapi:

Thats it! You now should be able to take snapshots of your virtual machines within XenCenter.

Here is a script to automate the process:

Xen Cloud Platform (XCP) –

Assigning A VLAN An IP Address

Once you create VLAN's on your XCP server you will find that the physical interface you used to be able to hit by an IP address no longer responds. This is because you have trunked the interface into multiple VLAN's and need to assign IP's to the VLAN interfaces rather than the physical interface.

In order to do this first remove the static IP from the physical interface if it was assigned one. Then do the following:

(Note this example uses eth1 as the physical interface the VLAN's are assigned to.)

First lets check the settings on the VLAN interface you need to assign an IP to.

This should return a list of the VLAN interfaces parameters for VLAN201

Notice the following parameters are all empty:

Lets assign an IP address to this VLAN:

Now we can confirm the setting by listing the parameters again:

Which should show:

Now try to ping the address it should respond again.

You can also set it to use DHCP instead of assigning a Static

address by using

Xen Cloud Platform (XCP) – Setting Up A VLAN

One very useful feature of XCP is the ability to setup VLAN networks for your virtual machines to use. This gives an administrator fine grained control on what network a machine belongs on. This will work as long as the network interface that is assigned to your VM's is plugged into a trunked port on a switch that has been setup with VLAN tags.

In order to start assigning VLAN's to a physical interface or pif in Xen terms do the following:

Get a list of the physical network interfaces and their corresponding UUID's:

This should return something similar to this:

This shows us the 2 network interfaces that are installed in the server. ETH0 is assigned as the management interface and ETH1 is assigned to the virtual machines.

Now we need to create a new network for our VLAN:

This creates a new network named network201 that corresponds to the VLAN tag 201 on the switch. When running this command you will get a uuid as output.

Now we need to assign the VLAN tag 201 to this network, and

bind it to a physical interface (eth1 in this case):

That's it the new VLAN201 is created and can be selected for use by your virtual machines. You can confirm its presence by doing:

Which now will display your 2 physical interfaces as well as the new VLAN201 interface:

Xen Cloud Platform (XCP) – Add A Disk To The Default LVM Volume Group

If you initially install XCP using only one disk you can still add more drives later if you find you need more room. XCP uses LVM storage to manage the partitions created for new Virtual Machines which allows you to extend the default volume group at your leisure among other things.

In order to add a new disk to your default volume group perform the following steps.

Install the new drive into the computer

Partition it using fdisk

After you create the partition change its type to Linux LVM:

Enter the menu option T to change the partition type:

Enter the menu option 1 to choose the first partition:

Enter the menu item 8e to choose Linux LVM:

Enter the menu item p to print the partition output to verify that it is set to Linux LVM:

Enter the menu item w to save the partition information:

Enter the menu item q to quit:

Now we can add the partition to the volume group.

Set the partition as a physical volume:

Find out the name of the existing volume group:

This should output something like this:

We are interested in this line:

Add the new physical volume to the existing volume group

Now you can verify that it worked by typing:

This should list the volume group and display it's current size, which should indicate the size of the old volume group + the size of the added disk.

Xen Cloud Platform (XCP) – Local ISO Storage

Repositories

XCP supports local ISO storage on an internal hard drive. It is not available as an option in the xsconsole – Disk and Storage Repositories menu. However you can set one up with a little bit of command line magic.

The best way to do this is by using a second hard drive in the server. Once installed use fdisk to create the partition.

Proceed to follow the prompts and create a new partition using the ext3 file system.

Once the partition is created format it:

Now make the directory that will be the mount point for the volume.

Mount it.

Then add the following entry in /etc/fstab so that the partition is mounted at boot.

Now let Xen know it's there.

That's it. You can now place ISO images in /var/opt/xen/iso_import and they should automatically show up in XenCenter under local storage. You will also be able to use them to install new virtual machines via the drop down menu.