



# **XEN-based server cluster with iSCSI- storage**

# cluster ?

- here: *not* a cluster for parallel computing!
- 2 or more cluster servers
- 1 storage for all cluster servers (san)

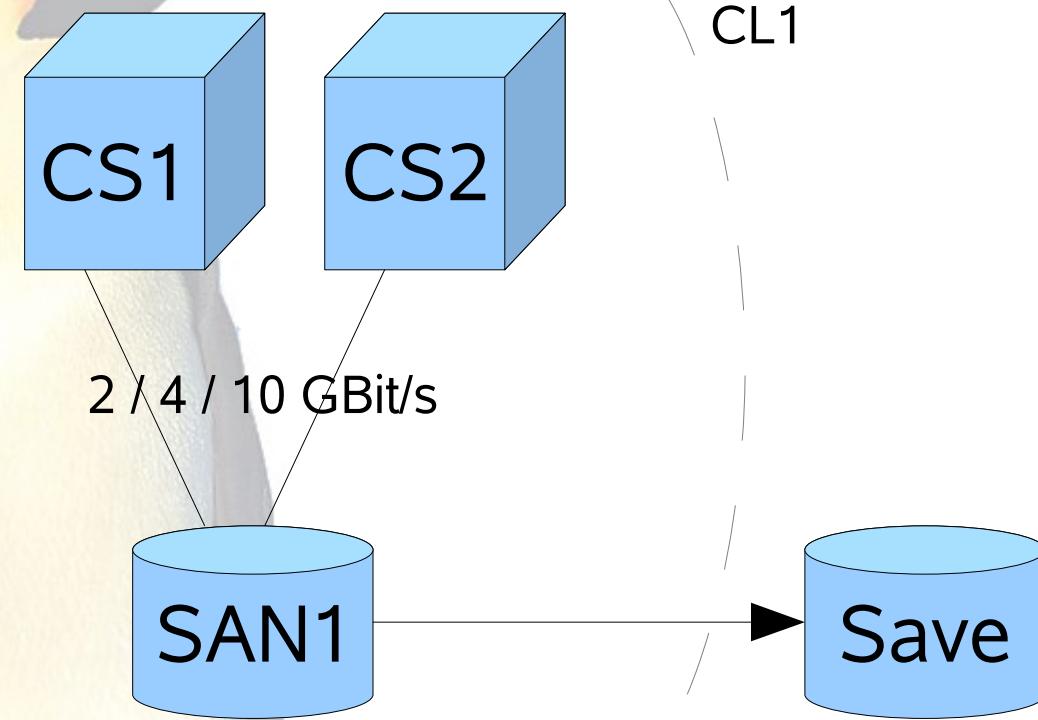


7 HE  
12 TByte  
12 CPUs  
48 Cores  
192 GB RAM  
10 GBit Switch

# targets of a virtualization cluster

- higher availability
- computing power as necessary
- consolidate physical servers
- better hardware usage
- flexible expansion
- save energy
- easier maintenance

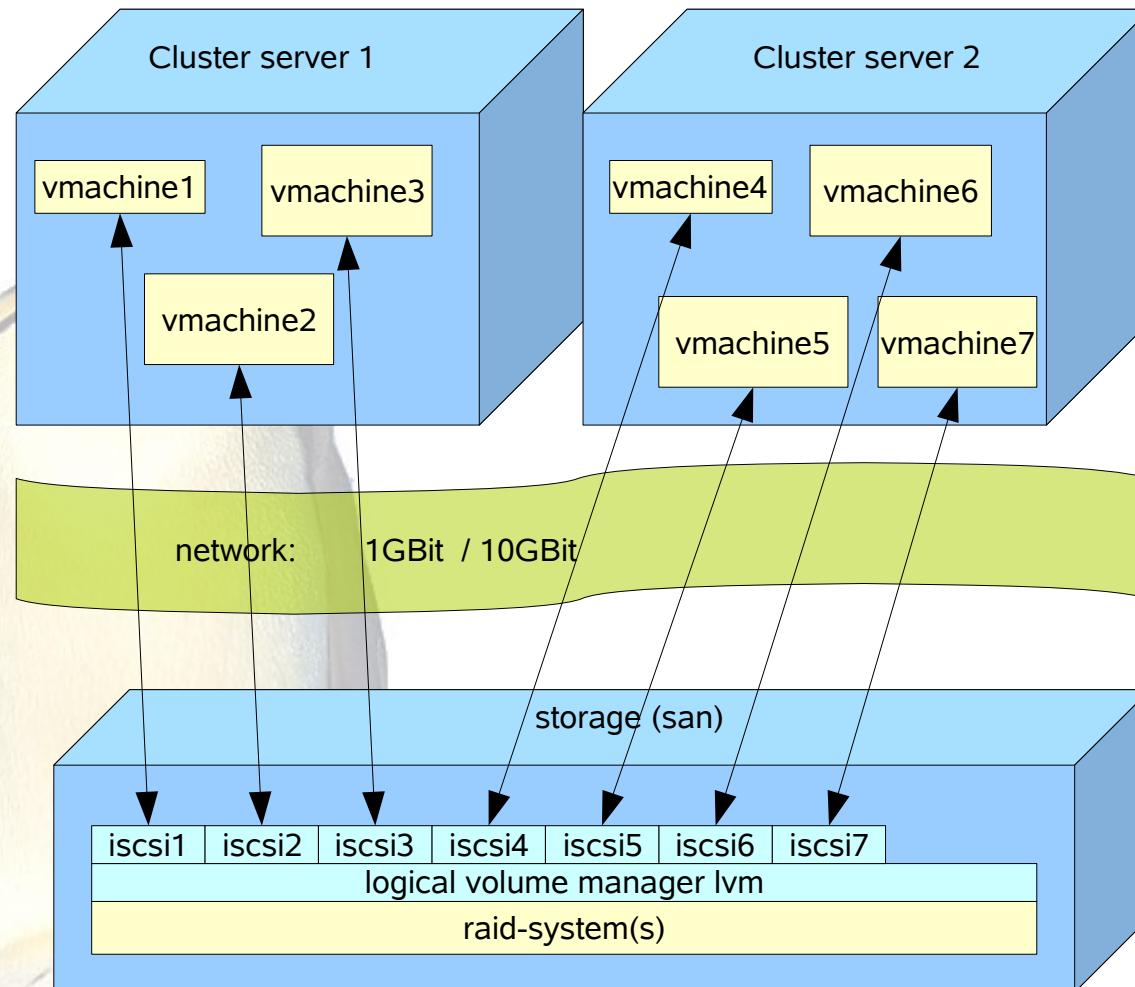
# structure of a cluster



# the cluster servers

- the cluster servers (cs) ( $>=2$ ) consist the virtual machines (vm) i.e. the computing power.
- they need a small mass storage only for the base operating system (XEN: dom0) or boot via iscsi

# structure of a cluster



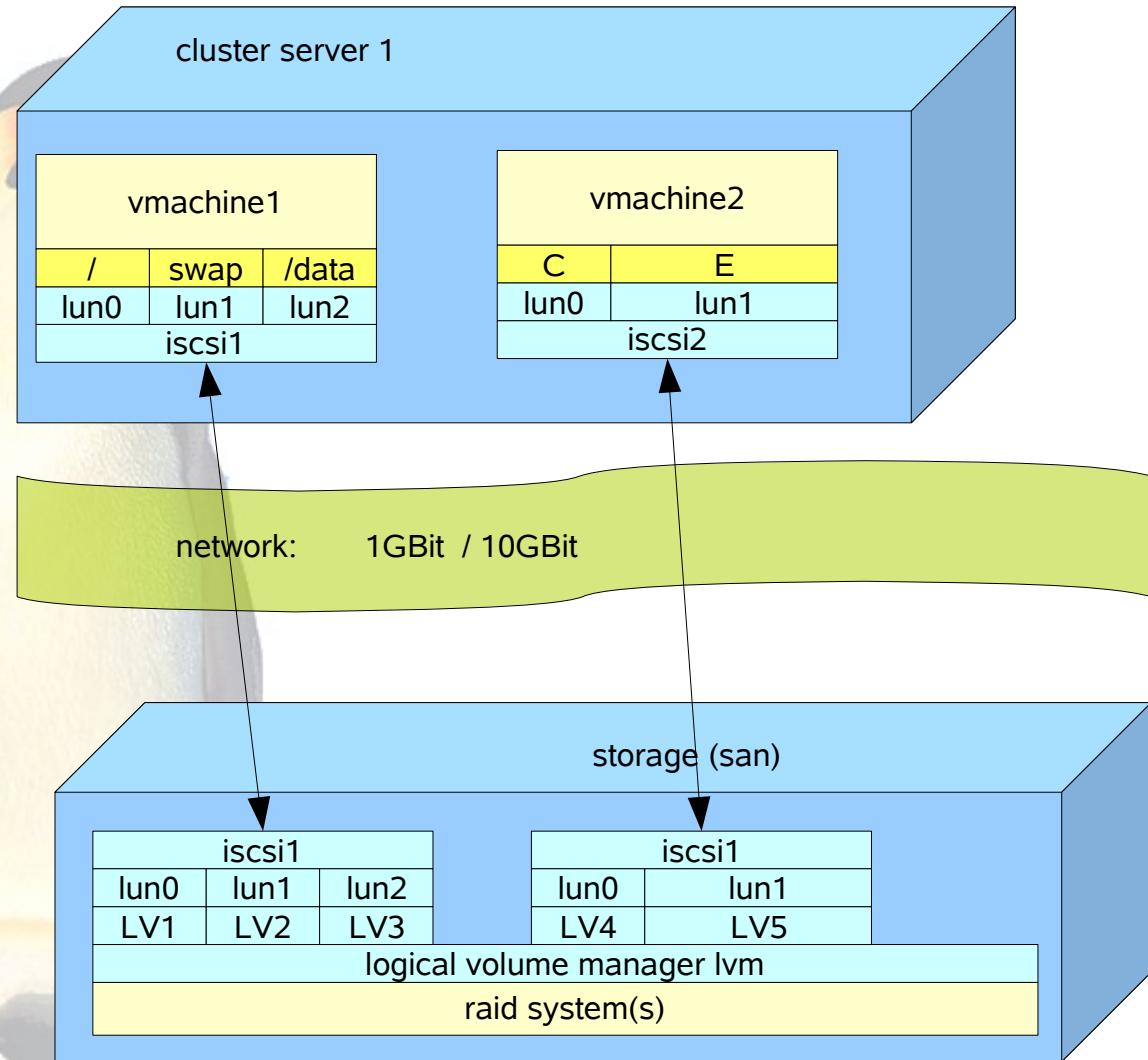
# virtual machines

- run on the cluster servers (use their memory and cpu)
- boot via iscsi from san („see“ only sata / ide disks)
- more disks, e.g. for application data, are possible
- live migration to another cluster server possible

# storage (san)

- architecture: disk → raid → lvm → iscsi
- the san publishes the storage as blockdevices in form of *logical volumes*
- logical volumes host the full os installation and maybe additional data areas.
- blockdevices are available as iscsi devices across the storage network

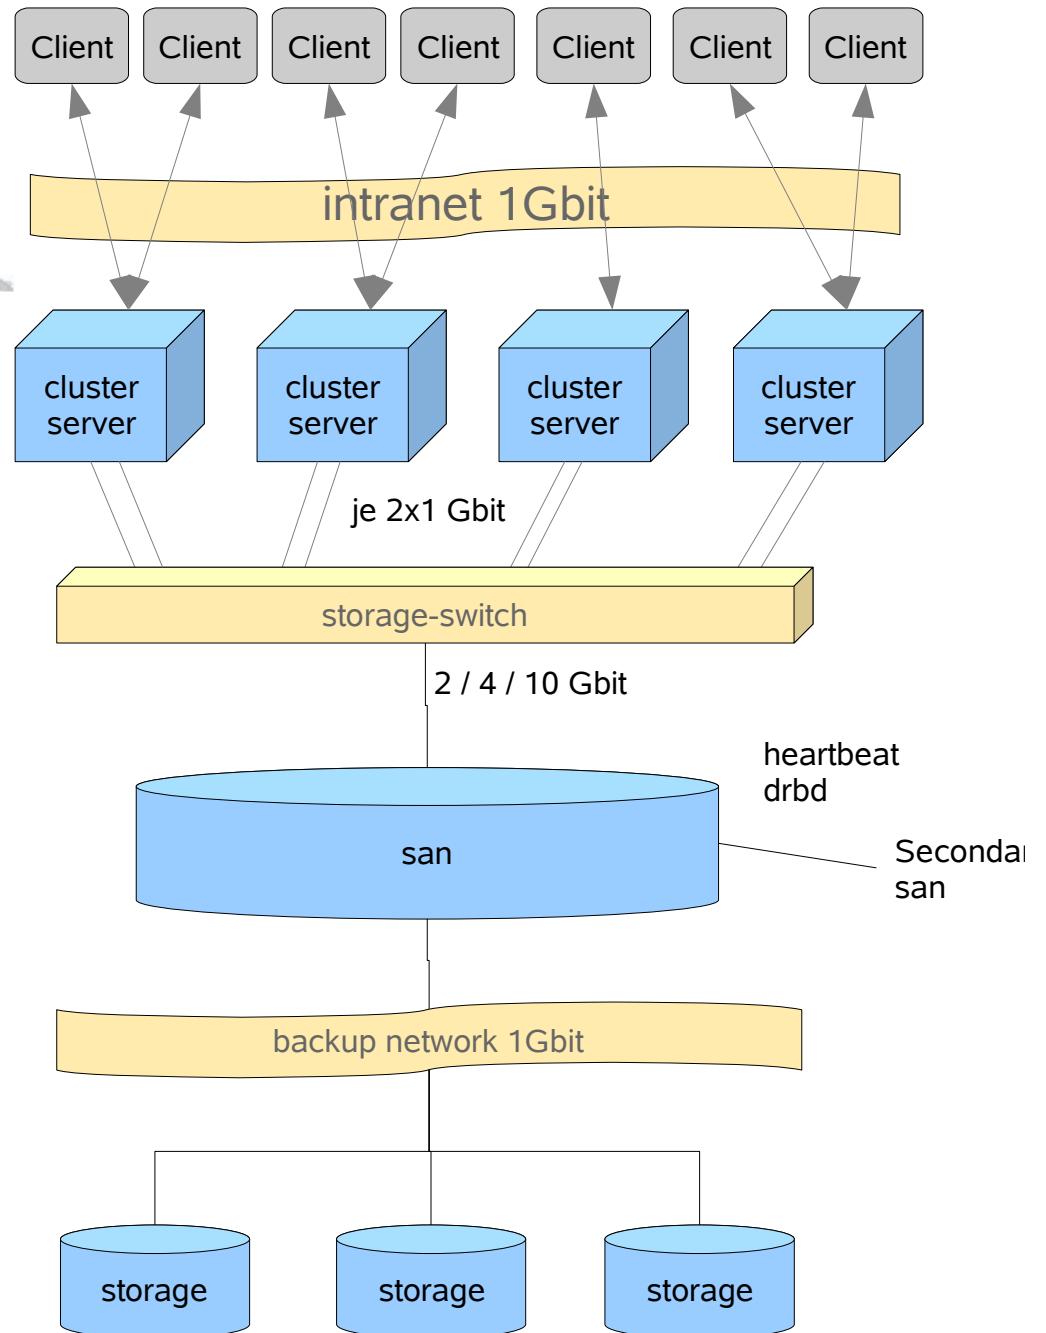
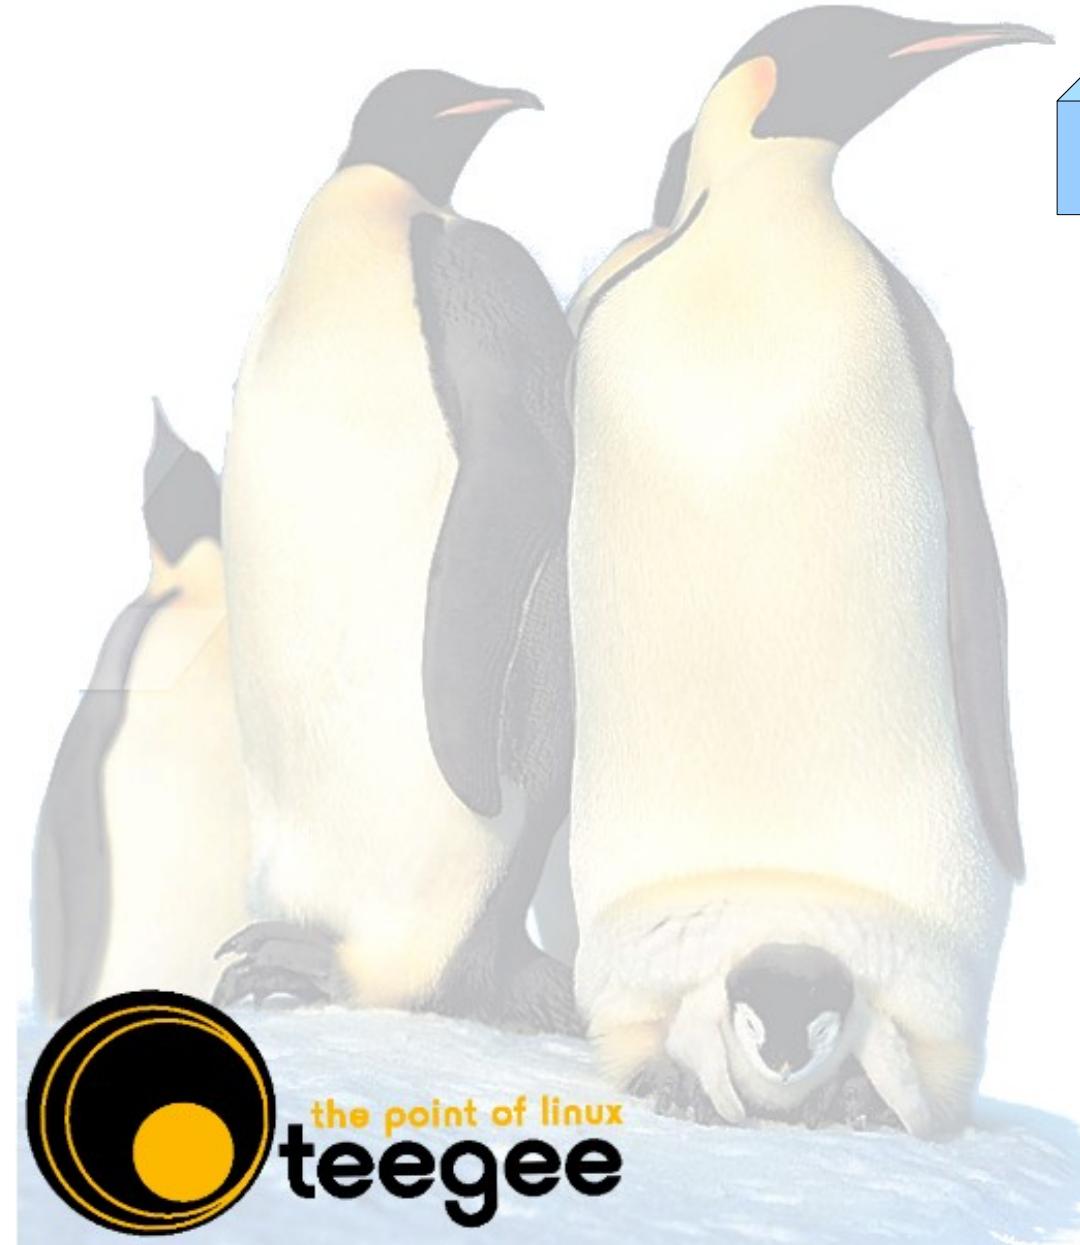
# cluster more detailed



# storage (2)

- because of the usage of lvm it's easy to increase / shrink the virtual machines
- you can increase lvm-systems by more disks / raids or even by external storage
- the management software laxCluster typically lives here
- the san is just a linux server

# 4 levels architecture



the point of linux  
**teegee**

# 4 levels architecture

- optimize your network by separating intranet (user-interactions), storage network and backup network
- longterm upscale strategy by adding more cluster servers
- virtual machines benefit from new, faster hardware by live migration or restart on a new cluster server

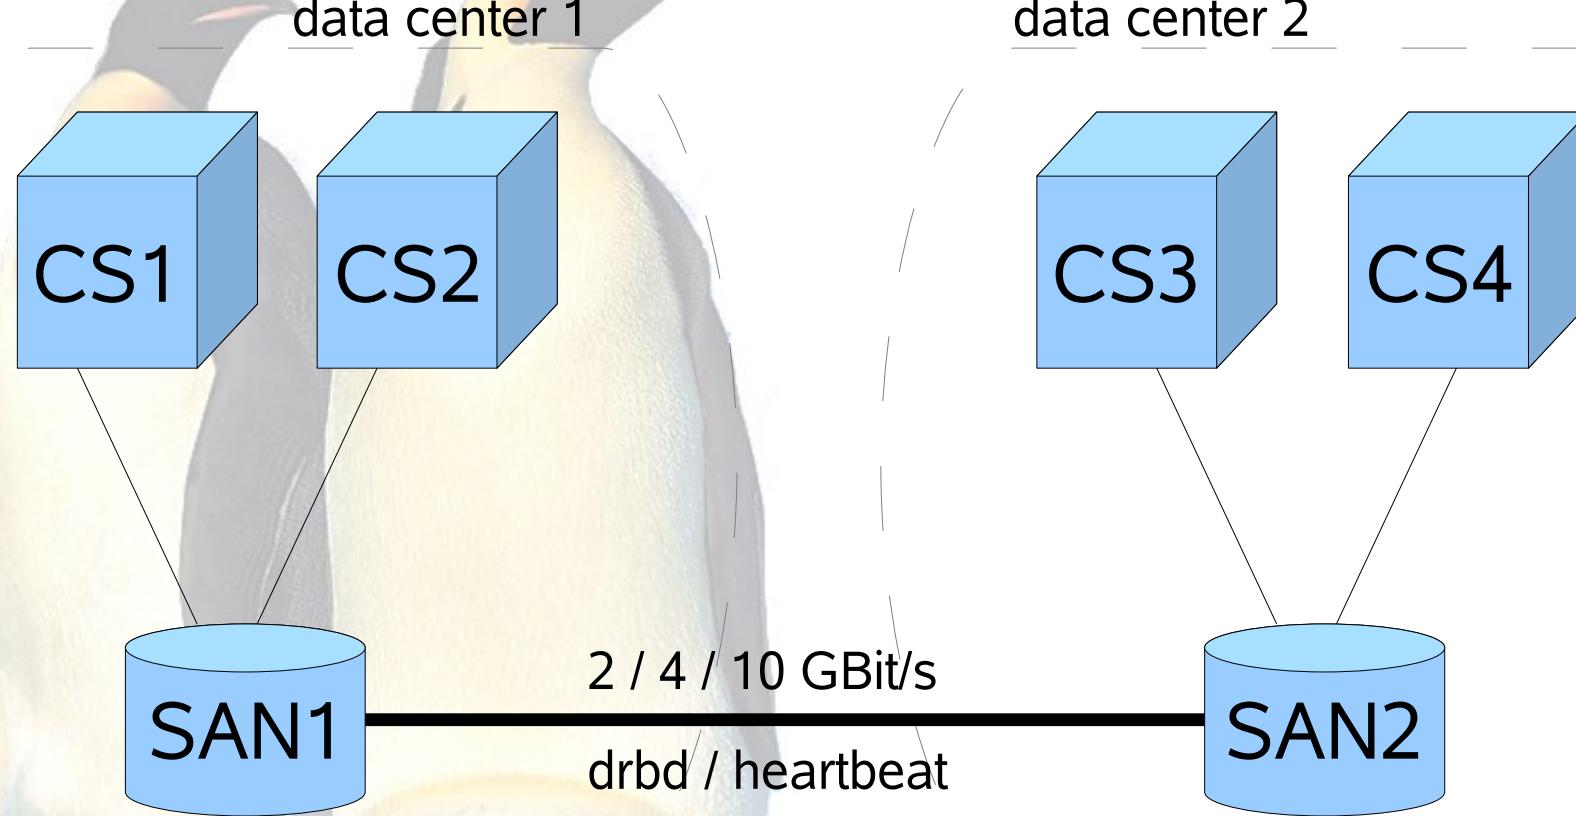
# 4 levels architecture

- migrate virtual machines because of security and rate
- backup the machines online
- start virtual machines because of maintenance offline on the san

# higher availability

- build up 2 clusters in separate data centers
- connect the sans by a highspeed link for realtime data exchange
- if one data center (cluster) fails the virtual machines can be started on the remaining cluster

# Higher availability (2)



# Software for virtualization



hypervisor-virtualization technologie XEN 3.2

- SuSE Linux Enterprise Server SLES10 with SP2
- OpenSuSE 11.0

# Xen

- runs on each linux compatible hardware
- GPL / open source
- strong industry support (Citrix, Microsoft, RedHat, Oracle)
- supports the virtualization hardware extensions from AMD and Intel
- runs unmodified operating systems (Windows)

# Xen

- exclusive usage of hardware by virtual machines
- supports internal subnets by routing / nat
- live migration of virtual machines (with Xen 3.2  
Windows too)
- delivers many informations about the state of the  
virtual machines

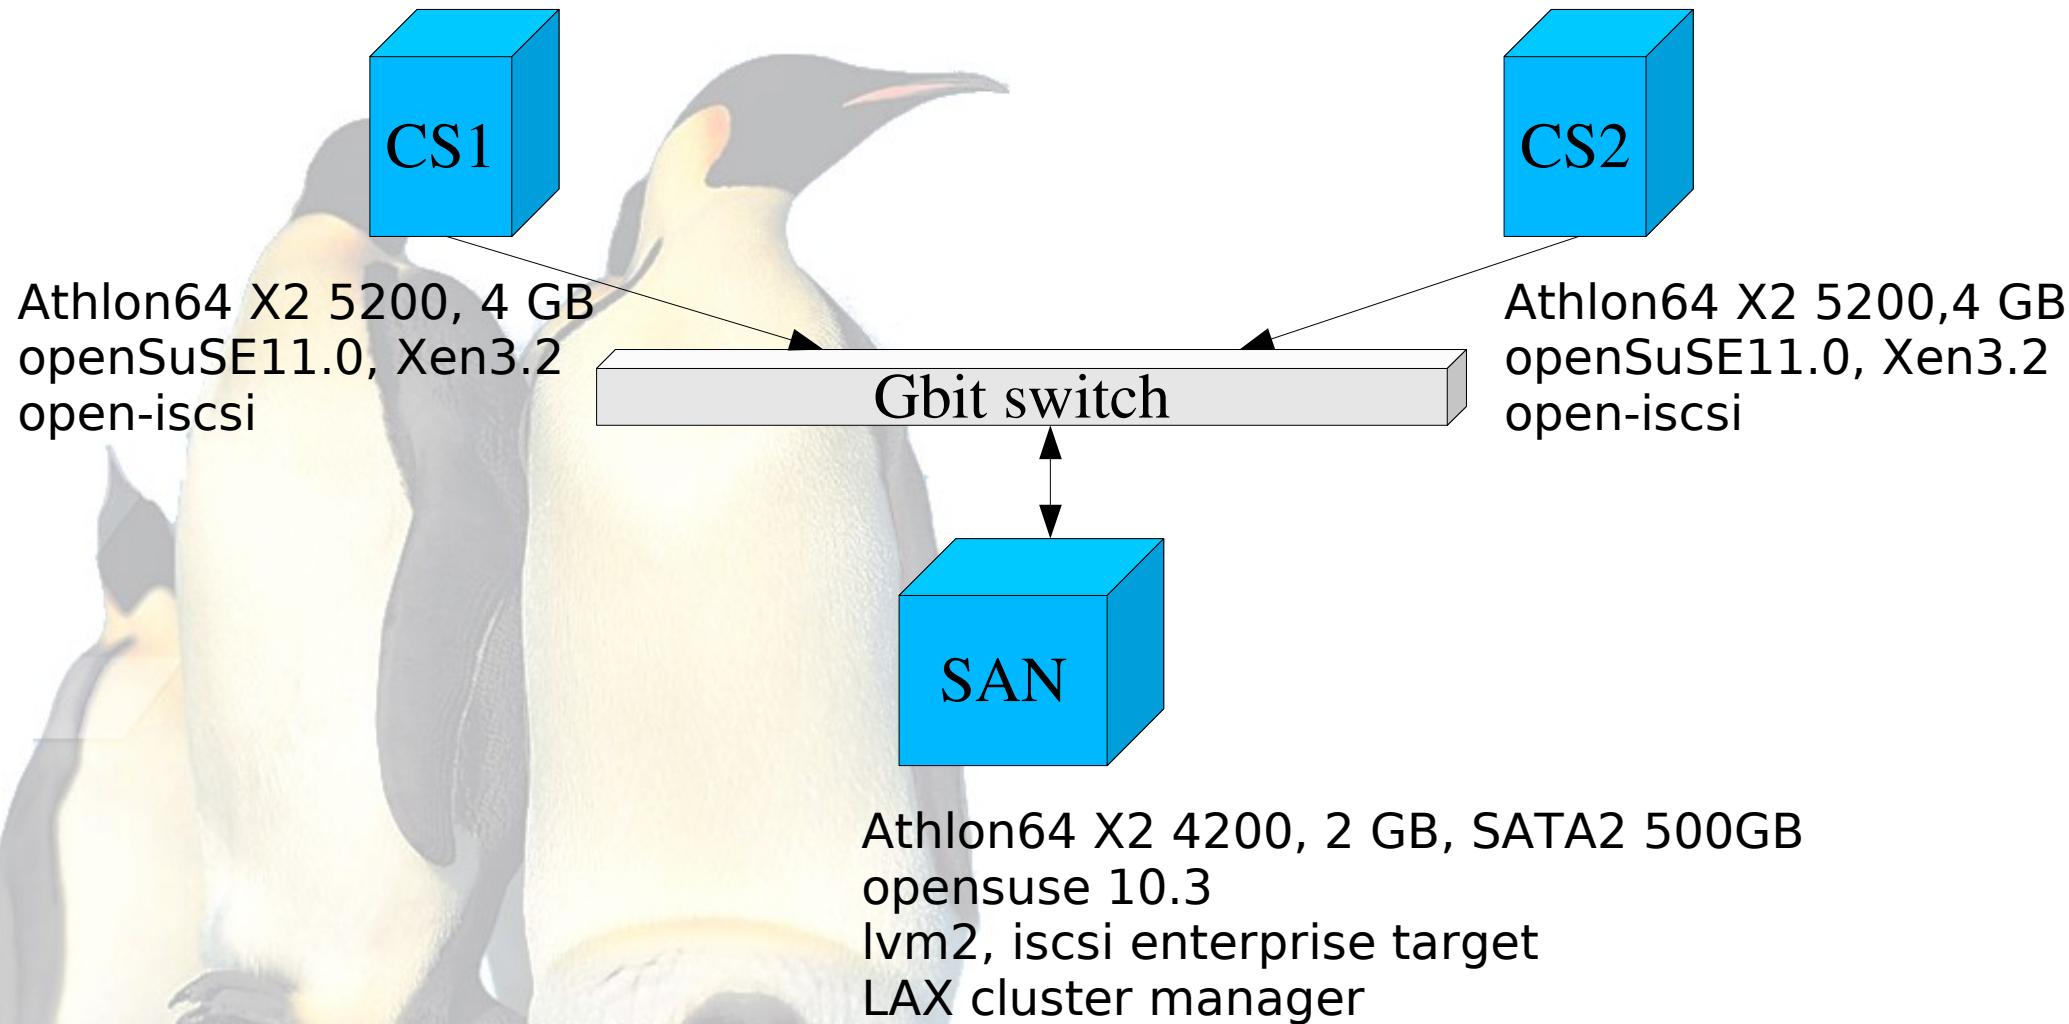
# Lax cluster manager

- bundles the technologies lvm, iscsi and xen
- Installes machines from templates and customizes them to usable virtual maschines
- manages active and inactive machines
- direct access to the virtual machines
- integrates the machines into the LAX network db

# backup

- virtual machines can be saved online by lvm snapshot technology
- the backups can be stored local or on remote hosts
- full - and difference backups possible

# test system cluster cl1



# LaxCluster - overview

LaxClustermanager

Cluster-Baum

- cl1
  - cs1
  - cs2

Eigenschaften Clusterserver

freier Speicher 1361 in MByte

CPU Cores: 2 Takt: 2493 MHz AMD Athlon(tm) 64 X2 Dual Core Processor 4800+

Eigenschaften

Plattenplatz (SAN): 165.34G LabelWin

cl1 cs1 1

virtuelle Maschinen

--	--	--	--	--	--	--	--	--	--	--

aktive virtuelle Maschinen inaktive

Maschine	Cserver	ID	Speicher	vCPU	Status	Zeit	Disk	MAC	IP-Adresse	Template	Info
etch1	cs2	1	512	1	-b----	5.1	6G	00:16:3E:5E:7C:B1	192.168.30.204	debian-etch	Beschreibung_erererer
p3	cs1	1	128	1	-b----	4.5	4G	00:16:3E:A1:2F:9A	192.168.30.240	debian-etch	Beschreibung_p3
probe1	cs2	8	128	1	-b----	0.0	4G	00:16:3E:EC:70:3B	192.168.30.241	suse103	probe1
probe3	cs1	3	1024	1	-b----	4.3	8G	00:16:3E:32:69:07	192.168.30.197	debian-etch	probe3
test1	cs2	2	256	1	-b----	8.3	4G	00:16:3E:31:DB:5D	192.168.30.201	suse103-update	test1
test2	cs1	4	128	1	-b----	7.5	4G	00:16:3E:75:75:39	192.168.30.202	suse103-update	test2
test3	cs1	5	128	1	-b----	7.3	4G	00:16:3E:B7:5F:62	192.168.30.203	suse103-update	testmaschine_3
ttt	cs2	3	256	1	-b----	4.8	4G	00:16:3E:B6:0F:CE	192.168.30.205	debian-etch	Beschreibung_ttt
winxp1	cs2	7	256	1	-b----	7101.5	6G	00:16:3E:10:D1:B7	192.168.30.211	wxp	Beschreibung_winxp1
winxp2	cs1	6	512	1	r----	6356.3	4G	00:16:3E:2D:94:55	192.168.30.212	wxp	Windows Testmaschine2

# LaxCluster – install a new machine

