

TSM on zLinux

Experience Report of a Migration Project:
TSM z/OS -> TSM zLinux

Peter Micke

RZNet AG

Kerpen, Germany

pmicke@rznet.de

Agenda

- Customer's description
- Starting Point
- Immediate measure
- Reasons for Migration
- Migration
- Experience

Customer's description

- Customer in finance area
- TSM is the Backup tool of choice
- Lots of applications running on zSeries
 - e.g. Lotus Domino on USS / TSM on zOS
- Customer has a well equipped zSeries environment (CPU, Disk, Tape)
- Disaster Recovery Concept based on two datacenters (~ 250 m distance)
- Growing number of Open System clients (Windows, Linux, both on Hardware or VMWare)

Technical starting point


- TSM 5.1.5.x on z/OS
- ~ 80 GByte Diskspace for Storagepools
- ~ 40 GByte TSM DB
- Storing of data in two VTS Systems (one in each datacenter), copies are done by TSM
- ~ 200 TSM Clients
- ~ 11 TB of Data (primary)
- ~ 12.500 logical volumes per VTS
- Incomplete Backup of Clients (w/o Bootdrives)
- Incomplete data mirroring

Customer has a bit of a headache ...


- Backup window is becoming larger and larger ...
- ‚Housekeeping‘ jobs run during the backup window (expire inventory, spacereclamation)
- CPU – usage is too high (~ 200 Mips)
- Restores take too long ...
- Customer says: Fileservers can't get restored within an acceptable timeframe
- Technical development doesn't show up for z/OS (e.g. LAN Free)
- Upgrade dependencies between z/OS and TSM

Hardware at starting point

RZA



RZB




Changes by introducing VTS PtP Mirroring


- ‚Scratch mount penalty‘ in immediate mode by delaying a scratch mount by appr. 80 s (ESCON)
- Backup window hits the ceiling!

Changes with introduction of VTS PtP Mirroring

RZA



RZB



VTS PtP

3494

3494

Immediate measures

- Breaking up the housekeeping jobs
 - Expiration, DB Backup and Reclamation were running as parallel processes
- Moving Tape storage pools from VTS to native 3590 H1A drives
- Simplification of TSM Storage hierarchy
 - Less storage pools, some of them using collocation

Further improvements / Options

- Migration to another TSM platform
 - > Open Systems (AIX, Solaris, (z)Linux, Windows)
- Changes in tape hardware
 - > 3592/LTO
- Licences costs(TSM / Third party vendors)

License costs

- TSM and Lotus Domino on z/OS -> 850 MIPS CPU load
 - Websphere -> 1600 MIPS expected CPU load
 - Replacement by 6 IFLs
 - Saving potential (license cost):
about 15 k€/Monat
- Migration -> zLinux using IBM 3592

Migration from one to the other ...

- Migration with only a minimum of data migration (Drying out Backup- and Archivepools within 100 days, real migration only for the archive data)
- Parallel use of both TSM environments (z/OS and zLinux)
- Use of „new“ disk areas
- Use of new tape drives and cartridges


Technical details

- Different zLinux partitions on z/VM
- Domino and TSM server use their own partitions
- Attachment of TSM disks via ESCON first, later on via FICON
- Attachment TSM Tape via FC

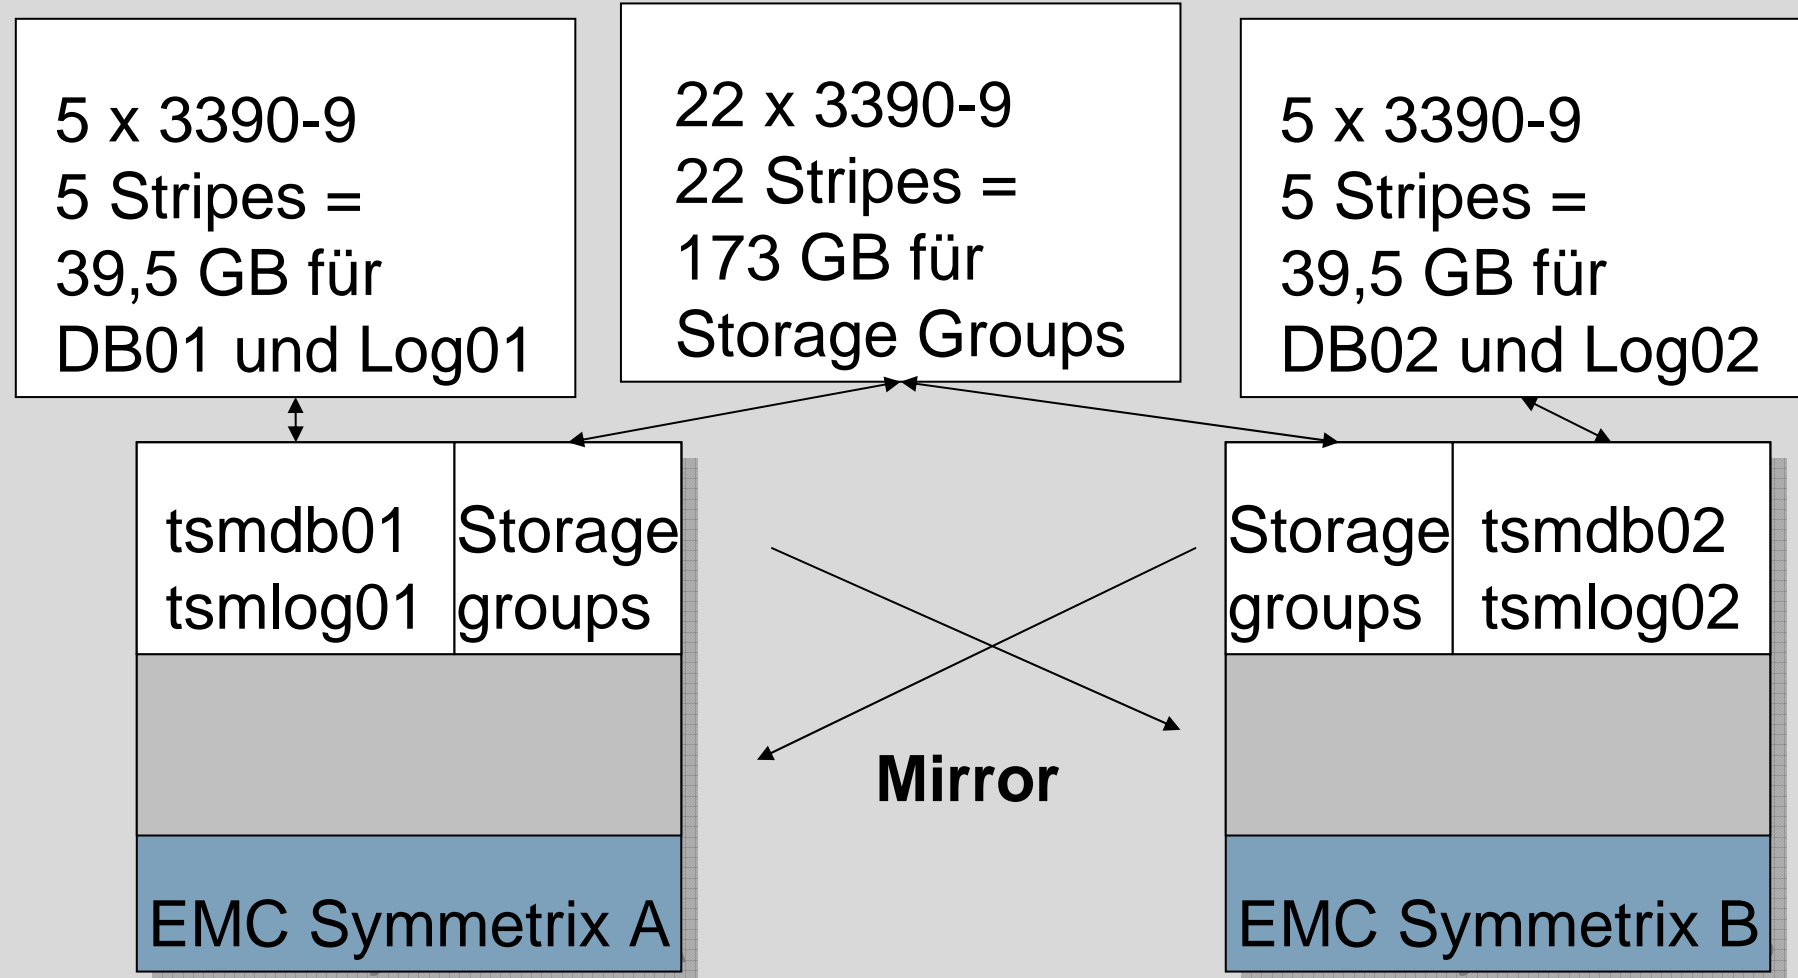
FC cabeling

RZA

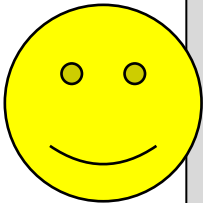
RZB




Verteilung DASD




K-Fall (Ein Plattensubsystem fällt aus, Spiegel funktioniert nicht)

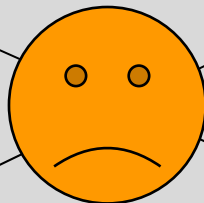
5 x 3390-9
5 Stripes = 
39,5 GB für
DB01 und Log01

22 x 3390-9
22 Storage Groups
173 GB für
Storage Groups




5 x 3390-9
5 Stripes = 
39,5 GB für
DB02 und Log02

tsmdb01	Storage groups
tsmlog01	Storage groups
EMC Symmetrix A	



Mirror

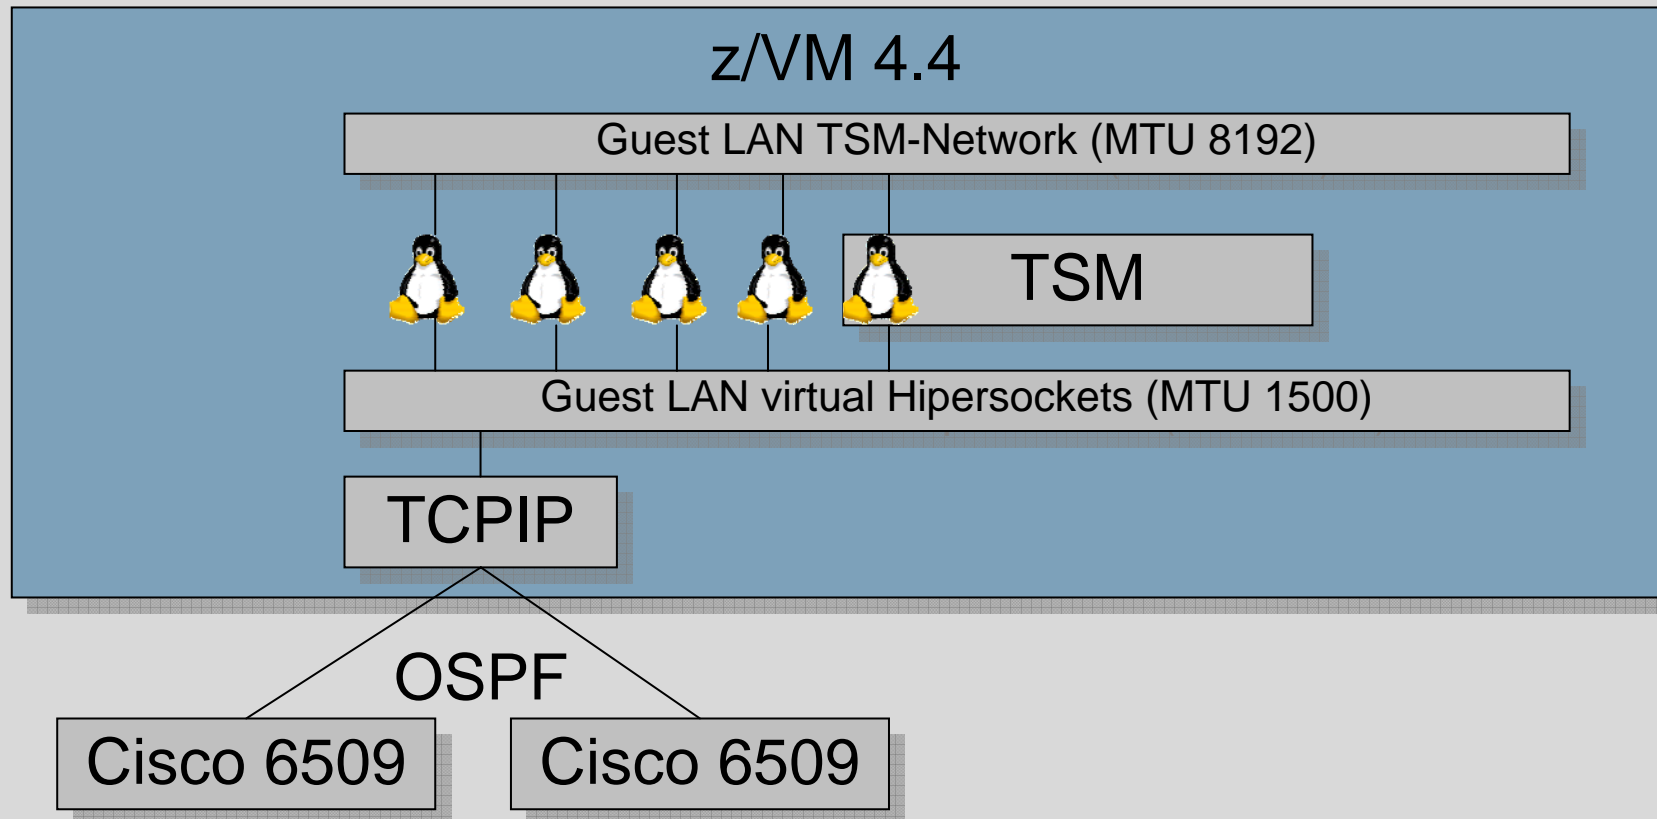
Storage groups	tsmdb02
Storage groups	tsmlog02
EMC Symmetrix B	



Recommendations: usage of DASD

- We distributed the disk volumes amongst many channels and disk arrays on two MC Symmetrics systems and used striping with the LVM.
 - Read IO with about 80 MB/s despite the usage of ESCON
- EXT3 performs better with large files.
- RAW-Devices are not supported by TSM on zLinux yet!!!
- TSM-DBv and TSM-Logv should reside on non striped disk devices.
 - We experienced a problem with a striped disk device and lost the logspace of the TSM server and had to do a point in time recovery...
- TSM is also able to affect other applications by intensive write I/Os (e.g. formatting DB volumes etc.)
 - DASDs for TSM should be attached through channels of its own and should use dedicated disk arrays.

Network attachment



Networkattachment

- The use of a dedicated Network as TSM Network is recommended:
 - Higher MTU-Size increases throuput
 - We experienced bottlenecks on the hypersocket network without using a dedicated network
- The usage of a dedicated OSA Adapter is also recommended.
 - OSPF is available on Linux via GNU Zebra

Kernel, devicedriver and zfc

- The Linux kernel level is determined by the requirements of the IBM tapedriver, we used SLES8 Service Pack 3.
 - Other packages, particularly glibc mustn't be on a higher version.
- The VM guest has to be restarted to use newly defined tape drives.
- Usage of one zfc device per tape drive is recommended.
- We didn't check path failover yet (we didn't have enough fibre cables between datacenters).

Performance

- Performance from network to TSM:
Same as before (data was compressed on the client),
improvement by using compression on the tape drives
(20 MByte/s, Imagebackup)
- Migration Disk to Tape: ~ 15 to 20 MByte/s per process (using
compressed data) (ESCON, FICON seemes to be the same,
we can't use FC attached disk yet)
- Copy from tape to tape, (Move Data): ~ 30 MByte/s
(compressed data), we still use 1 GBit/s FC
- Expiration and database backup ~ 2-3 times faster than on
z/OS

What happened later ... ?


- We experienced a bad network performance backing up Domino data from one VM guest to the other (TSM), starting at several MB/s going down to 50 kB/s
- With increasing backup load we suffered a heavily decreasing overall network performance

Important changes coming up ...

- Leaving VM in favour of using LPARS
- We changed ESCON attached EMC boxes with FICON attached ESS 800.
- We left disk and network performance problems behind!!!

Conclusion

- After the final changes we run a very stable TSM environment with ~ 250 Clients
- Running image backups of file servers and Windows boot drives on a regular base without any problems
- Performance is good, still allowing future growth (backup windows o.k., restores are o.k.)
- Still chances for improvement:
 - FC attachment of disks
 - Changing FC switches 1 GB/s -> 2 GB/s (Tape)



?

Questions?