CMS Dynamic Resource Provisioning

Claudio Grandi

INFN Bologna

CMS Workload Management

Central role of the glidein-WMS

- Appropriate plug-in's in the glidein factories allow accessing heterogeneous resources
  - Grid, Cloud, opportunistic resources
  - Credentials managed according to infrastructure's req's
- Current architecture requires the glidein running on the execution host to be able to contact the glidein-WMS to get the jobs to be executed
- The startd publishes hardware resources available on the resource for proper match-making
  - Single- and multi-core resources are allowed

The whole picture

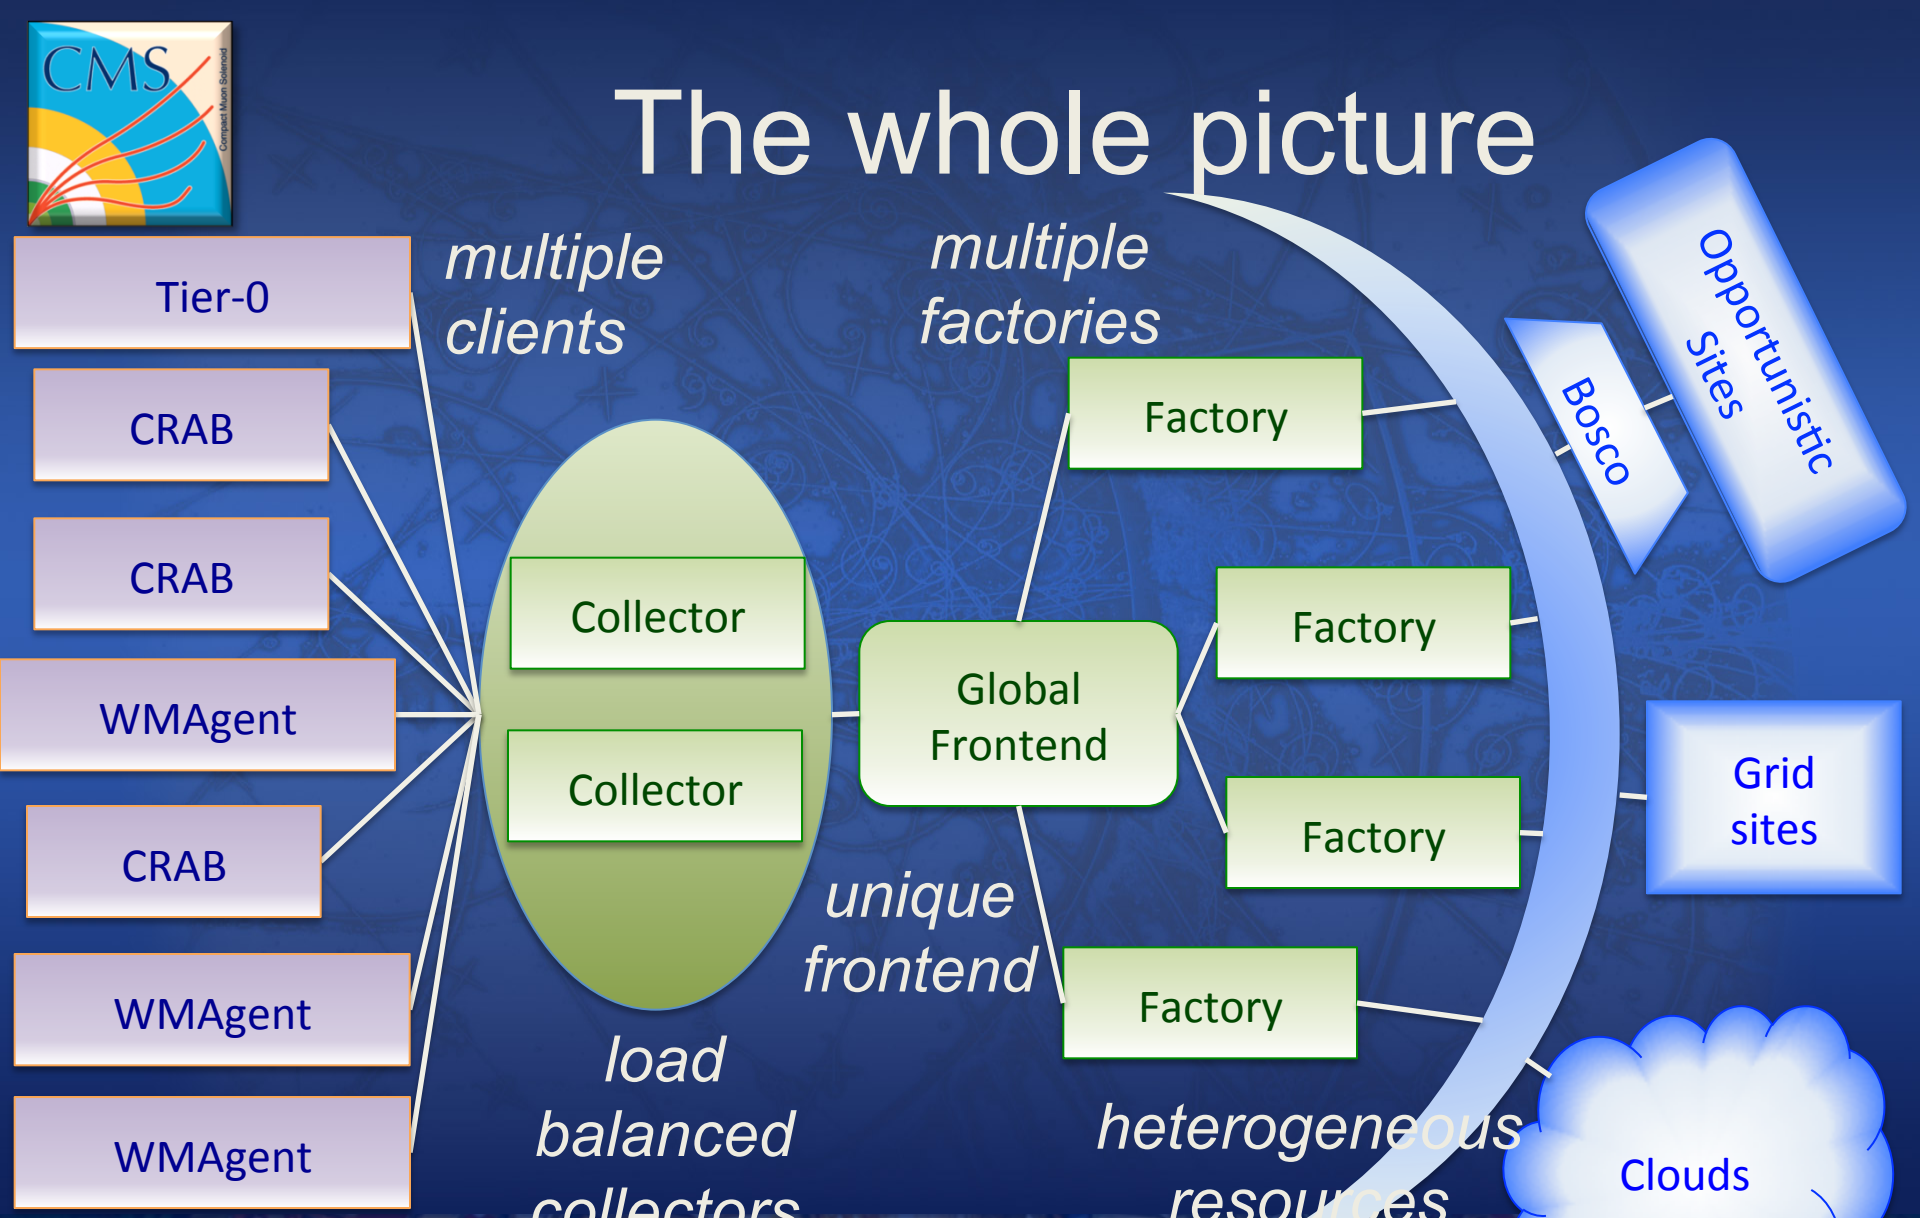

The core: the Glidein Factory

The Factory harvests batch slots on the Grid

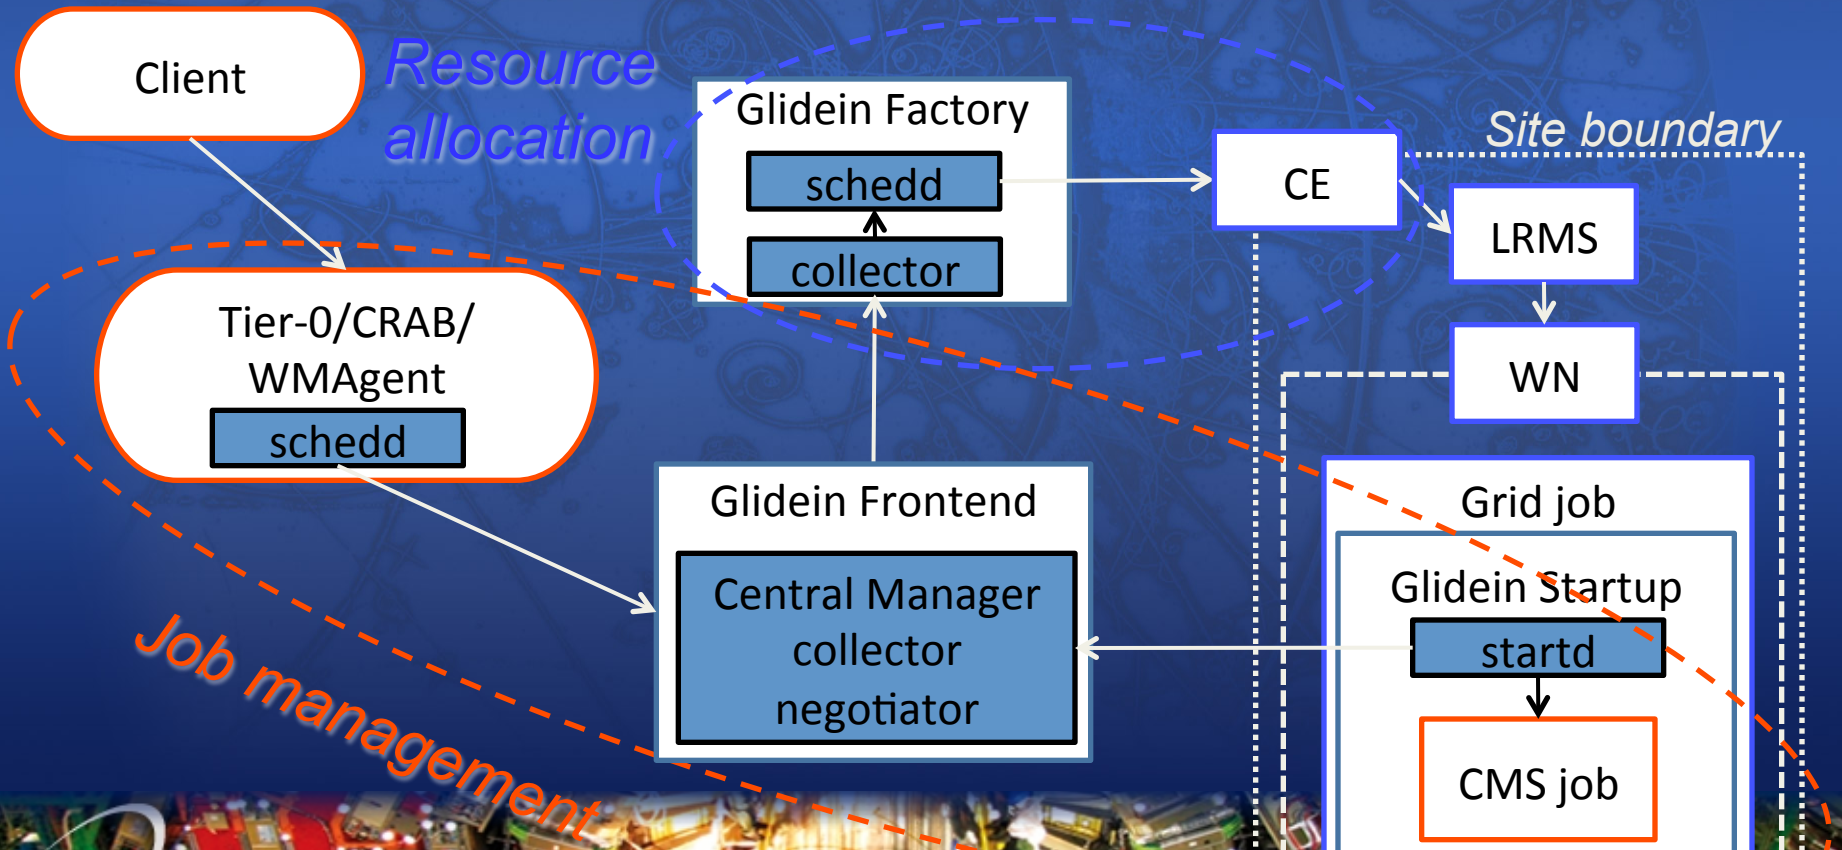

The core: the Glidein Factory

The Factory harvests whole machines on the Cloud

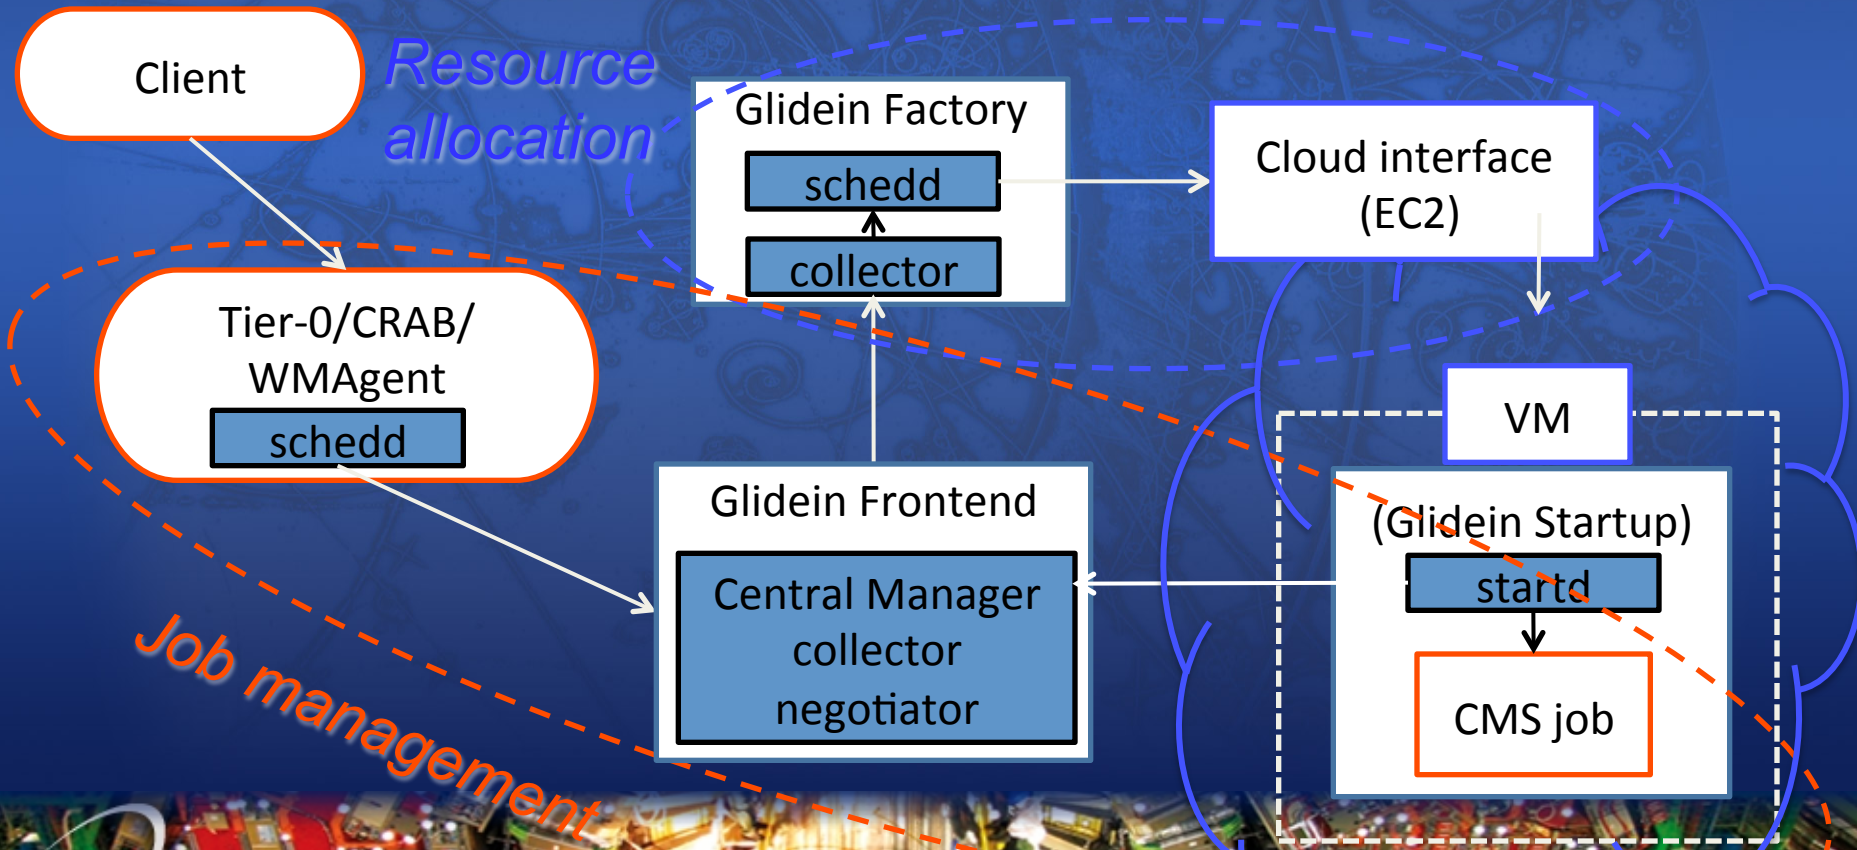

The core: the Glidein Factory

The Factory harvests batch slots on opportunistic resources

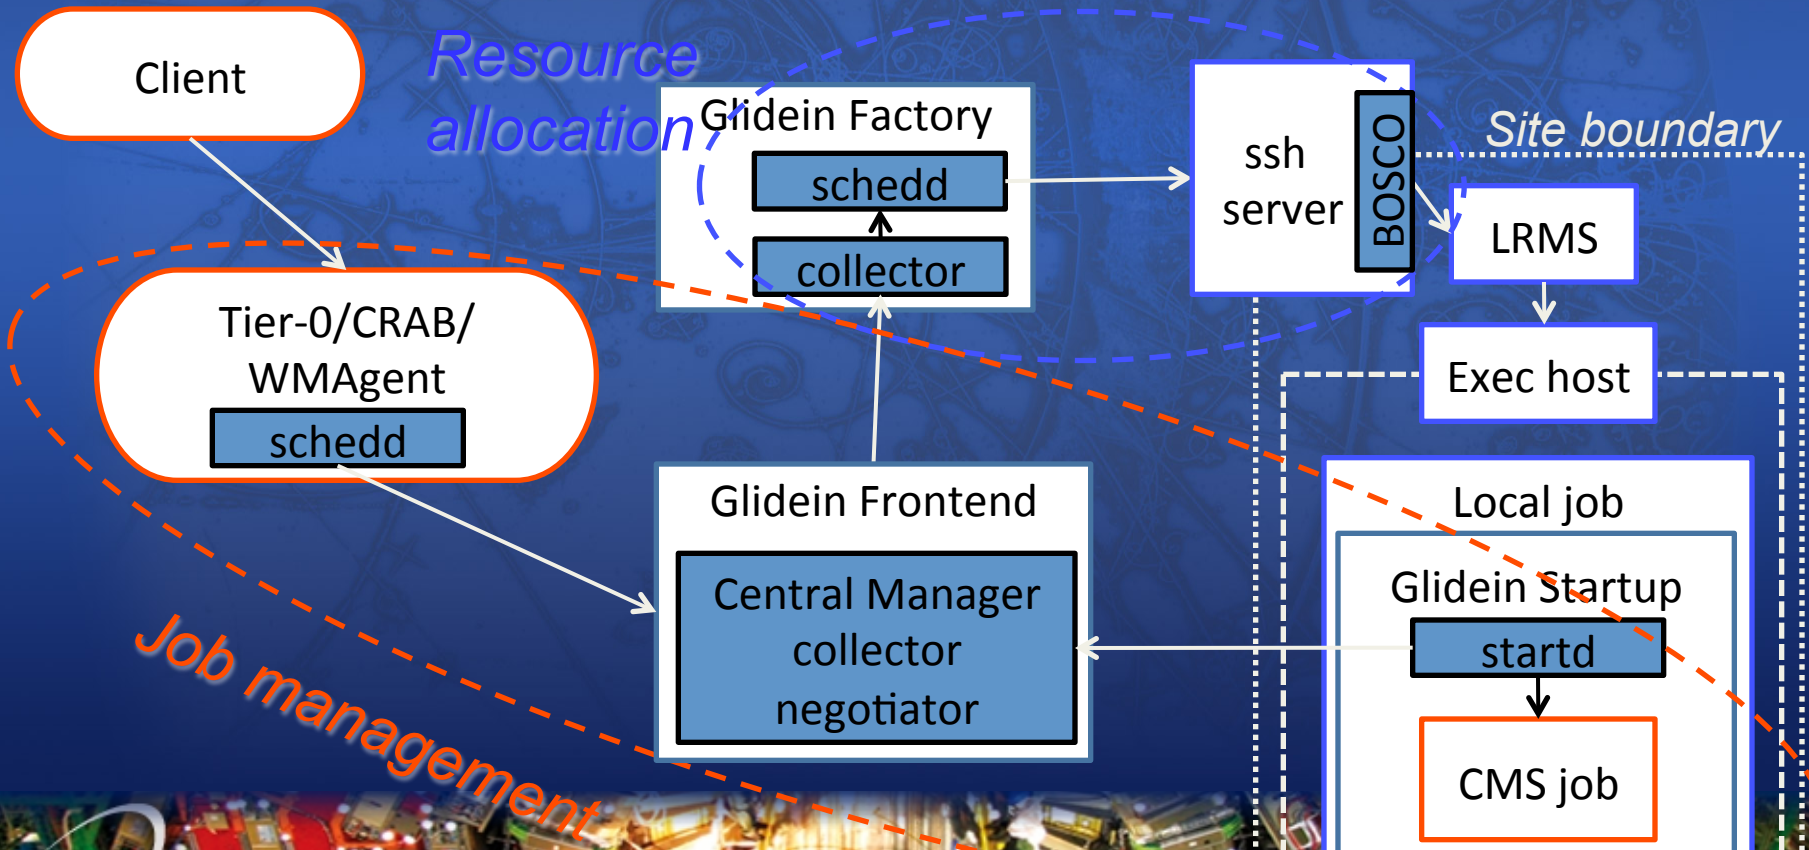

Configuration and software

- Configuration is retrieved from CVMFS
- Configuration is site dependent
  - Including read and write protocols, LFN to PFN translations, squid configuration for Frontier…
- Site name, architecture, OS, are
  - Discovered on a CMS Grid site
  - Included in the VM image on Clouds
  - In a Parrot wrapper (see Dirk's presentation)
  - Software is retrieved from CVMFS
  - Multiple architectures available

Site services

- All services can be accessed remotely with some performance penalty (depending on the network connection and the workflow type)
  - Remote CVMFS and Frontier squid servers
  - Read data through xrootd
  - Write data typically through gridftp
- Use Parrot for CVMFS if it needs to be in user space
- If local services are available/deployable they can be configured and used
- Still the glidein on the execution host needs to be able to contact the glidein-WMS

Checkpointing

- Currently no checkpointing is available
  - Try to keep jobs reasonably short on sites where preemption is possible (e.g. HLT farm)
- Work is going on to investigate checkpointing
  - at application level
  - of the VM on Clouds
- Models for single event processors are being investigated but no real work going on