



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

GlideinWMS

Marco Mambelli

Stakeholders Meeting

September 12, 2018

Overview

- Releases since last stakeholder's meeting
- Upcoming releases
- Current focus
- GlideinWMS roadmap
- Developers spotlight
- Reference slides
 - GlideinWMS Architecture
 - Quick Facts

Releases Since Last Stakeholders Meeting

- v3_4 moved from OSG upcoming to OSG production
- v3_4_1 about to be released (RC released)
 - Improved Singularity support including the recommendations from the meetings
 - provisioning and execution according to NEVER/OPTIONAL/PREFERRED/REQUIRED selections
 - images dictionary and selection
 - mount-points support
 - support for GPUs
 - Increased the number and coverage of the unit tests
 - Updated code and documentation removing references to Corral and GlideinWMS v2

Releases Since Last Stakeholders Meeting (cont)

- v3_4_1 (cont)
 - User Collector using shared_port:
 - changed the GWMS provided HTCondor configuration
 - supporting ranges in sinful string values for collectors and CCBs
 - Fixed and improved entry_sets support
 - Improved monitoring
 - The entry selection algorithm is a plugin
 - Last version supporting Globus GRAM and last version with multi-user Factory

GlideinWMS: Current Focus

- Move to single user Factory
- Containerization
 - Improve Singularity support and make it very easy for VOs to switch to using it via GWMS
- Improve stability
 - More automated testing & CI (pylint, pythoscope, futurize, unittest ...) is an ongoing focus
 - Developer's test infrastructure to connect to Factory ITB services for scale testing
 - Test of new features on different sites in OSG
 - External contributions should be production ready
- Minimize wastage of resources from over-provisioning and improve auto-discovery
 - Improve handling of multi-node jobs
 - Auto - estimate of expected resources when provisioning
 - Actively follow the requests and adapt as the request goes down
 - First phases addressed with Glideins tracking requests up and down
 - Consider "transactional provisioning"

Next Planned Release

- v3_5 planned for mid of October, for OSG upcoming
 - Dropping Globus GRAM support
 - Single-user Factory: all Glideins will run using the factory user (no more separate users per-VO)
 - Changes in the Factory
 - Documentation and tools to ease migration
 - Track jobs that spawn multiple nodes, e.g. HPC submission
 - Adjust Singularity support with feedback from early adopters
 - Marco Mascheroni (ITB Factory) and Edgar (GlueX Frontend) will start tests soon
 - Monitoring for frontend: store the number of Job restarts
 - Add a configurable limit to the rate of jobs running and fail the glidein if the rate is passed (waiting on HTCondor ticket #6698)
 - Improvements to Factory and Frontend tools, especially the ones easing Factory operation

GlideinWMS Roadmap

- Medium term (2018 – mid 2019)
 - Keep up with the scalability requirements
 - Investigate and incorporate new technologies like pandas dataframes, numpy, etc
 - Optimization of the interactions w/ HTCondor
 - Containerization
 - Singularity and other containers: integration with HTCondor provided solutions [#20811]
 - Outsource GlideinWMS functionalities to HTCondor
 - Work with the HTCondor team to provide some of the Frontend functionalities natively through HTCondor
 - Leaner & modular Frontend
 - Adapt to changes/introduction of Acquisition Engine by HTCondor
 - Dependent on the work that will be done in HTCondor in the future
 - Very thin GlideinWMS Factory
 - Support for new HPC sites with stricter policies (e.g. no outbound connection except gateways, MFA)
 - Depends on support from HTCondor.
 - Monitoring Modernization
 - Retire GlideinWMS monitoring pages
 - Move to grafana/graphite/elastic search based solution

GlideinWMS Roadmap

- Long term (> mid-2019)
 - Move to Python 3
 - Start moving the code after v3.5 or following release
 - Have Python 3 version (v3.7) parallel to Python 2 version by end of Summer 2019
 - Move to Decision Engine (DE)
 - Replace the Frontend with the Decision Engine
 - Make Glidein as a service capable of talking to multiple WMS middleware/frameworks

Developers Spotlight

Marco Mambelli – Recent focus

- Contacts w/ GlideinWMS users (CMS, OSG, FIFE)
- GlideinWMS 3.4.1 contributions
 - Improved Singularity support including the recommendations from the meetings
 - Management of tickets and cutting the release
- Plans for GlideinWMS 3.5
 - Follow-up on Singularity tests and adoption, changes from feedback
 - Track jobs that spawn multiple nodes
 - Focus on Frontend tickets
 - Review tickets

Lorena Lobato - My focus on the project

+ Documentation

- Review, remove obsolete references and update information from the GlideinWMS documentation + Remove Corral documentation
- GlideinWMS ticket review from 2010 to do a first valuation and clarification about them and an importance

+ Review & Testing

- Increase unit test coverage for GlideinWMS
- Testing robustness of configurable Glidein Variables which are int
- Add the possibility to skip idle removal per entry
- Add a variable for VOs to add extra singularity bind mount points
- Use *127.0.0.1* for CONDOR_VIEW_HOST in our config instead of "localhost"
- Review if the blacklisting script works for GlideinWMS frontend

+ Development

- Switch child collectors to shared_port
- Remove dependency on condor_root_switchboard
- Problems with the default 'frontend user in the Factory'

Lorena Lobato - My focus on the project

+ Development (II)

- Add a configurable limit to the rate of jobs running and fail the glidein if the rate is passed
- Error message when upgrading FrontEnd
- Should tarball installations be supported?

+ Summary

- Finished **v3.4.1 tickets** and starting with v3.5 ones
- 6 months review completed
- Starting to have more autonomy
- Suggested to do more documentation, specially when feedback
- Participating in other meetings (FIFE, HEPCloud)
- Trying to get more experience with EDX courses

Dennis Box

- Main Focus
 - unit testing/code stability
 - preparation for Python 3 migration
- Unit Testing
 - Upcoming release (3.4.1) has 36% python unit test coverage
 - Coverage run nightly on our CI system as part of unit tests
- Python3 Migration
 - Futurize has been run against code base in ‘nonaggressive’ mode by CI for about a year
 - Recently started running unit tests with python3
 - They failed, 0% coverage
 - Futurize missed some ‘low hanging fruit’ file()->open() etc

Dennis Box - 2

- Python3 continued
 - Aggressive futurize auto-upgrade and other minor changes brought pass rate to ~33% and coverage to 10%
 - These changes made Python2 tests fail
- Plans
 - Increase coverage of unit tests
 - Continue to investigate automation tools
 - Autopep8, testfixtures show promise
 - Migrate to a code base that supports both Python2 and Python3
 - futurize, six, other compatibility libraries should help

Marco Mascheroni – Items I have been working on that will go into 3.4.1

- Taking care of factory ops/CMS items
- Commissioning of metasites
 - Monitoring fixes (e.g.: popups not showing under certain circumstances)
 - Not reload temporary files upon upgrade/reconfig
 - Improve how subentries are picked for metasites
- Manual_submit_glidein tweaked after use in production contexts
 - Very handy script
- Disable idle glidein removal at entry level
- Documentation improvements (e.g.: work dir, conflicting parameters)

Marco Mascheroni - And for 3.5 will work on...

- Feedback items” as they come
 - Expect more metasites/tools tweaks
- Had a conversation with Brian B. and Jeff about automatic configuration generation from CRIC/OSG collector
 - Agreed we need a tool that will do a entry diff
 - Already started working on it
 - Intermediate human friendly files to “patch” what is generated automatically
 - Merge tool that generate the possible xml stub
- Scaling limitations in the frontend emerged during CMS scale tests [\[20302\]](#)


Questions/Comments

Reference Slides

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 grlideinwms-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.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
- 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

Default tabs not too useful

2. Click custom query for stakeholder or version roadmap

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

#	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