

'Life without ESD'

(proposal from DP)

Ueda I.

2011.02.02.

Introduction

As we have been discussing since the last S&C week, the data volume will become huge with a probable increase in the trigger rate = ESD cannot be put on T1 disks

This proposal from the Data Preparation et al. is still a draft under discussion

- There have been a series of meetings to gather information if and how people can adapt their life without ESD
- The feedback from the perf/phys groups are reflected in the draft

Baseline:

- Most of ESD do not stay on disk for long term
- No ESD archived on tape

ESD on disk

- No long term storage of ESDs on disk for bulk physics streams
- “Bulk stream” : 2 replicas on disk for short period for det/trig/perf studies
 - ▶ Approximate 10% of the full-year data
 - ▶ ESD from T0 : for 1 – 2 months
 - ▶ ESD from repro : a specified period (corresponds to 10%)
 - ▶ replicate to T1s = possible DESD production at T1s
 - ▶ No PD2P allowed
- Specific small streams : 2 replicas for all data all the time
 - ▶ ~ 35 Hz (incl. express 15 Hz)
 - ▶ full set for Nth processing, only the ZeroBias and CosmicCalo for N-1
 - ▶ PD2P allowed
- No ESD copy on TAPE
 - ▶ reprocessing from ESD abandoned = always from RAW

RAW on disk

One replica of RAW on disk

- No automatic replication allowed

Aim:

- investigation of “exciting discovery” events as fast as possible (<24h)
 - ▶ only required for data taken within the last year
 - ▶ ~ 10 – 100 events to look at in detail distributed randomly
 - ▶ need a user interface to turn a RAW event into an ESD etc.
- Also for event displays or detailed debugging of specific events

To be considered

- The number of replicas could be revised
 - ▶ The latency and operational burden for recovering a loss should be estimated
- Possible reduction of the RAW event size by a factor of 2
 - ▶ “gzip” or better packing

AOD/DESD

AOD

- at least 6 primary replicas for the physics streams
- additional secondary copies can be created depending on user demand and disk space constraints
- no need to keep the express stream AOD

DESD

- at least 4 primary replicas
- additional secondary copies can be created depending on user demands and disk space constraints
- only needs to be kept for the most recent (Nth) processing
 - ▶ DESD is primarily used for performance studies
 - ▶ A possible exception when used for physics analysis

Rough Estimate

The data volumes from run 167607 in 2010 (400 Hz, 11h)

200 days of pp running with 30% LHC up-time

Rough Estimate for 2011

The data sizes from run 167607 in 2010 (400 Hz, 11h)

200 days of pp running with 30% LHC up-time

	streams	1 copy	#rep	N	N-1	total
RAW	bulk	2.7 PB ₍₁₎	1			2.7 PB
ESD	bulk	2.8 PB	2 ₍₂₎	0.8 PB	0.8 PB	1.6 PB
ESD	small	200 TB	2	0.8 PB	0.1 PB	0.9 PB
AOD	all ₍₃₎	400 TB	6	2.4 PB	2.4 PB	4.8 PB
DESD	all	380 TB	4	1.5 PB	-	1.5 PB
total				5.5 PB	3.3 PB	11.5 PB

(1) full size (without 1/2 reduction) (2) x 60d (3) all but express

NB. some numbers need to be understood

Alternative scenarios

- 10 AOD and DESD primary replicas for N : **12.7 PB**
- 30% reduction in ESD size : 7.9 PB
- 20% increase in the DESD and Express stream : 9.6 PB
- 2 DESD for (N-1) : 9.5 PB
- LHC efficiency 50% : **14.7 PB**
- increase of 20% of the AOD : 9.8 PB
- 3 AOD of the (N-1) : **7.6 PB**

- The current (2010) scenario : **27 PB**
 - ▶ 2 ESD, 10 AOD and DESD for both N and (N-1),
 - ▶ all streams apart from the express stream,
 - ▶ and no RAW data on disk

Other than 2011 pp

2010 data

- reduce replicas — 4 PB
- delete rel15 ESD with a 3-month notice — 3 PB
- no need for RAW on disk as long as ESD on disk — no increase
- consider deleting periods A-C (only 300/nb)
- and older data (2009 and cosmics)

Heavy Ion data

- to be accounted

MC

- not yet