

A Highly Versatile Virtual Data Center Ressource Pool

Benefits of XenServer to virtualize services in a virtual pool

Stefan Bujack

A Highly Versatile Virtual Data Center Ressource Pool
Umeå, 27.05.09

- > The way to virtualization
- > Benefits of XenServer
 - Shared NetApp storage
 - High Availability
 - XenCenter GUI
- > Our current XenServer environment
 - The topology
 - Installation Methods for PVMs
 - Installation Methods for HVMs
- > Problems we encountered
- > Summary and future



The way to virtualization

- > The situation was that many services ran on dedicated servers
 - Waste of resources, no efficient usage of hardware
- > We saw virtualization as solution to save resources
 - Less hardware and network topology saves money
- > Opensource Xen was already in use
 - Lacks high availability
- > XenServer offered some features to handle a secure hosting of services
 - Cheaper than competitors
 - Roots in the Linux community



Benefits of XenServer

- > Creating virtual pools to optimize the allocation of VMs on the servers
- > Shared NetApp storage in a pool
 - Snapshots and templates
- > XenMotion
 - Zero downtime migration of PVMs from one server to another
- > High Availability
- > Workload Balancing and Live Scale Out
- > XenCenter GUI
 - Single instance management console
- > Windows, Linux and Solaris VMs
 - 32bit and 64bit
- > Support by hardware vendors
 - DELL, NetApp, ...



Shared NetApp storage

- > Shared iSCSI storage in a pool
- > Use of NetApp functions via direct Ontap access out of XenServer
- > Easy adding of additional disks and changing disk sizes on the fly
- > Secure storage for virtual disks through RAID DP and controller failover
- > Deduplication technique from NetApp
 - Drastically reduces image space
- > Moderate I/O rates to run VMs which don't have many parallel jobs

Front



Rear



High Availability

- > Heartbeat Cache: 1GB shared iSCSI storage per pool
- > Heartbeat storage caches the work flow data of the VMs in the pool
- > Possibility to set different protection levels for VMs
- > Important VMs will restart automatically on other hardware after crash
- > Free resources must be available on other servers to take over VMs

High Availability

Configuration

Pool HA enabled: Yes

Configured failure capacity: 1

Current failure capacity: 1

Heartbeating status

Virtual machine	Description	Protection level	Agile?
it-dcot-dev		Protected	✓
it-lcg3d-ctl	Instance to control ...	Protected	✓
spamfilter1		Protected	✓
spamfilter1-Clon...	NIICHT GLEICHZEIT...	Do not restart	✓
spamfilter2		Restart if possible	✓
spamfilter3		Restart if possible	✓
spamfilter4		Restart if possible	✓

Protect

Restart if Possible

Do Not Restart

New failover capacity

1

The number of server failures that could be tolerated in this pool given these HA protection levels is 1.



XenCenter GUI

- > The XenCenter provides an overview of all pools, servers and VMs
- > Everyday operations can be done via mouse click
 - ISO-mounting
 - Rebooting
 - Installation
 - Cloning
 - Disk operations
 - ...
- > Drag and drop of VMs from one hardware to another
- > System monitoring automatically shows operating parameters
 - CPU, memory, disk access and network bandwidth
- > Access to VMs via local console when ssh and rdesktop is not possible



The screenshot displays the XenCenter web interface. The left sidebar shows a tree view of the resource pool structure, including servers like 'it-xen011' and 'it-xen008', and various VMs such as 'bibwiki', 'wofwiki', and 'ippedmstea'. The main panel shows the 'DMZ-Pool-v5 Overview' with a table of resource usage for each VM.

Name	CPU Usage	Used Memory	Disks (avg / max KBs)	Network (avg / max KBs)
DMZ-Pool-v5 Default install of XenServer	-	-	-	-
it-xen011 Default install of XenServer	13% of 8 CPUs	96% of 16 GB	-	61/122
bibwiki	0% of 1 CPU	100% of 1 GB	16/16	0/0
ippedmstea	0% of 1 CPU	32% of 1 GB	4/4	0/0
it-xen	0% of 1 CPU	100% of 512 MB	6/6	0/0
wofwiki2	1% of 1 CPU	100% of 1 GB	16/16	0/0
wofwiki3	100% of 1 CPU	100% of 1 GB	76/76	92/92
wofwiki5	0% of 1 CPU	100% of 1 GB	24/24	0/0
wofwiki7	0% of 1 CPU	100% of 1 GB	6/6	0/0
wofwiki8	0% of 1 CPU	100% of 1 GB	3/3	0/0
wofwiki9	0% of 1 CPU	100% of 1 GB	1/1	0/0
wofwiki10	0% of 1 CPU	100% of 1 GB	6/6	0/0



Our current XenServer environment

18 Hardware servers

14x DELL Blade Server 1955

2x Intel Quad Core 2,5 GHz

16 GB memory

4x DELL R805 Server

2x AMD Quad Core 2,5 GHz

64 GB memory

90 VMs

CentOS 4.5 - 5.3

Debian 4 and 5

Oracle EL 5.X

RHEL 4.5 - 5.3

SLES 9 - 11

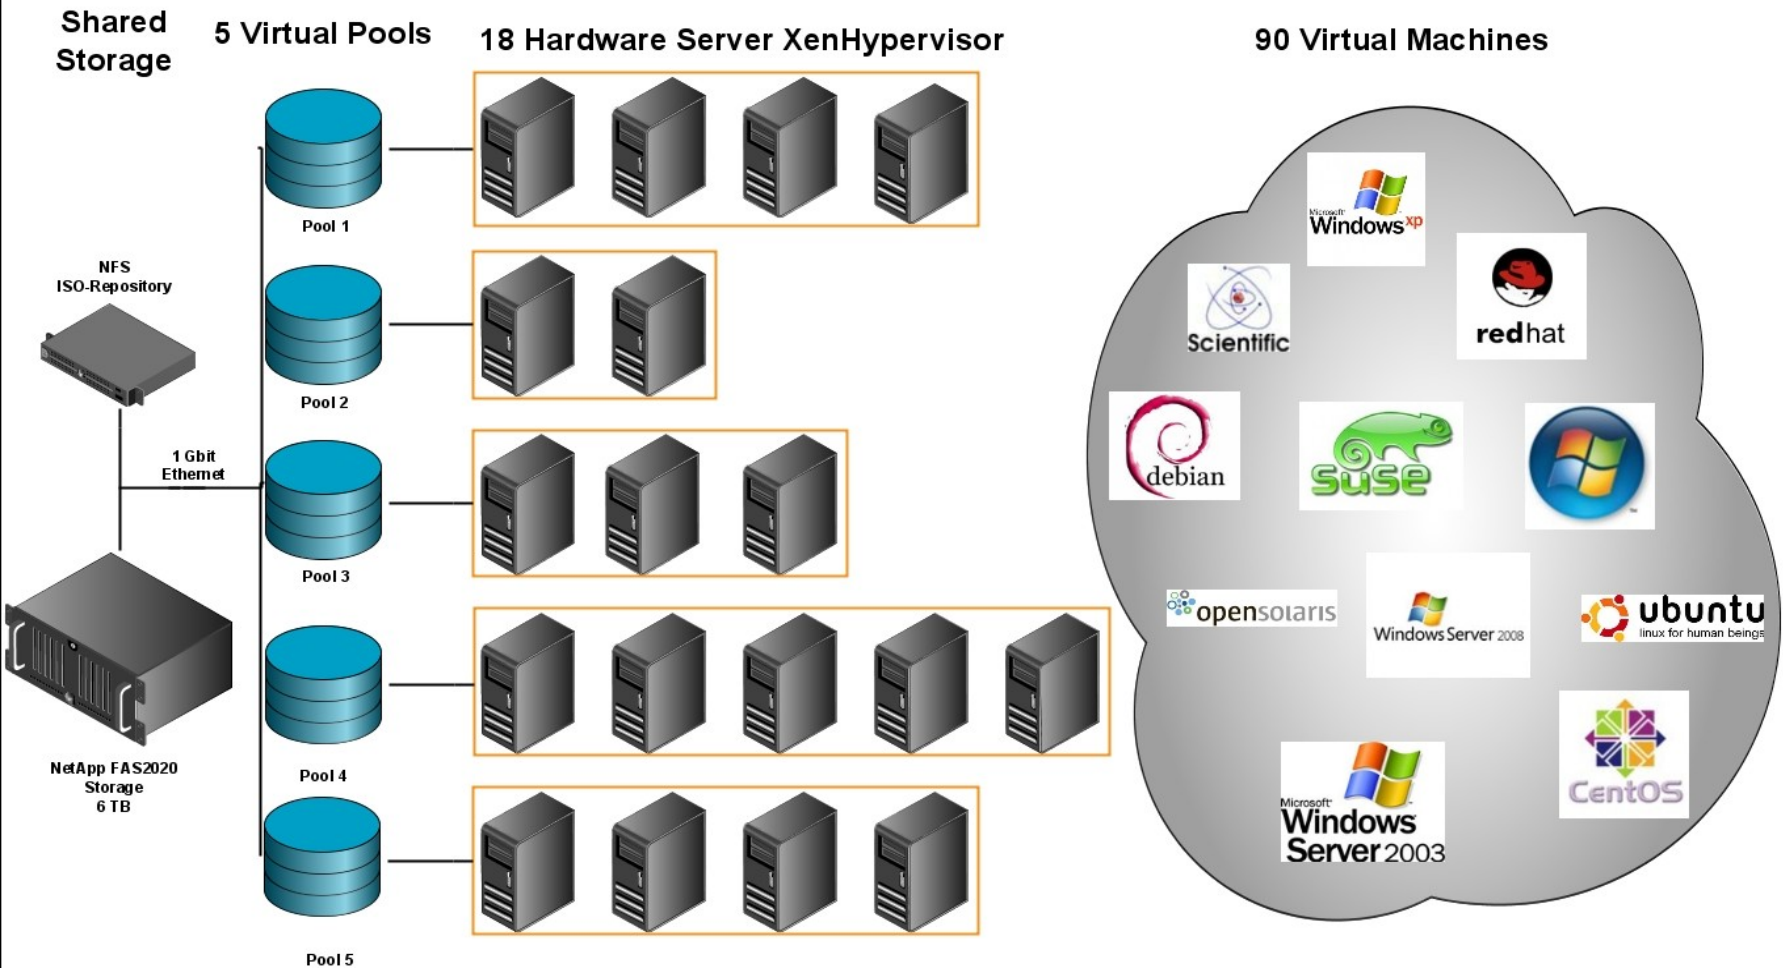
SL 4.5 - 5.3

Windows Server 2003 2008

Windows 2000 XP Vista



The topology



Installation Methods for PVMs

- > Paravirtualisation is our preferred approach
 - Required for HA features
- > Para-virtual machines (PVMs) need some modifications to function
 - Specific boot-loader, partition layout, kernel, Xen daemons and drivers
- > Citrix templates can be used to install preconfigured PVMs
- > Installation of Scientific-Linux-DESY PVMs possible by using the kickstart files and templates
- > Automatic installation of DESY Windows PVM via RIS or WDS
- > Converting a prototype VM to a template to create clones
- > Cloning VMs from template in 50sec until login
- > Take snapshot and recover VM from it in 1min until login



Installation Methods for HVMs

- > Only if PVM does not work
 - E.g. Fedora 11, Solaris
- > Hardware assisted virtual machines (HVMs) can be installed like any physical system via ISO or network boot
- > All major operating systems can be installed as HVM
 - As long as they fit the chip architecture
- > HVMs don't use Xen API
 - Can not use XenMotion and High Availability
- > Solaris can only be installed as HVM in XenServer right now



Problems we encountered

- > First version of XenServer 3 and 4 didn't fulfill the promises
 - No mature storage concept, no High Availability
- > Integration in our installation environment was not easy accomplished
 - The right use of kickstart files, templates and repositories is tricky
- > Finding the right storage for a virtual pool
 - NFS has small operation possibilities, iSCSI showed big timeouts in a VM
- > Not possible to mix different CPU architectures in a pool
 - Due to different virtualization implementations by AMD and Intel
- > The license model for XenServer has changed in april this year
 - The Server without NetApp storage and high availability options is free now
- > Citrix's support model differed from XenSource's
 - No direct support from Citrix, only via forum
 - Resellers are providing support for XenServer



Summary and future

- > We use XenServer in combination with a virtual network store
 - Consolidated 90 machines on 18 servers
 - Saved space, money, and energy
 - Gained high availability and disaster recovery
- > We virtualize a wide range of services on various operating systems
- > XenServer pools can easily be managed and extended
 - We run it with two technicians (part time), less than 0.5 FTE
- > We haven't virtualized all applicable services yet
 - We will grow bigger and host more services



Thank you for your attention

Questions?

