

SAM Database Usage

**Adam Lyon
CD/REX/Ops (DØ)**

**DØ DB Taking Stock Meeting
June 11, 2007**

DB Status

- ◆ Smooth operations, thanks to DB group and DBAs!
- ◆ Several problems worked out:
 - ❖ DB corruption 12/06-1/07; fixed by moving data off of problematic tray in disk array
 - ❖ Periodic severe performance degradation in 10g
 - Oracle analysis of tables caused favoring of full table scans over indices
 - Particularly bad for data_files, where many files start with the same character sequence (tar on this issue?)
 - Fixed by changing how reanalysis is performed
 - We now also set an Oracle Session parameter in DB server to heavily deweight the use of full table scans

DB Status 2

- ◆ D0ora2 performance problems
 - ❖ Intermittent 100% CPU usage
 - ❖ Increased memory on machine
 - ❖ We moved all v7 DB servers to linux machines
 - ❖ Complete review of cron-jobs on d0ora2, significantly reduced the number and frequency of jobs

Sam Status

- ◆ SAM runs well at DØ, CDF, and MINOS
- ◆ No major development for SAM-DH planned
- ◆ No schema changes for SAM-DH planned for future

- ◆ v5-v7 transition nearly complete
 - ❖ Some holdouts are hard to drop

- ◆ Helping to test SAN

Scope of the system

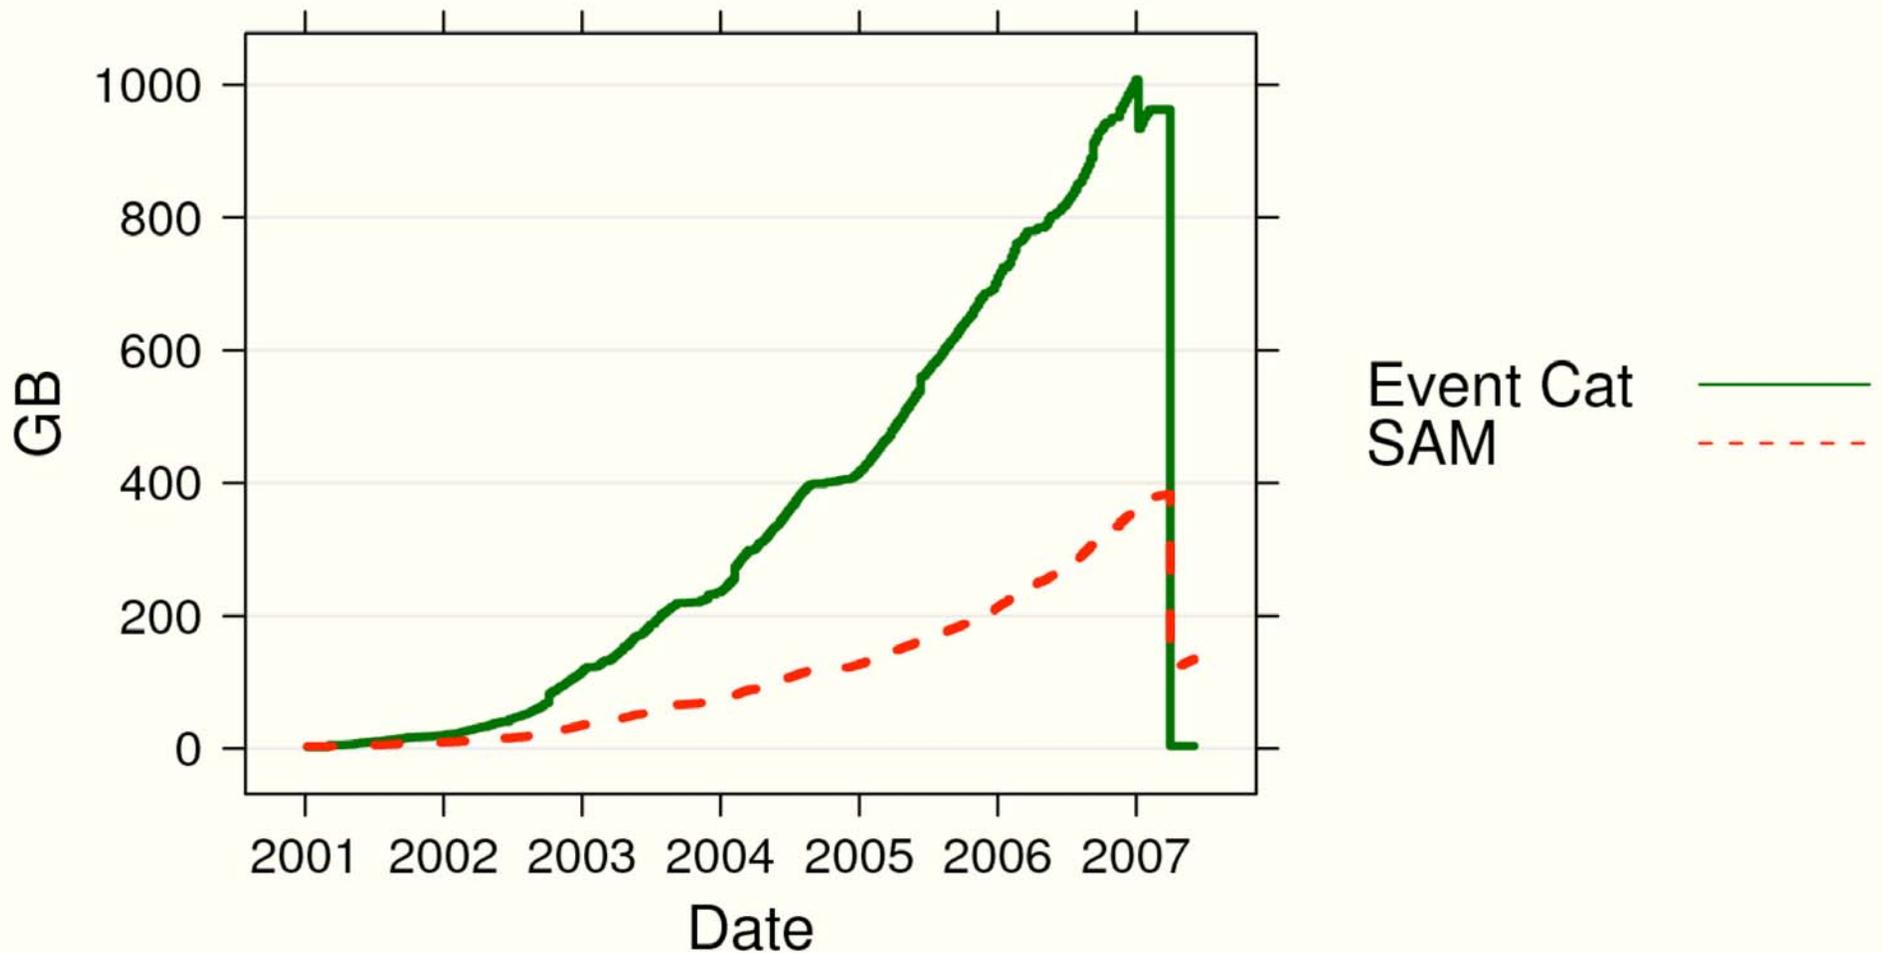
	7/7/06	6/9/07	Table
# files	10,626,703	17,332,216	data_files
# datasets	363,570	487,740	project_definitions
# snapshots	433,945	540,787	project_snapshots
# projects	564,608	835,492	analysis_projects
# processes	3,690,413	5,494,413	analysis_processes
# events	2,392,329,503	0	events

Reducing DB Storage

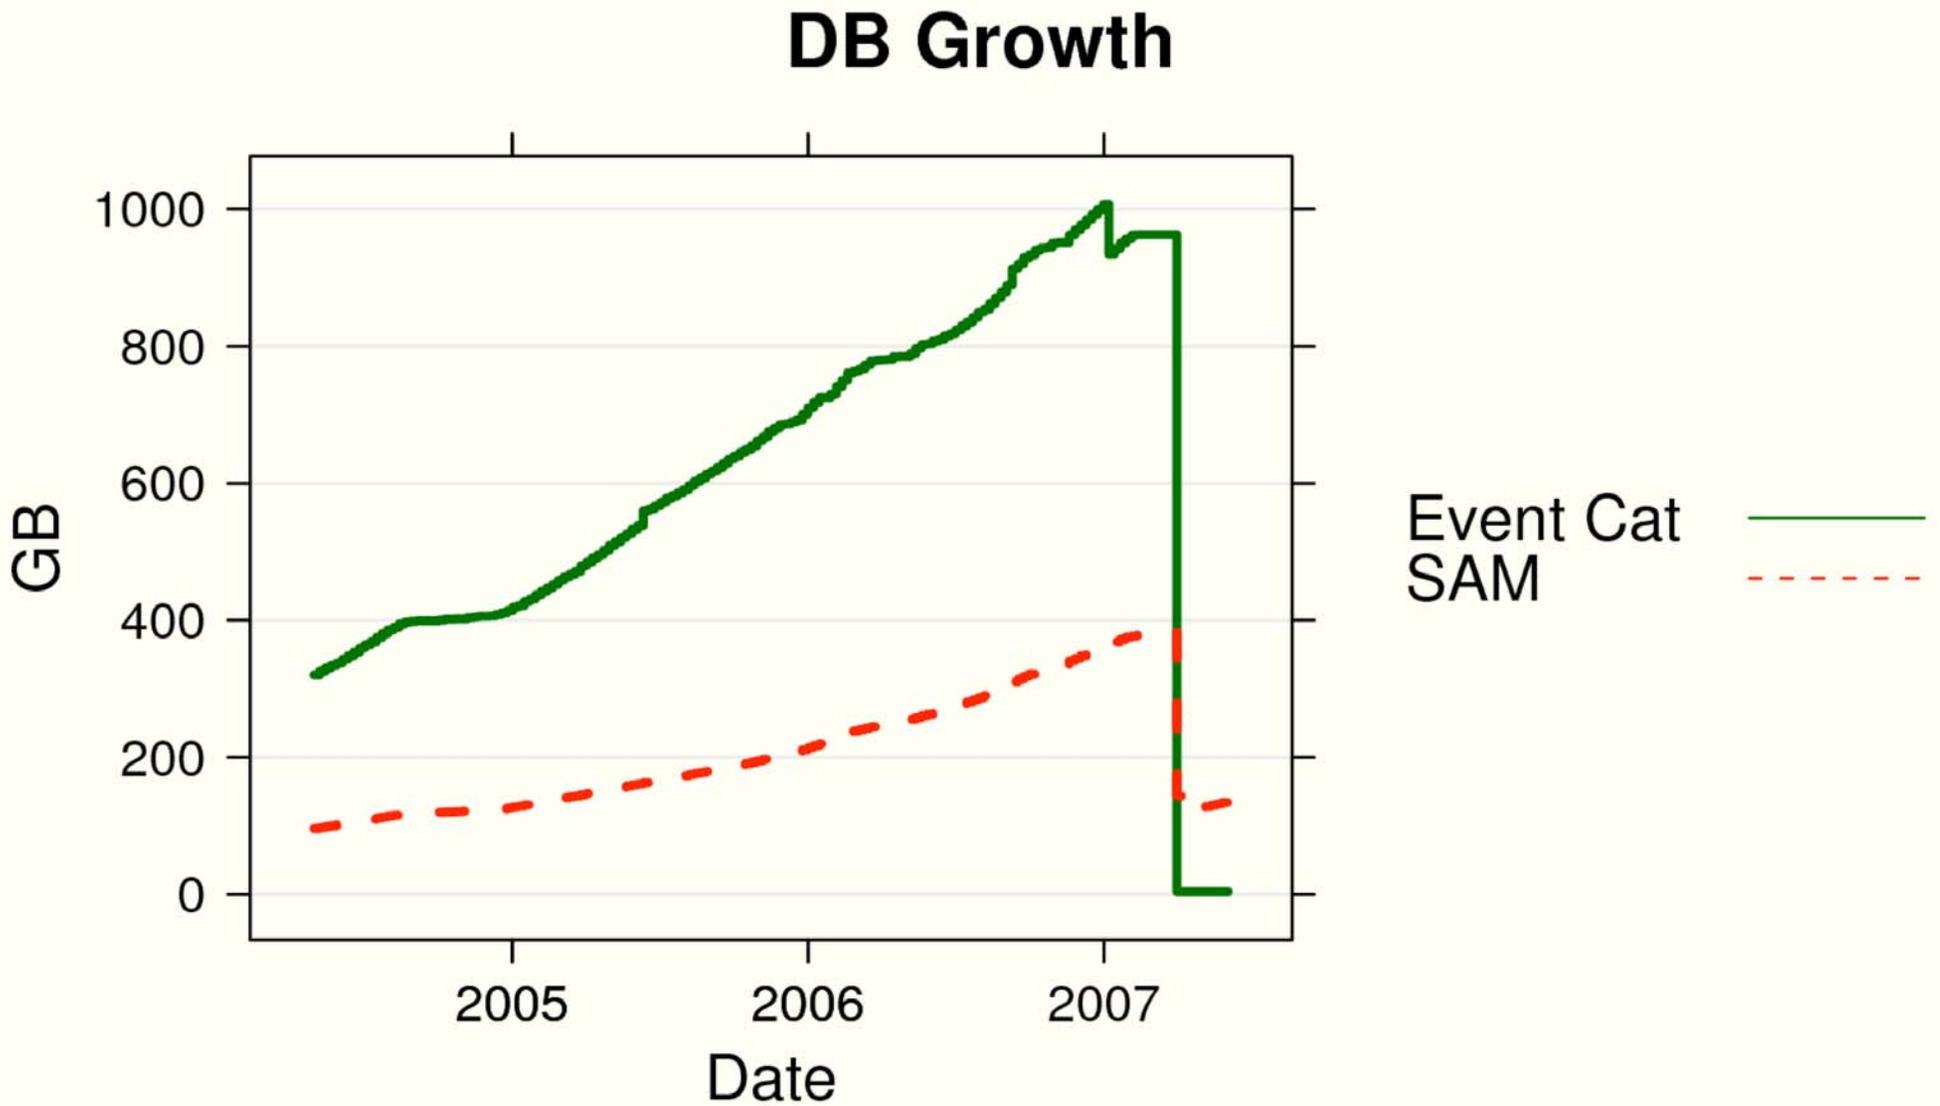
- ◆ Remove the events catalog
- ◆ Remove old data in cached_files (history of cached files since the beginning of time)
 - ❖ Keeping data from 1/2006 onward
- ◆ DB Size
 - ❖ **6/9/07: 138 GB total SAM DB (73% indices)**
 - ❖ 7/2/06: 282 GB total SAM DB (92% indices)
 - ❖ 7/2/06: 1.1 TB Total DB size

DB Size

DB Growth

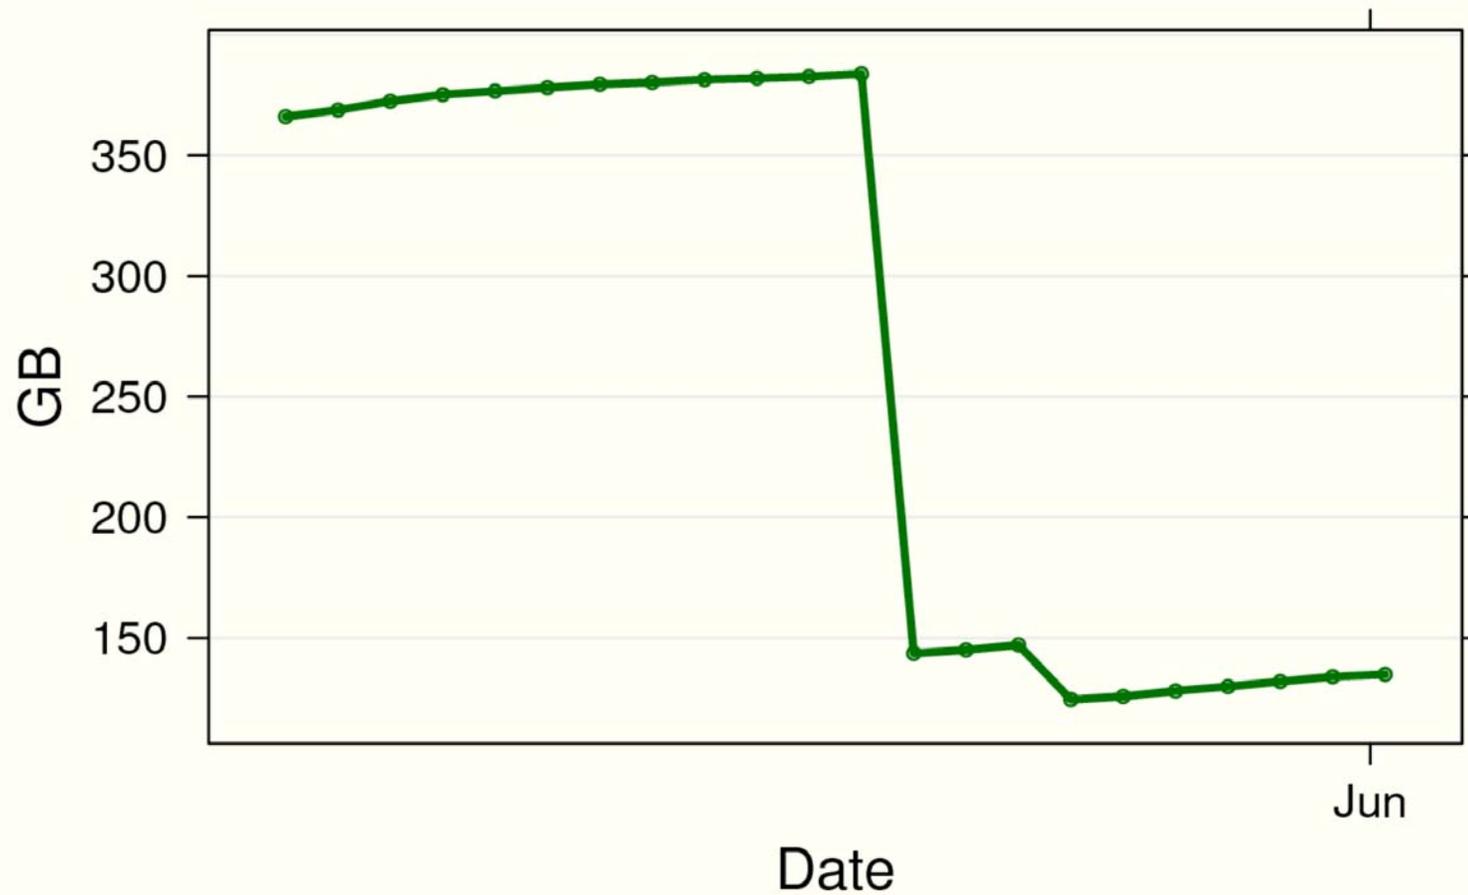


DB Size



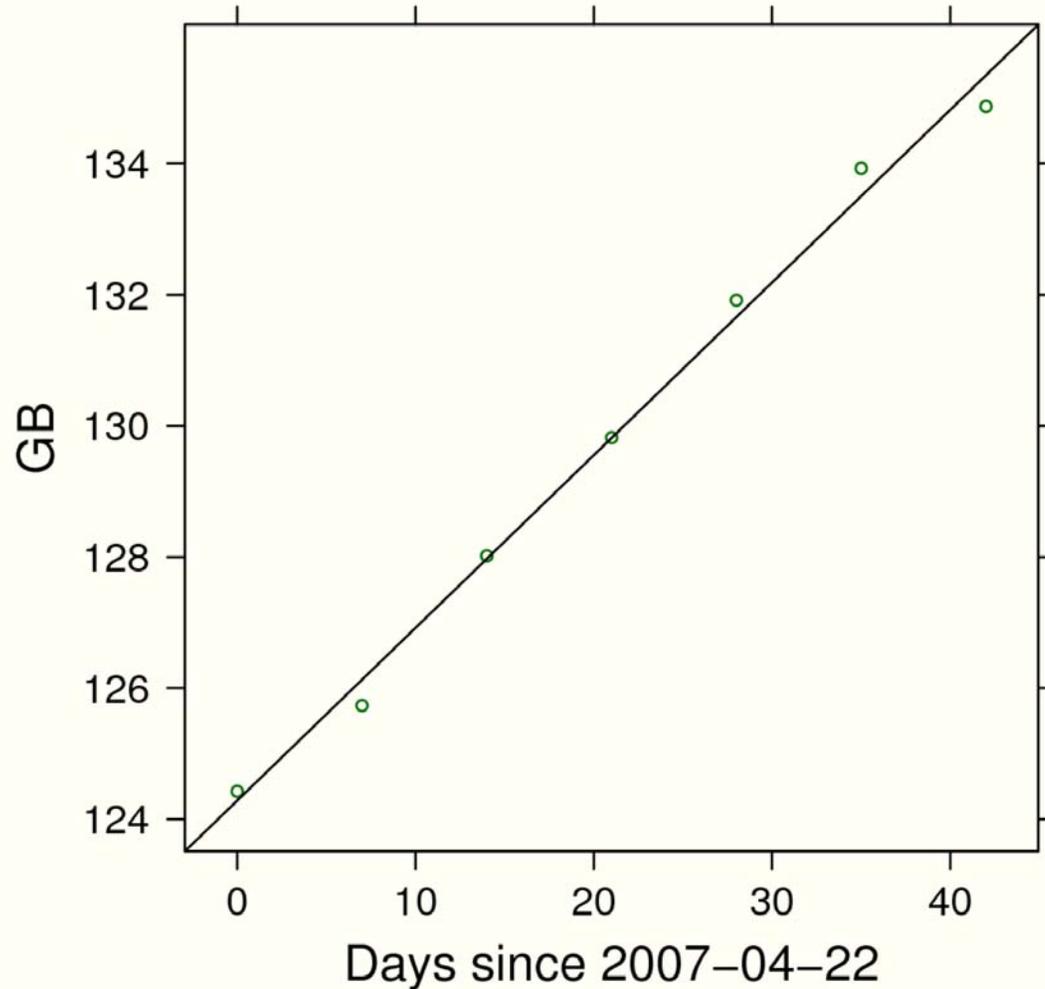
DB Size (SAM Only from 1/1/07)

DB Growth



Growth Rate

DB Growth Rate



◆ 0.26 GB/day added to DB

Predictions

- ◆ 94 GB/yr (assumes no shutdowns)

- ◆ 6/2008: 220 GB
- ◆ 6/2009: 316 GB
- ◆ 6/2010: 412 GB (assumes still taking data)



