KLOE Computing Update

Paolo Santangelo

INFN LNF

KLOE General Meeting

University of Rome 2, Tor Vergata

2002, December 19-20

index

1) DAQ machines
2) automated tape space (2003)
3) useful data for sizing computing hw
4) disk space for DST (2002)
5) processing power (2003)

2002 – 3.6 kHz DAQ – 1.6 kHz T3

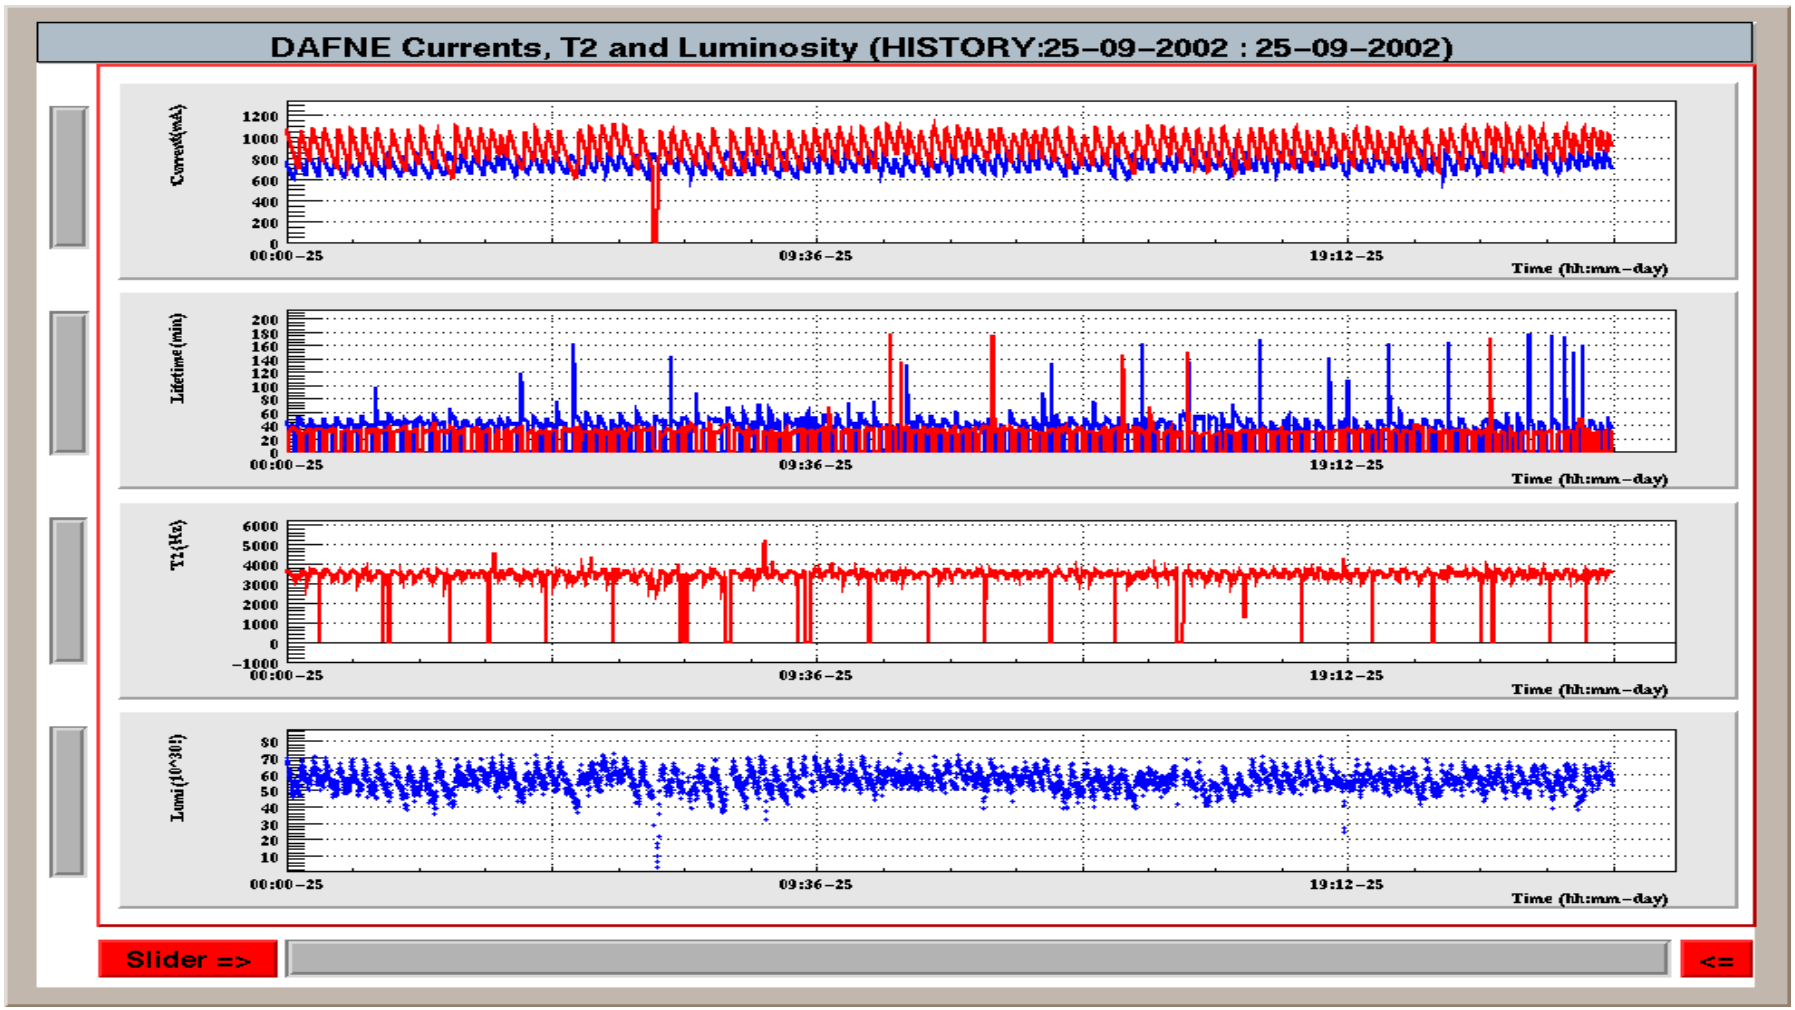

$L_{peak} \sim 7 \;\; 10^{31}\ cm^{-2}\ s^{-1}$ $\quad <L> \sim 5.4 \;\; 10^{31}\ cm^{-2}\ s^{-1}$ $\quad L_{int\ max} = 4.8\ pb^{-1}$ / day

on-line farm

- measured rates with 3 (4 way SMP) machines
  - up to 5 kHz of DAQ input
  - 2002 run conditions:
    - L3 filter cuts 3.6 kHz DAQ input to 1.6 kHz output
- Unix *fixed priorities* for *DAQ processes*
  - guarantee a smooth DAQ acquisition
- DAQ processes not limited by processor speed

on-line farm computers

1 run control
3 data acquisition
2 idle:
previously data servers
to be be reassigned soon
1 online calibration
1 data quality control
1 database server (DB2)

500 SpecInt95

caption: IBM F50 (4 way 166 MHz PowerPC)
IBM H50 (4 way 332 MHz PowerPC)

KLOE archived Data - October 2002

~~1999~~		~~raw~~	~~6 TB~~	GONE
2000	~20 pb^{-1}	raw	22 TB	x
		reconstructed	12 TB	x
2001	~180 pb^{-1}	raw	48 TB	s
		reconstructed	37 TB	s
2002	~288 pb^{-1}	raw	35 TB	
		reconstructed	29 TB	
total			183 TB	

tape library capacity is 200 TB + compression

upgrade of drives will increase capacity to 300 TB (ordered)

210 additional cartridges (ordered)

tape storage – 2003 technology

- KLOE has used a full generation of drive/media
  from 10 → 60 GB per cartridge

- what will be available next ?
  a new generation of drives and media (3Q 2003)
  higher track density (300 GB to 1 TB per cartridge)
- expected costs for the new generation ?
  cheaper tape drives
  more expensive cartridges
  total cost similar (in numbers of automated cartridges)

tape storage – 2003 upgrades

- capacity upgrade of present KLOE tape library
  - year 2002: 200 TB (full)
  - year 2003: 300 TB (not sufficient for 2003 data taking)
- more tape space is needed for 2003
  - next generation drives compatible with present library
  - no room for next generation cartridges in present library
  - upgrade of present library impractical

- conclusion
  - KLOE needs an additional tape library (2Q/3Q 2003)

KLOE data taking conditions and CPUs for data processing

year	trigger rate, Hz	ϕ + Bhabha Rate, Hz	data taking DAQ hours/pb^{-1}	data recon. hours*CPU/pb^{-1}	total Gb/pb^{-1}
2000	2100	77	33	970	1500
2001	2000	220	11	520	470
2002	1600	375	6.8	230	210
2003	2150	920	2.7	190	145
200x	5800	4600	0.6	167	115

extrapolated assuming 2002 background and trigger conditions

nominal processing power for concurrent reconstruction (in units of B80 CPUs) is 34, 70 and 300 CPU units for years 2002, 2003 and 200x respectively

these numbers do not include the sources of inefficiencies, MC production and concurrent reprocessing

data volume for data and MC samples

1 fb^{-1}	raw data	reconstructed	DSTs	MC files	MC DSTs
	115 TB	90 TB	10 TB	83 TB	20 TB

CPU power for data processing and MC generation

1 fb^{-1}	reprocessing from raw data	reprocessing from streamed data	MC ϕ decay
day CPU	**9600**	kpm 1440 ksl 1142 rad 198 bha 1440 **4220**	simulation 6650 reconstruction 5100 **11750**

these numbers do not include the sources of inefficiencies

- using 2002 background and trigger conditions
- all numbers refer to a sample of 1 fb^{-1}
- day CPU number are in units of B80 CPUs

disk space usage – 2002 hw

- DAQ **(1.5 TB)**
  5 strings - 300 GB each - RAID 1
  can buffer 8 hours of DAQ data at 50 MB/s
- 2 disk and tape servers **(3.5 TB)**
  12 strings - 300 TB each - RAID 1
  1+1 - reconstruction output
  5+5 - data staging for reprocessings or analyses
- AFS **(2.0 TB)**
  several RAID 5 strings
  user volumes and analysis group volumes

- total: **7.0 TB**

disk space - technology

- all KLOE disks are Directly Attached Storage (DAS)

  IBM SSA 160 MB/s technology

  flexible and yet simple architecture

- upgrades using DAS seem expensive

  check "*INFN Gruppo I*" price list (everything is cheap)

  other technologies must be selected (NAS, SAN, …)

- Network Attached Storage (NAS)

  cheap but usually too distributed

disk space - technology

- Storage Area Network (SAN)
  cost is intermediate between NAS and DAS
  storage partitioned to servers according to requirements

- present SAN technology is Fiber Channel (FC)
  FC technology is 10 years old and available only today

- future tape drives will be *FC attached*
  expected number of FC attachments is about 40
  this will require redundant SAN switches
  (overall balance increases however …)

disk space – 2002 / 2003 upgrades

- DST model requires fast access to data
  disk space sized according to DST sample
- estimated KLOE disk space for DSTs
  0.5 fb^{-1} → 5 TB
  1.0 fb^{-1} → 10 TB
- SAN Fiber Channel technology
  2002 / 2003 disk acquisitions
  at 10-15 keuro/TB
  initial investment
  later returns (?)

- 4 TB SAN FC disk space (ordered)

data reconstruction for 2002 data taking

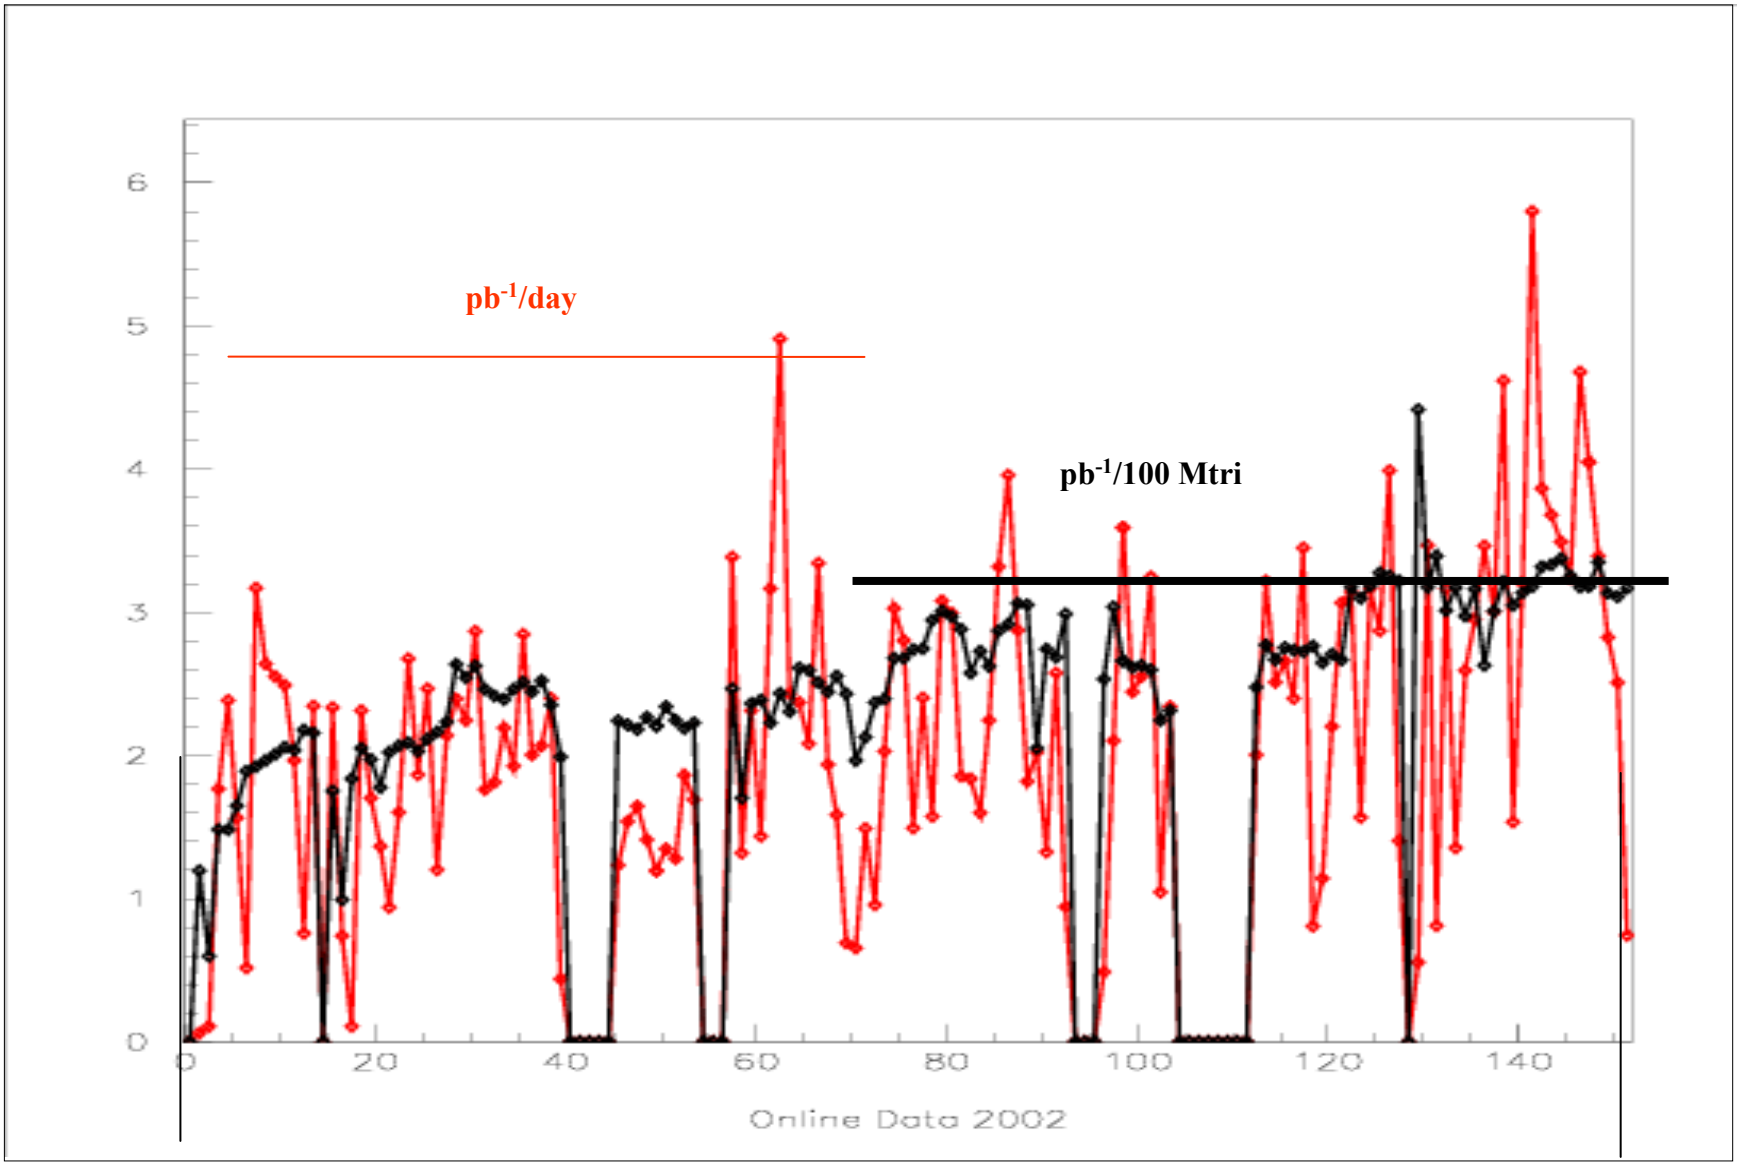

May, 3rd

Sep, 30th

trigger composition and reconstruction timings

	ϕ + Bha	background		
		filtered	tracked	total
triggers	4%	74%	26%	96%
reconstruction time	63 ms	1 ms	51 ms	14 ms
	16%	4%	80%	84%

year 2000
physics is a tiny fraction
computing used mainly for
tracking background events

triggers	11%	67%	33%	89%
reconstruction time	63 ms	1 ms	50 ms	17 ms
	31%	3%	66%	69%

year 2001
DAΦNE gives more physics

triggers	23%	78%	22%	77%
reconstruction time	63 ms	1 ms	33 ms	8 ms
	70%	3%	27%	30%

year 2002
physics is now 23 %
also computing is now used
for useful physics

data processing – 2002 hw

2 AFS clients (analysis)
8 montecarlo

700 SpecInt95
40 processors
0.8 kHz nominal reconstruction rate

4 AFS clients (analysis)
28 data processing

4900 SpecFp95
96 processors
4.5 kHz nominal reconstruction rate

caption: Sun Enterprise 450 (4 way 400 MHz Ultra Sparc 2)
IBM B80 (4 way 375 MHz Power3 II)

data processing – 2003 upgrades

- 2003 data taking

  luminosity 10^{32} cm^{-2} s^{-1}

  with 2002 background and trigger conditions

- doubles required CPUs for online production

  in units of present KLOE B80 CPUs

- processor performance

  evaluated on the basis KLOE specific benchmarks

  SPEC metrics, almost meaningless

conclusions

- KLOE computing ok for 2002 data taking
- uptime only constrained by external events
- 2003 data taking will require upgrades
  + 1 tape library (+1.5 PB)
  + 10 TB disk space
  + 80 CPUs (in units of B80 CPUs)

Backup Slides

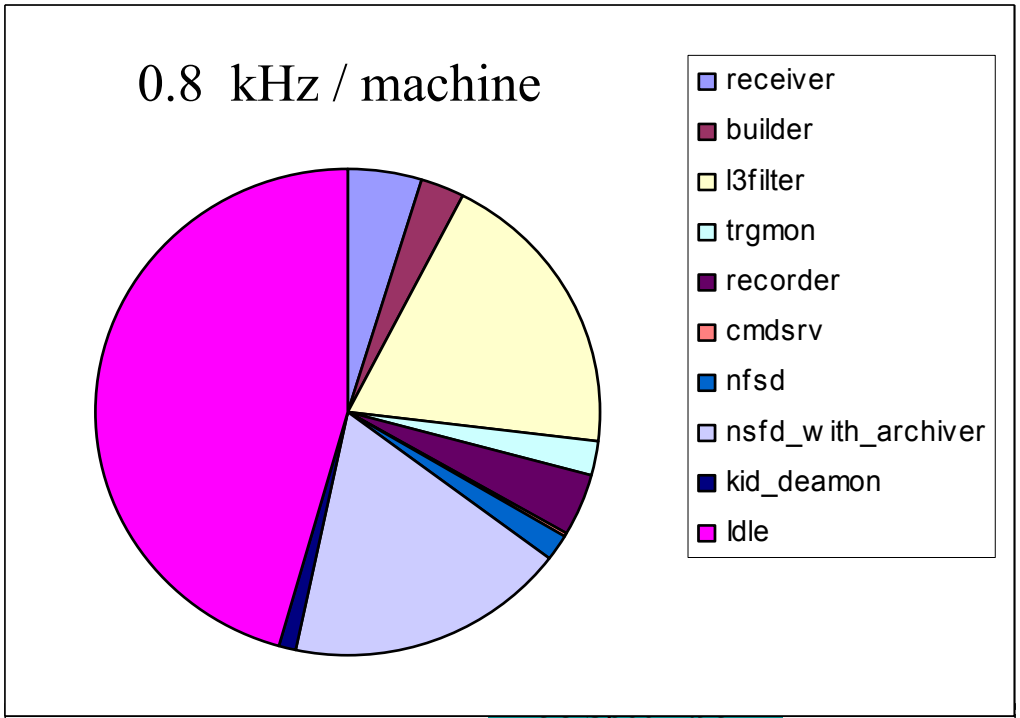

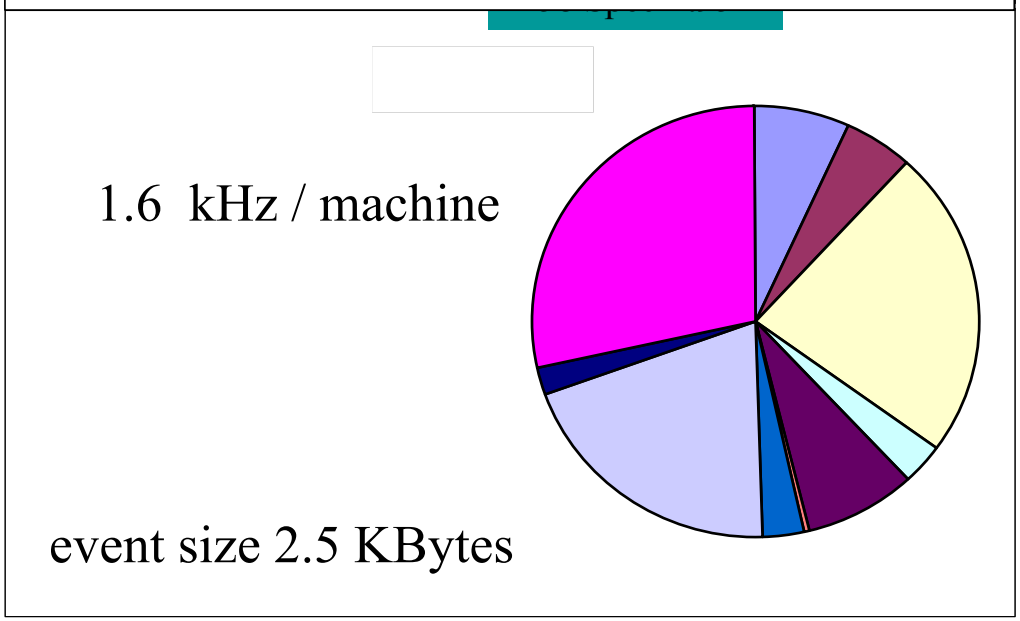

low rate
2.4 kHz DAQ input
3 computers

each computer
4 way SMP

data moving
simultaneous
with smooth DAQ

processes are
compatible
with processors

high rate
4.8 kHz DAQ input
3 computers

offline computing resources

IBM 7026-B80
4-way 375 MHz

MC & DST PRODUCTION
32 Sun CPU's

Sun Enterprise 450
4-way 400 MHz

RECONSTRUCTION
84 IBM CPU's

Tape/disk servers

ANALYSIS
8 Sun + 8 IBM CPU's

Local online disks: 1.5 TB
Data acquisition
Calibration work

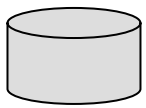

afs cell: 2.0 TB
User areas
Analysis/working groups

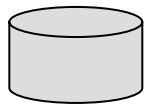

Managed disk space: 3.0 TB
1.2 TB: Input/output staging
Reconstruction
MC/DST production
1.4 TB: Cache for data on tape
DST's maintained on disk

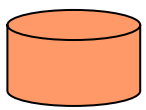

Tape library: 220 TB
5500 40GB slots
12 Magstar drives
14 MB/s each
to be upgraded to 60GB/cartridge

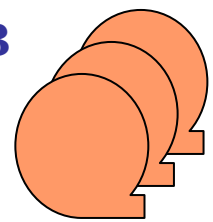

tape storage – 2002 hw

- tape library

  15 (+2) box long IBM 3494 tape library

  5,500 cartridge slots (5,100 cartridges)

  high-availability:

  - dual active accessors
  - dual library control (standby takeover)

- 12 tape drives

  14 MB/s IBM Magstar (linear, high reliability)

  presently 40 GB per cartridge (uncompressed)

  drive upgrade to 60 GB per cartridge (ordered)

- safe operations

  some cartridges mounted up to 10,000 times

reconstruction – year 2002

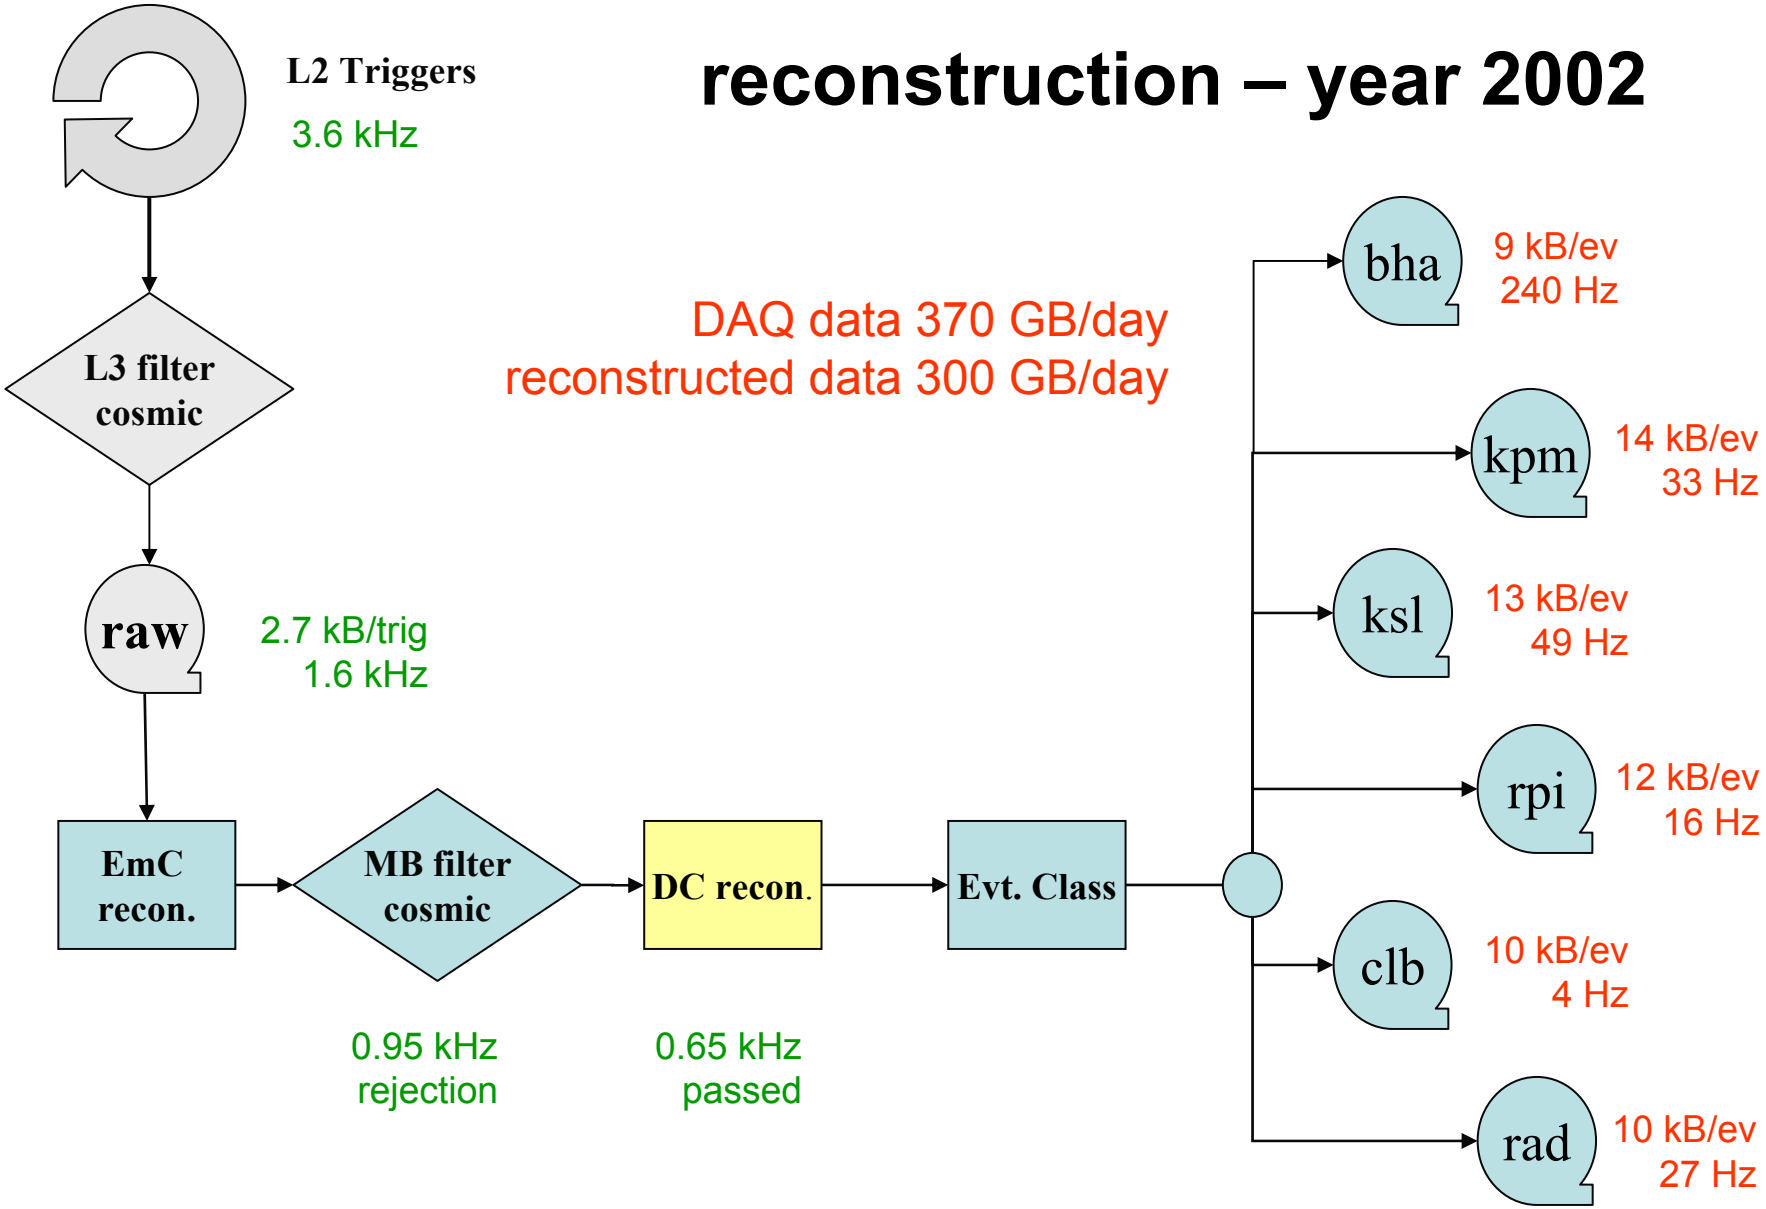

Processor Comparison for KLOE Tasks

	IBM Power 3 375 MHz	IBM Power 4 1 GHz	Sun ES 450 450 MHz	PentiumIII 1 GHz	Athlon XP2000+
ms/trigger					
data, full reconstruction	17	8	40	29	18

ms/event					
data, tracking only	24	12	66	53	32
MC-1 ϕ production	210	110	420	270	150
MC-2 $Ks \rightarrow \pi^+ \pi^-$	120	60	240	160	90
MC-1 reconstruction	70	35	170	130	76
MC-2 reconstruction	120	60	280	210	123