



AMI

AMI replication

Solveig Albrand



23/04/2008



S.A.

1

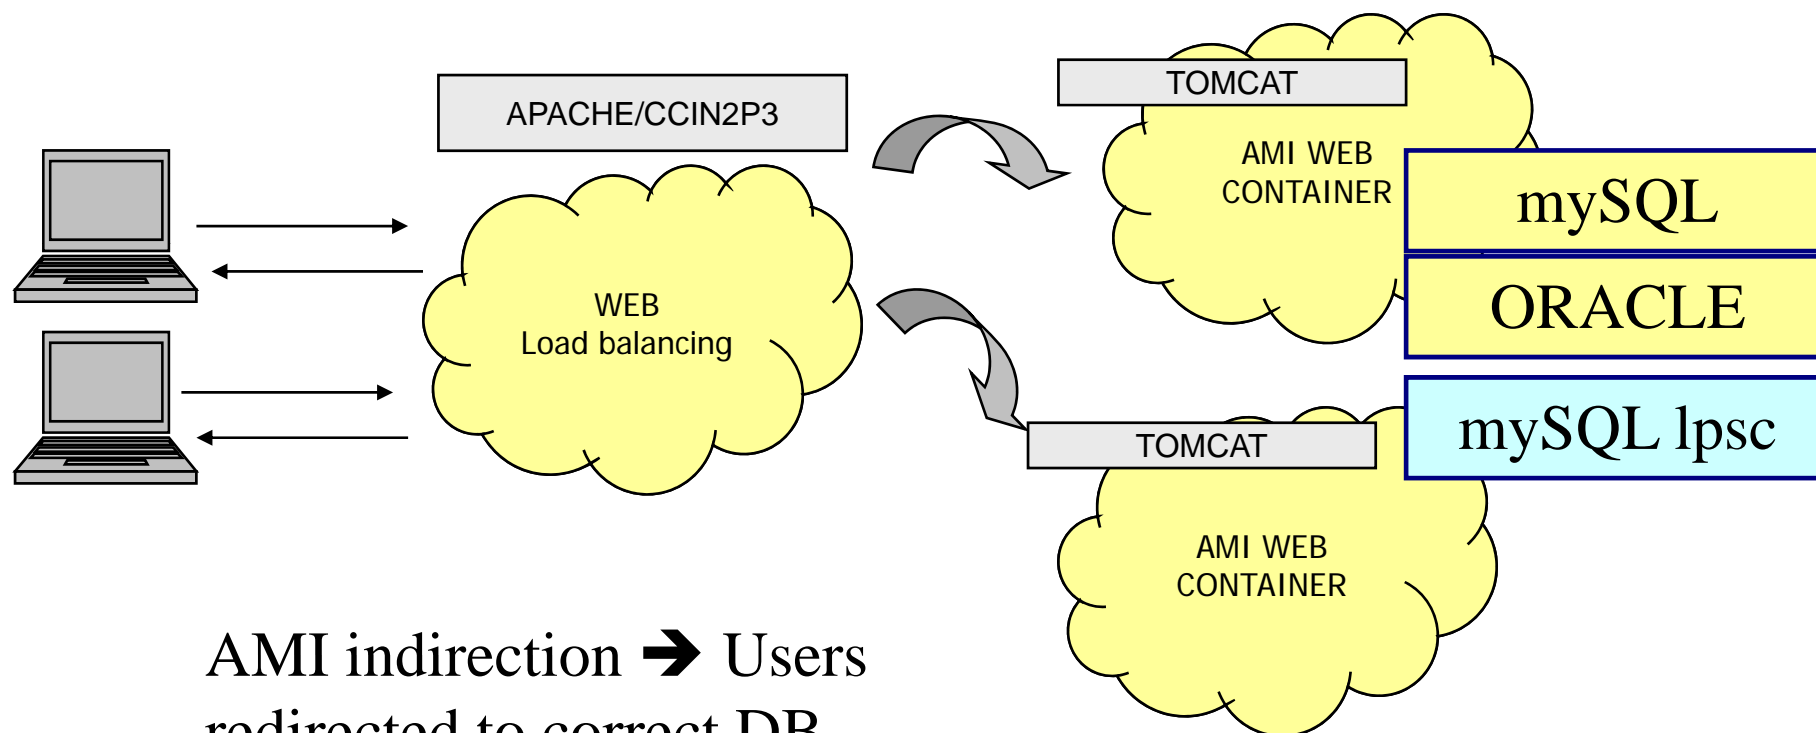


Current deployment of AMI and TC

Component	Location	
Dataset catalogues since 2007; Central dataset catalogues REPLICATION SOON	ORACLE @CCIN2P3 Lyon	data08, mc08,... reference tables + global dataset
Legacy dataset catalogues	mySQL@LPSC Grenoble	csc, mc11,....
Tag Collector	mySQL@CCIN2P3 Lyon	For reasons of performance
AMI Router. Connection indirection and user rights. REPLICATION LATER	mySQL@CCIN2P3 Lyon	For TC Performance, but should go to ORACLE in 2008



Servers: Load balancing / Fail over





Replication

- Unusual! Replication from Tier 1 to Tier 0
- Mechanism was tested in 2007. Currently waiting for new tests on “real” data.
- Only the datasets in catalogues started since Jan 2007 will be replicated at first.
- This means TC will NOT be replicated, and at first no server will be needed at CERN.



Replication

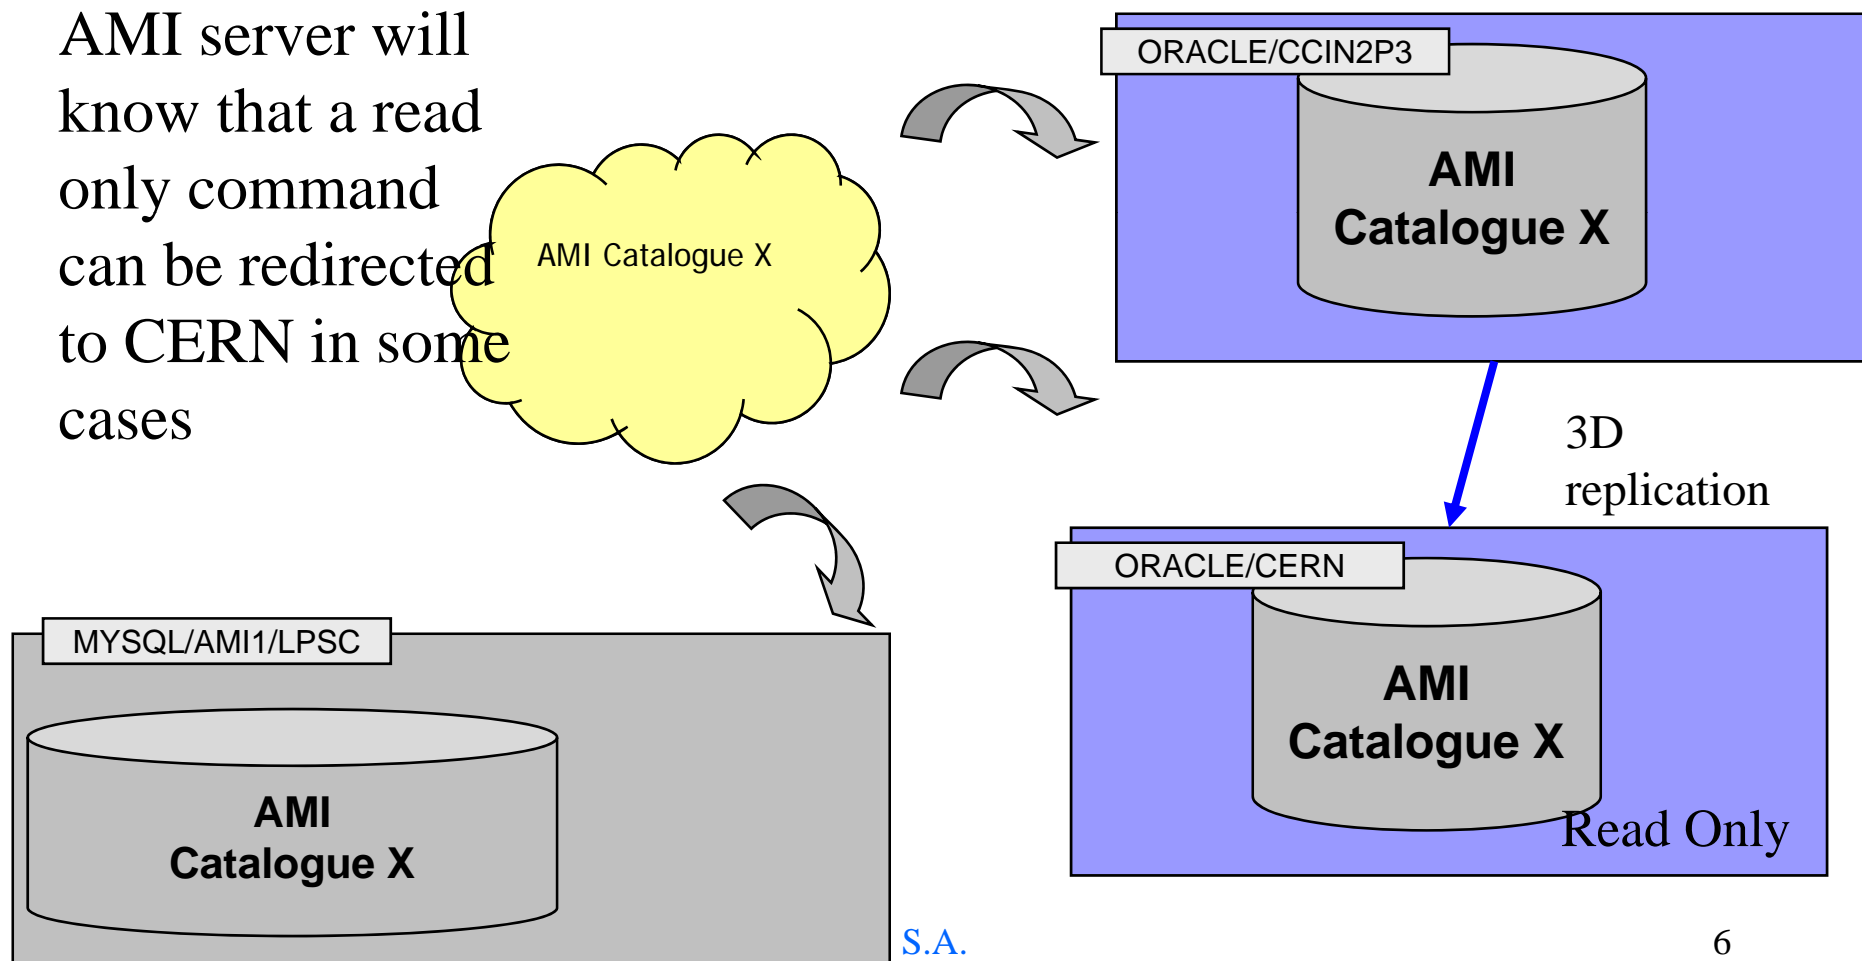
- Estimations of size are difficult.
 - Proportional to nFiles produced and nTimes data reprocessed.
- The size of AMI is almost negligible compared with the total.
 - Estimated to be 0.04 TB in 2010
 - AMI will have $\sim 0.5 \times 10^6$ transactions /day



How to use the replica?

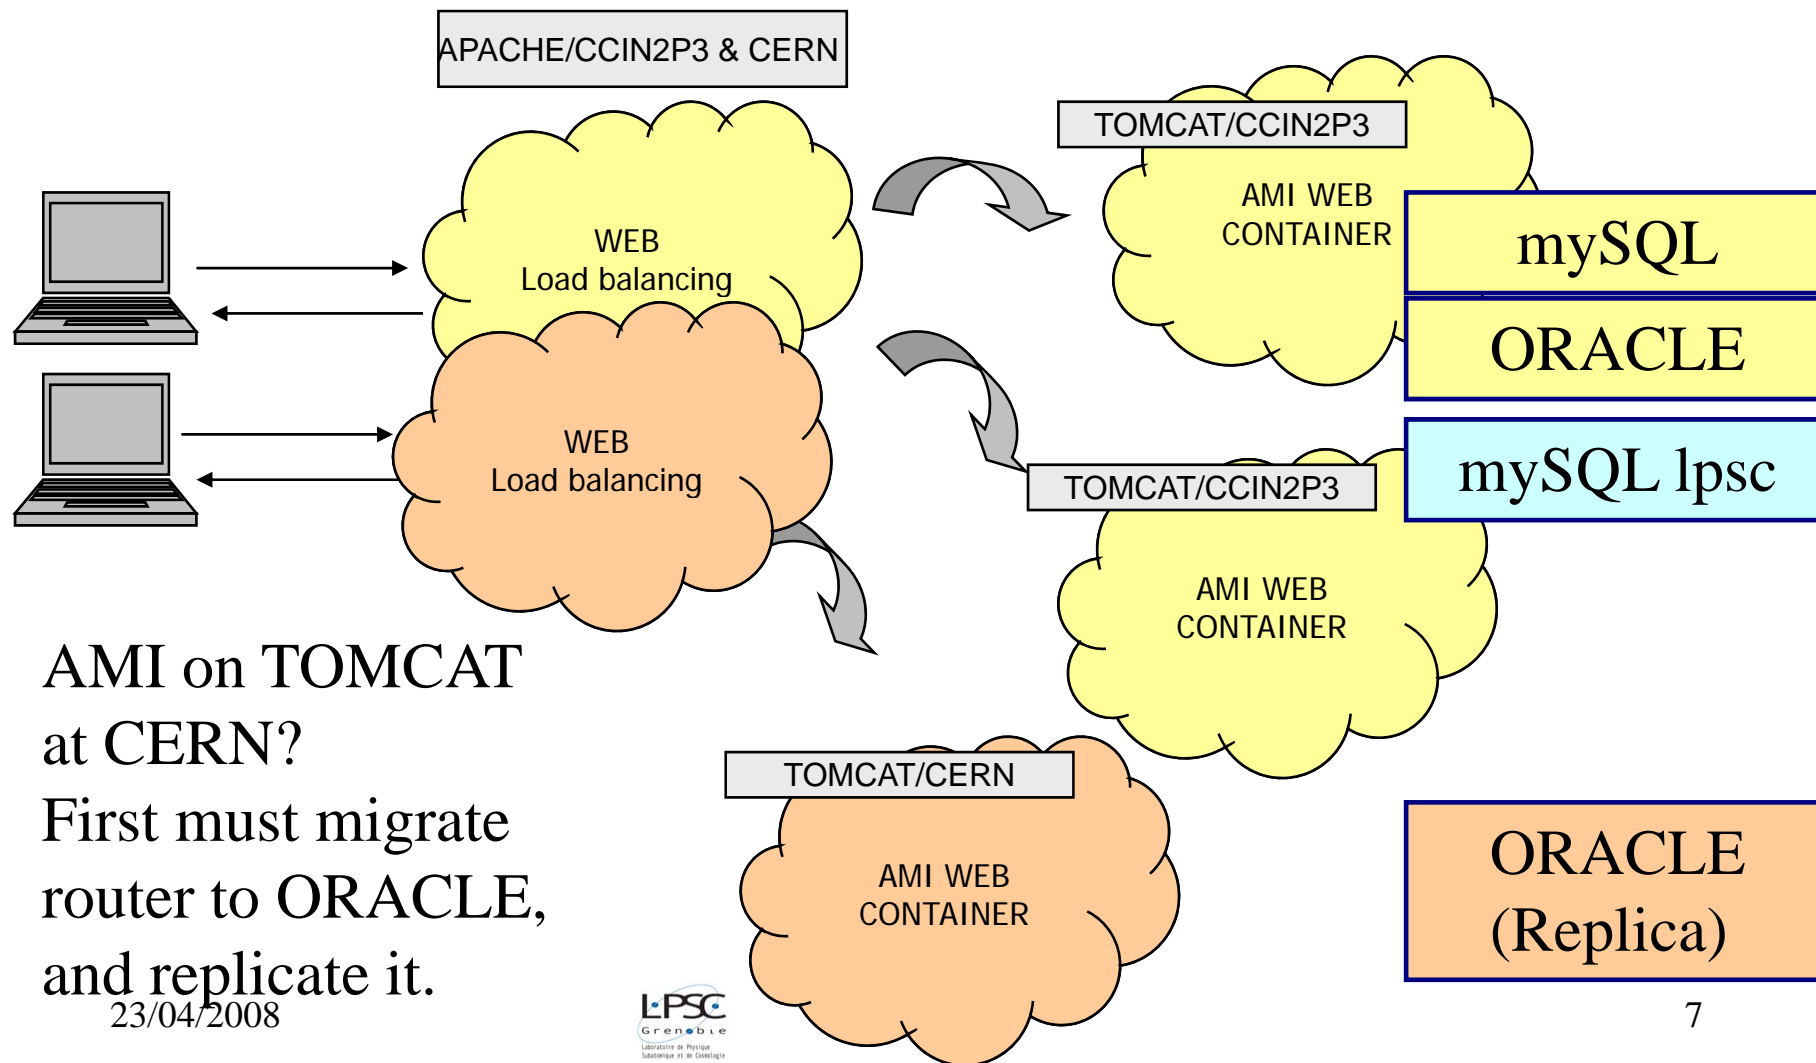
Try to manage use of replicas seamlessly

AMI server will know that a read only command can be redirected to CERN in some cases





End 2008?





Explanation

- No problems anywhere.
 - Use either front end.
 - Both servers know that a read only command on some catalogues can go to the CERN replica. Writes must go to Lyon.
- Problem with Lyon?
 - Only CERN front end up
 - Only reads available at CERN
- Problems everywhere?
 - Take coffee break.