

Persistence Management in ICE

Moreno Marzolla, Alvise Dorigo

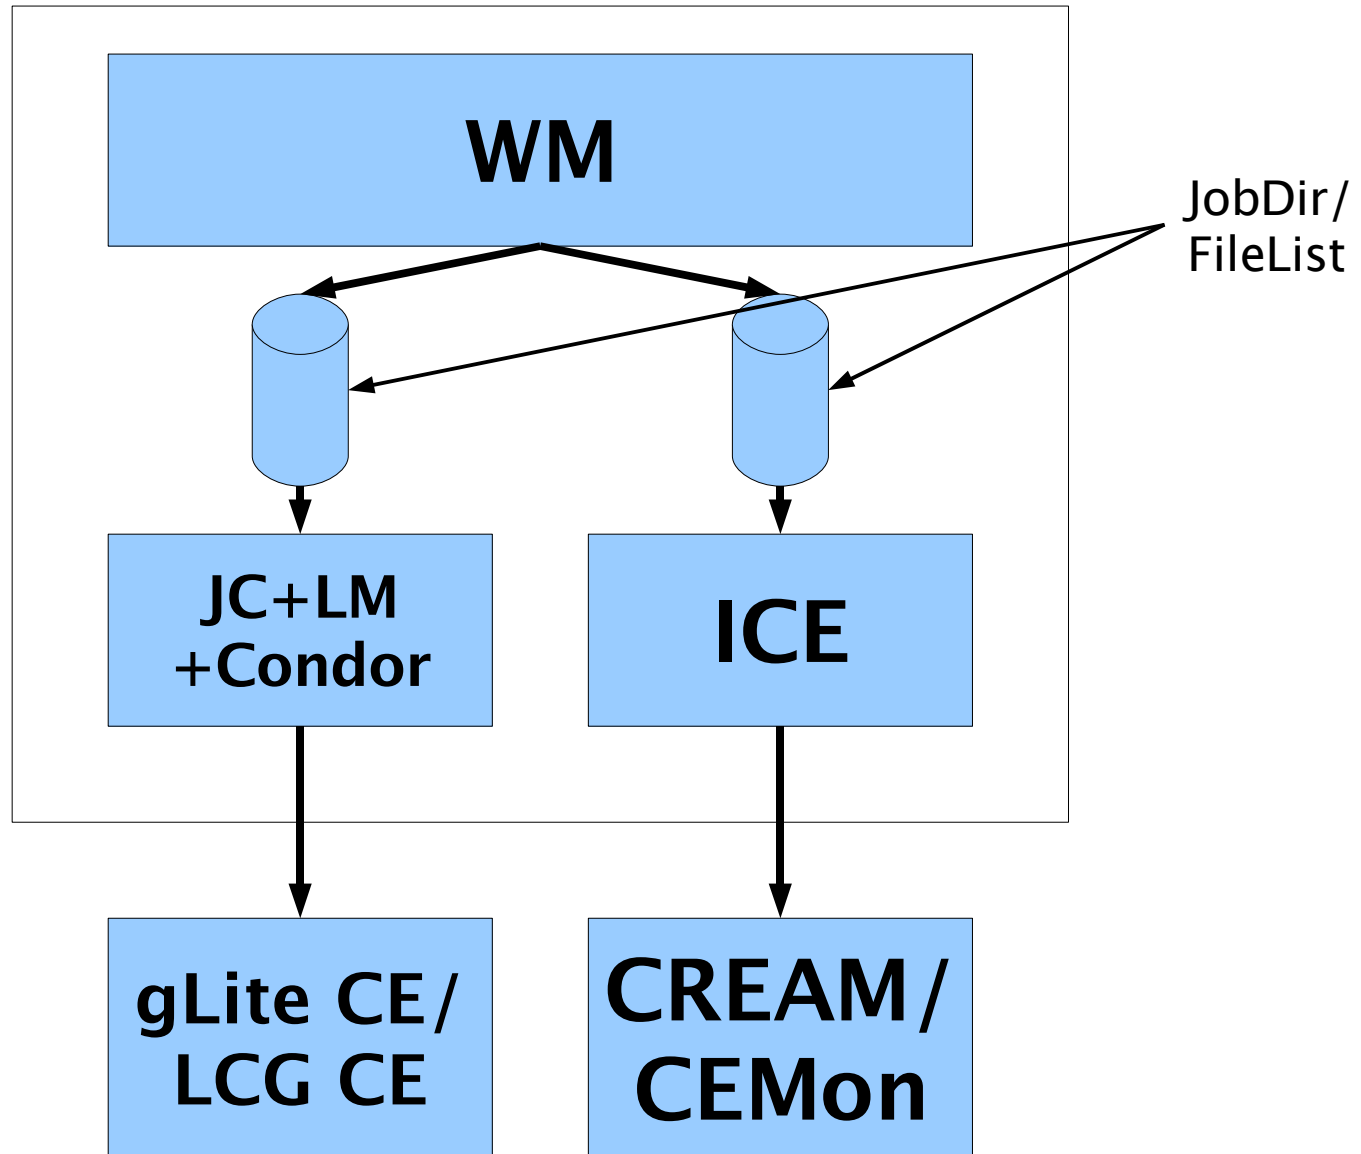
INFN Padova

{moreno.marzolla | alvise.dorigo}@pd.infn.it

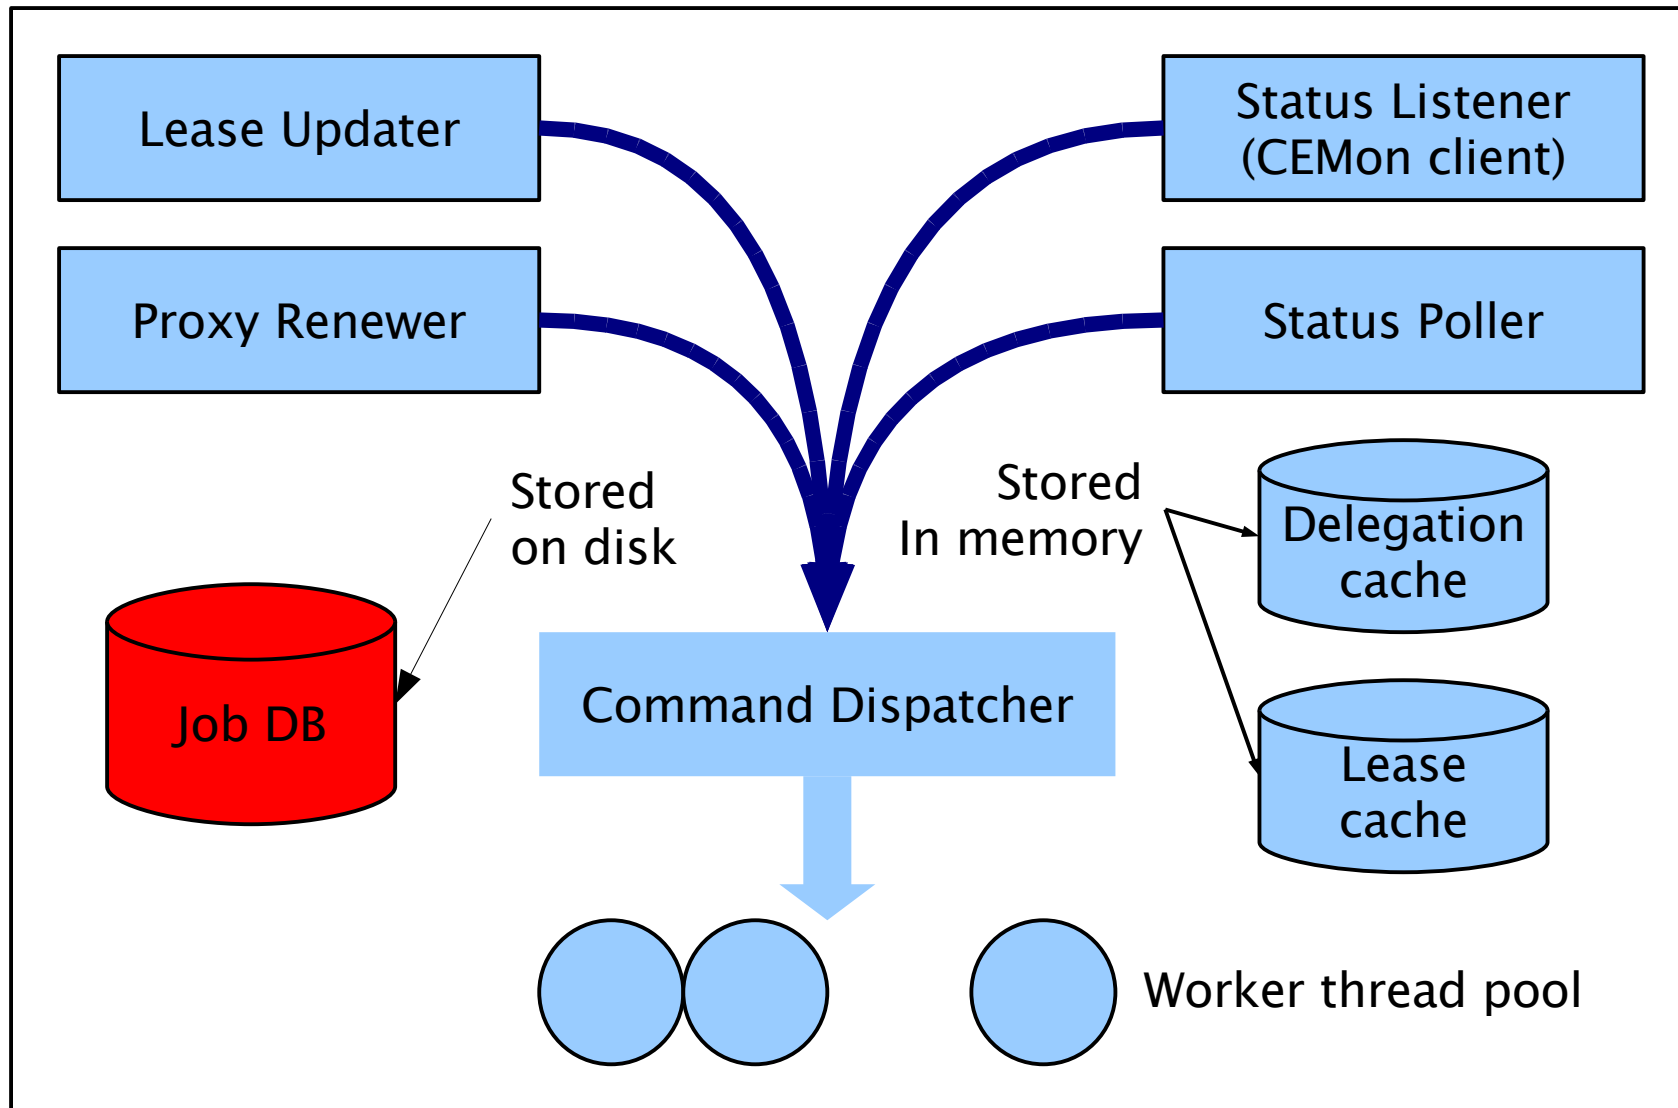
What is ICE

- Interface to Cream Environment
- Manages job submission/cancellation requests from the WM to CREAM CEs
- Takes care of all interactions with CREAM
 - Proxy Renewal
 - CREAM job lease management
 - Monitoring job status changes
 - LB job status updates

What is ICE



ICE internals



Persistency in ICE

- The only persistent structure is the job cache
 - Contains information on each “active” job (i.e., non-terminated jobs)
 - CREAM job id, Grid job id, lease id, JDL, CE id, job status, proxy cert path, LB sequence code, user DN, last time ICE received notifications about job status change, CREAM and CEMon endpoints...
 - See `org.glite.wms.ice/src/iceUtils/creamJob.h`

Persistency in ICE

- The job cache is stored using Berkeley DBs
 - <http://www.oracle.com/technology/products/berkeley-db/index.html>
- Simple, lightweight persistency mechanism using on-disk files
 - ICE persistency informations are kept in the directory specified with the `persist_dir` configuration parameter in section ICE
 - Default: "`${GLITE_LOCATION_VAR}/ice/ice_db`"

Persistency in ICE

- Queries performed by ICE
 - Lookup for a job with a given CREAM job id
 - Lookup for a job with a given GRID job id
 - Sequential sweep over the job cache
 - Removal of one element from the cache
- Still room for future optimizations
 - Search for all jobs with the same lease ID (to purge lease-expired jobs)
 - ...

Why BDB?

- Simple, lightweight, general-purpose, very well documented
- Embedded DB: no need to install/configure a new separate service
- It is installed by default on SLCX
 - BDBs are used for the RPM database in RPM-based linux distribution
- “It just works”[®]
 - In ~1 year of use, we had only one database corruption. It was our fault, and the recovery tools shipped with the BDB successfully recovered the database

Changing persistency mechanism in ICE: personal view

- MUST be carefully planned
 - NO “let's try to see if that works”
- Current solution works and has been tested in many nasty situations
 - Multiple unclean service restarts during ICE debugging and testing
- Time consuming
 - Estimated manpower for changing persistency layer in ICE ~1 to 2 person/months plus debugging
 - We NEED good written documentation: BDBs have been used without the need to write a single e-mail to the developers