



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

GlideinWMS

Parag Mhashilkar

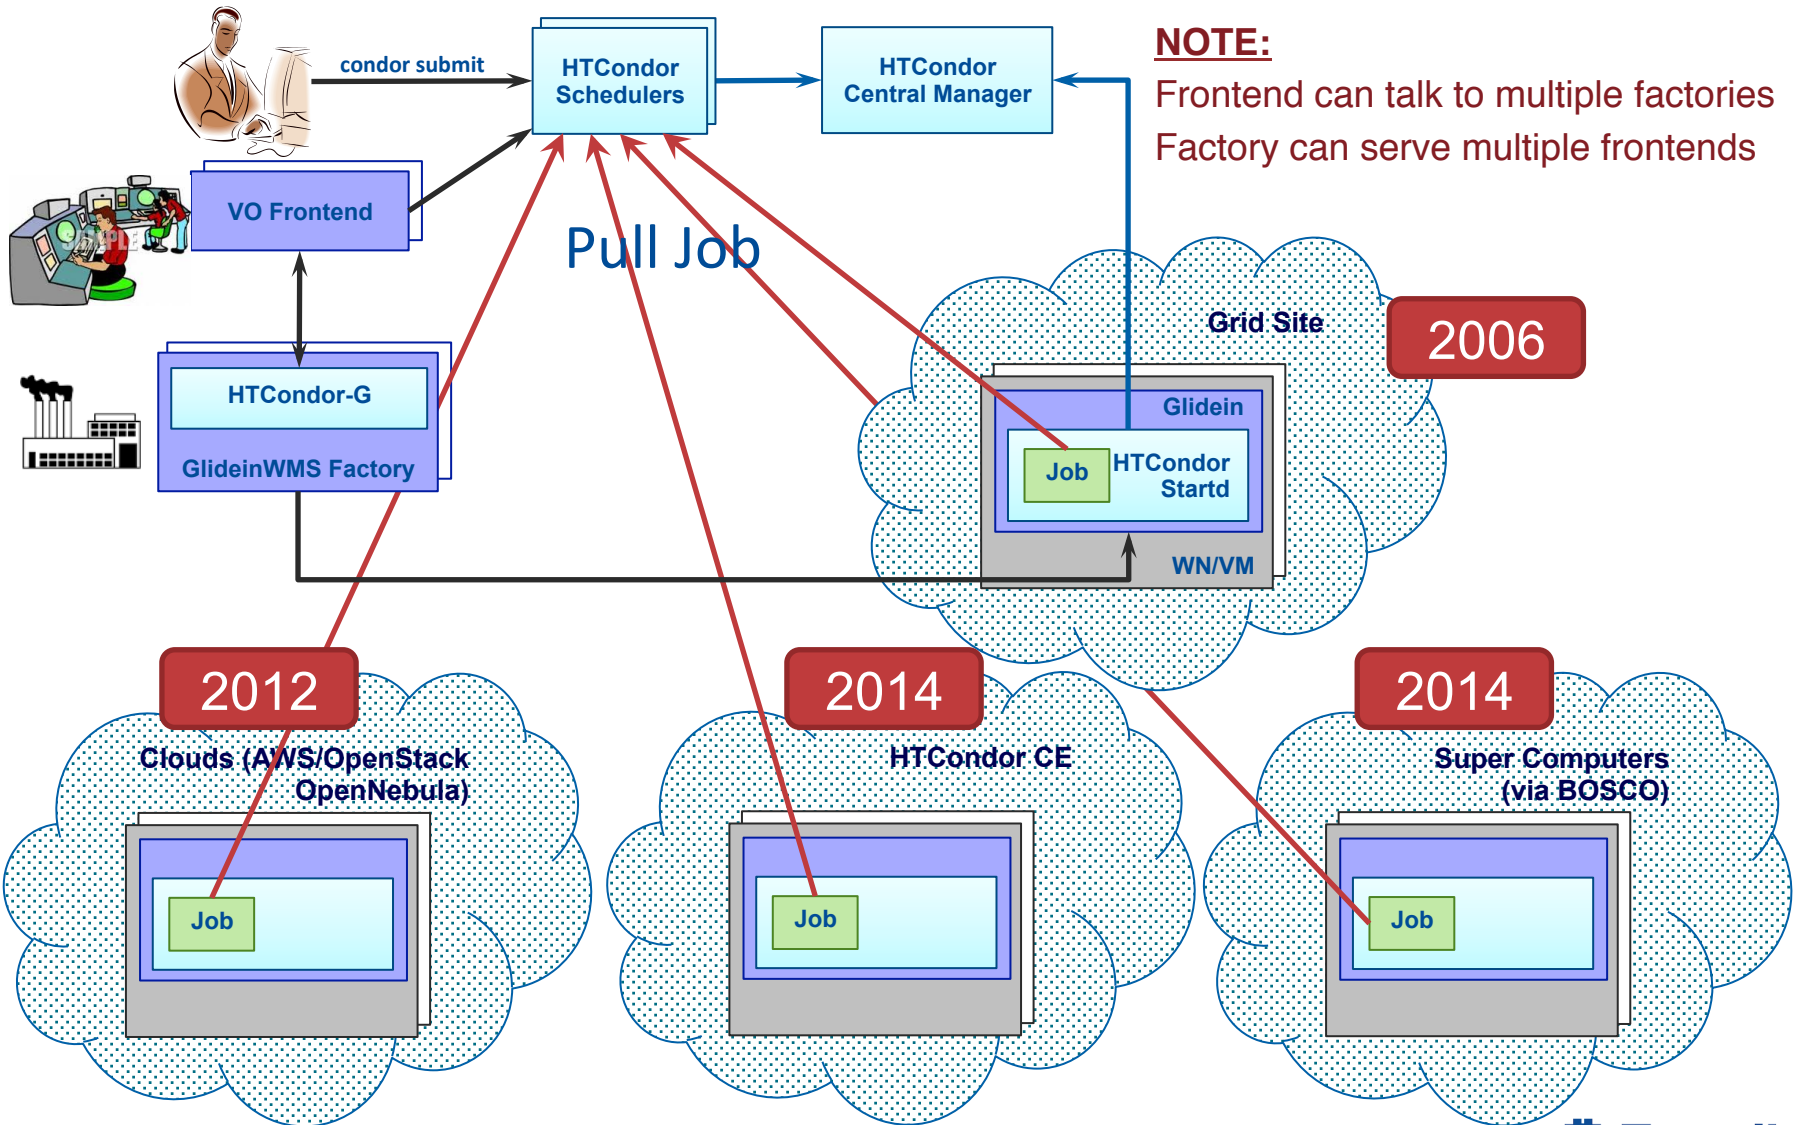
Stakeholders Meeting

May 15, 2015

Overview

- GlideinWMS Overview
- Updates since last stakeholder's meeting
- What's New?
- Stakeholder Input

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:
 - Project team reorganization in last few months
 - Project Lead transitioned from Burt Holzman → Parag Mhashilkar
 - Big thanks to Burt (Assistant head/Scientific Facilities coordinator) for leading the project for 7+ years and helping with the transition of project leadership, his continued guidance and support

Role	Resources	Effort (FTE)
Project Mgmt.	Parag Mhashilkar (0.15 USCMS)	0.15
Development & Support	Parag Mhashilkar (0.45 SCD) Marco Mambelli (0.8 SCD, 1 from June 2015) Hyunwoo Kim (0.5 SCD)	1.75 (1.95 June 2015)
Cloud Integration	Anthony Tiradani (0.2 USCMS)	0.2
TOTAL		2.1

Table: Current Resources & Roles

- Additional Code Contributions (Past year)
 - Jeff Dost (UCSD)
 - Igor Sfiligoi (UCSD)
 - Brian Bockelman (OSG/UNL)

Highlights Since Last Stakeholders Meeting

- Releases: v3.2.4 - v3.2.9
 - Total 9: Includes 3 high priority bug fix releases
 - Highlights of releases in extra slides
- Tickets/Issues Resolved
 - Features: 33
 - Bugs: 68

Milestones from last time:

Frontend Scalability

- Improvements released in GlideinWMS v3.2.4, v3.2.5, v3.2.6
 - Stakeholders: CMS, OSG
 - Frontend performs more tasks in parallel
 - Multiple HTCondor queries in parallel
 - Better utilization of multiple CPUs

- **Status: Complete**

Milestones from last time:

Better prevention of “Black Hole” workers

- Issue: <https://cdcv.s.fnal.gov/redmine/issues/6309>
- Stakeholders: OSG
- 3 Common failure nodes
 - Insufficient validation of worker nodes
 - Resolution: Add validation scripts to identify the problem
 - Worker nodes start experiencing problems after job start
 - Issue: <https://cdcv.s.fnal.gov/redmine/issues/2409>
 - Will be in GlideinWMS v3.2.10
 - Failures specific to type of user jobs
 - Solutions mostly VO specific
- **Status: In Progress**

Milestones from last time:

Factory/Frontend Configurability

- **Challenge: Preserve backward compatibility as much as possible to minimize configuration and code changes**
- Stakeholders: CMS, OSG
- Solution will be in several stages
- Frontend
 - First stage: Pluggable policy configuration
 - <https://cdcvs.fnal.gov/redmine/issues/6309>
- Factory
 - First stage: Extract & make entry configuration pluggable
 - Prototyped by Jeff Dost
 - <https://cdcvs.fnal.gov/redmine/issues/8437>
- **Status: In Progress**

Milestones from last time:

Aggregate Monitoring

- We need to pull together the monitoring across multiple factories and multiple frontends
- Stakeholders: OSG, CMS
- Proposal: A server that makes the hostnames and URLs of existing monitoring and aggregates the output
 - Dmytro Kovalskyi (UCSD/USCMS) started working on this in Q4 2014.
- **Status: Stalled (after lost resource)**
- **Other monitoring requests from OSG & CMS**
 - **Project will continue to address other monitoring improvements in upcoming GlideinWMS releases**

Milestones from last time:

“Why is my job not running”?

- <https://cdcv.s.fnal.gov/redmine/issues/4989>
- Stakeholders: OSG, CMS
- Working on a tool - functionality similar to ‘condor_q -analyze’
 - Tool is partially functional
- **Status: Stalled (No new updates)**

New Milestones Achieved

- New high impact milestones/requests to the project after the stakeholders meeting
- Support additional resource types (CMS/OSG/FIFE)
 - HTCondor CE: v3.2.4
 - Allocations on Leadership class machines via BOSCO: v3.2.6
- Support 200K+ jobs scale (CMS)
 - Support shared ports for HTCondor daemons: v3.2.7
 - Support CCB configuration separate from User collector: v3.2.9
- Native support for fail over for VO Frontend (CMS/FIFE)
 - Support Frontends running on HA (master-slave) mode: v3.2.9
- Better support for Multi Core glideins (CMS/OSG)
 - Several features/bug fixes between v3.2.4 - v3.2.9

New Milestones Achieved

- **Simplify operations (OSG/CMS)**
 - Changes between v3.2.4 - v3.2.9
 - Several monitoring enhancements
 - New tools to aid in operations
 - External contributions
 - Thanks to external contributors!

Support Structure

- Support Mailing list: glideinwms-support@fnal.gov
- Issues tracked in redmine issue tracker
 - <https://cdcv.s.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
 - Entire development team is responsible for support
- Project Management
 - Project Status reported monthly at CS Project status meetings!
 - At the request of computing management
 - Project management absorbed into the project effort

Tracking Stakeholder Requests 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 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

#	Tracker	Target version	Status	Priority	Subject	Assignee	Updated	Labels
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 Mhashikar	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 Mhashikar	04/13/2015 03:21 AM	OSG
8437	Feature	v3_2_10	New	Normal	Split entry configuration from main factory config	Parag Mhashikar	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-g-analyze	Burt	02/03/2015 02:44 PM	OSG

What's Brewing?

- Production Series (v3.2.x)
 - Series to mostly focus on
 - High impact bug fixes
 - High impact features that do not break backward compatibility
 - Improve support for Multi Core glideins
 - Monitoring enhancements
 - Support entries O(600+)
 - Next release v3.2.10
 - Initial Roadmap:
https://cdcvms.fnal.gov/redmine/projects/glideinwms/issues?query_id=53
 - Tentative Release: End of July 2015

What's Brewing?

- Development Series (v3.3.x)
 - Usually production quality
 - But new features maybe in unpolished state
 - We try to maintain backward compatibility
 - **Disclaimer:** May break backward compatibility for some features
 - Primary Focus (One Facility/CMS/OSG)
 - Support different EC2 features in GlideinWMS Support manageable solution for complex VO provisioning policies
 - Factory/Frontend Configurability
 - Solution in multiple stages
 - Refer to one of the previous slide
 - Will be declared production after polishing and hardening

What's Brewing?

- v3.3 in the planning stage
 - Initial Roadmap: https://cdcvs.fnal.gov/redmine/projects/glideinwms/issues?query_id=26
 - Timeline: Tentatively 2-3 months
 - Focus
 - Support different EC2 features in GlideinWMS
 - Spot pricing
 - Regions
 - Availability Zones
 - Support manageable solution for complex VO provisioning policies
 - <https://cdcvs.fnal.gov/redmine/issues/6309>
 - Extract policies from the VO Frontend configuration
 - Make policies pluggable

Future Releases: Stakeholder Impact & Feedback

- GlideinWMS on SL/RHEL 5 platform
 - No support from next production release - v3.2.10
 - **Question: Will this be an issue for any stakeholder?**
 - No support from next development release - v3.3
- Support for GlideinWMS v2.7 series
 - v2.7.2 released on Sep 10, 2013
 - Available through OSG 3.1 only and in maintenance mode
 - We would like to discontinue support for GlideinWMS v2.7 series
 - **Questions:**
 - **Any stakeholder still using them?**
 - **Do you need project's assistance in migrating to v3.2.9+?**
- Support for SL7/RHEL 7 platform
 - **Question: Any stakeholder looking at SL7/RHEL7 seriously?**

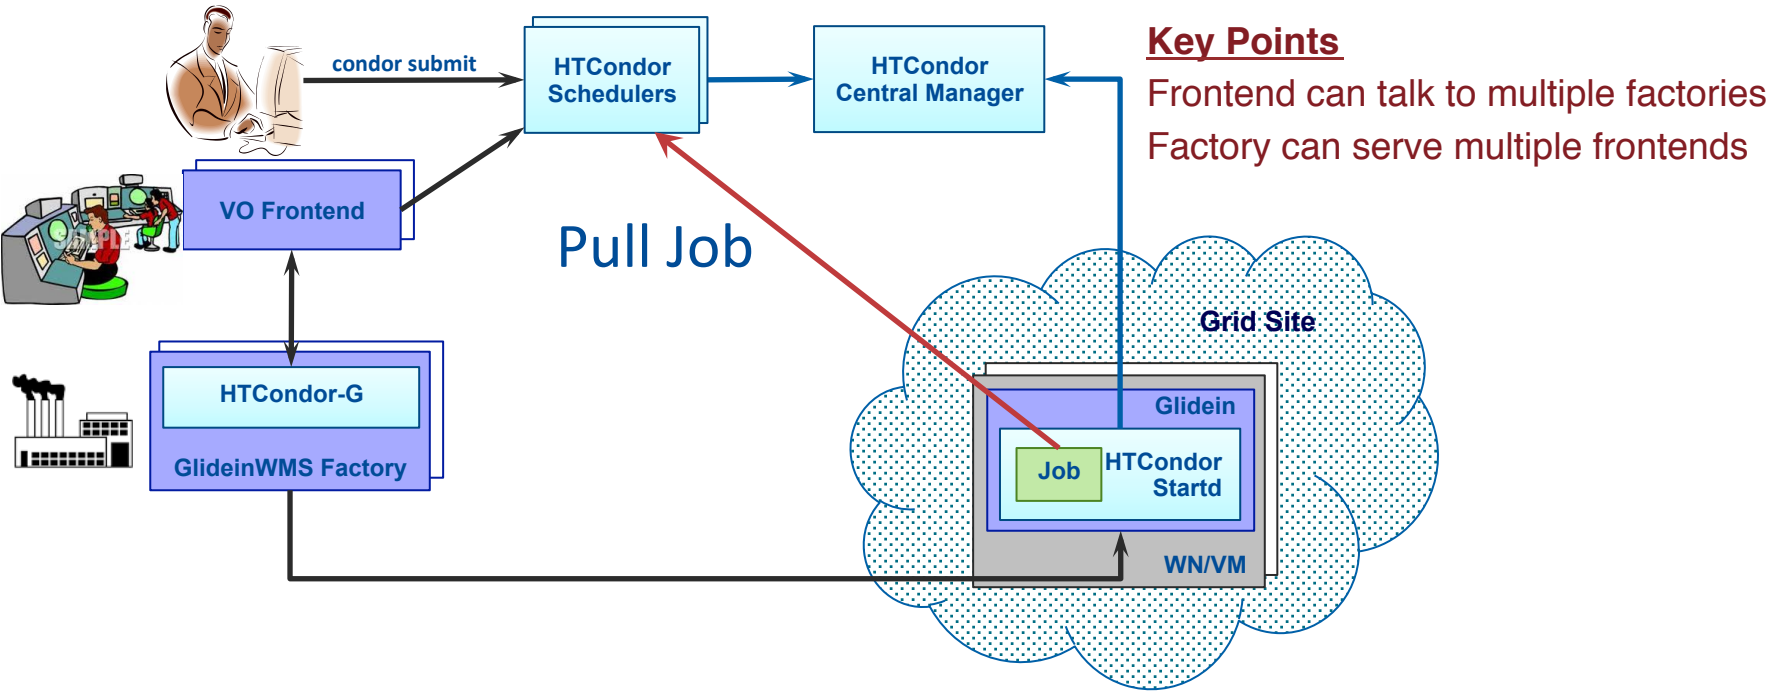
Stakeholders Input

EXTRA SLIDES

Why GlideinWMS?

- Computing resources in various infrastructure types
 - Scientific computing is not restricted on-site batch clusters
 - Grid Computing is over decade and half old
 - Computing Clouds will play major roles in future
 - Big players like CMS & FIFE are actively pursuing clouds as a solution to support peak demands
 - Additional computing power is also available in the form of allocations on the Leadership class machines/Super computers like Gordon @ UCSD
- We need a Workload Management System, (specifically GlideinWMS) to -
 - Provide manageable solution to handle diversity in the resource types
 - Provision resources on different infrastructure types based on policies
 - Provide uniform access to the computing resources
 - Shield the user jobs from heterogeneity of resources and failures on the worker nodes and provide them a batch system (HTCondor) friendly interface
 - [...]

GlideinWMS



Key Points
Frontend can talk to multiple factories
Factory can serve multiple frontends

GlideinWMS Releases - Key Features

v3_2_4 - April 14, 2014

- Added support for HTCondor-CE
- Several performance improvements and better CPU utilization by frontend
- Added global total idle limits and limits on total idle glideins requested by frontend
- Bug Fix: Fixed partitioning of multi-core glideins

v3_2_5 - May 19, 2014

- Frontend scalability improvements
- Added administrative commands for frontend `fetch_glidein_log`, `glidein_off` and `enter_frontend_env`

v3_2_5_1 - June 23, 2014

- Bug Fix: Fixed an issue with the `factory_startup` template

v3_2_6 - July 28, 2014

- Added support for BOSCO
- Added new tool to purge old glideins
- Monitoring improvement
- Allow for separation of Factory collector and CondorG collector

GlideinWMS Releases - Key Features

v3_2_7 - October 14, 2014

- Glideins now support shared port
- Bug Fix: Frontend now correctly provisions multicore glideins if the GLIDEIN_CPUS is configured for the entry

v3_2_7_1 - November 05, 2014

- Set USE_SHARED_PORT to get around the issue with HTCondor 8.2.3

v3_2_7_2 - November 06, 2014

- Bug Fix: Set MASTER.USE_SHARED_PORT instead of USE_SHARED_PORT to avoid secondary collectors using the shared port daemon

v3_2_8 - December 30, 2014

- Support HTCondor GANGLIAD monitoring
- Added failed glidein/idle/running/total core statistics to frontend monitoring
- Bug Fix: Glideins do not mail admins when HTCondor daemon crash
- Bug Fix: Fixed FD leaks in factory

v3_2_9 - May 08, 2015

- VO Frontend supports a master-slave HA mode
- Frontend supports configuring CCBs separate from the User Collector. This lets the infrastructure to scale more than 150K jobs in user pool
- Updated the dependency of GlideinWMS to HTCondor v8.2.2