



---

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

---

## **GlideinWMS**

Marco Mambelli

Stakeholders Meeting

March 11, 2020

# Agenda

---

- Status update
- Communications
- Important changes
- GlideinMonitor
- Roadmap
- (Reference slides)

# Project status updates

---

- v3.6.1 is current OSG production (OSG 3.4 and 3.5)
- Two new releases in OSG development
  - v3.6.2 RC2, final release later this week
  - v3.7 RC1, final release next week
- Project effort unchanged
  - 2.50 FTE
  - Management and 4 developers
  - Everyone involved also in other projects: Factory Ops, HEPCloud, Jobsub, FIFE support, Landscape

# Communication

---

- Communication
  - Please review periodically your tickets/priorities  
<https://cdcvns.fnal.gov/redmine/projects/glideinwms/wiki/RoadmapSummary>
  - How can we further improve communication
    - Should we participate in any other meetings?
    - Communicating priorities?
- Support
  - Glideins do not renew x509 proxy (v3.4.x v3.5.x v3.6, v3.6.1) - hotfix available
  - Broken GCE and AWS submission in v3.5.x, 3.6, 3.6.1 (fixed in 3.6.2) – hotfix available

# Important changes

---



- All factories should have migrated to single user factories ( $\geq 3.6$ )
- In 3.6.2
  - Add requirement for HTCondor Python binding
- Scheduled for 3.6.3
  - TAR files distribution
- Scheduled/**Planned** for 3.7.1
  - Drop GExec support
  - **Default to shared port in User Collector**

# GlideinMonitor

- Web application to view Glidein log files
  - user interface
  - tools for quick searches and to decode log content
  - efficient managed archive of log files
  - framework to add log processing
- Two components
  - Indexer (glideinmonitor-indexer)
  - Web server (glideinmonitor-webserver)
- Available via RPM (osg-development)
- Code on Github
  - <https://github.com/glideinWMS/glideinmonitor>

[https://docs.google.com/document/d/1PqmQ\\_-JtqK472DT-Y0D08yU1m-lPeFvrp8n4x9f5QE/edit?usp=sharing](https://docs.google.com/document/d/1PqmQ_-JtqK472DT-Y0D08yU1m-lPeFvrp8n4x9f5QE/edit?usp=sharing)

GWMS Factory Job Logs Home Goto GUID

## Factory Monitoring Job View

Filters below alter the data in the table

Click search once you have narrowed the query

Timestamp From  Timestamp To

Show  entries

Entry Name dropdown: entry\_HCC\_US\_Omaha\_cri, entry\_HCC\_US\_Omaha\_crane\_gpu, entry\_ITB\_FC\_CE2\_mc4, entry\_ITB\_FC\_HTCE1, entry\_Lucille\_CE, entry\_ITB\_FC\_CE2

JobID	FileSize	Timestamp	FrontendUsername	InstanceName	EntryName	Maste
job.100.0	18061	2015-11-05T09:28:13-06:00	user_frontend	glidein_gfactory_instance	entry_HCC_US_Omaha_crane_gpu	False
job.101.0	18031	2015-11-05T09:43:31-06:00	user_frontend	glidein_gfactory_instance	entry_HCC_US_Omaha_crane_gpu	False
job.102.0	18027	2015-11-05T09:59:38-06:00	user_frontend	glidein_gfactory_instance	entry_HCC_US_Omaha_crane_gpu	False
job.103.0	18011	2015-11-05T10:09:37-06:00	user_frontend	glidein_gfactory_instance	entry_HCC_US_Omaha_crane_gpu	False
job.104.0	0	2015-11-05T10:09:37-06:00	user_frontend	glidein_gfactory_instance	entry_HCC_US_Omaha_crane_gpu	False

GWMS Factory Job Logs Home Goto GUID

Job 900 Time: 2019-07-20T23:40:15-05:00

### General Information

Timestamp	1563684015
FileSize (err + .out)	196219
Entry Name	entry_ITB_FC_CE2
Instance Name	glidein_gfactory_instance
Frontend Username	user_frontend
GUID	user_frontend@glidein_gfactory_instance@entry_ITB_FC_CE2@job.462237.0

### Data Files

Full Logs

900.out →

900.err →

900.tar.gz →

900.json →

### Condor Logs

Master Log →

Startd Log →

**Starter Log →**

StardHist Log →

XML Description →

### Log Search

Output & Error Log Combined

error

ERROR\_GEN\_PATH /scratch/12131510.1.long/home\_crm09\_017103526/CREAM017103526/glide\_n31060/main/error\_gen.sh

error\_gen.j7j90R.sh: OK


# Completed new features and fixes

---

- 3.6.2
  - OS independent file transfer via “pychirp” (OSG/CMS)
  - Automatic Factory configuration via CRIC (CMS)
  - Reduced the Frontend queries to the collector (CMS)
  - HTCondor GPU monitor enabled (OSG)
  - Fixed x509 proxy renewal
  - Fixed submission to AWS and GCE
- 3.7
  - More robust Glideins using token-auth (OSG)
  - New logging channel
    - Standard API
    - Thread-safe
    - Allows reporting while the Glidein is running
    - Custom endpoint

# Marco Mascheroni – query optimization

- CMS [study](#) on collector queries
  - GWMS frontend is doing 100 queries per minute
- Two types of queries
  - Frequent fast queries to get the information about the schedds
  - Less frequent heavier query to get the status of the pool
- The frequent queries are a good candidate for caching
  - Schedd information rarely change



- Worked on a **patch** that significantly **reduces** the number of small **queries**
- **Tested** in the CMS ITB frontend
  - Number of queries dropped from 2.2k to 250
- Already **applied in production** as a hotpatch.
  - Will be included in 3.6.2



# Marco Mascheroni

Autoconf project

- Finalized and included the work on the automatic generation of factory entries
  - Included new *OSG\_autoconf* tool that generated entry configuration from the OSG collector
- Includes new *gfdiff* tool that does a smart diff between entries
- Will keep iterating once the tools are used in production by operations
- More details in the [CHEP2019 proceeding](#)

Bugfixes

- Fixed an issue with pilot proxies expiring prematurely
  - Could be confused with a pilot regularly reaching the draining time because of the particular landscape of the resources
  - Caltech noticed because their pilots are one week long and ran on 160 cores machine
- Adjusted condor\_q custom format after 3.6 single user migration
- Fixed a bug in the factory monitoring that broke client reports (all zeroes)

Entry Name		Status:									Requested:		Client Monitor:						
		Running	Idle	Waiting	Pending	Staging in	Staging out	Unknown	Held	Running cores	Max glideins	Idle	Claimed cores	User run here	User running	Unmatched cores	User idle	Registered cores	Info age
Glow_US_Syracuse_condor-ce2	↑	1191	1061	1002	53	0	0	6	4123	1191	41007	75	1080	0	0	0	0	0	0
Glow_US_Syracuse_condor-ce3	↑	919	77	21	54	0	0	2	860	919	2581	69	751	0	0	0	0	0	0
OSG_US_IIT_litce1	↑	823	75	0	75	0	0	0	112	823	1902	92	825	0	0	0	0	0	0
OSG_US_UConn_gluskap_op	↑	796	91	0	91	0	0	0	102	796	3202	100	791	0	0	0	0	0	0

???

# Bruno Coimbra - Introduction

---

- Python drop-in replacement for condor\_chirp (PyChirp)
  - It's now called simply condor\_chirp and ships with glideins
  - The script itself wraps a python client extended from htchirp
  - Also implemented the gwms\_python wrapper
    - Finds the most recent version of Python
    - Adds GlideinWMS python scripts to PYTHONPATH
- Started working on SciTokens for GlideinWMS
- Started working on dropping tarball installation support from GlideinWMS

## Dennis Box – in v3.7

---

- HTCondor Token authentication available in 3.7 (development) release
- Requirements
  - Condor version 8.9.1+ on Frontend and Glidein tarball
  - Factory and Frontend both 3.7+
- Behavior
  - Frontend detects condor version of Factory entry point tarball
  - If OK, generates and encrypts condor token, advertises to Factory
  - Factory decrypts and stores Frontend generated token, re-encrypts and passes to glidein
  - Glidein uses token to communicate with Frontend Daemons

## Dennis Box – coming next

---

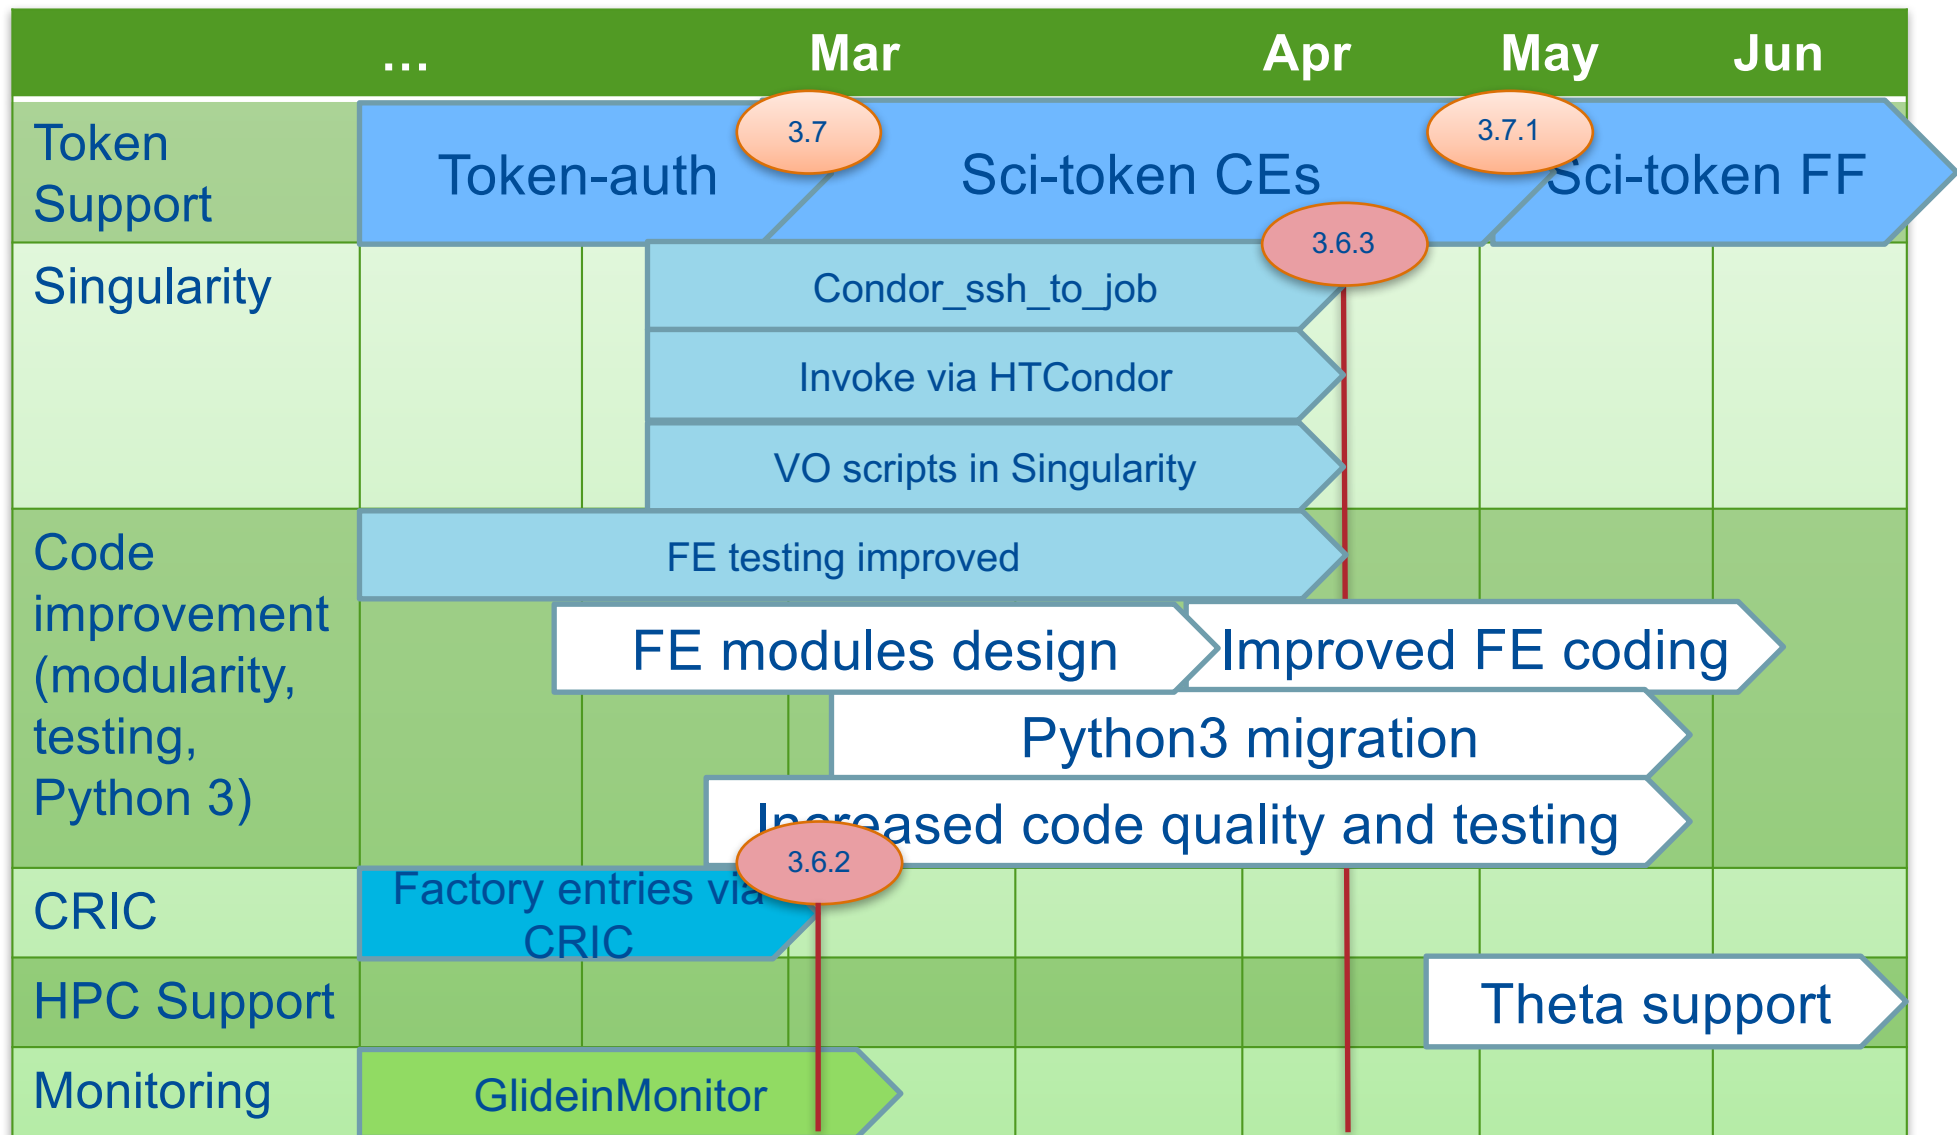
- Currently working on SciToken integration
- Parts that work so far:
  - Self signed + generated SciTokens authenticate with local installation of HTCondor-CE-4.0+
  - condor\_ping authenticates self-signed tokens against daemons in more complicated (frontend, factory) installations
  - CredMon, WLCGTokens can be made to authenticate with minimal 'personal condor'
- Remaining: everything else.

## Coming next

---

- Use of token authentication (OSG) – Dennis and Bruno
  - Support sites with sci-token
  - Use of tokens to authenticate Factories w/ Frontends
- Singularity support (CMS/OSG/Fermilab) – Marco Mambelli
  - Simplified code by Singularity invocation via HTCondor
  - Simplified troubleshooting by allowing ssh to job in Singularity
  - Allow VO test/setup scripts inside Singularity
- Improve modularity and code quality (HEPCloud)
  - Improve modularity of Frontend to include in DE Framework
  - Broaden and streamline testing
  - Migration to Python3
  - Expand, simplify and automate testing
- Leadership Computing Facilities support (HEPCloud/CMS)

# Roadmap overview



Dennis Box

Bruno Coimbra

Marco Mambelli

Marco Mascheroni

Interns

Not assigned



# Summary

---

- Project mostly on track
- Development effort is limited and frequently pulled away
- Update to 3.6.2 (end of March)
  - Cloud submission and Pilot x509 fixes
  - pychirp and Automatic configuration via CRIC
  - Less collector queries from the Frontend
- Development version 3.7
  - Improved logging
  - Token-auth for Glideins
- Top priorities
  - token support
  - VO tests in Singularity
  - Frontend restructuring
  - Use of LCF computers

---

# Reference Slides

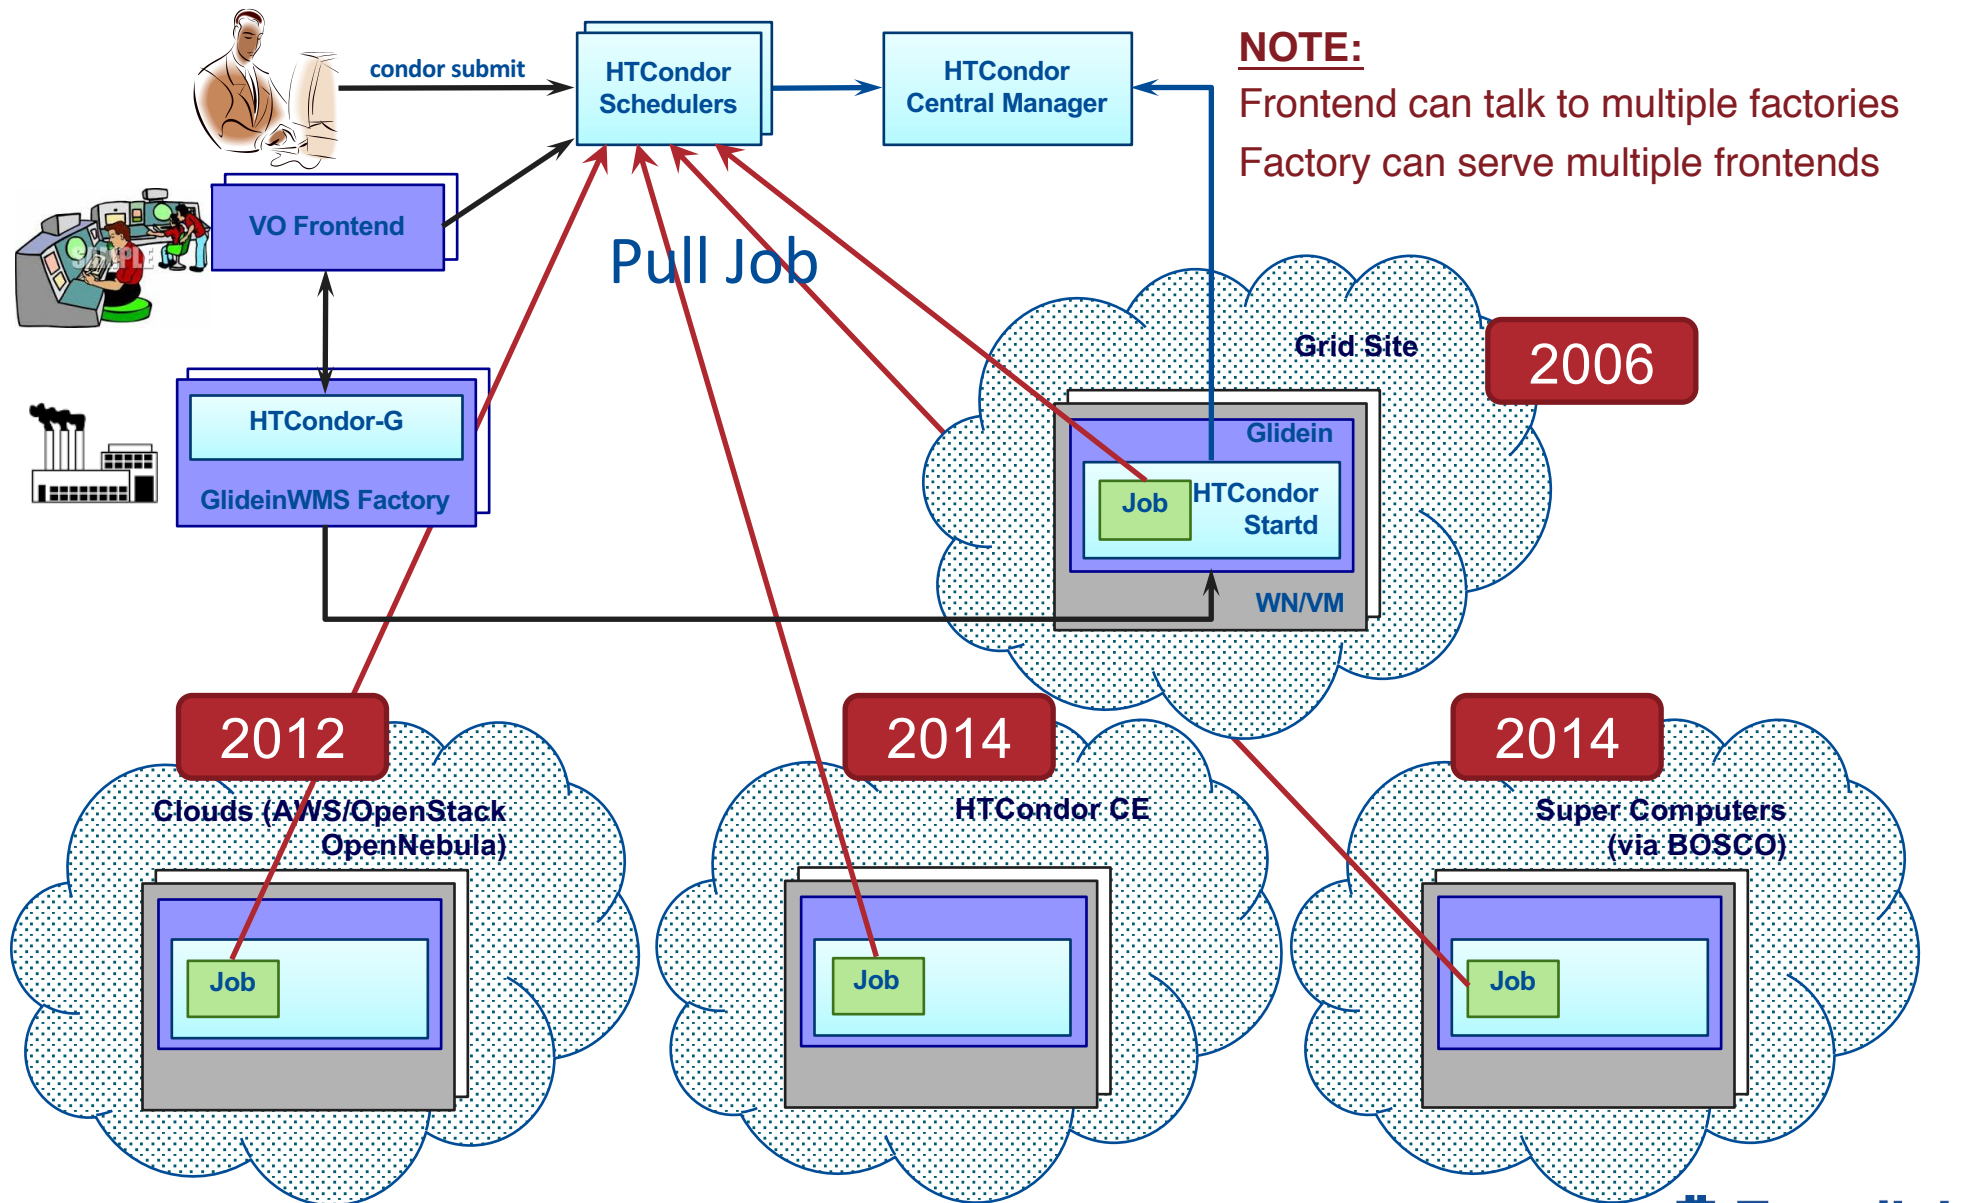


# Move to single user Factory

---

- 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://cdcvcs.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

glideinWMS

Overview Activity Roadmap **Issues** New Issue Gantt Calendar Documents Wiki Files Repository Code reviews Settings

- OSG Tickets

Filters Options

Apply Clear

#	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

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