Update on the LIGO Data Analysis System

Kent Blackburn / Stuart Anderson
LIGO Laboratory
California Institute of Technology

LSC Meeting
LIGO Livingston Observatory
March, 2002

Outline

- Summary from E7 Run
- LDAS Timeline
- Highlights for LDAS Releases
- Announcements to LDAS Users
- Hardware Update – Stuart Anderson

Update on LDAS Software

(Kent Blackburn)

Preparation for E7 Run

- Ran with version 0.0.23 (alpha code).
- Missing a major sub-system (diskCacheAPI).
- 3 weeks of testing prior to December release.
- Shown to have 99.7% reliability on Caltech development and test hardware platforms.
  - Primarily due to unresolved multithread problem.
  - Lesser issue with communication timeouts in parallel computing components.

Summary - E7 Run

	Hanford LDAS	Livingston LDAS	MIT LDAS	CIT-TEST LDAS	TOTAL
Total Jobs	63600	48775	280	915	113570
Database Rows	4188188	2789132	1062	2096	6980478

- LDAS up entire run: Dec. 28th, 2001 - Jan. 14th, 2002
  - Approximately one job every 10 seconds (averaged).
  - Approximately five rows every second (averaged).
- Greater than 90% of jobs completed successfully
  - LHO roughly 92%; LLO roughly 95%; Not checked at CIT or MIT.

LDAS Release Timeline

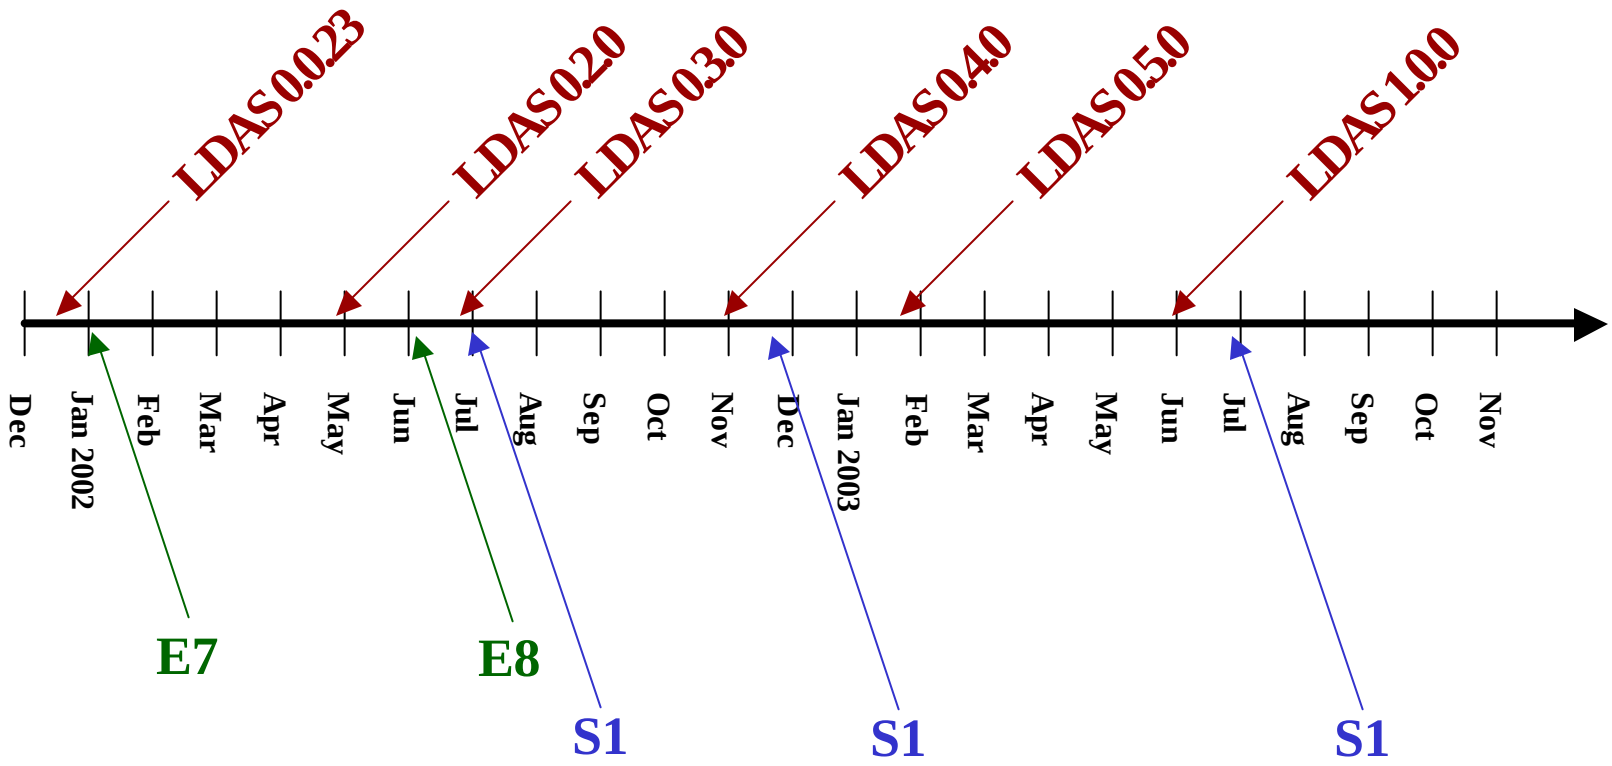

LDAS 0.2.0 Highlights

- First beta version - major features in place.
- Fixes for majority of issues identified in E7.
- FrameCPP C++ I/O library now Version 5.
- New diskCacheAPI.
- Added signal resampling to frameAPI.
- Added *interpolation()* action.
- New *usertag* database option in dataPipeline.
- Extended support for producing reduced frame files in frameAPI.

LDAS 0.3.0 Highlights

- Speed up MPI initialization with persistent LAM daemons.
- Implement TCL “channels” to manage data sockets.
- Refine support for Frame 5 specification.
- Refine RDS frame production.
- Extend dataConditionAPI signal processing.
- Extensive testing in preparation for S1.

LDAS 0.4.0 & 0.5.0 Highlights

- Port archive technology to diskCacheAPI.
- Integrate calibration solution(s).
- Profile & apply C++ optimizations.
- Port code to next generation of Linux, compiler and CPU hardware.
- Add security authentication.
  - Leverage off GriPhyN collaboration.
  - Demonstrated as layer over LDAS at SC01.

LDAS 1.0.0 Highlights

- Continue profiling of LDAS.
- Continue optimization of LDAS.
- Extend signal processing tools.
- Introduce new user commands.
- Integrate Grid data publishing interface.
  - Also demonstrated in prototype at SC01.
- Exhaustively test in preparation for first full release and S3.

Important Announcements

- Two new tables in database for “DSOs”.
- Intermediate() action becoming output().
  - Existing dataPipeline commands must be updated.
- Shared LDAS User Accounts being removed.
- All LDAS User Accounts must be renewed for the LDAS release used by first Science Run.
  - Reapply starting in early June.
- LDAS will migrate this year to use a secure authentication.
  - User accounts will be reissued again for that release.

Update on LDAS Hardware

(Stuart Anderson)

LDAS S1 & S2 Configuration

“Increase computational capacity over E7 and investigate advanced storage configurations but delay full compute farm deployment until S3.”

	SAN (TB)	IDE (TB)	CPU (GHz)	Tape (TB)
LHO	12		60	18
LLO	6		60	9
CIT	3	50	60	100
MIT	1	TBD	35	TBD
DEV	1		15	1.5
TEST	1		21	

SAM-QFS vs HPSS

- SAM-QFS advantages
  - Simplicity/reliability
  - Media import/export
  - Cost allows for use at observatories
  - Disaster recovery (GNU TAR)
  - Metadata performance (x1000)
  - Single vendor solution (server, software and OEM storage)
- HPSS advantages
  - Few man year experience
  - Free at Caltech
  - 35TB successfully stored