Improving Security in the ATLAS PanDA System

Douglas Smith for the ATLAS Collaboration

CHEP, Taipei 2010

PanDA

- PanDA is built on a pilot-based architecture:
  - Pilots are submitted to grid sites, w/o interruption, from a central managed system. Users don't need to worry about these pilots. Pilots check the environment, create a simplified consistent environment for user jobs , and provides for monitoring capabilities.
  - In parallel, users submit their jobs to a central service (PanDA server). Jobs are not submitted to the Grid sites directly.

PanDA

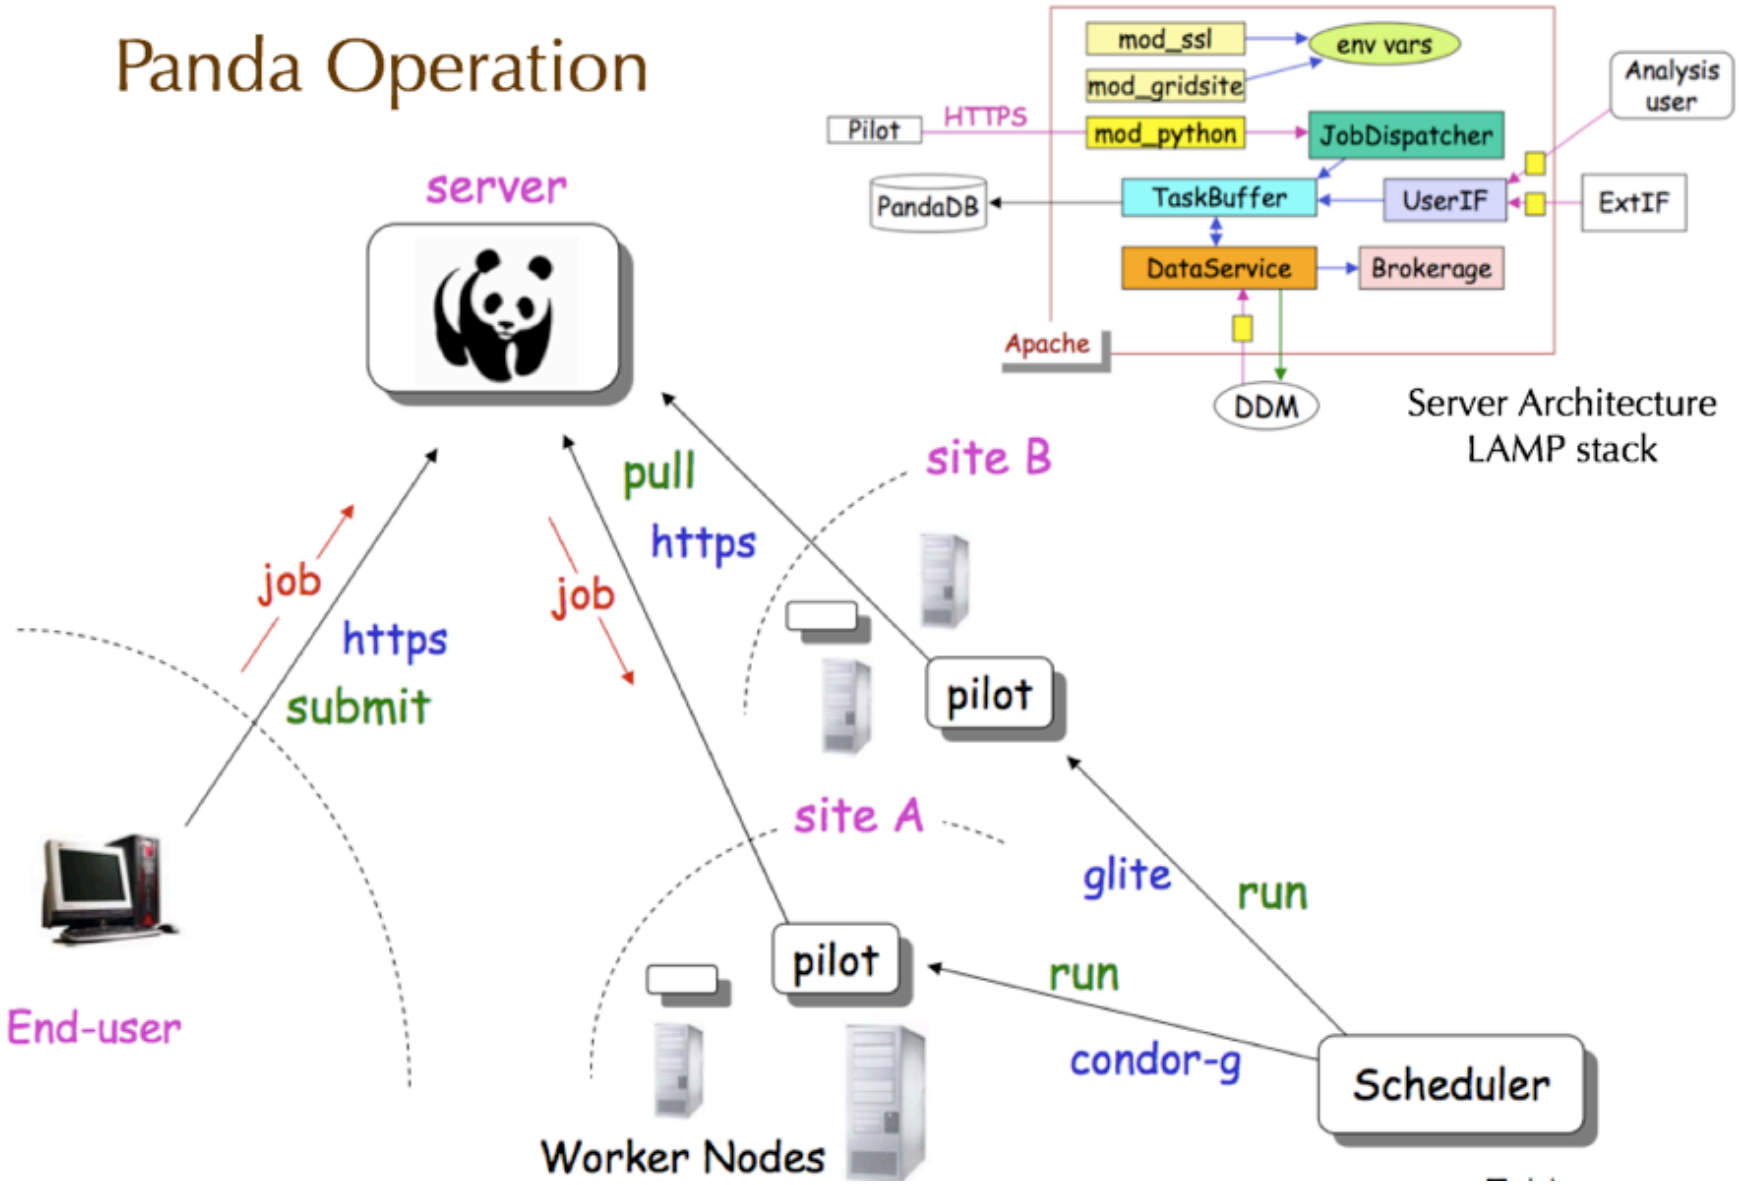
Panda Operation
server
pull
https
job
https
submit
job
site B
site A
pilot
pilot
glite
run
run
condor-g
Scheduler
End-user
Worker Nodes
mod_ssl
env vars
mod_gridsite
Pilot
HTTPS
mod_python
JobDispatcher
Analysis user
PandaDB
TaskBuffer
UserIF
ExtIF
DataService
Brokerage
Apache
DDM
Server Architecture
LAMP stack

Models of pilot architecture

- User-pilot: users submit their own pilots, one per job, with their own credentials

- Multiuser-pilot: pilots are submitted from a central service, under special credentials, and are valid for all user jobs.
  - PanDA is a multiuser-pilot based system.

Advantages of the pilot architecture

- Users do not need to worry about brokering or scheduling.
- As soon as one pilot arrives to an operative WN a job is picked up. Jobs are not waiting on queue on a site while there are free resources on other sites.
- Users do not need to concern themselves with Grid interfaces. PanDA presents them with an unified mode of access to Grid resources.

Disadvantages of the pilot architecture

- Accounting and traceability are more difficult.
- There is a risk on submitting pilots with special credentials for all end user jobs:
  - The end user job has for a while all privileges: including writing (and deleting) files, databases, etc.
  - The end user job can steal the pilot credential.

Security

To address some of the security risk in the pilot model, some security mechanism have been implemented:

- Communication between server and client is https, with authentication based on grid certificates
- Client <-> server validation, payload validation to come
  - Will prevent hijacking PanDA pilots to run unauthorized jobs
- Authorizing pilot submitters pilots themselves (beyond simply holding valid certificate) with time-limited token
- Identity switch with gLExec

gLExec

- gLExec is the “grid” version of the apache suexec.
- It changes the (UNIX) identity of the current process based on grid proxies certificates.
- The executing process will be changed from the general pilot credentials to the submitting user credentials:
  - Not allowed to modify entries in data bases restricted to the user
  - Not allowed to write/delete files in the Storage Elements restricted to the user
  - No rights to read the pilot proxy

Implementation

- When an user submits a job to the PanDA server, a copy of his/her credentials are delegated to a credentials repository (a MyProxy server)
  - Only the PanDA pilots are allowed to retrieve that proxy
  - The credentials are stored with a random key only known by the PanDA server
- The information needed to retrieve that proxy is stored in the PanDA server data base, and is provided to the pilot when a new job is requested. The pilot now knows how to retrieve the end user proxy.

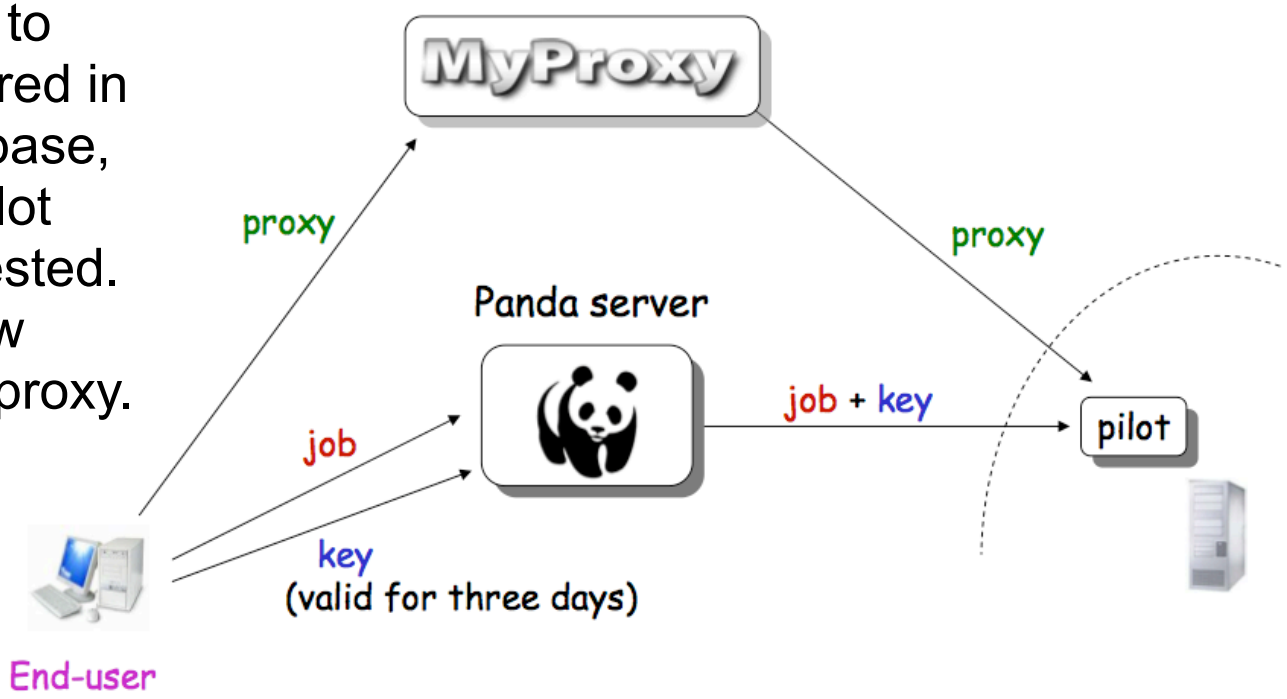

Some security mechanisms

- Is the retrieved proxy the one supposed to be?
  The pilot asks the MyProxy server for a proxy with a name (logname). We use the end user's DN as this logname. The match between this logname and real DN in the retrieved proxy is checked.

- Can anyone switch identity?
  Only credentials with Role=pilot will be allowed to invoke gLExec and switch the identity.

- Could a stolen pilot retrieve all end users' proxies?
  The same pilot only can ask for jobs to the Panda server during a limited time. A compromised pilot cannot get the required information to retrieve too many end users' proxies from the MyProxy server.

Security Service Challenge 4

- Security service challenges are run in Europe by the EGEE security team (now EGI security team)
  - Test the security procedures at sites and their ability to perform diagnostics on grid jobs.
- Challenge jobs are sent to sites:
  - These are allowed to run for some time
  - Site is alerted to ‘suspicious’ behavior, e.g., outbound contacts with certain IP address ranges and asked to investigate
  - Site is expected to follow best practice security guidelines in dealing with the incident.

Security with pilots

- Security is complicated in a pilot job environment:
  - Submitter of the payload is not usually the submitter of the pilot which is running the job
  - Given the dominance of pilot job processing by LHC experiments the security team asked ATLAS to help run SSC4
    - We saw this as an opportunity to test our own procedures and see how sites would react to a problem with jobs submitted via PanDA
- For this challenge we:
  - Setup a special pilot factory using an SSC4 certificate
    - To allow sites to ban job submission with this DN without disturbing production operations
  - Allowed a special ‘rogue user’ to submit jobs into PanDA which would be picked up by these pilots
    - Rogue payload was a mini-botnet connected back to a controller in NIKHEF

SSC4 in practice

- Discovered that many sites were able to easily use the PanDA monitor to obtained detailed information about jobs without our help
  - We were able to quickly liaise sites didn't manage this to provide job information
- We were able to provide the user's executed code to the site to assist in their investigations
  - This cold just be bootstrap code, but it provides valuable diagnostics
- We exercised ATLAS security incident handling procedures several times
  - This provided valuable training for experts

ATLAS Operational Security Lessons

- Improved the ATLAS security contact list membership
  - Now directly involves ADC experts, PanDA experts, central services team, VOMS administrators.
  - Easier to track and deal with incidents.
- Improved security procedures
  - Triage of security incidents is now the responsibility of the ADC Manager on Duty
  - Security handling instructions in twiki
- Improved security tools
  - Simple backlisting of a suspected rogue user in PanDA
    - Cancels all current jobs and prevents new submissions
    - This is faster and more easily reversed than certificate revocation (though more limited in scope)