



---

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

---

## **GlideinWMS**

Marco Mambelli, Dennis Box, Bruno Coimbra, Marco Mascheroni,  
Namratha Urs, Maia Boyd

Stakeholders Meeting

July 14, 2021

# Agenda

---

- Status update
- Important changes
- Project Progress and Roadmap
- Developers' spotlight
- Stakeholders Reports or Presentations
- Roundtable
- Open discussion

## Project status updates

---

- v3.6.5 is in OSG production (OSG 3.4 and 3.5)
- v3.7.4 in OSG upcoming (OSG 3.5) - production
- v3.9.2 (Python 3) in OSG contrib
  
- Project effort unchanged
  - 2.15 FTE
  - Management and 4 developers
  - Everyone involved also in other projects: Factory Ops, HEPCloud, Jobsub, FIFE support, Landscape
  - Namratha Urs and Maia Boyd joined the team for the Summer

# Important changes

---



- 3.7.5
  - Make sure that the condor configuration changes are applied (look for possible .rpmnew files and merge them)
- 3.9.2
  - Drop of BDII support
  - The Python code in the configuration must be valid Python 3
    - The match expression is validated
    - We are unable to control the type of the variables, check the type of factory/job match attributes!
- And please review periodically your tickets/priorities  
<https://cdcvs.fnal.gov/redmine/projects/glideinwms/wiki/RoadmapSummary>

# GlideinWMS Releases



- 2 active release series in OSG 3.5
  - (Production, also in OSG 3.4, closed)
  - Development – more disruptive changes (tokens, logging) [OSG]
  - Python 3, soon also in OSG 3.6 [HEPCloud]
- Fixes in the lowest possible series
- Features propagated only to greater series

## Release series

---



- 3.7/3.8 main release in OSG 3.5
  - OSG 3.5 upcoming repo
  - Python 2
  - Support GSI and JWT (with proper HTCondor)
  - How do you plan to use it?
    - GSI, Tokens (SciTokens/WLCG), both?
    - Which HTCondor version?
- 3.9 main release in OSG 3.6
  - OSG 3.5 contrib repo and OSG 3.6 (production)
  - Python 3
  - No GSI in OSG 3.6
  - Status of Python 3 adoption?

## 3.7 Series Status - Dennis Box

---

- GlideinWMS 3.7.5 plans (end of July)
  - IDToken improvements
    - Incorrect condor version check removed
    - Naming convention: EntryName vs Glidein\_Site
    - TRUST\_DOMAIN queried from condor configuration instead of hard coded
    - Reduce complexity of per-site token revocation
  - SciToken improvements
    - Fixed problem where SciTokens overwriting each other in factory storage
    - Provide framework for admin provided scripts to generate SciTokens (as opposed to assuming they live at place specified in configuration)
  - Documentation improvements
    - Document external schedd to frontend IDToken authentication

## 3.7 Series Status - Dennis Box

---

- GlideinWMS 3.7.5 (continued)
  - Setup CVMFS via cvmfs exec when available
  - Jwt monitoring artifacts only created when asked for in config
  - Custom execution points for VO scripts
  - Fix GLIDEIN\_SINGULARITY\_BINARY\_OVERRIDE



# Complete token implementation

---

- Planned for 3.7.6 / 3.8
- Support token-only, GSI only and token+GSI configurations
  - Clear configuration
  - Clear (main) documentation
  - No errors
  - Explicit configuration (choose GSI vs Tokens)
  - **Fallback?**
- Satisfy OSG and other VOs requirements
  - **What is still missing?**
  - **We need VO feedback**

## 3.9 Series (Python 3) Status and plans – Bruno C

---

- GlideinWMS 3.9.2 released
  - The release has been validated and is now available in the OSG 3.5 contrib repository.
  - Reminder: this release drops BDII support from `manageFactoryDowntimes.py` and the new brings `check_python3_expr.py` script.
- GlideinWMS 3.9.3
  - There still are a few encoding issues in 3.9.2 that will be fixed in the next release.
  - Other changes will include code style and performance improvements for Python 3.
  - We're still accepting requests for this release.

# Bruno Coimbra

---

- Custom execution points for VO scripts
  - The project to implement multiple execution points for the VO scripts on the glidein is still ongoing. I expect to be able to focus most of my time to it now.
- Direct Submission Application
  - I'm working with our new intern Maia to extend LeRayah's project that allows custom scripts to be submitted directly to an entry from a factory. This will help Maia on testing her changes for her own project.
- Current Focus
  - Improvements to the glidein\_startup script.
  - Address the open issues for the 3.9 branch.
  - Preparing the 3.9.3 release.

## Marco Mascheroni

---

- Pilot logs monitoring
  - Removed the need to run the indexer by adding a dedicated html page that skip internal database lookups
  - Added possibility to get a “public” link to a specific pilot log in order to share them with people who do not have access to the toolt.
- Completed attribute expansion feature
- Fixed bug related to identification of CENTOS 8 systems when `CONDOR_OS="auto"` is used

# Namratha Urs

---

Graduate Student in Computer Science at University of North Texas

Developer, GlideinWMS (Returning FCSI Intern)

Feature addition to provision CernVM-FS on sites that do not provide it locally

- Ship `cvmfsexec` with the `glidein`
  - Downloadable on-demand
  - Determine the most reliable option to provide CVMFS given the worker node environment
- Start singularity in unprivileged mode to allow user jobs to run (singularity + CVMFS)

Status:

- Completed end-to-end testing of the CVMFS mount/unmount workflow
- Fixed the bug found in cleanup script invocation as part of un-mounting CVMFS — in v3.7.5
- Work in Progress items:
  - Download CVMFS utilities only for entries that need it (via an `<attr>` switch) — looking at v3.7.5
- Future plan:
  - Add mode 3 of `cvmfsexec` (using `cvmfsexec` directly)
    - Requires major changes to the `glidein` (additional effort to ensure all the required variables are passed to the subprocess)
  - Conditional download of CVMFS utilities (per each OS and CVMFS config repository) — planned for v3.7.6

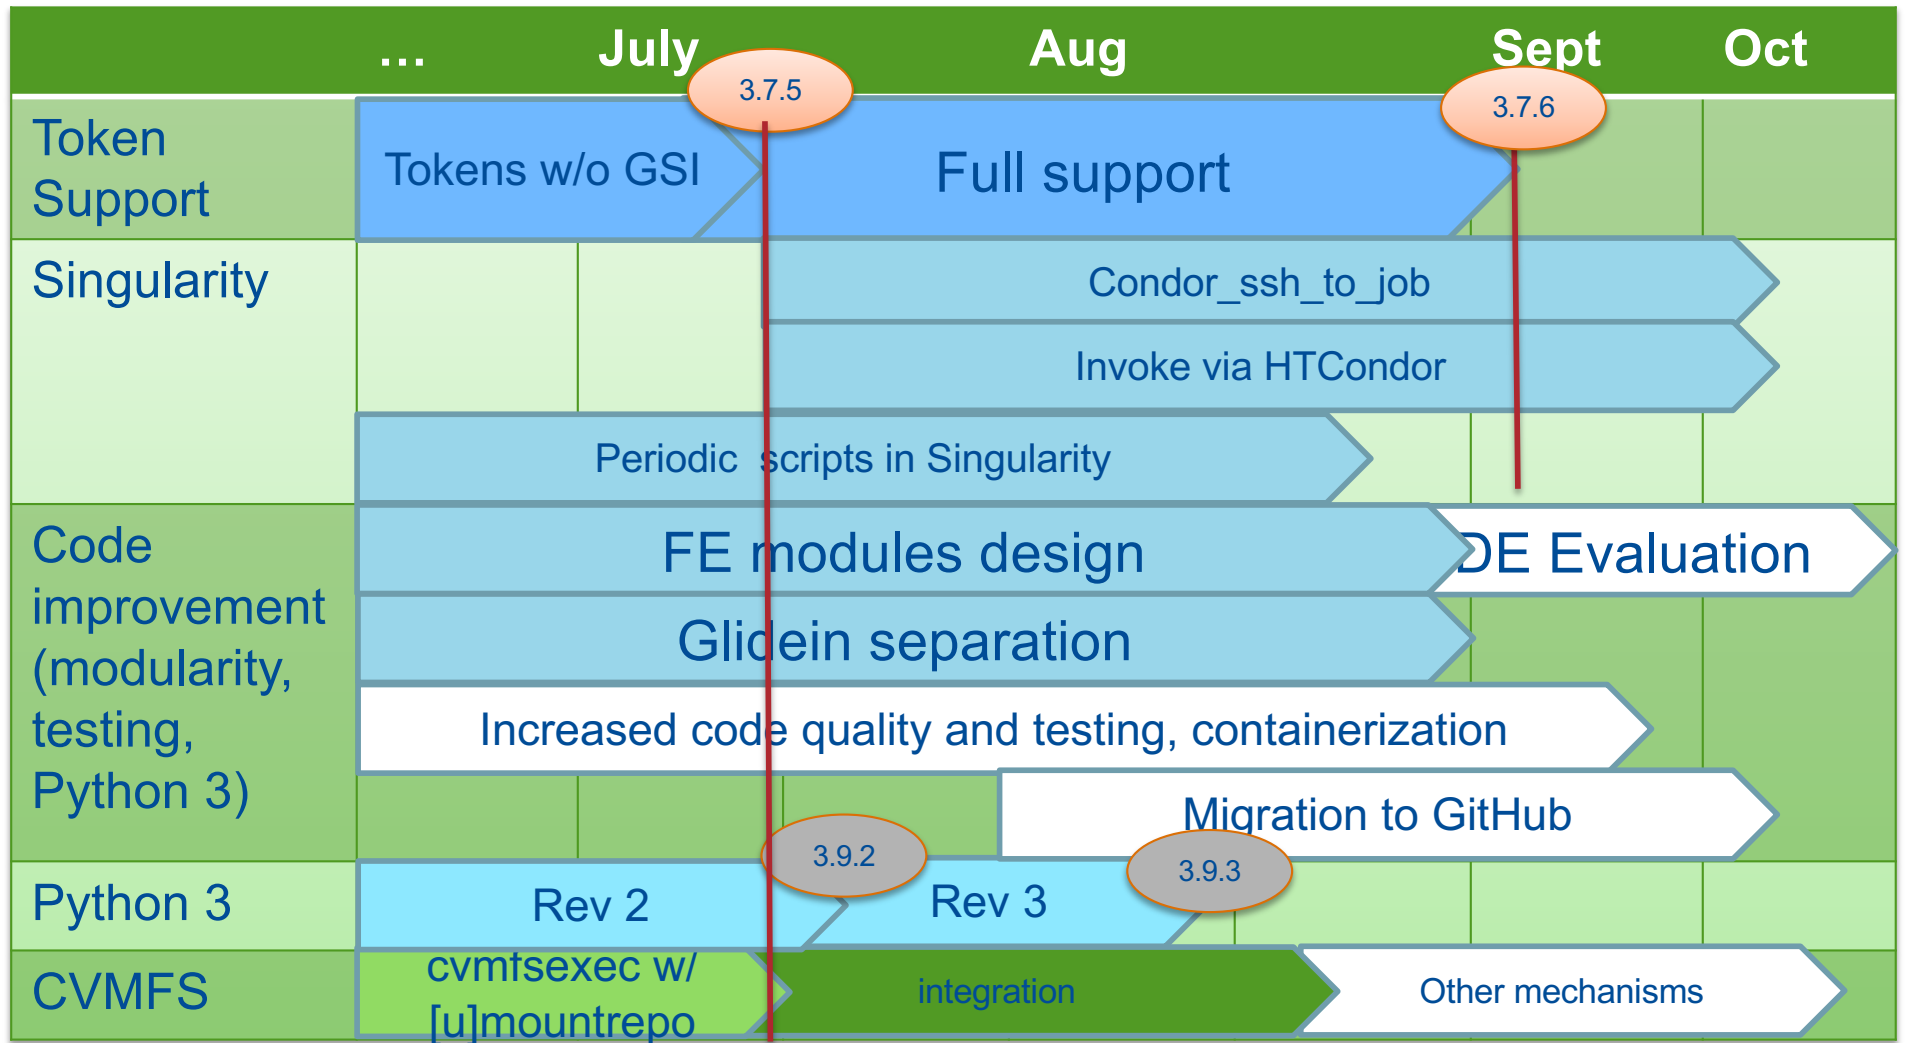
# Maia Boyd - GlideinWMS Intern

---

Incoming third-year at the University of Chicago studying Computer Science  
GlideinWMS Intern

- Improving Direct Submission Application
  - I have been working to extend a testing application that allows custom scripts to be submitted directly to an entry from a factory. I have modified the application to make it more portable and I have removed unnecessary files and variables.
  - This application will allow me to test different versions of Glidein.
- Evaluate Alternatives to Bash
  - I will explore and evaluate languages such as Python, Go and Nim as alternatives to bash for the glidein. I will be evaluating languages based on their portability and their ability to simplify the Glidein code.
- Rewriting glidein\_startup.sh
  - After selecting one or two languages, I will rewrite the glidein\_startup.sh code in those languages. During this process, I will consult developers and literatures on writing robust and postable code.
- Testing New Code
  - After rewriting the Gildein code, I will test the new code in a distributed environment possibly on Grids, Clouds and Supercomputers.

# Development Roadmap overview



Dennis Box

Bruno Coimbra

Marco Mambelli

Marco Mascheroni

Interns

Not assigned



## Roadmap in detail – Plans for 3.8 / 3.7.6 (mid Sept)

---

<https://cdcvs.fnal.gov/redmine/projects/glideinwms/roadmap>

- Close 3.6 series (attribute expansion in configuration)
- Complete token authentication (OSG) – Dennis
  - Fully working token implementation
  - No errors if no GSI
- Singularity support (CMS/OSG/Fermilab) – Marco Mambelli
  - Work on sites without Singularity or CVMFS installed – Dave, Namratha
  - Simplified code by Singularity invocation via HTCondor
  - Simplified troubleshooting by allowing ssh to job in Singularity
  - Ongoing collaboration w/ OSG – Mats
- Ease operations
- Containerization (OSG/Fermilab/FactoryOps)
  - Factory and Frontend in containers
  - Ongoing collaboration w/ Fermilab/OSG/FactoryOps – Shreyas, Diego



# Roadmap: modularity and code quality (HEPCloud/OSG)

---

- Improvements of Python 3 version - Bruno
  - Better code structure and packaging
  - Remove additional obsolete code
- Glidein as 3<sup>rd</sup> component (DE, containerization) – Marco Mmb
  - Separate packaging
  - Easier and more robust support for custom scripts - Bruno
- Improve modularity of Frontend (DE) - Marco Mambelli
  - Configuration
  - Business logic <https://drive.google.com/file/d/1KTThCsx4ZGjxwqe-gfZtonT3eGJamPed/view?usp=sharing>
  - Communication w/ Factory
- Integration with HEPCloud/DecisionEngine (HEPCloud/CMS)
- Expand, simplify and automate testing

# Summary

---

- Development focus
  - Tokens
- Expect more development also for
  - Glideins
  - Singularity
  - Frontend modularity

---

# Extra Slides

---

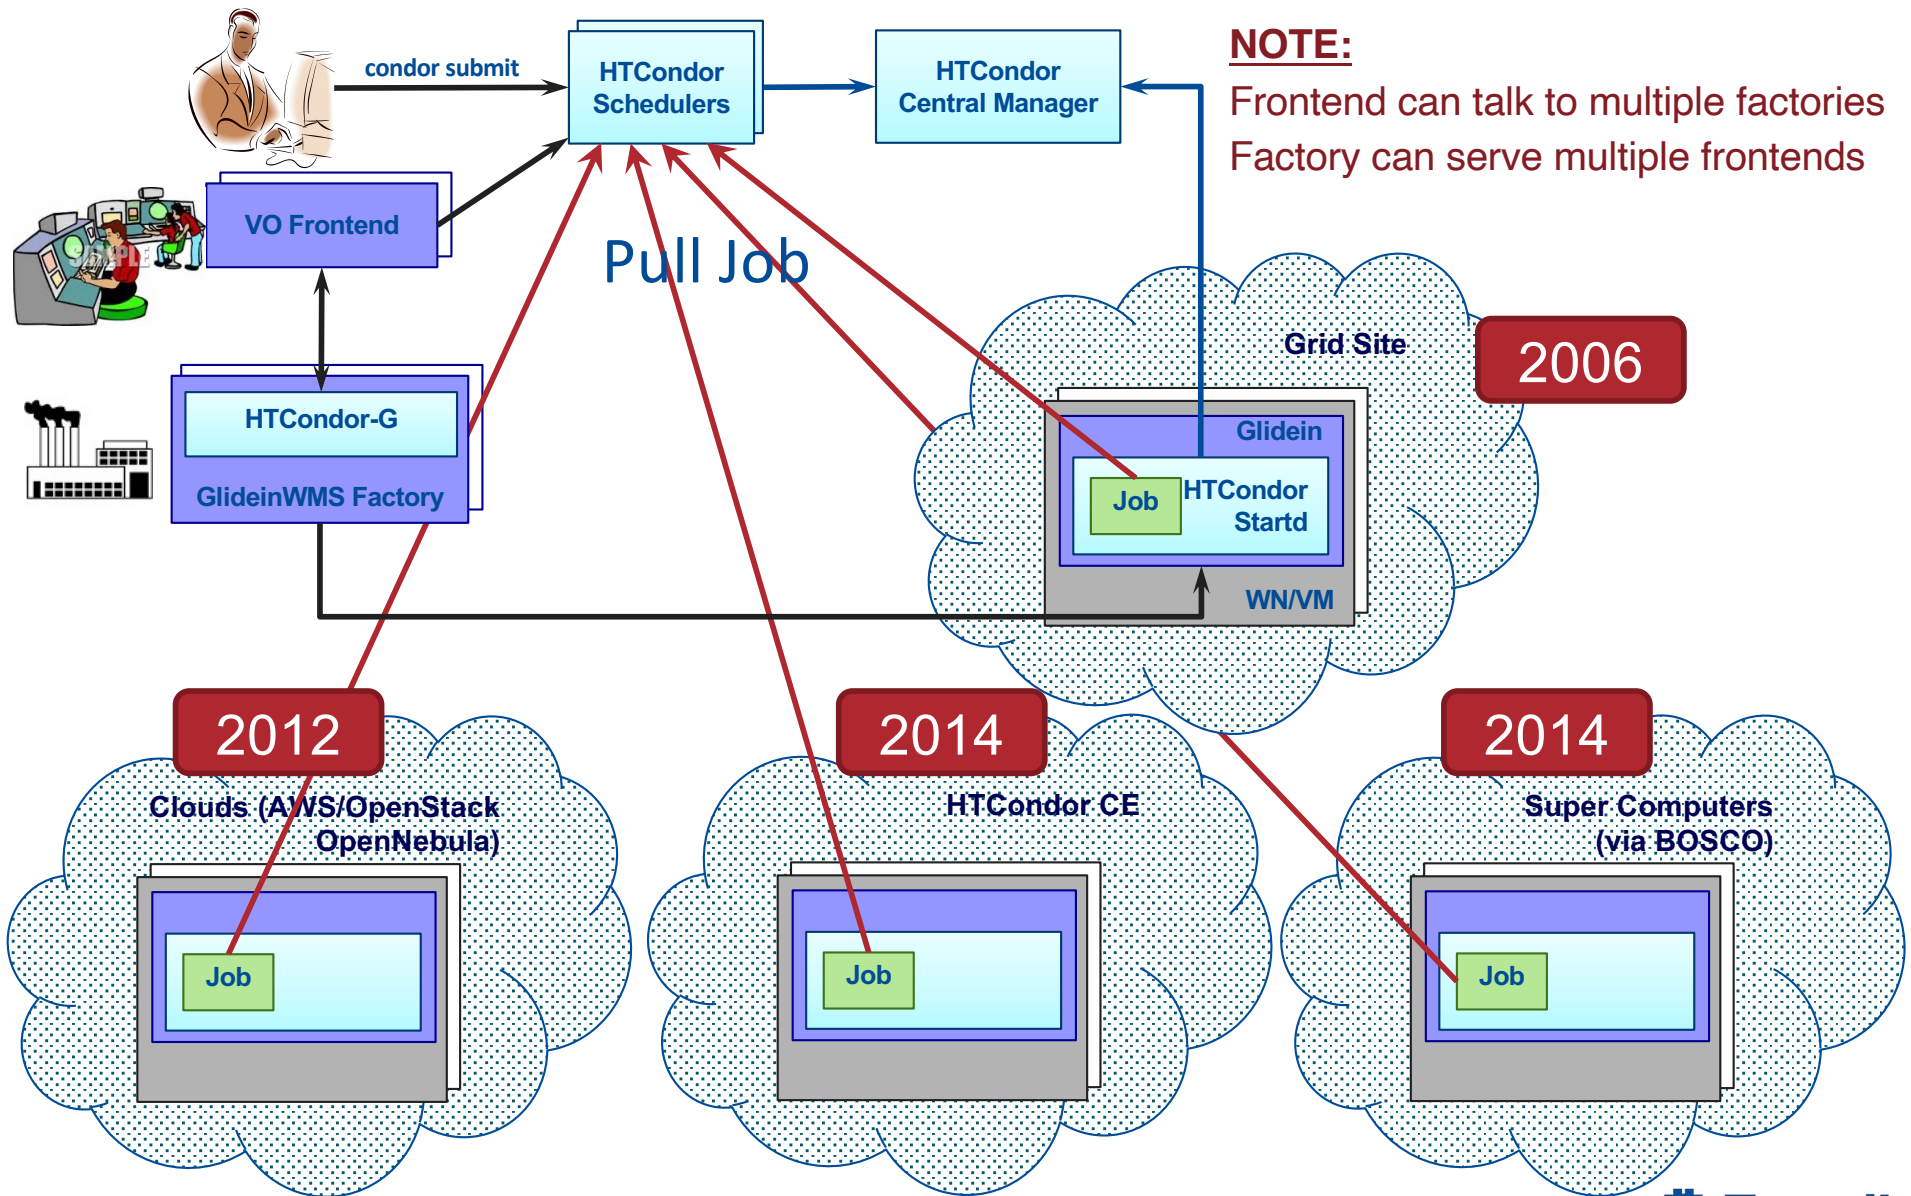
# Reference Slides

# Move to single user Factory

---

- Starting from 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.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	<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

2. Click custom query for stakeholder or version roadmap

Issues

- View all issues
- Summary
- Calendar
- Gantt

My custom queries

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