

We make sure



CentricStor

Concept, Design and Architecture

Klaus Schaefer

Oktober 2006

We make sure



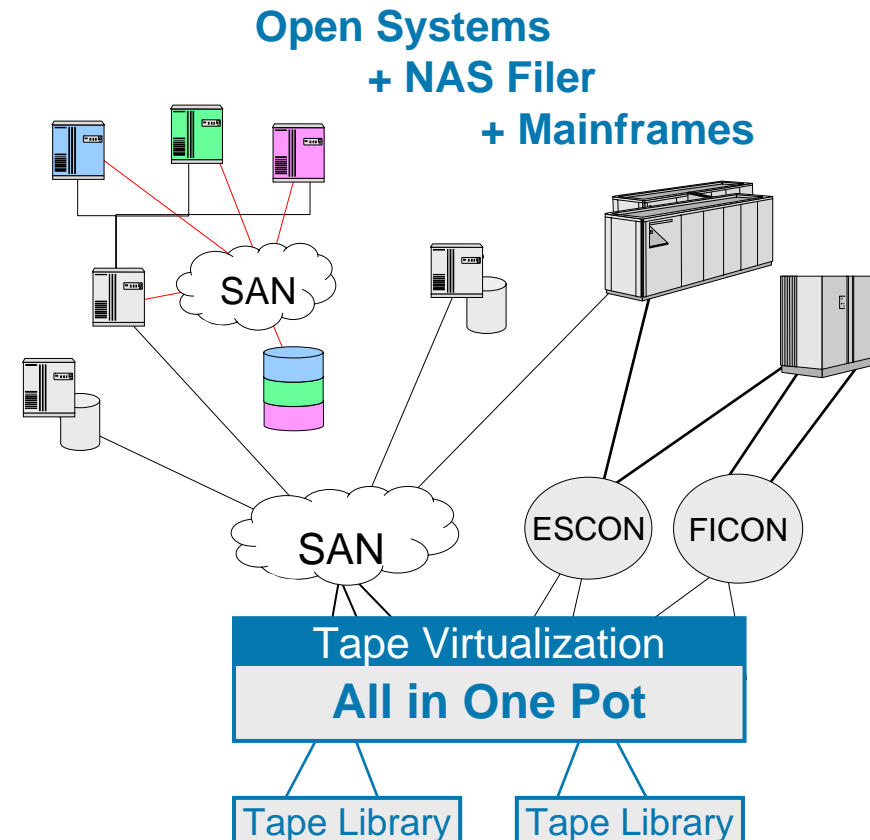
concepts

We make sure



Nearline storage - what customers want:

- Faster data access
- Enhanced throughput
- High degree of parallelism
- Shared usage of resources
- Increased nearline capacity
- Full utilization of media
- Reduction in physical devices
- Better utilization of floor space
- Improved availability
- Technology independence
- Disaster recovery



Warum (noch) Tape für Backup?

Vorteile:

- ✓ Preiswertes Medium
- ✓ Ruhendes Medium
- ✓ Unbegrenzte Speicherkapazität
- ✓ Transportierbar
- ✓ Nearline und Offline Storage
- ✓ Langzeitarchiv



Herausforderungen:

- Backup Performance
- B/R beeinflusst Produktivbetrieb
- Restore Performance
- Zuverlässigkeit
- Managementkosten
- Tape Drive/ Library Kosten
- Kosten für Medien

D2D2T means: management of two- / threefold solutions



- Business Apps
- Data base App
- ...

- Space management
- Migration scheduling
- ...

- Tape management issues
- robustness
- streaming
- ...



D2D2T means: management of two- / threefold solutions

Data protection solution has to be changed

Data migration has to be implemented additionally



- Business Apps
- Data base App
- ...

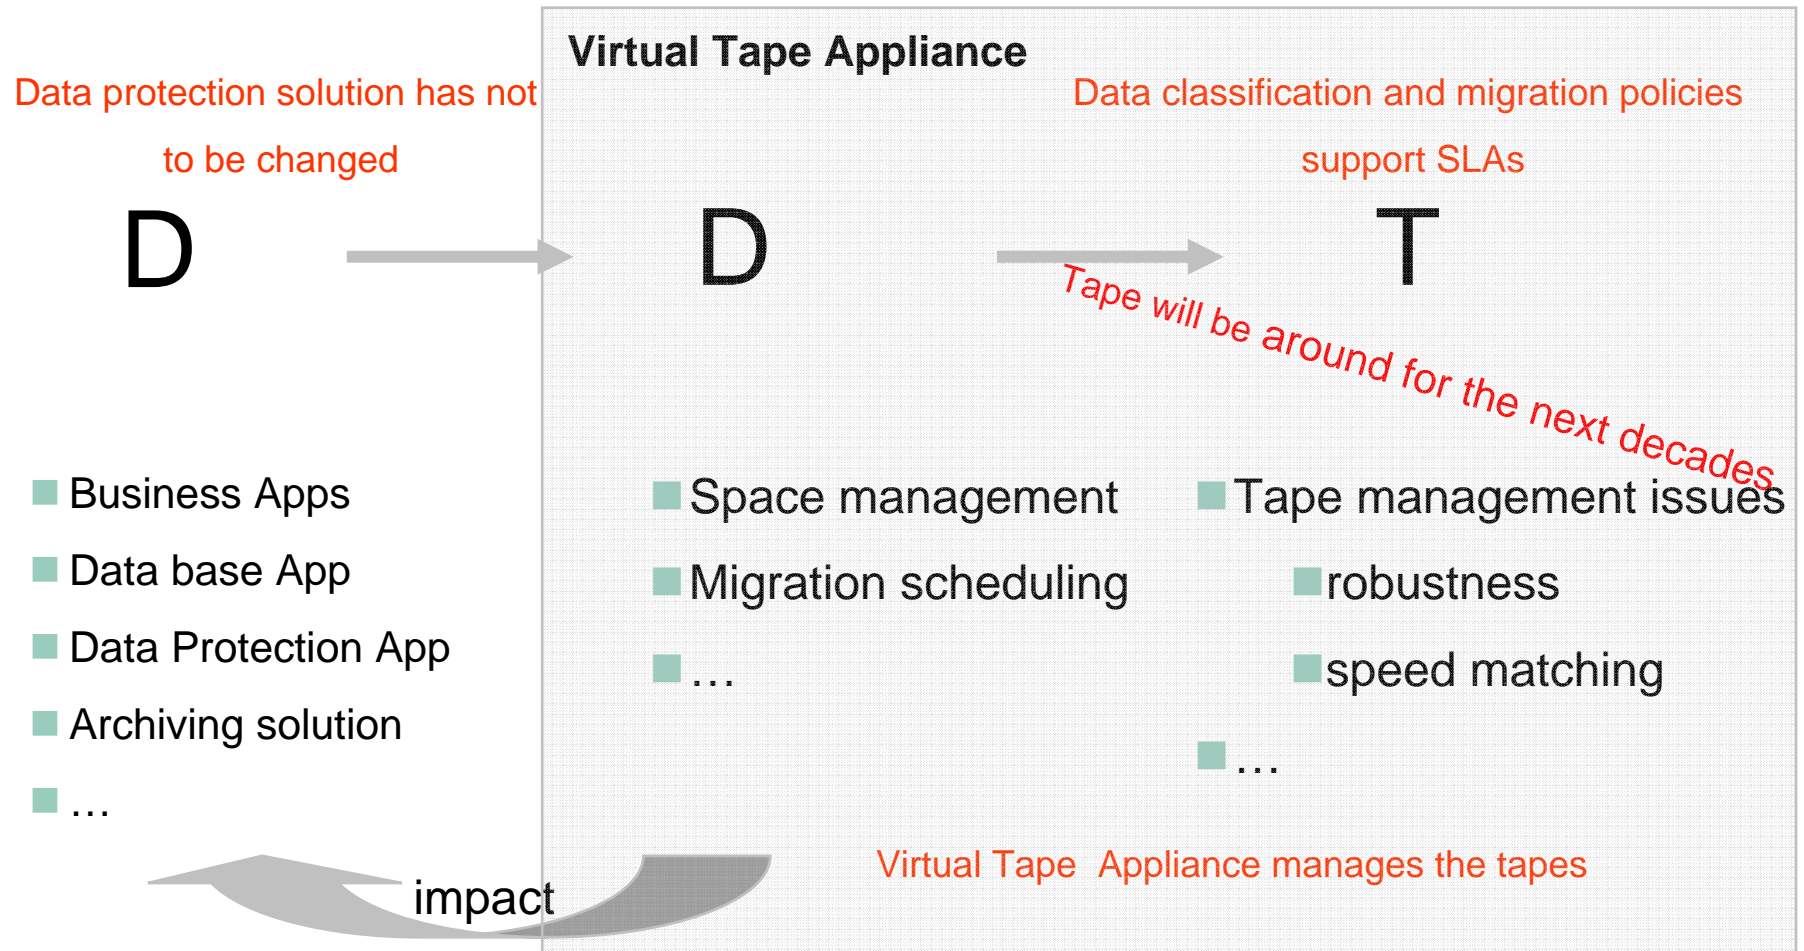
- Space management
- Migration scheduling
- ...

- Tape management issues
- robustness
- streaming
- ...

twofold architecture has to be managed




Virtual Tape Appliance manages the complexity



CentricStor – Virtual Tape Appliance

- Efficiency by virtualization
 - heterogeneous tape storage consolidation
- Maximizes data security with enhanced disaster protection
- Universal connectivity
 - across different host platforms
 - across different tape systems and technologies
- Integrates smoothly into existing data center concept
 - For all major backup applications
 - Host applications continue to run unchanged



We make sure



Inside


CentricStor VTA – major specification

Integrated Universal Processors (ISP)

- Core Component acting as ICP, IDP, VLP from 1 (VTC) to 16

Scaleable Components of CentricStor

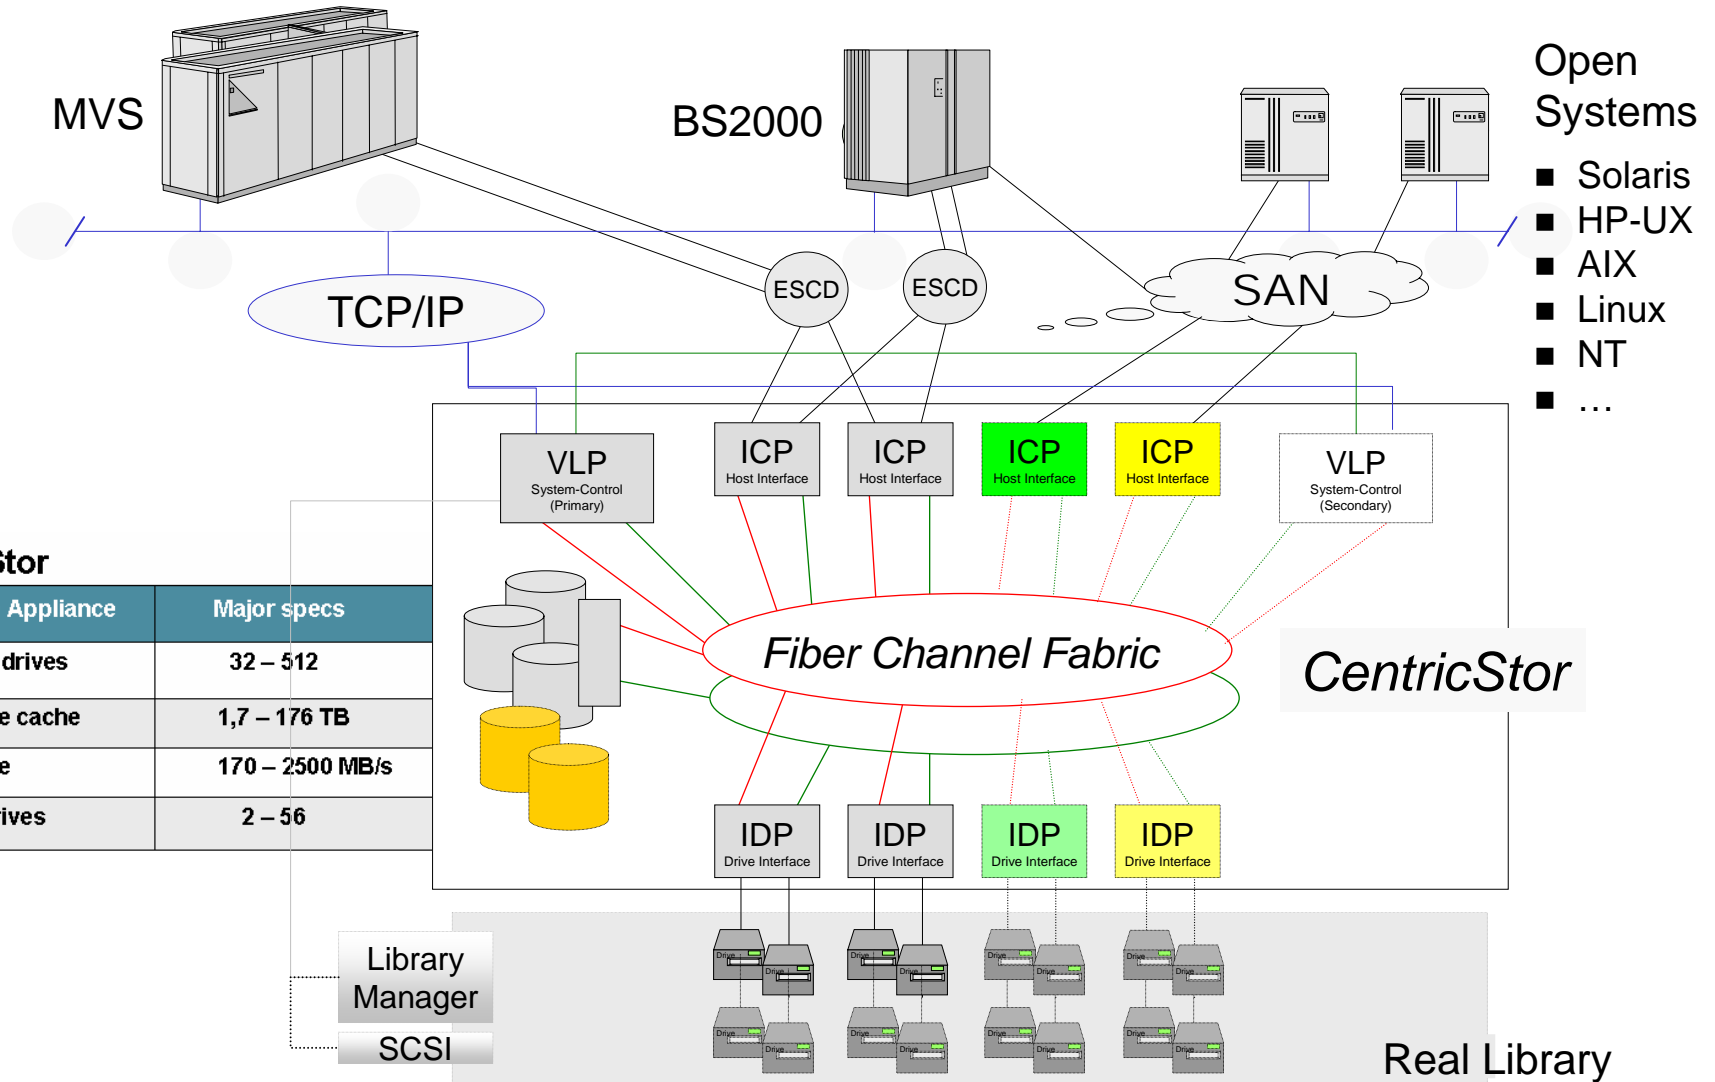
- ICP = Integrated Channel Processor
Numbers of channels and virtual drives
- IDP = Integrated Device Processor
Numbers of real drives
- VLP = Virtual Library Processor
2 for failover
- Tape Volume Cache (TVC)
Capacity for virtual volumes
- Switched Fabric (SF)
two for redundancy and workload balancing



We make sure



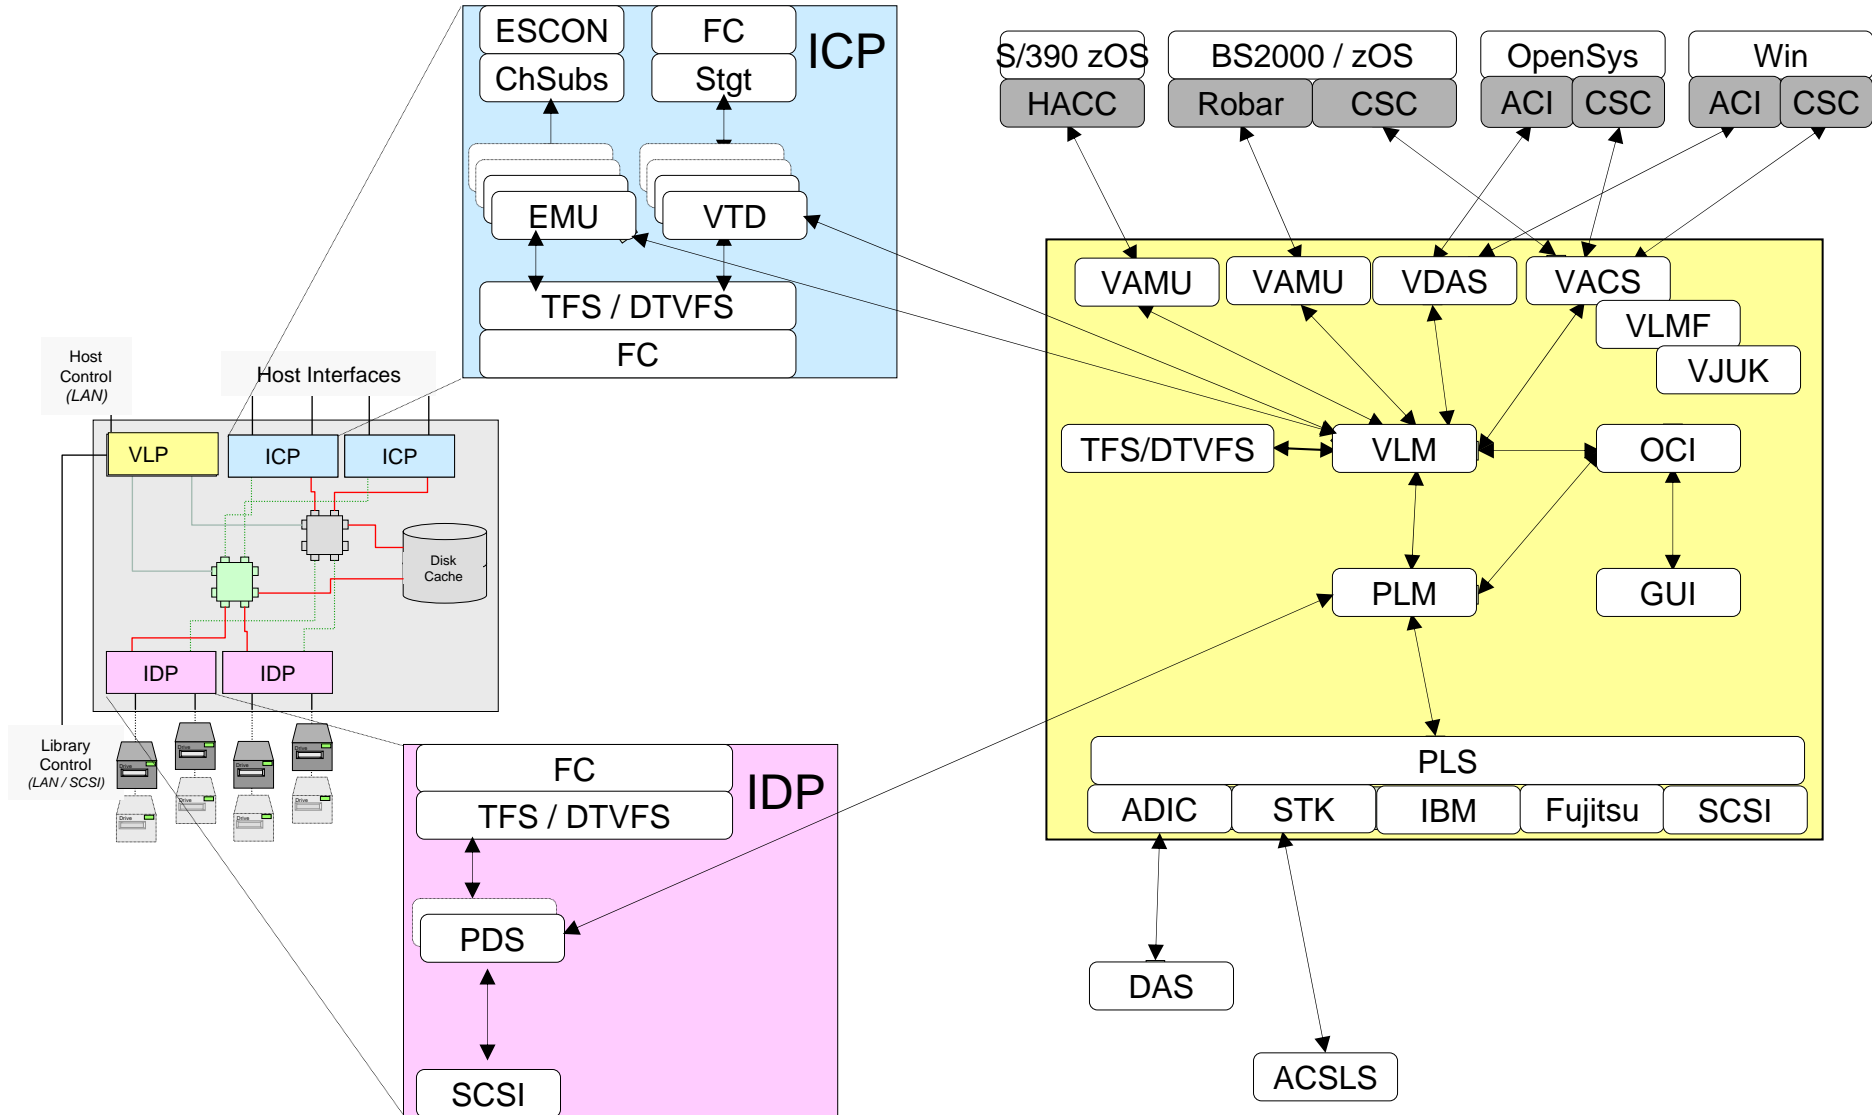
Hardware Architecture



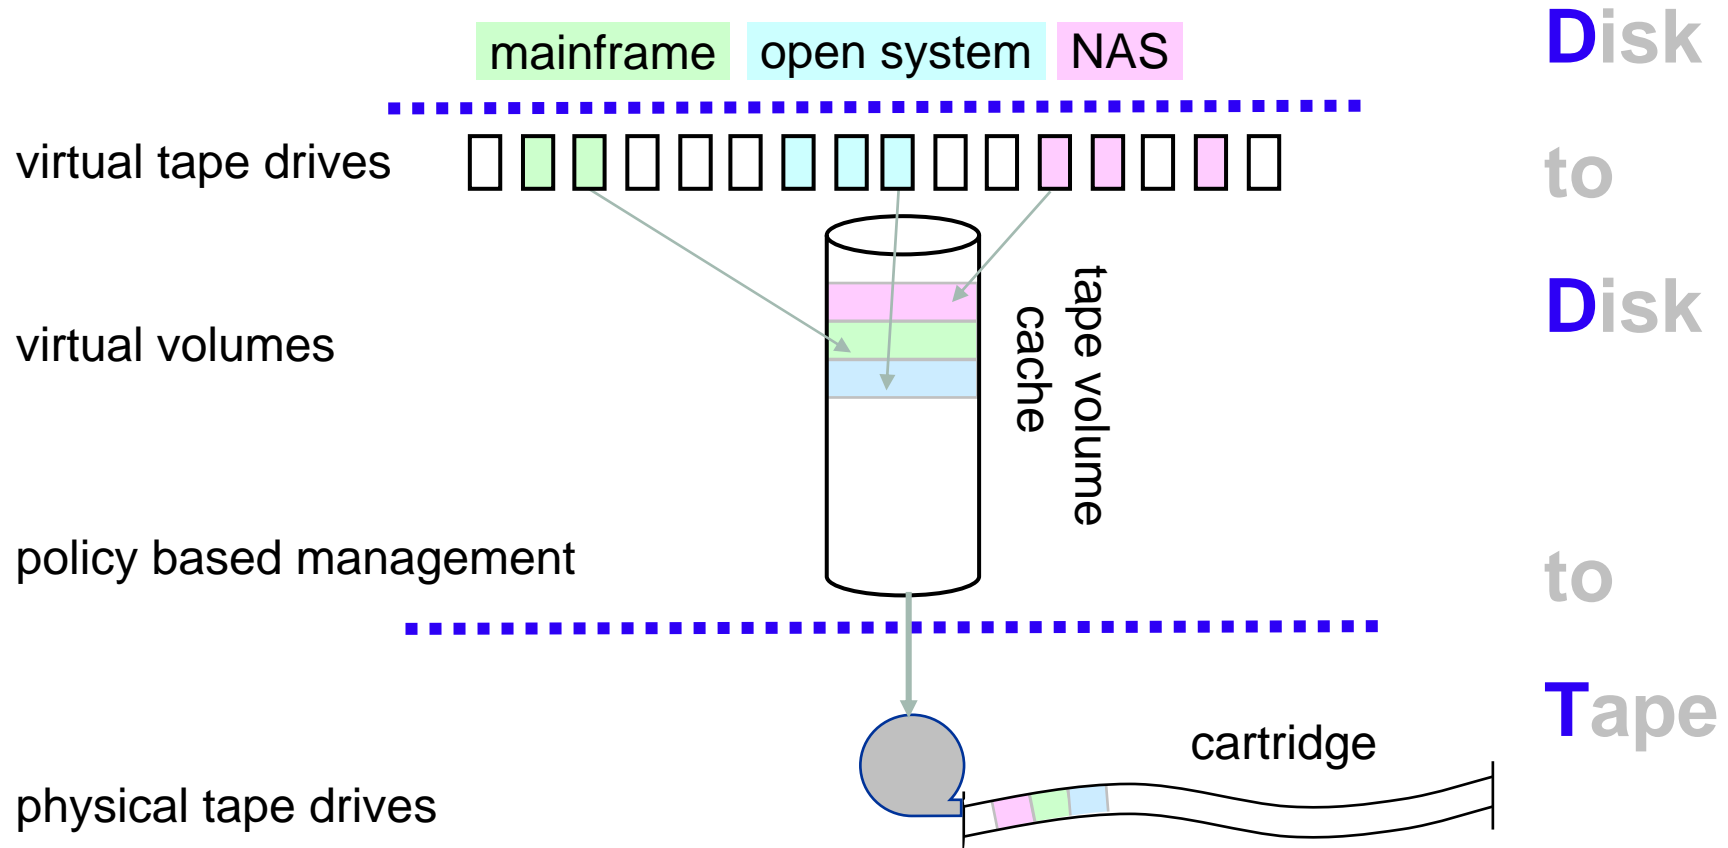
CentricStor

Virtual Tape Appliance	Major specs
Virtual tape drives	32 – 512
Tape volume cache	1,7 – 176 TB
Performance	170 – 2500 MB/s
Real tape drives	2 – 56

Software Architecture




How does it work ?




Reorganization of a physical Volume (1)

Writing PV's in Stacked-Volume-Manner



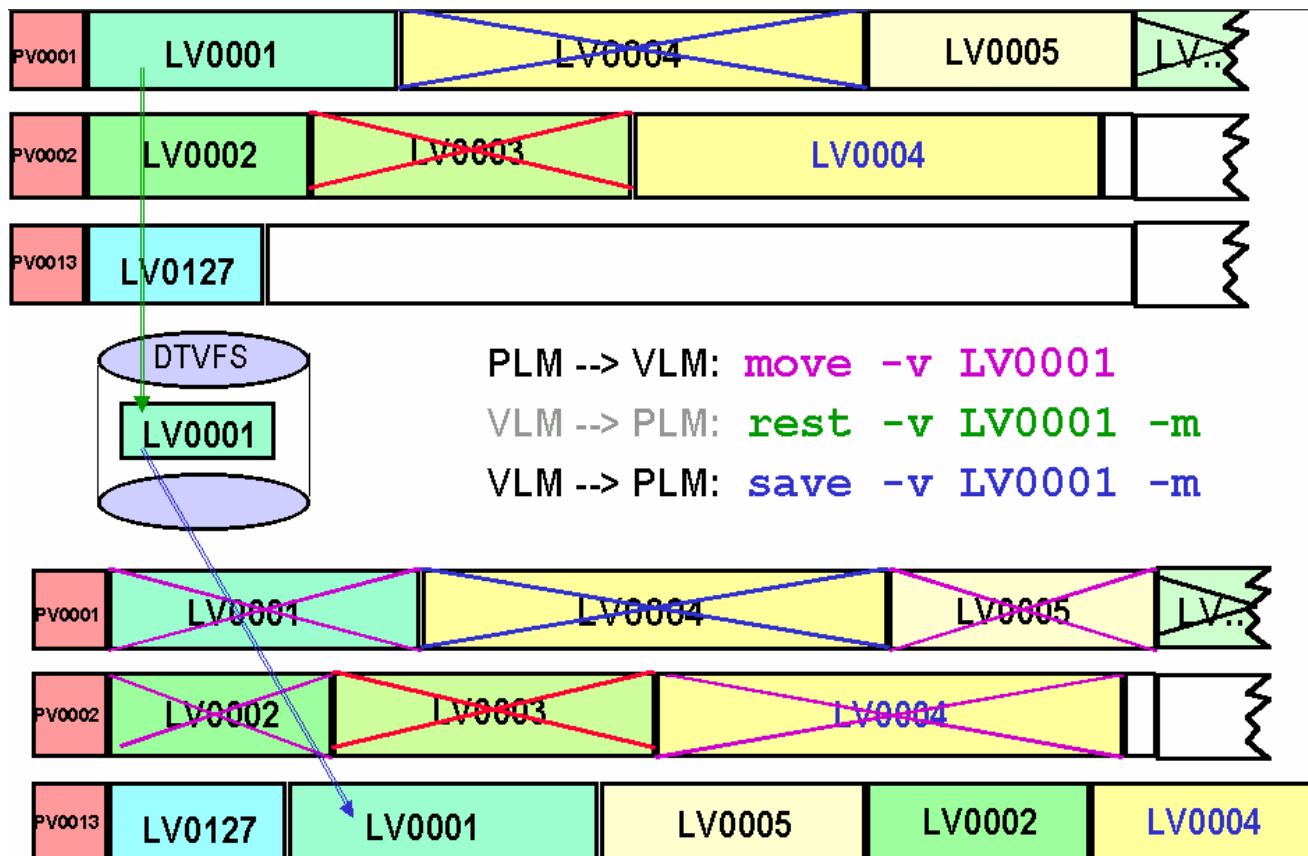
VLM -> PLM: **delete -v LV0003**

VLM -> PLM: **save -v LV0004**



Reorganization of a physical Volume (2)

Reorganization of Stacked-Physical-Volumes



Reorganization of a physical Volume (3)

When starts the reorganization ?


- Through a user command
- Time based, controlled by the user
- »Hard« from the system
- Absolute during emergency by the system
- Faulty physical volumes
- Refreshing of physical volumes

All parameter for the reorganization are specific configurable for each physical Volume Group (PVG).

Reorganization of a physical Volume (4)

 = Softlimit

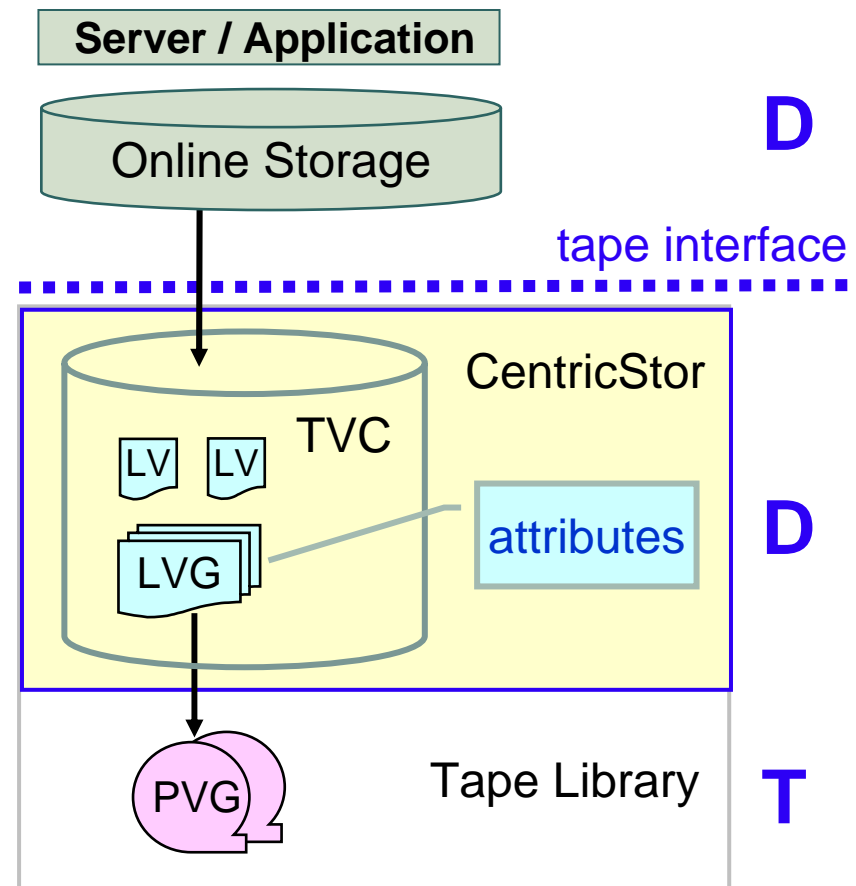
 = Hardlimit



Elements of the virtual tape solution

Realizing D2D2T backup

- D2D2T - disk-to-disk-to-tape
- Virtual Tape Appliance
 - all-in-one box
 - use disk as tape
- TVC - Tape volume cache
- LV - Logical volume
- LVG - Logical volume group
- PV - Physical volume
- PVG - Physical volume group



Volume groups

Clear separation of data

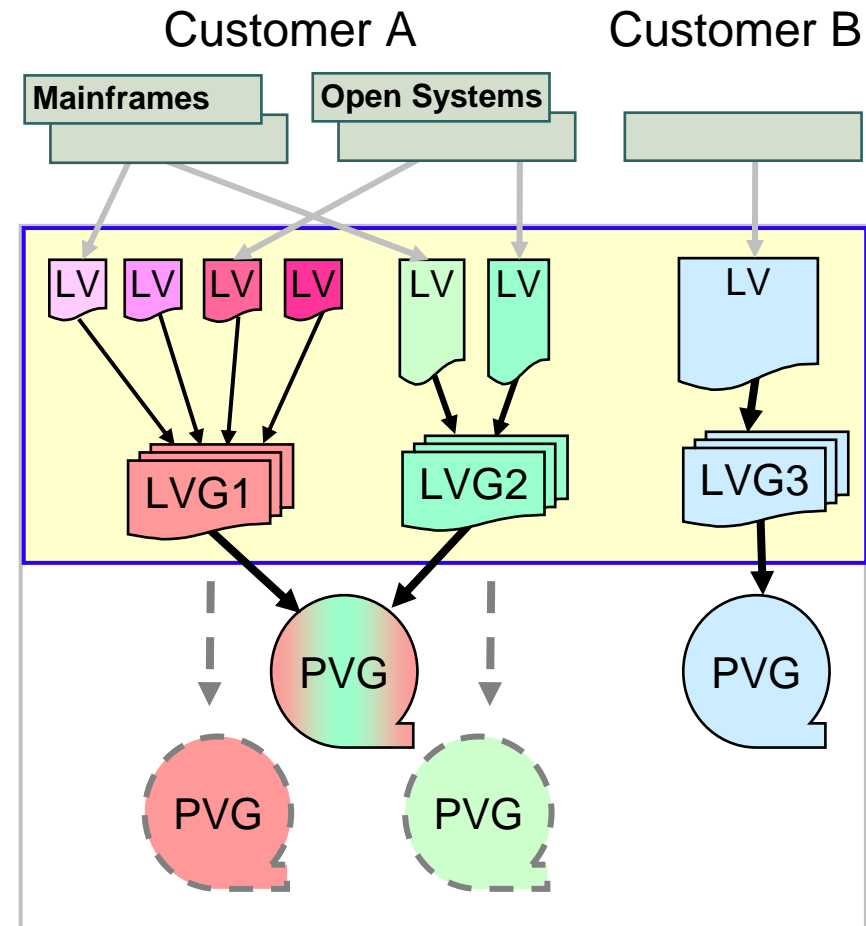
- different systems
- different applications
- different customers
- different value

Logical volume groups

- different service levels

Physical volume groups

- different tape cartridges



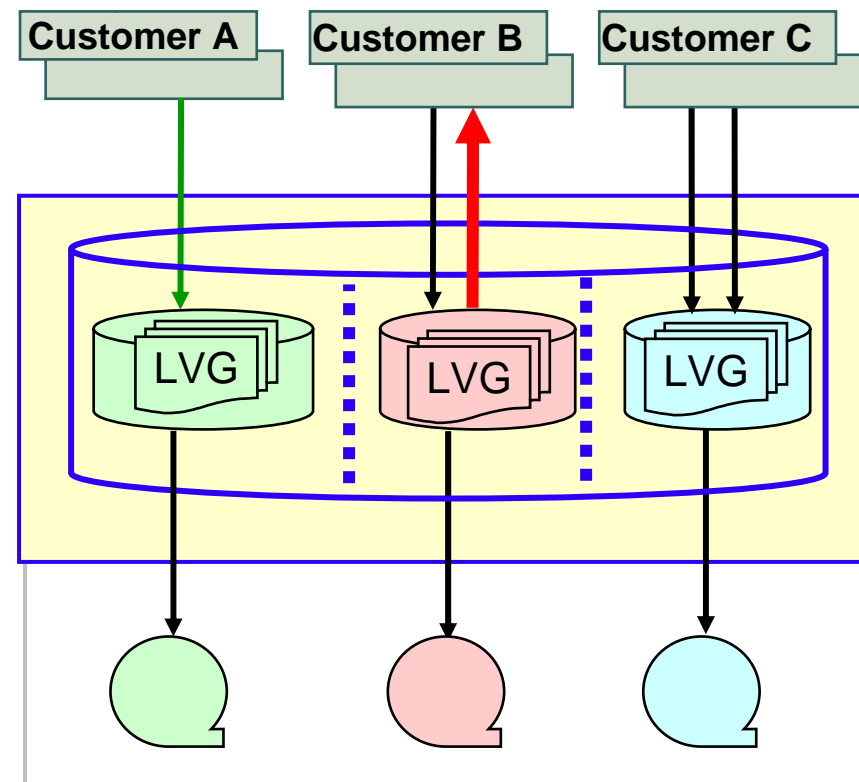
Cache partitioning and disk cache residency

Flexibility for service providers

- cache partitioning
 - strong separation of application / customer data
 - implementation of different SLA

- disk cache residency
 - guarantees fast recovery
 - data immediate out of the cache
 - an attribute for LVG (max. 1TB)

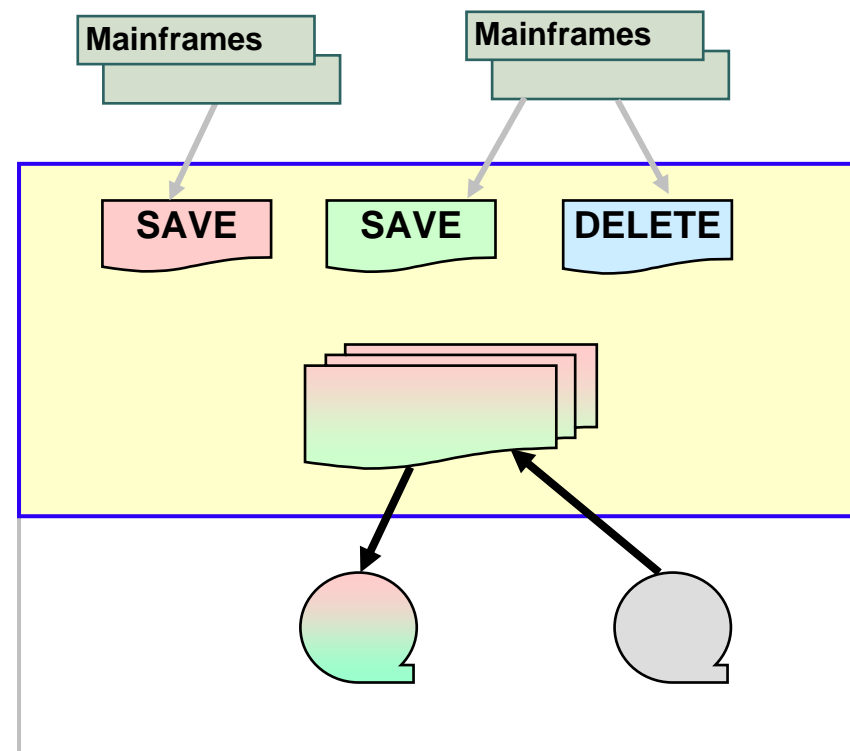
- Increased scalability of disc cache
 - in performance: 2,5 GB/sec
 - in price: S-ATA and/or FC disks
 - in capacity: 176 TB disk cache (native)



Continuous tape operation

Tape handling without server load

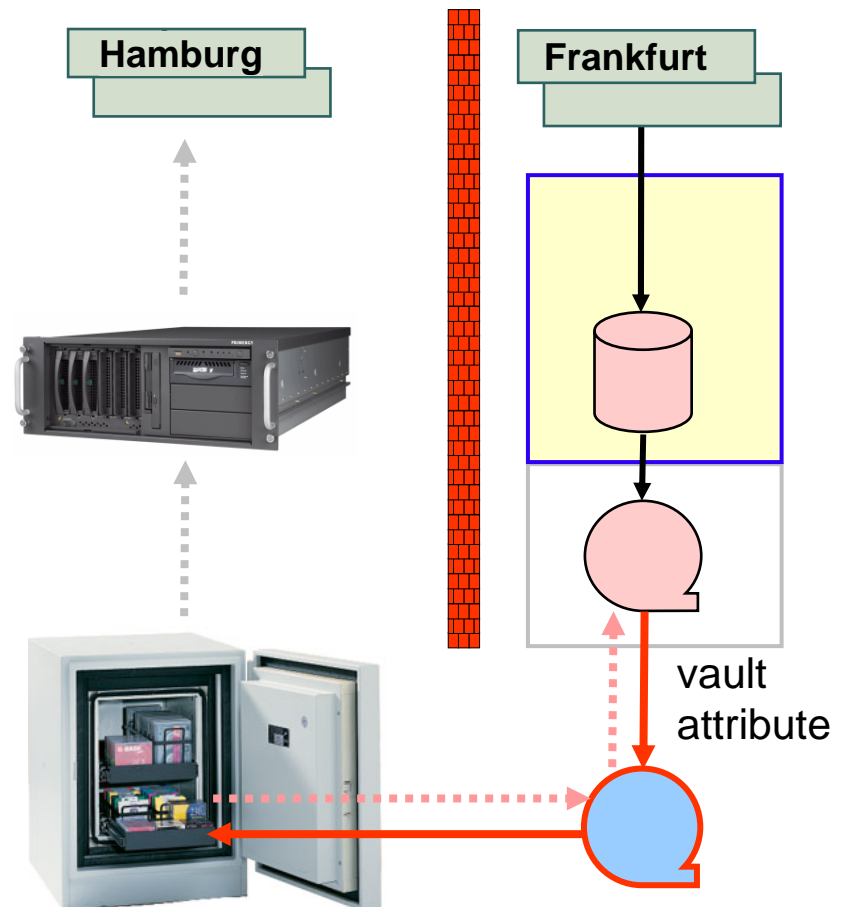
- data written into cache is already compressed
- data from cache
 - gets seamless copy onto tape
 - immediate copy of disk to tape
 - as soon as the server is ready
 - old data in cache will be replaced by new one (LRU = last recently used)
- tape reorganization
 - if data is expired
 - economic usage of free capacity
- tape refresh
 - old tape has to be rewritten
 - new technology generation



Vault tape

Tape on a secure location

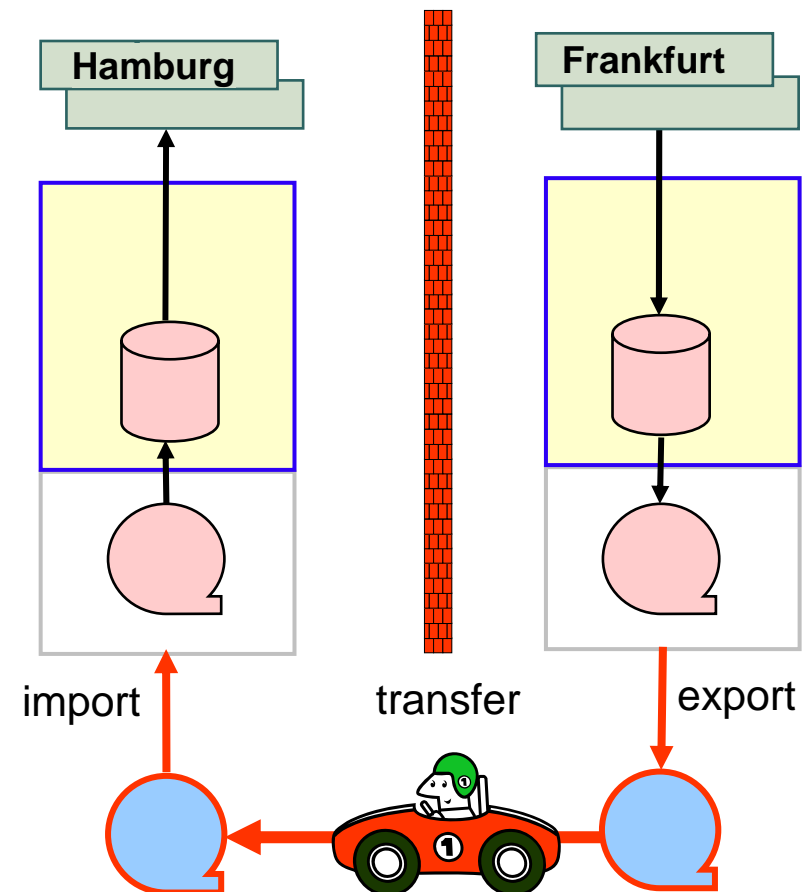
- PVG with vault attribute
- LVG assigned to this PGV'vault
- cartridges with PVG'vault can be removed from the tape library and transferred to a vault location
- PGV'vault keeps marked in the catalogue and can be retrieved
- in addition with dual save this is a third copy
- physical tape also always readable with CentricStor Virtual Tape Controller (VTC)



Transfer tape

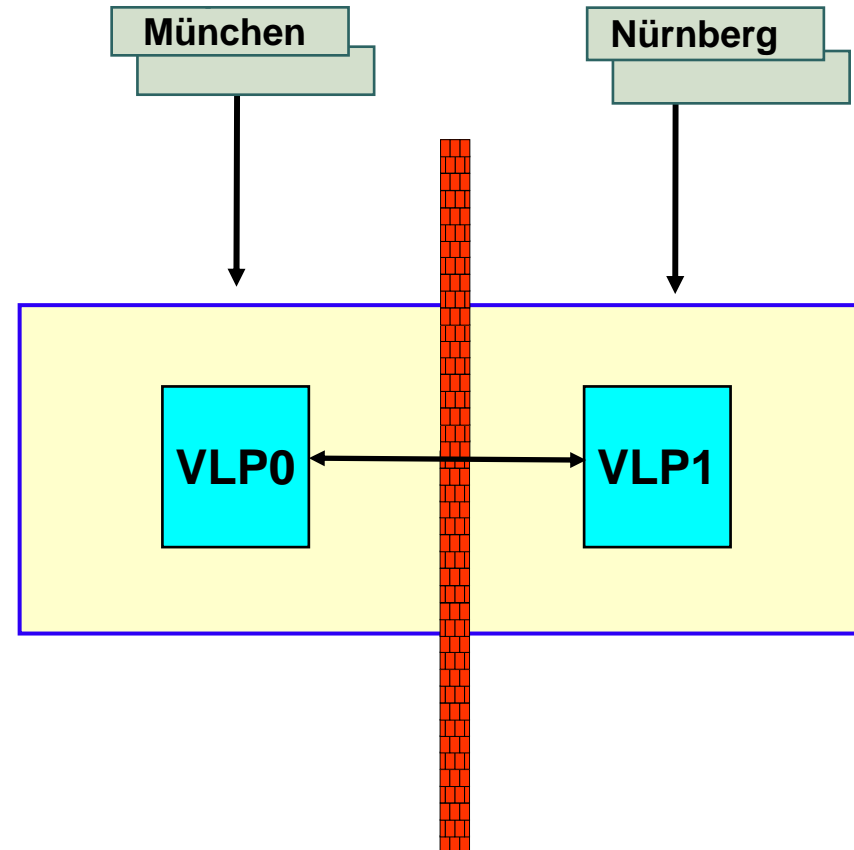
Export & Import of tapes

- copy LVG on PVG'transfer
- export -> cartridges no longer under control of the system
- tape exchange between two CentricStor systems
- import -> can be reinstalled in any other CentricStor system
- physical tape also always readable with CentricStor Virtual Tape Controller (VTC)



VLP Failover

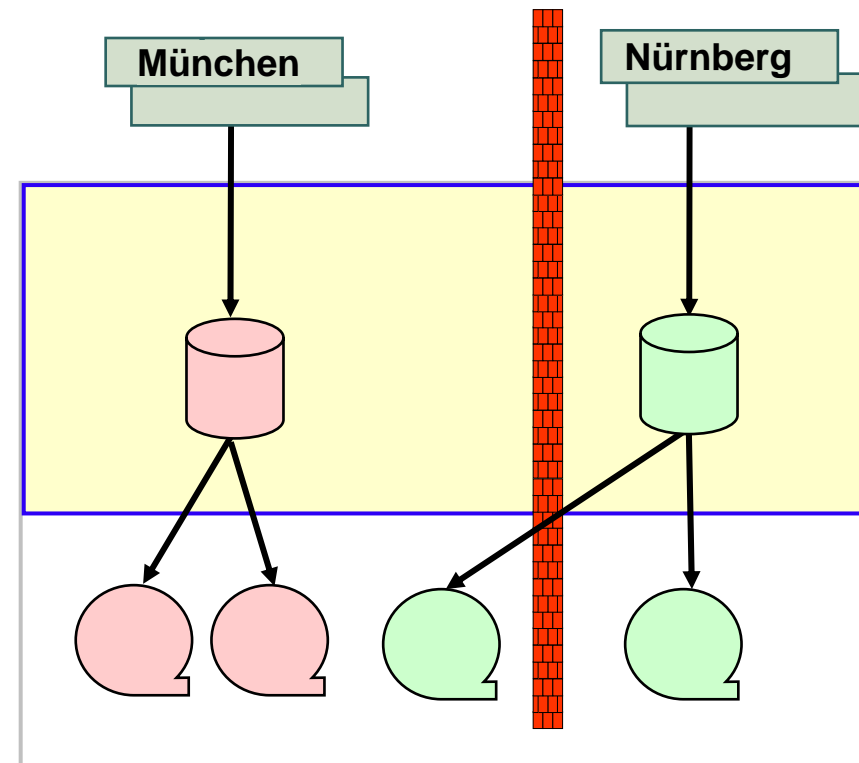
- automatic failover
- VLP1 controls active VLP0
- If VLP0 fails, VLP1 becomes VLP0



Dual save

Physical tape redundancy

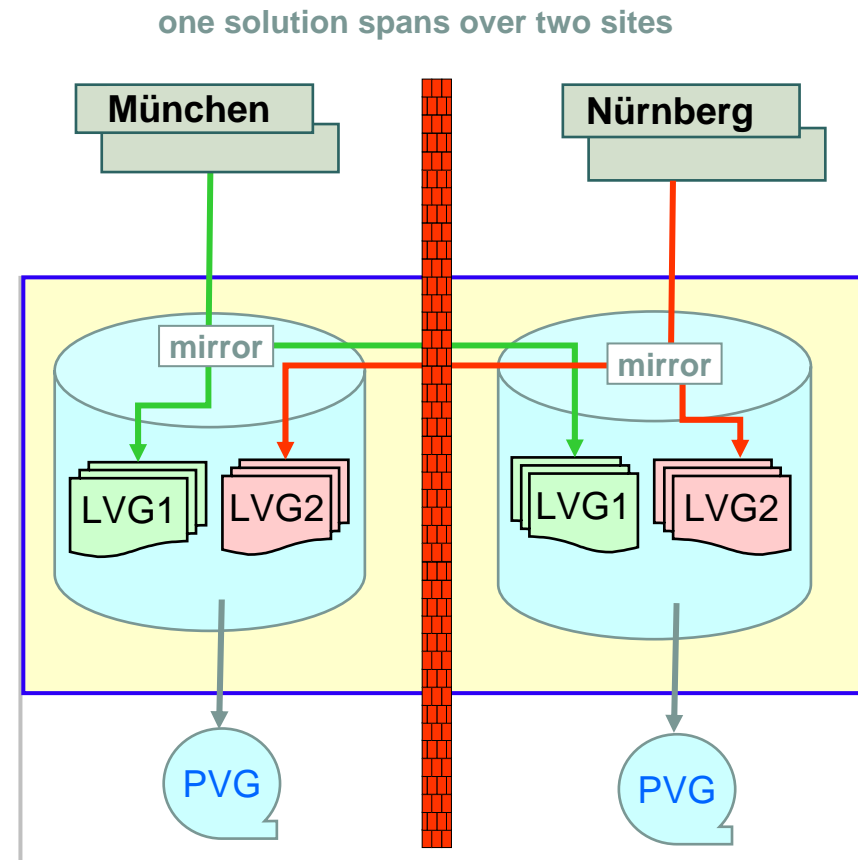
- 1:1 copy
- dual save
 - two tapes be written simultaneously
 - no server load
 - no network load
- remote dual save
 - copy of tape to another location
- two tapes prevent from
 - physical destruction of data
 - human errors



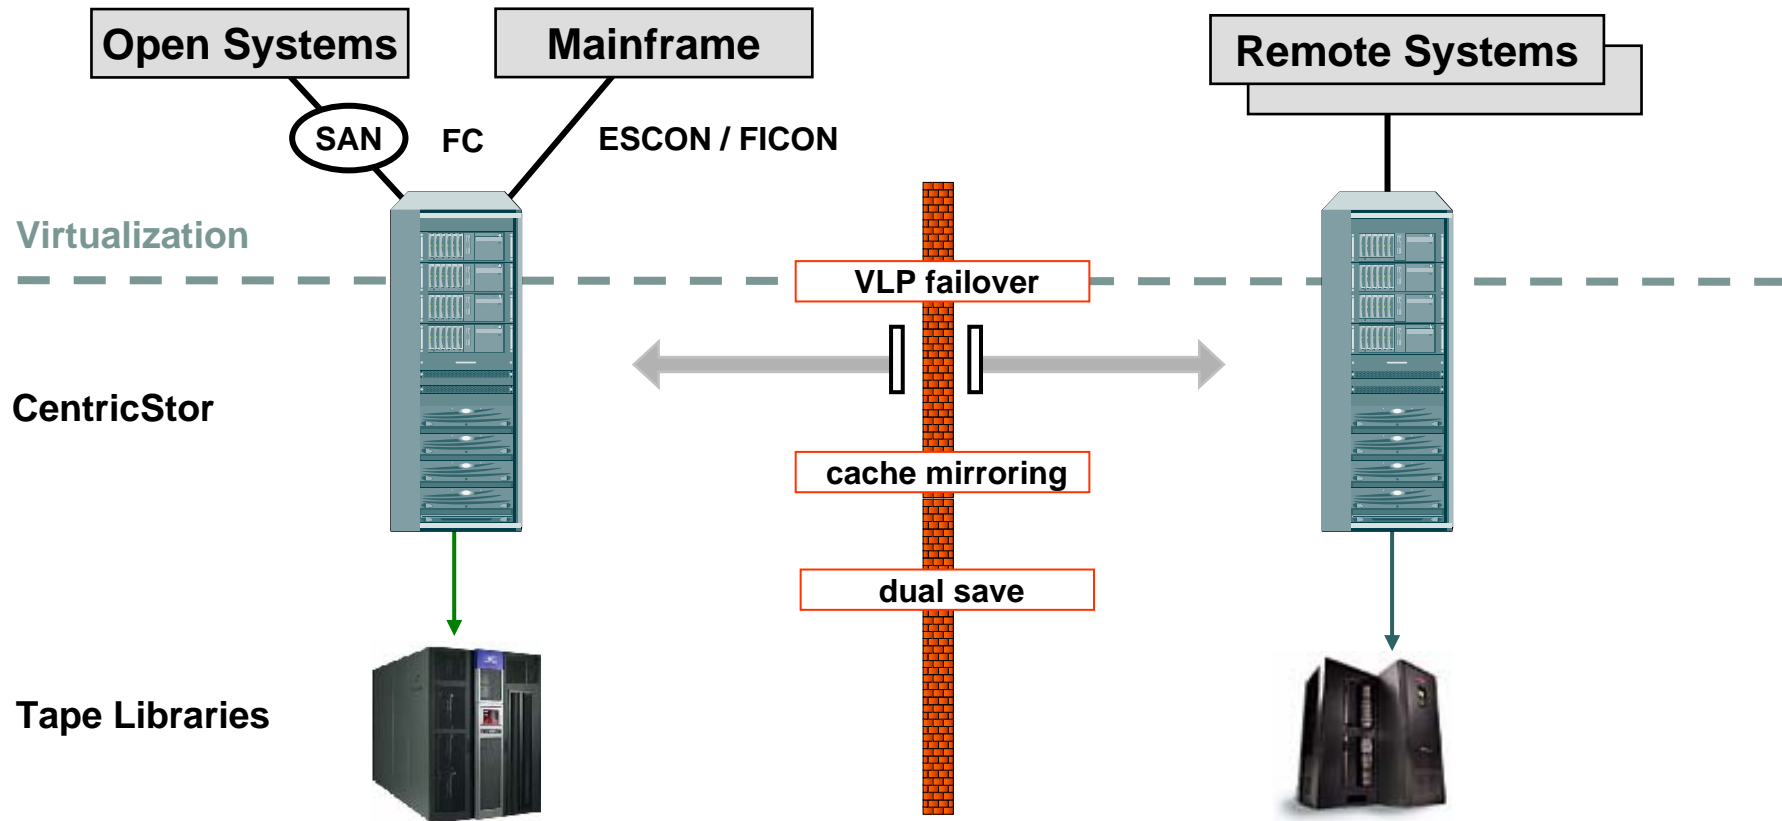
Cache mirroring

Data protection against catastrophes

- real time mirroring
 - for the complete cache
 - implemented in software
 - RAID system independent solution
- data transferred to a second site
 - recovery based on synchronous mirrored data (100% identical)
- extreme reduction of recovery time
 - less complexity than tape recovery



High availability configuration



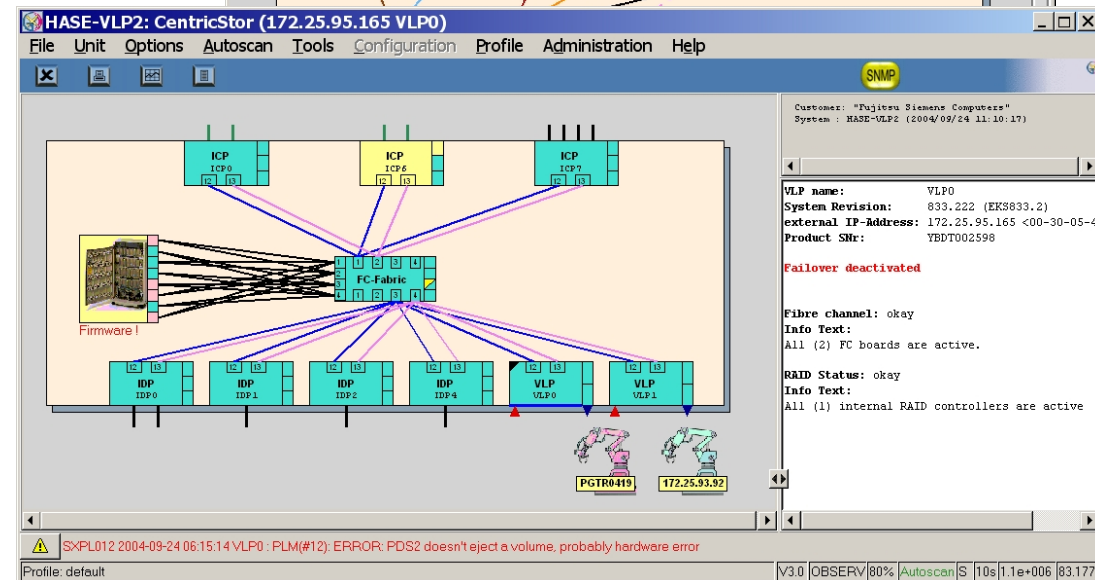
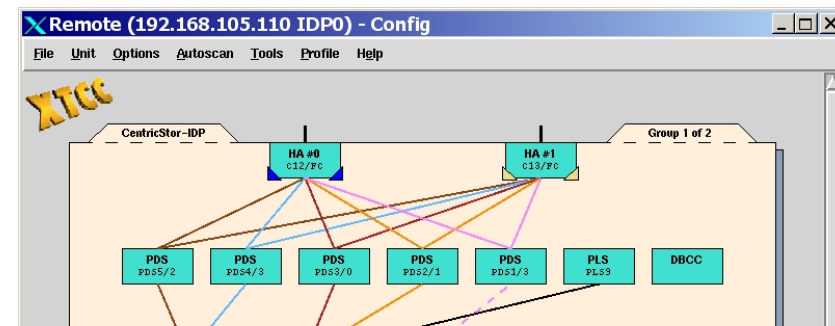
We make sure



CentricStor Administration

Global eXtended Control Center (GXCC)

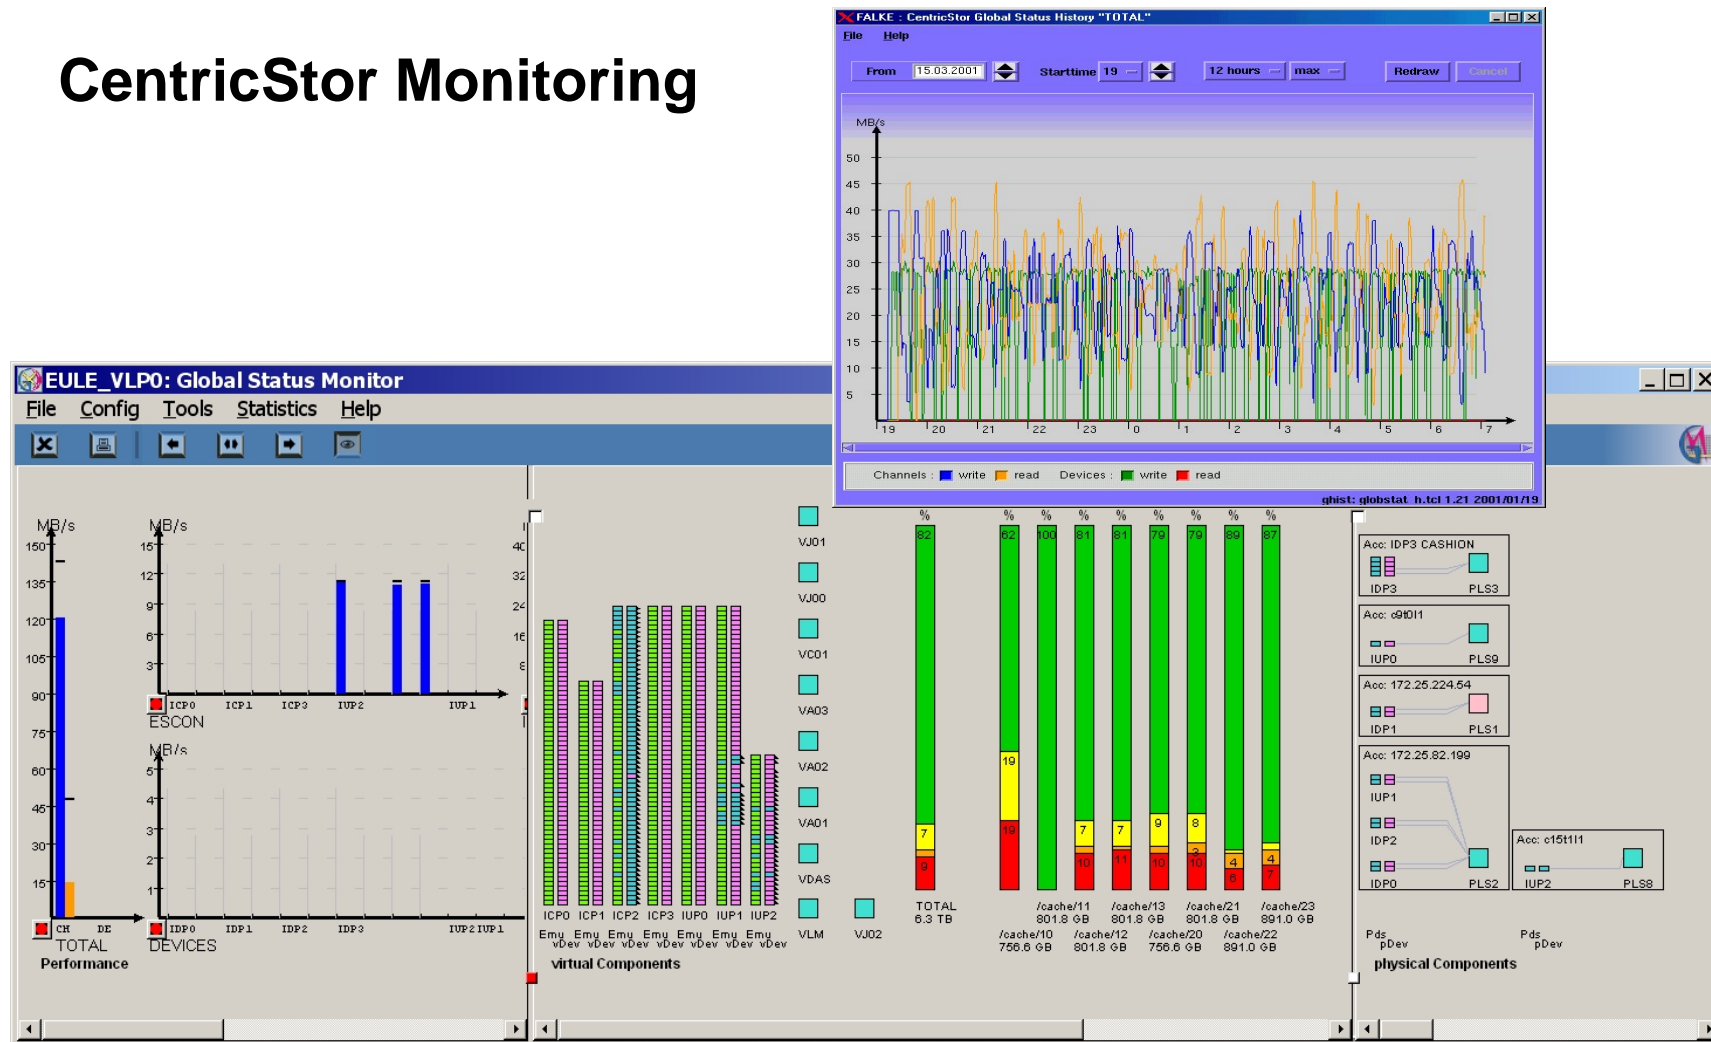
- Installation and configuration
- Current operation status
- Performance statistics
- Also remote via LAN



We make sure



CentricStor Monitoring



We make sure



Summary

Summary

System

- capacity on demand, accounting
- management, remote
- high availability 7x24
- no single point of failure, redundancy
- automatic VLP failover

Logical volumes

- flexible volume size
- logical volume groups
- dual save / remote dual save
- future attributes (WORM,

Tape volume cache (disk)

- cache partitioning
- cache residency
- cache mirror
- cache configuration
(flexibility, scalability, performance, cost)

Physical volumes / physical tape

- physical volume groups
- refresh / reorganization
- export of tape / vault location
- encrypted data on tape
(security of data in transit)

We make sure



Fujitsu Siemens Computers' vision for CentricStor

With CentricStor - based on unique virtualization technology - **we make sure** that our customers reach two goals:

- choose freely among the most cost-effective Nearline storage systems
- achieve a superior level of protection for their data and efficiency of management throughout the entire information lifecycle.

We make sure



Customer Short List



“Do More with Less”

We make sure



**Fujitsu Siemens Computers
Enterprise Storage**

We make sure



questions and answers