

# Disk storage procurements

Helge Meinhard  
CERN-IT  
29-Nov-2007

- Most (> 90%) capacity used as temporary data buffer between tapes and CPUs
  - Sequential access of modest number of streams
  - 7'200 rpm SATA disks fine
  - Most cost-effective solution: “NAS-style” boxes
    - Connected via single Gigabit Ethernet
    - Integrated PC servers with PCI RAID cards
    - 16...24 disks, 5...6 TB usable
      - More capacity not sensible, as bandwidth does not scale up with capacity
      - Highest capacity disks not adequate, as multiple streams push for more spindles

- Some database storage required as well
  - Metadata for physics data on tape and/or on “NAS-style” boxes
  - Engineering, administrative, financial data
- Chosen solution: Oracle RAC
  - Running on a number of clusters
    - Mixture of FC SAN and multi-GigE network topologies
    - For the former, a number of RAID arrays with FC uplink is required

- Minor storage requirements
  - Home directories (AFS)
  - Backup (TSM)
- Database-type random access
- RAID arrays with FC uplinks to SAN infrastructure well suited

- More than 5 PB usable (to be) ordered
- 4.5 PB as 800 big “NAS-style” boxes with 16...24 disks of 500 GB
- Some smaller “NAS-style” boxes with 8 disks of 500 GB
- 80 RAID arrays, 8...24 disks, 150 GB...500 GB/disk

- **First step: Market Survey**
  - Qualify companies to receive invitations to tender for 12 months
  - Criteria: commercial rather than technical
  - Member-state origin
- **Second step: Invitation to tender**
  - Either detailed component specifications, or constraints and capacity/performance reqs
- **Lowest compliant bid wins**
  - No account for added value
- **Large orders require Member States' approval**