

AutoPyFactory Recap, Status, Roadmap
US ATLAS Technical Planning Meeting
Chicago, Aug 1, 2016

John Hover <jhover@bnl.gov>

Jose Caballero <jcaballero@bnl.gov>

Recap

- Pilot factory for PanDA; source code in PanDA org GitHub repos.
- Python, single-process, multi-threaded, automated, object-oriented daemon.
- Reliable, stable, and generic platform for development of new functionality.
- “Everything is a plugin” modular architecture makes new features easy and safe.
- Chainable scheduler logic plugins allow “algorithms from config,” to cover novel/unusual needs.
- When HTCondor [WMS, batchsubmit, batchstatus] plugins are used, leverages all dev effort, auth mechanisms, and resource targets (CEs, Clouds, SSH) HTCondor provides.
 - But HTCondor is ***not required*** if you 1) don’t need persistence or 2) you’re willing to write the functionality into the plugin.
- RPMs packaged and published to YUM repos by OSG (as part of PanDA).
- Reporting to ATLAS APF Monitor [by Peter Love]: <http://apfmon.lancs.ac.uk/mon/>

Status

- Used by ATLAS worldwide for PanDA pilot submission (version 2.3 and 2.4)
 - Centrally from CERN and BNL to grid sites in Europe and US
 - Locally into Openstack VM pools (IN2P3, Sim@P1, BNL)
 - Locally into CloudScheduler-provisioned Condor pool(s) at TRIUMF
 - Locally into Tier 3s? Doug B.?
- Used by ATLAS for VM provisioning (and pilot submission).
 - Amazon project.
- Used by BNL for VM provisioning for Biology clusters.
- Currently being set up for VC3 project at U. Chicago:
 - VM provisioning
 - SSH vc3-pilot submission
 - Grid/CE vc3-pilot/glidein submission

APF Roadmap I: v2.5 October 2016(?)

v2.5 (definite)

- Fully automated remote factory configuration and service restart. (ATLAS and VC3)
- Bosco (Condor-G submission via SSH) integration with full authentication account profiles, i.e. different targets use different SSH keypairs. (VC3)
- Auth profiles include X509, SSH keypairs, and Cloud tenant keys--all used uniformly.
- Fully automated factory installation, config, and restart integration within VC3.
- Enhance idle job cancellation functionality to allow more flexible responses.
 - Eventually we will want the option to cancel pilots/VMs that haven't started yet, so we can submit more elsewhere.
- Remotely deployable no-Condor non-root package, for MFA/SSH head nodes.
 - Runs on HPC or cluster head node, submitting into local batch system.

v2.5 (tentative)

- Panda payload+pilot submission via SSH [contingent on job delegation]
 - Same as what aCT is doing, but would require “no-heartbeat” mode.
 - This could work even on resources without outbound network: just SSH.

APF Roadmap II: v3.0 February 2017(?)

v3.0 (definite)

- Policy-driven provisioning and/or job submission. (VC3, HEPCloud)
 - I.e. system consumes a machine-parsable description of desired provisioning policy. Run up to W on resource A, or if it is full (for some definition of full) before reaching X jobs, spill over to resource B (up to a maximum of Y jobs). And don't run more than Z jobs, per month, in aggregate on resource B. After the Z limit is reached, submit to resource C.

V3.0 (tentative)

- Encapsulation within library (e.g. libapf.py) for direct in-process embedding within larger frameworks (e.g. HEPCloud) [contingent on HEPCloud project]
- We can definitely implement additional integration with PanDA to cover most of the (non-HPC) Harvester concerns. Some of these are already VC3 and HEPCloud requirements.

References

Recent APF update for ATLAS TIM:

<https://docs.google.com/presentation/d/1VyW3GAYESYEG3U-XZ4QuyVcaPV88YPqpm-LyQJ11mDw/edit?usp=sharing>

Recent APF orientation for VC3 group:

<https://docs.google.com/presentation/d/1JZluZ1xpbp50TmtnsWa4kEPuASRccAvyWHzECeA6WNA/edit?usp=sharing>

r