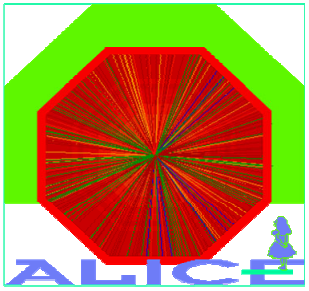

Evaluation of Computing Technologies for the ALICE Data Acquisition

Klaus.Schossmaier@cern.ch

CERN PH-AID

Visitors from UK Department of Trade and Industry (DTI)

30 June 2004

DAQ Architecture

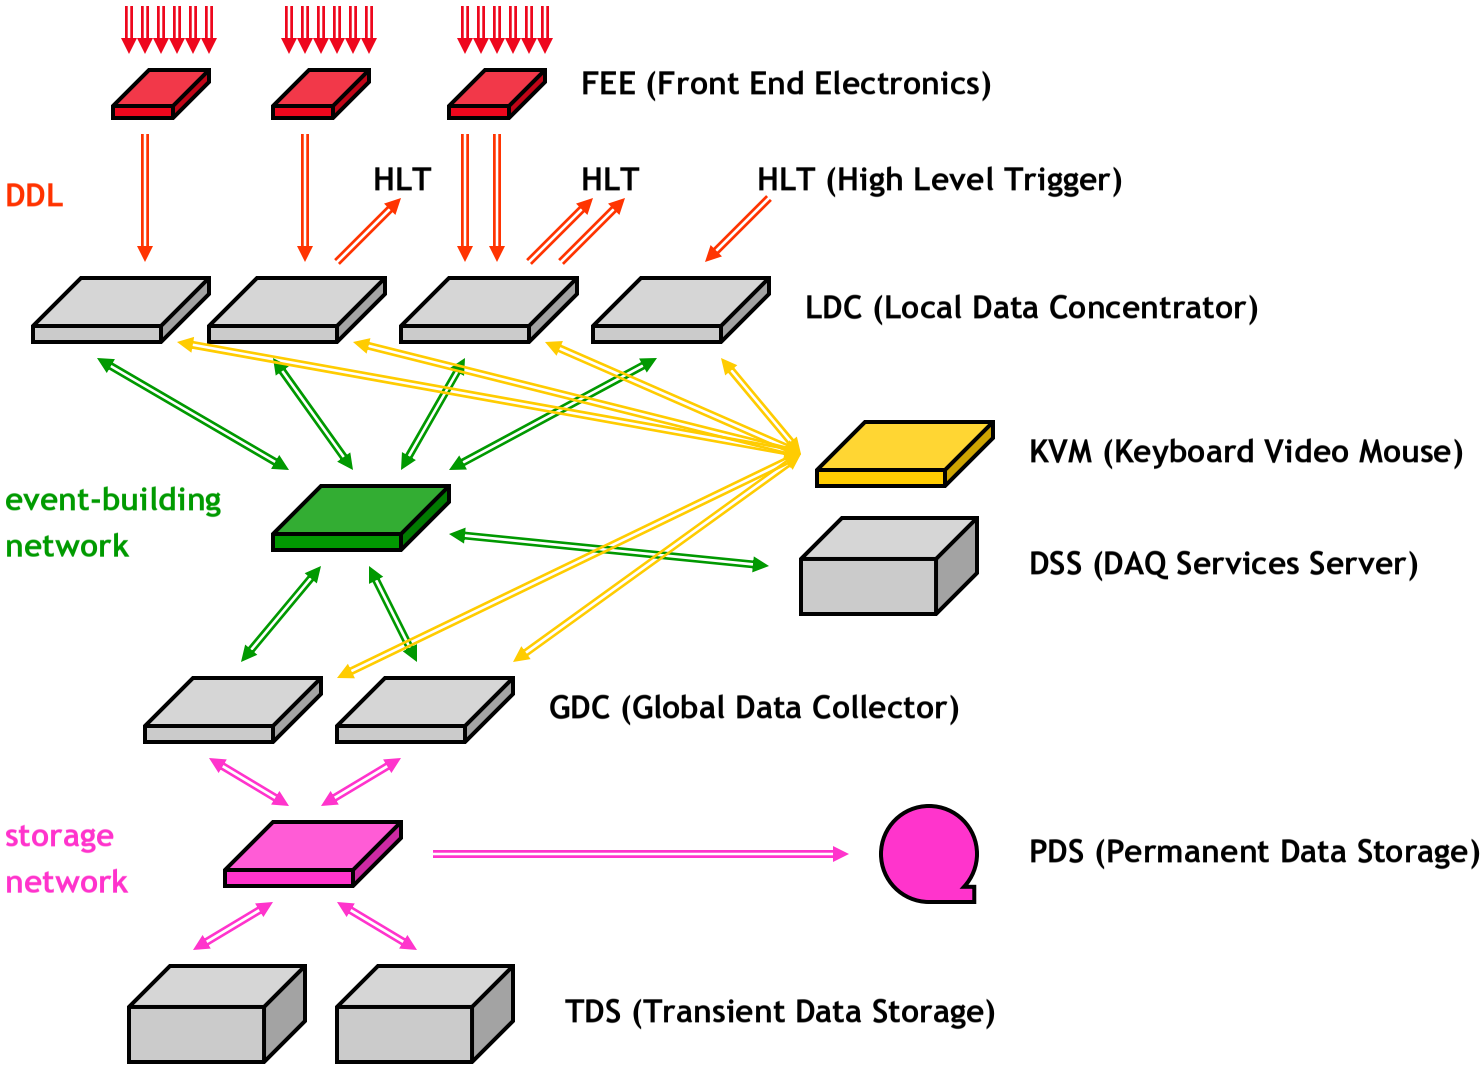

DAQ Installations

DAQ System	DDL	LDC+GDC+DSS	Location
Development + Evaluation	2	10	AID lab CERN
DDL Development	2	2	DDL lab CERN
Detector Groups	1	1	Institutes
Test Beams	2	4	Test Beam areas
ALICE Data Challenges	0	80	IT CERN
Reference Setup	6	7	AID lab CERN
Test + Commissioning	25	15	LHC Point 2
Final System	422	254	LHC Point 2

DAQ Examples

Reference System:

to be finished end of 2004

Test Beam for TPC Sector Test:

May 2004, 4x1 system, 533 GB

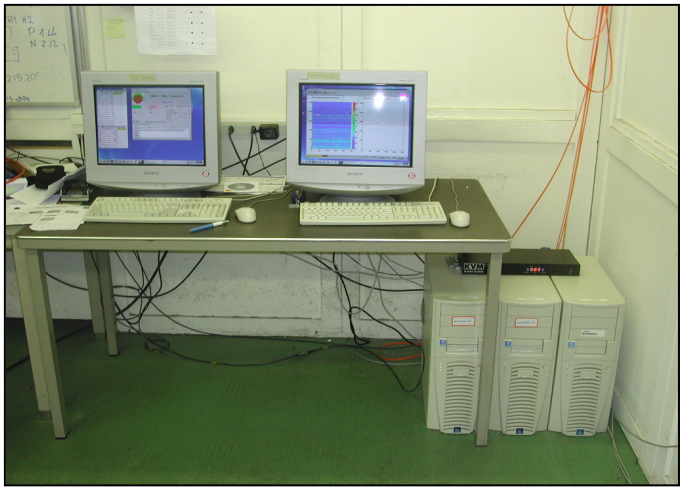

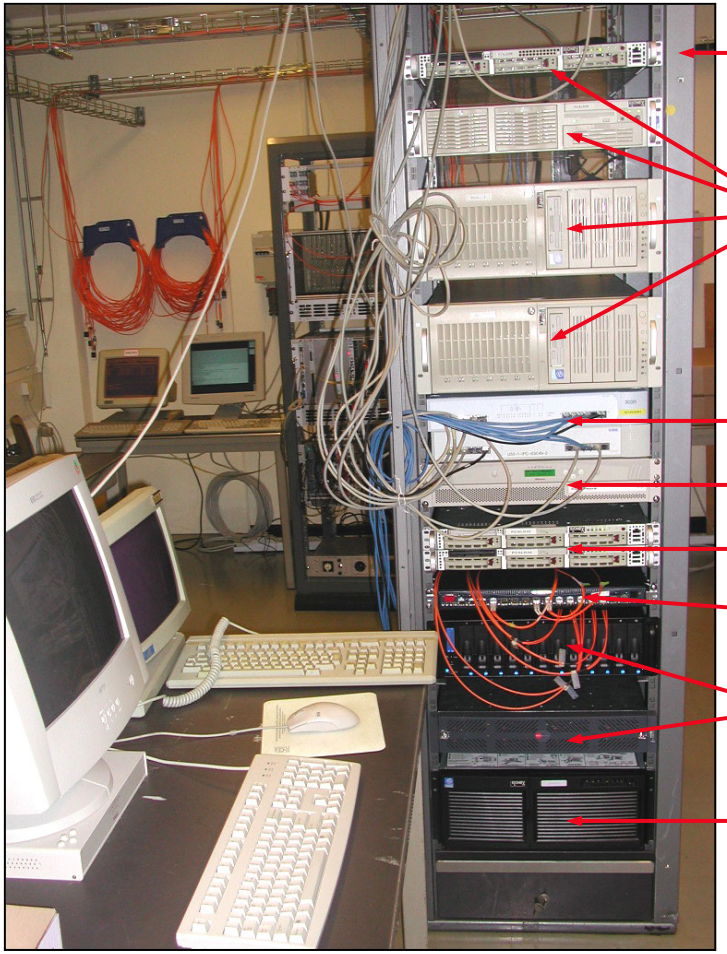

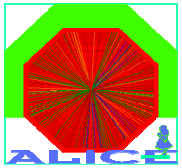

Computing Elements

- Platforms
  - Single/Dual/Quad CPU
  - Rack-mount and VMEbus form-factor
  - PCI or PCI-X slots for RORCs and FC adaptors
  - Linux operating system (Red Hat)
  - DATE software including custom drivers
- Networking
  - Gigabit Ethernet adaptors/switches
  - Fiber Channel adaptors/switches
- Data Storage
  - Disk Array (RAID level 0 or 5)
- Infrastructure
  - KVM (keyboard, video, mouse) system
  - PDUs (power distribution units)

Linux Platforms

NetServer

Pentium II and III

VMEbus SBC

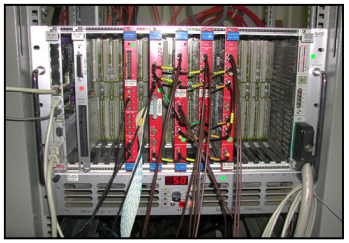

Dual Xeon

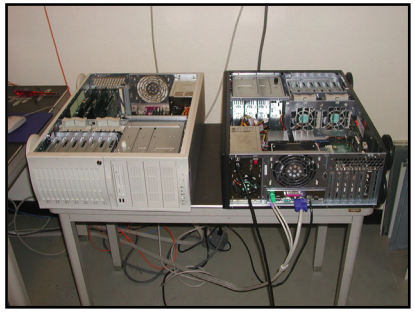

Quad Xeon

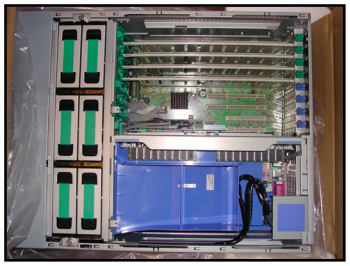

AthlonMP

Opteron

Itanium

32-bit Machines

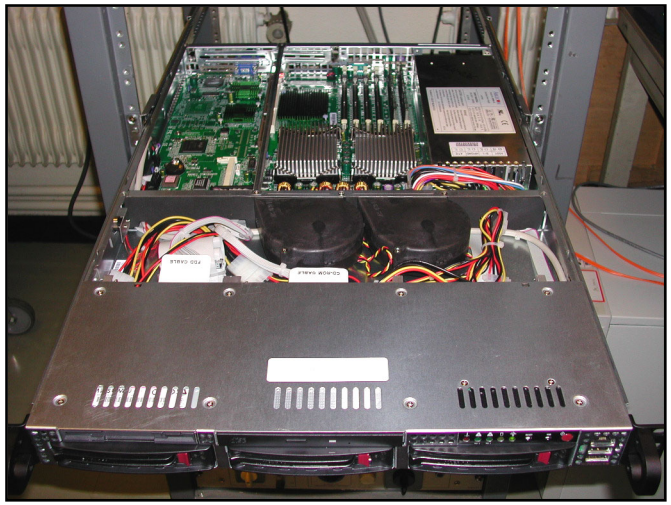

Supermicro X5DPR-iG2 [Broadberry]

CPU: dual Xeon 2.4 GHz, 512 KB
Chipset: Intel E7500, 400 MHz FSB
Memory: 1 GB, 266 MHz DDR SDRAM
Disks: 80 GB IDE
Gb Ethernet: 82546EB, RJ45 ports
Graphics: ATI Rage XL, 8 MB
Slots: 1x PCI-X 64/66, 1x PCI-X 64/133

Linux:
Red Hat 7.3 ➽

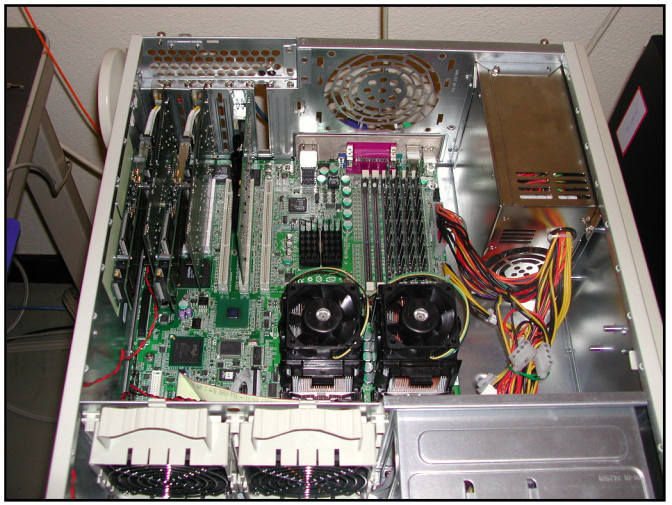

Intel Server Board SE7501HG2 [Elonex]

CPU: dual Xeon 2.66 GHz, 512 KB
Chipset: Intel E7501, 533 MHz FSB
Memory: 2 GB, 266 MHz DDR SDRAM
Disks: 120 GB IDE
Gb Ethernet: 82546EB, RJ45 ports
Graphics: ATI Rage XL, 8 MB
Slots: 2x PCI-X 64/100, 1x PCI-X 64/133, 3x PCI 32/33

Most likely used as LDC and GDC

64-bit Machines

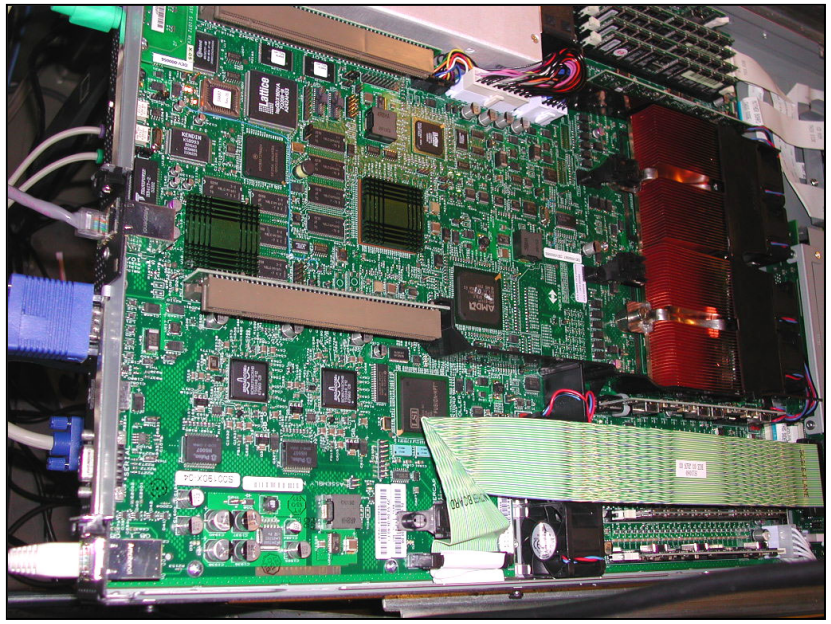

2-way 1U server [Newisys 2100]

CPU: dual Opteron 244 (1.8 GHz), 1MB L2 cache
Chipset: AMD 8131 PCI-X bridge, AMD 8111 Southbridge
Memory: 8 GB, 266 MHz DDR SDRAM (in banks)
Disks: 36.7 GB ULTRA320 SCSI disks
Gb Ethernet: 2x BCM5703X, RJ45 ports
Graphics: Trident Blade 3D, 8 MB
PCI-X slots: 1x 64/100, 1x 64/133
Periphery: PS/2, USB, CD-ROM, floppy, etc.
Service Processor: web server for administration

32-bit Linux: Red Hat 7.3 ➳
64-bit Linux: UnitedLinux 1.0 ➳
CERN Enterprise Linux 3 ➳

x86-64 architecture:
- seamless 32-bit and 64-bit mode
- tools gcc 2.96 (32-bit), gcc 3.2.1 (32- and 64-bit)
- DATE runs under Red Hat Linux 7.3
- ROOT and mySQL runs under United Linux 1.0

Presentation: CERN Computing Seminar 3rd Sep 2003

VMEbus Computer

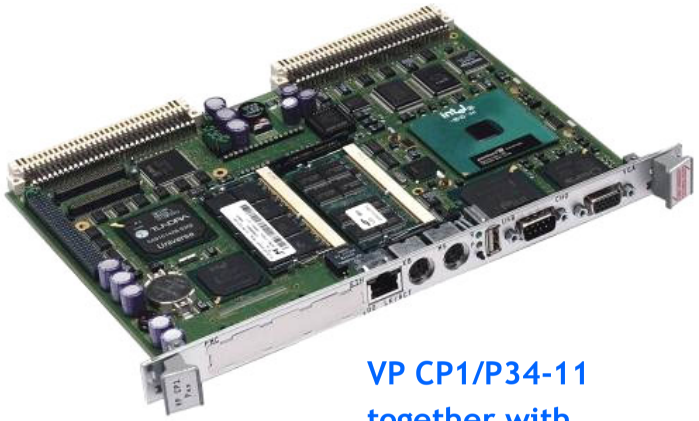

VP CP1/P34-11
together with
DS MSS/IFP-14

Concurrent Technologies Plc. (CCT):
One-slot VME64 single board computers
6U form factor, PC-compatible

CPU: Pentium III @ 850MHz
Chipset: 440BX/ZX/DX
Memory: 256 MB, 100 MHz ECC
Fast Ethernet: 82559ER, RJ45 port
Graphics: CT69030, 4 MB
Expansion: PMC slot, no PCI slot
Storage module: 10 GB disk, floppy drive
Linux: Red Hat 6.2 and 7.2 and 7.3 ➳

VMEbus access A16/24/32, D8/16/32, BLT:
- Tundra Universe II chip
- linuxvme-1.10.01driver
- single-cycle rate: 3.5 MB/s read, 14 MB/s write
- DMA64 rate: 43 MB/s read, 44 MB/s write

Reports: ALICE-INT-2001-034, ALICE-INT-2000-019

Benchmarking CPUs

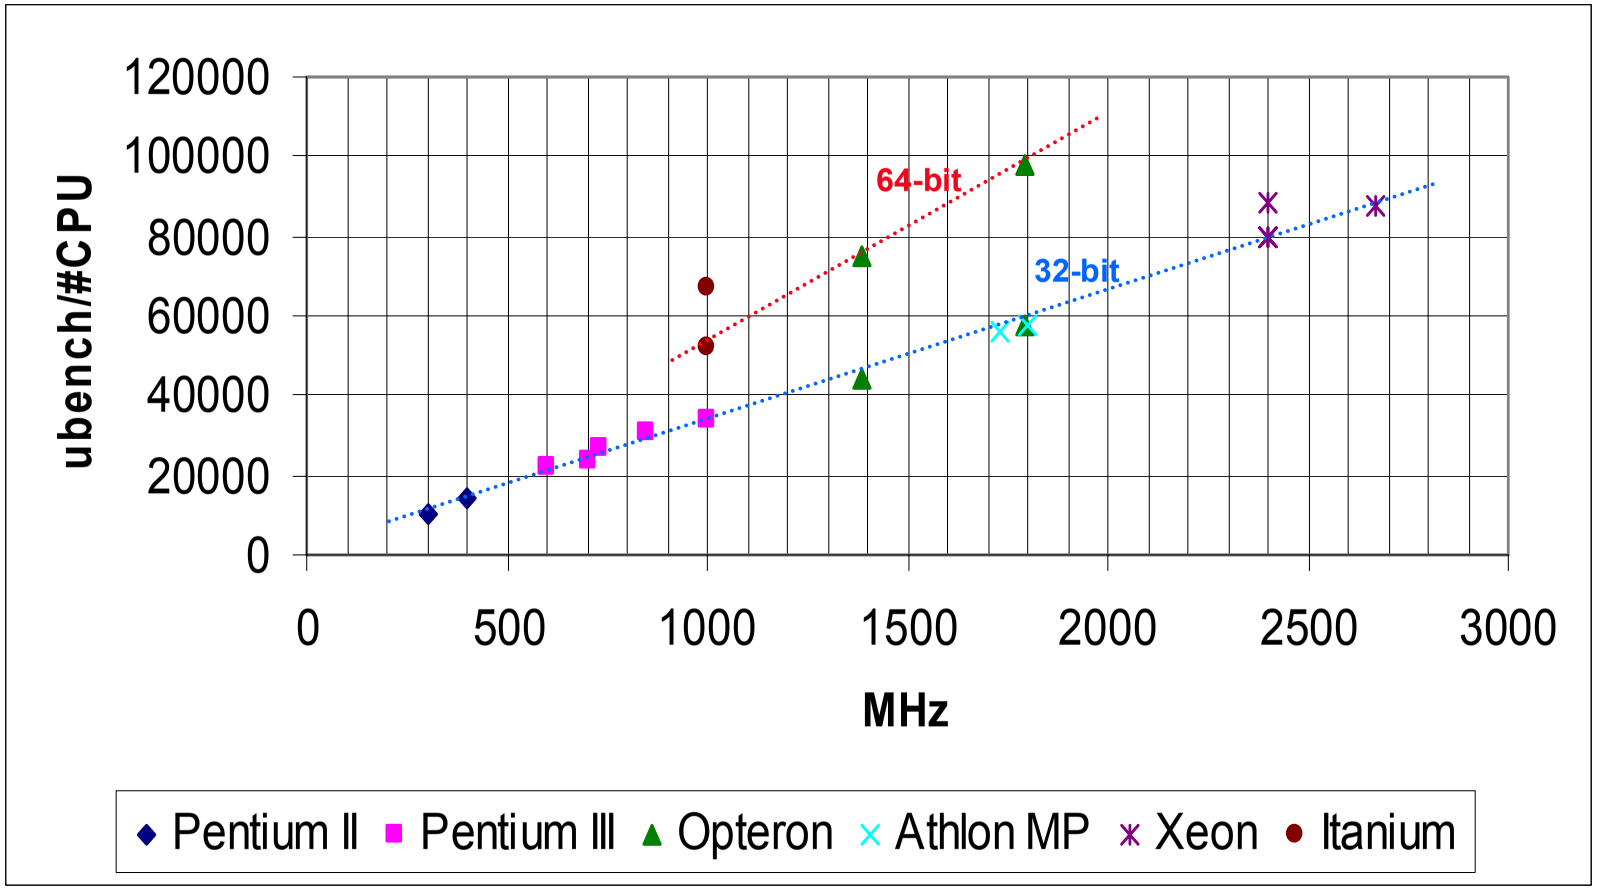

Benchmarking Memory

1x Stream:

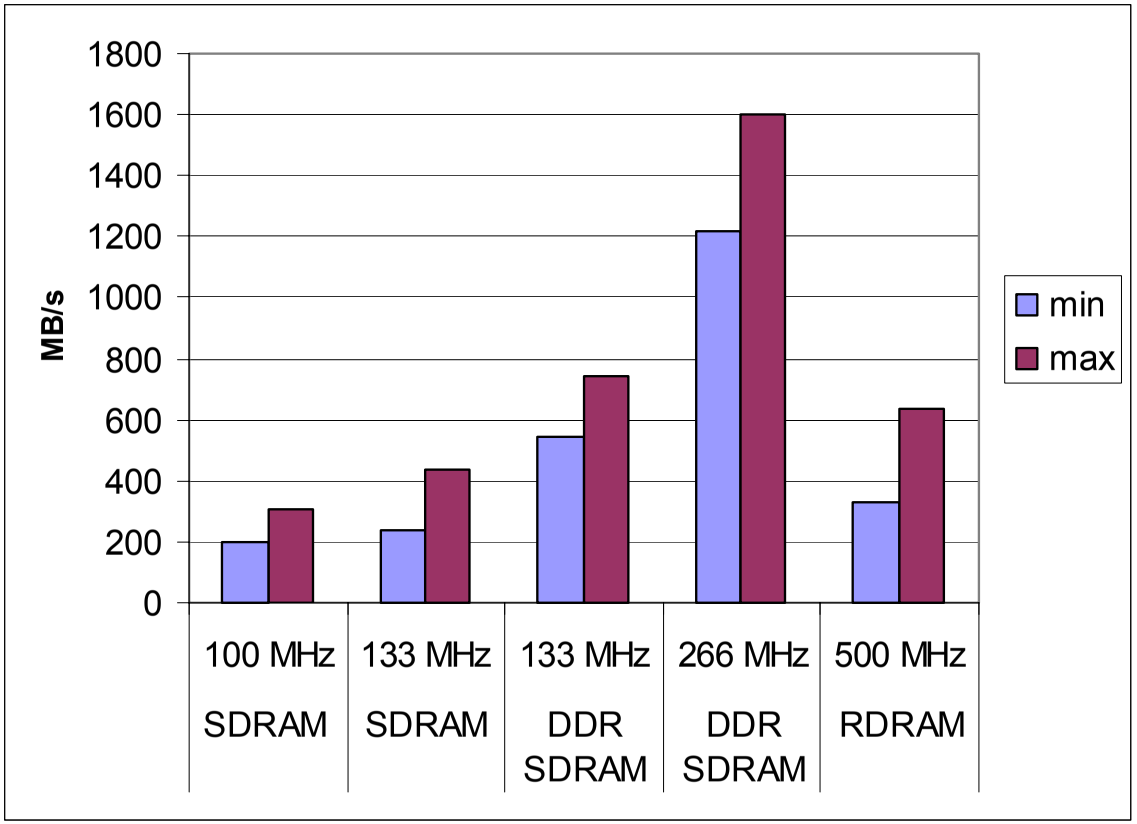

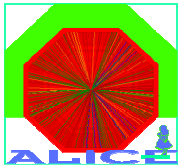

Network Infrastructure

- Network technologies
  - Fast/Gigabit/10Gigabit Ethernet for event-building aggregate bandwidth: 2.5 GB/s
  - Fiber Channel (FC) for storage network aggregate bandwidth: 1.25 GB/s
- System issues
  - Network controller
  - Linux kernel + driver
  - Network switch
- Evaluation
  - Transfer rate (message size, TCP socket size)
  - CPU load (% CPU per MB/s)
  - Long-term stability

Benchmarking Ethernet

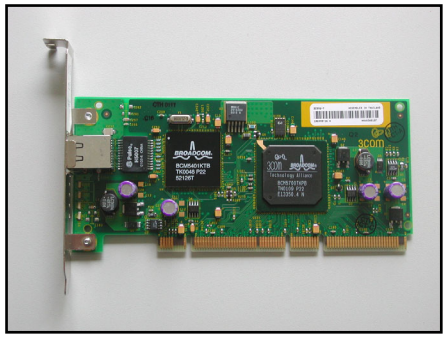

3C996 (BCM5700 controller)

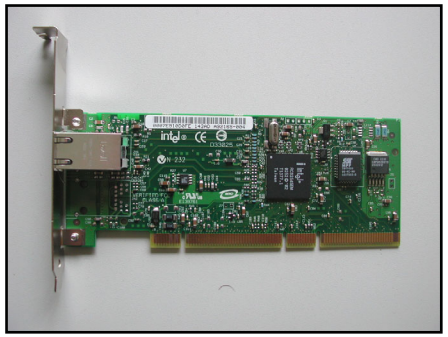

Pro/1000 MT server (82545EM controller)

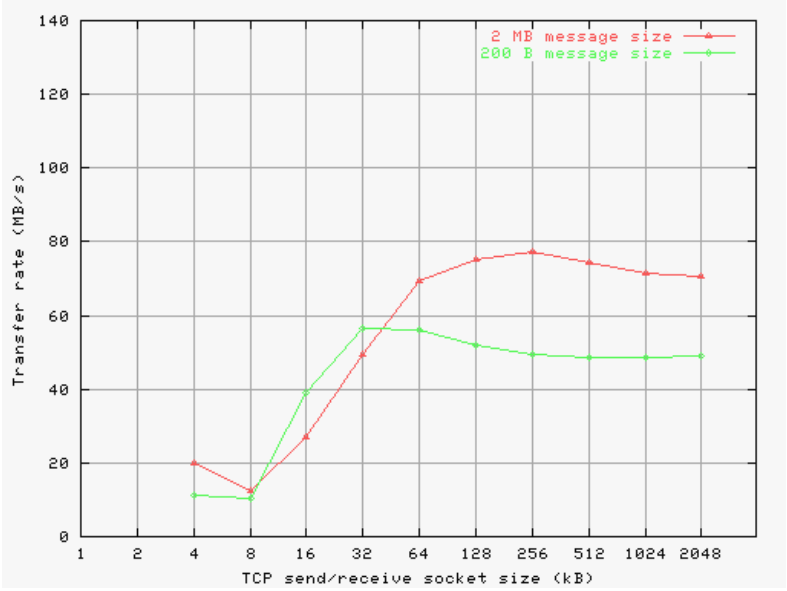

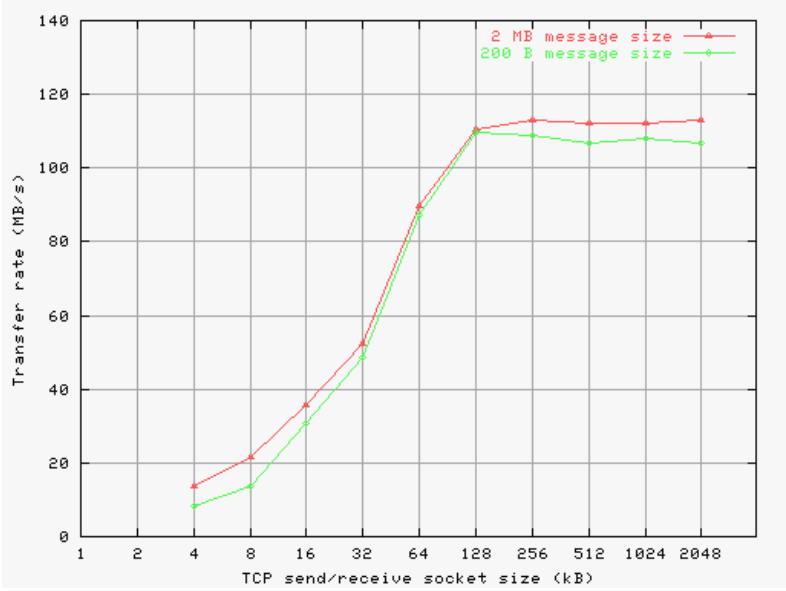

Readout Performance

Test Setup

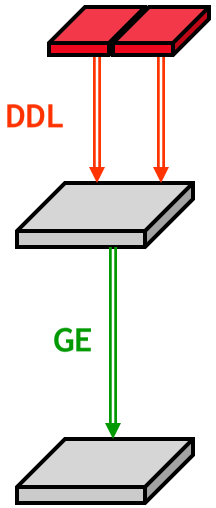

2x FEE:
- Front-End Emulator
- 100 MB/s

1x LDC:
- Intel Server Board SE7501HG2
- 2 pRORC + 2 DDL + 2 FEE
- 100KB pages, 250MB physmem
- Red Hat Linux 7.3

1x GDC:
- Supermicro X5DPE-G2
- 64MB IPC for event building
- Red Hat Linux 7.3

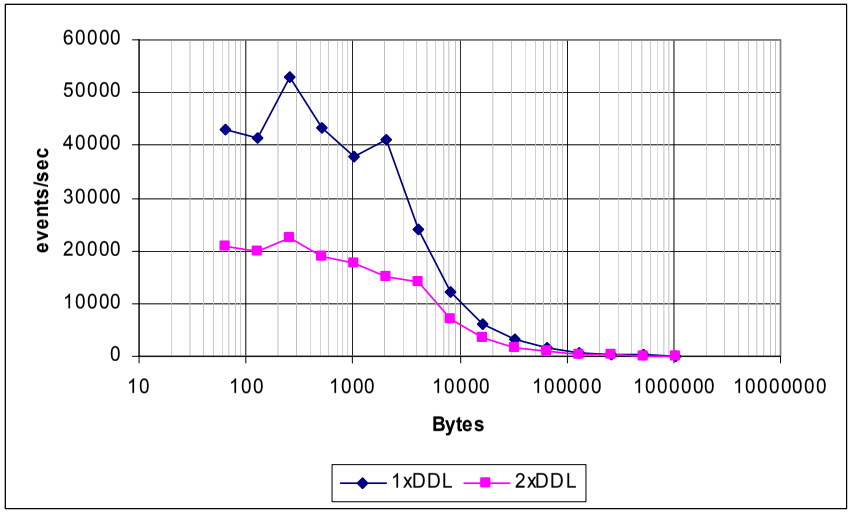

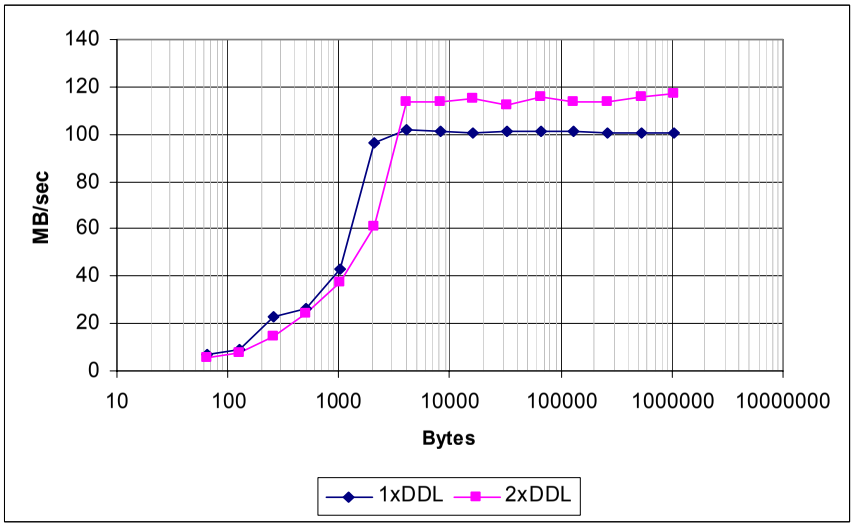

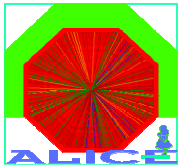

Transient Data Storage

- **Connection technologies**
  - DAS (Direct Attached Storage): scaling problems
  - NAS (Network Attached Storage): not conclusive
  - SAN (Storage Attached Network): good results with FC
- **Disk Arrays**
  - Product from DotHill: 12x 36.7 GB based on FC disks
  - Product from Infortrend: 12x 300 GB based on IDE disks
  - Going to test products based on SATA technology
  - Dual 2 GBit/s ports, rack-mount
  - Configuration: RAID level, file system (ext2, xfs), etc.
- **Fiber Channel Fabric**
  - Host Bus Adaptor (HBA): products from QLogic
  - Switch: product from Brocade with 15 ports

Benchmarking Disk Arrays

Infortrend IFT-6330-12F2D

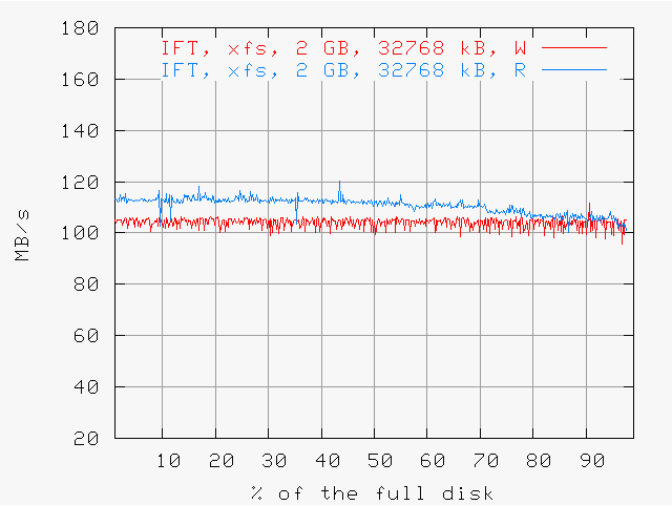

DotHill SANnetII 200

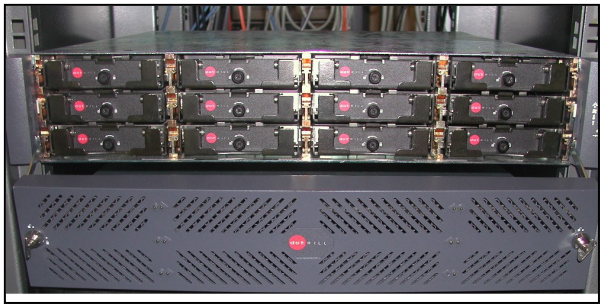

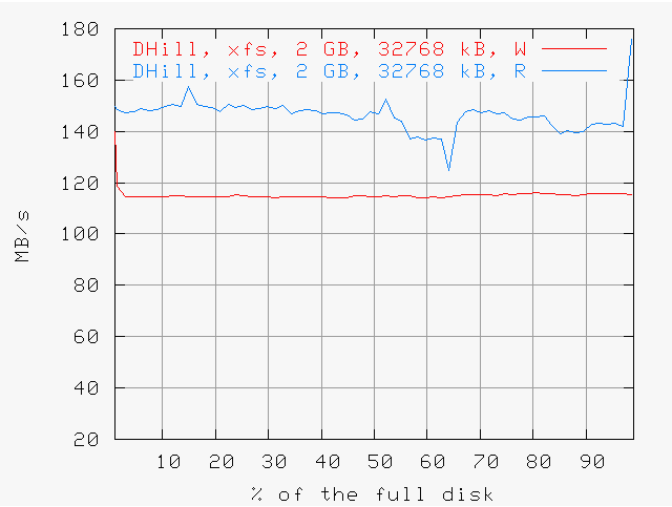

KVM System

Raritan Paragon UMT2161

2 users, 16 ports

daisy chaining …

2nd user (over IP)

1st user

Cat5 cables

Computer Interface Module (CIM)

Keyboard
Video
Mouse

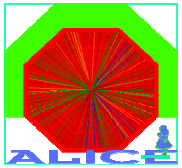

Résumé

- ALICE DAQ is based on standard technology
  - Technology is driven by the market
  - Performance increases by almost constant price
  - Choosing from different vendors
  - Fast evolution cycles need to be watched
- Thorough evaluation of computing elements
  - Doing functional tests and benchmarking
  - Working close together with the manufacturer
  - Ongoing task with documentation
- Purchasing
  - Need to follow the CERN regulations
  - Installation happens in phases
- Potentials
  - Integration, management, surveys