



---

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

---

## **GlideinWMS**

Marco Mambelli

Stakeholders Meeting

January 9, 2019

# Overview

---

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

## Next Planned Releases

---

- No release since the last stakeholders meeting
- We have 2 releases close to completion
  - v3.4.3 w/ bug fixes and minor features, for OSG production, expected in the next couple of weeks
  - v3.5 w/ single-user Factory and some other features, for OSG upcoming, planned for mid February

## Next Planned Release, v3.4.3

---

- v3\_4\_3 planned in two weeks, for OSG production
  - Hardening of shell scripts (linting, review)
  - Adjusted some glitches in 3.4.1/2 (upgrade controls work also if there is no Factory, improved some help messages)
  - Some changes to Singularity thanks to the feedback from NOVA (improved site troubleshooting)
  - Fixes to a couple of bugs highlighted by the interactions w/ HEPCloud
    - Frontend not recognizing entries in downtime
    - Stale running and held Glidein numbers reported in Factory classads
    - Print a warning when the Factory configuration contains conflicting attributes
  - Factory scripts improvements (more robust and better messages)

## Next Planned Release, v3.5

---

- v3\_5 planned for mid February, 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
  - Monitoring for Frontend: store the number of Job restarts
  - Improvements to Factory and Frontend tools, especially the ones easing Factory operations
  - Added a configurable limit to the rate of jobs running and fail the glidein if the rate is passed (waiting on HTCondor ticket #6698)

# GlideinWMS Roadmap

---

- Medium term (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 of the documentation to Jekyll
    - Use of templates will ease page maintenance
  - Stronger adoption of Github
    - Redmine, especially the tickets, currently works well
  - Move to Decision Engine (DE)
    - Support Frontend and 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 3.4.3 contributions
  - Singularity follow-ups
  - Add the possibility to disable completely Glidein removal
  - Stale running and held glidein numbers reported in factory classads
  - Focus on Frontend tickets
  - Management of tickets and cutting the release
- GlideinWMS 3.5 contributions
  - Follow-up on Singularity tests and adoption
  - Track jobs that spawn multiple nodes
- After
  - Monitoring improvements
  - Singularity support improvement (easy testing scripts), other changes from feedback

# Lorena Lobato - My focus on the project

---

## + Review & Testing (different GWMS versions)

- Release code gives the wrong help message
- Frontend upgrade is failing if it is unable to determine the version of the Factory
- Unit Tests review
- The factory seems to ignore the configuration values in the files in the config.d directory w/ entry configurations
- Remove really old files from reconfig
- Automatically remove glideins after walltime
- Testing robustness of configurable Glidein Variables which are int
- Improve the way condor\_jdl dict is populated for metasites
- Testing GlideinWMS 3.4.2 + 3.4.3
- Opened a long-term tickets to list all the possible issues

# Lorena Lobato - My focus on the project

---

## + GlideinWMS 3.4.3 contributions

- Potential bug in 3.4.2 frontend--not recognizing entries in downtime.
- Problems with the default 'frontend' user in the Factory
- Removal of support Globus GRAM GT2/GT5 as gridType
- Removal of dependency on condor\_root\_switchboard
- Create GlideinWMS RPMs

## + What I am working right now

- Review if the blacklisting script works for GlideinWMS frontend
- Error message related to entry in the Factory logs
- Should tarball installation be supported?
- Gather requirements to have security alerts GWMS dependencies in the GitHub repository

# Marco Mascheroni

---

- Items included in 3.4.3
  - Fixes and improvements
    - Metasites reconfiguration failures
    - Fixed another case of EntryGroup process leaks
    - “Entry level” attributes ignored when global one are present and const attribute is discordant
  - Factory ops feedback
    - Remove old files from reconfig
    - Automatically remove glideins after the walltime is hit
    - Manual\_submit\_glideins improvements: usability and automation
  - Testing, documentation, tickets reviews, improved error messages
- Working on...
  - Configuration generation from CRIC
    - In the process of validating generation script (using the gfdiff one)
  - Other smaller items as required

# Dennis Box

---

- Code quality and testing remains focus

## Containerized CI example

Source: <https://github.com/ddbox/gwms-test>

CI build: <https://travis-ci.org/ddbox/gwms-test>

Hub: <https://cloud.docker.com/u/dbox/repository/docker/dbox/gwms-test>

Example usage in our CI system

[https://buildmaster.fnal.gov/job/gwms-run-test/ws/146/146\\_results.html](https://buildmaster.fnal.gov/job/gwms-run-test/ws/146/146_results.html)

22 minute run time, relatively easy to find logs and coverage reports

Above CI report also runs on Travis-ci

Size looks right, haven't been able to offload artifacts back to github

This is supposed to be possible

Compare to our 'Legacy' CI

[https://buildmaster.fnal.gov/job/glideinwms\\_ci/711/](https://buildmaster.fnal.gov/job/glideinwms_ci/711/)

3 hr 35 m run time, coverage report only available for last build

# Thomas Hein - GlideinWMS Monitoring System

---

- GlideinWMS provides monitoring on both a Factory and Frontend level using RRD Databases and XML Files
- Monitoring for RRD is being updated directly in the code in various files with no easy way to add additional monitoring systems
- The goal of this project is to replace anything RRD/XML specific with a monitoring class where new monitoring “modules” can simply tap into the class
- RRD and XML will be rewritten into “modules” and still collect the very same data it did before
- InfluxDB will be added as an additional module as an example
- Currently, the frontend is complete with this change and the factory is nearly complete
- After the factory, documentation will be written on usage

---

# Questions/Comments

---

# Reference Slides

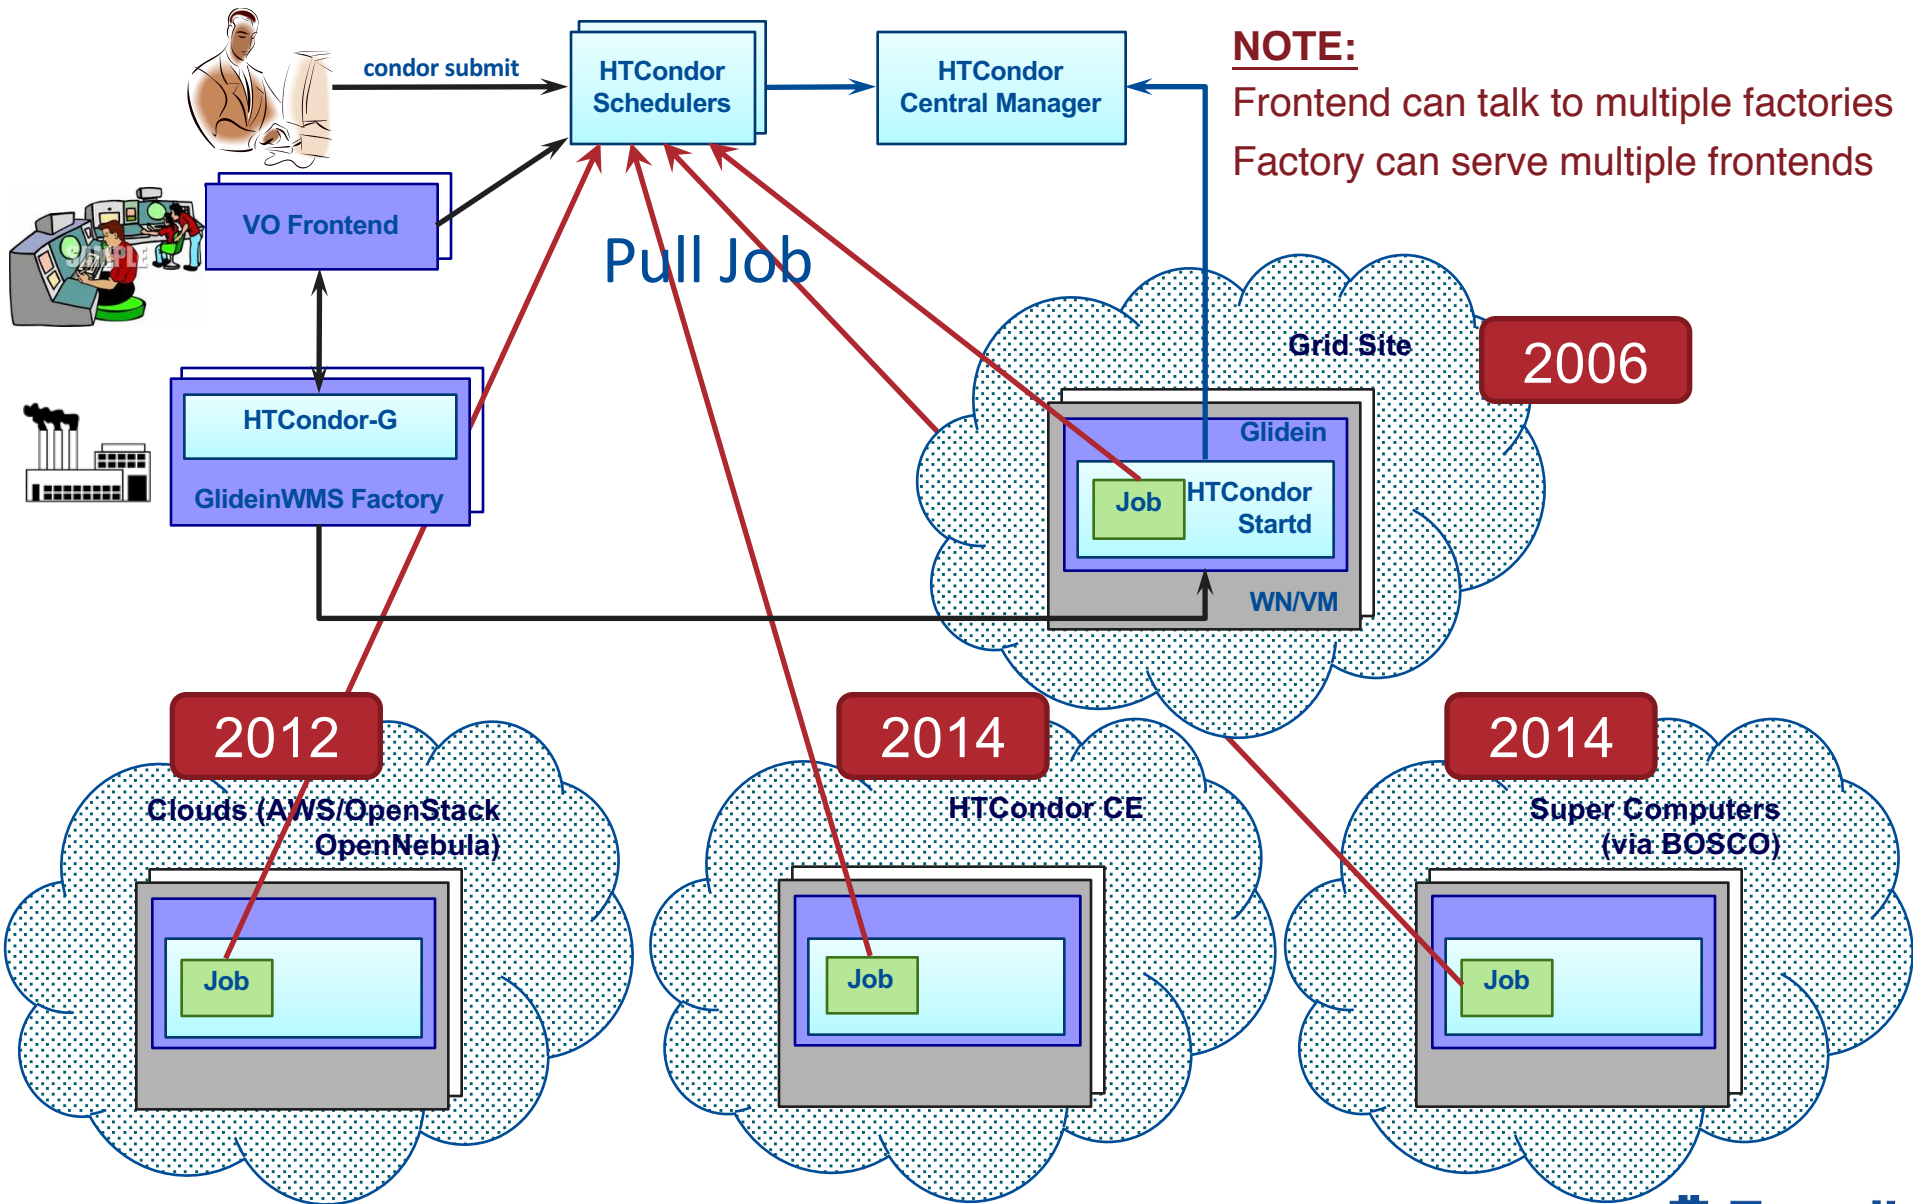


# Move to single user Factory

---

- 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  $\geq$  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://cdcvs.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	<a href="mailto:glideinwms-support@fnal.gov">glideinwms-support@fnal.gov</a>
Stakeholders	<a href="mailto:glideinwms-stakeholders@fnal.gov">glideinwms-stakeholders@fnal.gov</a>
Release Announcements	<a href="mailto:glideinwms-support@fnal.gov">glideinwms-support@fnal.gov</a> <a href="mailto:cms-dct-wms@fnal.gov">cms-dct-wms@fnal.gov</a> <a href="mailto:glideinwms-stakeholders@fnal.gov">glideinwms-stakeholders@fnal.gov</a>
Future Release plans	See next slide
Discussions	<a href="mailto:glideinwms-discuss@fnal.gov">glideinwms-discuss@fnal.gov</a>
Code commits	<a href="mailto:glideinwms-commit@fnal.gov">glideinwms-commit@fnal.gov</a> 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

My custom queries

- Issues Awaiting My Feedback
- Unclosed 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	OSG
<b>v3_2_10</b>								
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
<b>v3_2_x</b>								
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 Holzman	02/03/2015 02:44 PM	OSG