


Offline resource allocation




M. Moulson, 11 February 2003

Discussion

Outline:

- *Currently mounted disk space*
- *Allocation of new disk space*
- *CPU resources*






DST space requirements

DST	Luminosity processed (pb ⁻¹)	Current output size (GB)	Event size (KB)	Events per pb ⁻¹ (K)	Output size for 460 pb ⁻¹ (GB)	Output size for 1 fb ⁻¹ (GB)
dkc	34*	107	4.80	680	1500	3260
dk0	460	1371	3.22	971	1440	3130
d3p	80	67	3.12	280	80†	170†
drc	460	657	2.78	538	690	1500
drn	460	130	2.24	132	140	300
Total		2332			3850	8360

Notes: * *Includes some duplicate DST's at different DBV levels*

† *Assuming DST's produced for 20% of runs*



MC DST space requirements


Active production requests:

Type	Events requested	Event size (KB)	Total size (GB)
$K_S \rightarrow \text{all}, K_L \rightarrow \text{all}$	400 M (400 pb ⁻¹)	4.1	1640
$\phi \rightarrow \text{all}$	300 M (100 pb ⁻¹)	3.4*	1020

Notes: * Scaled from $K_S K_L$ assuming .dst size proportional to .mcr size

Need for space could substantially increase:


- 2001-2002 MC needs not satisfied by active requests
 - Impossible to estimate future needs (2003 and beyond)
-



Currently mounted disk space

Label	Capacity (GB)	Group	Description
/recalled	540	DST	DST latency
/recalled1	540	PROD	Input staging for datarec
/recalled2	180	DSTPROD	Input staging for DST production
/recalled3	540	DST	DST latency
/recalled4	540	USER	General user recall area
/datarec	540		Output staging for datarec/DST prod.
/analysis	180		Log files, histograms for datarec
/data/farm	1340		Data acquisition

Recover 540 GB for DST's by reallocating /datarec1, but only temporarily



Currently available AFS disk space

Area	Capacity (GB)	Used (%)
user	180	
cpwrk	195	87
ecl	50	49
emc	130	85
kaon	170	89
kwrk	150	42
mc	90	6
phidec	300	77
recwrk	30	71
trg	100	71
trk	90	61

Total AFS capacity: **1745 GB**

kloeafs1 $3 \times 240 + 105$ GB

kloeafs2 $3 \times 240 + 190$ GB

Free space: 146 GB (8%)

Redistribution:

Not much to be gained

Usage patterns:

Files tend to accumulate

Capacity of users' PC's saturated?

AFS suitable for analysis of Ntuples?


PAW/HBOOK very slow if

Ntuple not entirely cached

Mainly a problem for Ntuple-to-Ntuple analysis

Need some NFS user/group disk?

Allocation of new disk space



In a nutshell:

Immediate needs for DST's: 3.9 TB (460 pb⁻¹)

Immediate needs for MC DST's: 2.7 TB (700 Mevts currently planned)


Not necessary to have 100% in cache, but having DST's in cache greatly accelerates analysis work

Current space for DST latency: 1.1 TB → 1.6 TB if PROD reallocated

4 TB disk space newly acquired, mounted on fibm0a/0b (SAN/FC)

Options for use:

- If all allocated to DST's: 75% of DST's in disk cache
 - Expand AFS user/group storage space
Move current recalled disks to AFS in 240 GB chunks
 - Create some NFS user/group storage space
-



Reconstruction and CPU consumption

year	trigger rate, Hz	Luminosity (ub ⁻¹ /sec)	Integrated luminosity (pb ⁻¹)	φ + Bhabha Rate (Hz)	Processing time CPU hours/pb ⁻¹	Processing time on 76 CPU's (days)
2000	2100	9	24	77	970	5
2001	2000	24	170	220	520	48
2002	1600	42	290	375	230	36
2003	2150	100	540 (62.5 d)	920	190	56

200x	5800	510		4600	167
------	------	-----	--	------	-----

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