

BaBar's Production

- 22 billion events (~540/fb) recorded
- ~9 billion events reconstructed
- >10 billion events simulated with latest reconstruction code ... so far
- 24 major reprocessings

- Total data volume collected and produced > 6 PetaBytes
- Legacy data volume ~2 PetaBytes (800TB of raw data, 1.2 PB from the last data reprocessing)

BaBar Collaboration Caps Meeting Week with 400th Scientific Publication

by Lauren Knoche

The BaBar Collaboration reached another milestone Tuesday—just in time for celebration during the [group's meeting](#), which ends today at SLAC. The collaboration published its 400th paper Tuesday, less than nine years after publishing its first in 2001. That's an average of one publication per week, every week, for nearly nine years straight.

"I do not know of any other collaboration that has achieved such a production rate of outstanding quality science in particle physics, it is really something rare," said BaBar spokesperson Francois Le Diberder.

The [milestone paper](#) was published online Tuesday and appears in the November 1, 2009 issue of *Physical Review D* (Volume 80, Number 9). The study examines differences in the rates at which subatomic particles called B^+ mesons and their antiparticle partners, B^- mesons, decay to related particles called "charm" and "strange"



(Photo by Brad Plummer.)

BaBar's Data at SLAC

- Legacy data (raw data + latest reprocessing/simulation, physics skims and user data) stored in HPSS on 1 TB T10K tapes
- Old raw data tapes (9940a/b 185 GB) stored away for safe keeping in building 33
- Disk buffer (Kan servers; SUN x4540s)
540 Tbytes
- Reconstructed/simulated/skimmed data accessed via XROOTD
- Data staged onto BaBar's archival system as requested by users' jobs

How/Why BaBar Computing became Distributed

- Funding that would have been needed to satisfy the computing needs for doing it all at SLAC was too much.
- Needed resources made available at the participating institutions
- Developed mechanisms to coordinate and combine these resources to satisfy the needs.

BaBar's Data Production, Storage, and Use is Distributed



◆ SLAC

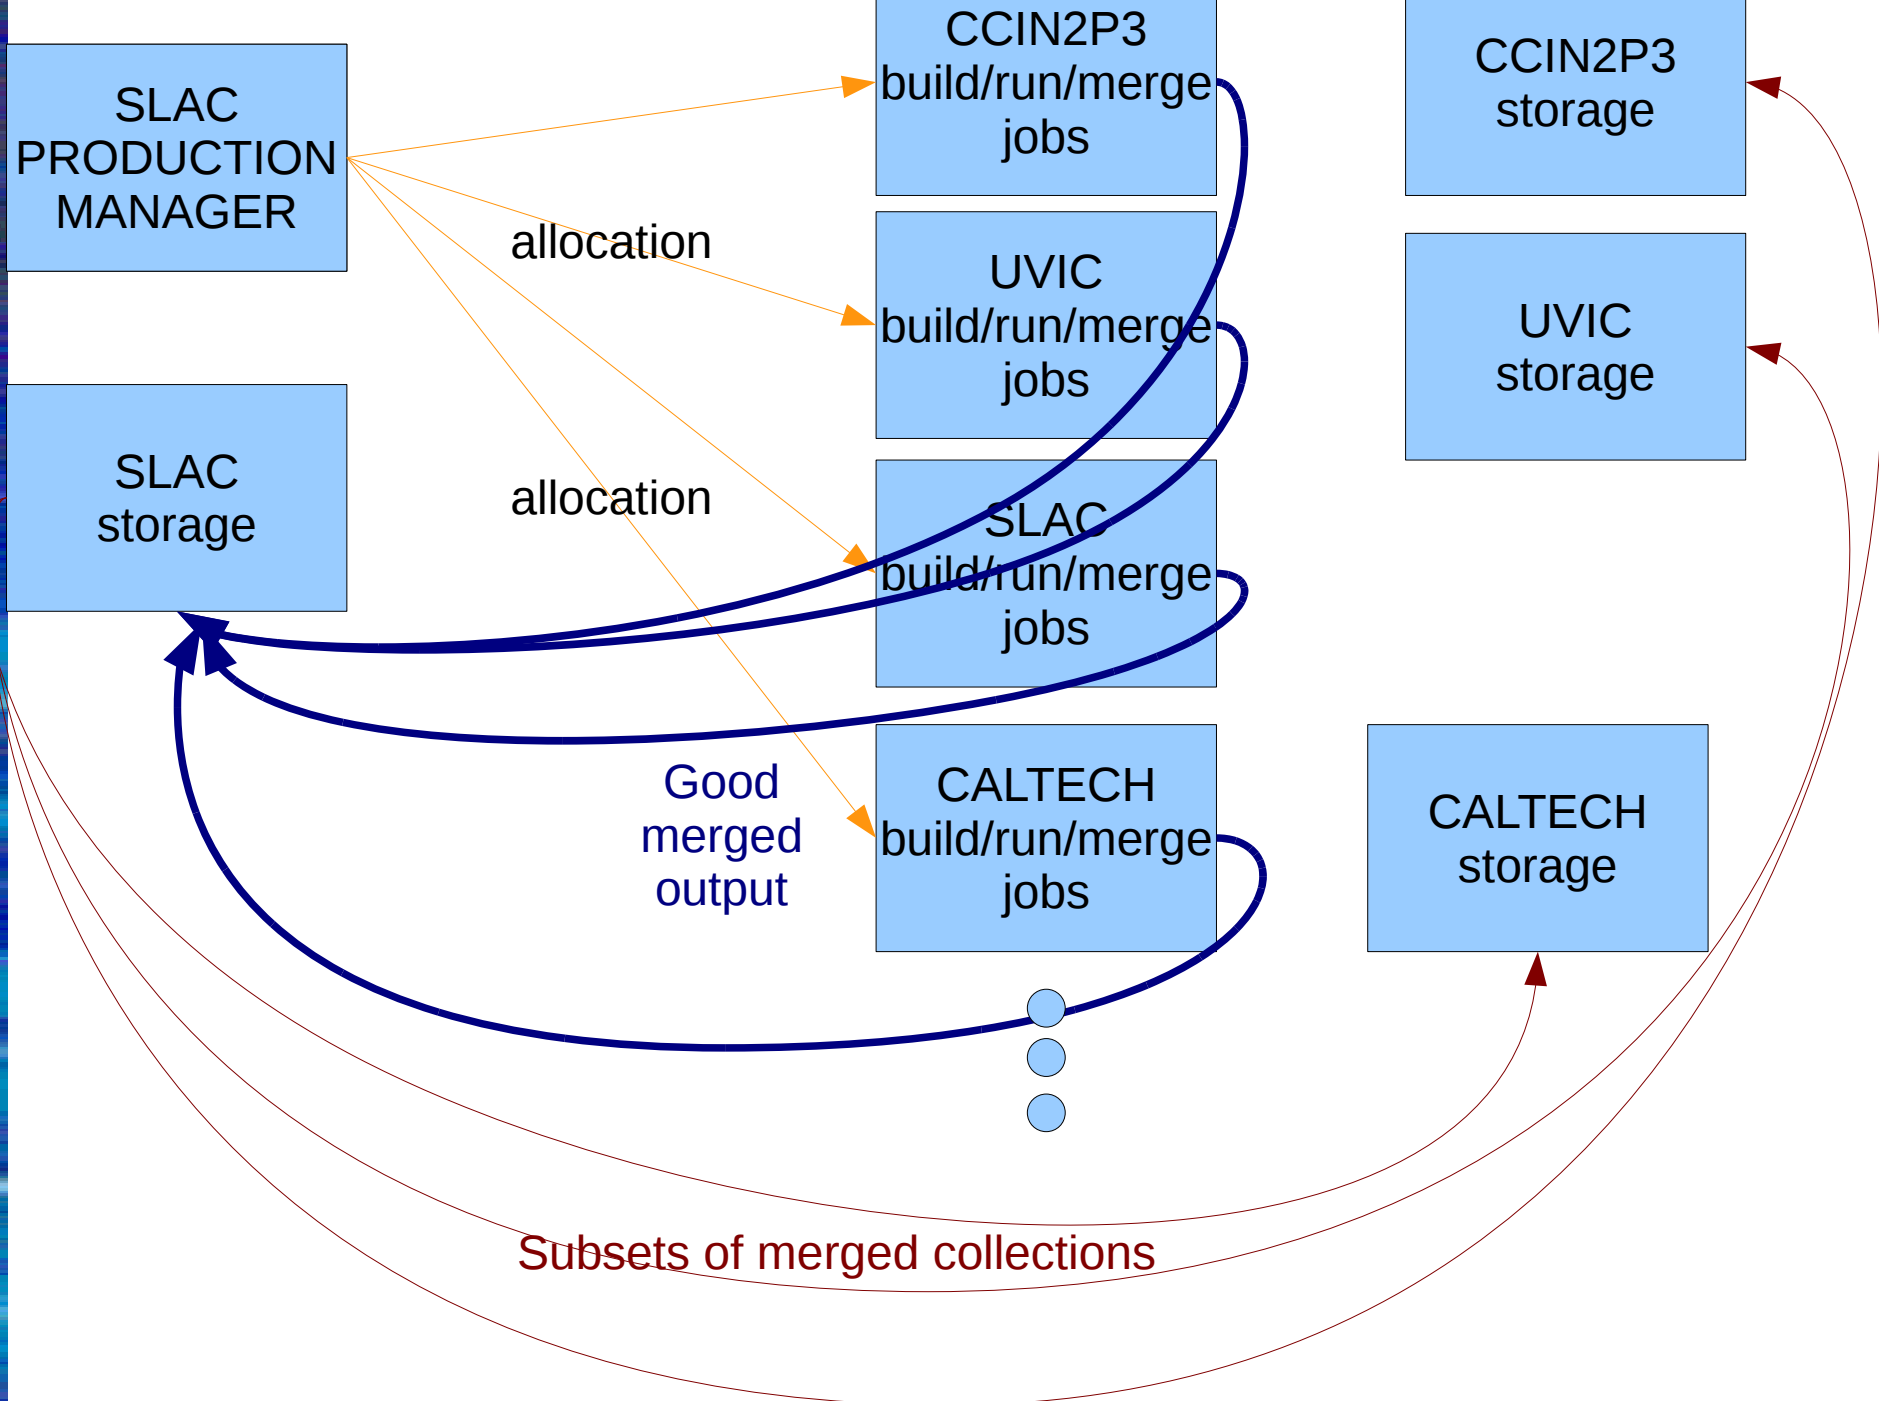
+

data sent where needed by Analysis Working Groups:

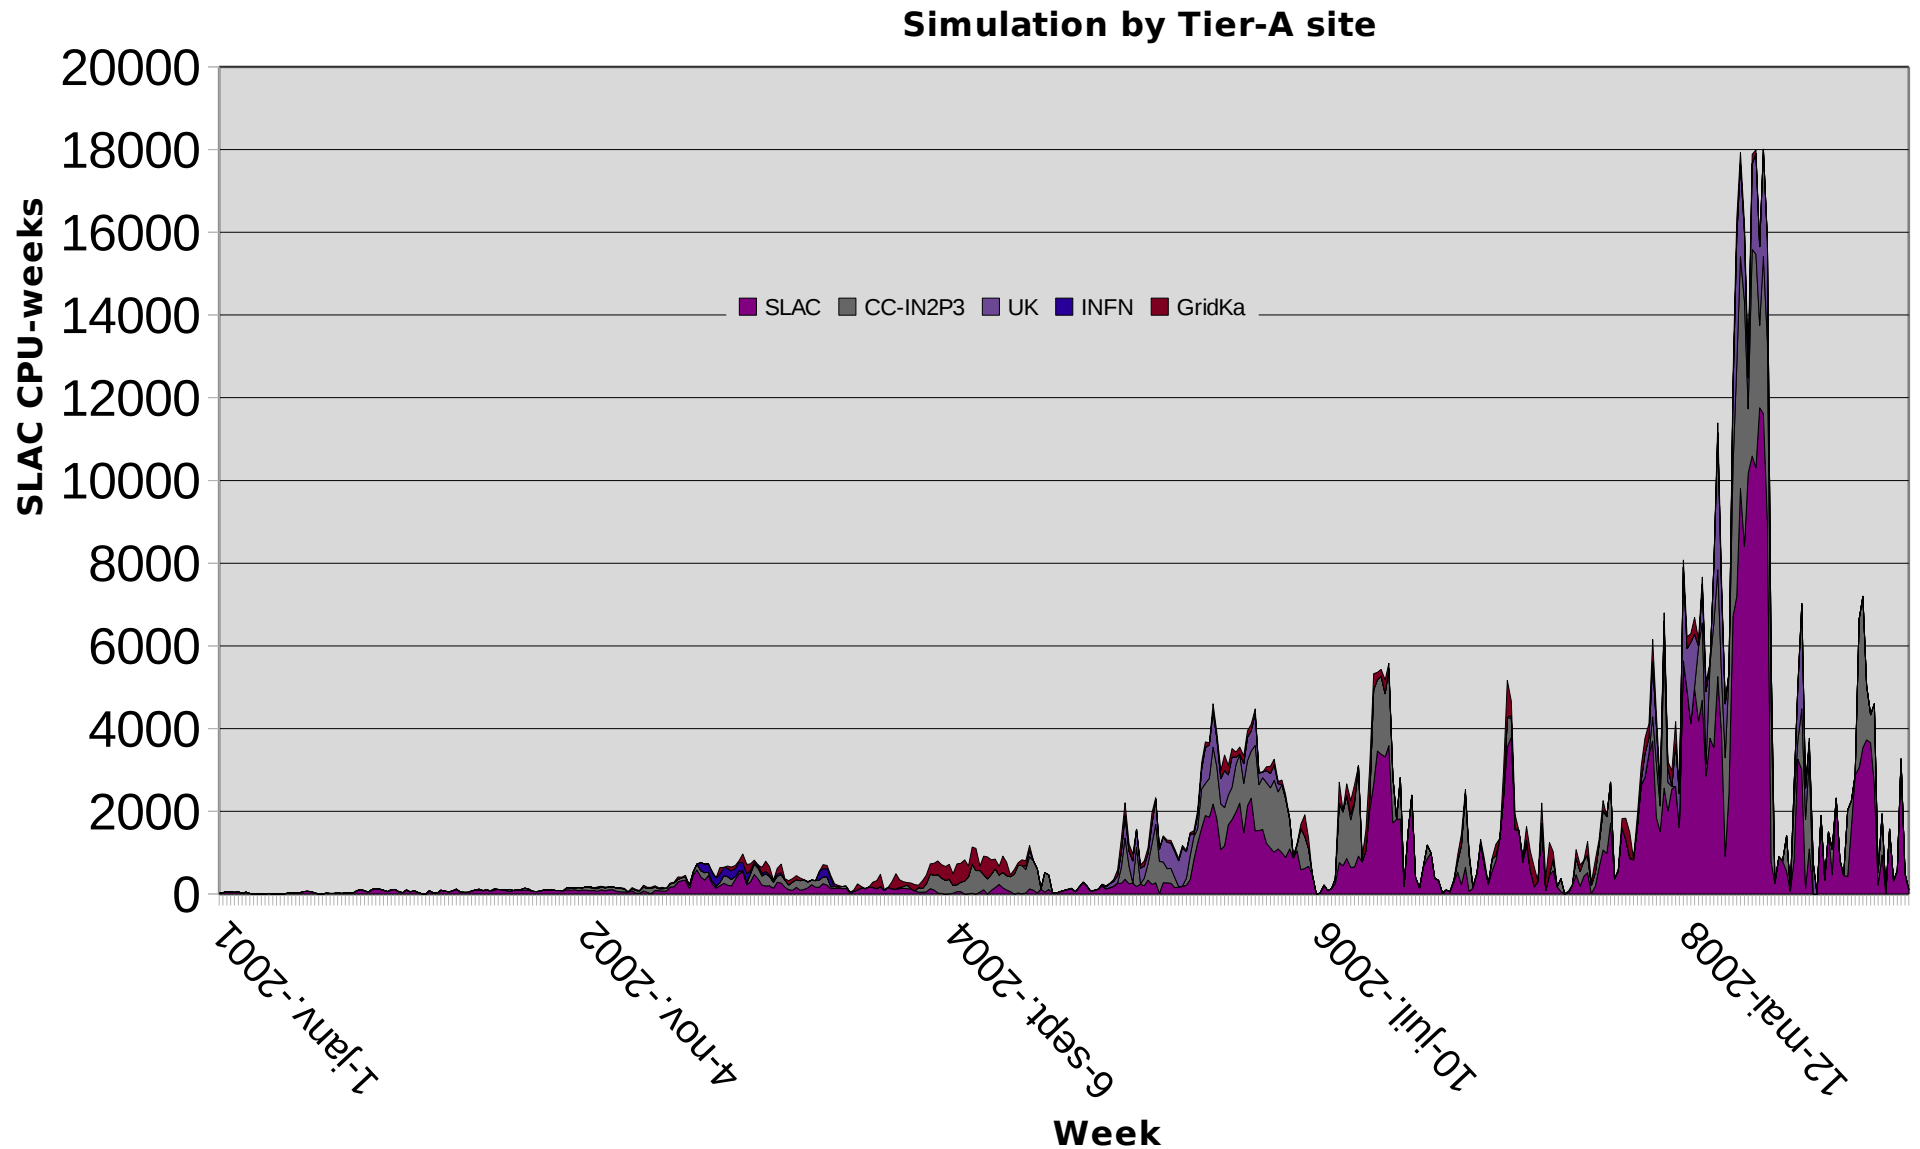
- ◆ CCIN2P3 (simulation++, AWGs++, skim production and backup of Legacy Data)
- ◆ CNAF (simulation & AWGs)
- ◆ GRIDKA (skim production & AWGs)
- ◆ UVIC (simulation (including CLOUDS) & AWG)

Production from these sites plus ~15 universities + RAL (former TierA) PADOVA continues to serve as the raw data backup

BaBar Simulation Production



Contributions to BaBar Simulation Production from the TierA sites



BaBar Distributed Simulation Production

- Sites are validated against equivalent SLAC validation production
 - ~4000 histograms (detector, PID, analysis variables, etc...)
 - Intel vs. Intel, AMD vs. AMD
- Production remains distributed although an eventual collapse back onto SLAC is foreseen.
- Specialized productions mostly done at SLAC for maximum control. However, having multiple sites has been very useful when several varieties of production needed to be done at the same time.

- Reached 850 million events produced per week

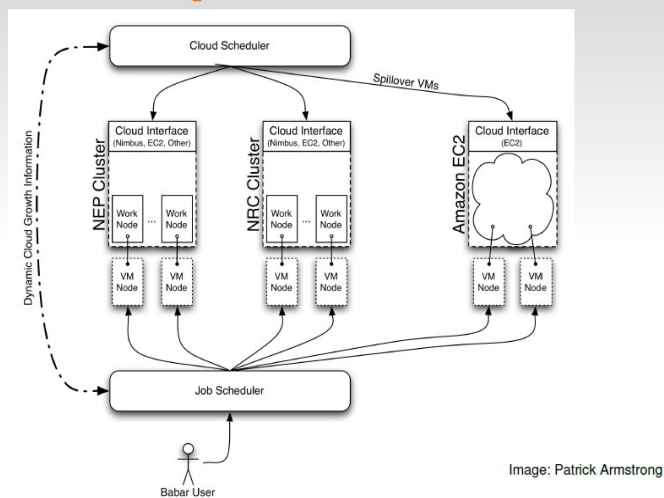
ALL:

Week Beginning	TOTAL	caltech	ccin2p3	cu-boulder	fzk	gridx1	infn	infngrid	osu	ral	slac	slac2	tud	udo	uofl	utd	utenn	uvic2	westgrid
21-9-2008	191970	0	8488	0	0	0	2218	744	832	3720	86896	63608	2648	3808	0	6424	0	12536	48
14-9-2008	846163	6889	168816	20216	4736	0	6312	1144	14016	126496	206264	221994	4288	15208	0	8352	0	33528	7904
7-9-2008	796994	11898	165760	7352	17392	0	3736	7006	9168	77976	217416	219392	2000	14418	0	1176	0	35920	6384

Last Modified: Last Modified: Tue Sep 23 02:18:38 PDT - slightly behind last week (some loss due to validations)

Since last year BaBar simulation production has also been using clouds

The Implementation



May 26, 2009 (DPLTA)


Kyle Fransham - BaBar Virtualization

10



Completed SP11 Jobs

Completed Jobs by Site



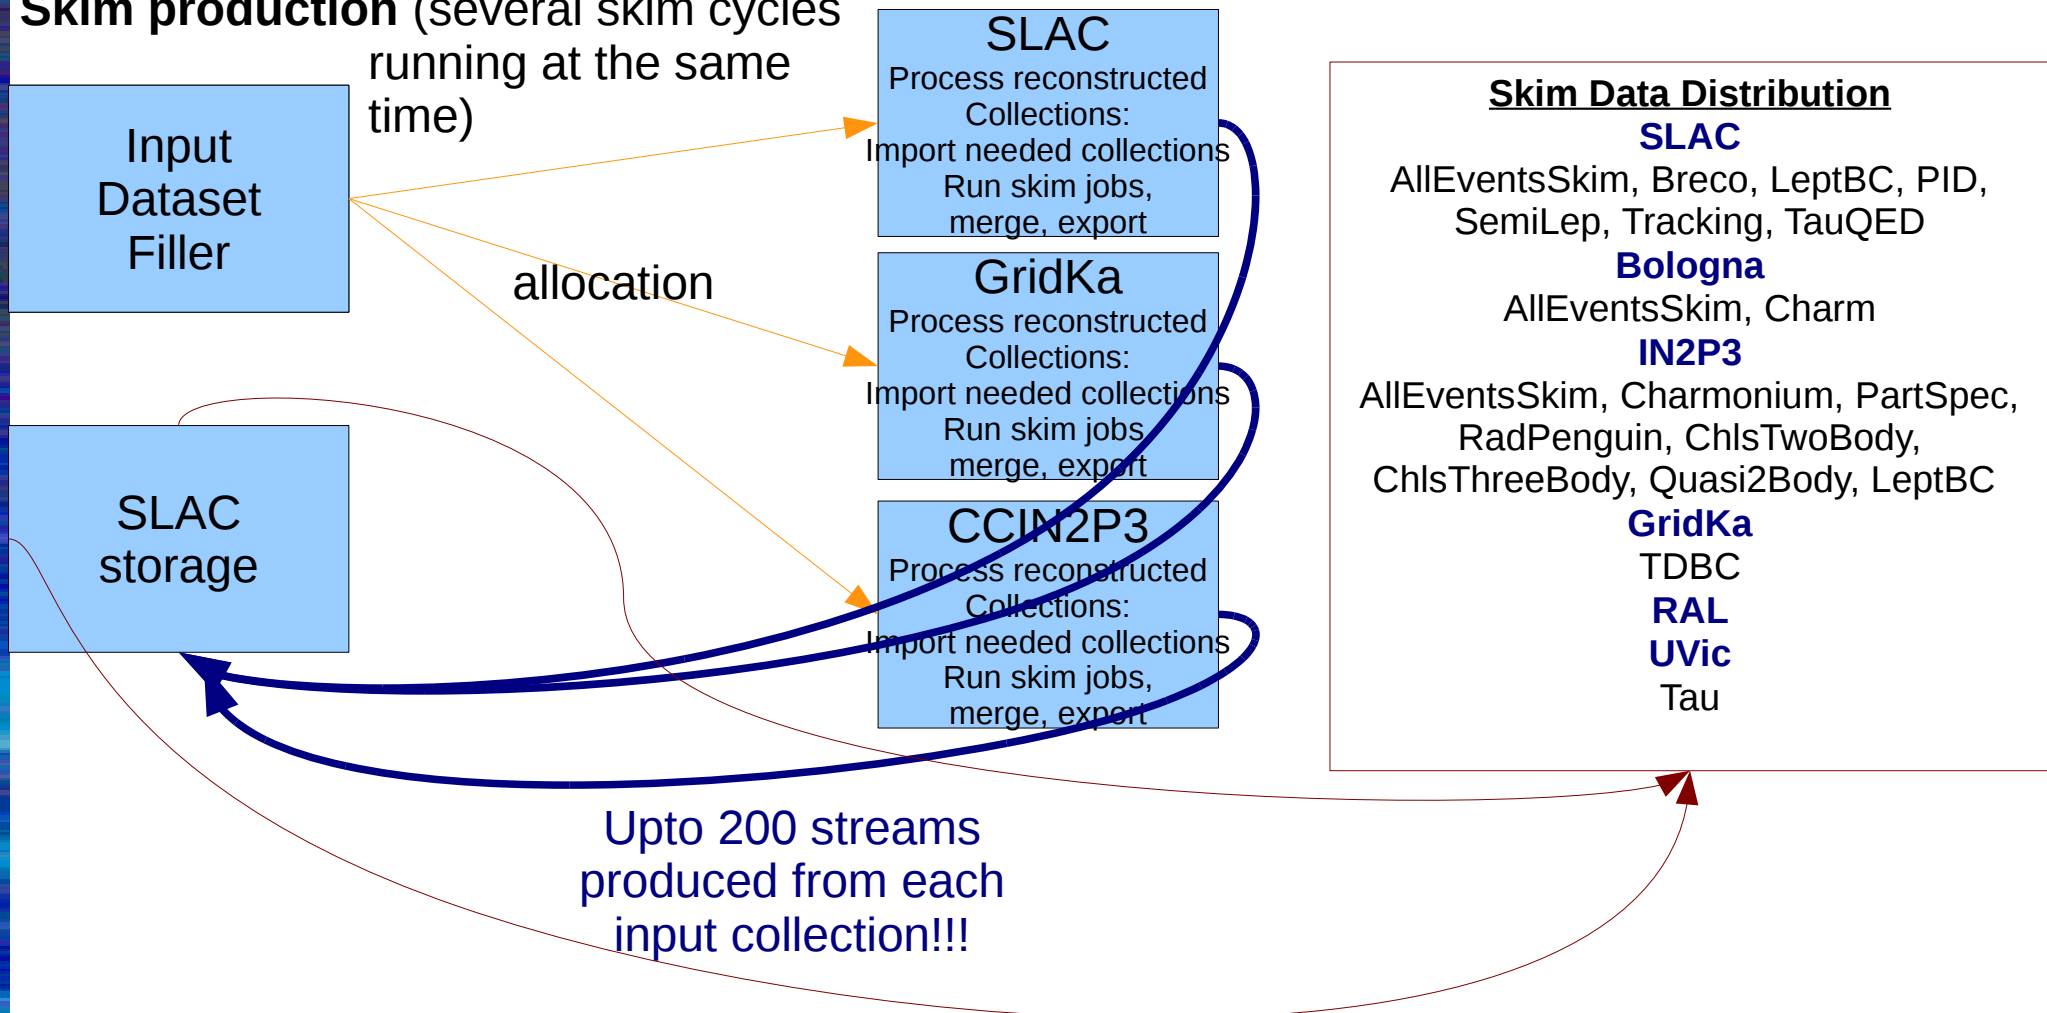
Kyle Fransham



University of Victoria **NRC-CMRC**

7

Skim production (several skim cycles running at the same time)



All streams must be perfectly produced, merged, exported, stored and bookkept

Streams needed by working groups exported to the working groups host site

Amount of data to be skimmed at a site depends on its recent good production rate

Skimmed data can contain all, part or none of the input event data and point to the events in the input event data for everything else.

Detector Data Reconstruction

- 14 farms split between Padova and SLAC

IR2-OPR-BBK
DATA CHECK

SELECT RUN

Run 7: details
 81221 - 81498 (Apr 07)
 80064 - 81220 (Mar 07)
 79009 - 80063 (Feb 07)
 77785 - 79008 (Jan 07)
 77320 - 77784 (Dec 07)

Run 1-7: summary
 Run 7
 Run 6
 Run 5
 Run 4
 Run 3
 Run 2
 Run 1

USER GUIDE

IR2 Summary (help)			
	# of runs	Events	Lumi (1/pb)
All runs	272		
All colliding runs (with)	200	175250087	2352835.7
Good runs	200	175250087	2352835.7
Bad runs	0	0	0.0
Detector Studies	0	0	0.0
Cosmics	15	9019066	0.0
Calibration	52	10353637	0.0
Other	5	646326	0.1

OPR Summary (help)			
	# of runs	Events	Lumi (1/pb)
All processed runs	198	175247571	2352816.0
No-go DETAILS	2	2516	19.7
Good runs	195	174696764	2347232.5
Bad in IR2 but good in OPR (salvaged)	0	0	0.0
Already bad in IR2	0	0	0.0
Last PC & ER not 'done'	0	0	0.0
Last PC not 'done' DETAILS	1	6276	82.5
Last ER not 'done'	0	0	0.0
PC & ER done, DQG Not set	0	0	0.0
PC & ER done, DQG Bad / Flawed	2	544531	5501.0

BBK Summary (help)			
	# of runs	Events	Lumi (1/pb)
AllEvents (good)	195	54983469	2347232.5
BkgTriggers (good)	195	56382	2347232.5
TriggerStream (good)	195	13070493	2347232.5

SKIM Summary (help)			
	# of runs	Events	Lumi (1/pb)
R22d : AllEvents Skim	0	0	0.0

Detector Data Reconstruction

To: OPROper-hn@slac.stanford.edu
From: "marco pappagallo" <marco.pappagallo@ba.infn.it>
Subject: Padova daily summary for 12 June 2008

Hi,
ER5 has been updated.
So we are processing Run2 with ER1,ER2,ER4,ER5,ER6,ER7,ER8 and ER9.
No problem at all today.

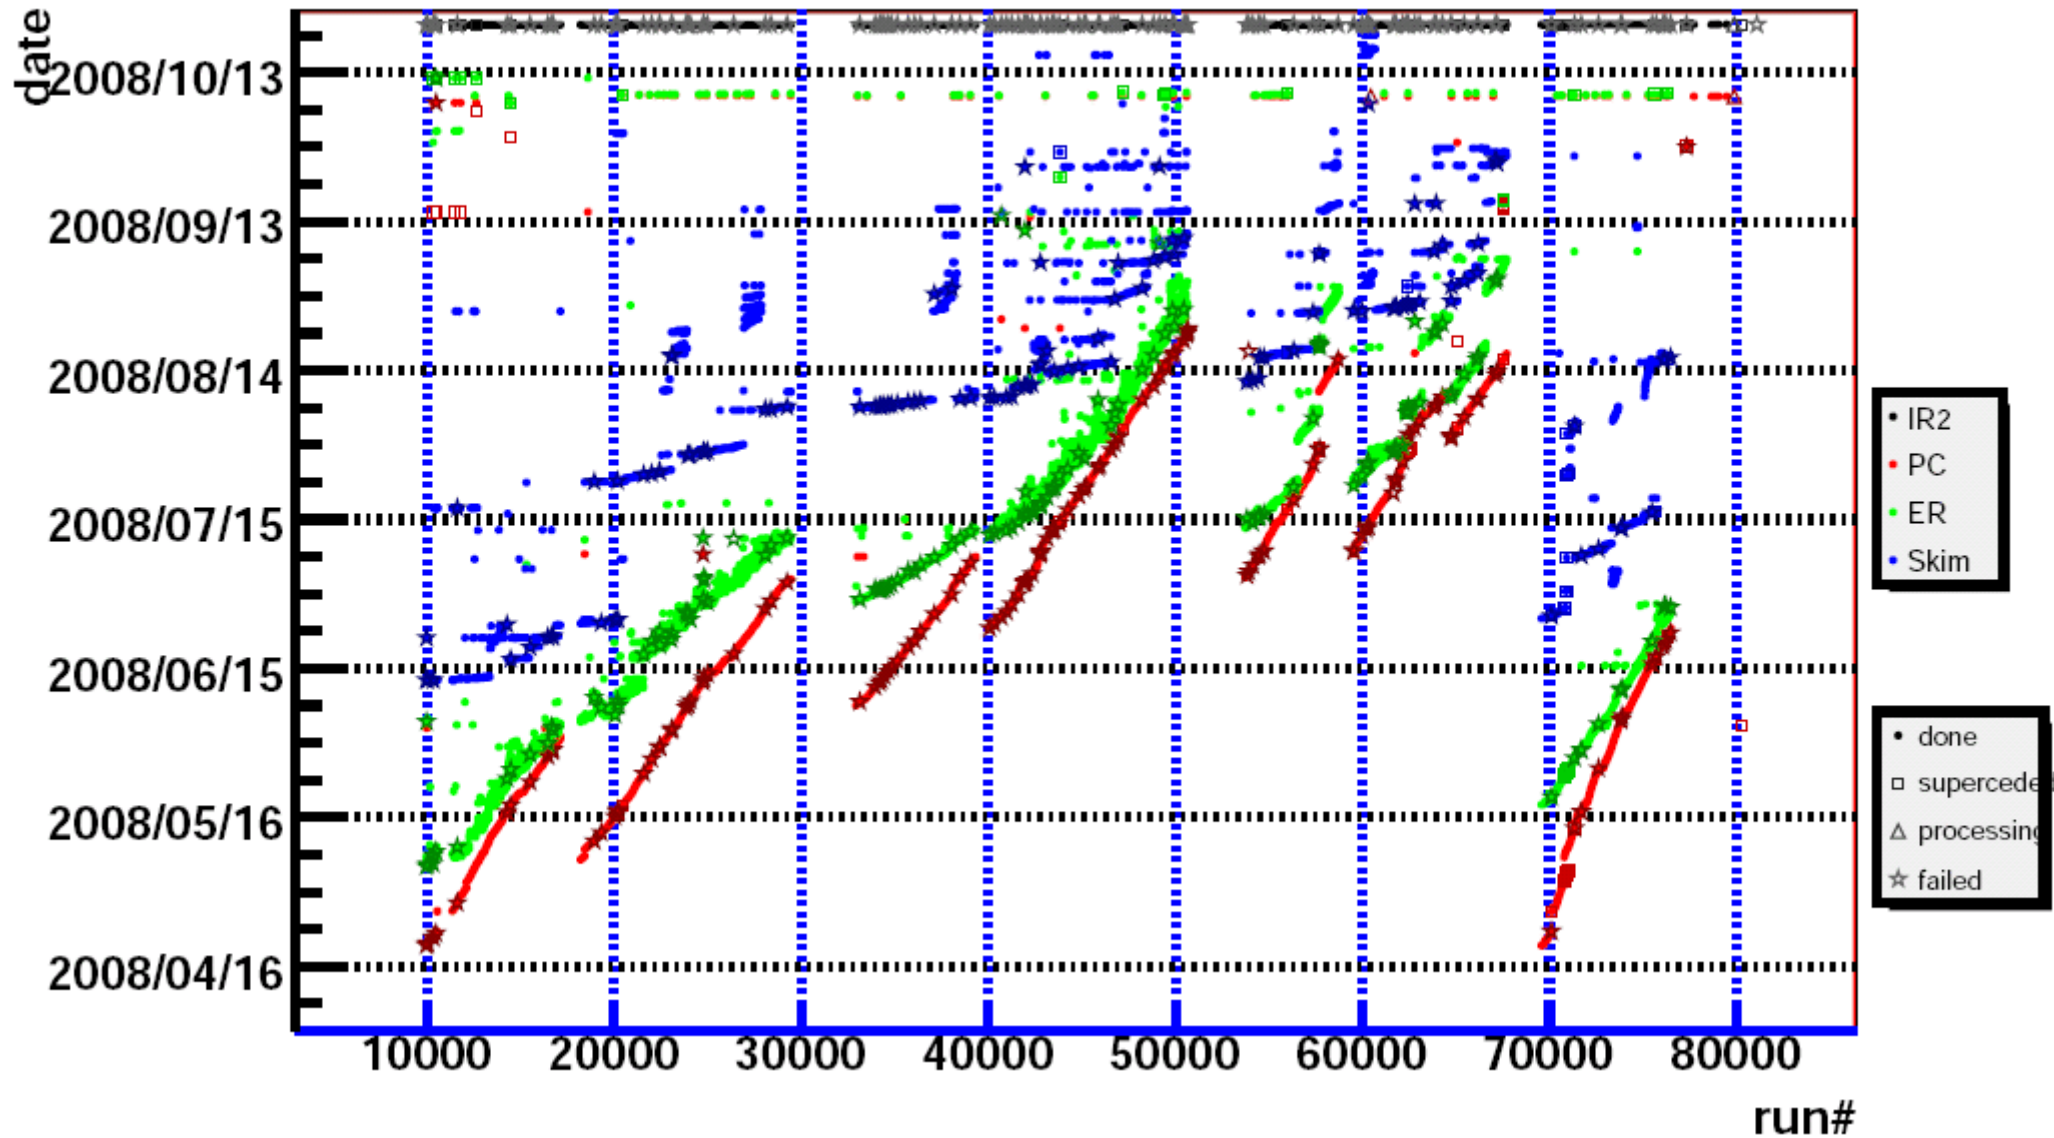
Ciao
Marco

PR Summary for runs which finished processing between
11-Jun-2008 17:00:00 --> 12-Jun-2008 17:00:00

	ER9	ER8	ER7	ER6	ER5	ER4	ER3	ER2	ER1
Proc. Runs:	55	17	12	18	2	23	0	20	18
Proc. Events/10 ⁶ :	26.8	8.4	6.1	9.4	0.4	7.3	0.0	8.8	10.9
Average Rate [ev/s]:	331.0	99.9	109.4	111.8	40.9	85.3	0.0	105.8	129.3

Detector Data Reprocessing Chain

ir2 to pc to er to skim production (Fri Oct 24 09:43:03 2008)



done using 3 PC farms at SLAC, augmented ER capacity at Padova (ER1-->ER9) and 5 ER farms at SLAC (ER10-->ER14) (all++ of BaBar's allocation of Black Box #1)

Following production: Ex: the Simulation Production Database

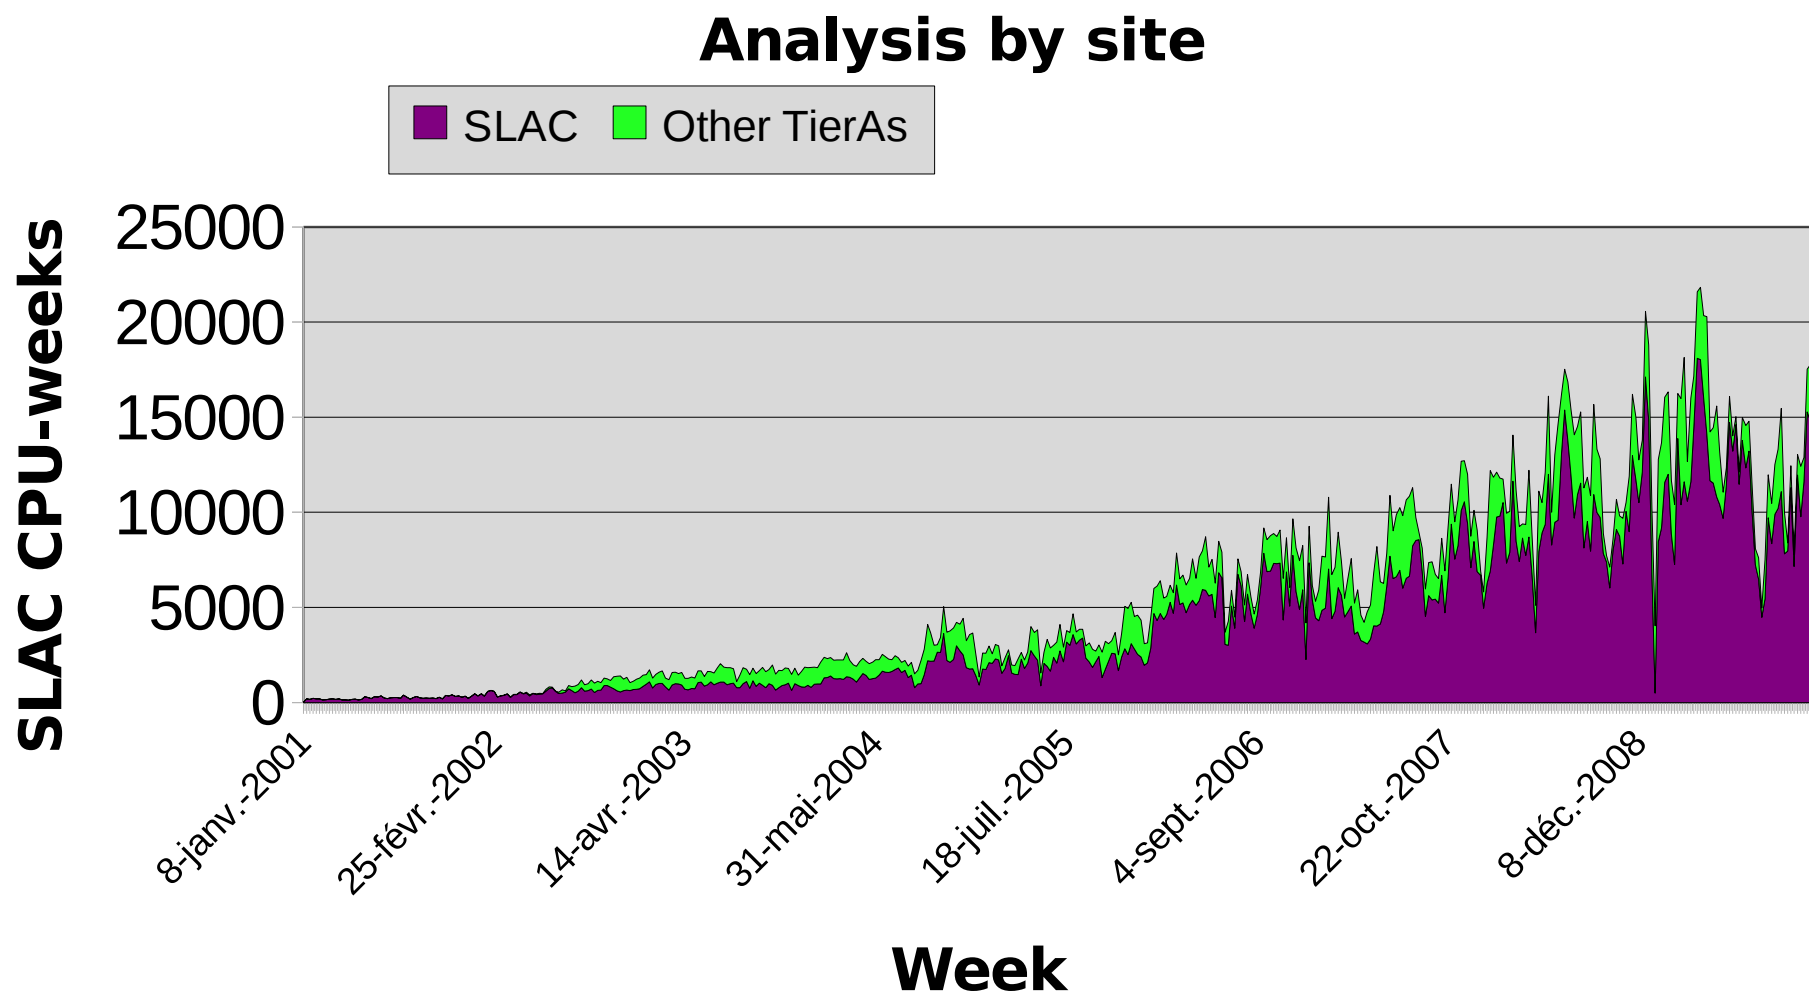
```
$ BbkSPUser --modenum 10200 modenum run prod_req_updated release --release 24\* run_status run_events  
--summary location merge_status
```

MODENUM	#RUN	PROD_REQ_UPDATED	RELEASE	RUN_STATUS	RUN_EVENTS	LOCATION	MERGE_STATUS
10200	20	16-APR-2010	24.3.5d	finished	1000	uvic2	bookkept
10200	16	16-APR-2010	24.3.5d	finished	2000	uvic2	bookkept
10200	10	16-APR-2010	24.3.5d	finished	3000	uvic2	bookkept
10200	13	16-APR-2010	24.3.5d	finished	4000	uvic2	bookkept
10200	10	16-APR-2010	24.3.5d	finished	5000	uvic2	bookkept
10200	11	16-APR-2010	24.3.5d	finished	6000	uvic2	bookkept
10200	6	16-APR-2010	24.3.5d	finished	7000	uvic2	bookkept
10200	22	16-APR-2010	24.3.5d	finished	8000	uvic2	bookkept

Making Data Available to Users

Datasets tagged for physics analysis use only
after milestone of validated production level
achieved

Analysis Activity



Analysis Working Groups distributed among the TierA sites hosting needed skimmed data.

Backup of the BaBar data

- Raw data at Padova
- Simulation data at CCIN2P3
- BaBar legacy data recently migrated to new media at SLAC
- Copy of all BaBar legacy data being copied to CCIN2P3



Summary

- BaBar data management still involves distribution of data to each of the TierA sites for analysis work and production of data at several sites
- In the future the fraction of analysis and production work done at SLAC will increase as remote sites start to focus on new projects
- Starting in 2012, all new analysis and most production work will be handled by the archival system with input data being staged in from HPSS if not already on the system.

Origin of BaBar's Distributed Computing

- Funding that would have been needed to satisfy the computing needs for doing it all at SLAC was too much.
- Needed resources made available at the participating institutions
- Developed mechanisms to coordinate and combine these resources to satisfy the needs.
- *BaBar would not have succeeded without this.*

-
- Responsibility for core computing (CPU & disk) provision divided as ~2/3 SLAC, ~1/3 non-US
 - Tier-A countries delivered their shares largely on an in-kind basis at their Tier-A sites, recognized at 50% of nominal value
 - BaBar operating and online computing costs split ~50-50%

Data Files database (BookKeeping)

```
$ BbkUser --dbname bbkr24 --run 77686 --summary --dse_type PR collection events_in events gbytes dse_status
COLLECTION                                     +EVENTS_IN +EVENTS +GBYTES DSE_STATUS
=====
/store/PR/R24/AllEvents/0007/76/24.1.3c/AllEvents_00077686_24.1.3cV00      225598 106251 1.0 2
/store/PR/R24/AllEvents/0007/76/24.2.1f/AllEvents_00077686_24.2.1fV00      225598 108574 1.0 2
/store/PR/R24/AllEvents/0007/76/24.3.5d/AllEvents_00077686_24.3.5dV00      225598 125897 1.1 1
/store/PR/R24/AllEvents/0007/76/24.5.7/AllEvents_00077686_24.5.7V00        225598 125896 1.0 0
/store/PR/R24/BGFilterSkim/0007/76/24.2.1n/BGFilterSkim_00077686_24.2.1nV00 225598 12462 0.1 2
/store/PR/R24/BkgTriggers/0007/76/24.1.3c/BkgTriggers_00077686_24.1.3cV00  225598 129 0.0 2
/store/PR/R24/BkgTriggers/0007/76/24.2.1f/BkgTriggers_00077686_24.2.1fV00  225598 129 0.0 2
/store/PR/R24/BkgTriggers/0007/76/24.3.5d/BkgTriggers_00077686_24.3.5dV00  225598 129 0.0 1
/store/PR/R24/BkgTriggers/0007/76/24.5.7/BkgTriggers_00077686_24.5.7V00    225598 129 0.0 0
/store/PR/R24/TriggerStream/0007/76/24.1.3c/TriggerStream_00077686_24.1.3cV00 225598 31817 0.3 2
/store/PR/R24/TriggerStream/0007/76/24.2.1f/TriggerStream_00077686_24.2.1fV00 225598 31984 0.3 2
/store/PR/R24/TriggerStream/0007/76/24.3.5d/TriggerStream_00077686_24.3.5dV00 225598 46820 0.4 1
/store/PR/R24/TriggerStream/0007/76/24.5.7/TriggerStream_00077686_24.5.7V00  225598 46856 0.3 0
=====
Totals                                     2932774 637073 5.3
```

Future of BaBar Distributed Computing

- BaBar computing to be concentrated in the archival system
BUT
 - With virtualization being used, the ability to use CLOUDs has naturally evolved
 - Archival system is designed to be reproducible at other institutions
 - Strong support will continue from some of the TierA sites

