


Tivoli Storage Manager at CERN


David Asbury CERN/IT/FIO

HEPIX, Rome, 6 April 2006



Reasons for backup

- ◆ **Accidental corruption of files/mail/dbs**
- ◆ **Storage failure, disks, raid arrays(!)**
- ◆ **Need to get a server running again**
- ◆ **Fire, water, power, cooling problems**
- ◆ ***Backup is very boring until***



CERN Policy on Data

◆ Home Directories

◆ AFS

AFS volume backup -> Castor

◆ Windows DFS

TSM

◆ Mail

◆ Microsoft Exchange

Only Servers

◆ Databases

TSM
TSM

◆ Unix group & project servers

TSM

◆ Experimental Data

Castor

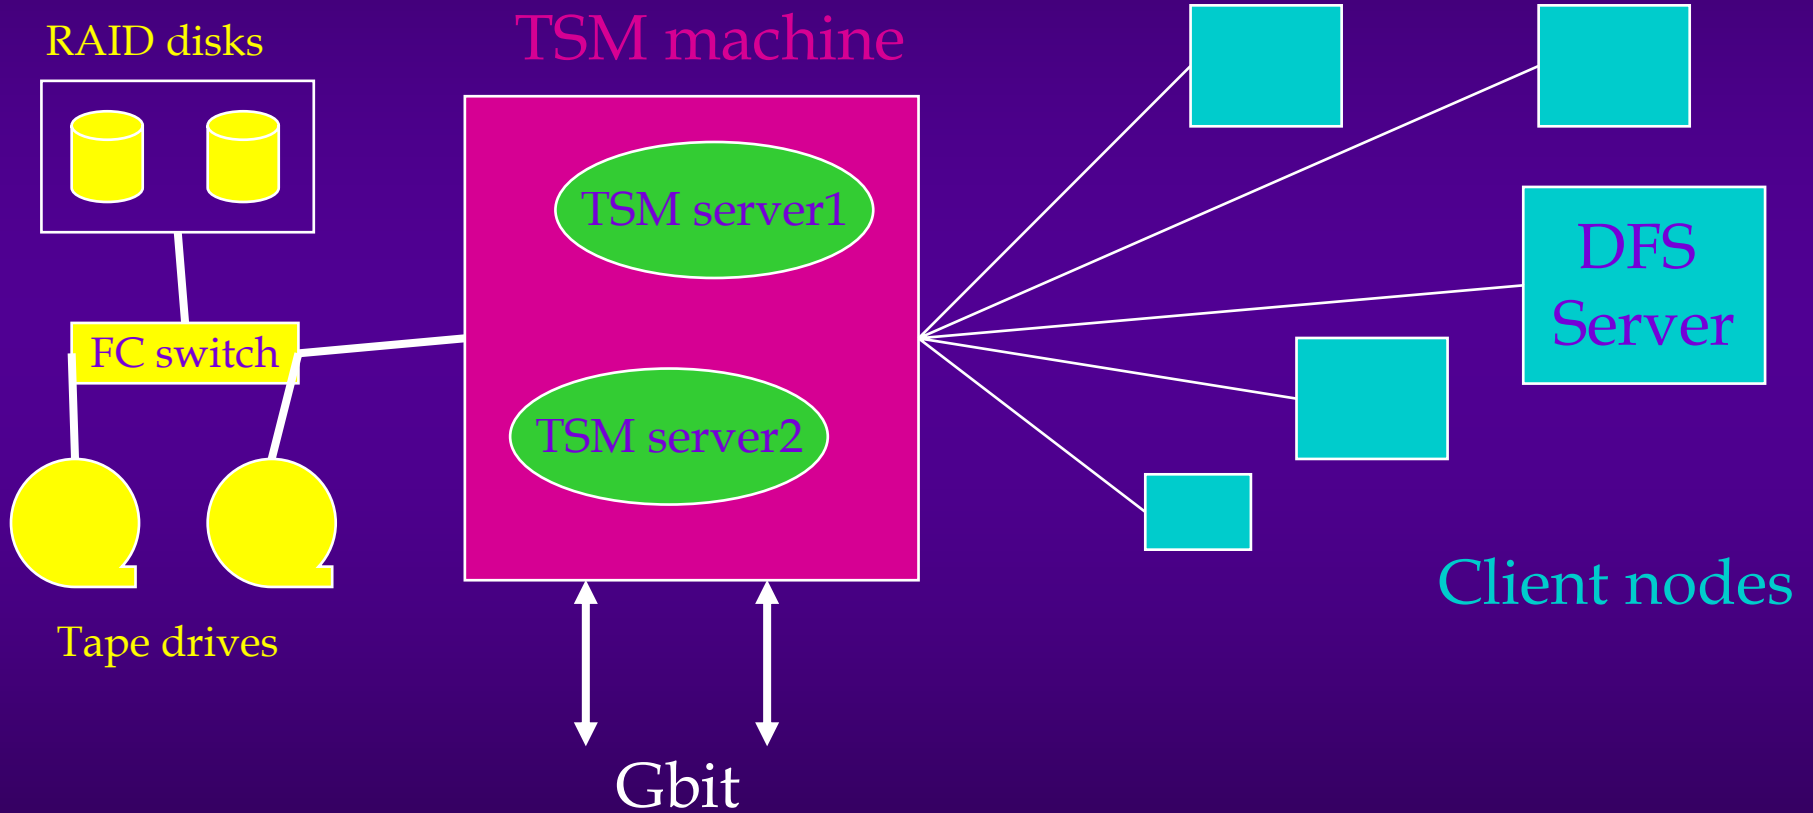



Why use TSM?

- ◆ **Good experience over 15 years**
 - ◆ **Backup & Archive available**
 - ◆ **Collocation supported (so restores are faster)**
 - ◆ **Multiple copies in different locations**
 - ◆ **Very tuneable and scales to a large service**
 - ◆ **TDP available with RMAN for dbs (incremental)**
 - ◆ **Good monitoring software – Servergraph**
 - ◆ **Cost saving from using single product (stop Legato)**



Terminology






Service Revived in 2003

- ◆ **2 x old IBM F50 machines**
- ◆ **Added modern IBM p630 machine in 2003**
- ◆ **FC switch to attach disks & tapes**
- ◆ **Infortrend disk RAID arrays**
- ◆ **Added IBM p5-550 machine (in another building), connected by FC**

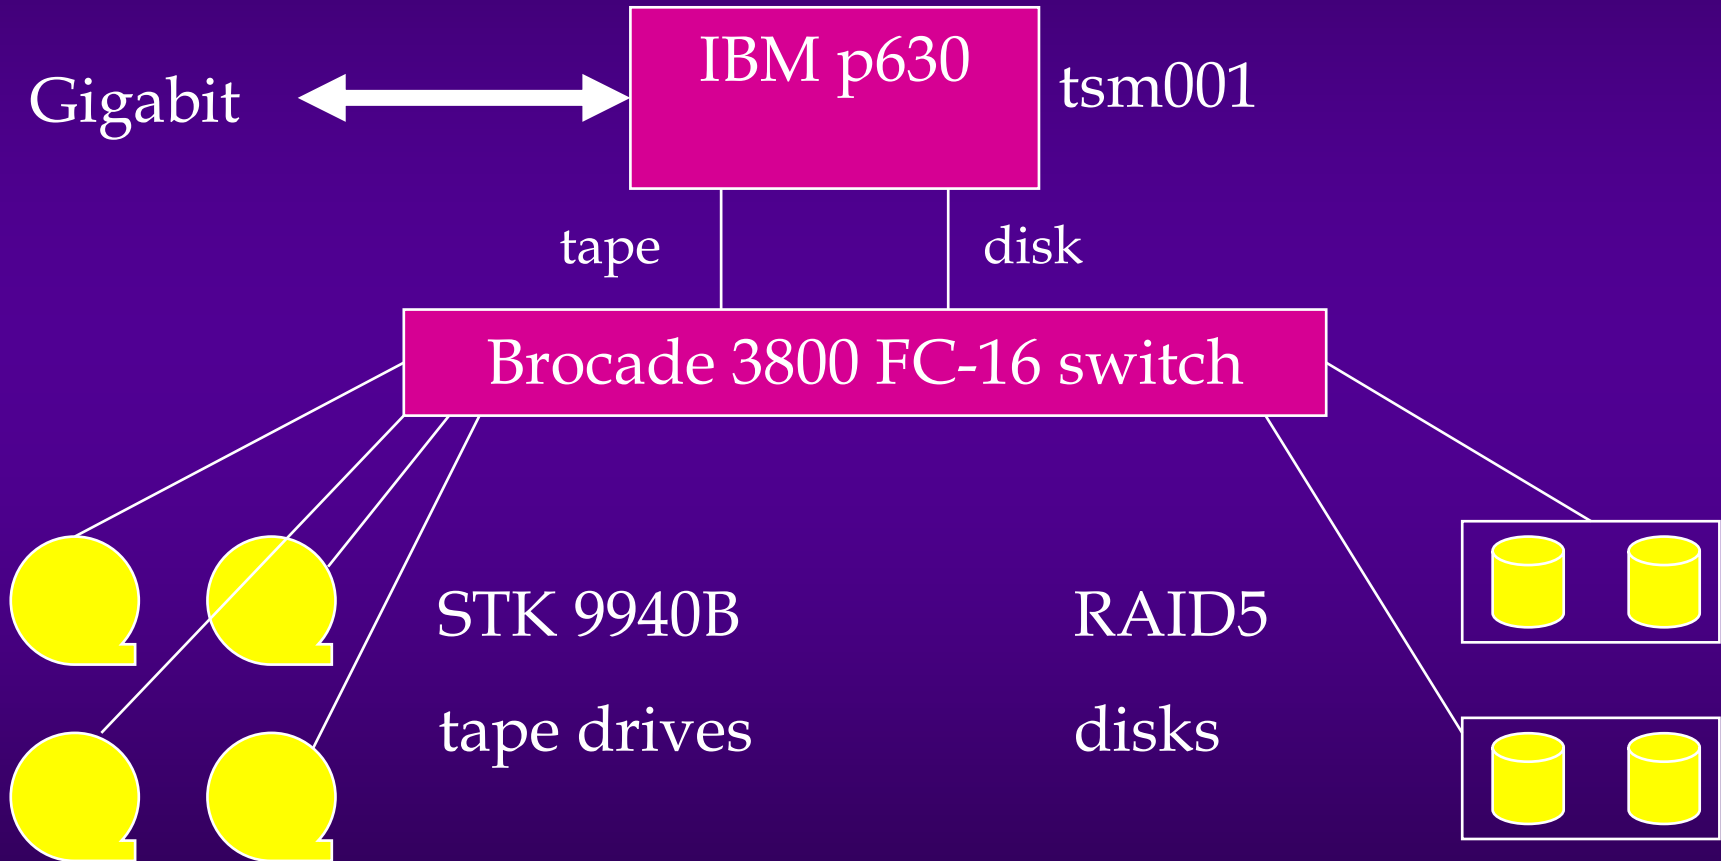



New Machine hardware

- ◆ **IBM p5-550, AIX 5**
 - ◆ **2 x cpus 1.6 GHz**
 - ◆ **8GB memory**
 - ◆ **2 fiber-channel HBAs (disk, tape)**
 - ◆ **2 Gigabit interfaces**
- ◆ **Brocade 3800 fiber-channel switch 16-port**
- ◆ **2 Infortrend SATA disk RAID trays**
 - ◆ **5.5TB usable RAID5 staging space**
- ◆ **1 Infortrend FC disk tray (for TSM dbs)**
- ◆ **5 STK 9940B tape drives (200GB)**

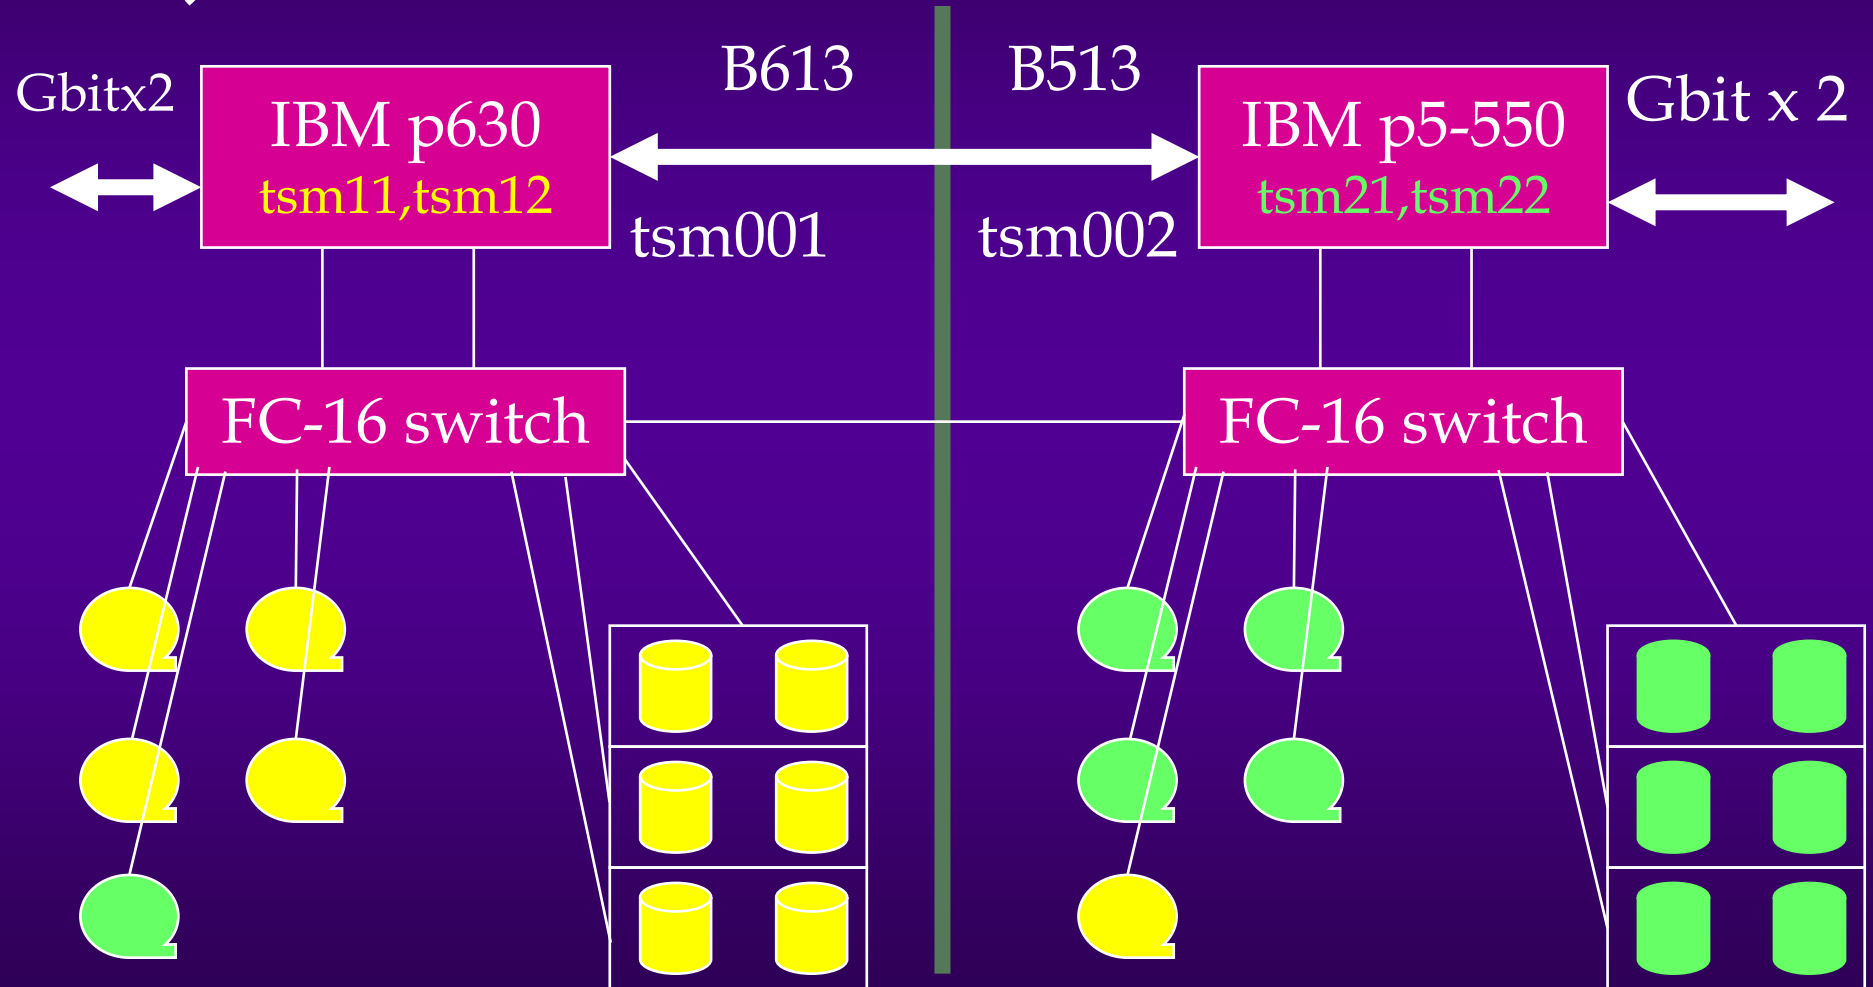



1st Machine Config





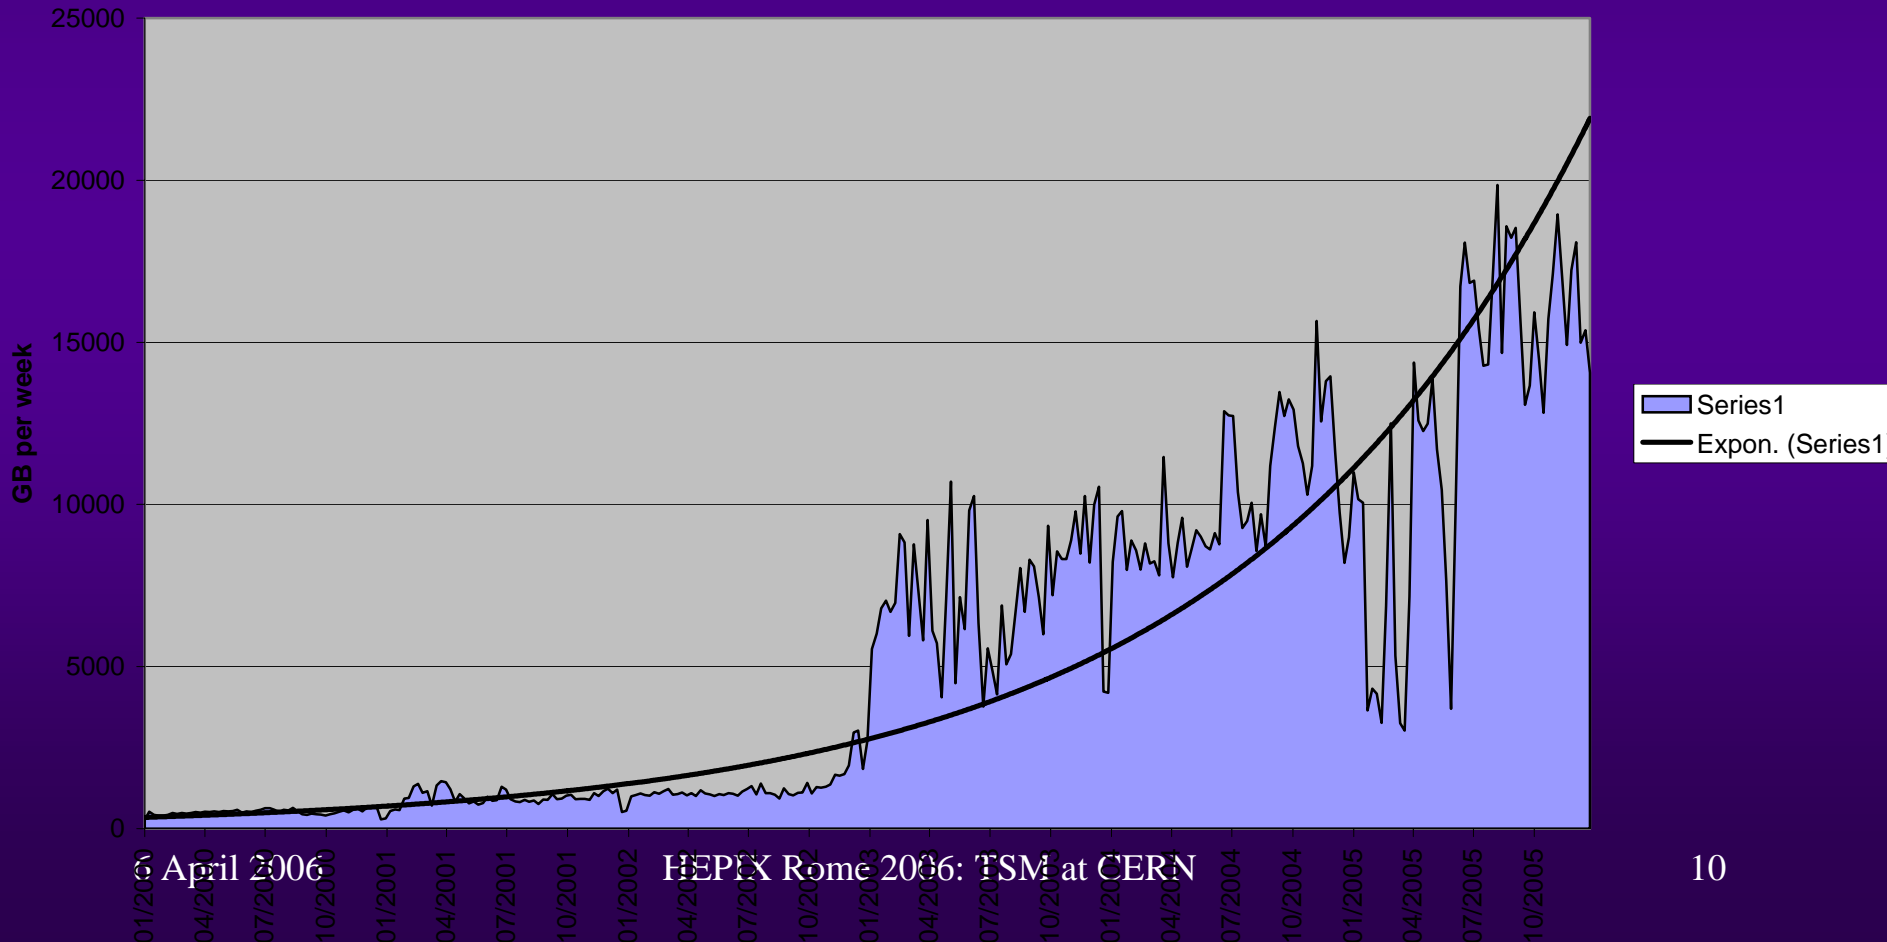
2nd Backup machine






Growth!

All servers GB per week





Monitoring Software

◆ Servergraph/TSM

- ◆ **Monitors all licensed TSM servers**
- ◆ **Mature product, understands TSM well**
- ◆ **Web interface - extendable**
- ◆ **Email warnings**
 - ◆ To users about failed backups
 - ◆ To admin about server and hardware problems
- ◆ **Capacity Planning**




TSM on a Linux machine

- ◆ Reduce capital costs
- ◆ Easier to prevent scalability problems
- ◆ Being tested in CERN
- ◆ Recent version supports STK ACSLS



TDP for ORACLE

- ◆ **Tivoli Data Protection for Oracle**
- ◆ **Allows incremental backups with RMAN**
- ◆ **Important for future Physics meta-dbs**
- ◆ **~60 databases in CERN now**

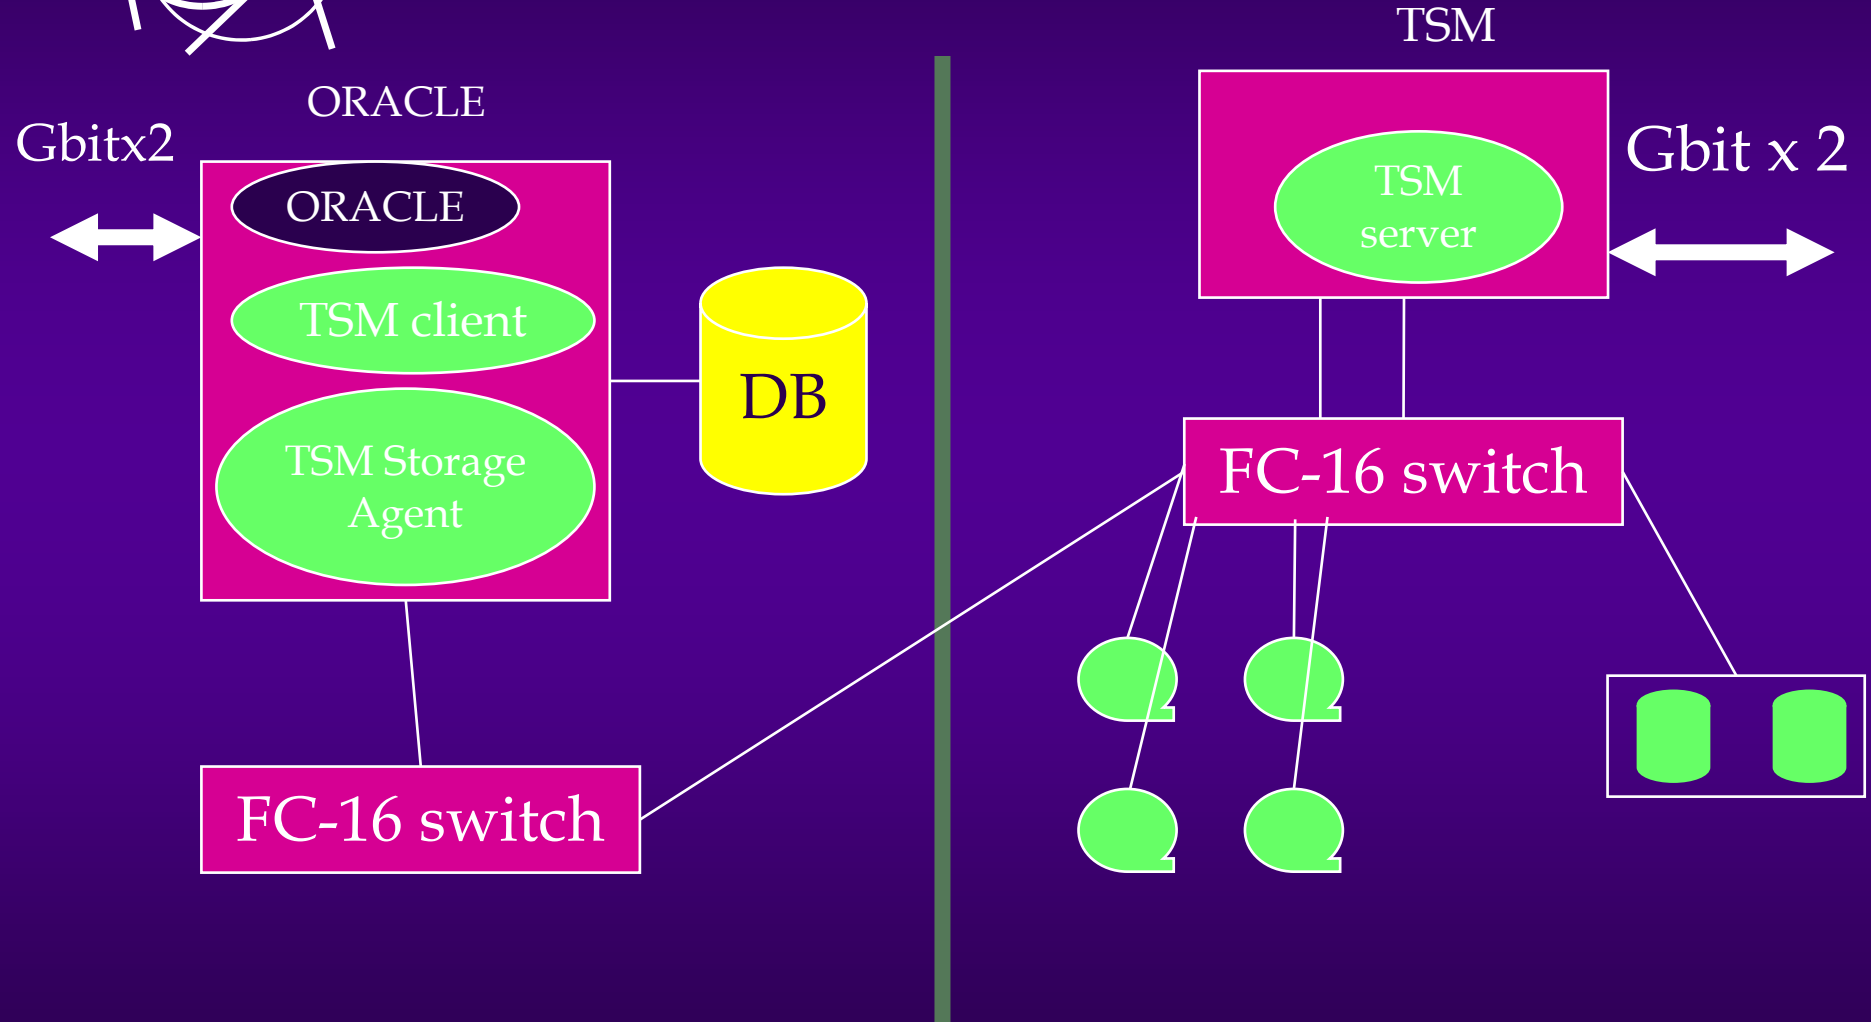



LANless Backups

- ◆ Client to TSM tape via fiber-channel
- ◆ Avoids TCP/IP overhead in client & server
- ◆ Avoids bottlenecks in LAN
- ◆ Needs available tape drives (as before)
- ◆ Still need staging space for REDO logs
- ◆ High rate of metadata from LHC



LANless Backup







Disk contention effects

“normal” operations

Continuous migration!


stk3-stk52a5,mount
stk3-stk52a6,mount
stk3-stk52a7,mount
stk3-stk52a8,mount






Scalability

- ◆ **Management of TSM database**
 - ◆ **IBM recommend not letting db grow too big**
 - ◆ **db on tsm1 is already too big (> 100GB)**
 - ◆ **Daily dump of the db takes >3 hours**
 - ◆ **Expiry of old items is already time limited**
- ◆ **DB very active unless**
 - ◆ **Client uses journaling**
 - ◆ **Incremental by date**



Server architecture validated

- ◆ **IBM server and fiber channel work well**
 - ◆ **Should allow expansion for some time ahead**
 - ◆ **Can increase workload by adding disks & tape drives**
 - ◆ **Can activate 2nd Gigabit interface when needed**
 - ◆ **Fiber channel switch aids problem solving**



Questions?