



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

GlideinWMS

Marco Mambelli

Stakeholders Meeting

November 14, 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_2 released on October 26 (3.4.1 on 10/18)
 - Improved Singularity support including the recommendations from the meetings
 - provisioning and execution according to NEVER/OPTIONAL/PREFERRED/REQUIRED selections
 - Singularity 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_2 (cont)
 - User Collector using shared_port:
 - Changed the GWMS provided HTCondor configuration
 - Supporting ranges in sinful string values for collectors and CCBs
 - Easy transition supporting both shared and distinct ports
 - 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
 - (Only 4.2) Added a test of the Frontend configuration allowing some configurations (UC shared port and new Singularity features) only if the connected Factories support them

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 in 2 weeks (end of November), 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
 - Edgar (GlueX and GLOW VO) started testing
 - Added scaling factor for Glidein limits in Factory Entries
 - Added ability to keep unused Glideins around (for site troubleshooting)
 - Some bug fixes and code hardening (metasite reconfiguration, space handling in shell scripts, review of shell scripts)
 - Improvements to Factory and Frontend tools, especially the ones easing Factory operation

GlideinWMS Roadmap

- Medium term (2018 – mid 2019)
 - Ticket review (December) and Code review (January)
 - 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]
 - Advertise array of resources (e.g. Multiple platforms)
 - 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
 - Focus on Frontend tickets
- GlideinWMS 3.5 contributions
 - Follow-up on Singularity tests and adoption
 - Track jobs that spawn multiple nodes
- Plans for GlideinWMS 3.5.1
 - Singularity support improvement (easy testing scripts), other changes from feedback

Lorena Lobato - My focus on the project

+ GlideinWMS 3.4.1 contributions

- 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
- Error message when upgrading the frontend
- Switch child collectors to shared_port
 - Sinful strings in the collector cannot have spaces
 - Collector and CCB strings cut a coma
 - Configure User Pool secondary collectors to answer also on different ports

+ GlideinWMS 3.5 contributions (ongoing)

- Review if the blacklisting script works for GlideinWMS frontend
- Removal of support Globus GRAM gt2/gt5 as gridType
- Removal of dependency on condor_root_switchboard
 - HTCondor team provided a script to ease the transition
- Problems with the default ‘frontend user in the Factory’

Lorena Lobato - My focus on the project

+ Plans for GlideinWMS 3.5.1

- Add a configurable limit to the rate of jobs running and fail the glidein if the rate is passed
 - HTCondor ticket #6698
- Gather requirements to have security alerts GWMS dependencies in the GitHub repository
- Should tarball installations be supported?

+ Review & Testing (different GWMS versions)

- 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
- Improve the way condor_jdl dict is populated for metasites
- Testing GlideinWMS 3.4.1 RC + 3.4.2

Dennis Box

- Focus continues on code quality and testing
 - Containerized gwms testing to simplify setup, usage
 - Publicly available on docker hub, github

```
mac-128665:~ dbox$ docker run dbox/gwms-test help
Linux f322d28c8d0c 4.9.93-linuxkit-aufs #1 SMP Wed Jun 6 16:55:56 UTC 2018 x86_64 x86_64 x86_64 GNU/Linux
Scientific Linux release 7.5 (Nitrogen)
glideinWMS test container
some commands:
    help: this message
    new_gwms_branches N: show branches that have been changed in the last N days
    gwms_quick_tests <branch>: execute the fast-running unit tests on <branch>
    run_gwms_coverage <branch>: run a full coverage test on <branch>
    run_gwms_pylint <branch>: run pylint and pycodestyle on <branch>
```

- Checkin to github triggers test, build on travis-ci.org
- Also tested on in-house CI (buildmaster.fnal.gov)
 - Works on build nodes that have docker installed

Marco Mascheroni – Items I have been working on that will go into 3.5

- Fixes and improvements
 - Found and fixed a bug in metasite reconfiguration
 - Several improvements on the manual_submit_glidein
- Factory ops feedback
 - Automatically remove running glideins after walltime
 - Add a scaling factor for all glideins limits in the entries (i.e.: num_factories in the config)
- Configuration generation from CRIC
 - Finalized gfdiff script: will go in 3.5 (print differences between two entries)
 - Generation script: will do a first round of test before including in the release
- Transitioning script for condor_switchboard removal
 - Script to change old glideins usernames to single “gfactory” user
- Testing, documentation, review

Marco Mascheroni – And for 3.5.1 will work on...

- Configuration generation from CRIC
 - First prototype of the generation script ready
 - Collaborating with the CRIC team to make sure needed information is taken
 - More info at my [talk](#) at the CMS offline and computing week
- Factory ops feedback
 - Remove really old files from reconfig
 - Clean up configuration files from disabled entries
 - Remove really old files from reconfig
- Scaling limitations in the frontend emerged during CMS scale tests [\[20302\]](#)

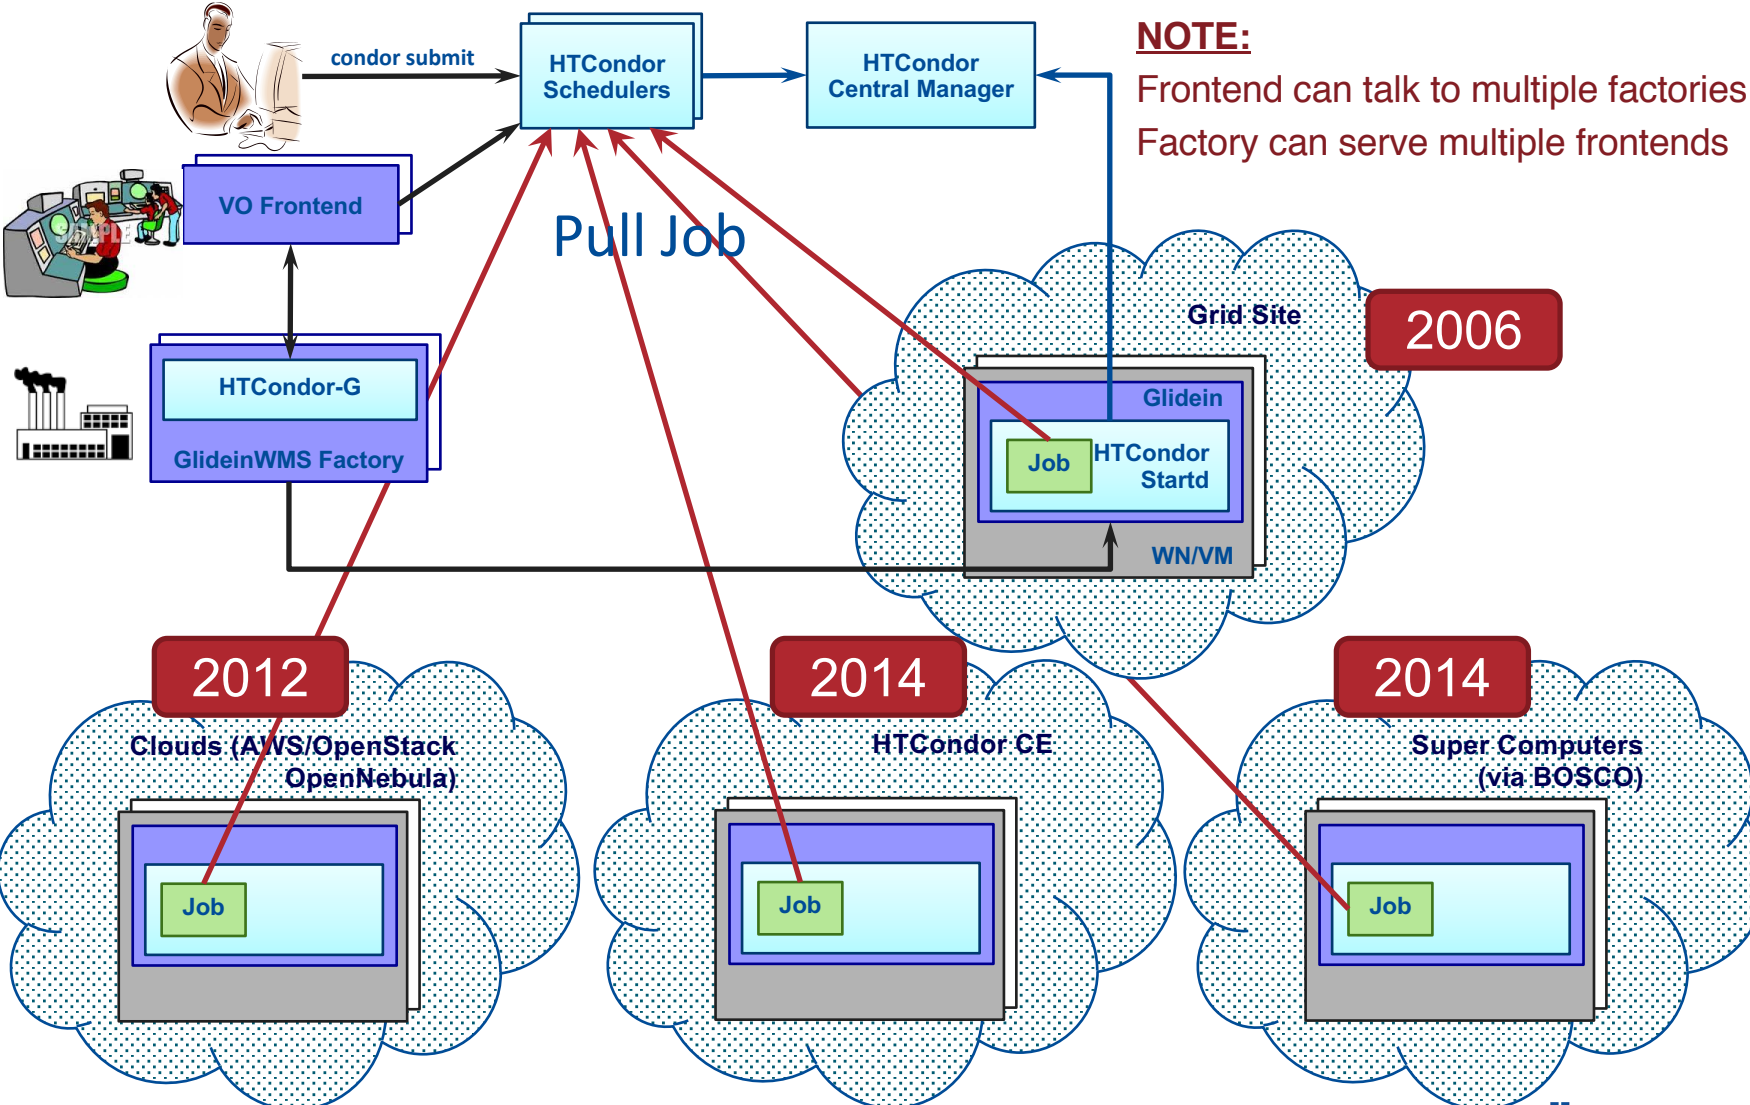
Questions/Comments

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.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

Default tabs not too useful

2. Click custom query for stakeholder or version roadmap

The screenshot shows the Redmine interface for the glideinWMS project. The 'Issues' tab is selected. The main content area displays a table of issues. The right sidebar contains a 'My custom queries' section with a list of queries, including 'OSG Tickets', 'Stakeholder Requests', and 'v3.2.10: Roadmap'. Annotations include a red box around the 'Active' and 'Roadmap' tabs, a green callout box pointing to the 'Issues' tab, and a green oval around the 'OSG Tickets' query in the sidebar.

#	Tracker	Target version	Status	Priority	Subject	Assignee	Updated	Category
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 Mhashikar	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 Mhashikar	04/13/2015 03:21 AM	OSG
8437	Feature	v3_2_10	New	Normal	Split entry configuration from main factory config	Parag Mhashikar	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 Holzman	02/03/2015 02:44 PM	OSG