



TM & © Nelvana

Developments in BaBar simulation

Douglas A. Smith, Peter Kim, Wilko Kroeger
Stanford Linear Accelerator Center

CHEP 2007

Sep. 5, 2007, Victoria, BC, Canada


For the BaBar computing group



TM & © Nelvana

Simulation Production in BaBar


- In current HEP experiments there is a need for a large amount of simulated events to be produced. (In BaBar 3 times data for B-pair, equal to data for continuum.)
- The size and scope of meeting this need now lies beyond the ability of any one computing site: ave. production cycle is ~2M jobs, ~1000 cpu years, ~100TB of data; needed in less than one year for use.
- In BaBar this is now a distributed computing problem, to ~20 sites.
- See D. Smith, *et al.*, Talk 339, CHEP 2004, D. Smith, *et al.*, Talk 299, CHEP 2006 -- for more details.



TM & © Nelvana

BaBar history


- The first computing model included that data storage would be into Objectivity databases. Was put into use at the beginning around 1998.
- As the experiment continued this was shown to not be scalable to needed effort, a computing model 2 (CM2) plan was started.
- The resulting CM2 included that data storage will be in root files. Developed through 2003 -- put into use in early 2004. See P. Elmer, Talk 502, CHEP 2004 – for more details.
- At the time, data was put in root files, but the conditions database was kept in Objectivity.
- Was considered too much effort at the time to completely remove Objectivity use.



TM & © Nelvana

Root-based condition data


- After CM2 was in production, people worked on migration of condition data. Condition data is produced into a MySQL database, then saved into root files for distribution and use.
- These condition data root files are saved in the BaBar event store, and info. kept in the Bookkeeping, with a dataset, just like any other data.
- For simulation jobs, there is no difference. Handled by startup controls, once job knows which tech. to use to get condition data, the rest of the job runs the same.
- (Un)Fortunately the jobs run the same (cpu use, data rate unchanged).
- See I. Gaponenko *et al.*, talk 352, CHEP 2006 for more details.



TM & © Nelvana

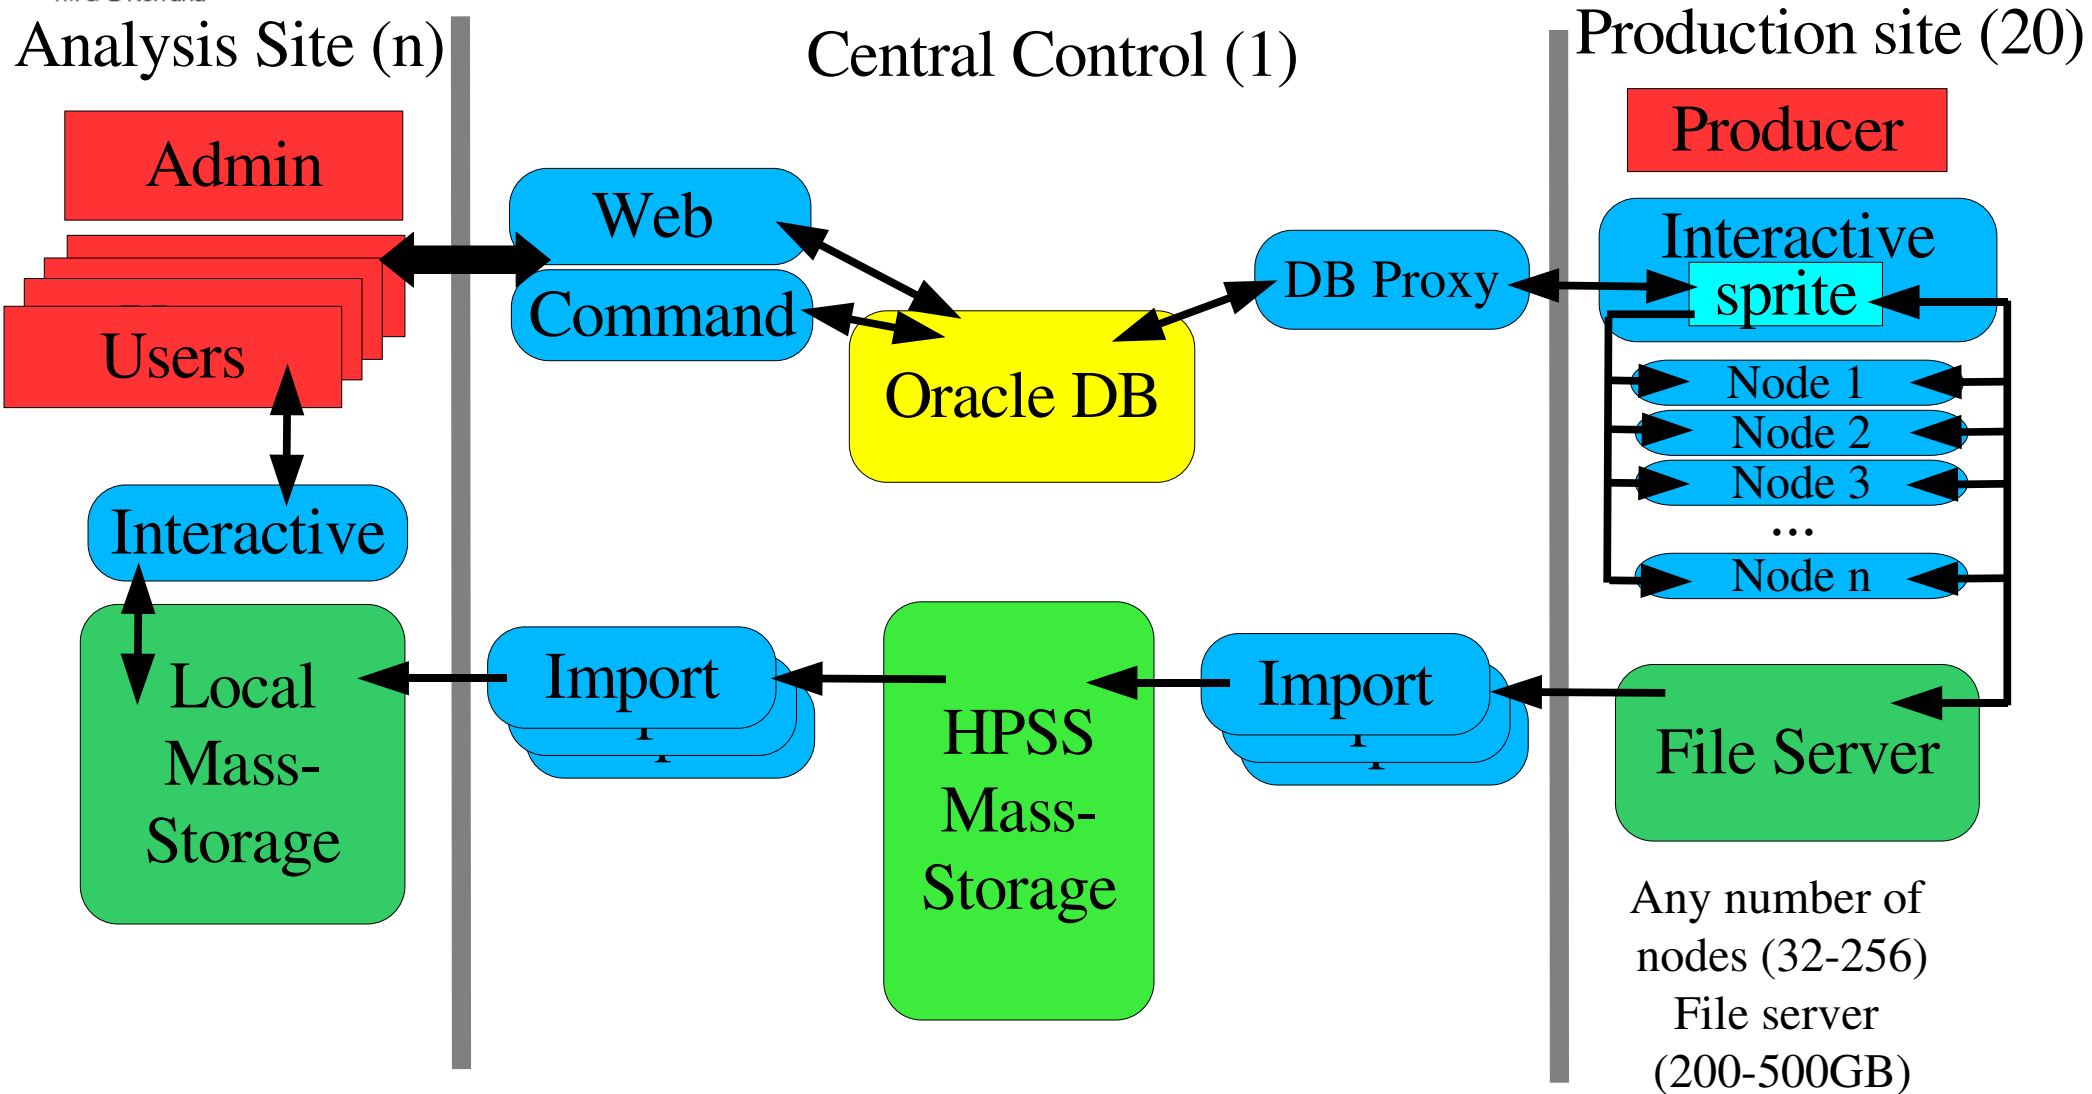
Simu. Production (SP) cycles


- Production in BaBar divided into cycles, and numbered.
- SP5 was last CM1 production cycle, was converted in CM2.
- CM2 production cycles are now SP6, SP8 and SP9 (SP7 was not done due to production and release schedules).
- Production cycles are for certain data run cycles (6 so far in BaBar).
SP5 used for Run 1, 2 and 3; SP6 used for Run 4; SP8 used for Run 1-5; SP9 used for Run 6.
- SP9 first to not use Objectivity databases.
- Luminosity of each data run cycle determines size of requests for simulation cycle.



TM & © Nelvana

Cartoon of production






TM & © Nelvana

Production Numbers

Cycle	Jobs (M)	Events (B)	CPU time (y)	Data (TB)	Files (k)
SP5	1.5	2.0	415	28.6	81
SP6	1.6	2.9	987	44.8	65
SP8	6.2	12.1	2060	173.6	390
SP9	0.4	1.3	250	18.1	45

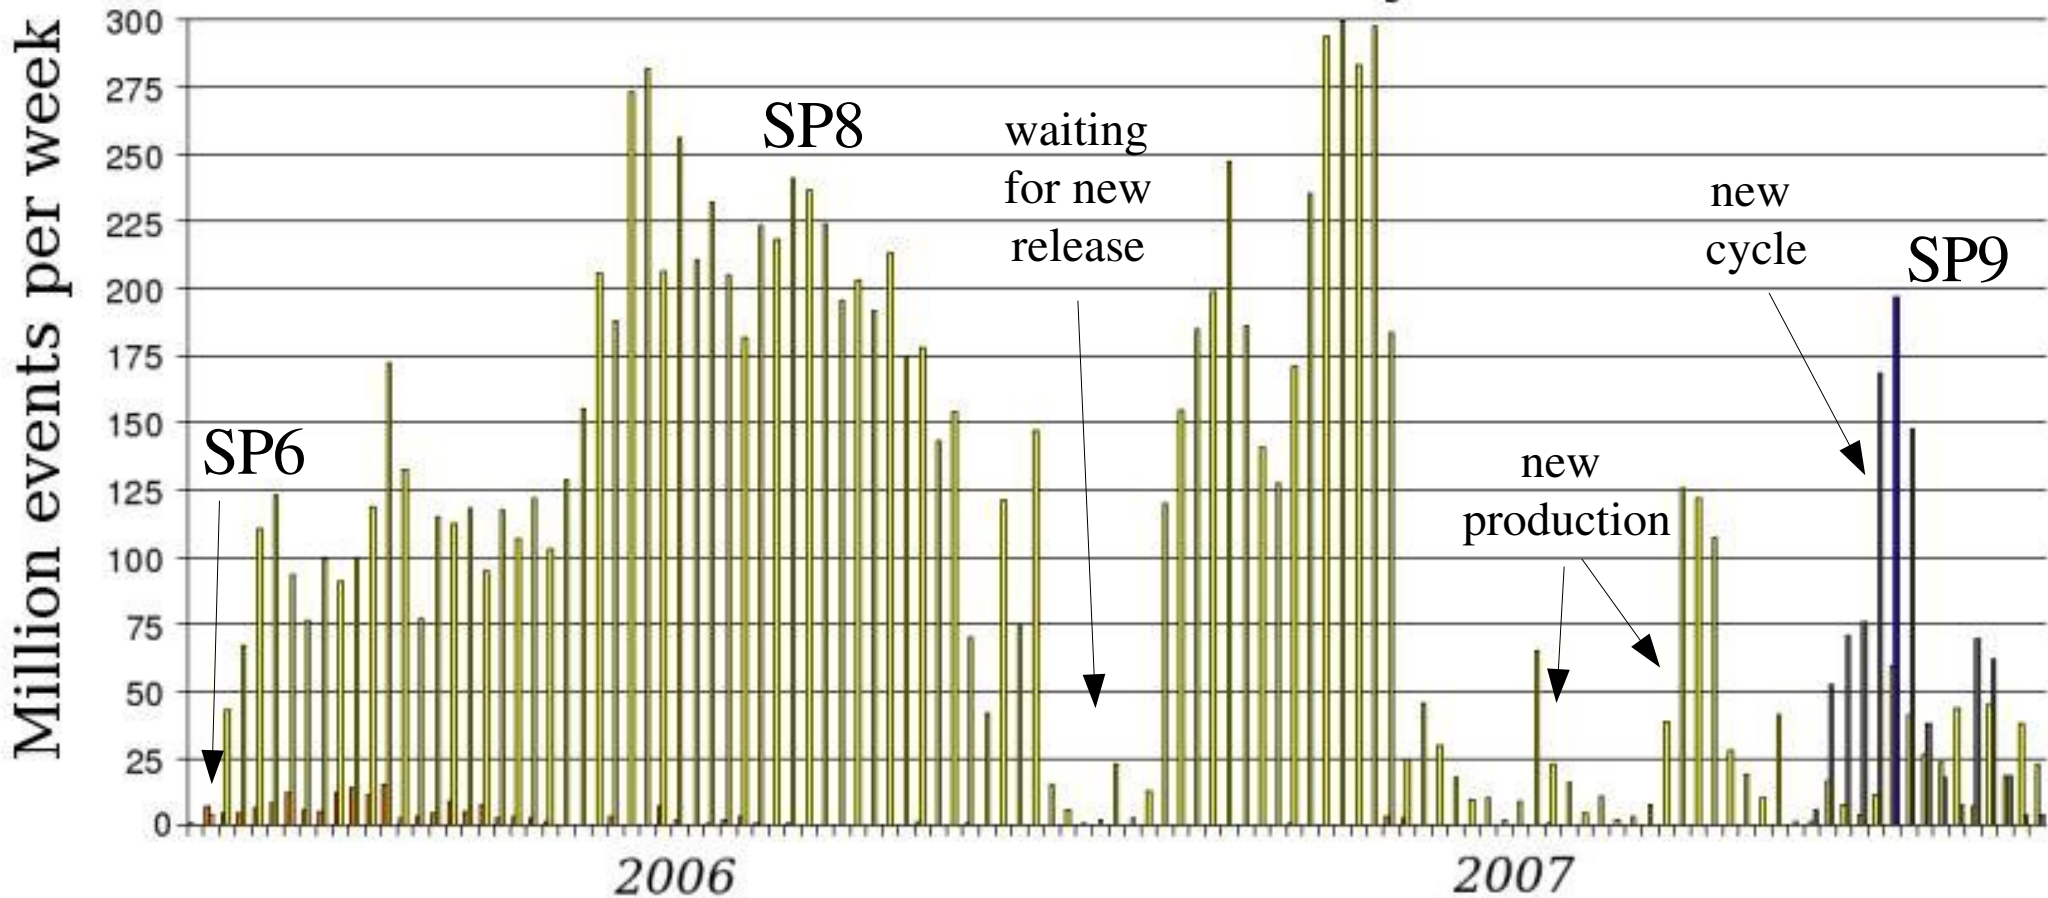
- A couple nice things to notice in these numbers
 - Numbers are fairly large, but not so bad if you have ~2000 cpus, and 200TB storage.
 - From SP6 to SP8 amount of CPU time reduced, even though significantly more events produced. Due to farm upgrades to better cpus (and about 20% due to software changes).




TM & © Nelvana

Production by week

Simulation Production by week





TM & © Nelvana

Results of changes


- Simplified system
 - no objectivity setup now, less to do.
 - Conditions data just like other data, stored as root files in bookkeeping.
 - Distributed the same as other BaBar data.
- Smaller size, less storage, but not a huge effect, about half the size of the Objectivity database. Now needs about 20-30GB on disk.
- Needed only small changes to production system, and had no adverse effect on production.
- BaBar production is finally free of Objectivity use, and things are working well.



TM & © Nelvana

Ideas of future

- Package conditions for storage on worker nodes
 - Conditions data now served from xrootd event-by-event (which works well), but for some sites it would be interesting to try shipping cond. data with job to save on network traffic.
- Split conditions into subsets of data.
 - Each production job will only access a fraction of the cond. data in database. Could split up database into ~60 smaller databases, a job would only need one of these. Each database could be ~500MB in size, and then be shipped with job to worker node.



TM & © Nelvana

Production status

- Things are working well for BaBar simu. production.
- SP8 was a complete reprocessing, producing more data than before in BaBar. But the system worked well to scale to the needs, resulting in 12B events, in a little over a year.
- Transition to Objectivity free running was painless, after years of painful development, and SP9 produced needed data for analysis of the latest BaBar run cycle, and on time for results shown this summer.
- Life without a database is much better.