

CD-Doc xxxx

Run II, Tapes and 2008 Onward

Rob Kutschke, CD/IDS
April 29, 2008

Outline

- 1) Density Migration.
- 2) Recent Tape Usage.
- 3) Data Protection Policy

Notation

- Per Gene's request:
 - 1 kB = 1024 Bytes
 - 1 MB = 1024 kB
 - 1 GB = 1024 MB
 - 1 TB = 1024 GB
 - 1 PB = 1024 TB
- Migration vs Duplication:
 - Duplication:
 - At end of process two catalog entries are present.
 - Migration:
 - At end of the process only one catalog entry is present.
 - Current D0 LT02-LT04 is using duplication.

Migration Driven by Lists of Files

- Example: 4x density increase for 8 input volumes.
 - Default behaviour creates 2 to 4 output volumes.
 - Input volumes: $i_0, i_1, i_2, i_3, \dots, i_7$; Output volumes: o_1, o_2
 - List 1: Good files from (i_0, i_1, i_2, i_3) to o_1 .
 - List 2: Good files from (i_4, i_5, i_6, i_7) to o_2 .
 - List 3: Obsolete files from (i_0, i_1, i_2, i_3) , if any, to o_3 .
 - List 4: Obsolete files from (i_4, i_5, i_6, i_7) , if any, to o_4 .
 - Both CDF and D0 say that they prefer not to migrate their deleted/obsolete files.
- **With existing code, must all files be in same file family?**

Untangling Mixed Volumes

- Example:
 - CDF wrote file families to separate sets of volumes.
 - Easy to replace a file family with a new version.
 - But: leaves unused space at the end of some volumes.
 - CDF recaptured this space so single volumes may now contain two (**several?**) file families.
- Rick says that it is OK if these mixed volumes are copied as is.
- QiZhong says that D0 does not have such volumes.

Ongoing D0 LT02-LT04 Migration

- Migration station:
 - CPU + local disk dedicated to migration.
 - \$ supplied by the lab.
- 1 or more input LT02 drives read in parallel.
 - One LT04 output drive writes files as they come in.
 - At completion rewind LT04.
 - Readback LT04.
 - Compute and compare checksums on all files.
 - **Input drives idle during readback.**
- For 9940A has about 13x density increase!
 - So output drive will not be very busy unless there are a lot of input drives.

Technology Parameters

	Cartridge		Speed (MB/s)	EOSL
	Size(GB)	Cost(\$)	Nominal	
9940A	60		10	6/30/2009
9940B	200		30	12/31/2010
LT02	200		30 (35?)	?/?/2012
LT03	400	\$38.+	40/80	> 5 years
LT04	800	\$96.+	60/120	> 5 years

+ per slot tax: \$47. + \$7./year maintenance.

Our infrastructure supports only 60 MB/s.

Miscellaneous Parameters

(Assumed technology independent)

6 s	Time to read past a file mark.
8 s	Time to write: EOF, EOF, backspace and start to write next file.
2 s	Time to write: EOF and start to write next file.
120 s	Time to access/mount/unmount/reshelve tape.
95%	Usable fraction of nominal capacity (guess based on CDF 9940B usage)
240 s	Rewind time (guess but not a time driver).

CDF Files in 9940B

Size (TB)	# Distinct Volumes	# Files	Data Type	Last Date
721.5	4,386	841,384	RECONSTRUCTED	2007-04
707.2	3,883	814,000	RAW	2007-10
144.2	1,123	151,268	SIMULATED_RECONSTRUCTED	2007-02
138.1	1,741	207,143	SIMULATION	2007-02
58.3	340	48,645	STANDARD_NTUPLE	2007-01
50.5	381	80,709	ROOT_NTUPLE	2006-08
31.0	220	38,960	STRIPPED_RECONSTRUCTED	2007-02
11.2	85	11,573	STANDARD_NTUPLE_MC	2007-01
8.5	52	8,517	TOP_NTUPLE	2006-11
5.3	48	4,988	TOP_NTUPLE_MC	2006-12
4.6	27	4,035	B_NTUPLE	2006-11
3.7	96	28,989	GENERATOR_LEVEL	2005-12
0.3	2	395	NEW_RECONSTRUCTED	2007-04
0.3	13	295	HBOOK_NTUPLE	2007-01
0.3	6	283	STRIPPED_SIMULATED	2007-01
1885.0	10,533	2,241,184	Sub total	
91.4	506	112,287	RAW - in D0 silo	2006-08
1976.4	11,039	2,353,471	Grand Total	

CDF 9940B to LT04 Migration

READ TIME:

799.5 Drive-days
 15.3 Drive-days
 163.4 Drive-days
 978.3 Drive-days

Time to read all data at max read speed.
 Time to fetch/mount/umount/replace tapes.
 Fixed per file overhead.

Total read time.

2811

Number of LT04 volumes needed

WRITE TIME (SLOW VERSION):

199.9 Drive-days
 217.9 Drive-days
 7.8 Drive-days
 199.9 Drive-days
 163.4 Drive-days
 3.9 Drive-days
 792.8 Drive-days

Time to write all data at maximum write speed.
 Time to write "EOF. EOF backspace" at end of file.
 Rewind time?
 Time to read all data at maximum read speed.
 Time to space past file marks on read.
 Time to mount/unmount the LT04 tape volumes
 Total: no contingency or efficiency.

WRITE TIME (FAST VERSION)

 54.5 Drive-days
 629.4 Drive-days

Time to write "EOF" at the end of each file.
 Total time in the variant model.

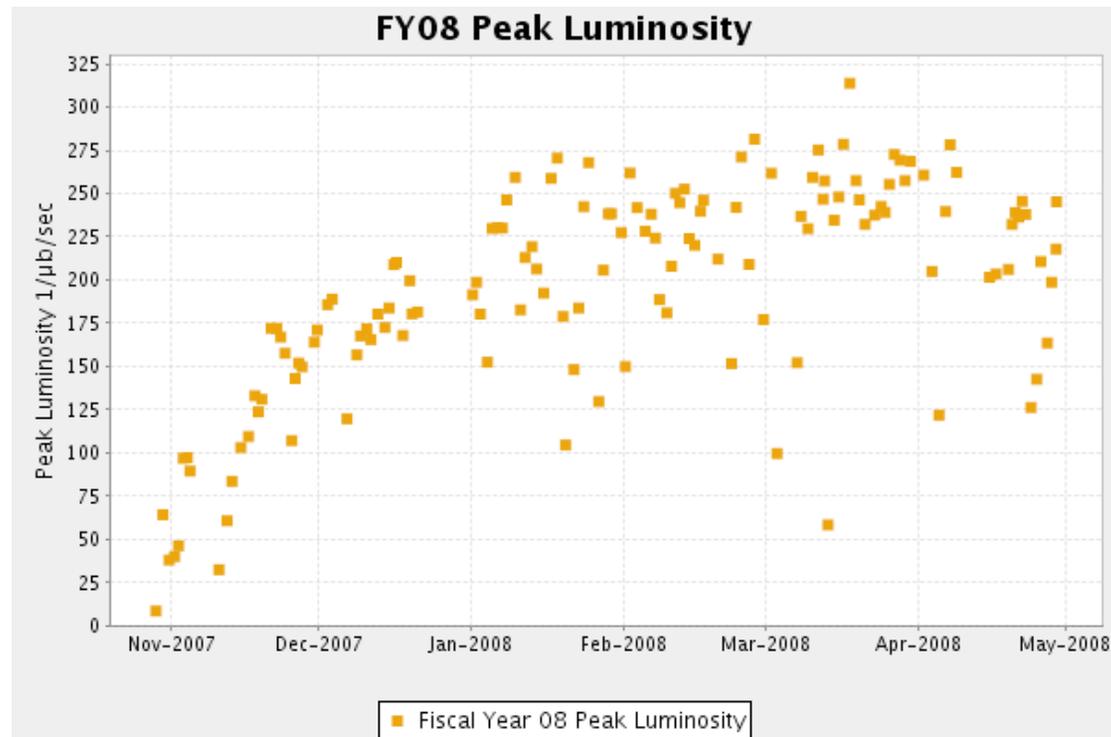
\$275,478
 \$132,117
 \$0
 \$0
 \$407,595

Cost of the LT04 media.
 Cost of the slots in the library
 Cost of additional 9940B movers?
 Cost of additional LT04 movers?
 Total cost.

Conclusions About Migration

- Fixed per-file time-costs are important:
 - Change (EOF, EOF, backspace) to EOF: gain 14%.
- There only really big winner is to perform
- If the disk buffer space is there, 1.25 reading LT02's can keep up with a LT04 writing!
 - Without EOF change.
- Add 25% contingency and the following will get the job done in 2 years:
 - 1.68 LT02 for reading
 - 1.36 LT04 for writing (without EOF change).
 - This implies enough disk buffer for input to get ahead of output during the readback for checksum.

FY08 Peak Lumi from AD



<http://www-bd.fnal.gov/pplot/today/DataSummaryTables.html>

Guess that Jan 1, 2008 onward is a good model for rest of FY: about 240.E30 initial lumi.

104/119 days so far have non-zero recorded events.

CDF Tape Usage in 2008 (All LT03)

	TB	Distinct Volumes	Files
Jan	88.9	272	87137
Feb	50.0	191	58278
Mar	113.6	334	105555
Apr	63.9	219	61628
Total	325.4	897	312598