

Fakultät für Physik

Klaus Steinberger • Ralph Simmler • Susanna Maurer

Redundant Server-Cluster with XEN Virtualisation





- Requirements
- The Hardware
- Cluste Suite and Xen
- Cluster Quircks and Hints
- XEN Quircks and Hints
- Future plans




■ Requirements

- Number of Users is growing as well as their requirements, so we need redundancy
 - Survive Hardware failures as much as possible
 - Minimize downtimes for maintenance
- Scalable
- Good Cost Factor
- Must support Linux/Unix Workstations -> NFS (V4)
- Must support Windows Workstations -> CIFS (Samba)
- A bunch of Services needed :
 - DHCP
 - Directory Services (Novell, AD)
 - DNS
 - Printservice
 - And more




■ Possible Solutions:

- Cluster + FC San Storage
 - We have already some Expertise with SAN
 - Redundant FC/Storage Solutions are available with good pricing
 - High performance
 - 4 Gbit -> 400 Mbyte/sec even more with „multibus“
 - Further Expansion is relative cheap.
- Cluster + Iscsi Storage
 - Some Expertise with Software Target
 - Network ist cheap
 - But! Only few redundant Solutions available
 - The redundant solutions are always fully mirrored, and performance suffers both from 1GBit Ethernet and mirroring latencys!
- Cluster + Netapp -> Storage very expensive, expansion probably also very expensive, far out of our budget



- OS: SL5.1 with Xen
- Alle Storage Paths are redundant
- RX300 was our old webserver
 - Will be replaced soon
 - Runs out of Service
 - Trouble with migration
- 3 Ethernet Connections
 - 1 GB/s Production Net
 - 100 Mbit/s Clustercommunication and IPMI
 - 100 Mbit/s LOM



- OS: SL5.1 64bit with Xen
- At the time the Dell's arrived we didn't plan to cluster.
- RX300S3 was later added, less space (2HE), less power requirements, but much faster!
- Anything is redundant except controller in our old T6100's
- Important data on Provigo
- 3 x Ethernet per Node
 - 1 Gbit/s Production Net
 - 1 Gbit/s Clustercommunication
 - 100 Mbit/s LOM (Drac4 on Dell, iRMC on FSC)



- Three node cluster, Quorum „2“ - 1 node can fail
 - 2 Node is possible too
 - special flags needed
 - fencing devices must be on same net as Cluster Communication
 - we used it through migration phase
 - More nodes are possible, quorum must be 50% of nodes + 1
 - No Quorum disk necessary
- We don't use luci/ricci – handediting of cluster.conf
- Fencing via IPMI – we don't use FC fencing
- We use rgmanager resource just for VM's
- On failure of a node the VM's will be restarted on the other nodes – short failure time < 1 Minute
- Clvmd handles the logical volumes



- Nearly all Services run in independent VM's
 - Advantages:
 - Migration is possible since SL 5.1 (5.0 did restart)
 - Maintenance on a Server could be done without Downtime
 - Strong Separation of Services reduces Dependencies
 - Additional Servers could be setup very fast.
 - Windows is possible with HVM
 - Rgmanager restarts failed VM's reliable (both on Node and VM failure)
 - Disadvantages:
 - More RAM necessary (at least 200MByte per VM)
 - The available RAM could only be used $n-1$ – but RAM is cheap nowadays
 - Some more Diskspace necessary (per Instance min. 4 Gbyte)
 - More Servers to maintain (even if they are virtual)
- Services running on dom0:
 - iSCSI Target - Failover is done on Client side with multipathd
 - NTP Server (DMZ) – no failover necessary, timekeeping probably more exact



- There is just one LUN from every RAID Set
- Anything else is done with LVM
- Images of VM's are Logical Volumes
- Fileservices:
 - We do not use clustered Filesystems – bad experience with GFS at that time
 - In most cases a Logical Volume containig a ext3 FS is connected to a VM
 - In some cases (most important filesystems – home etc.) LVM is setup inside the VM again, so we can use snapshots (not yet in production as we have some troubles with TSM Client and snapshots)
 - NFS and Samba must reside on same VM
 - Scaling of NFS is done by distributing „shares“ over VM's



- A Compute cluster was setup with GFS some years ago (SL 4.x)
- First setup of GFS was slow, slow, slow
 - Df was painful slow
 - Tivoli backup sometimes needed more than 2 days or even died
 - Big files were somewhat ok
- Most Tuning has to be done on creation!
 - Glock tuning is important – solved our backup problem
http://people.redhat.com/wcheng/Patches/GFS/readme.gfs_glock_trimming.R4
 - Number of Resource Groups must be reduced on large filesystem
(Tbyte range)
- GFS suffers with small record/filesizes



■ The Xen Instances we run (Homedirectory Server):

- DHCP -
- DNS - Domain Nameserver
- EDIR - Edirectory Replica(s)
- LTSP - Bootserver for Thinclients
- NFS(1-n) - NFS Server(s)
- LOGIN - Login Server for Thinclients and external Access
- LICENSE - Flexlm Server (MAC cloned from old one)
- CUPS - Cups Print Server
- More coming



■ The Xen Instances we run on our DMZ cluster

- Xen-static - A Bunch of Websites with Static Pages (also Fiona)
- Xen-tikiwiki - Some tikiwiki Sites we have
- Xen-typo3 - Some Typo3 Sites
- Xen-mediawiki - Wiki Sites
- Xen-Elog - E-Log Sites
- Xen-scientific - Mirror for Scientificlinux
- Homepages - User Web Homepages
- Svn - Subversion Repository
- Lists - Mailman List
- OTRS - OTRS Trouble Ticket System
- Webmail - Horde Webmail
- Mailin - Incoming Mail
- Mailout - Outgoing Mail
- IMAP - Cyrus Imapd (will move onto cluster next weekend)



- Clustercommunication should be on separate LAN
 - We had a flapping cluster with cluster communication on same net as xenbr
- Switch reboot breaks cluster
 - Our HP Procurve need some minutes for rebooting, this breaks the cluster!
 - Have at least a dumb switch at hand for planned reboots of switch
- Clvmd sometimes has a lock hung
 - On some occasion needed a complete reboot of cluster
 - hopefully better with latest clvmd
- No snapshot with clvmd yet (maybe with 5.2 or 5.3 ...)
 - Solution A: Export a SAN Lun to domU and make snapshot on storage
 - Solution B: Export a logical Volume to domU, make a Volume Group inside domU, and make snapshots there



■ Why Ballooning is not good

- Dom0 sometimes does not balloon out fast enough – Migration failed
- Heavy network traffic could kill the network driver – any domU loses communication! (problem was at least in 5.0)
- Solution: Limit dom0 Memory!

```
Grub: „kernel /xen.gz-2.6.18-53.1.14.el5 dom0_mem=1024M“  
xend-config.sxp: „(dom0-min-mem 1024)“
```

■ Reboot of a Cluster Node: 5.1 always tries migration

- Modify xend-config.sxp to allow migration! Otherwise vm's will be in migrate state endless.

```
xend-config.scp:  
(xend-relocation-server yes)  
(xend-relocation-port 8002)  
(xend-relocation-hosts-allow '^localhost$ ^localhost\\.localdomain$  
  ^yourhosts$)
```



- Some network daemons did fail after migration:
 - After migration NFS Connections start hanging (not stale! just hung)
 - Default for domU network: flipping receive path
 - Use copying path instead:
 - Add 'extra = „xennet.rx_copy“' to vm definition
 - But 2 times slower as flipping (3 Gbit/s instead of 6Gbit/s on RX300S3)
 - I found that by occassion – nowhere documented!
- How to make migration faster
 - Migration runs over same net as cluster communication – use Gbit Ethernet
 - Use live migration – just a few seconds unresponsive even with 100MBit
 - Add „-l“ to line 338 of /usr/share/cluster/vm.sh
 - err=\$(xm migrate -l \$OCF_RESKEY_name \$target 2>&1 | head -1)



■ Bonding did not work

- With SL 5.0 we had no success with Bonding – xenbridge did loose packets
- Mileage could be different with 5.1, we hesitate to test it on production cluster

■ PVM driver for Windows domU do not work

- With Windows HVM networking is painfully slow (especially TX path)
- Both the Novell and the emerging GPL PV drivers were not able to hide the boot device – doubled C: drive – fatal filesystem corruption
- Maybe SL 5.2 will solve that???
- Possible Alternative: Use of XEN3.2 hypervisor from xen.org
(There is an CentOS RPM)



■ Read Starvation

- Observed this problem since years, heavy writes did starve our NFS Servers
- First impression was some NFS Clients are the reason (old RH9 nodes?)
- But real reason is the default „CFS“ Completely (Un?)Fair Scheduler
- TUV seems to do benchmark tests just with Oracle, CFS seems to be good for Oracle, but is definitely bad for any other interactive workload
- Solution: Switch to „noop“ or better „deadline“
Add „elevator=deadline“ to kernel commandline
- For test you could switch online:
echo „deadline“ > /sys/block/sda/queue/scheduler (with SL 5.x)
- Response times got noticeable better with „deadline“



■ Multipath /Qlogic driver

- Don't forget to change parameters for qllogic driver or you will definitely loose data on failover:

```
options qla2xxx qlport_down_retry=1
```

- Also the following defaults in multipath.conf are useful:

```
defaults {  
    user_friendly_names yes  
    default_features "1 queue_if_no_path"  
}
```

■ In the NFS Servers fix the „fsid“

- Linux NFS calculates fsid from Major/Minor Number
- In LVM enviroments Minor Number changes! Serverreboots will led to stales!
- Use „fsid“ option in exports
(don't do it on already exported fs! different magic numbers!)
- Alternatively: „lvchange --persistent y --major m --minor n“



- Cluster in Garching will be upgraded to 10GE sometime
- We are currently planning a new Server Cluster
 - Homedirectories for whole faculty of physics – 2500 Users
 - Initially 20 TByte
 - Probably be again a FC Storage, currently inviting offers
 - Cluster Structure will be same as the running systems
 - 10GE at least in a second phase
- Maybe now a clustered filesystem? Suggestions?
- Hopefully we could solve the Windows PV driver problem, as we also need to add some Windows VM's (for example AD)



- Questions ?