

Cloning Linux Images on VM

Kris Van Hees – Sine Nomine Associates

Session 5546/9370

March 8th, 2002

Decide what you want...



- Determine major directory hierarchies that can be moved to separate minidisks (e.g. /usr, /opt, ...):
 - Writable files in the hierarchies can be moved to /var/local/<original location>, and original location can be made a symbolic link.
 - Some hierarchies are trickier: /bin, /lib, ...
- Determine the VM configuration for clones: what disks to link (RO) vs. what disks the clone will own (RW).

Two scenarios are now possible: RO or RW root fs!

Using a RW root file system: Setup



- Create the CP directory entry for the clone (use DIRM ADD ... LIKE ... to make life easier):
 - Add links to the RO minidisks owned by the master (including a 191 with a PROFILE EXEC to do any setup work and IPL the root minidisk).
 - Add RW minidisks for root, swap and /var.
- Use DDR (or CMSDDR) to create the initial root, swap and /var minidisks for the clone.
- Start the clone (XAUTOLOG, ...)

Using a RW root file system: Configuration



- The clones have a 191 minidisk that is a RO link to a disk at the master.
- Configuration information can easily be stored in a CMS NAMES format file (easy to parse in Linux):
 - Hostname (= VM userid?)
 - IP address, netmask, gateway
 - Disk list (if non-standard)
- Use the cmsfs toolset to access the configuration data.

Using a RW root file system: Finalizing the clone



- Execute `/etc/rc.finalize` from `/sbin/init.d/boot` (if found):
 - Use `hcp` to identify the VM (Q USERID)
 - Read configuration file with `cmscat`
 - Configure Linux
 - Remove `/etc/rc.finalize`
 - Reboot

Using a RO root file system: Setup



- Root file system contains RW and RO portions that are needed at boot time (before mounting other file systems):
 - /dev and /etc are RW
 - /bin and /lib are RO
- Two options:
 - Use statically linked utilities at boot time.
 - Use the initrd feature in Linux.

Using a RO root file system: initrd



- Initial root file system at boot time, loaded as ram disk.
- Only needs to contain `/dev` and `/etc`.
- After executing `/linuxrc` (in `initrd`):
 - New root is created as `/`.
 - Old root is mounted as `/initrd`.
- RO root should contain links to `/initrd/dev` and `/initrd/etc`.
- Changes to `/dev` and `/etc` will be lost – use the master to do that.

Using a RO root file system: Side-effects (good and bad)



- The clone cannot make permanent changes to /dev and /etc.
- Updating configuration, boot parameters, kernel, ... can be done on the master (clones just need to reboot).
- Clones are almost identical to the master.
- More flexibility can be gained from separating pre-mount configurations from post-mount configurations:
 - /etc/foo -> /initrd/etc/foo [pre-mount]
 - /etc/bar -> /var/local/etc/bar [post-mount]

Contact information



Kris Van Hees

Sine Nomine Associates

kvanhees@sinenomine.net

+1 703 723 6673