



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

GlideinWMS

Marco Mambelli

Stakeholders Meeting

November 13, 2019

Overview

- Project updates since last stakeholders meeting
- Completed and Upcoming releases
- GlideinWMS roadmap
- Developers spotlight
- Reference slides
 - GlideinWMS Architecture
 - Quick Facts

Project Updates Since Last Stakeholders Meeting

- Announcements
 - Lorena Lobato Pardavila leaving the team
 - Bruno Coimbra joining
 - GlideinWMS v3.6 released September 25, in OSG production
 - GlideinWMS v3.6.1 RC released November 12
- Project Effort (2.50 FTE)
 - Project Management: 0.15 FTE
 - Development & Support: 2.35 FTE
- Temporary effort
 - 1 on call collaborator, limited effort, Thomas Hein

Project Updates Since Last Stakeholders Meeting

- Communication
 - Please review periodically your tickets/priorities
<https://cdcvs.fnal.gov/redmine/projects/glideinwms/wiki/RoadmapSummary>
 - How can we further improve communication
 - Should we participate in any other meetings?
 - Communicating priorities?

- Support
 - Incompatibility with the HTCondor configuration in OSG 3.5 (fixed in 3.6.1)

Action Items from Previous Stakeholders Meeting

Action Items	Status
Add a roadmap overview	done
Add a GPU cluster to the ITB Frontend/Factory in Fermicloud	started

Completed and Next Planned Releases

- Released
 - GlideinWMS v3.6 released September 25, in OSG production, renaming of v3.5.1
 - OSG and CMS production Factories are still v3.4.6

- We have 3 releases in the pipeline
 - v3.6.1 in the production series for OSG 3.4 and 3.5
 - v3.6.2 in the production series mid December
 - v3.7 in OSG upcoming, in 1 week, release candidate out

GlideinWMS Roadmap – dropping support for...

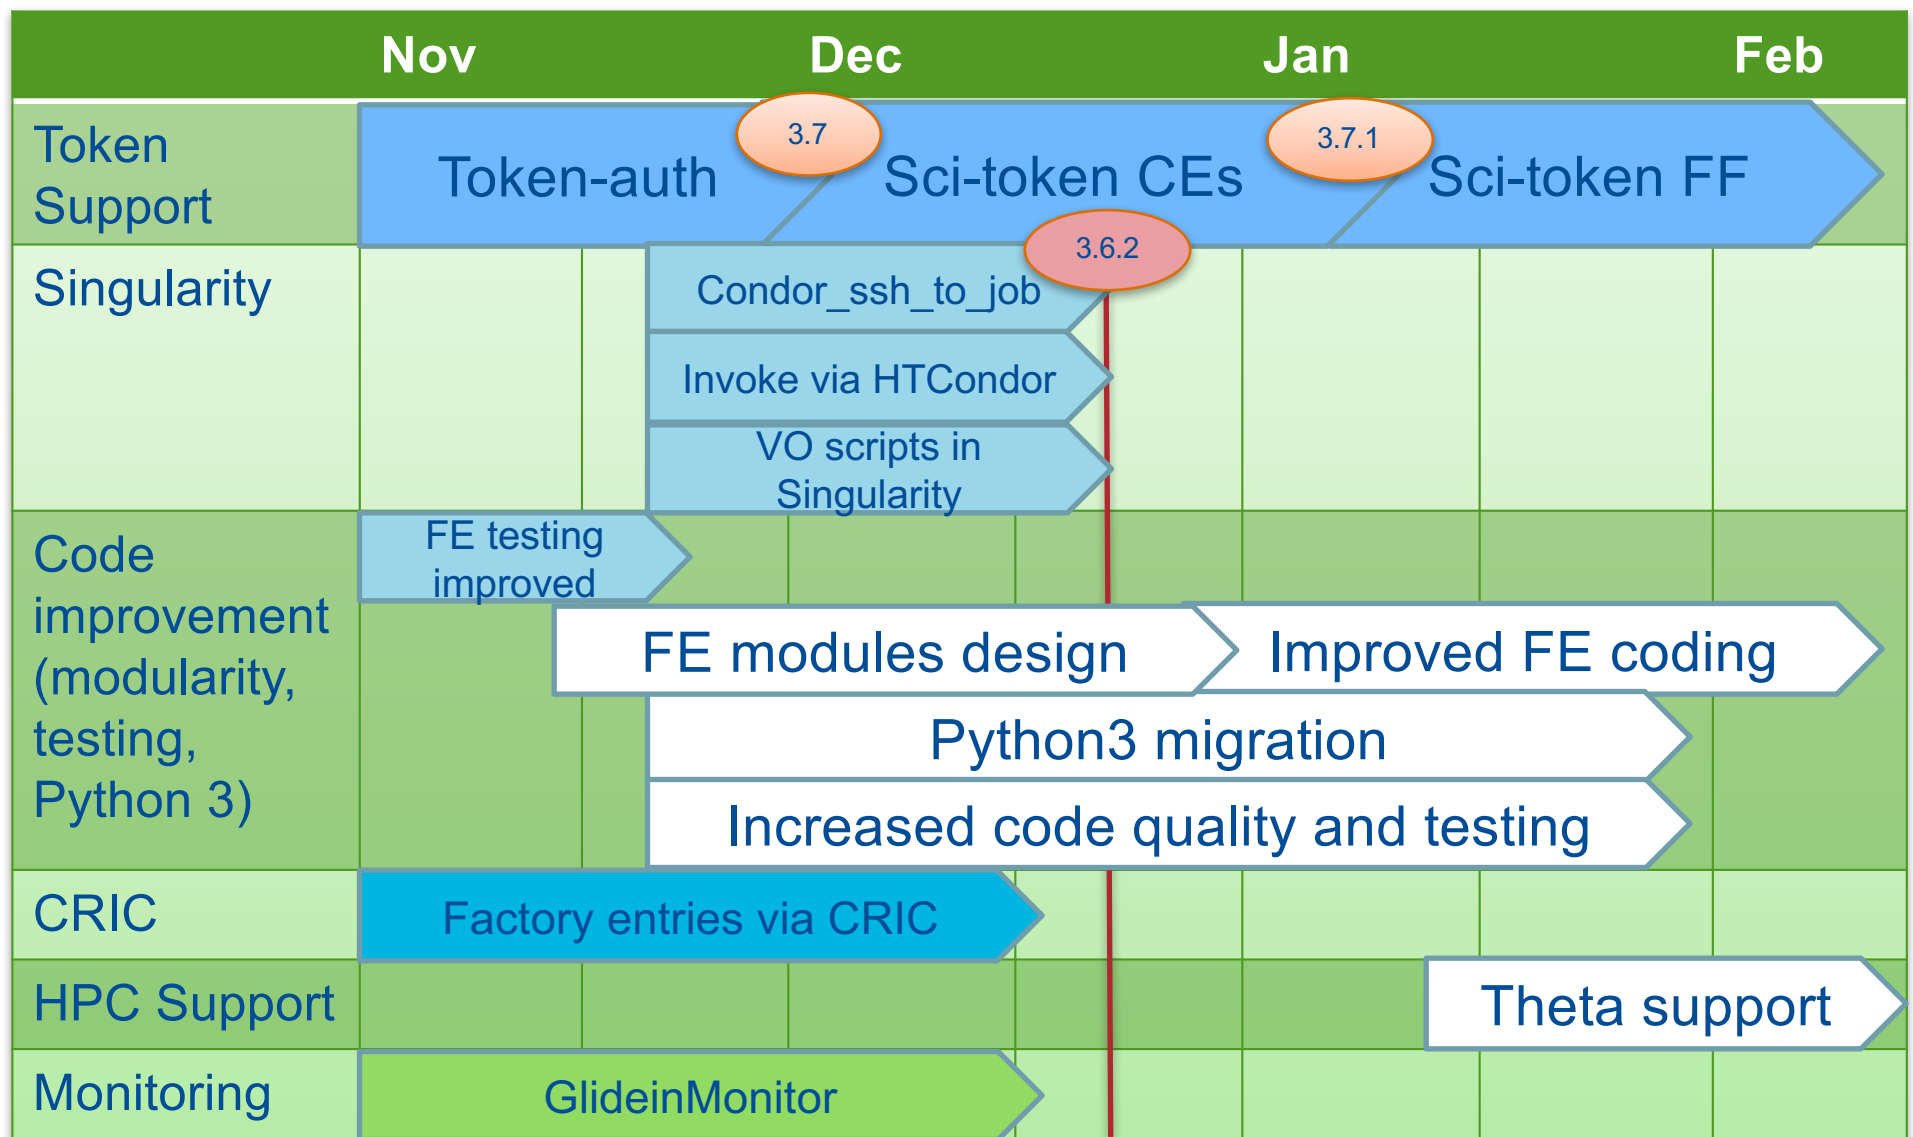
- Scheduled for 3.6.2
 - TAR files distribution
 - Add requirement for HTCondor Python binding
- Planned for 3.7.1
 - GIExec
 - Separate User collector ports (only shared port)

GlideinWMS Roadmap – high priority

- Migration to single-user Factory – Marco Mascheroni
 - Improvement of Factory tools
- Use of token authentication (security without x509 certificates) – Dennis Box
Collaboration w/ HTCondor and OSG
 - Use token-auth to authenticate Glideins
 - Support sites with sci-token
 - Use of tokens to authenticate Factories w/ Frontends
- Singularity support – Marco Mambelli
Collaboration w/ HTCondor
 - Hardening of Singularity and expanding use-cases
 - Having HTCondor invoke Singularity
 - Support condor_ssh_to_job
 - Allow VO test/setup scripts inside Singularity
- Automatic Factory configuration generation, via CRIC (3.6.2) – Marco Mascheroni
- Improve modularity and code quality (especially of Frontend)
 - Improve modularity to include in DE Framework
 - Broaden and streamline testing
 - Migration to Python3
 - Expand, simplify and automate testing

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

Roadmap overview



Dennis Box
Bruno Coimbra
Marco Mambelli
Marco Mascheroni
Interns
Not assigned



GlideinWMS Roadmap - other

- Move main repository to GitHub
- Monitoring Modernization
 - Contributions of Summer interns projects
 - Support standard logging for Glidein and VO scripts (3.7)
 - Extend logging and improve reliability (3.7)
 - GlideinMonitor
 - Move to grafana/graphite/elastic search based solution
 - Retire GlideinWMS monitoring pages
- Collaborate with HTCondor team to support new HPC sites with stricter policies (e.g. no outbound connection except gateways, MFA)

GlideinWMS Roadmap – other (cont)

- Deploy GlideinWMS in containers
- Move processing in HTCondor
- Collaboration w/ HTCondor
 - Auto-clustering to decide about provisions
- Modernize configuration
 - Move to YAML
 - More modular, orthogonal, better default handling
 - Re-evaluate upgrade/reconfig mechanisms
- Move of the documentation to Jekyll
 - Use of templates will ease page maintenance

Completed Release, v3.6.1

- v3_6_1 OSG OSG 3.4 and OSG 3.5, RC now
 - Added compatibility w/ HTCondor 8.8.x in OSG 3.5
 - Monitoring pages use https if available
 - Improved search and testing of Singularity binary
 - UnsetLD_LIBRARY_PATH and PATH for jobs in Singularity
 - Updated documentation links and Google search
 - Improved CI testing
 - Stop considering held limits when counting maximum jobs in Factory
 - Bug fix: Fix Factory tools (entry_rm, entry_q and entry_ls) to be more verbose with single user Factory
 - Bug fix: Removed hardcoded CVMFS requirement for Singularity
 - Bug fix: Improve diagnostic messages when rsa.key file is corrupted

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

Next Planned Production Release, v3.6.2

- v3_6_2 OSG 3.4 and 3.5, expected mid December
 - Automate the generation of factory configuration via CRIC
 - Allow a Frontend to run in parallel w/o affecting the Factory
 - Adopt Singularity mechanisms provided by HTCondor
 - Support condor_ssh_to_job to Singularity jobs
 - Support to run VO scripts within Singularity
 - Adding shell scripts checking to CI
 - Dropping TAR files distribution

https://cdcv.s.fnal.gov/redmine/projects/glideinwms/issues?query_id=182

Next Planned Development Release, v3.7

- v3_7 OSG 3.5, expected in two weeks
 - Support HTCondor token-auth for Glideins
 - Improved Glidein logging
 - Improved Glidein scripts
 - Adding shell scripts checking to CI
 - Dropping TAR files distribution

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

Developers Spotlight

Marco Mambelli

- Discussions to revise priorities and work on roadmap
- Team support
- Development topics
 - Ability to run two Frontend in parallel
 - Singularity support and improvement
 - Changes for 3.6.2
 - Progress on invocation via HTCCondor

Marco Mascheroni

- Trip to Fermilab with Jeff to discuss factory operations topic (minutes [here](#))
- Worked with James on the CHEP [talk](#) about automatic generation of factory entries with CRIC
- Adapted tools used by operations to better work with single user factory
 - New condor_q custom output format to show frontend user as opposed to gfactory
 - fename classad added to replace the Owner when you want to select/show a specific VO
- Stop considering held limits when counting maximum jobs
 - Addresses big sites with lot of opportunistic resources that might suddenly disappear
- Improved diagnostic and error messages in case of (not so) rare file corruption instances
 - Happening in production factory because of I/O issues caused by weekly fstrim
- Revised mechanism to manage restrictions on singularity images and allow non standard locations
- Upgraded the gentle (not so gentle) pilot draining mechanism
 - Now it allows site admin to schedule a downtime in advance (in 3.6.2)

Lorena Lobato

- Participation in the releases of GlideinWMS candidates
- HTTPs Support now available for GlideinWMS monitoring pages
- Single-user factory exhaustive testing
- Regular code reviews
- Discussions with Factory operators
 - Current GlideinWMS testing reliability
 - Multiple (independent) user collectors per frontend
- FIFE ITB Frontend management

Lorena Lobato

- Blackhole detection interaction with HTCondor team
 - [7328: Make knob for the startd to drop final machinead into log on shutdown](#)
 - [#7329: Keep history of updates to machine ads similar to how job ads work](#)
- Grace Hooper Celebration – Interviewed women for different computing projects at SCD
 - Next year GlideinWMS summer intern - Naw Safrin Sattar
 - Streamline complex workflows on HPC project
- Will have a new role in another department– Leaving GlideinWMS project
 - Knowledge Transfer – Bruno Coimbra

Bruno Coimbra - Introduction

- You might already know me
 - Past CMSWeb Operator at CERN (Cat-A)
 - FERRY developer
 - Distributed Computing Support (ongoing)

- Starting at new projects
 - GlideinWMS
 - Landscape

Bruno Coimbra - Current and Future Work

- Learning
 - Going through documentation and learning about the basics
 - Deploying a frontend/factory basic setup to fully understand the system
 - Configuration files still look less cryptic now
- Starting
 - Lorena is helping me with the transition
 - Started working on pychirp (Drop-in replacement of condor_chirp in Pure Python)
 - What comes next:
 - Drop tar ball installation of Factory and Frontend
 - Blacklist improvements
 - Get involved in new activities that might come up

Dennis Box

- Focus has been on Condor TOKEN Auth/ GlideinWMS Integration for 3.7 release
 -
 - A 'proof of concept' implementation was presented and discussed with HTCondor team members 10/22/19 internally reviewed by GWMS team 10/28/19
 - Suggested changes
 - Reduce workload for frontend operators
 - Autodetect if frontend supports tokens
 - Autodetect if factory entry points support tokens
 - By default, use token_auth if both are true
 - Generate distinct token per entry point
 - Makes blacklisting easier to manage

Dennis Box

- Proof of concept was done with minimal code changes
- Code changes are needed to implement desired behavior
- DONE:
 - Default frontend token generation (with tests) completed
- TODO:
 - Pass encrypted tokens to factory by default
 - Code changes to glidein running on CE
 - Blacklisting compromised tokens
 - Documentation for operators

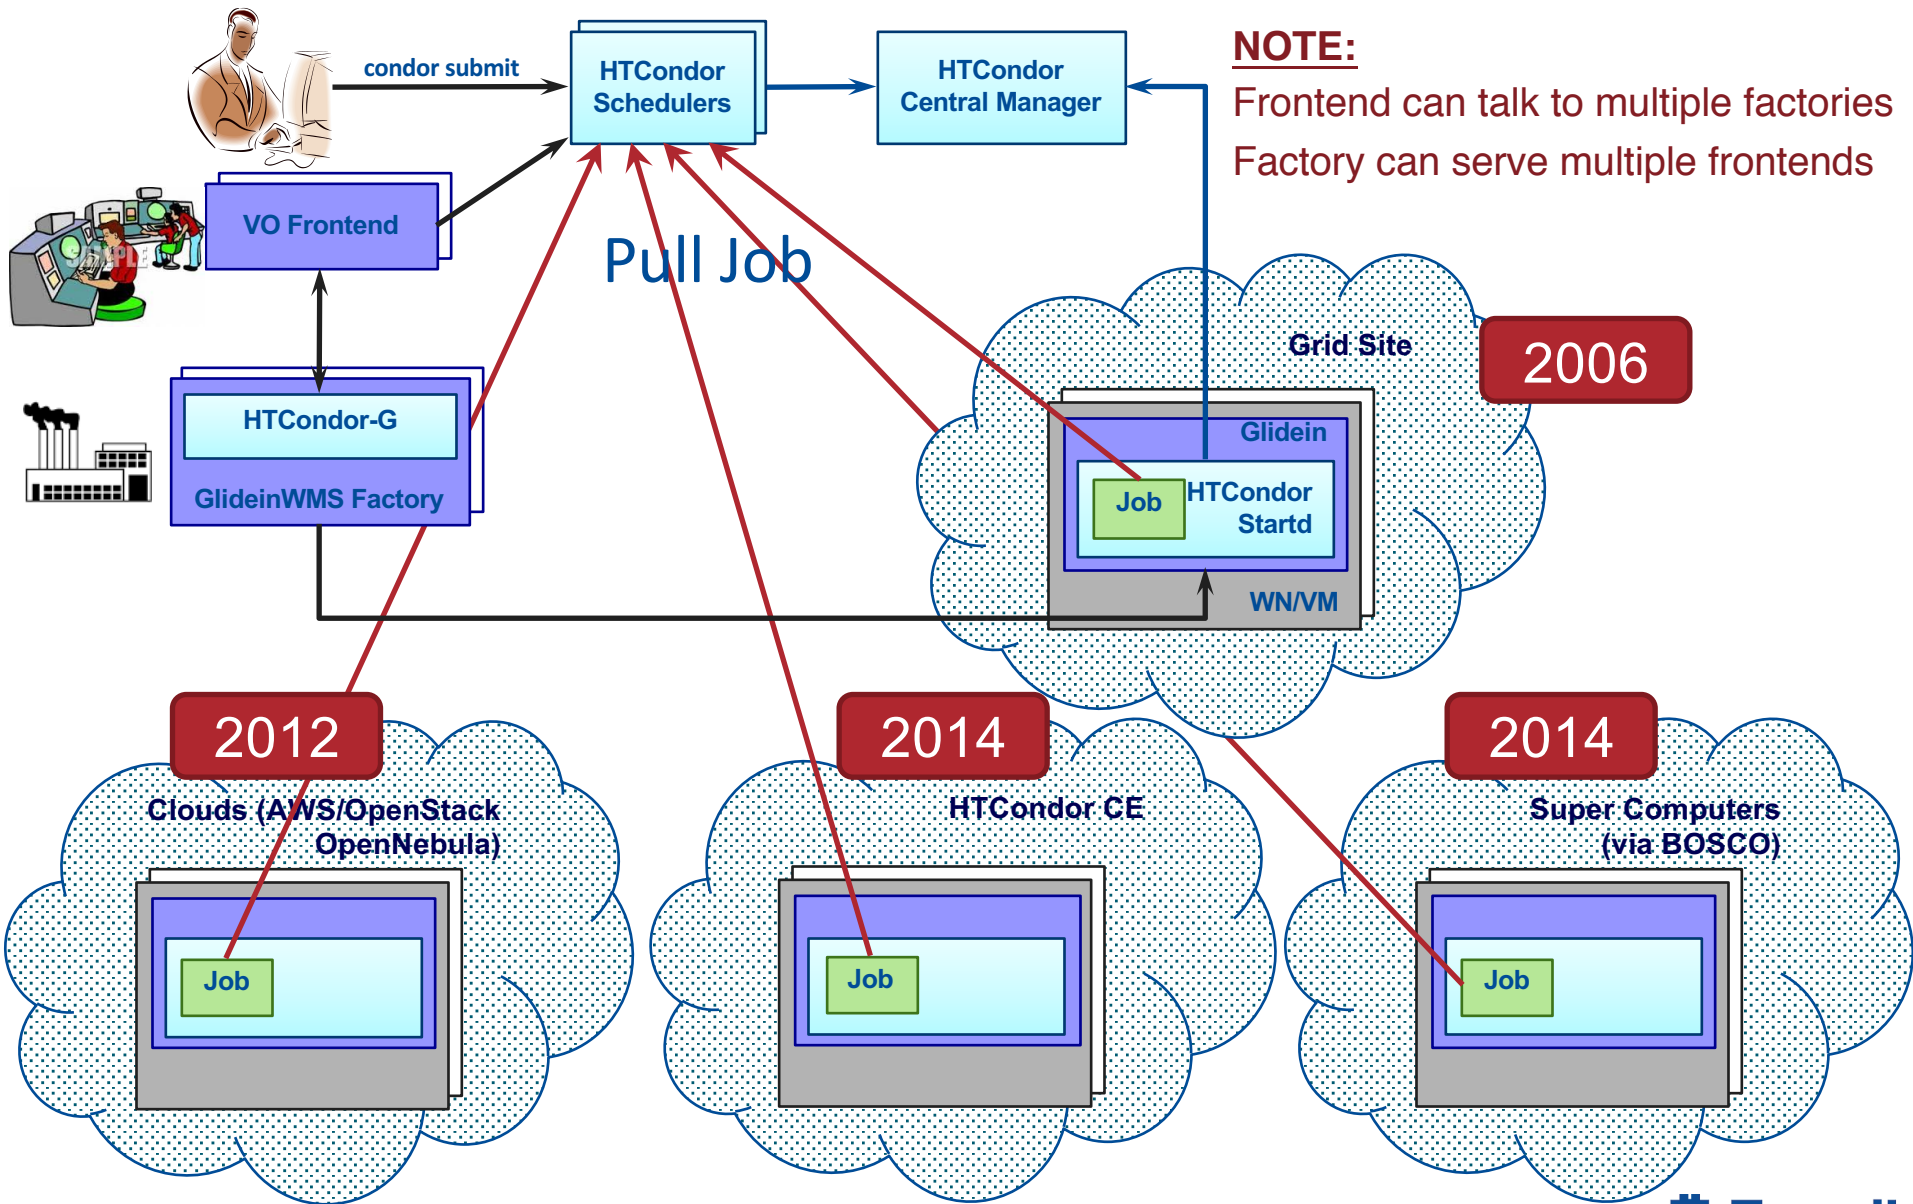
Questions/Comments

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 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	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

The screenshot shows the Redmine interface for the glideinWMS project. The 'Issues' tab is selected, displaying a list of tickets. A search bar at the top right contains 'glideinWMS'. On the right sidebar, the 'My custom queries' section is visible, with 'OSG Tickets' and 'Stakeholder Requests' highlighted by green circles. The main table lists tickets with columns for ID, Tracker, Target version, Status, Priority, Subject, Assignee, and Updated.

#	Tracker	Target version	Status	Priority	Subject	Assignee	Updated	OSG
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
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