




gLExec and MyProxy integration in the ATLAS/OSG PanDA Workload Management System



J. Caballero^a, J. Hover^a, M. Litmaath^b, P. Nilsson^c, M. Potekhin^a, T. Wenaus^a, X. Zhao^a
^aBrookhaven National Laboratory, ^bCERN, ^cUniversity of Texas at Arlington

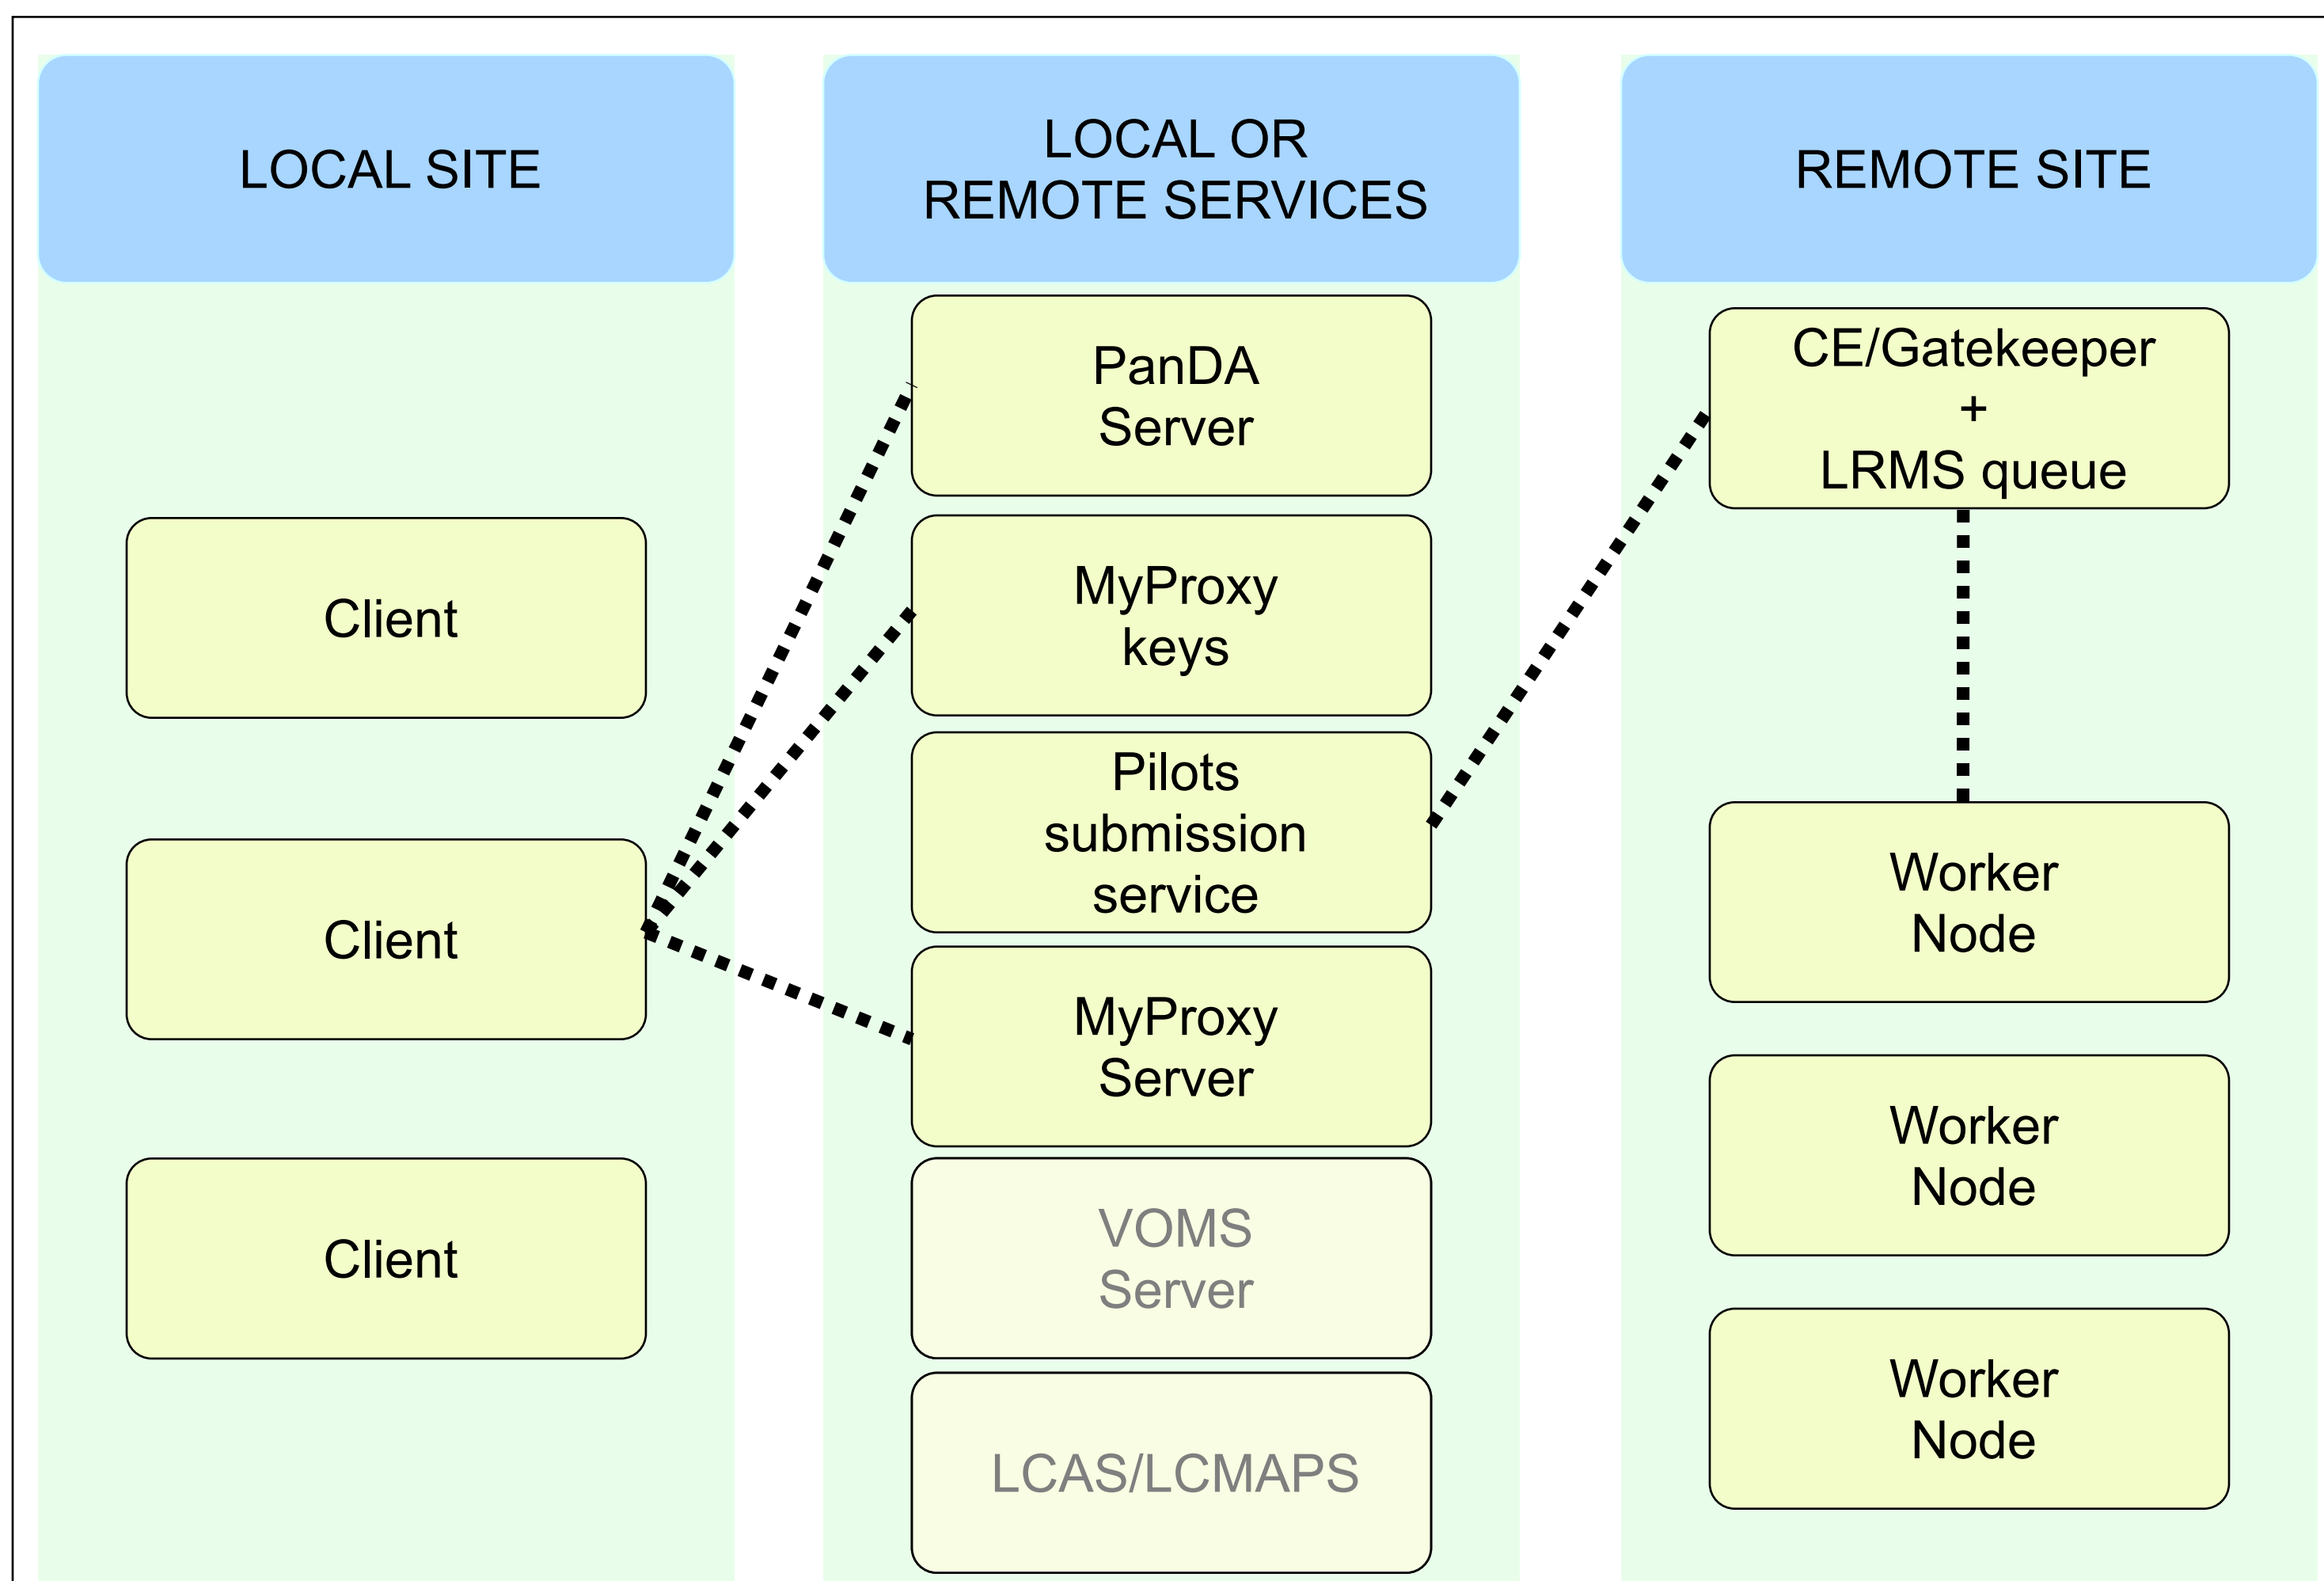
Abstract

Worker nodes on the grid exhibit great diversity, making it difficult to offer uniform processing resources. A pilot job architecture, which probes the environment on the remote worker node before pulling down a payload job, can help. Pilot jobs become smart wrappers, preparing an appropriate environment for job execution and providing logging and monitoring capabilities.

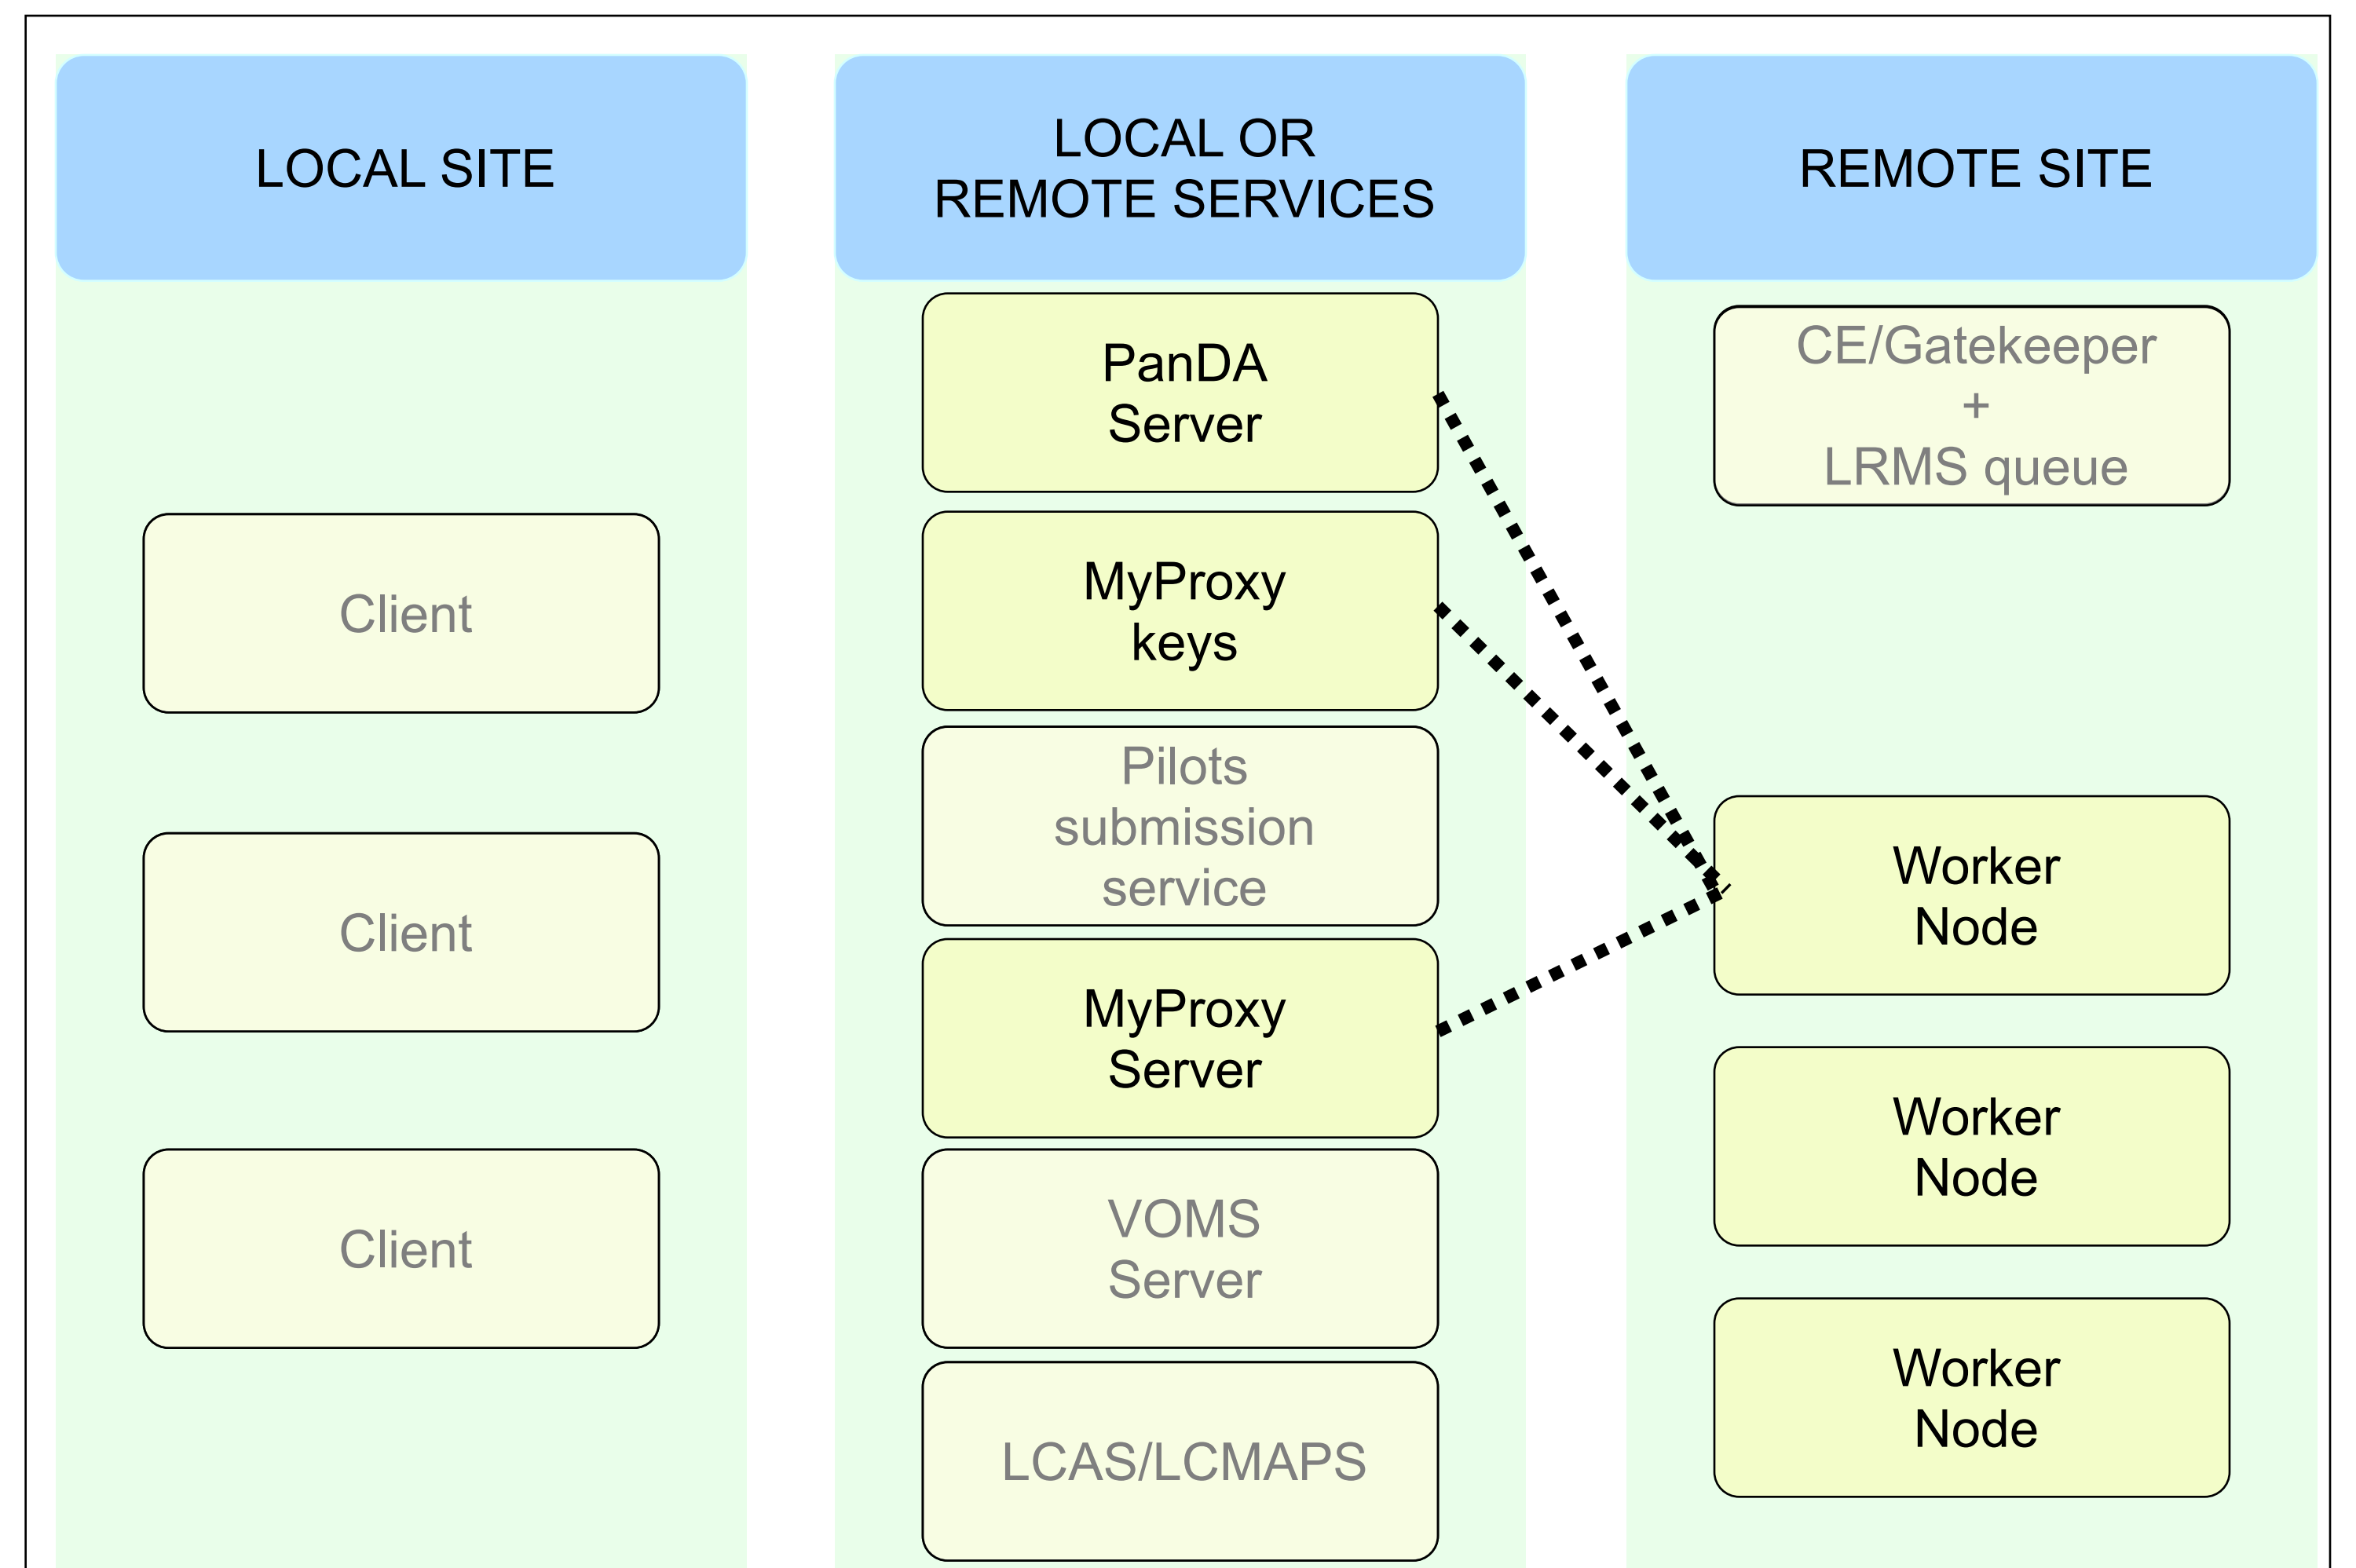
PanDA (Production and Distributed Analysis), an ATLAS and OSG workload management system, follows this design. However, in the simplest (and most efficient) pilot submission

approach of identical pilots carrying the same identifying grid proxy, end-user accounting by the site can only be done with application-level information (PanDA maintains its own end-user accounting), and end-user jobs run with the identity and privileges of the proxy carried by the pilots, which may be seen as a security risk.

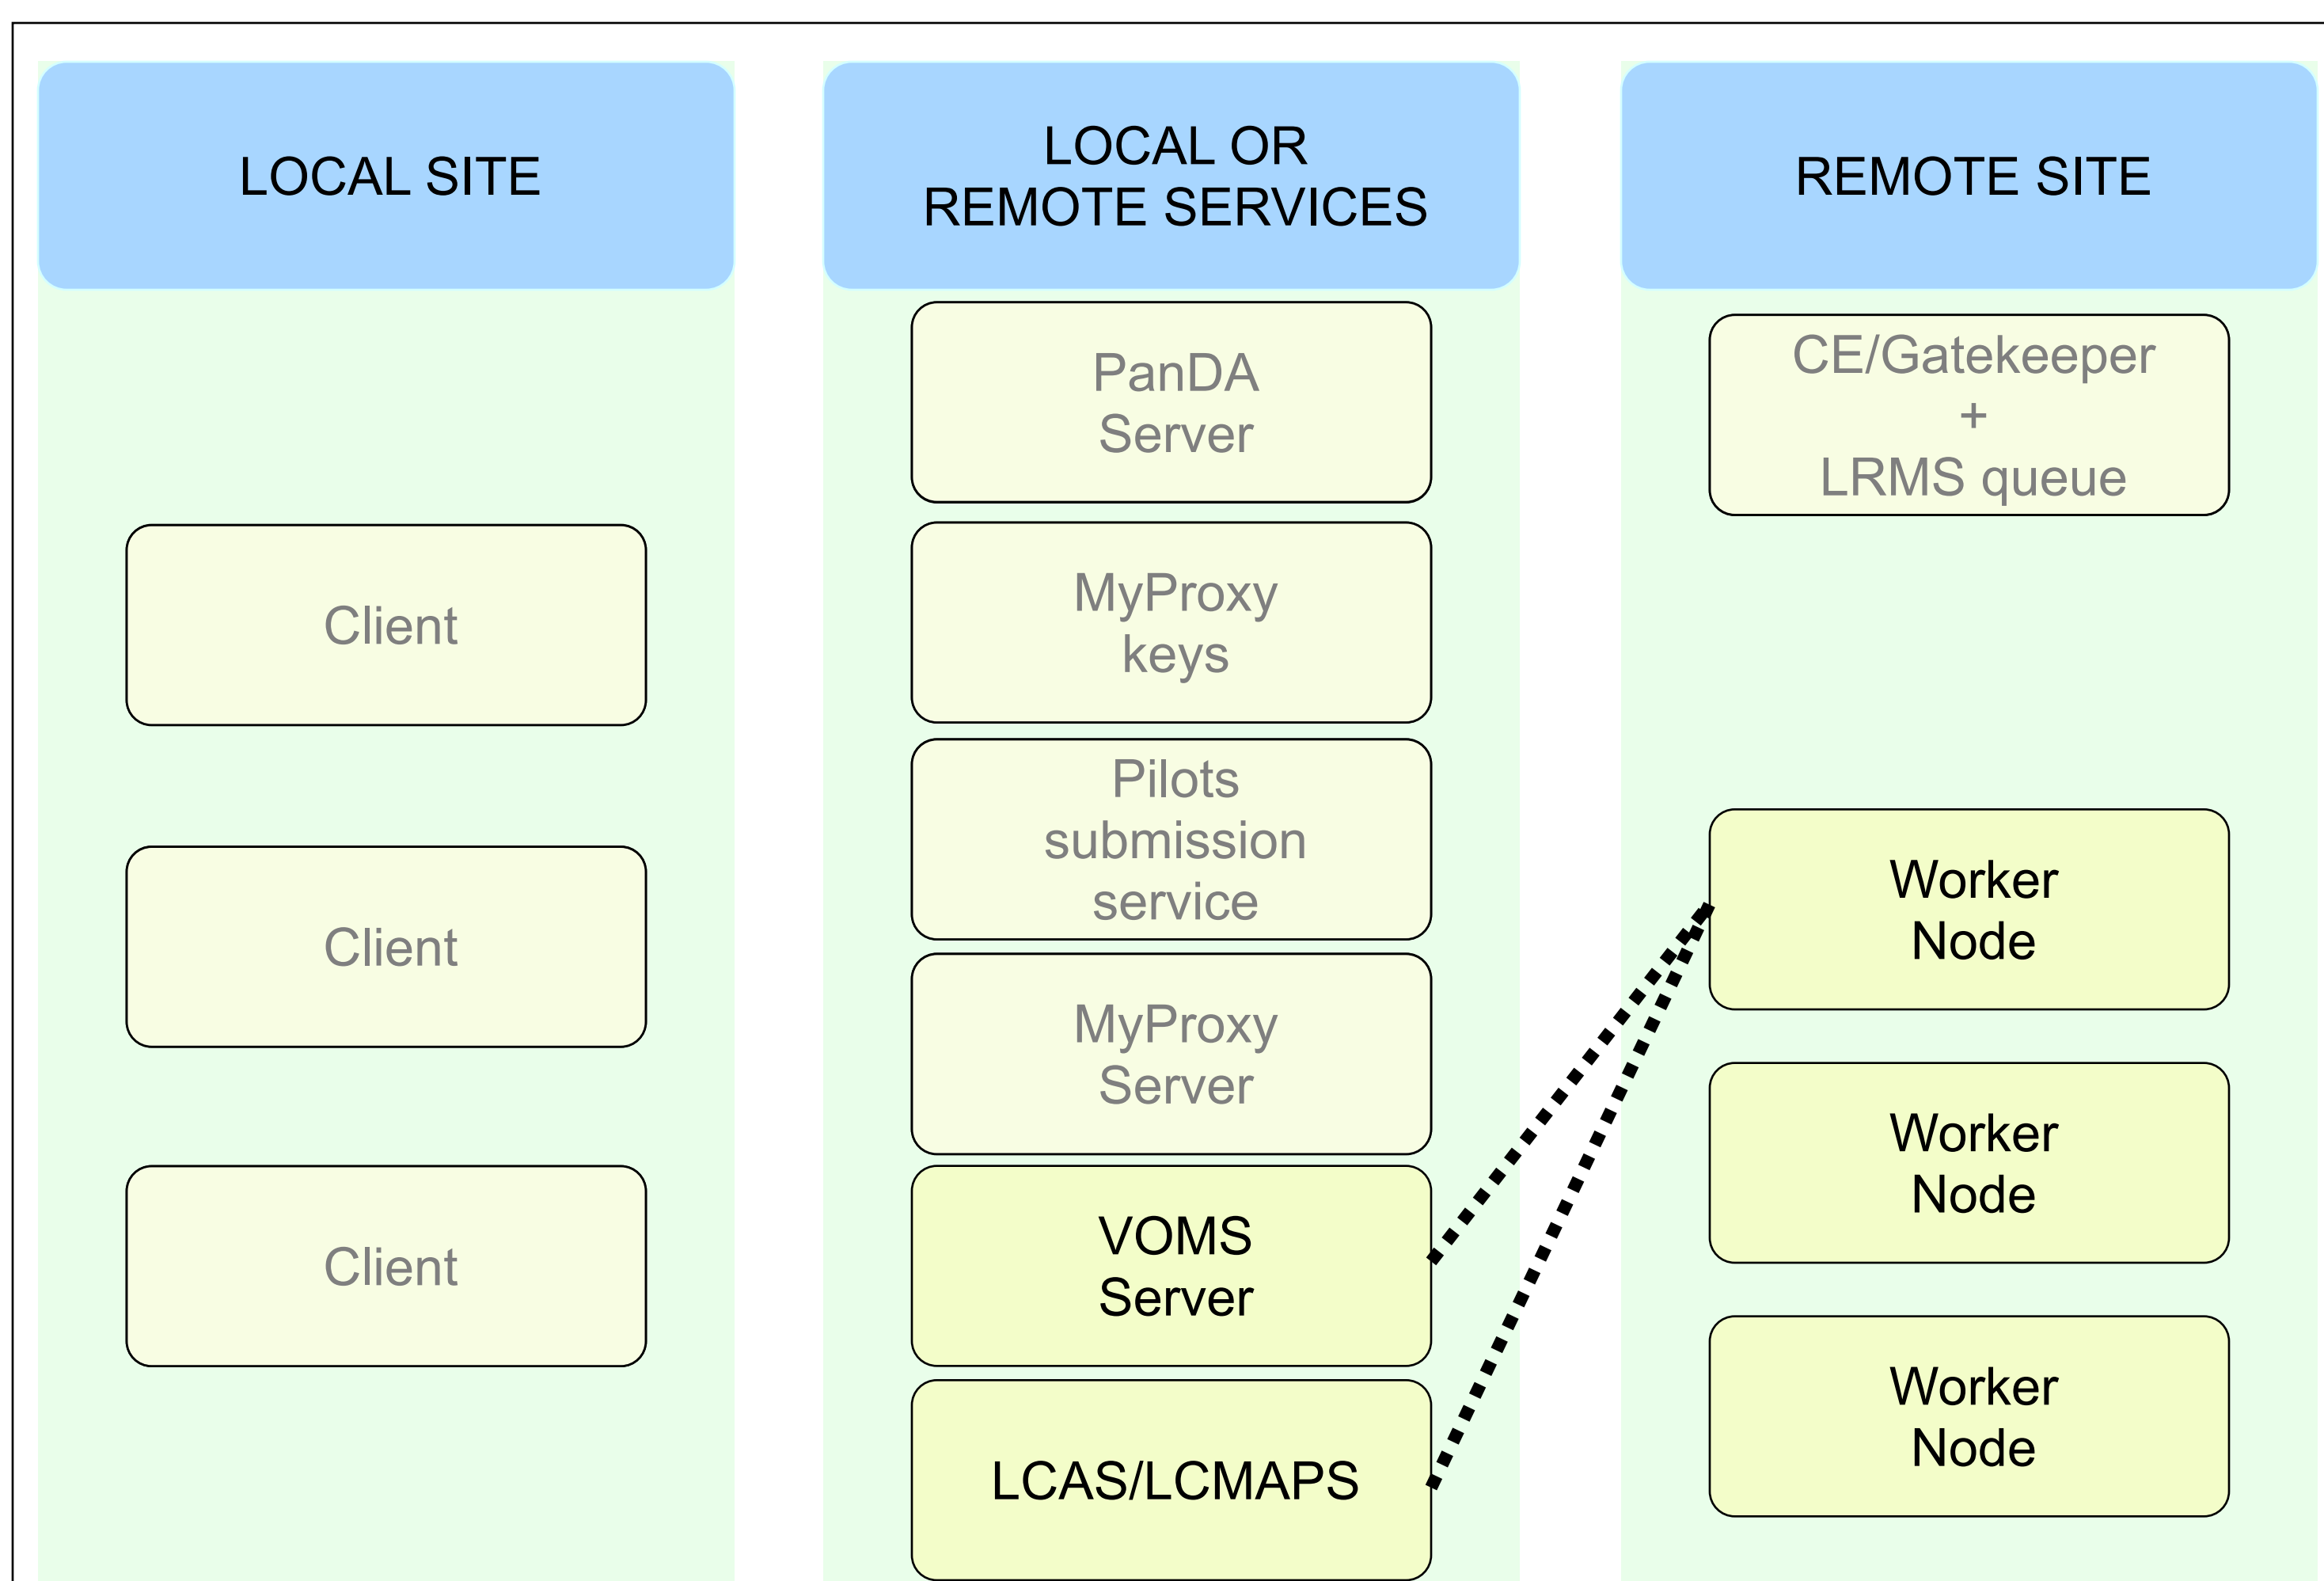
To address these issues, we have enabled Panda to use gLExec, a tool provided by EGEE which runs payload jobs under an end-user's identity. End-user proxies are pre-staged in a credential caching service, MyProxy, and the information needed by the pilots to access them is stored in the PanDA DB. gLExec then extracts from the user's proxy the proper identity under which to run.



When a client submits a new analysis job to the PanDA server, the existence and lifetime of his credentials in a MyProxy server are checked. New credentials are generated and delegated if needed. An associated random key, needed for the retrieval, is also created and stored. Only pilots with special VOMS roles are allowed to retrieve stored credentials. **ALL COMMUNICATIONS ARE SECURE!** Simultaneously new pilots are submitted to sites.

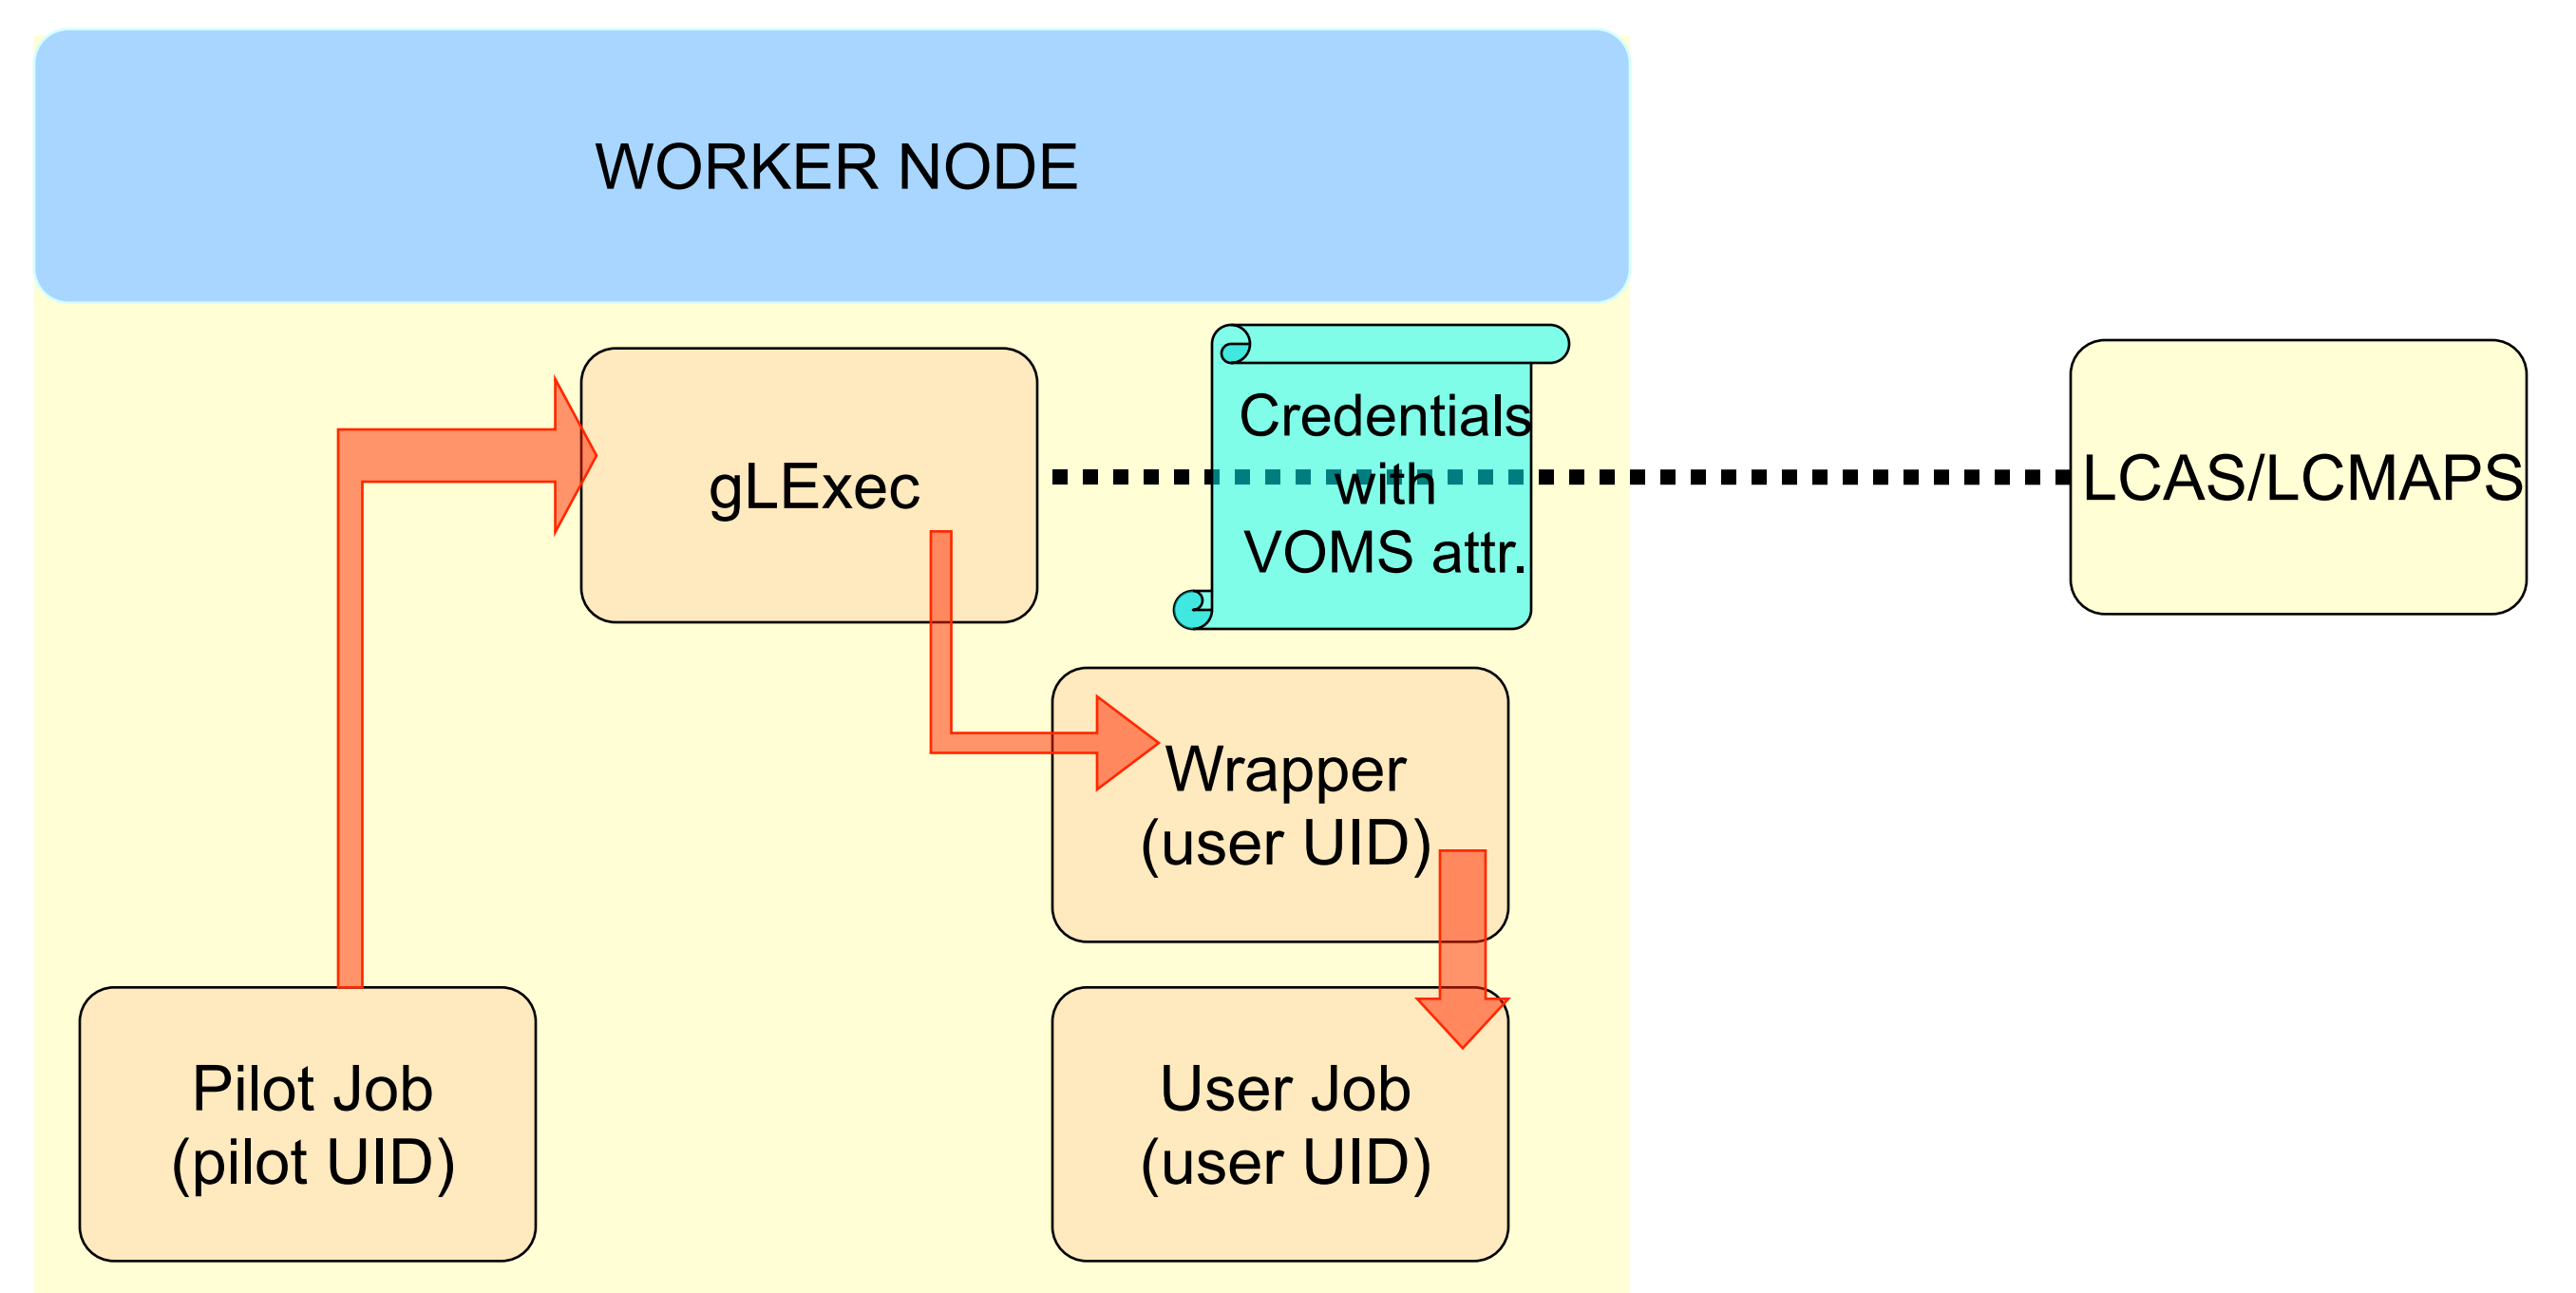


When a pilot picks up a new analysis job from the server, all information needed to retrieve the delegated proxy is obtained. The end user credentials can then be retrieved. Given its short lifetime (configurable), the pilot is also responsible for its periodical renewal. **ALL COMMUNICATIONS ARE SECURE!** A tokens-based mechanism protects against a compromised pilot querying the servers to retrieve many user credentials.



If retrieved credentials do not bring VOMS extensions they are added on the WN querying a VOMS server. gLExec is invoked. gLExec can log the user credentials or can operate in setuid mode: it changes the current UNIX user id based on results from LCAS/LCMAPS mapping (root privileges are required). After execution, the user credentials on the WN are destroyed.

Execution



After the uid switch, the pilot environment has been replaced by the end user's one. The current execution location has moved to the end user HOME directory. gLExec has been integrated to execute an intermediate wrapper:

- The vanished environment is reconstructed.
- Other pre-post processing tasks can be performed.
- The wrapper invokes the analysis job.

Deployment status

OSG

gLExec is integrated and functioning in OSG middleware stack.
 PanDA pilot use of gLExec in OSG environment is validated.
 MyProxy service in support of PanDA/gLExec is deployed and operating.
 No US ATLAS sites have adopted gLExec as yet, so no US production deployments.

EGEE

gLExec is still being integrated/tested in the EGEE middleware stack.
 EGEE deployment will follow (expected for spring 2009).
 gLExec usage in Panda analysis production on EGEE will follow that.

For more information contact:

Jose Caballero
 Brookhaven National Lab
 PAS Group, Physics Department
 11973, Upton, NY
 jcaballero@bnl.gov