

BLU Meeting Notes: Wed, Mar 21, 2007

Linux Soup X: Red Hat 5 Virtualization

Written by Martin Owens

iSCSI – Targets available for RH4U2, RH5, RH6 using a combination of open/linux-iscsi;

Virtualisation, Utiliation normal hardware and configuration advantages; Support for changing hardware around, maintenance of hardware and using software that may not be quite suitable for the hardware platform; Use less hardware by running more than one platform. Virtualisation requires hardware extensions in the processor and host operating system.

- 1) Single Kernel Image (SKI), good and bad points where a multiple virtualisation platforms share the same kernel. Speed advantages but at the cost of flexibility
- 2) Full Virtualisation, Everything translated e.g. Vmware, mvs; the software controls all the platforms, penalty of about 10% speed to guests but the advantage is that the guests are not aware they are running in a virtual environment and hardware incompatibilities are mitigated by the virtualisation. Host kernel runs in privileged position.
- 3) Para-virtualisation, Each platform cooperatively talks to the virtualisation system to improve performance. The down point is that it wont work with unsupported guest platforms. e.g. Xen.

PAE – older 32 cpu can only address 4GB of RAM, PAE allows you to have more RAM by turning it into a paged system controlled by the cpu. Can check the capability via /proc/cpuinfo or through intel website when you are looking to prospectively buy hardware. Mostly used for para-virtualisation. Older versions of Linux will support non PAE systems.

A lot of laptop makers have hardware which support VTX/I (virtualisation technology, x86/Itanium) but are required to be switched on in the BIOS, unfortunately most don't provide a way to do this. Some laptops provide full VT support and this isn't required to be switched on.

Paused migration where 1 guest is paused, moved and then re-enabled. So long as both hosts have access to the same nfs disk the migration goes well.

Live migration uses the current running system to copy over the guest to another machine while running. The machine is brought to a lower performance in order to finish off the move so as not to incur a race condition with new memory and io updates on the old host server. There is a small gap where the old guest is disabled and the new guest is brought back. In effect a video could be running on the guest platform while the migration is taking place and the video will continue to run without stopping.

No real GUI yet for virtualisation when doing migrations, although much is planned since all of the tools currently take advantage of cli tools so the gui tools should be too far away. (compete with vmware administration tools?)

(note) some fears that HV has a security flaw where the privilege escalation allows possible root kits, new virtualisation machines should probably be installed by the owner for security.

With the configuration file for the live virtualisation you can copy the configuration to create a new virtualisation guest by changing minimal parameters. (no tool yet to automate this task)

Hard drive sectioning can use an entire disk or a simple file. Simple files allow LVM control including expanding the disk as required; allowing you to kick start a guest by over committing the disk the beginning of the guest file is stored on (feature/bug).

When the guest crashes logs can be found in /var/xen/ allowing you standard ways of understanding what may have gone wrong. When machines are running in full the server can run a small good kernel in order to dump logs to the system or even network options.

The limitations of the guest and host server software is not to do with the software but the software support. it's possible to overcome marketing limitations with enough self support.

GFS – clustering file system where virtual guests can access the clustered file system. Better performance than nfs and includes locking to stop each computer sharing the same space from corrupting each other, gfs is more like a mountable disk that is elsewhere; nfs is more like a file communication which is mounted like a partition but the requests are not handled like disk io (maybe confused with iSCSI check).

New redhat version has some new tools that we request be shown, including compvis 3D desktop.