
TSM on Linux

The Backup Scenario at University of Cologne

Claus Kalle

kalle@uni-koeln.de

28. September 2005

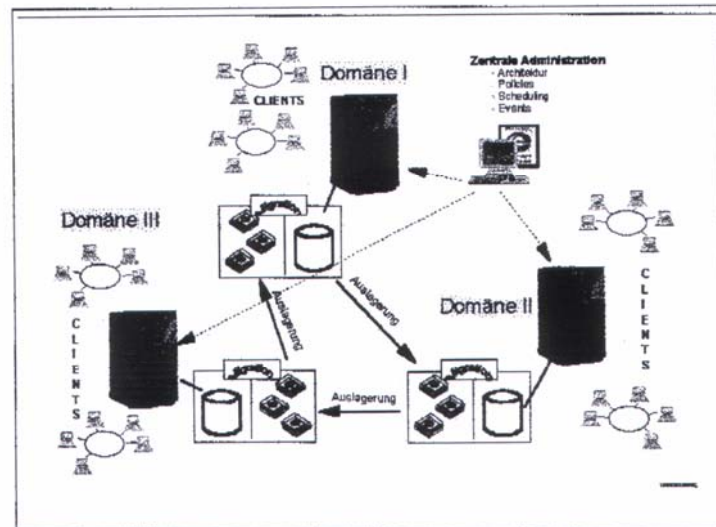
Oxford University TSM Symposium

Zentrum für Angewandte Informatik der Universität zu Köln (ZAIK)



Original Situation (1997)

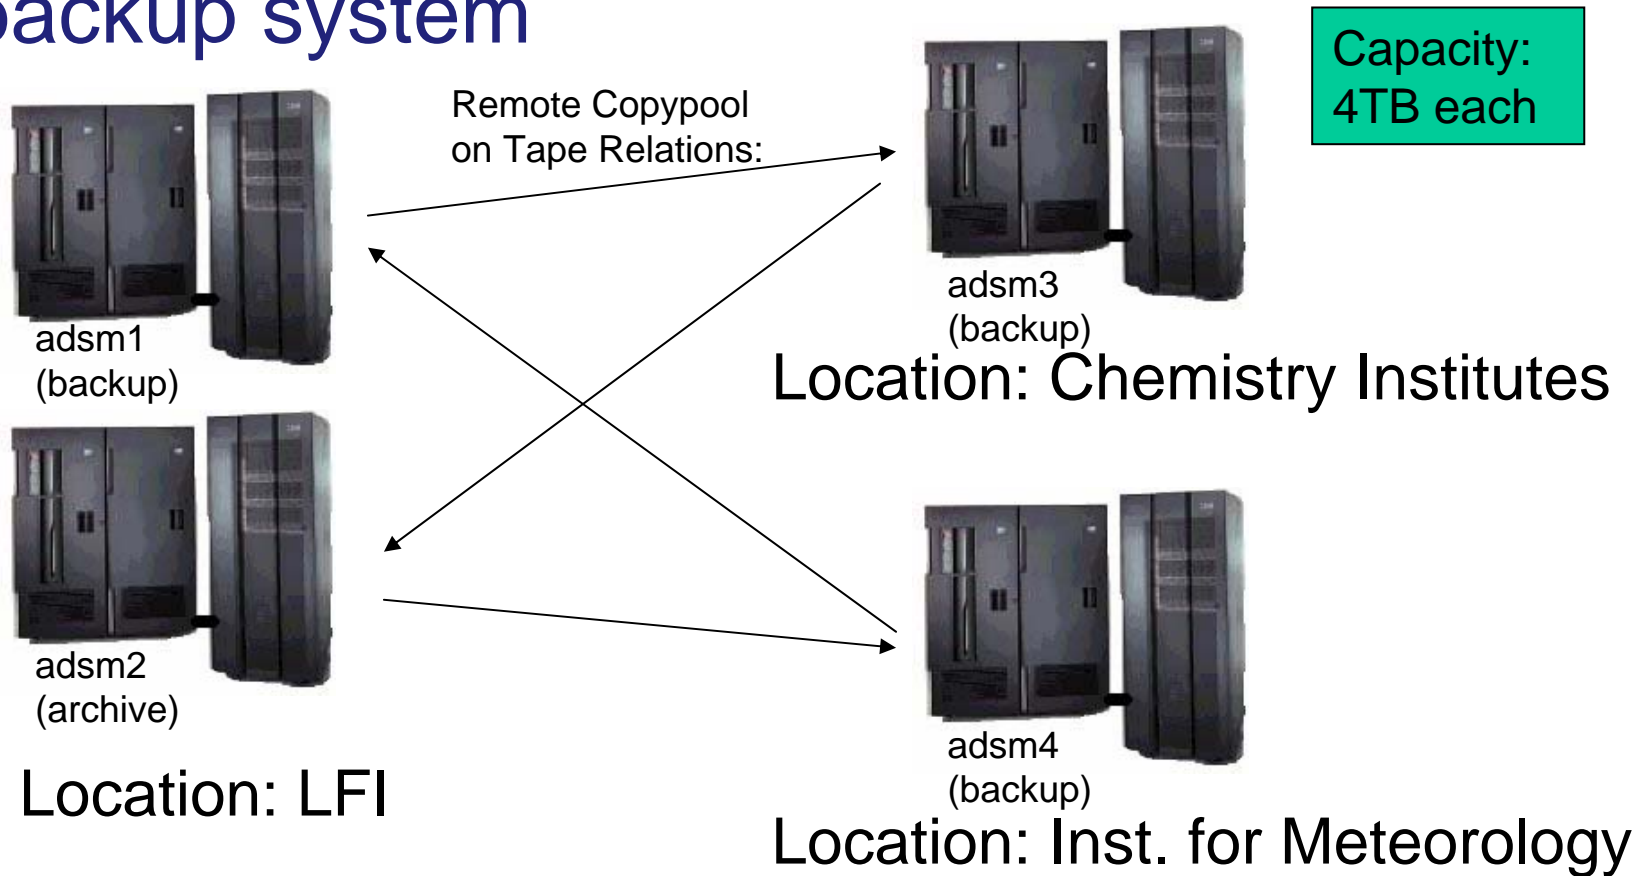
- Legato Backup System became too small..
- Asked for 3x 3466-NSM-C30 in 1998



- Applied for disaster safe concept long before 9-11

UofC ADSM start configuration

- Got 4x3466-NSM-C30 for a campus-wide backup system

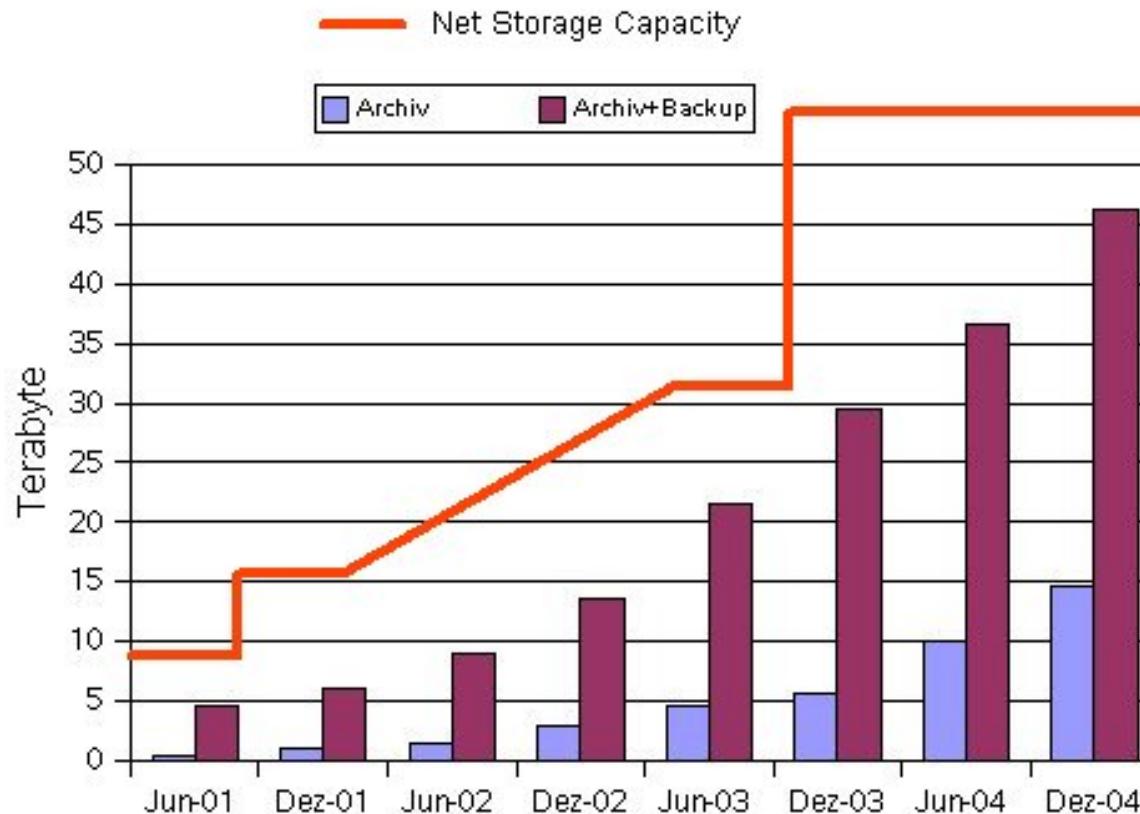


The Life of a TSM Solution

- Incremental hardware upgrades ...
 - 2000/2001: Tape Quality: density/length 3590-J/K, 3590-E1A-XL drives
 - 2003: Tape Storage: adsm1-3: each +1xS10
- New ADSM-Versions, then Tivoli/TSM
- ... Finally Decomposition of 3466-NSM product:
 - IBM H50 2CPU 768 MB SSA disks 12x9,1GB AIX 5.2
 - 3494 Robot
 - Tivoli TSM Server Version 5, Release 2, Level 2.5
 - Management on our own (DBSIZE, CACHE, HW, ...)

Growing Usage Numbers

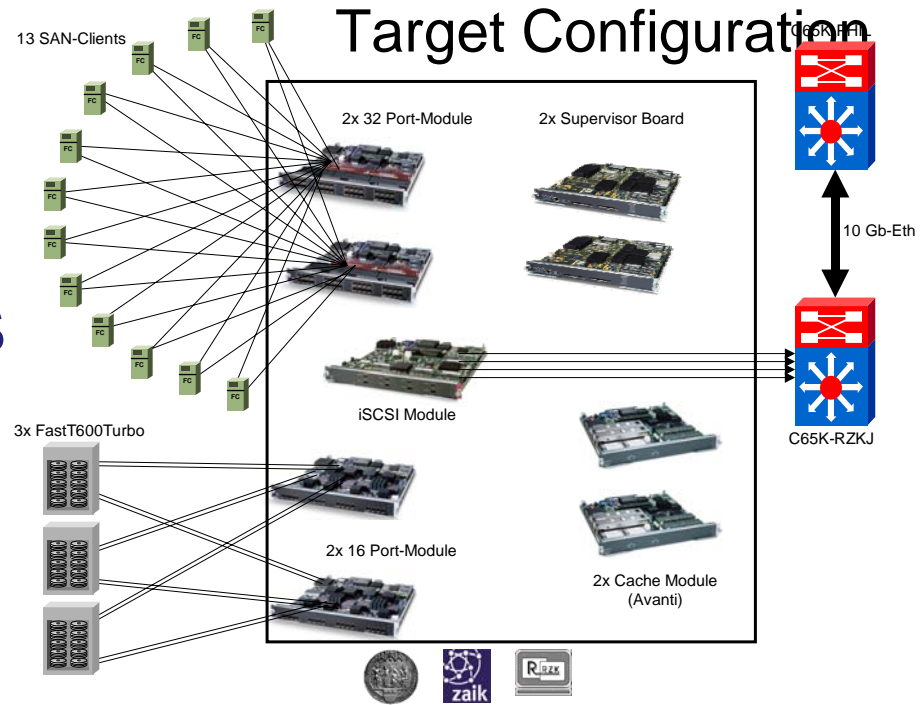
- Backup more popular than Archive




- More than 700 nodes registered

2004: dramatic increase in storage volume

- SVC-Installation
- End 2004: 80 TB
- Now: 8 SVC nodes



In-Band-Virtualization



2004: Linux-Cluster giving more Cha(lle)nges

- Attached to SVC/SAN with iSCSI
- HPC-Users producing lots of data




- 128+1 Nodes Type SUN V20z
1 rackunit each
in 5 Racks (32/32/1/32/32)

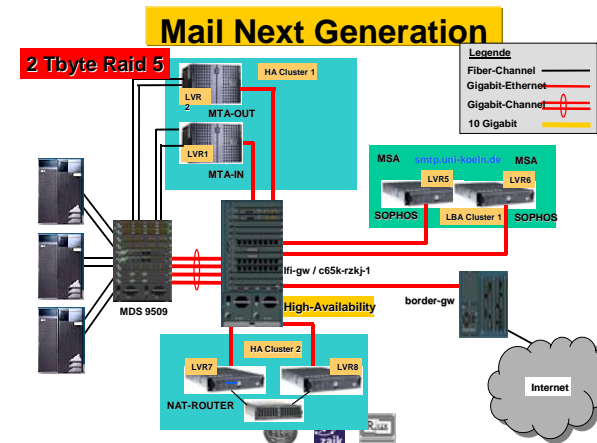
- Per Node:
2x AMD Opteron 2.2GHz
4 GB RAM
2x 36GB SCSI-Disk
2x Gigabit-Ethernet
1x Infiniband-Interface (10 Gbps)
Serviceprozessor
etc.

2004: Email-Server Migration to Linux

- Dell Poweredge 6450 installed begin 2001, 4xPentiumIII/XEON@900MHz, 8 GB memory, 4 PCI buses



- Not suitable any longer
 - for H/A-Environment (1 power cable only)
 - Anti-Virus-Performance
- 2x4-way Servers „available“



Planning in 2004


- Old TSM-solution will hit capacity- and performance-limit
- Strategic Direction: Linux (Redhat)
- 2xDell PE6450 for free
- 3592-J Tape Units available
 - More than 7x capacity
 - Faster: >40MB/sec
- Plan to build Linux-based TSM-Environment

Put New Technology into 3494

- **Hardware-Change at 3494 of adsm1 and adsm3:**
 - add 2 drives 3592-J into existing libraries (L12 frame)
 - add another S10 frame (120TB)
 - Upgrade OS/2 machine in robot
- **Setup on Dell PE 6450 (adsm1n, adsm3n):**
 - reuse Transtec-low-cost raid (2TB) as TSM-cache
 - TSM-Linux-Server Software available in 2004
 - IBMtape driver for 3494 under Linux
- **ease of setup/installation with IBM-Partner
CARUS**

Configuration of Linux-based TSM-Server

- adsm1n hardware



Configuration of Linux-based-TSM-Server

- Software
- Redhat Linux Advanced Server 3.2

```
[root@lvr2 root]# uname -a
Linux lvr2.rrz.uni-koeln.de 2.4.21-27.0.2.ELsmp #1 SMP Wed Jan 12 23:35:44 EST 2
005 i686 i686 i386 GNU/Linux
[root@lvr2 root]#
```

- IBMtape 1.5.9
- Act as working plans dictate
- ->Success!

First experiences

- No problems with TSM-server-software
- No problems with IBMtape-driver
- Performance of Low-Cost RAID insufficient
 - 3592-J does not stream well
- 2 Tape Units insufficient
 - Reclaim does not work without manual intervention
 - Normal operations get blocked by reclaim

Enhancements

- Use SVC VDISK as TSM cache at adsm1n
 - 1 TB SVC VDISK in image mode backed up via GB-Ethernet-LAN in 15 hours (~20 MB/sec)
- Plans for integration of another 2x3592-J in both robots
- Plan for usage of IBM DS4100 (28x250GB SATA) as cache for adsm3n

Alternative Backup Technology

- FLASHCOPY
 - Feature of 2145-SVC: point in time copy
 - Journaled ext3 Linux-Filesystem
 - Journaled SUN-ufs
 - Scripts to be written for each application
 - Alternate destination VDISK
 - Make new flashcopy relations each time
 - umount/mount flashcopy VDISKs
 - Used for mail-next-generation: 1,2 TB in <16h
 - Used for AFS-Home-directories: 1 TB in 18h
- Very short time-to-recover relative to TSM usage

Current Backup Situation at UofC

- Homes(AFS), 1TB/320GB:
 - alternating flashcopy (Mon/Wed/Fri)
 - AFS buta on Linux (laatsch@uni-koeln.de) every other day
 - Initially incremental TSM file backup
 - BC/TTR: 2h30, (reboot, change LUN alloc + salvage (1h10) + attach 45000 volumes (0h35))
- Email(cyrus etc.), 1,5TB/210GB:
 - alternating flashcopy (Sat/Sun/Tue/Thu)
 - TSM incremental over 5 million files in ~7h every day
 - BC, TTR: 1h (reboot, change LUN alloc/fsck/start cyrus).

Current Backup Situation at UofC

- Archive data, 10TB/10TB:
 - AFS buta on Linux, every other day
- HPC-Linux-Cluster clio, 1TB iSCSI:
 - TSM incremental 3,5 Million files in 3h30.
- Others:
 - DAS, OS-Disks, Workstations/PCs
 - ...seem to fit in schedule...
- The above mentioned BC/TTR numbers have not been tested!!!

Expectations

- Capacity growth in archive area
 - Digital age („It's all digital now...")
 - Reuse 3590-K tapes in adsm2/adsm4 (~35TB)
- Capacity growth in backup area
 - High-performance parallel cluster file servers
 - SVC-based VDISKs via iSCSI used in campus
 - Will have to qualify data for backup
- New Techniques (LAN-free backup, server-free backup, ...) will reduce upcoming bottlenecks
- Introduce/enhance storage request justification/management