



Managed by Fermi Research Alliance, LLC for the U.S. Department of Energy Office of Science

GlideinWMS

Parag Mhashilkar
Stakeholders Meeting
June 28, 2017

Overview

- Updates since last stakeholder's meeting
- Upcoming releases
- Reference slides
 - GlideinWMS Architecture
 - Quick Facts
 - Releases since last stakeholders meeting

Highlights Since Last Stakeholders Meeting

- Announcements
 - New Hire - Process in final stages
- Project Effort (2.85 FTE)
 - Project Management: 0.15 FTE
 - Development & Support: 2.70 FTE
 - Thomas Hein: Summer student through Extended Target program at FNAL
 - Working on code migration to Python 3 migration
- Communication
 - Regular participation of CMS in the weekly GlideinWMS meetings
- Support
 - HEPCloud uses GlideinWMS factory as the provisioner
 - Increasing use of BOSCO to submit to NERSC has highlighted several issues within BOSCO/HTCondor
 - We are working with the HEPCloud team and HTCondor team to address these issues

Highlights Since Last Stakeholders Meeting

Year	Production Series (v3.2)	Development Series (v3.3)
2016	3.2.12 - 3.2.16 versions=7, features=27, bugfixes=53	v3.3 - v3.3.1 versions=2, features=5, bugfixes=1
2017	v3.2.17 - v3.2.19 versions=3, features=20, bugfixes=25	v3.3.2 versions=1, features=1, bugfixes=2

- Releases: (Details in Reference Slides)
 - v3_2_19: May 30, 2017
 - Improvements to scaling down glideins when job pressure drops
 - Per-frontend unique idle job counts – Older than X hours
 - Log number of activation/claims per glidein
 - Support newer versions of HTCondor, new configuration formats and tools
 - Automatic CPU detection specially for glideins running in the PBS
 - Several bug fixes
 - v3_3_2: March 24, 2017
 - Based on v3_2_18
 - Allow multiple remote directories for BOSCO submissions
 - Fixes issues affecting glidein submission to AWS

Upcoming Releases – Milestones for series v3.2

- Primary Focus of Production Series:
 - High impact bug fixes and features that do not break backward compatibility

v3_2_20 – Tentatively by second half of July 2017

- Singularity support
- HTCondor 8.6 support (works with 3.2.19, jobs are running, but there is an error observed in the HTCondor logs)

v3_2_21+

- Continue improving the tracking of needed resource during ramp-down to reduce waste
- Balancing requests to factory entries that are aliases or meta sites
- Scale factory to O(600+) entries
- Move to full use of HTCondor shared port daemon to reduce the port used in communication and because HTCondor is pushing in that direction
- Improve troubleshooting. Why is my job not running?
- Modernize the code (Python 2.6/2.7 standards) eliminating old syntax and making easier the porting to Python 3 (long term)

Upcoming Releases – Milestones for series v3.3

- Primary Focus of Development Series:
 - Production quality but some features maybe experimental
 - Milestones and timelines expected to be primarily driven by the HEP Cloud Decision Engine requirements
- Next Release: v3.3.3

v3_3_3 – Tentatively by end of July 2017

- Based on 3.2.20
- Manage HPC jobs that spawn multiple nodes (currently running via work-around and not correctly counted in provisioning)
- Plan for moving v3.3 series to production late summer/early fall 2017
 - Final v3.3 release becomes v3.4

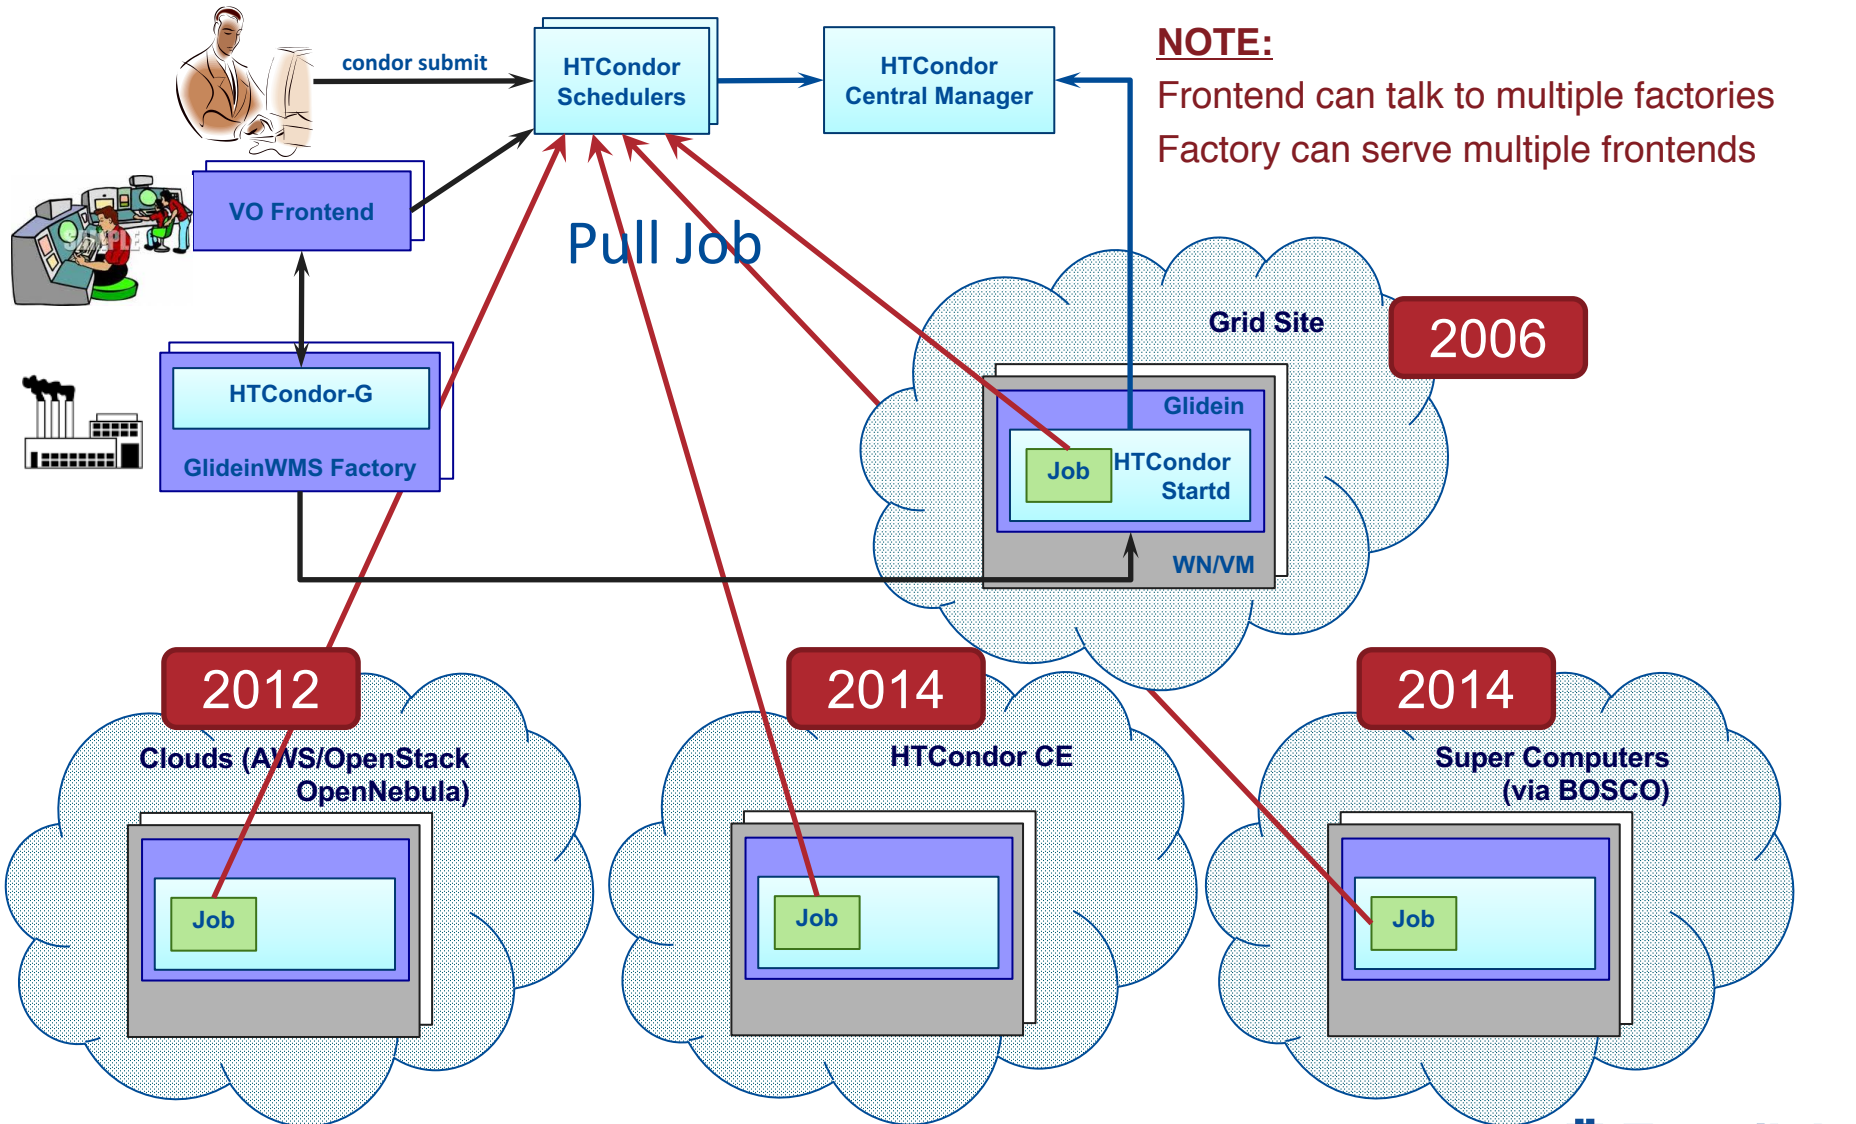
Other Milestones

- Improved support for HPCs
 - BOSCO
 - Edge services
- Improvements to troubleshooting of jobs
 - Retrieve pilot and custom logs
 - From Clouds & BOSCO
- Revisit condor_annex for bulk provision of resources
- Improve Factory and Frontend monitoring
 - Grafana
 - HTCondor's MetricsD

Questions/Comments

Reference Slides

GlideinWMS



GlideinWMS: Quick Facts

- GlideinWMS is an open-source product (<http://tinyurl.com/glideinWMS>)
- Heavy reliance on HTCondor (UW Madison) and we work closely with them
- Effort:

Role	Resources	Effort (FTE)
Project Mgmt/Lead	Parag Mhashilkar (0.15 USCMS)	0.15
Development & Support	Parag Mhashilkar (0.20 SCD) Marco Mambelli (1 SCD) Hyunwoo Kim (0.5 SCD) Dennis Box (0.5 SCD) Marco Mascheroni (0.5 CMS - Contractor)	2.70
TOTAL		2.85

Table: Current Resources & Roles

- Additional Code Contributions from collaborators

Quick Facts: Releases & Support Structure

- Releases
 - Issues tracked in redmine issue tracker
 - <https://cdcv.s.fnal.gov/redmine/projects/glideinwms/issues>
 - Categorized and prioritized based on impact, urgency and requester
 - Issues are now associated with respective stakeholders
 - Issues are assigned based on developer's expertise and other workload
 - Roadmap for upcoming releases available in redmine (See reference slides)
 - SCM
 - All releases are version controlled and tagged
 - <http://glideinwms.fnal.gov/doc.prn/download.html>
 - Release notes & history
 - <http://glideinwms.fnal.gov/doc.prn/history.html>
- Support
 - Entire development team is responsible for support

Quick Facts: Project Status & Communication Channels

- Project meeting: Wednesdays 10 – 11 am
 - Technical discussions & status updates
 - Regular stakeholder participation
 - Contact Parag Mhashilkar if you need invite for this meeting
- Quarterly Stakeholders Meeting
- Project Management
 - Project Status reported monthly at CS Project status meetings

Area of Interest	Mailing Lists
Support	glideinwms-support@fnal.gov
Stakeholders	glideinwms-stakeholders@fnal.gov
Release Announcements	glideinwms-support@fnal.gov cms-dct-wms@fnal.gov glideinwms-stakeholders@fnal.gov
Future Release plans	See next slide
Discussions	glideinwms-discuss@fnal.gov
Code commits	glideinwms-commit@fnal.gov Twitter Tag: @glideinwms

Tracking Releases in Redmine

1. Visit the redmine issues tab for GlideinWMS or the URL

Default tabs not too useful

2. Click custom query for stakeholder or version roadmap

Issues

My custom queries

Custom queries

- CMS Tickets
- FIFE Tickets
- OSG Tickets
- One Facility Tickets
- Operations Tickets
- CMS Master Ticket
- Igor Sfiligoi's Tickets
- Issues Awaiting Developer Feedback
- Issues by Age
- John Weigand's Tickets
- List With Due Date
- My Open Issues
- Not Closed + Grouped by Assigned To
- Not Closed + Grouped by Target Version
- Open Bugs: High Priority
- Series v2.7.x
- Series v3.2.x
- Series v3.x
- Stakeholder Requests
- Tickets with no version
- Unassigned Tickets
- v3.2.10: Roadmap
- v3.2.9: Open Issues
- v3.2.9: Road Map
- v3.3: Road Map

#	Tracker	Target version	Status	Priority	Subject	Assignee	Updated	OSG
2409	Feature	v3_2_10	New	High	Support validation scripts to run periodically (not just at pilot start)	Marco Mambelli	05/13/2015 01:26 PM	OSG VO
6897	Bug	v3_2_10	Feedback	High	Partitionable glideins not accounted for correctly	Parag Mhashilkar	05/12/2015 04:42 PM	OSG
7927	Feature	v3_2_10	New	High	Per-frontend unique idle job counts	Marco Mambelli	04/13/2015 04:16 AM	OSG
4904	Feature	v3_2_10	Assigned	Normal	Specify a default for GWMS_XSLT_PLUGIN_DIR	Marco Mambelli	04/01/2015 01:43 PM	OSG
4846	Feature	v3_2_10	Assigned	Normal	Include GWMS_XSLT_PLUGIN_DIR setting in generated {factory,frontend}_startup	Marco Mambelli	04/01/2015 01:42 PM	OSG
7807	Feature	v3_2_10	New	Normal	Support CONDORCE_COLLECTOR_HOST for OSG sites	Parag Mhashilkar	04/13/2015 03:21 AM	OSG
8437	Feature	v3_2_10	New	Normal	Split entry configuration from main factory config	Parag Mhashilkar	04/27/2015 12:13 PM	OSG, CMS
4587	Bug	v3_2_x	Assigned	Normal	OSG: Automate / fix HTCondor-CE issues	Burt Holzman	02/03/2015 02:44 PM	OSG
6732	Bug	v3_2_x	New	Normal	Glideins report 100% success rate even if startd fails to start/run a matched job.	Marco Mambelli	03/04/2015 04:47 PM	CMS, OSG
3203	Feature	v3_2_x	Assigned	Normal	OSG: condor g -analyze	Burt	02/03/2015 02:44 PM	OSG

GlideinWMS Releases - Key Features

v3_2_18 – May 30, 2017

- NOTE: If using HTCondor 8.6.x you have to disable the collector use of shared port
- Added counters for Idle jobs older than X hours
- Enabling GWMS_XSLT_PLUGIN_DIR by default for glideinwms-vofrontend-standalone rpm
- Linked Frontend monitoring from Factory monitoring
- Improved glideins scale down by adding a timeout, GLIDEIN_IDLE_LIFETIME
- Log number of activation/claims per glidein
- Several documentation improvements, including reconfig and upgrade operations in SL7 RPM installation
- Bug Fix: reduced the number of file descriptors used per Entry in the Factory
- Bug Fix: Factory entries submitting glideins even after hitting the limit
- Bug Fix: Errors with HTCondor 8.5/6 upgrade: SUBSYS.LOCALNAME.* warning triggered by GWMS htcondor configuration, Glidein sent unintentionally to multiple schedds on the Factory
- Bug Fix: Clarified attribute types and fixed globbing behavior
- Bug Fix: RPM verification fails when config files are changed
- Bug Fix: GLIDEIN_CPU settings "node" and "slot" supported also in Frontend, better documented and improved to compensate for PBS misconfiguration
- Bug Fix: Fixed submission to AWS which uses key_pair as auth_method

GlideinWMS Releases - Key Features

v3_3_2 – March 24, 2017

- Includes all features and bug fixes released in v3_2_18
- Allow multiple remote directories for BOSCO submissions
- Bug fix: Submit attributes in entry configuration are now transmitted to AWS VM
- Documented vm_id_fname and vm_type_fname
- Bug Fix: Fixed submission to AWS which uses key_pair as auth_method