



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

GlideinWMS

Marco Mambelli

Stakeholders Meeting

May 8, 2019

Overview

- Completed and Upcoming releases
- GlideinWMS roadmap
- Developers spotlight
- Reference slides
 - GlideinWMS Architecture
 - Quick Facts

Completed and Next Planned Releases

- Released
 - GlideinWMS v3.4.5 was released on April 17 and is released in OSG 3.4.28. This follows GWMS 3.4.2
- We have 1 release close to completion
 - v3.5 w/ single-user Factory, HTCondor started Singularity, for OSG upcoming, now planned for end of May. Delayed by 3.4.5 and changes in HTCondor handling of Singularity

Completed Release, v3.4.5

- v3_4_5 in OSG production (OSG 3.4.28, released 4/25)
 - Fixed Error preventing the Frontend to match jobs
 - Singularity improvement (include system files, OSG distributed binary)
 - Propagate to Factory and glidein submission attributes controlled by FE (HEPCloud)
 - Multi-node jobs accounting (CMS, OSG)
 - Fixed Glidein not killing HTCondor processes (OSG, CMS)
 - Joint effort w/ Diego (CMS) and Eric (Purdue) and OSG
 - Fix problems with Factory monitoring when there are no Frontends (HEPCloud)

https://cdcv.sfnal.gov/redmine/projects/glideinwms/issues?query_id=26

https://cdcv.sfnal.gov/redmine/projects/glideinwms/issues?query_id=53

Completed Release, v3.4.5 - NOTES

- Major Singularity features were introduced in GlideinWMS 3.4.1
 - To use them all factories and frontends need to be $\geq 3.4.1$
- HTCondor configuration changes announced in the release notes. Do not ignore that.
 - Are integral part to providing some functionality
 - Those are the tested configurations
- Enables shared port, allowing to require only port 9618
 - To ease the transition to shared port, the User Collector secondary collectors and CCBs support both shared and separate, individual ports
 - VOs started testing shared port usage. Update the User Collector configuration!

See also NOTES DETAIL in the Reference Slides and

<https://opensciencegrid.org/docs/release/3.4/release-3-4-28/>

Next Planned Release, v3.5

- v3_5 delayed to end of May, for OSG upcoming
 - Dropping Globus GRAM support
 - Single-user Factory
 - Invoke Singularity via HTCondor
 - Condor now allows custom parameters that will allow this
 - Will allow condor_ssh_to_job if unprivileged Singularity is used
 - Black hole prevention
 - Automate the generation of factory configuration via CRIC
 - Frontend matching performance improvement

https://cdcvs.fnal.gov/redmine/projects/glideinwms/issues?query_id=186

GlideinWMS Roadmap – dropping support for...

- Scheduled for 3.5
 - GRAM GT2/GT5
- Planned for 3.6 (possibly some 3.5.x - Summer)
 - GLExec
 - Separate User collector ports (only shared port)
- Planned for 3.7 (late Summer - 3.6 will be in parallel until late Fall)
 - Python2
 - Is it OK to move to support only Python 3 by the fall?

GlideinWMS Roadmap – high priority

- Move to Python 3
 - Branch with Python 3 migration
 - Have a Python 3 version version in OSG upcoming by late Summer 2019
- Factory supporting multiple frontend like services
 - Decision Engine support started in 3.4.4
- Collaboration with HTCondor
 - Black hole prevention (3.5)
 - Singularity invocation (3.5)
 - Use of tokens (security without x509 certificates)
- Automatic Factory configuration generation, via CRIC (3.5)

<https://cdcv.s.fnal.gov/redmine/projects/glideinwms/wiki/RoadmapSummary>

GlideinWMS Roadmap - other

- Monitoring Modernization
 - Retire GlideinWMS monitoring pages
 - Move to grafana/graphite/elastic search based solution
- Collaborate with HTCondor team to support new HPC sites with stricter policies (e.g. no outbound connection except gateways, MFA)
- Move of the documentation to Jekyll (Summer program)
 - Use of templates will ease page maintenance
- Deploy GlideinWMS in containers

Developers Spotlight

Marco Mambelli

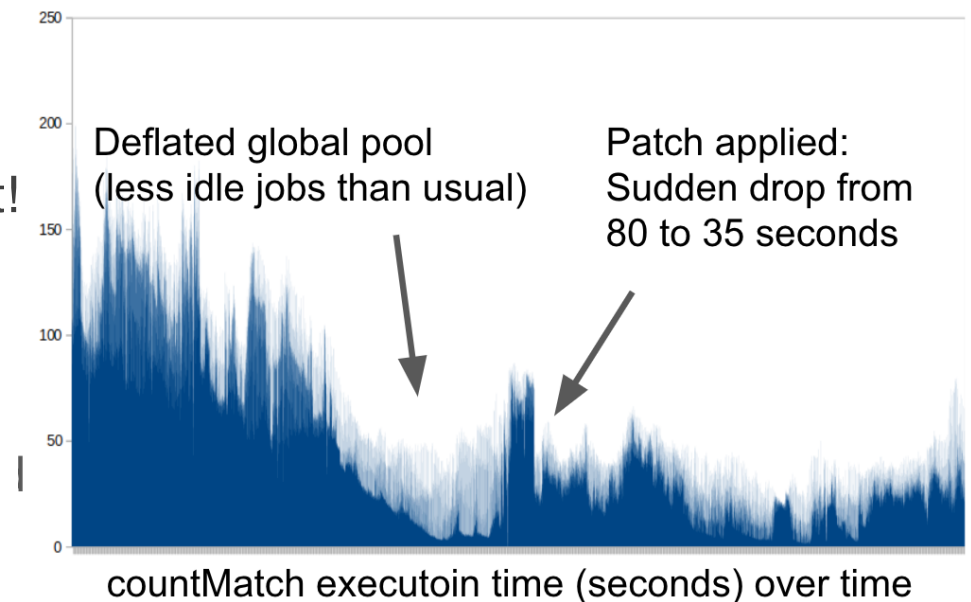
- Lead FIFE-Containers working group (Fermilab)
- Return to the monitoring discussion
- Joint effort to solve HTCondor not being killed in PBS clusters
- Monthly code discussion and challenge of the month
- Summer projects
 - Monitoring
 - Improved Glidein functionality (error reporting)
 - Migrating documentation to Jekyll
- Development topics
 - Singularity
 - invocation via HTCondor in 3.5
 - Easy VO scripts for testing and setup

Lorena Lobato

- Attributes controlled by Frontend can be now propagated to the glidein submission - HEPCloud ([GWMS 3.4.4](#))
- Removal of dependency of condor_root_switchboard: Single user Factory ([expected for GWMS 3.5](#))
- Investigating periodic scripts using prefix inconsistently
- Working on blacklist & blackhole detection
 - Interaction with HTCondor team to support the integration of new stats that will help to identify blackholes.
 - Using FIFE team as use case
 - More complete information in the logs
 - Solution for blacklist script and preventive measures to avoid back-hole effects

Marco Mascheroni - CMS scale tests: frontend improvements

- Matching function proved to be a critical point during last CMS scale tests
 - Wrote code to save a snapshot data structures, and used to retrieve real production data
 - Used production data and cprofile to individuate parts of the code that needed improvements
 - Cached arithmetical operation in inner loop previously executed $O(J_2^*E)$, and now executed $O(J^*E)$ [J=Job clusters, E=Entries]
- Profiling showed an execution time more than 50% faster
- Patch applied in production
 - Improvements immediately evident!
- More improvements needs major refactoring (is it worth considering the code has already been replaced In the decision engine?)



Marco Mascheroni - Other activities

- Started rolling out in production automatic generation of factory xml from CRIC
 - Verified it works in ITB on UCSD entry
 - Adding other entries (plan to have ~20 by July)
- Added the possibility to ignore entries in downtime when calculating frontend pressure
 - Cause of frontend low pressure calculations
- Presented recent developments in gliedeinWMS at Hepix 2019. See [here](#)

Dennis Box

- Dennis Box
- Code quality, testing, github migration
 - Containerized CI – using github, travis-ci, docker-hub exclusively
 - Source for CI tests at
 - <https://github.com/ddbox/gwms-test>
 - Checkins to github cause a CI build
 - <https://travis-ci.org/ddbox/gwms-test>
 - CI build loads containers to
 - <https://hub.docker.com/r/dbox/gwms-test>
 - CI build also loads test artifacts, reports, etc, to
 - <https://github.com/ddbox/gwms-test/documents/test-results>
 - Html artifacts do not display properly here,
 - Need to be forwarded to a web server
- CI → CD Idea:
 - build RPMs at CI stage
 - 'smoke test' them for basic functionality
 - using existing scripts

Questions/Comments

Reference Slides

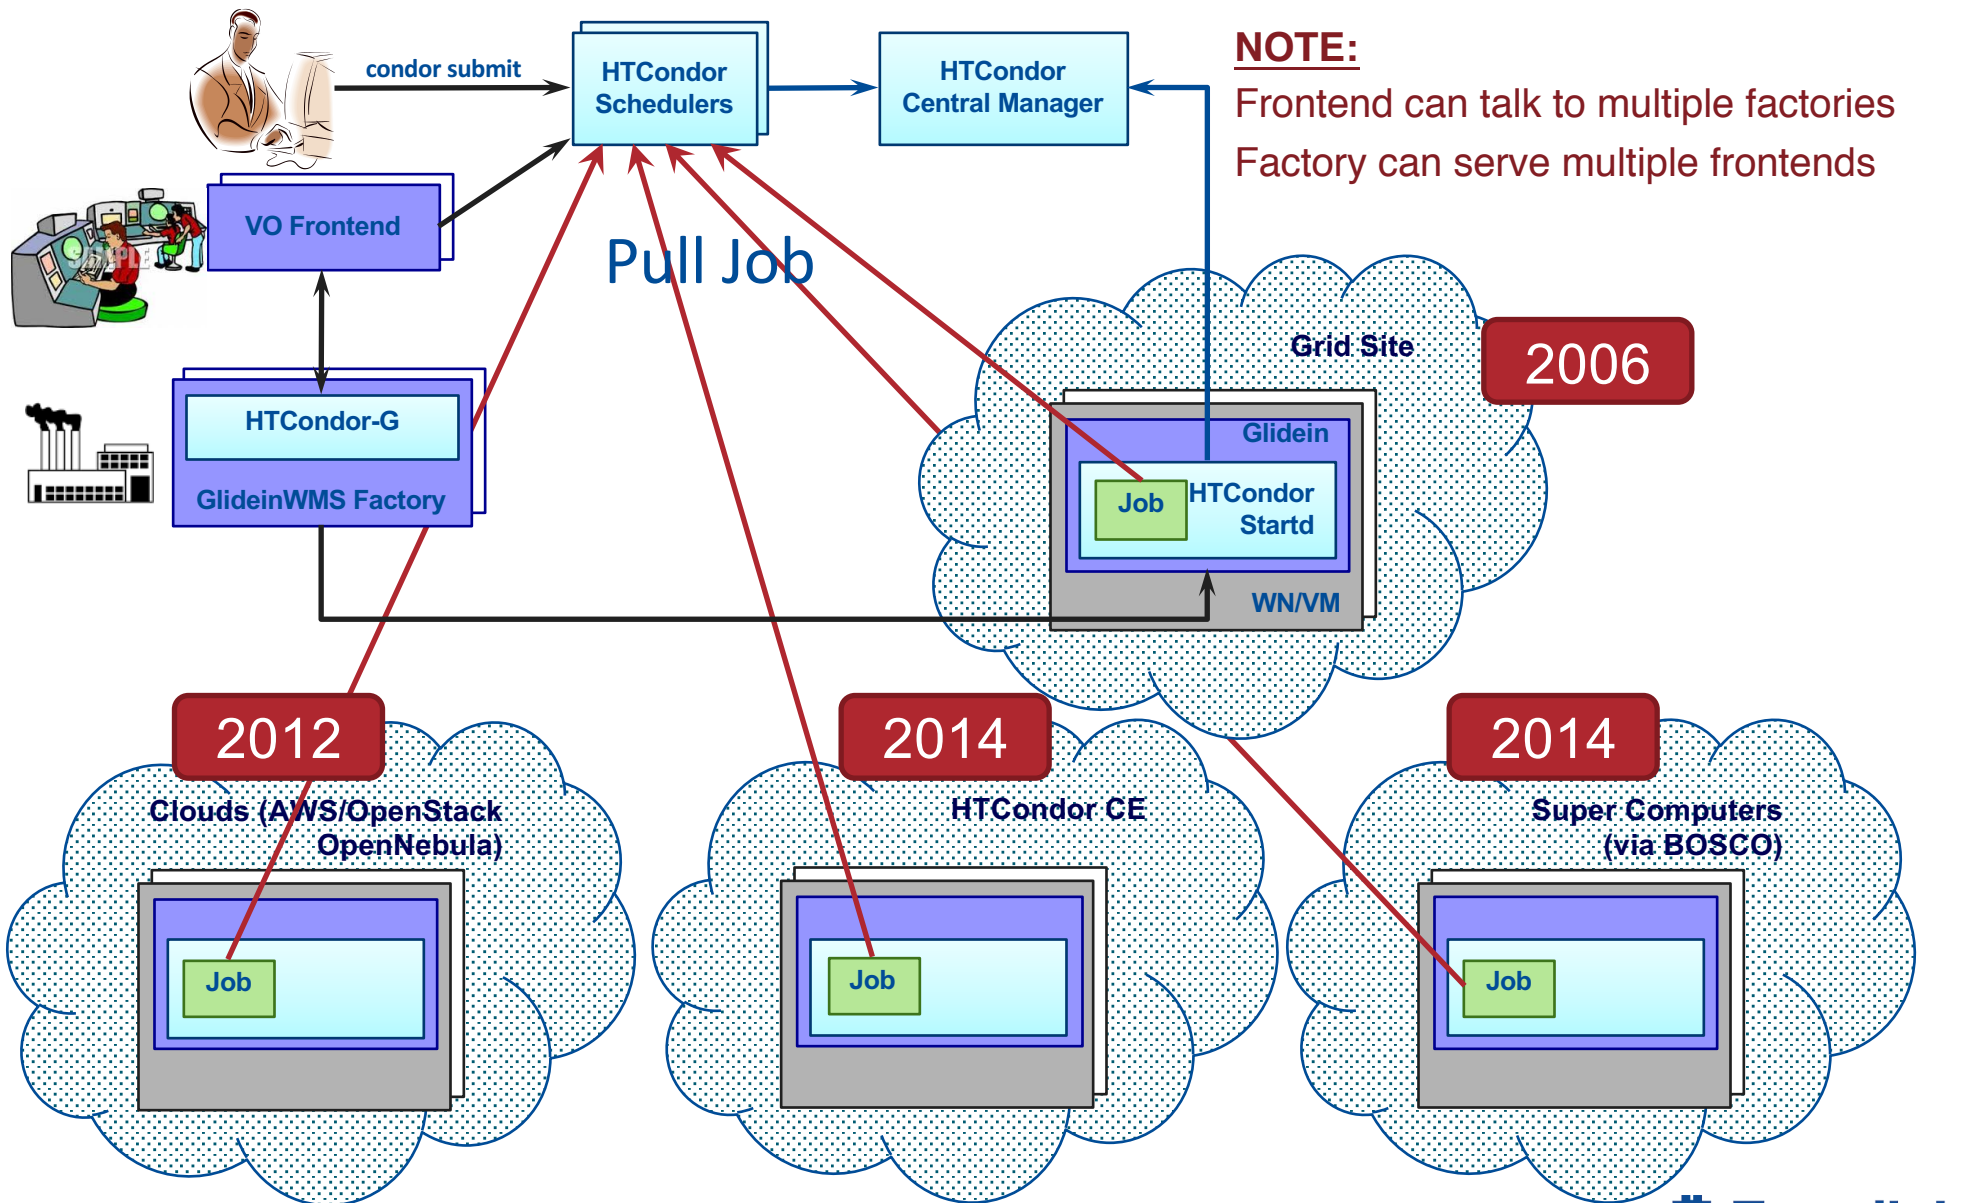
Completed Release, v3.4.5 – NOTES DETAIL

- For new Singularity features introduced in GlideinWMS 3.4.1, all factories and frontends need to be $\geq 3.4.1$.
 - OSG GlideinWMS factories are running at least 3.4.1
 - If some of the connected Factories are $\leq 3.4.1$ you will see an error during reconfig/upgrade if you try to use features that require a newer Factory. To start using Singularity via GlideinWMS, see:
 - <https://glideinwms.fnal.gov/doc.prd/frontend/configuration.html#singularity>
 - <https://glideinwms.fnal.gov/doc.prd/factory/configuration.html#singularity>
 - https://glideinwms.fnal.gov/doc.prd/factory/custom_vars.html#singularity_vars
- Upgrades may require merging /etc/condor/config.d/*.rpmnew files and a restart of HTCondor (check /etc/condor/config.d). Or updating of your separate HTCondor config
- Enables shared port, allowing to require only port 9618. To ease the transition to shared port, the User Collector secondary collectors and CCBs support both shared and separate, individual ports. To start using shared port, change the secondary collectors lines and the CCBs lines (if any) in /etc/gwms-frontend/frontend.xml, changing the address to include the shared port sinful string:
 - `<collector DN="/DC=org/DC=opensciencegrid/O=Open Science Grid/OU=Services/CN=gwms-frontend.domain" group="default" node="gwms-frontend.domain:9618?sock=collector0-40" secondary="True"/>`
 - Replacing gwms-frontend-domain with the hostname of your GlideinWMS frontend. See the [GlideinWMS documentation](#) for details.

Move to single user Factory

- Will be in the next release, v3.5
- All Glideins will run using the factory user, no more separate users per-VO
 - Currently different VOs (Frontend groups) can use different users to improve isolation
- It is safe
 - The HTCondor team assured us that once we remove Globus GRAM support, the other Gridmanager clients cannot decide which file to retrieve from the Factory (it is HTCondor on the Factory deciding what to send), so will be safe to run as a single user
- The directory structure will remain the same
 - Only the ownership will change
 - Your log files will be in the same place
- Migration:
 - GWMS will provide instructions and tools to ease it: change the files ownership, ...
 - if you use HTCondor < 8.7.2 you can upgrade GWMS when convenient for you
 - if you need HTCondor >= 8.7.2 (including 8.8) we recommend to upgrade
 - but if you want to delay the change to 3.5 you can still do that if you are comfortable in using the glideinwms-root-switchboard RPM that we built and tested, but is not supported by OSG.

GlideinWMS



Quick Facts: Releases & Support Structure

- Releases
 - Issues tracked in redmine issue tracker
 - <https://cdcvns.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.prd/download.html>
 - Release notes & history
 - <http://glideinwms.fnal.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
- Stakeholders Meeting every two months
- 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

2. Click custom query for stakeholder or version roadmap

glideinWMS

Overview Activity Roadmap **Issues** New Issue Gantt Calendar Documents Wiki Files Repository Code reviews Settings

OSG Tickets

Filters

Options

Apply Clear

#	Tracker	Target version	Status	Priority	Subject	Assignee	Updated	
v3_2_10 7								
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
v3_2_x 12								
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 analyze	Burt	02/03/2015 02:44 PM	OSG

Issues

View all issues

Summary

Calendar

Gantt

My custom queries

Issues Awaiting My Feedback

My Unassigned Issues

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