HTCondor-CE

The Open Science Grid

A framework for large scale distributed resource sharing

addressing the technology, policy, and social requirements of sharing

- OSG is a consortium of software, service and resource providers and researchers, from universities, national laboratories and computing centers across the U.S., who together build and operate the OSG project. The project is funded by the NSF and DOE, and provides staff for managing various aspects of the OSG.
- Brings petascale computing and storage resources into a uniform grid computing environment
- Integrates computing and storage resources from over 50 sites in the U.S. and beyond

Current OSG Resources

- OSG has more than 80 participating institutions, including self-operated research VOs, campus grids, regional grids and OSG-operated VOs
- Provides about 80,000 CPU-days per day in processing
- Provides 1 Petabyte per day in data transport

Challenge: Job Management

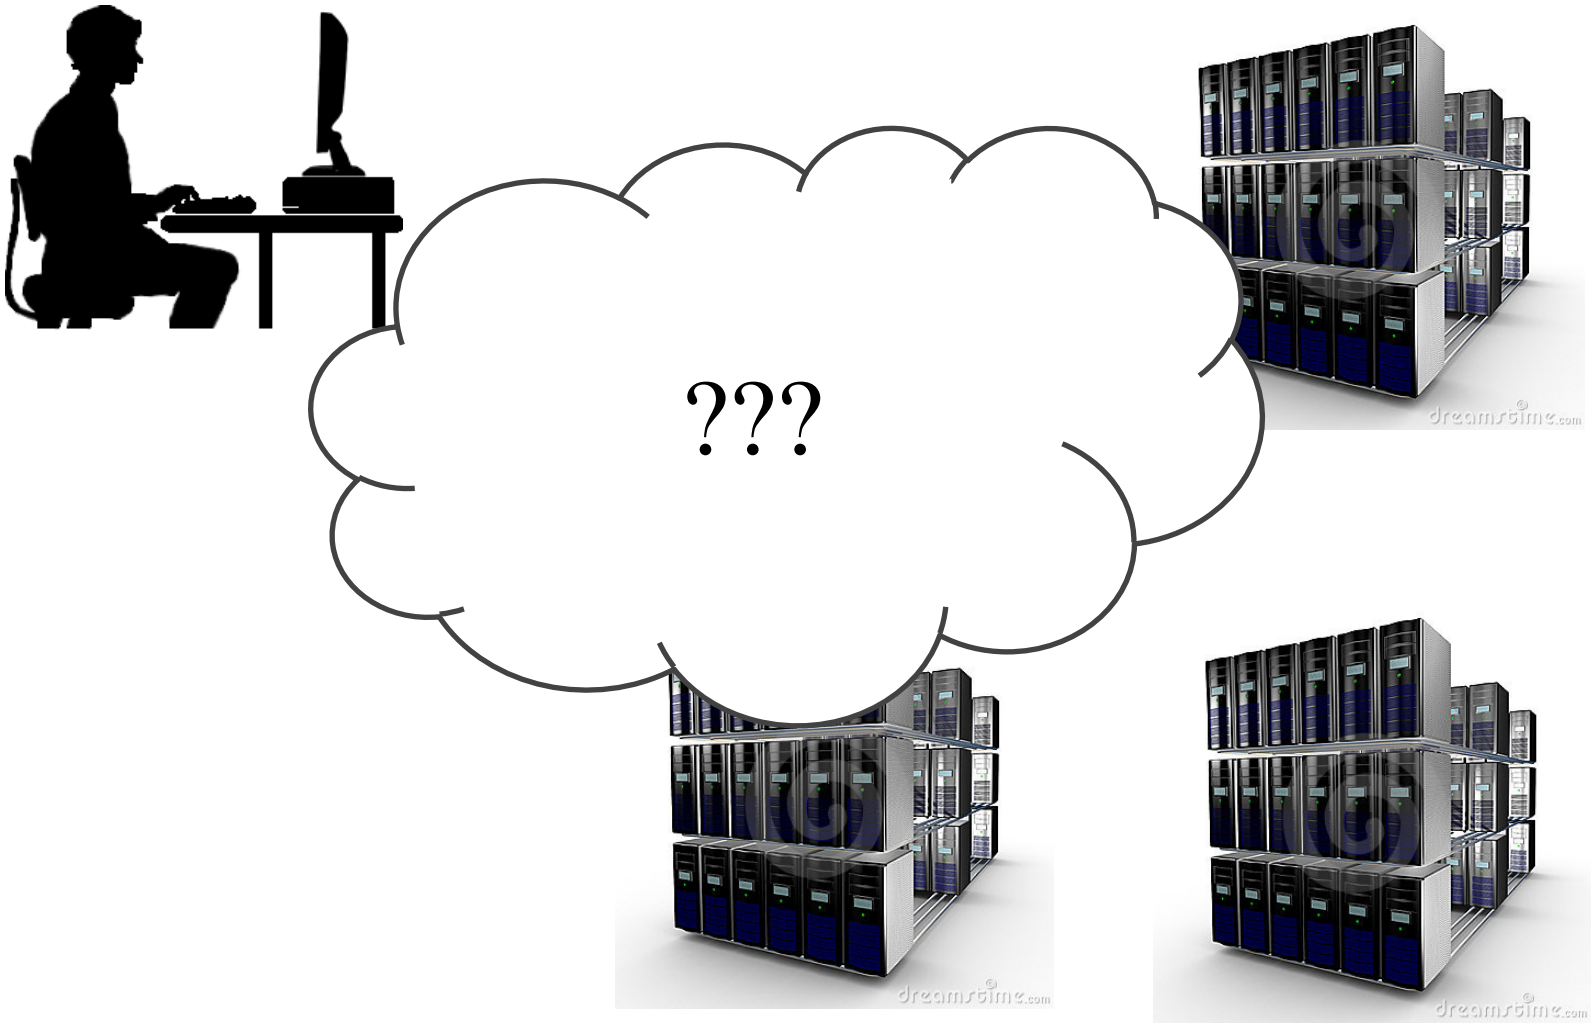

HTCondor-CE

- Special configuration of HTCondor
- Sits on CE* of each cluster
- Allows…
  - Remove job submission and management
  - Strong authentication (GSI/VOMS)
  - Logging and monitoring
  - Scalability
  - Work with existing batch systems

* Compute Element: gateway to a cluster

HTCondor Building Blocks

- HTCondor-C
  - Submit jobs from one HTCondor scheduler to another (user's machine to CE)
- Job Router
  - Transform jobs (localize jobs at CE)
- Blahp
  - Submit jobs to batch system (PBS, SGE, etc.)

How It Works

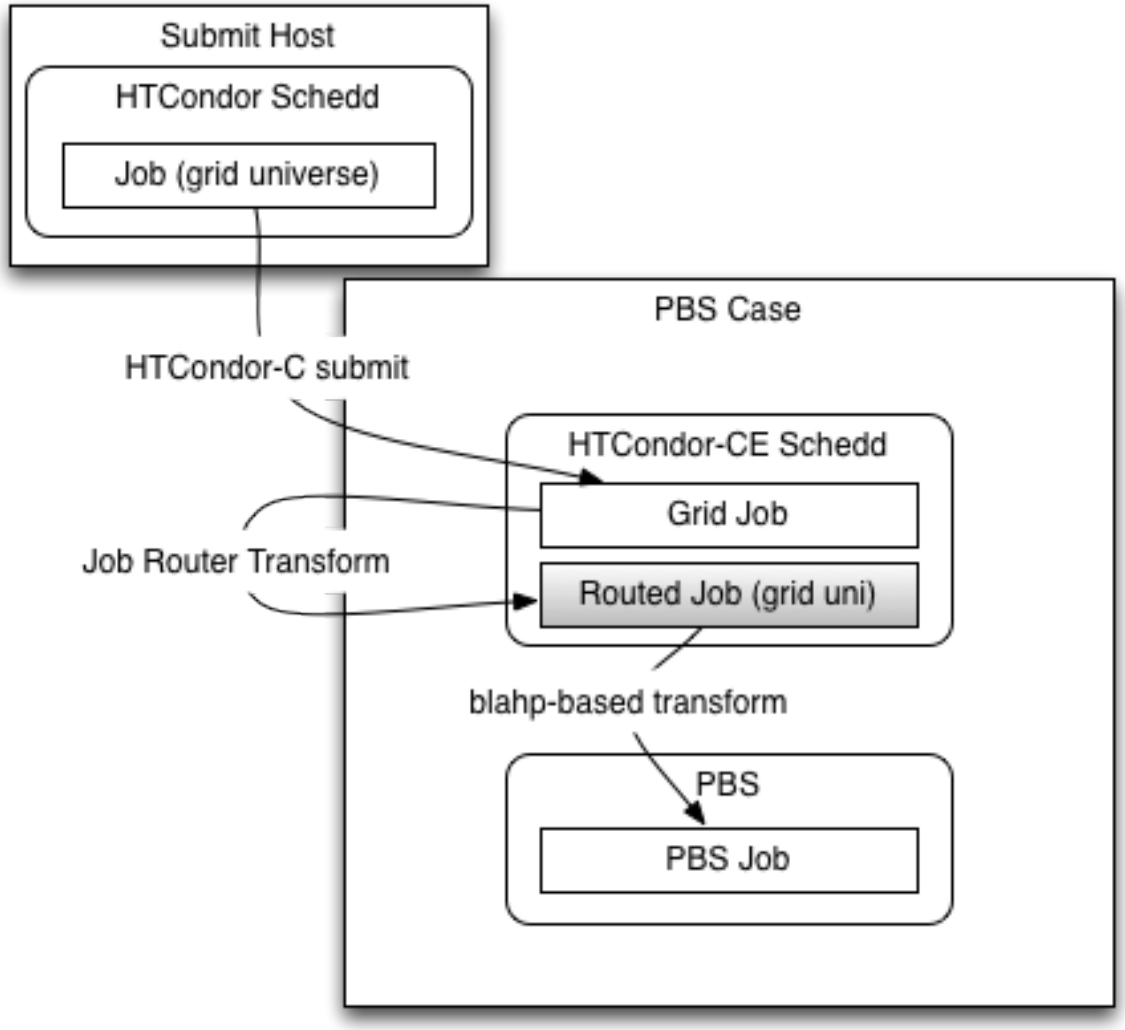
Submit Host
HTCondor Schedd
Job (grid universe)
PBS Case
HTCondor-C submit
HTCondor-CE Schedd
Grid Job
Job Router Transform
Routed Job (grid uni)
blahp-based transform
PBS
PBS Job

HTCondor-CE Work List

- Harden
- Scale
- Audit
- Transform via local policy
- Troubleshooting tools
- Package

Harden and Scale

- Blahp
  - Improved file cleanup
  - Better error messages on failure
  - Handle errors more gracefully
- HTCondor-C
  - Handle more submitters and jobs without making schedd unresponsive to others

Security Audit Log

- Record actions by the user that affect the job queue
  - Submission, removal, modification
- Record how the user was authenticated
- Record job credential files
- Time-based rotation

condor_ping

- Diagnose communication problems
- Detailed diagnosis of failures
  - Can you connect to the server?
  - Can you authenticate with the server?
  - Are you authorized by the server?

condor_ping Example

```
% condor_ping -verbose write
Remote Version:                 $CondorVersion: 7.9.6 ...
Local  Version:                 $CondorVersion: 7.9.6 ...
Session ID:                     nostos:31671:1367353167:9
Instruction:                    write
Command:                        60021
Encryption:                     none
Integrity:                      none
Authenticated using:            FS
All authentication methods:     GSI,FS
Remote Mapping:                 jfrey@cs.wisc.edu
Authorized:                     TRUE
```