



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

GlideinWMS

Parag Mhashilkar

Stakeholders Meeting

November 15, 2016

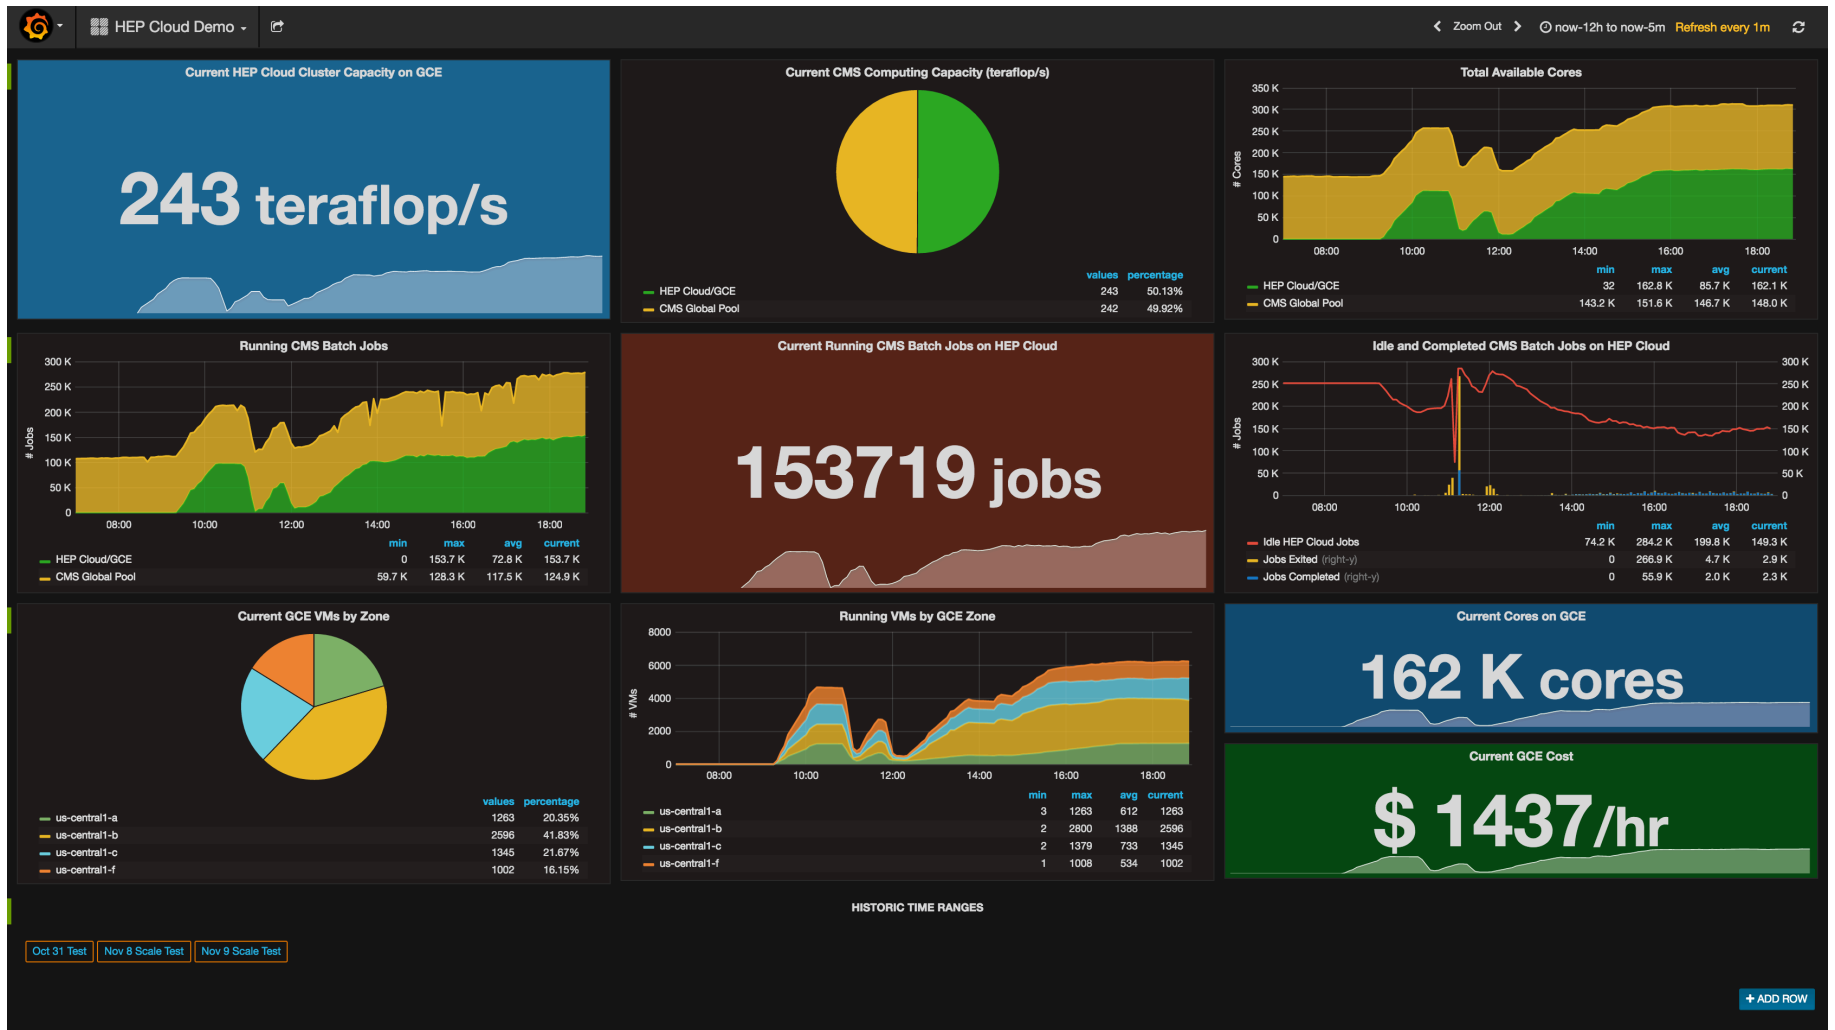
Overview

- Updates since last stakeholder's meeting
- Upcoming releases
- Reference slides
 - GlideinWMS Architecture
 - Quick Facts
 - Releases since last stakeholders meeting

Decade of Glidein-ing!!!

Sep 19, 2016 marked 10 successful years since first official GlideinWMS commit was registered in SCM!!!

Super Computing: HEP Cloud – Google Demo Scale



Highlights Since Last Stakeholders Meeting

- Releases: (Details in Reference Slides)
 - v3_2_15: August 17, 2016
 - Monitoring improvements
 - Auto detect worker node draining and adapt
 - Ability to keep idle glideins at site at all times
 - v3_2_16: October 21, 2016
 - Machine and Job feature attributes now available in the glidein startd's classad
 - VO Frontend can now control BOSCO credentials
 - Ability to launch glidein at site manually
 - Several bug fixes that address stability and accounting issues
 - v3_3_1: October 25, 2016
 - Based on v3_2_16
 - v3_3: August 30, 2016
 - Based on v3_2_15
 - Support for Google Compute Engine
 - AWS cloud related requests from HEPCloud
 - Allow updating AWS credentials in frontend without need to reconfig/restart the service
 - Improve frontend policy configuration

Highlights Since Last Stakeholders Meeting

- Communication
 - Regular participation of CMS and OSG Factory Ops in the weekly GlideinWMS meetings
- Support
 - Usual support load with intermittent peaks
- Project Effort (3.1 FTE)
 - Project Management: 0.15 FTE
 - Development & Support: 2.95 FTE
 - Parag Mhashilkar: **-0.3** (0.75 ➔ 0.45)
 - Dennis Box (New): **+0.5**

Milestones

- Previous Milestones
 - “Why is my job not running”?
 - Scheduled for ~~v3.2.16~~ v3.2.17+
- New Milestones
 - Support for Google Compute Engine
 - Complete: Scheduled for v3.3
 - Support frontend to provide BOSCO credentials
 - Complete: Scheduled for v3.2.16
 - Support advertisement for Machine Job Features (MJF) in startd’s classad
 - Complete: Scheduled for v3.2.16
 - Support draining a glidein if the worker node is marked for draining
 - Complete: Scheduled for v3.2.16

Upcoming Releases - Production Series (v3.2.x)

- Primary Focus of Production Series:
 - High impact bug fixes and features that do not break backward compatibility

v3_2_17+ - Tentatively early January 2016

- Frontend monitoring to store number of job restarts
- Support SL7 sysctl scripts
- Advertise payload info in the glideins classad on glidein termination
- Collect performance statistics for the factory and frontend services for health monitoring
- Balancing requests to factory entries that are aliases or meta sites
- Scale factory to O(600+) entries
- “Why is my job not running?”

Upcoming Releases - Development Series (v3.3.x)

- Primary Focus of Development Series:
 - Production quality but some features maybe experimental
 - Factory/Frontend Configurability
- Next Release: v3.3.2
 - Driven by stakeholder requests

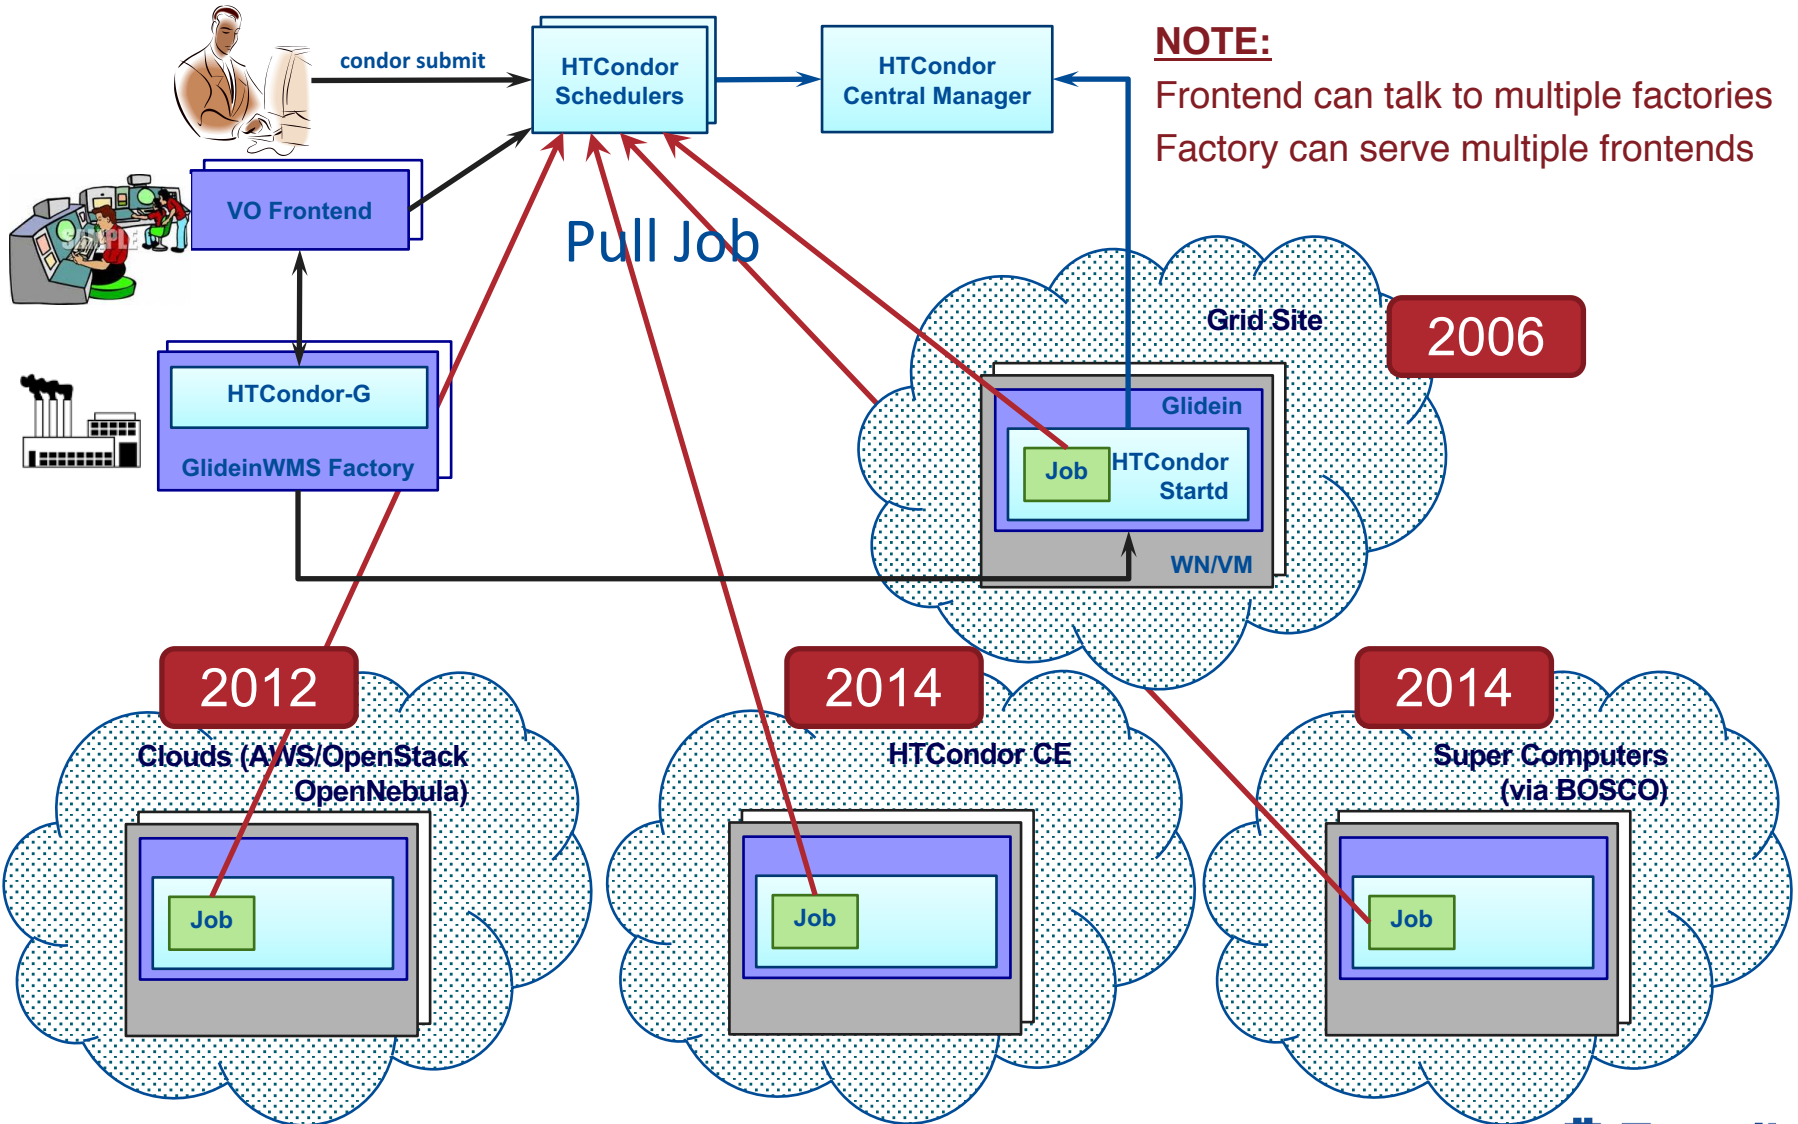
v3_3_2 – Not scheduled yet

- Investigate use of HTCondor's autocluster feature to improve frontend performance

Questions/Comments

Reference Slides

GlideinWMS



GlideinWMS: Quick Facts

- GlideinWMS is an open-source product (<http://tinyurl.com/glideinWMS>)
- Heavy reliance on HTCondor (UW Madison) and we work closely with them
- Effort:

Role	Resources	Effort (FTE)
Project Mgmt/Lead	Parag Mhashilkar (0.15 USCMS)	0.15
Development & Support	Parag Mhashilkar (0.45 SCD) Marco Mambelli (1 SCD) Hyunwoo Kim (0.5 SCD) Dennis Box (0.5 SCD) Marco Mascheroni (0.5 CMS - Contractor)	2.95
TOTAL		3.1

Table: Current Resources & Roles

- Additional Code Contributions (Past year)
 - Jeff Dost (UCSD)

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: Mondays 3-4pm
 - Technical discussions & status updates
 - Regular stakeholder participation
 - Contact Parag Mhashilkar if you need invite for this meeting
- Quarterly Stakeholders Meeting
- 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

Default tabs not too useful

2. Click custom query for stakeholder or version roadmap

Issues

My custom queries

- Issues Awaiting My Feedback
- My Unresolved 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
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-g-analyze	Burt Holzman	02/03/2015 02:44 PM	OSG

GlideinWMS Releases - Key Features

v3_2_16 – October 21, 2016

- Information about job and machine features are now available in glidein startd's classad
- New tool `manual_glidein_startup` now lets you manually start a glidein for a given factory entry and frontend group
- BOSCO entries can now accept credentials provided from the frontend
- Bug Fix: Glidein is now correctly auto detect RHEL6 and RHEL7
- Bug Fix: Dagman and Schedd universe jobs are not counted against max jobs running
- Bug Fix: An entry in downtime will now show up in the glidein status
- Bug Fix: Added `condor-python rpm` as a dependency
- Bug Fix: Complex credentials (`vm_id+vm_type`) are now correctly interpreted
- Bug Fix: Frontend is now more resilient to transient errors communicating with the HTCondor daemons
- Bug Fix: Slave frontend now correctly looks up the master frontend when using `condor-python` bindings
- Bug Fix: `create_condor_tarball` now correctly includes required globus libraries from `lib/condor` that are loaded by HTCondor at runtime
- Bug Fix: Frontend now correctly calculates the `ReqMaxGlideins` for single core glidein entries
- Bug Fix: Numerical data in glideclient classad is not quoted to preserve the data type
- Bug Fix: Frontend will not request any glideins at entry that is in downtime

GlideinWMS Releases - Key Features

v3_2_15 – August 17, 2016

- Factory will now automatically remove unrecoverable glidein jobs with forcex if they are held for 20 times or more
- Several X509 related attributes from the matching job are now also available in the glidein's startd's classad
- Prefix for attributes created by the periodic scripts is now customizable
- If the glidein detects that a worker node is marked for draining and if it is approaching the drain time, it will now kill the user job
- Frontend can be configured to request idle glideins at all times irrespective of the jobs in the queue
- Bug Fix: Fixed a bug in the frontend downtime command in the script /etc/init.d/gwms-frontend
- Bug Fix: Frontend now correctly considers group credentials before frontend's global credential
- Bug Fix: Installing and upgrading GlideinWMS rpms now correctly trigger httpd and htcondor reload commands in case of EL7
- Bug Fix: Fixed a bug where frontend would crash with pickling error while using htcondor-python bindings
- Bug Fix: Factory now correctly remembers pilots submitted using RFC proxy
- Bug Fix: Fixed a bug where running reconfig or upgrade command on a frontend service would crash

GlideinWMS Releases - Key Features

v3_3_1 – October 25, 2016

- Includes all features and bug fixes released in v3_2_16

v3_3 – August 30, 2016

- Includes all features and bug fixes released in v3_2_15
- Support native configuration of EC2 spot prices and AZ in the entry
- Support Google Compute Engine CE. Requires glidein cloud vm rpms v2+
- Support frontend policies specified in external python file
- Support changes to VM ID and VM Type without need to reconfig/upgrade frontend service