DEF CON 24 Presentation: VLAN Hopping, ARP Poisoning, & Man-in-the-Middle Attacks in Virtualized Environments

---

DefCon 23 Presentation: Exploring Layer 2 Network Security in Virtualized Environments

---

DefCon 23 Video Demo: MAC Flooding on Citrix XenServer 6.2 with Open vSwitch 1.4.6

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

---

Exploring Layer 2 Network Security In Virtualized Environments – DerbyCon 4.0

I gave a talk this past weekend on part of my Ph.D. dissertation research at the DerbyCon 4.0 “Family Rootz” Computer Security conference in Louisville, KY. Take a look at the following video to view the talk in its entirety. The rest of the conference videos are available here.

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) – 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):

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

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