

Virtualization and High Availability with Linux

Matteo Redaelli
2005-07-14

XEN virtual machine monitor

“Xen is a virtual machine monitor for x86 that supports execution of multiple guest operating systems with unprecedented levels of performance and resource isolation. Xen is Open Source software, released under the terms of the [GNU General Public License](#). We have a fully functional ports of Linux 2.4 and 2.6 running over Xen, and regularly use it for running demanding applications like MySQL, Apache and PostgreSQL. Any Linux distribution (RedHat, SuSE, Debian, Mandrake) should run unmodified over the ported OS.”.



Xen supporters are:



Operating System and Systems Management



Hardware Systems



<http://www.cl.cam.ac.uk/Research/SRG/netos/papers/2004-xen-ols.pdf>
<http://www.cl.cam.ac.uk/netos/papers/2005-xen-may.ppt>

Setup

SO	Linux Debian 3.1 (sarge)
Hostname	sarge
IP	Dhcp (151.10.39.195)
SW	Xen v. 2.0.6 http://www.cl.cam.ac.uk/Research/SRG/netos/xen/
/etc/xen/xend-config.sxp	(xend-port 8000) (xend-event-port 8001) (xend-address 'localhost') (console-port-base 9600) (console-address 'localhost') (network-script network) (vif-bridge xen-br0) (vif-script vif-bridge) (vif-antispoof no) (block-file block-file) (block-enbd block-enbd)
/boot/grub/menu.lst	...

	<pre> title Xen 2.0 / XenLinux 2.6 kernel /boot/xen-2.0.gz dom0_mem=131072 module /boot/vmlinuz-2.6-xen0 root=/dev/hda1 ro console=tty0 ... </pre>
/etc/xen/auto/sarge1	<pre> kernel = "/boot/vmlinuz-2.6.11-xenU" memory = 64 name = "sarge1" vif = ['mac=AA:00:00:58:4F:06, bridge=xen-br0'] disk = ['file:/home/matteo/xen_images/sarge1,sda1,w', 'file:/home/matteo/xen_images/sarge1_swap,sda2,w', 'file:/home/matteo/xen_images/sandisk,sda3,w'] root = "/dev/sda1 ro" dhcp="dhcp" hostname= "sarge1" extra = "2" </pre>
/etc/xen/auto/sarge2	<pre> kernel = "/boot/vmlinuz-2.6.11-xenU" memory = 64 name = "sarge2" vif = ['mac=AA:00:00:2E:5D:7E, bridge=xen-br0'] disk = ['file:/home/matteo/xen_images/sarge2,sda1,w', 'file:/home/matteo/xen_images/sarge2_swap,sda2,w', 'file:/home/matteo/xen_images/sandisk,sda3,w'] root = "/dev/sda1 ro" dhcp="dhcp" hostname= "sarge2" extra = "2" </pre>

Install the sw as suggested by the UserGuide (with “cd xen-2.0-install && ./install.sh” and grub configuration). Reboot del server and check if new kernel (“Linux sarge 2.6.11.10-xen0”) is running.

Following the UserGuide,

- I created the file system for the virtual machines from files
- I installed the virtual machine “sarge1” with the command “debootstrap”
- I created the second virtual machine “sarge2” with a simple copy file (“cp -a”) of sarge1

The result is this:

```

sarge:/home/matteo/xen_images# ls -l
total 1435020
-rw-r--r-- 1 root root 209715200 2005-06-30 16:13 sandisk
-rw-r--r-- 1 root root 524288000 2005-06-30 16:04 sarge1
-rw-r--r-- 1 root root 104857600 2005-06-29 19:55 sarge1_swap
-rw-r--r-- 1 root root 524288000 2005-06-30 16:04 sarge2
-rw-r--r-- 1 root root 104857600 2005-06-29 19:55 sarge2_swap

```

Notes:

- I set the MAC address of all network cards of virtual machines due to get the same ip address from DHCP server after each reboot of the servers. Otherwise the MAC address is randomly generated.

- The “sandisk” file (system) will be used to simulate a file system on a Storage Area Network (SAN): it will be mountable (read/write) from servers “sarge1” and “sarge2” but it will be mounted exclusively only from server that run apache (Section “linux-ha”, test2)

Start the xen service with
`/etc/init.d/xend start`

And start all the virtual machines with
`/etc/init.d/xendoms start`


To see all running virtual machines
`xm list`

To start a single a virtual machine (sarge1) run
`xm create /etc/xen/auto/sarge1`

and to stop it
`xm shutdown`
or “shutdown -h now” inside sarge1

Linux-HA

“The Linux-HA project is widely used, and is an important component in many interesting High-Availability solutions. We currently have a few thousand installations up in production in the real world. We also work well with the [LVS \(Linux Virtual Server\)](#) project and expect to continue to collaborate with them in the future, since our goals are complementary. Interest in this project is growing very rapidly. These web pages are averaging many thousands of hits per day, and we see several thousand downloads of the software per month”



“The Heartbeat program is one of the core components of the Linux-HA (High-Availability Linux) project. Heartbeat is highly portable, and runs on every known Linux platform, and also on FreeBSD and Solaris. Ports to other OSes are also in progress.”

“Heartbeat is the first piece of software which was written for the Linux-HA project. It currently performs death-of-node detection, communications and cluster management in one process. It does the first of these two fairly well. However, its cluster management functions leave a lot to be desired.”

Linux-ha is used all over the world: see <http://www.linux-ha.org/SuccessStories>

Setup

	Sarge1	Sarge2
OS	Linux Debian 3.1 (sarge)	Linux Debian 3.1 (sarge)
IP	151.10.39.231 (via dhcp)	151.10.39.234 (via dhcp)
SW	Heartbeat (v. 1.2.3) http://www.linux-ha.org/	
/etc/hostname	sarge1	sarge2
/etc/hosts	127.0.0.1 sarge1 localhost 151.10.39.234 sarge2	127.0.0.1 sarge2 localhost 151.10.39.231 sarge1

Test1 (vip + network disk)

/etc/ha.d/haresources	sarge1 151.10.39.246 Filesystem::/dev/sda3::/sandisk::ext3
/etc/ha.d/ha.cf	logfacility daemon node sarge1 sarge2 keepalive 1 deadtime 10 bcast eth0 #ping 151.10.39.1 auto_failback off respawn hacluster /usr/lib/heartbeat/ipfail
/etc/ha.d/authkeys	auth 1 1 sha1 d8f5bad01e76ab90ae864e1f2d31679e

To put a node in standby you can run

/usr/lib/heartbeat/hb_standby

Test2 (apache)

/etc/apache2/sites-available/default	Alias /sandisk/ "/sandisk/" <Directory "/sandisk/"> Options Indexes MultiViews FollowSymLinks AllowOverride None Order deny,allow Deny from all Allow from 151.10.0.0/255.255.0 ::1/128 </Directory>
/etc/ha.d/haresources	sarge1 151.10.39.246 apache2 Filesystem::/dev/sda3::/sandisk::ext3

Disable the automatic start of apache with

```
cd /etc/rc2.d/ && mv S91apache2 _S91apache2
```

and create a test file /sandisk/index.html

```
<html>  
<body>  
Hello from an HA Apache installation  
</body>  
</html>
```