

# Data infrastructure and services at CSC

for FUUG 5th Mar 2009



# Main components

- **2 x SAN**
- **1 x USP 100**
- **2 x AMS 1000**
- **2 x AMS 2500**
- **1 x USP VM**
- **1 x Storagetek SL8500 robot**
- **1 x ADIC i2000 robot for backups**
- **NFS server clusters ( 2 x internal + 2 x external)**
- **Archive server**
- **support/administration/test servers**
- **Murska + Louhi internal work areas**

# SAN

## ➤ **Storage Area Network**

- 2 separate networks for redundancy
- 8 or 4Gbit/s links to hosts
- 2x or more Trunk links between switches
- 10+10 Brocade Switches

## ➤ **SAN attached storage**

- No service breaks
- Scalable architecture (size, performance)
- 2 Symmetric path to disk
- Excellent interoperability
- Dedicated software support available
- 100 % usability



# AMS1000

- **FC + SATA disks**
- **Troughput 1520MBytes/s**  
=> 12Gbit/s
- **300 000 iops (Burst) per AMS**
- **Cache: 2 x 8 GB per AMS**
- **~850 Disks**
- **AMS1 Disks:**
  - 32% FC 300 GB 10k
  - 27% SATA 500 GB 7.2k
  - 41 % SATA 750 GB 7.2k
- **AMS2 Disks:**
  - 36 % FC 300 GB 10k
  - 47 % SATA 750 GB 7.2k
  - 17 % SATA 1 TB 7.2k





# AMS2500

- SAS+SATA disks
- Troughput 4800MBytes/s  
=> 38Gbit/s
- 900 000 iops (Burst) per AMS
- Cache: 2 x 8 GB per AMS
- up to 2\*480=>960 Disks
- AMS2500-1 Disks:
  - 68% SAS
  - 32% SATA
- AMS2500-2 Disks:
  - 68 % FC
  - 32 % SATA



# USP VM

- **Storage virtualization**
  - AMS 1000-1
  - AMS 1000-2
  - AMS 2500-1
  - AMS 2500-2
  - future storage systems
- **transparent migration to new storage**
- **96PB external storage capacity**
- **Cache 12GB**
- **No internal disks**



# USP 100

- **2 000 000 iops**
- **Heavy duty databases**
- **Expensive storage**
- **Library systems**
- **No downtime due to upgrades or maintenance**
- **Cache:24 GB**
- **Disks:**
  - 70% FC 146 GB 10k
  - 30% FC 146 GB 15k



# Storagetek SL8500

- 4 Tape drives (up to 448)
- 500/1470 Tape slots (up to 70 000 slots)
- 56 PB of storage capacity if needed
- Longterm Archival solution



# ADIC i2000 Robot

- **600 slots (max 3000 slots)**
- **4 tape drives**
- **Dedicated to Backups**
- **Capacity 240TB**
- **EMC Legato NetWorker backup software.**  
**All backups use Ethernet networks**



# NFS server clusters

## ➤ **SUN T5220 Servers (External)**

- Project disks
- Metawork
- Home directories
- Application Disk areas
- NFS services
- ZFS filesystem (no 2TB limit)

## ➤ **Sun Cluster 3.2**

## ➤ **SUN v240 Servers (internal)**

- Internal home directories
- Internal Application Disk area
- Internal Veppi
- Possibility to create internal project disk area
- NFS + CIFS services
- ZFS filesystem (no 2TB limit)



# Archive server

- **SUN V880 server**
- **Controls SL8500 Robot**
- **SAM-FS archive filesystem**
- **Slow seek 0-3min (if not in disk pool)**
- **fast throughput 80MB/s (from tape)**
- **400GBytes per file limitation**
- **Choke on small files (over NFS)**
- **Creates 2 copies of data**
- **“Green” storage**



# Support/administration/test servers

- **SUN (sparc) Solaris**
- **HP Blade (Intel x86\_64) Linux**
- **Filesystem compare tests**
- **Storage IO tests**
- **Failover testing**
- **Monitoring**






# Storage Requirements

- **Timeframe**
- **Size**
  - 100G-100TB ?
- **Speed**
  - SATA/FC/SSD ?
- **Cost per TB**
- **Power consumption**
- **Service uptime?**
- **target OS (Linux/Win/Solaris)**
- **Filesystem ( ZFS/XFS/EXT3/NTFS/FAT)**
- **lifecycle of the service**



# IOPS or SPACE ?

- **Input/Output operations per second**
  - Depends on read/write
  - block size
  - Used disk type
  - Access pattern (random vs streaming)
- **1 SATA Disk can achieve**
  - up to 120MB/s sustain read (streaming type data)
  - ~70 IOPS
- **1 LUN from SAN storage**
  - True Capacity (WYSIWYG)
  - IO power of multiple disks
  - Controller Cache to improve IO performance
  - Redundant



1 SATA disk




RAID5+1 over 10K FC



RAID5+1 over 7.2 SATA



RAID5+1 over 15K FC



# Power consumption 6TB of data

## ➤ **“Green way”**

- one LTO3 Tapedrive consumes 12W
- Writing speed 80 - 120 MB/s  $\Rightarrow$  21h
- Takes 15 tapes
- Total Power consumption 0,25 kWh

## ➤ **Online disk storage**

- one 500 GB SATA Disk consumes 6-12W  
15 Disks (incl. hotspare + filesystem)  
 $\Rightarrow$  180W/6 TB
- Total power consumption 4.2 kWh per day

# Questions?

Want to work with this gear?

