



---

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

---

## **GlideinWMS**

Marco Mambelli

Stakeholders Meeting

September 18, 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
  - GlideinWMS v3.4.6 released August 8, in OSG testing, eligible for production
  - GlideinWMS v3.5.1 released September 17
  - Seeking stakeholders input for future GlideinWMS releases • Dropping support for GT2/GT5, Glexec, python 2 • See Marco's talk for details •
- Project Effort (2.80 FTE)
  - Project Management: 0.15 FTE
  - Development & Support: 2.65 FTE
- Temporary effort
  - 1 Summer Interns and 1 on call collaborator

# Project Updates Since Last Stakeholders Meeting

---

- Communication
  - Problem in patch for Singularity
    - An early release of a patch via email caused problems
    - Added procedure for patches
  - How can we further improve communication
    - Should we participate in any other meetings?
    - Communicating priorities?
  
- Support

# Action Items from Previous Stakeholders Meeting

Action Items	Status
Send a reminder about the agreed plans for HTCondor binding requirement, drop of tar distributions, shared ports becoming a default	Done
Discussion about CREAM support in HTCondor, OSG and GlideinWMS	In progress
Discussion about GlideinWMS in containers: deployment and state	To do
Ask Edgar about access to resources to test MPI jobs	In progress
Start collaboration between Edgar and Thomas, to coordinate the monitoring effort	Done
Discussion about the GLIDEIN_Custom_Start	Done
Discussion about publishing the Glidein Logs	In progress

# Completed and Next Planned Releases

---

- Released
  - GlideinWMS v3.4.6 released August 8, in OSG testing, eligible for production
  - GlideinWMS v3.5.1 released September 18
  
- We have 2 releases in the pipeline
  - v3.4.7 production series OSG 3.4, dropped in favor of 3.5.x series in OSG 3.4 (HTCondor 8.8 support in the Factory)
  - v3.5.2 in the production series for OSG 3.4 and 3.5, end of October.
  - v3.6 in OSG upcoming, mid October

## Completed Release, v3.4.6

---

- v3\_4\_6 OSG production, released August 8, soon OSG 3.4
  - Fix problem with DNs including commas
  - Fix Factory compatibility w/ older 3.4.x Frontends
  - Singularity support fixed and improved
    - Fixed Debug options causing Singularity invocation to fail
    - Better GPU support
    - More robust work-dir creation
  - Document and expand multi-node Glidein
  - Site-customized pilots
  - Simplify usage of manual\_glidein\_submit
  - Backport: GlideinWMS proxy renewal service broken for Xenon
  - Fixing chkconfig lines on proxy renewal

<https://glideinwms.fnal.gov/doc.prd/history.html>

## Completed Release, v3.5.1

---

- v3\_5\_1 OSG OSG 3.4 and OSG 3.5, September 18
  - Include 3.4.6 features
  - Improved documentation and scripts to migrate Factories from 3.4.x
  - Improved manual\_glidein\_startup
  - Advertise if a Glidein can use privileged or unprivileged Singularity
  - Added release lifetime and compatibility statements
  - Streamlined and documented release testing

[https://cdcvs.fnal.gov/redmine/projects/glideinwms/wiki/ReleaseTestingMatrix\\_3\\_5\\_1\\_rc1](https://cdcvs.fnal.gov/redmine/projects/glideinwms/wiki/ReleaseTestingMatrix_3_5_1_rc1)

[https://cdcvs.fnal.gov/redmine/projects/glideinwms/issues?query\\_id=53](https://cdcvs.fnal.gov/redmine/projects/glideinwms/issues?query_id=53)



## Next Planned Release, v3.5.2

---

- v3\_5\_2 OSG 3.5, expected end of October
  - Black hole prevention in Glideins
  - Automate the generation of factory configuration via CRIC
  - Adopt Singularity mechanisms provided by HTCondor
  - Support condor\_ssh\_to\_job to Singularity jobs
  - Fix Factory monitoring when interacting with Decision Engine
  - Factory and Frontend monitoring under https
  - 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=182](https://cdcv.s.fnal.gov/redmine/projects/glideinwms/issues?query_id=182)

## Next Planned Release, v3.6

---

- v3\_6 OSG upcoming, expected for mid of October
  - HTCondor token-auth for Glideins

# GlideinWMS Roadmap – dropping support for...

---

- Scheduled for 3.5.2
  - TAR files distribution
  - Add requirement for HTCondor Python binding
- Planned for 3.6 (possibly some 3.5.x - Fall)
  - GIExec
  - Separate User collector ports (only shared port)
- Planned for 3.7 (Fall- 3.6 will be in parallel until Spring 2020)
  - Python2
  - Is it OK to move to support only Python 3 by the fall?

# GlideinWMS Roadmap – high priority

---

- Use of token authentication (security without x509 certificates)  
Collaboration w/ HTCondor and OSG
  - Use token-auth to authenticate Glideins (3.6)
  - Support sites with sci-token (3.6.1)
  - Use of tokens to authenticate Factories w/ Frontends
- Singularity support  
Collaboration w/ HTCondor
  - Improving singularity support (Unprivileged Singularity, more robust site support, better logging, ...)
  - Adding new features used by VOs (libraries, robust GPU support, `condor_chirp` ...)
  - Having HTCondor invoke Singularity
  - Support `condor_ssh_to_job`
  - Allow VO test/setup scripts inside Singularity
- Automatic Factory configuration generation, via CRIC (3.5.2)
- Factory supporting multiple frontend like services
  - HEPCloud/Decision Engine support started in 3.4.4
- Modernize and simplify code. Broaden and streamline testing

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

# GlideinWMS Roadmap - other

---

- Move to Python 3
  - Branch with Python 3 migration
  - Have a Python 3 version in OSG upcoming by mid Fall 2019
- Monitoring Modernization
  - Contributions of Summer interns projects
    - Support standard logging for Glidein and VO scripts (3.5.2)
    - Extend logging and improve reliability (3.5.3)
    - 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

---

# Developers Spotlight

# Marco Mambelli

---

- Singularity support and improvement
- Joint effort to solve HTCondor not being killed in PBS clusters
- Monthly code discussion and challenge of the month
- Summer projects
  - GlideinMonitor
  - Improved Glidein logging
- Development topics
  - Singularity support and improvement
    - Invocation via HTCondor in 3.5.2
    - Easy VO scripts for testing and setup



# Marco Mascheroni

---

## Items I have been working on that will go into 3.5.1

- Single user factory: written script that checks if the factory is 3.5 ready
  - Upgrading from 3.4 requires to change ownership of jobs and log/proxy directories
  - It runs at startup and checks that (1) directory ownership have been changed to gfactory, and (2) that all the jobs belongs to gfactory
- Detected and fixed a case when the factory could not be restarted after a file corruption
- Better documentation for manual\_glidein\_startup (aka glideins in a vacuum)
  - Allows sites to start glideins directly on the WN
- Custom pilots included in 3.4.6
  - Add possibility of customizing the pilot start expression on the WN
  - Production 3.4.5 factories were already patched for CMS
- Currently working on:
  - Better handling of constant parameters
  - Improving gentle/hard draining of resources as for stakeholders feedback

# Lorena Lobato

---

- Participation in the releases of GlideinWMS candidates
- Mentor/support summer interns
- More reinforcement in testing: proposed ideas + improving documentation
- Code review - Python bindings + context managers
- FIFE ITB Frontend to test configurations changes and containers – singularity
- Blackhole detection (**expected for 3.5.2**)
  - Interaction with HTCondor team
- Will have new role – operations (20%)
  - ITB Frontend and production – Access Factory

## Dennis Box

---

- Condor TOKEN Auth/ GlideinWMS Integration
  - OSG is moving from GSI to JWT SciToken auth.
  - Condor 8.9.2 (development) supports JWT based Condor TOKEN Auth, and SciToken auth with some fiddling.
- Phase 1
  - Upgrade working factory/frontend to Condor 8.9.2
  - Convert frontend auth to TOKEN
  - Modify running glideins to phone home using TOKEN
  -
- Phase 2
  - Modify Factory to accept and auth with tokens from FE
  - Incorporate SciTokens into Factory/Frontend.

# Dennis Box

---

- Progress

Phase 1 close to completion

Done mostly with config file changes

Tokens currently generated by hand, location pointed to in frontend.xml

One new script uploads to CE from frontend, ensures libraries and tokens in correct places prior to proceeding with authorization

Redmine tickets have more detail

Use token auth for GWMS

<https://cdcvs.fnal.gov/redmine/issues/23092>

Condor 8.9.2 config changes

<https://cdcvs.fnal.gov/redmine/issues/23278>

# Leonardo Lai, summer intern

---

- Graduate student at Sant'Anna School, Pisa, Italy
  - BSc in Computer Engineering
  - MSc (currently) in Embedded Computing Systems
- As summer intern, I work on the development of a new *flexible, reliable and secure logging channel* for glideins
  - Rationale:** collect info from glideins, most useful when they fail
  - Current shortcomings:** unstructured logs; report available only at the end of Glidein's life; missing info for killed ones
  - ✓ **Step 1:** shared interface for Glidein bash scripts to log info.  
Logs are stored in JSON format
  - ✓ **Step 2:** flush logs to remote http server(s) upon request
  - ⚙ **Step 3:** token authentication with server, configurable options
- Other activities: static analysis, testing, refactoring of code

# Leonardo Lai – Log file example

```
[
  {
    "uuid": "a8fd0e49-7a46-4ae6-b188-fc1c5256707e",
    "name": "gfactory_instance",
    "factory": "gfactory_service",
    "client": "fermicloud093-fnal-gov_OSG_gWMSFrontend.main",
    "client_group": "main",
    "schedd": "schedd_glideins4@fermicloud070.fnal.gov"
  },
  {
    "invoker": "glidein_startup.sh",
    "pid": "16070",
    "timestamp": "2019-08-19T17:02:06-05:00",
    "severity": "info",
    "type": "text",
    "filename": "",
    "content": "Remote logging has been setup. Server address: http://fermicloud152.fnal.gov:80"
  },
  {
    "invoker": "glidein_startup.sh",
    "pid": "16070",
    "timestamp": "2019-08-19T17:02:17-05:00",
    "severity": "warn",
    "type": "text",
    "filename": "",
    "content": "curl not installed on this machine"
  },
  {
    "invoker": "glidein_memory_setup.sh",
    "pid": "29801",
    "timestamp": "2019-08-19T17:02:18-05:00",
    "severity": "debug",
    "type": "file",
    "filename": "logging_utils.source",
    "content": "begin-base64 644 -H4sIAGocW10AA+MCAJMG1zIBAAAA===="
  }
]
```

- GlideinMonitor
  - Tagging, Indexing and archiving Glidein Logs
  - Serving logs with a client side viewer
  - Some key features
    - Parses output, error, & condor logs
    - Minimal dependencies
    - Low overhead on server
    - REST API available
    - Log search for output, error, and condor logs
  - Demonstrator, future development
    - Logs retrieval (rsync, http?)
    - Data retention model
    - Sanitize logs
    - Security model
    - Packaging (containers?)

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: 08/13/2019 1:47 AM Entry Name: entry\_HCC\_US\_Omaha\_cr...  
Search

Copy Excel Print

Show 10 entries

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

Open

Full Logs	Condor Logs
900.out →	Master Log →
900.err →	Startd Log →
900.tar.gz →	Starter Log →
900.json →	StardHist Log →
	XML Description →

#### Log Search

Output & Error Log Combined

error Search

ERROR\_GEN\_PATH /scratch/12131510.1.1ong/home\_crm09\_017103526/CREAM017103526/glide\_n31060/main/error\_gen.sh  
error\_gen\_37j98R.sh: OK

---

# 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.

# Dennis Box

---

- Condor TOKEN Auth/ GlideinWMS Integration
  - Rationale:
    - OSG would like to move from GSI authentication to a JWT based SciToken scheme
  - Current Status:
    - Condor 8.9.2 (development series) supports JWT based Condor TOKEN Auth out of the box.
    - Also supports SciToken auth with inclusion of proper plugins/libraries
    - GlideinWMS documentation already refers to using tokens for authentication
      - Requires configuration file changes, no examples given.
      - Has never been integrated with Condor Tokens as they have just been introduced.

# Dennis Box

---

- Step 1: Convert GWMS Frontend to use Condor 8.9.2/TOKEN AUTH

<https://cdcvs.fnal.gov/redmine/issues/23092>

Start with a working frontend/factory

Upgrade condor to 8.9.2

8.9.2 config file changes

<https://cdcvs.fnal.gov/redmine/issues/23278>

Create a master password file and token

```
condor_store_cred -f /etc/condor/passwords.d/password
```

```
condor_token_create >> /etc/condor/tokens.d/admin.token
```

```
Change SEC_DEFAULT_AUTHENTICATION_METHODS  
from FS,GSI => TOKEN,GSI; condor_reconfig
```

Result: FS auth completely replaced with TOKEN on frontend  
as verified by condor\_ping

# Dennis Box

---

- Step 1: (continued) Note that if you were using FS to do condor commands previously they now fail
- - `$ condor_submit testjob.singularity.jdf`
  - Submitting job(s).
  - ERROR: Failed to create proc
  -
- Export `X509_USER_PROXY=(your_proxy)` for GSI auth or
- `condor_token_fetch` or `condor_token_create` for TOKEN auth
- `condor_submit`, `condor_q`, etc work as they did before.
- 
- 
-

# Dennis Box

---

- Step 2: hand configure a glidein on CE to run jobs for frontend using TOKEN auth

Configure entry points on CE to use Condor 8.9.2 tarball (so we can authenticate with TOKEN)

Submit some jobs that start glideins on the CE

Copy one of the running glidein directorys on the CE to a new directory you can edit

```
Cp /tmp/glide_0S1zVj /tmp/glide_token_cfg
```

```
Cd /tmp/glide_token_cfg
```

Edit glidein\_config and condor\_config, changing all instances of glide\_0S1zVj to glide\_token\_cfg

Main/condor\_start.sh glidein\_config should start a glidein that connects to frontend. Verify with condor\_status

## Dennis Box

---

- Step 2 (continued): You can start a correctly configured GSI glidein. Make it a TOKEN authenticating glidein

You need some tools and libraries that do not (yet) come with the glidein. Copy them over from the frontend by hand

`condor_token_list`, `condor_token_fetch`

`LibMunge.so`, `libSciTokens.so`

condor config changes best go to an include file

`condor_config_startd_cron_include` so they dont get overwritten when starting a new glidein

`condor_token_fetch -name frontend.ip:9618`

When you get it right `condor_ping` shows you are authenticating via TOKEN

See the redmine tickets for actual settings


## Dennis Box

---

- Step 3 Configure factory and frontend so they create a glidein configured like the hand configured example in step 2.
  -
- Will involve code changes on both frontend and factory
- Frontend.xml already has 'token' as an auth type but I haven't been able to make any attempt to use it pass a reconfig, or find a working example.
- Looking at Leos branch right now, he has coding changes to use tokens for various things and I want to make my changes compatible



# 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

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

#	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 Holzman	02/03/2015 02:44 PM	OSG

Issues

- View all issues
- Summary
- Calendar
- Gantt

My custom queries

- Issues Awaiting My Feedback
- Unlinked 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