



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

GlideinWMS

Parag Mhashilkar
Stakeholders Meeting
December 07, 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


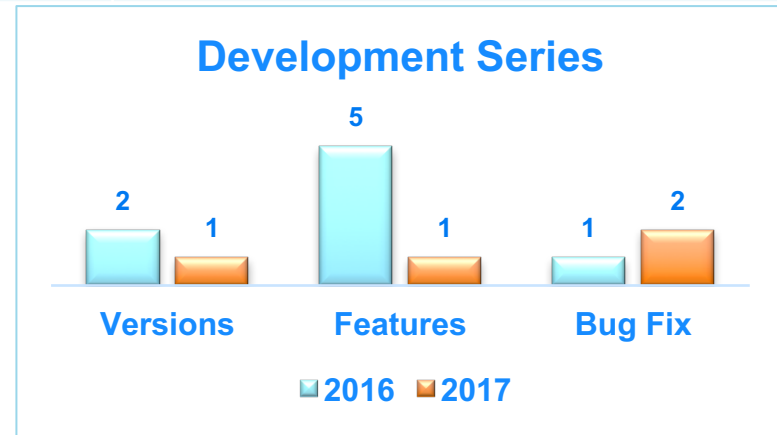
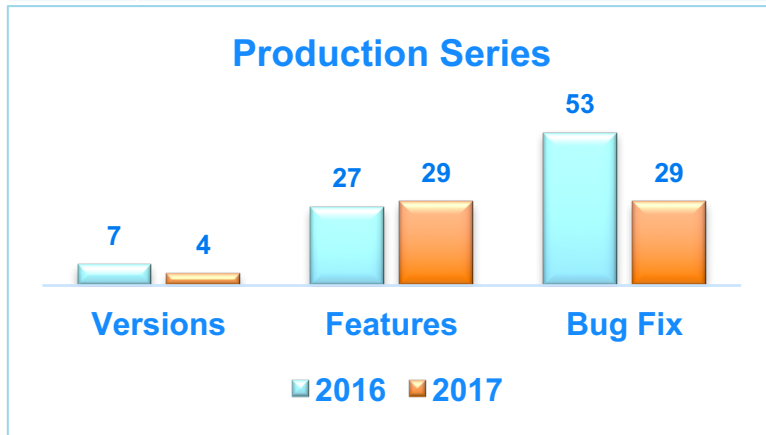
- Announcements
 - New Hire: Lorena Lobato Pardavila is expected to join on January 22, 2018
- Project Effort (2.60 FTE)
 - Reduction in FTE in 2017
 - Trying to adapt with new hire and increasing fraction of current team members
 - New members need time to ramp-up on the project
 - Project Management: 0.15 FTE
 - Development & Support: 2.45 FTE (↓0.25 FTE)
 - Marco Mambelli: 1 FTE
 - Dennis Box: 0.75 FTE (recent change: +0.25 FTE until March 2018)
 - Marco Mascheroni: 0.5 FTE (CMS contractor until December 2017(?))
 - Parag Mhashilkar: 0.2 FTE
 - Hyun Woo Kim: -0.5 FTE

Highlights Since Last Stakeholders Meeting

- Communication
 - Regular participation of CMS and Factory Operations in the weekly GlideinWMS meetings
 - GlideinWMS people started to attend CMS SI and CMS-HTCondor meetings
 - How can we further improve communication
 - Should we participate in any other meetings?
 - Communicating priorities?
- Support
 - Increased pull requests using GitHub from the contributors
 - GitHub started as the mirror repository only
 - We need to internally work on direction here
 - Till then we will be monitoring these requests better in future

Highlights Since Last Stakeholders Meeting

Year	Production Series	Development Series
2016	3.2.12 - 3.2.16	v3.3 - v3.3.1
2017	v3.2.17 - v3.2.20	v3.3.2



Highlights Since Last Stakeholders Meeting

- Releases:
 - v3_2_20: November 15, 2017
 - Support for Singularity
 - Monitor cores usage statistics in factory monitoring
 - Factory to handle 1500 entries on startup
 - Better support for SL7's init reload command
 - Remove dependency on Globus clients & libraries
 - Several bug fixes

Detailed Release Notes:

<http://glideinwms.fnal.gov/doc.prd/history.html>

<http://glideinwms.fnal.gov/doc.dev/history.html>

GlideinWMS Roadmap

- 2018 – 2019
 - Minimize wastage of resources by over-provisioning
 - Consider site physiology
 - Actively follow the requests
 - Improve support for HPC
 - Bulk provisioning
 - New Cloud providers as required
 - Microsoft Azure (?)
 - Improve support for sites with tighter security restrictions
 - Example: Support for shorter proxy lifetime
 - Benefit from HTCondor improvements
 - Python 3 migration readiness
 - Modernizing monitoring
 - Make troubleshooting easier
 - Stay aligned with the HEPCloud project

HL-LHC

- HL-LHC scales predicted to be ~ 50 times the present scale
 - Biggest impact on storage and network IO
 - How will this impact WMS?
 - Will it result in 50 times the number of jobs? Maybe not?
 - CMS Production
 - Progress in multicore jobs: 8 core jobs as default
 - CMS Analysis
 - Started with nano AOD for analysis.
 - If successful, event size 1KB v/s 30KB
 - Implies faster computation
 - Faster and bigger machines
 - Possibility to move to whole-node scheduling & glideins

HL-LHC: GlideinWMS Roadmap

- Improve scalability
 - Benefit from and support enhancements in HTCondor where ever applicable
 - Investigate and incorporate new technologies like pandas dataframes, numpy, etc
- Better support for commercial and allocation based resources: Clouds & HPC
 - Bulk provisioning
 - Investigate direct submission
 - Edge services, specially for HPC
 - HEPCloud
 - GlideinWMS to benefit from lessons learned in HEPCloud
 - Minimize unused resources
- Better Communication & Collaboration
 - Benefit from and support enhancements in HTCondor
 - Active communication with the stakeholders to understand their direction

Upcoming Releases – Milestone for Series v3.2

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

v3_2_21 – Mid December 2017

- Frontend to request the removal of unused glideins
 - Work will be split into smaller chunks
- Balancing requests to factory entries that are aliases or meta sites
- Support for unprivileged singularity

v3_2_22+ – Timeline for the releases to be decided based on stakeholder needs & input

- Improvements to glideins removal and meta-sites as required
- Replacement for condor_switchboard
- Handle proxies with limited/short lifetime
- Monitor number of job activation/claims per glidein
- Code modernization - getting ready for Python 3

Upcoming Releases - Milestone for Series v3.3

- Primary Focus in the Development Series:
 - Production quality but some features maybe experimental
 - Milestones and timelines expected to be primarily driven by the HEP Cloud Decision Engine requirements
- Plan for moving v3.3 series to production in 2018
 - Delayed from earlier planned summer/early fall 2017
 - Final v3.3 release becomes v3.4

v3_3_3 – Tentatively January 2018

- Based on 3.2.21
- Manage HPC jobs that spawn multiple nodes (currently running via work-around and not correctly counted in provisioning)

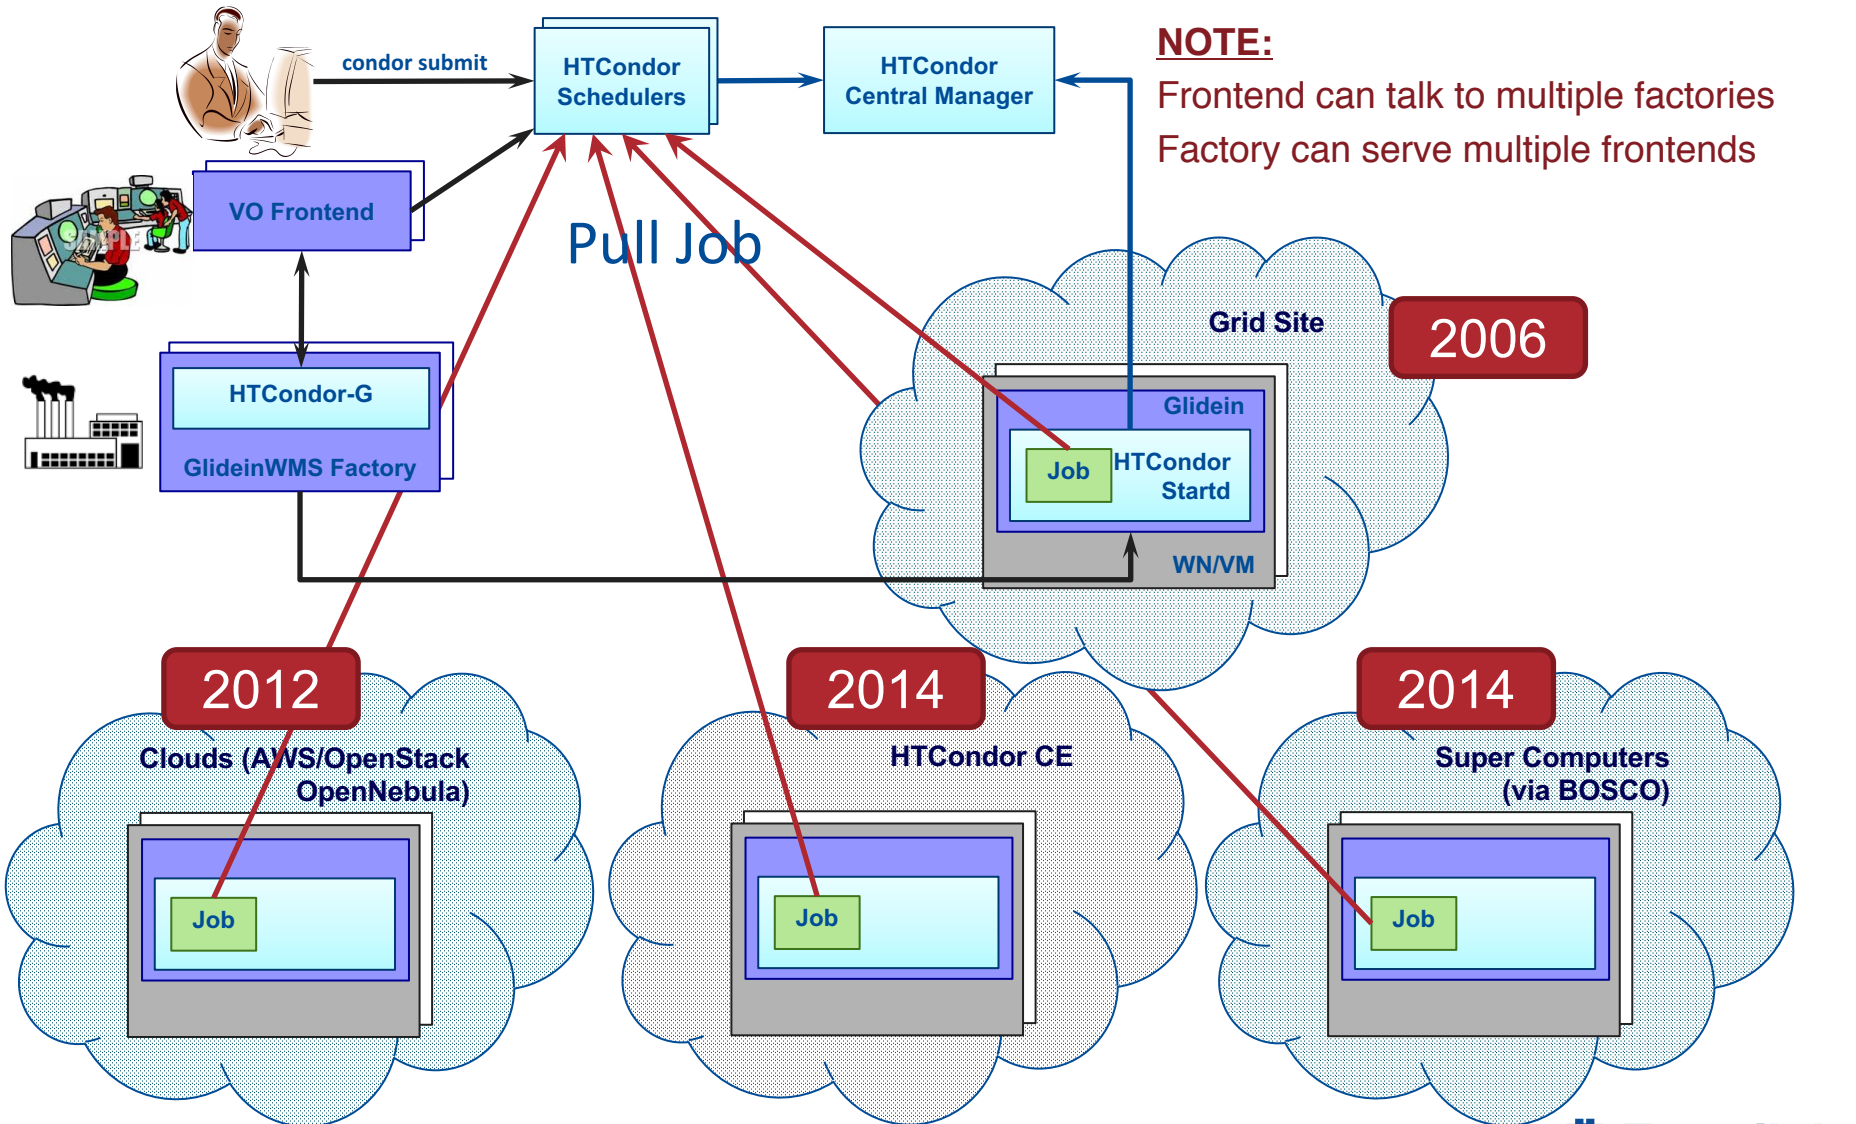
Other Milestones

- Previous Milestones
 - “Why is my job not running”?
 - Scheduled for ~~v3.2.19+~~ v3.2.22+
 - Other requests take priorities
 - Not enough development effort to focus on this ticket
 - Continuous Integration
 - More unit tests – ongoing activity
 - Automate deployment
 - Automate testing

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) Dennis Box (0.75 SCD) Marco Mascheroni (0.5 CMS - Contractor)	2.45
TOTAL		2.60

Table: Current Resources & Roles

Quick Facts: Releases & Support Structure

- Releases
 - Issues tracked in redmine issue tracker
 - <https://cdcv.sfnal.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.sfnal.gov/doc.prd/download.html>
 - Release notes & history
 - <http://glideinwms.sfnal.gov/doc.prd/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