



Tevatron Data Preservation: Part I

Bodhitha Jayatilaka, Ken Herner
Fermilab

DPHEP/DASPOS Workshop
CERN
22 March 2013




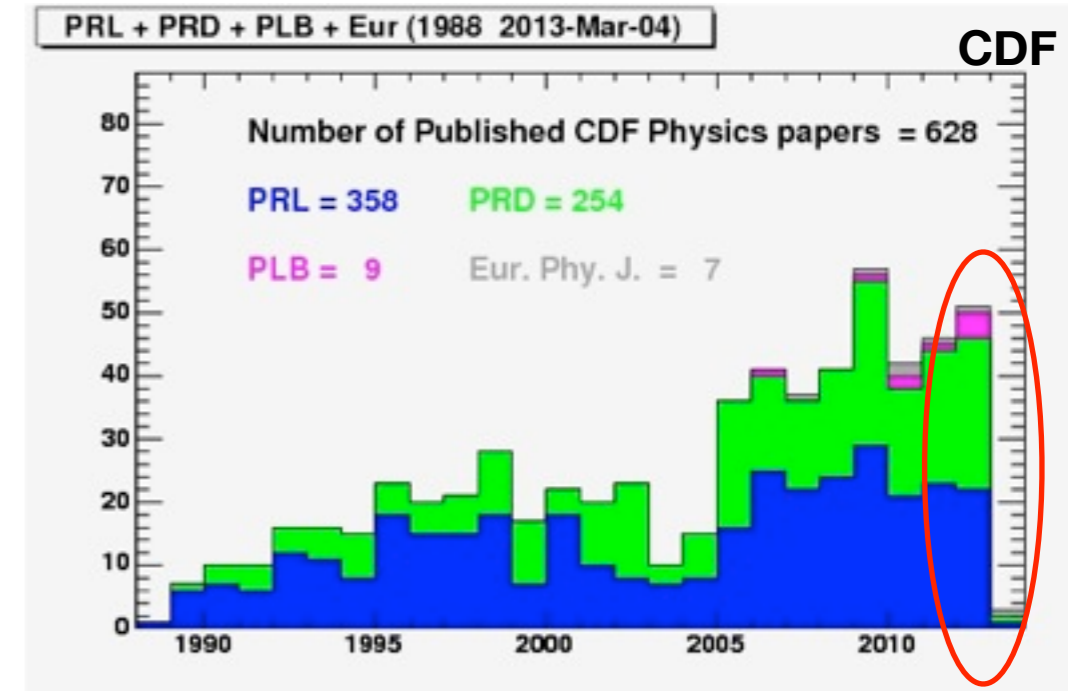
Outline

- **Part I** - this talk
 - Introduction - Tevatron
 - Fermilab Computing Sector data preservation project
 - Storage media and migration of Tevatron data
 - Data preservation progress at CDF
- **Part II** - K. Herner
 - Data preservation progress at DØ
 - Tevatron data access mechanisms
 - Long-term database access for Tevatron experiments

The Fermilab Tevatron



- 1.96 TeV proton-antiproton collider
- Ceased operations 30 Sep 2011
 - $\sim 12 \text{ fb}^{-1}$ ($\sim 10 \text{ fb}^{-1}$ recorded)/expt
- Unique data
 - Unique initial state vs LHC (e.g., asymmetry measurements, flavor physics)
 - Multiple energy collisions (300 GeV, 900 GeV, 1960 GeV)
 - Well-understood experiments: “legacy” precision measurements (e.g., M_W , m_t)
- Continued physics output
 - Expect long tail at both experiments
 - ~ 15 papers published+accepted in 2013 to-date





- Data preservation efforts at **both experiments** and Fermilab **Computing Sector (CS)**
 - Experiment efforts borne out of 2011-2012 **task forces for DP planning**
 - CDF offline operations now largely involved in data preservation efforts
 - CS efforts coordinated through **CS Data Preservation Project**
- CS Data Preservation Project funded for **two years** from US Dept. of Energy
 - Two dedicated physicists: Ken Herner (DØ), Bo Jayatilaka (CDF)
 - Technical project lead: Joe Boyd
 - Oversight in form of project manager (Rob Kennedy)
 - Steering committee: senior CS scientists and other stakeholders
 - Project sponsor is Rob Roser (Scientific Computing Div. head)

Fermilab Computing Sector DP Project



- Goal: maintain **full analysis capability** (level 4) through **2020** (end of support for Scientific Linux [SL]6)
 - Seek common solutions between experiments where possible
 - Efficient division of effort between experiments and CS
 - Experiments primarily responsible for code, documentation
 - CS primarily responsible for infrastructure (e.g. job submission, databases)
- Does **not** cover (as of now)
 - Accessibility of data by non-collaborators
 - Analysis capability beyond 2020
- Draft of project charter in circulation
 - Includes 6 month, 1 year, and 2 year goals

A draft project charter document for the Run II Data Preservation Project. The document is titled "Fermilab Computing Sector Project Charter" and "Run II Data Preservation Project". It includes the Fermilab logo, version information (Version 0.00, 2/18/2013), a DocID # field, and a signature line for the Project Sponsor (Rob Reser) with a date field.

Fermilab

Computing Sector
Project Charter

"Run II Data Preservation Project"

Version 0.00
2/18/2013

DocID # _____

PREPARED BY:
Robert D. Kennedy

CONCURRENCES:

Project Sponsor


Date



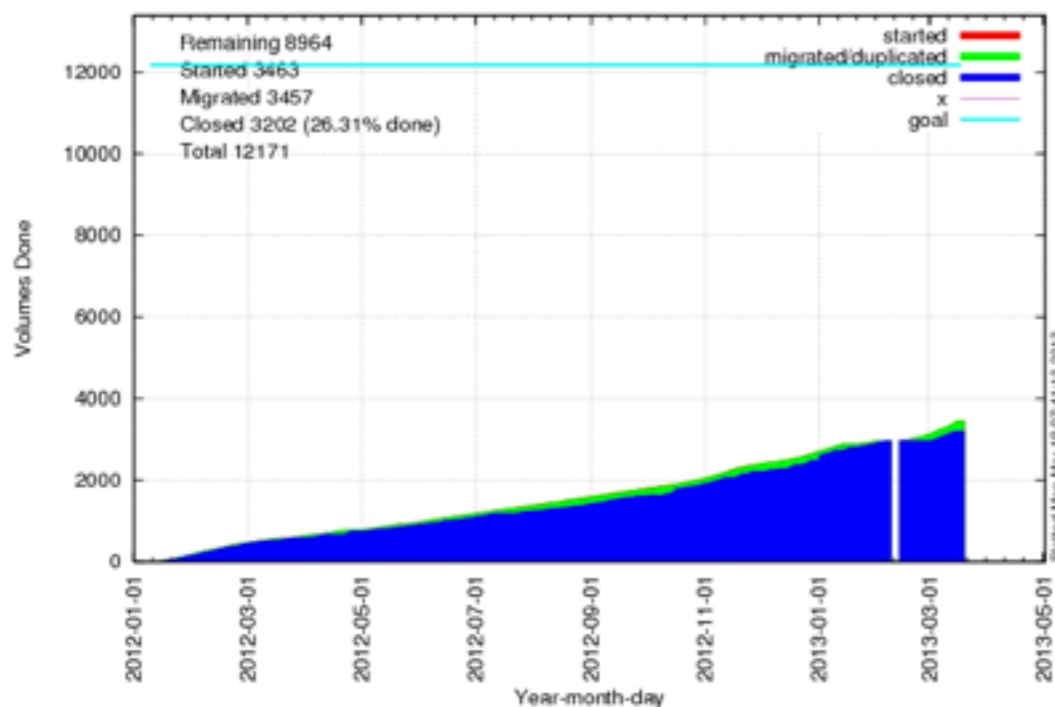
Mass Storage at Fermilab

- ~65 PB of tape storage in use at Fermilab
 - Employs LTO4 and T10K tape technologies
- Currently migrating LTO4 tapes to T10KC
 - ~5K out of 25K CDF+DØ tapes migrated
- Currently separate tape servers (enstore) and caching (dCache/SAMCache) for CDF/DØ
 - Investigating merging some/all of these resources by 2015


64.7 Petabytes on Tape 3/1/2013



Migration/Duplication summary accumulated for LTO4 on CDFen



Petabytes of Data on Tape



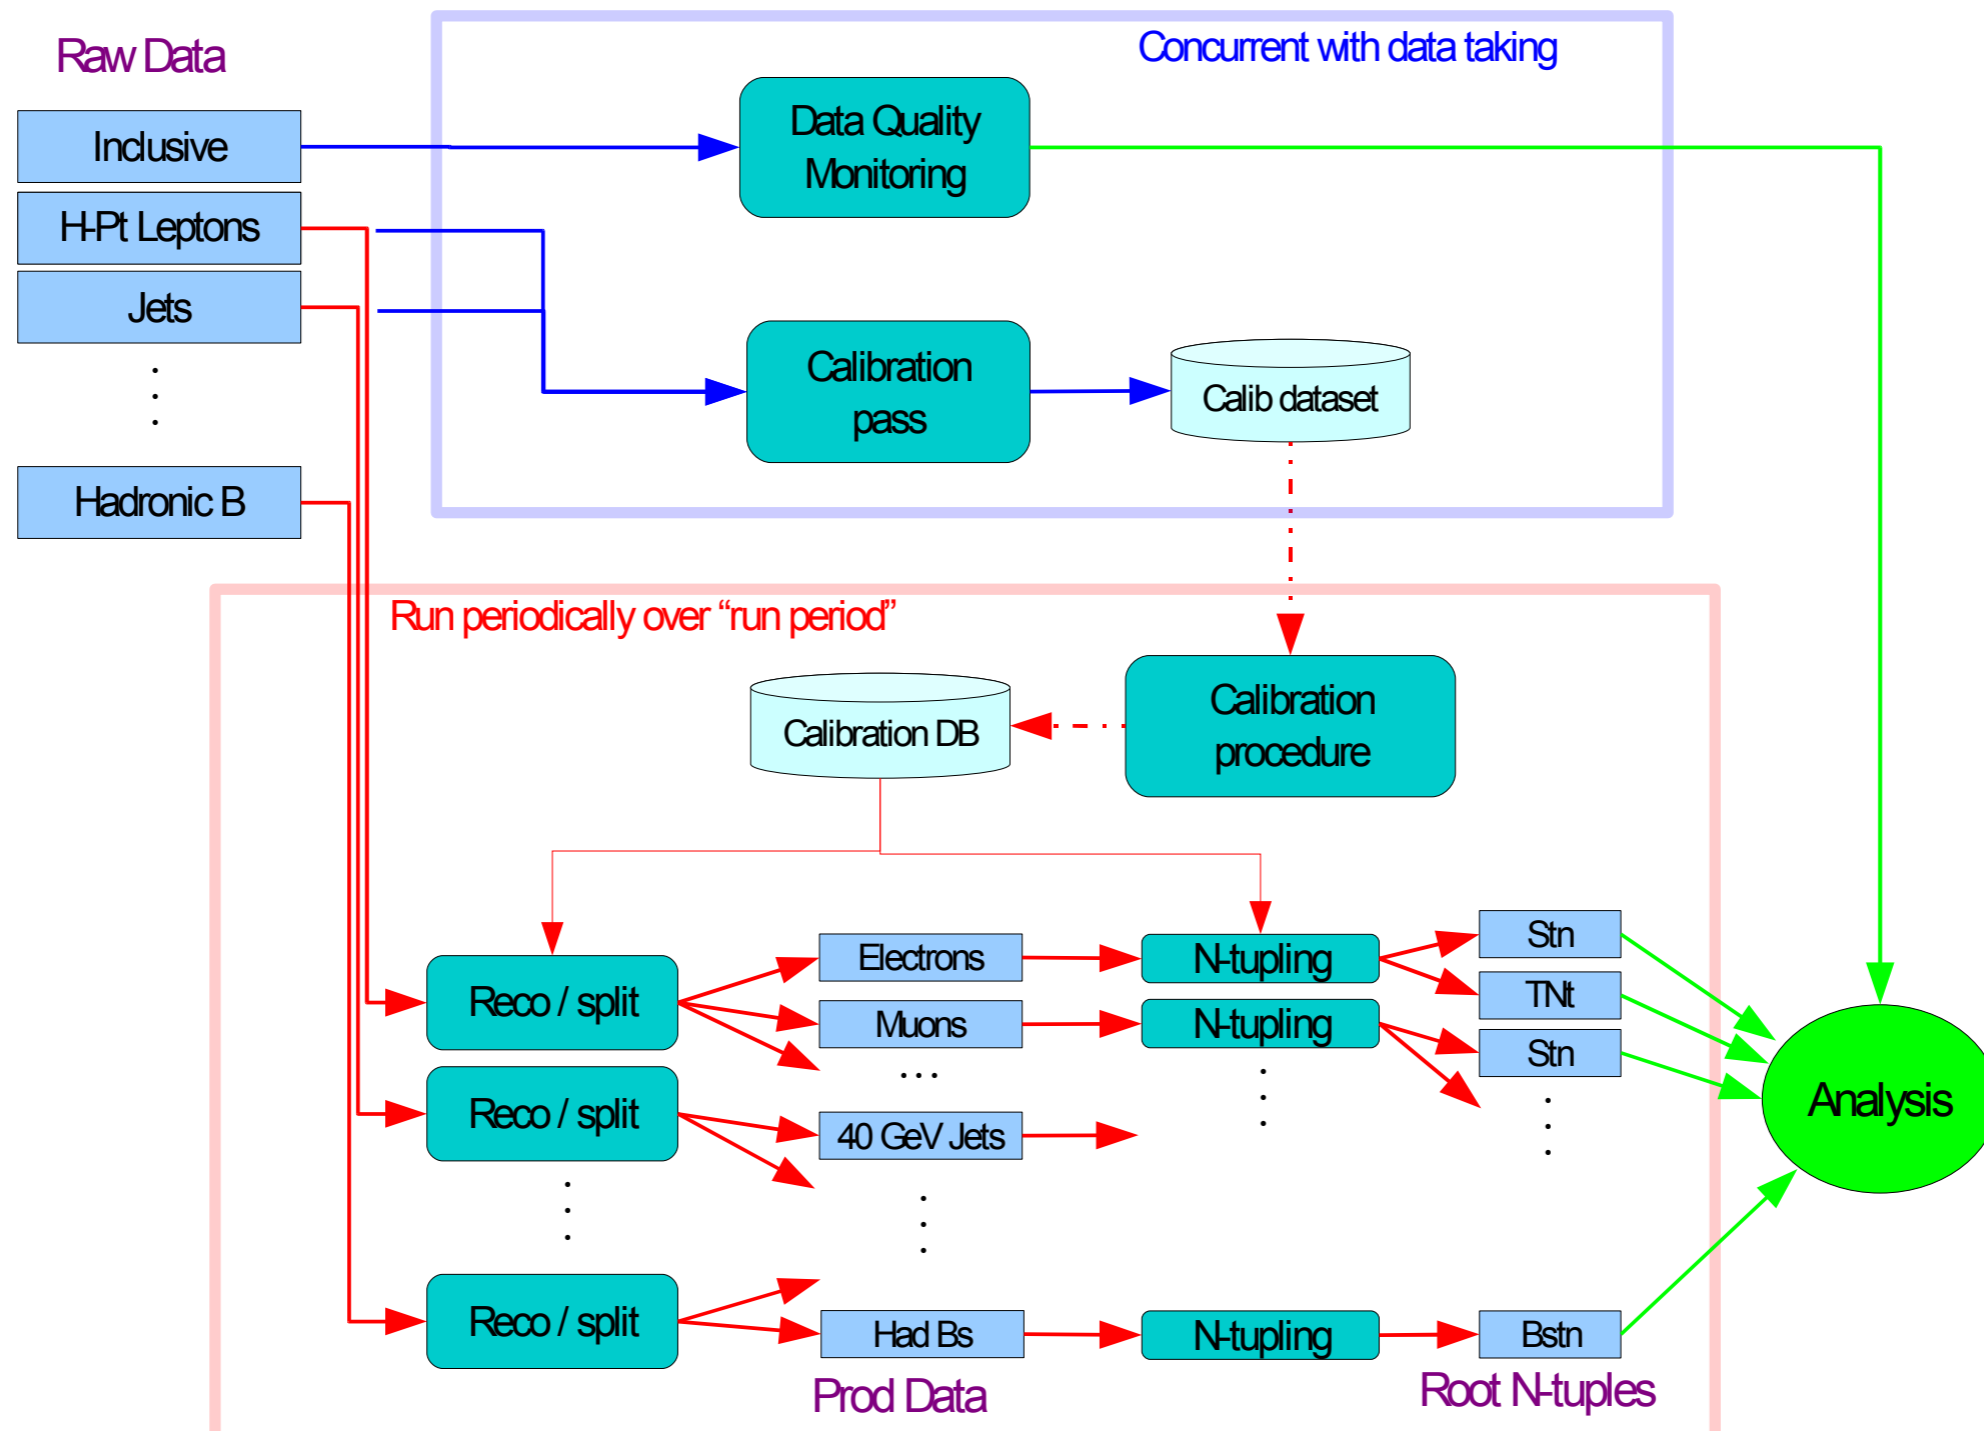


Tape Migration and Data Loss Protection

- After a tape is migrated each file on the original list is read back
 - Recalculates Adler32 checksum and compares against original
 - Checks for corruption in copy process
- Original tapes are kept as long as possible
 - Only LTO3 tapes have been erased/reused so far
 - LTO4 originals are still available
- Periodic reads from each tape
 - First file, last file, random file in middle
 - Checks for degradation and wear
- All **non-statistical data** are stored in Oracle database
 - *e.g.* calibration

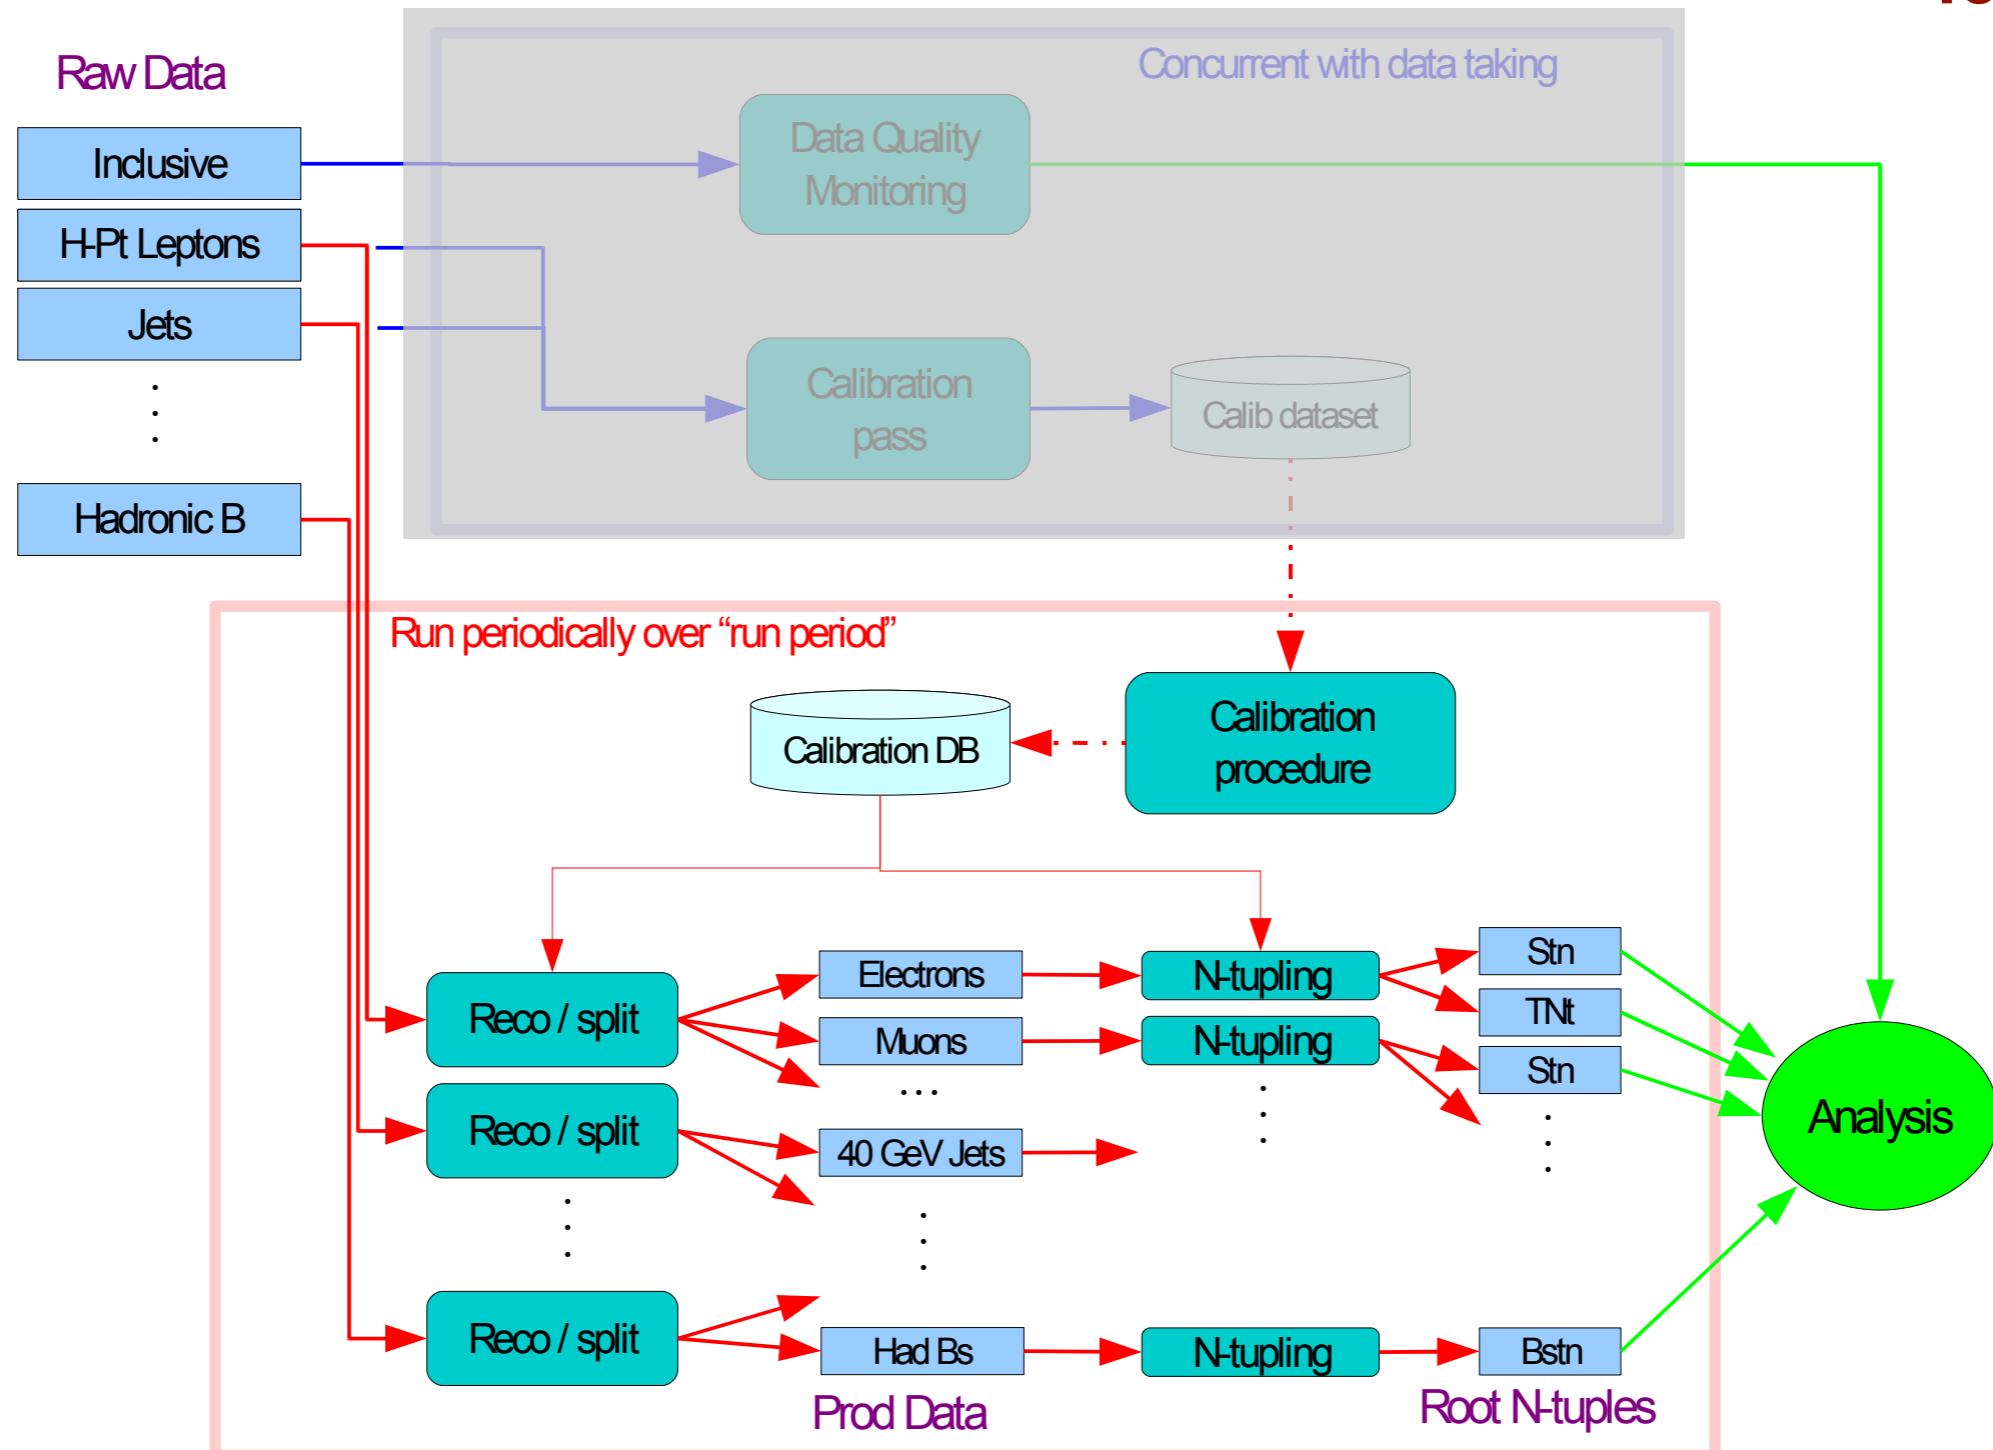
The CDF Analysis Chain: Collider Data

Up to September 2011



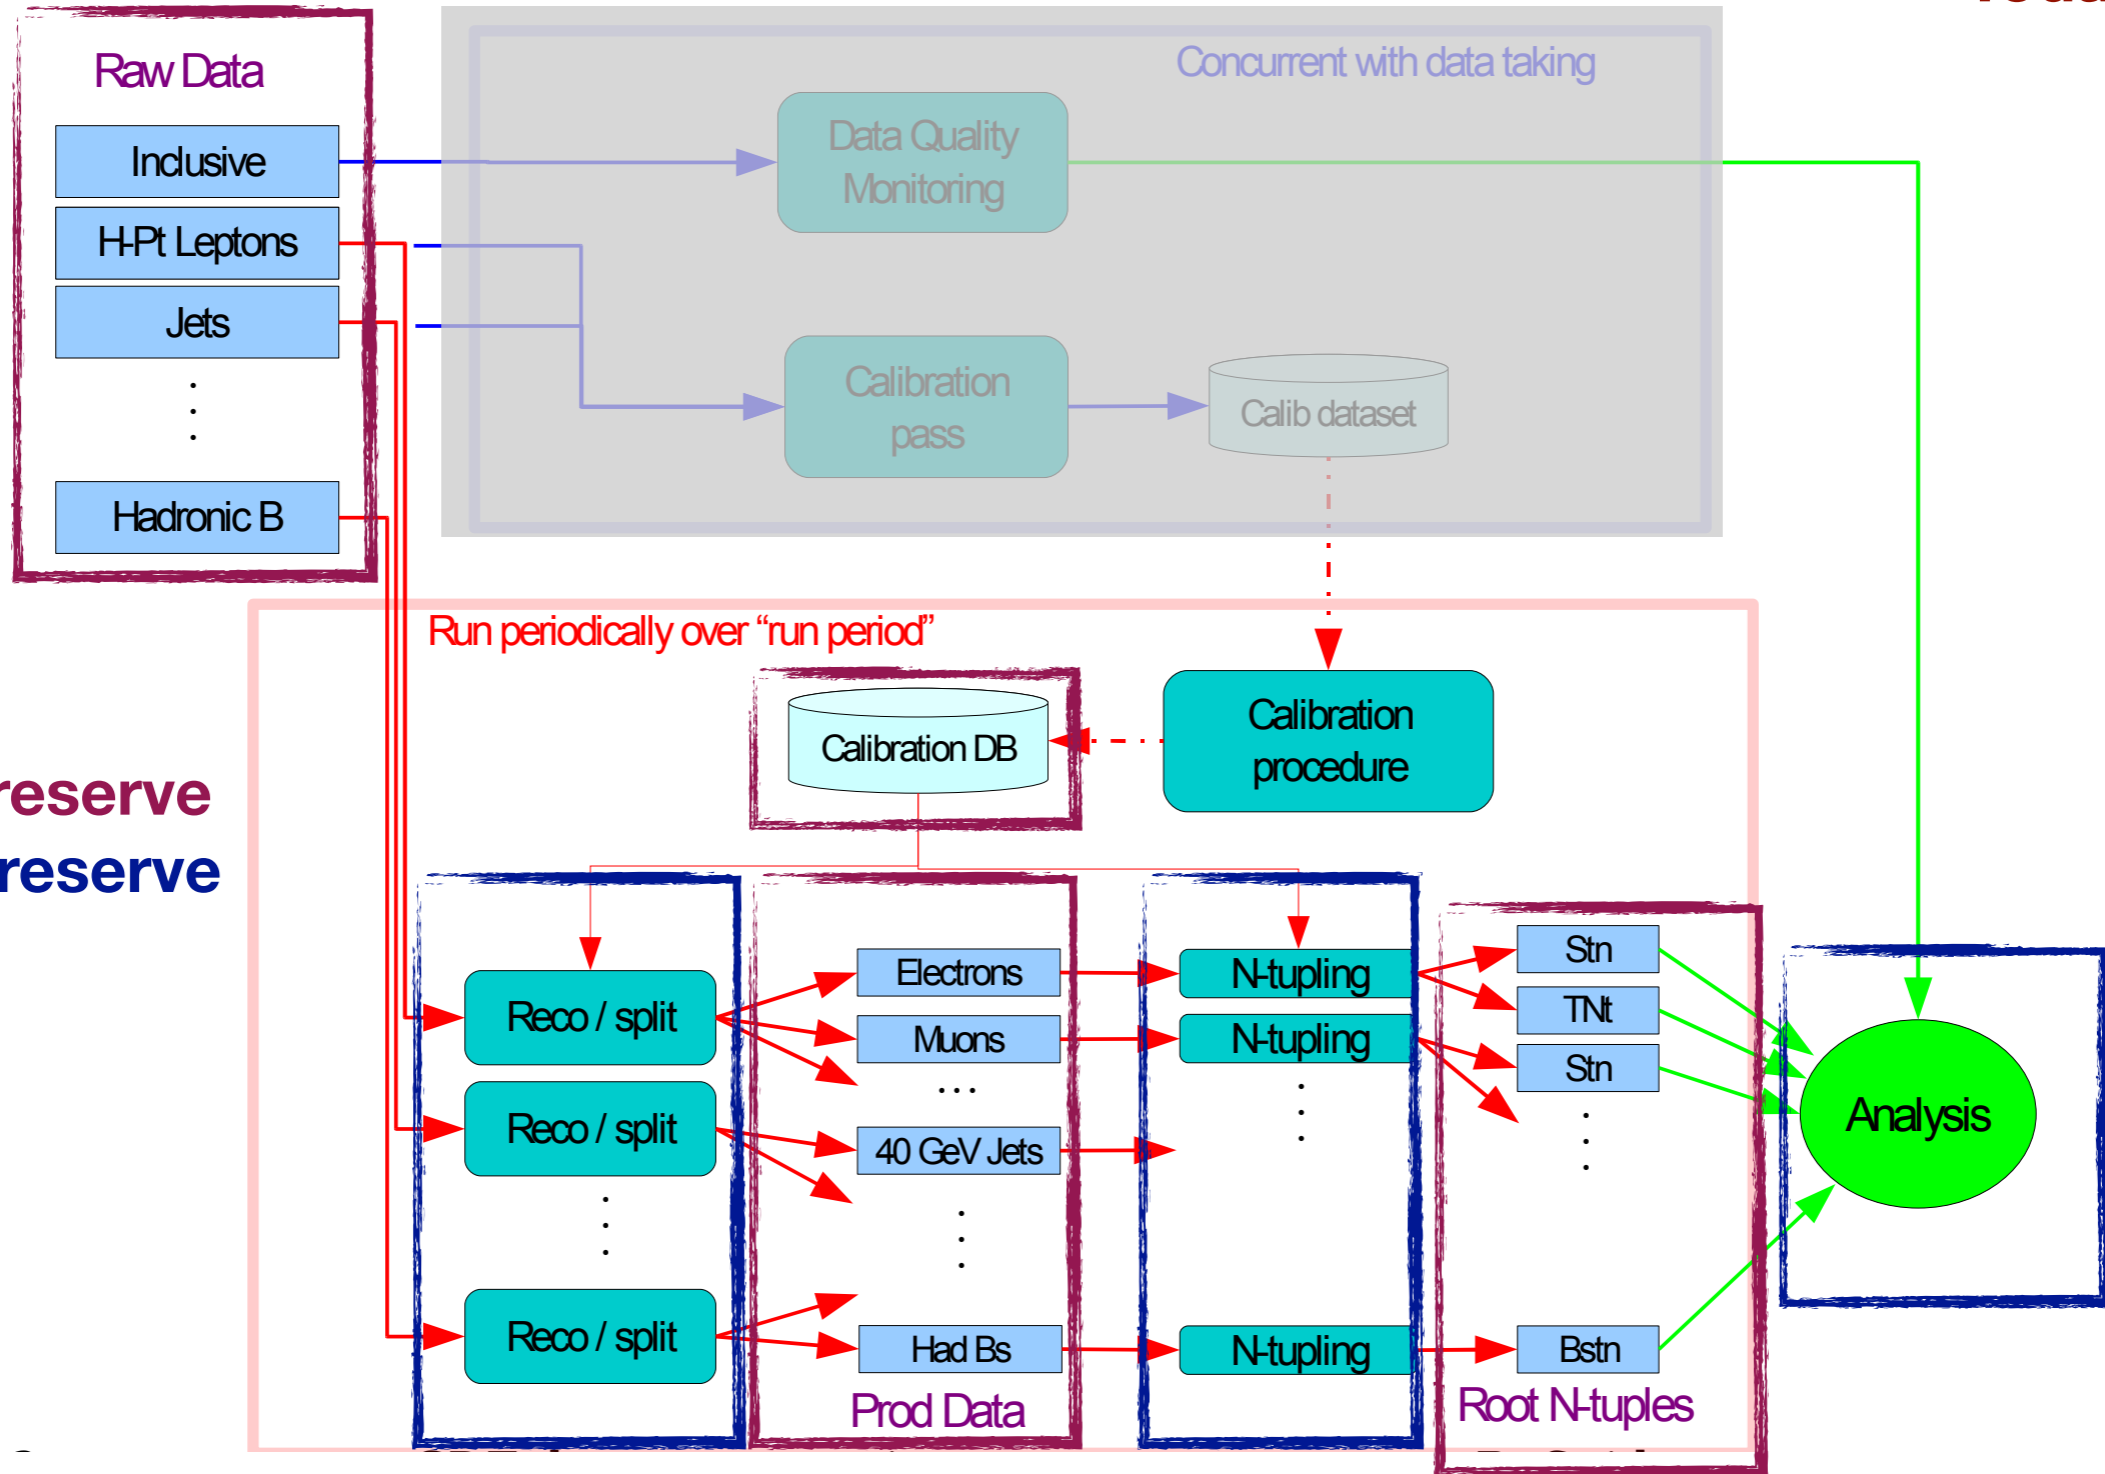
The CDF Analysis Chain: Collider Data

Today



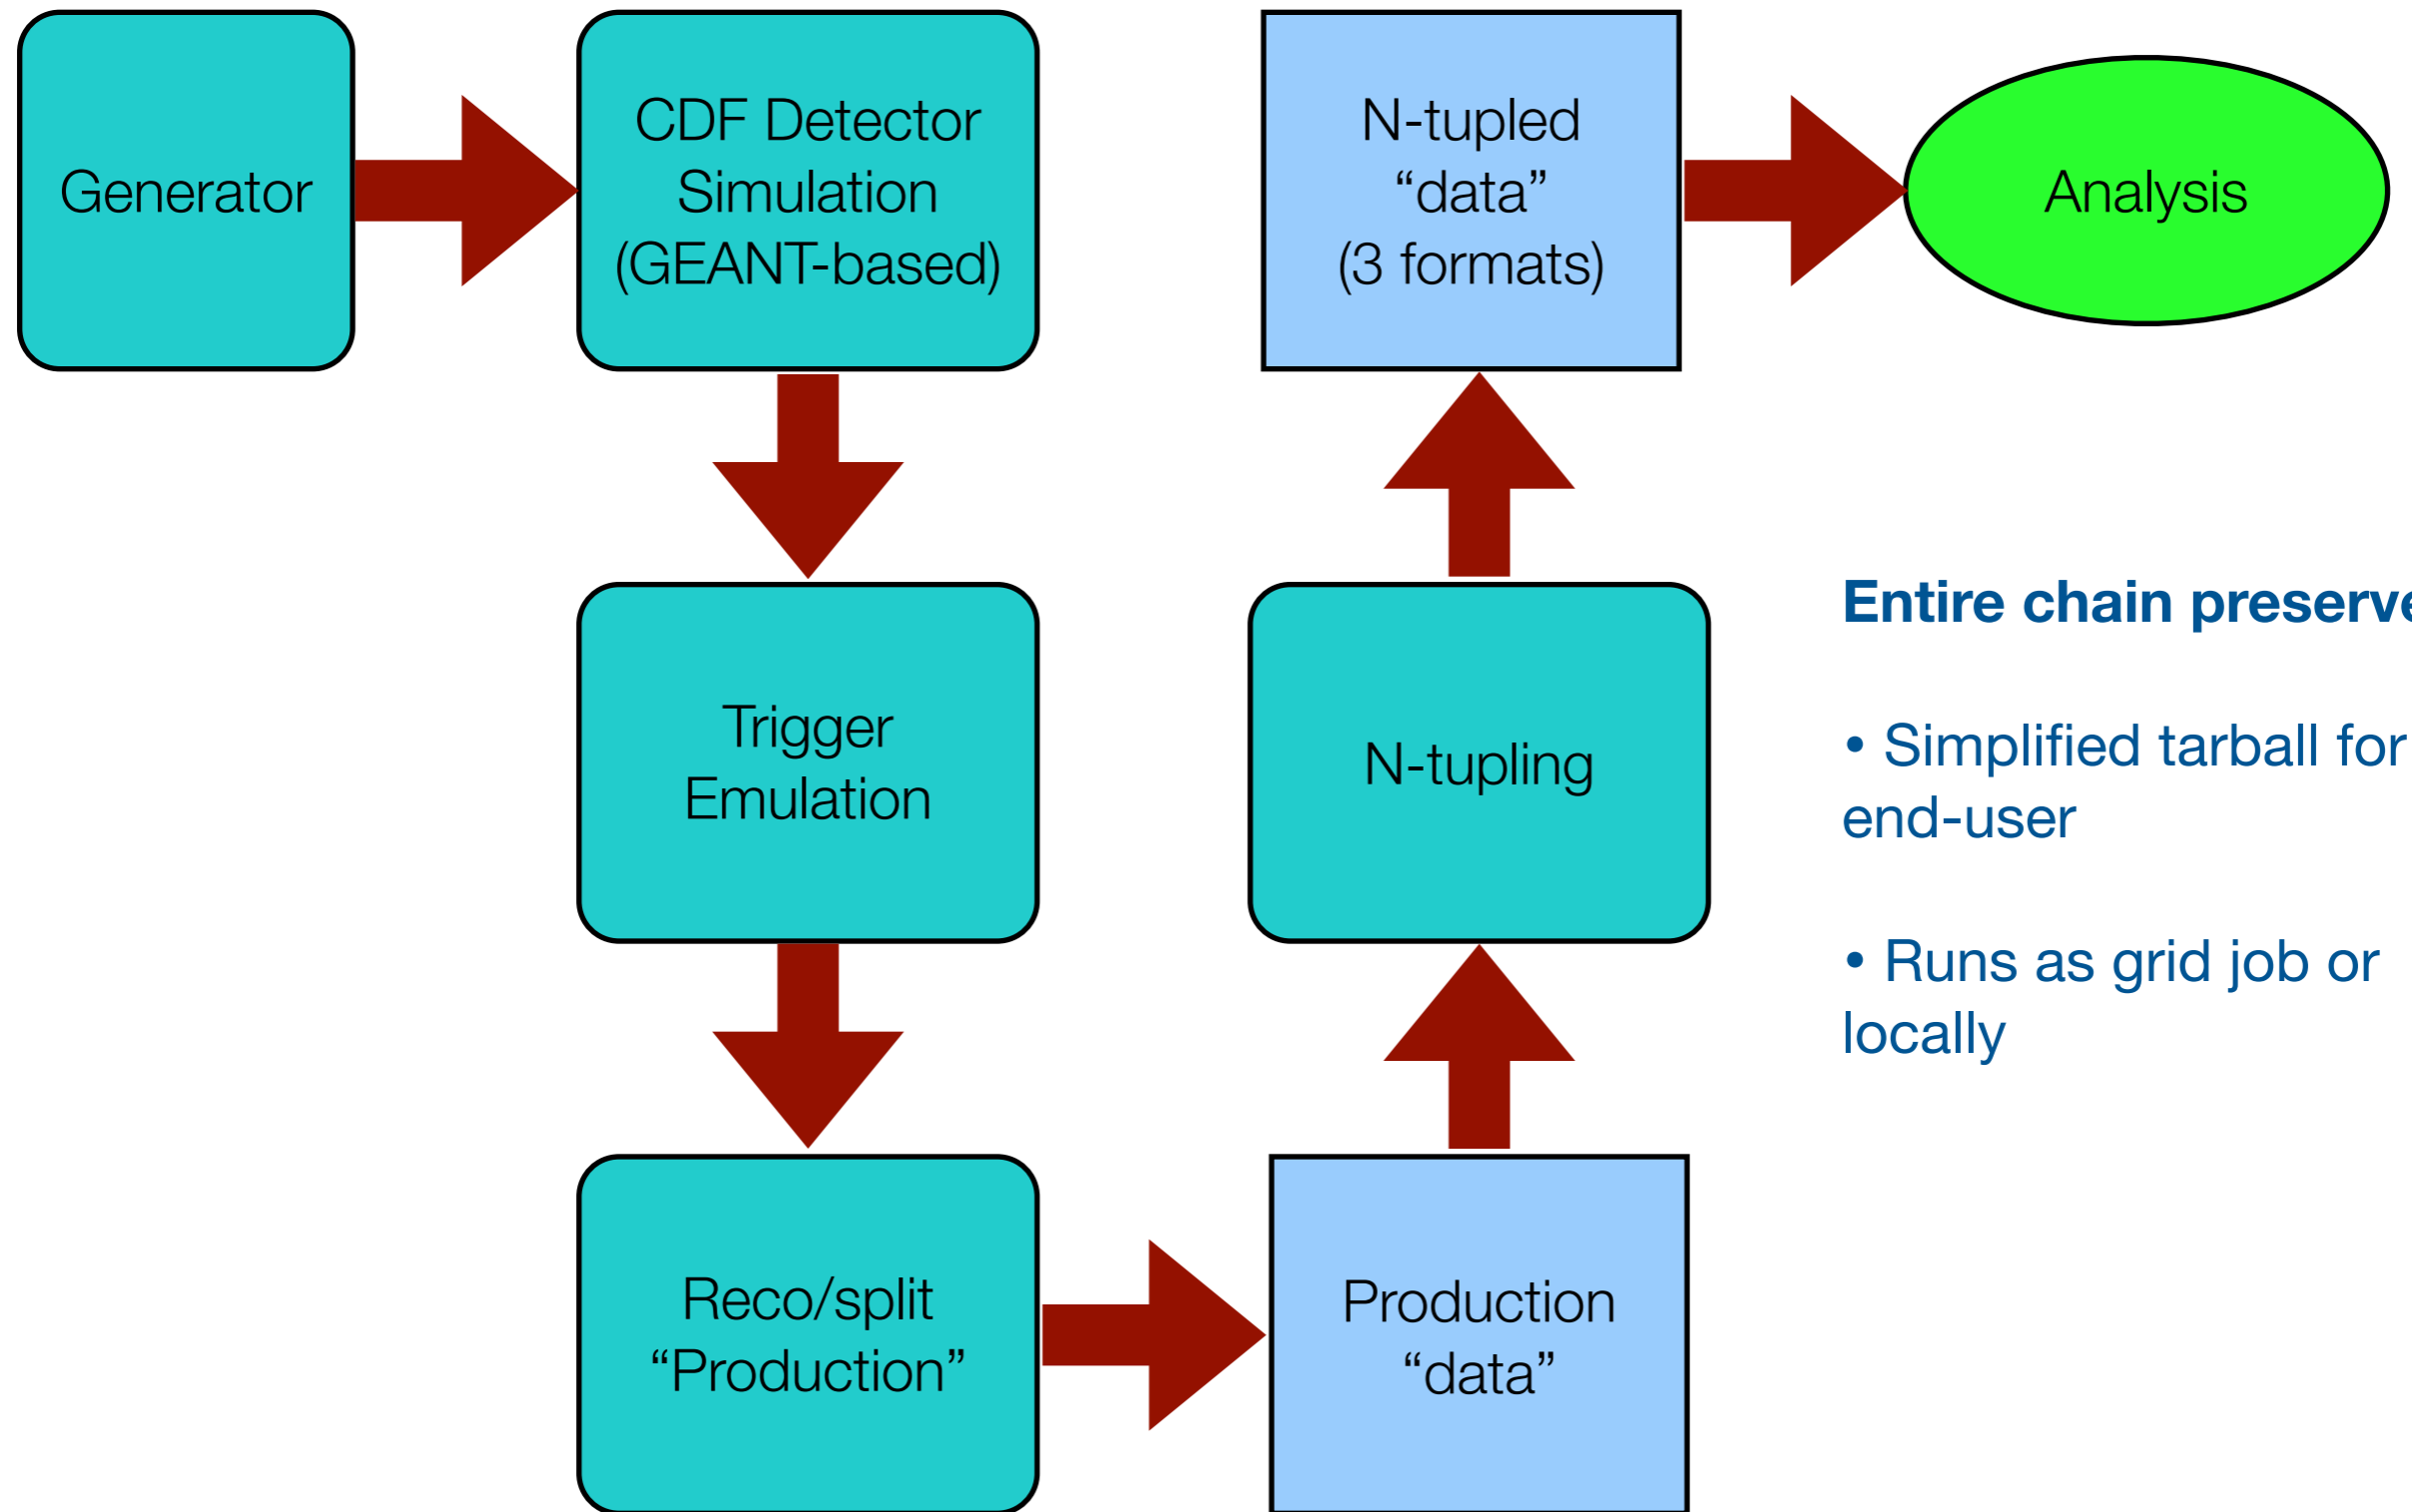
The CDF Analysis Chain: Collider Data

Today



Data to preserve
Code to preserve

The CDF Analysis Chain: MC Simulation



Entire chain preserved

- Simplified tarball for end-user
- Runs as grid job or locally

The CDF Data

- ~10 PB of archived data at CDF consisting of **raw** data (~2 PB), **reconstructed** (“production”) data (~4 PB), **ntupled** data (~1.5 PB) and **Monte Carlo simulation** (~2PB)
 - Data preservation strategy includes continued ability to create new CDF-simulated MC samples
- CDF Italian institutions plan to migrate a subset of the data (raw+ntuples) to CNAF as part of future analysis platform
 - Copied data will be accessible by user jobs running on EuroGrid
 - Using the GARR (Italian R&E) network for copying
 - First feasibility tests performed last summer: successful copy to CNAF and retrieval from tape of a small dataset
 - Project received first approval by INFN at the end of September





CDF Analysis Software

- CDF software in a set of **CVS** packages organized into centralized releases
 - User analysis code also kept in CVS repository: backed up/archived
- Last **major** release in ~2005 (“Gen6”)
 - Minor updates in following years, including **SL5 compatibility**
 - Uses **compatibility libraries** built on SL3/4
- Currently developing/testing a **legacy release** for CDF
 - Gen6-based release that is **built on SL5** (no compatibility libraries), executable on SL6. Then, built on SL6.
 - Complete build and validation planned by **end of summer 2013**
- CDF long term data preservation strategy relies on legacy release being executable through 2020 (on SL6)
 - Virtualized **build server** will be maintained for OS/kernel-related patching that is needed



CDF Analysis Software: Tools

- Additional tools are built and maintained using **ups/upd** (Unix Product Support/Distribution)
 - Includes tools such as root
- Examining complete list of CDF ups products and determining what can be dropped
 - Products that have not been used in years and likely won't be
 - Products that are redundant with OS-distributed versions (e.g., python)
- Others that have not been rebuilt in a while to be rebuilt on SL5/6

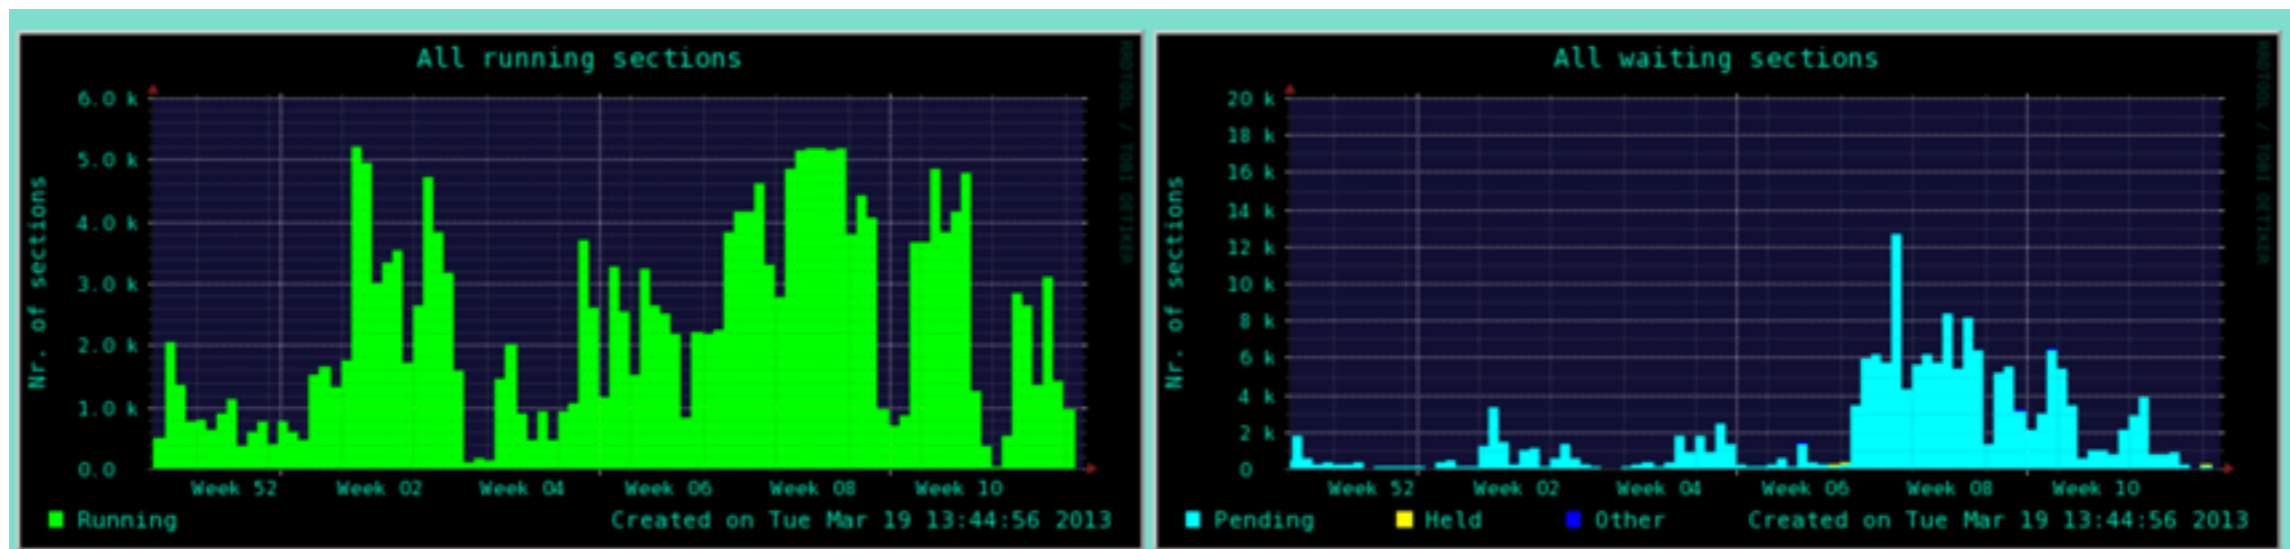
Product	Shortlist?	Type	Note	Version	Recompiled?	Fate
HTMLgen	No	Python generating HTML		v2_1	NO	Drop
blt	No	graphing w/tcl/tk	Does anybody use this?	v2_4u Ignore?		Drop
boost	No	Utility	Used (Weiming; others?)	New version		TBD
cafclient	Yes	Infrastructure	sets up an old version of both python and krb5_libs	use current	NoOP	keep current
cdf_utilities	YES	release-independent	various handy scripts	use current	NA	keep current
cdfdab	Yes	Infrastructure	NULL flavor	V1_1	NA	
cdfdb_data	No	Infrastructure	NULL flavor	use current	NA	
cdfsoft2_products	No	install tool	used to make sure necessary products are installed created when a release is ready for distribution	should match cdfsoft2 version	NA	
cern	Yes	Infrastructure		2006b	YES	

Example

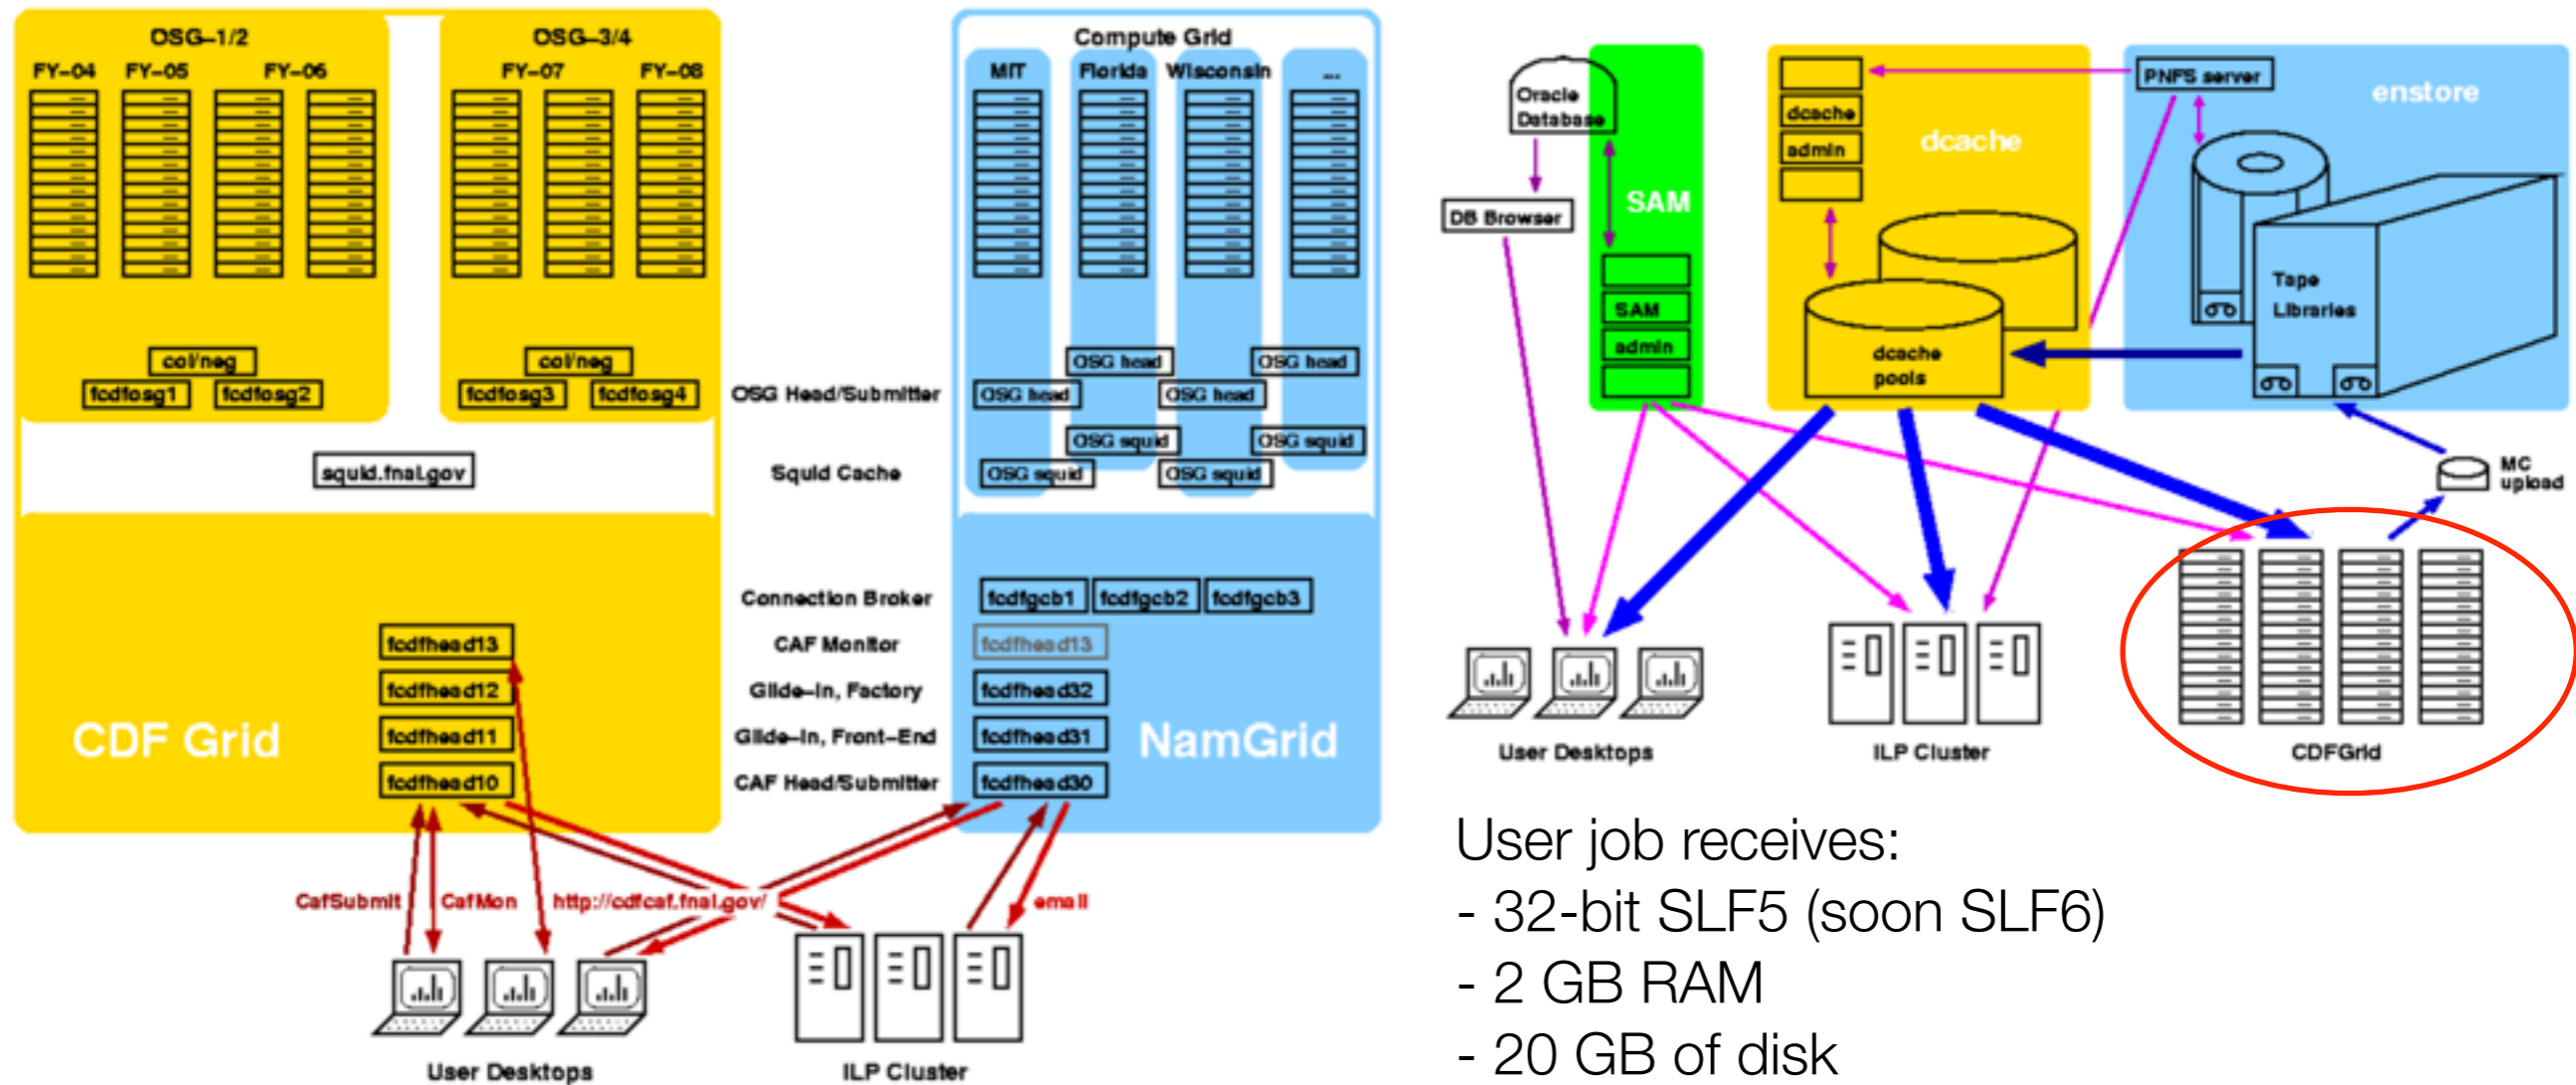
Job Submission at CDF

- CDF general computing on Central Analysis Farm (CAF)
 - Uses **Condor** with **glidein** WMS
 - Condor also used by many other Fermilab experiments (e.g. USCMS Tier1)
- CAF provides a generic interface for the end-user, interfacing to one of three sites
 - CDFGrid: Primarily dedicated resources at Fermilab
 - NamGrid: Dedicated resources off-site (e.g., MIT, KISTI) and opportunistic resources at Fermilab
 - EuroGrid: Opportunistic resources at CNAF and LCG Tier 1 sites
- Usage still fairly substantial

CDFGrid, past 3 months



Job Submission at CDF (cont.)



- Dedicated resources will dwindle
- Foresee eventually moving to **opportunistic use of Intensity Frontier computing**
 - Can recreate necessary hardware environment using **virtualization**
 - Condor/grid infrastructure already in use will require little change
 - End-users would utilize same tools

Documentation

- Much CDF knowledge preserved in a patchwork of web pages, wiki pages, and electronic logbooks
 - As much of this as possible being preserved
 - Would like to migrate wiki-based systems to static HTML
 - E-logs from detector operations utilizing e-log system still in use by Fermilab accelerator division
- ~950 thesis records for CDF+DØ in Inspire
 - Fermilab library has electronic versions for nearly all Run 2 theses



2000 CDF E-Log -- Eve shift, Wed Apr 26, 2000

DAQ Ace	Monitoring Ace	CO
Steve Hahn, flunkie		

2012 CDF E-Log -- Day shift, Tue Jan 17, 2012

SciCo	DAQ Ace	Monitoring Ace	CO	(Operations Manager)
Steve Hahn	Steve Hahn	Steve Hahn	Steve Hahn	Steve Hahn

CDF Fast Navigator

CDF News	CDF Forums	Online			
Shift tools	Beam status	MCR e-log	Shift e-log	All e-logs	Run summary
		Ace HELP	CO	SciCo	Expert phone list
Operation	B0 Home	Ops Manager	Ops e-log	Weekly Plan	CDF Safety
					Training records
Detector	Silicon / COT	Calorimeter	CLC	Forward Detectors	Run IIa / Run IIb
	Rad Monitoring	Muon	TOF		
Trigger	Trigger Home	L2 / L3	Trigger WG	B Trigger	Exotics Trigger
Data quality	DQM Home	Goodrun lists	Consumer Slides	Consumer Home	Physmon

actually Steve'

2000

Start of Shift Notes: 2012 - best year ever! (and last)

- Tue Jan 17 10:16:38** This e-log is now officially closed. For more details, try the "Make Entry" button in the navigation bar at the top. - Steve Hahn
- Tue Jan 17 15:01:41** comment by...AntiSteve -- That button provides no additional informations
- Tue Jan 17 15:51:29** comment by...Satyajit (CDF ops) -- How come I don't get to utter the last word?!
- Thu Jan 26 12:01:17** comment by...Farrukh --

Ha ha I have the last word

- Thu Feb 9 12:44:29** Hmm, have restored old code (yes, that means entries can be made again) till I figure out why new code caused problems accessing e-log. (entry outside this shift's time range) - Steve Hahn
- Wed Mar 21 13:35:26** Hello! (entry outside this shift's time range) - AntiSteve

2012

Tevatron Impact Symposium, 6/11/2012

Startup
Misc
TOTAL:
- Steve Hahn

CDF-1

Parameters of Colliding Beam Detector

R. Diebold, A. Tollestrup, T. Collins, S. Ecklund, J.K. Walker
23 Jan '78


Constraints

1. Kissing scheme for beams.
2. Maximum energy of MR/ED beams = 150/1000 GeV for pp and 1000/1000 for \bar{p}
3. Conventional magnets for the MR normal operation.
4. Low β achieved without loss of part of the 50m long straight.

(Note: R. Diebold paper #1 of Summer Study assumes 46m is free.)

- First 1000 CDF Notes previously only on paper now scanned and archived electronically
- Remaining ~10k notes to be migrated to Inspire
 - CDF Inspire account obtained and test note uploading/retrieval underway
 - Hope to have archive migrated by end of the year

Conclusion (Part I)



- CDF Data Preservation Task Force report released in September of 2012
- Implementation of this plan now well underway both by experiment and Fermilab CS
 - Fermilab CS project committed to demonstrating functionality of data preservation system within 2 years time
 - More Tevatron-wide issues will be discussed in Ken's talk next
- Human factor is most difficult
 - Expertise is mostly gone already and continues to dwindle
 - Time is of essence!