Fermilab | U.S. DEPARTMENT OF ENERGY | Office of Science

GlideinWMS

Marco Mambelli

Summer Student Guidelines

May 2020

Outline

- Scientific computing
- GlideinWMS
- HTCondor
- Resources
- Monitoring
- Links
- Demo

- **Scientific computing**
- GlideinWMS
- HTCondor
- Resources
- Monitoring
- Links
- Demo

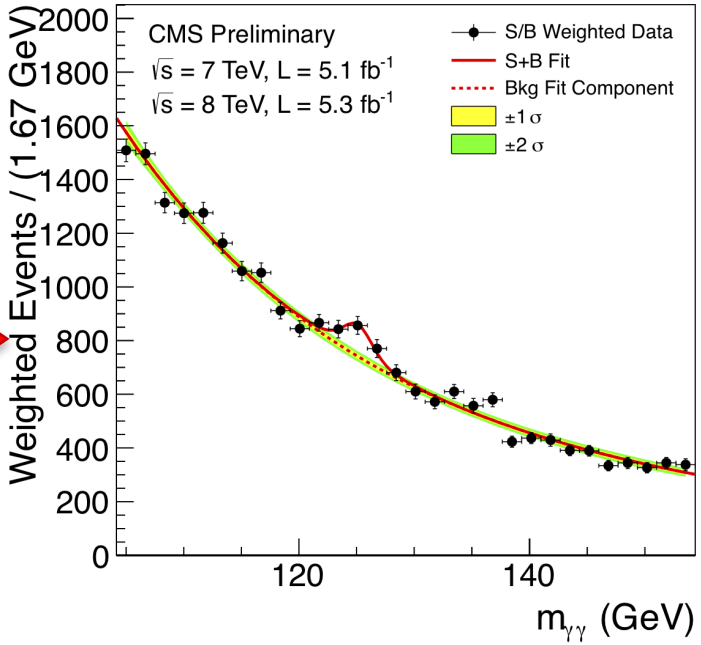
CMS Preliminary
√s = 7 TeV, L = 5.1 fb-1
√s = 8 TeV, L = 5.3 fb-1
S/B Weighted Data
S+B Fit
Bkg Fit Component
±1 σ
±2 σ
Weighted Events / (1.67 GeV)
2000
1800
1600
1400
1200
1000
800
600
400
200
0
120
140
mγγ (GeV)

Fermilab

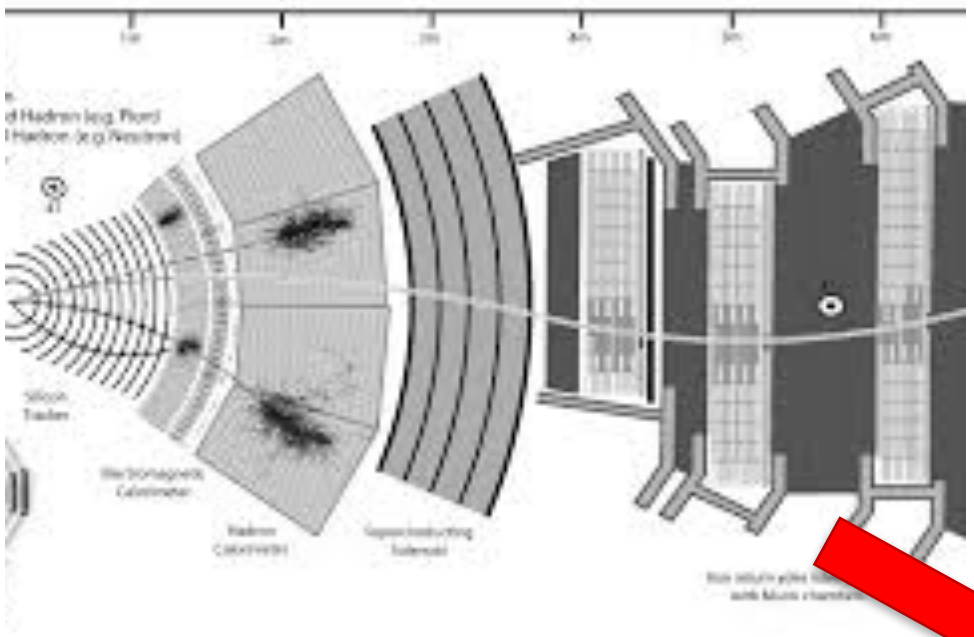

HEP experiments require computing!

- Accelerator and detector simulations
- Data reconstruction
- Data analysis

When one computer is not enough

- ✓ Supercomputer
- ✓ Cluster (Batch Systems)
- ✓ Grid
- ✓ Cloud

Computing resources

- Supercomputer
  - Special purpose computer fine tuned to achieve elevated number of operations per second
- Cluster (batch system)
  - Collection of parallel or distributed computers which are interconnected among themselves using high-speed networks
  - Local Resource Manager (or batch system) is the software managing the computers in the cluster (e.g. PBS, SLURM, HTCondor, SGE, LSF, ...)
- Grid (e.g. Open Science Grid)
  - Combines computers from multiple administrative domains to reach common goals, to solve tasks
  - System that coordinates resources which are not subject to centralized control, using standard, open, general-purpose protocols and interfaces to deliver nontrivial qualities of service

Computing resources (cont)

- Cloud
  - Refers to both the applications delivered as services over the Internet and the hardware and system software in the data centers that provide those services
  - aka Elastic computing, available or paid only when used
  - Software as a Service (SaaS) is a kind of services where in many users can make use of the software hosted by the service provider and pay only for time its being used (Workday, Slack)
  - Platform as a Service (PaaS) provides a high-level integrated environment to design, build, test, deploy and update online custom applications (Amazon orchestration, Google AE)
  - Infrastructure as a Service (IaaS) refers to the services provided to the users to use processing power, storage, network and other computing resources, to run any software including operating systems and applications (AWS, Google CE, Fermicloud)

Basic concepts – Service Models

Services provided to the users to use processing power, storage, network and other computing resources, to run any software including operating systems and applications (AWS, Google CE)

HTC problem: Growing needs and Trends

- High Throughput Computing
  - use of many computing resources over long periods of time to accomplish a computational task
- Need for more resources
  - Scale to more jobs
  - Access more resources
  - Simplify the management
- Less structured resources and infrastructure
  - Multiple organizations
  - Different systems
  - Less infrastructure
  - Different authentications

- Scientific computing
- **GlideinWMS**
- HTCondor
- Resources
- Monitoring
- Links
- Demo

GlideinWMS

GlideinWMS is a pilot based resource provisioning tool for distributed High Throughput Computing

- Provides reliable and uniform virtual clusters
- Submits Glideins to unreliable heterogeneous resources
- Leverages HTCondor
  - Provides HTCondor pools
  - Uses HTCondor capabilities

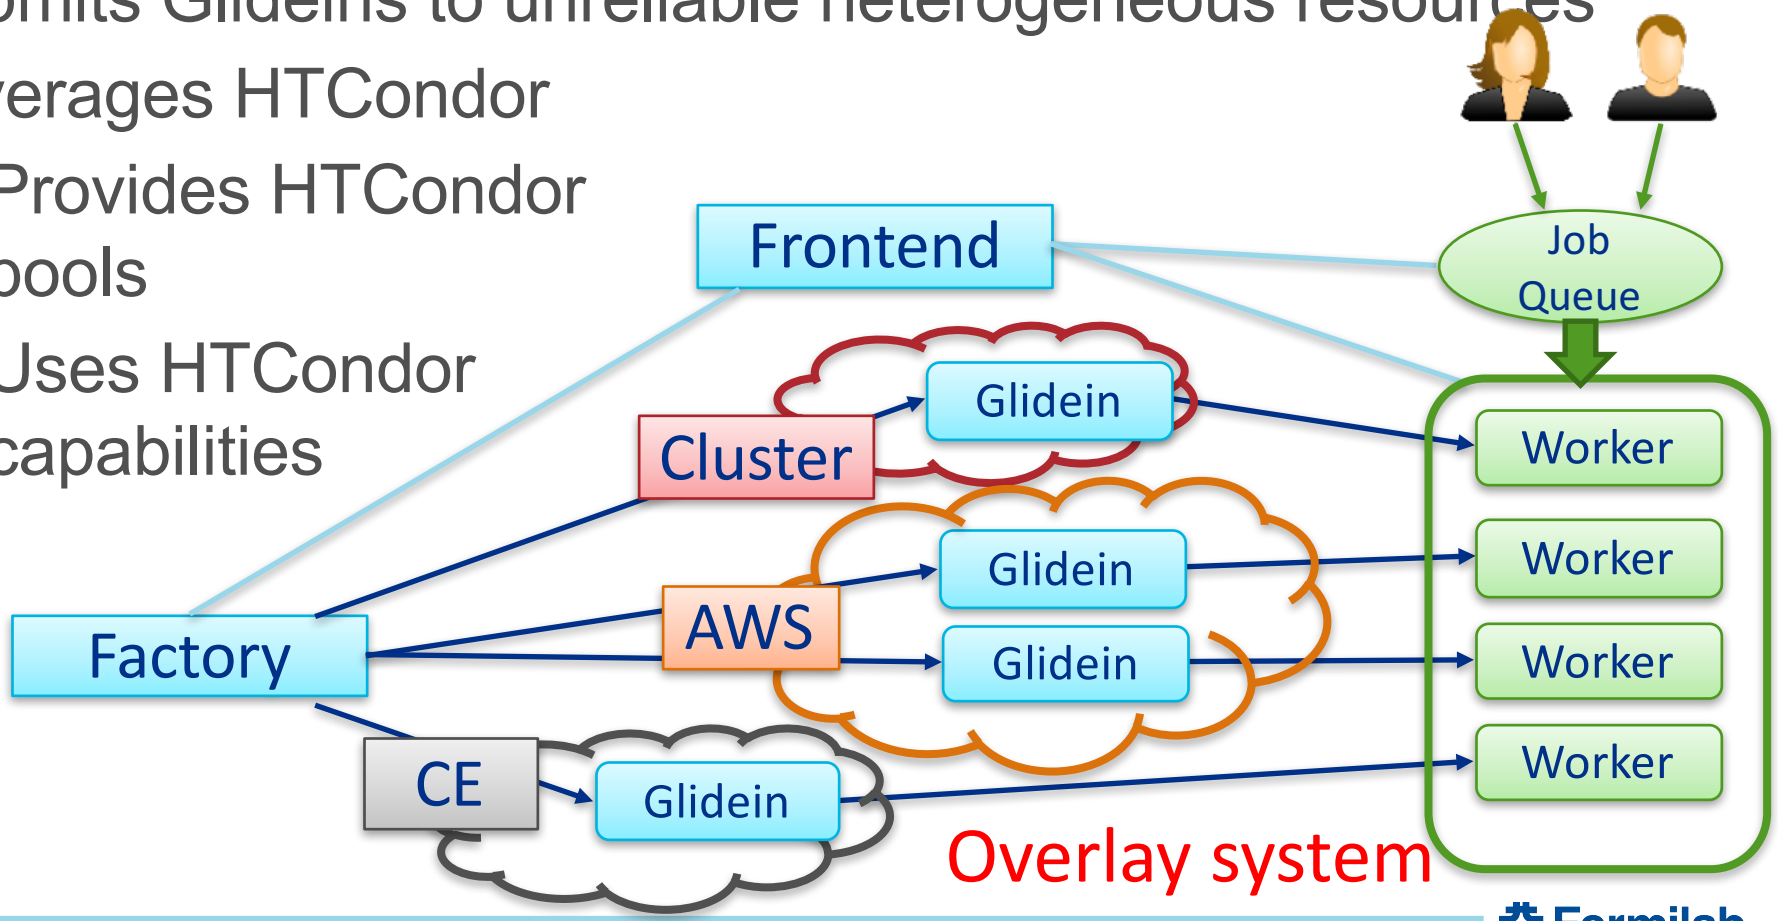

Fermilab

Glidein: node testing and customization

- Scouts for resources and validates the Worker node
  - Cores, memory, disk, GPU, …
  - OS, software installed
  - CVMFS
  - VO specific tests
- Customizes the Worker node
  - Environment, GPU libraries, …
  - Starting containers (Singularity, …)
  - VO specific setup
- Provides a reliable and customized execute node to HTCondor

Factory

- A Glidein Factory knows how to submit to sites
  - Sites are described in a local configuration
  - Only trusted and tested sites are included
- Each site entry in the configuration contains
  - Contact info (hostname, resource type, queue name)
  - Site configuration (startup dir, OS type, …)
  - VOs authorized/supported
  - Other attributes (Site name, core count, max memory, ...)
  - Glideins can also auto-detect resources
- Configuration can be auto-generated (e.g. from CRIC), admin curated, stored in VCS (e.g. GitHub)
- Condor does the heavy lifting of submissions.

Factory: Supported resources

- Remote or local clusters:
  - Can have batch systems other than HTCondor: PBS, SGE, Slurm, all supported.
- Grid sites (CREAM, ARC, HTCondor-CE)
- Hosted CEs
- Commercial cloud (AWS, Google)
- Open Source Cloud (OpenStack, OpenNebula)
- HPC sites
  - Uses an ssh-based system to ssh into HPC sites and submit directly from their login nodes.

Frontend

- Monitors jobs to see how many Glideins are needed
- Compares what entries (sites) are available
- Requests Glideins from the Factory
- Requests Factory to kill Glideins if there are too many
- Pressure-based system
  - Works keeping a certain number of Glideins running or idle at the sites
  - Glideins requests are gradual to avoid spikes and overloads
- Manages credentials and delegates them to the Factory.

GlideinWMS components

FRONTEND	FACTORY
- Controlled by VO Operators - Main task is to look for user jobs and ask the Glidein Factories to provide glideins, if needed - Decides which glidein Factory should submit the pilot Jobs and how many of them - Configuration xml file: /etc/gwms-frontend/frontend.xml	- Controlled by Factory operators (OSG) - Advertises itself, listen for requests from Frontend clients and submit glideins - Can handle multiple kinds of glidein - Configuration xml file: /etc/gwms-factory/glideinwms.xml

GLIDEIN
- Requested by the Frontend, launched by the Factory and joins the virtual cluster - Property configured execution node submitted as a Grid job. - Defines how multi-core glideins should split their resources - Pilot jobs may run multiple user Jobs - glidein_startup.sh configures and starts the condor startd daemon

Distributed

- N-to-M relationship
  - Each Frontend can talk to many Factories
  - Each Factory may serve many Frontends
- Multiple User Pools
- High Availability replicas

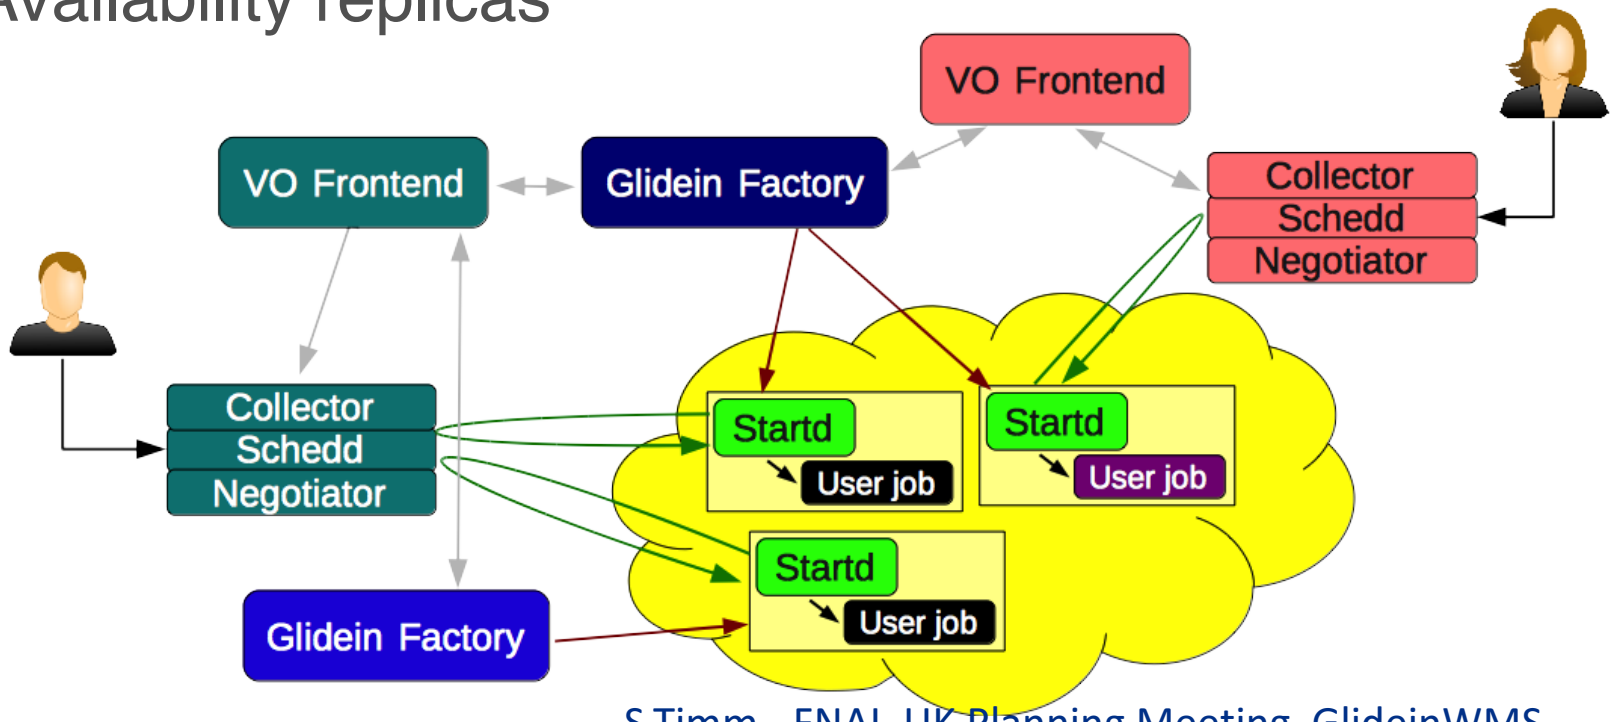

S.Timm - FNAL-UK Planning Meeting GlideinWMS

Frontend configuration example

```
<frontend downtimes_file="frontenddowntime" advertise_delay="5" frontend_name="vofrontend-v2_4" loop_delay="60">
  <log_retention >
    <process_logs >
      <process_log extension="info" max_days="7.0" max_mbytes="100.0" min_days="3.0" msg_types="INFO" backup_count="5" compression="gz" />
      <process_log extension="debug" max_days="7.0" max_mbytes="100.0" min_days="3.0" msg_types="DEBUG,ERR,WARN" backup_count="5" />
    </process_logs >
  </log_retention >
  <match match_expr="True" start_expr="True" policy_file="/path/to/python-policy-file">
    <factory query_expr="True">
      <match_attrs />
      <collectors>
        <collector DN="/DC=org/DC=doegrids/OU=Services/CN=factory-server.fnal.gov" comment="" factory_identity="factoryuser@factory-
        server.fnal.gov" my_identity="frontenduser@frontend-server.fnal.gov" node="factory-server.fnal.gov:8618" />
      </collectors>
    </factory>
    <job comment="" query_expr="(JobUniverse==5)&&(GLIDEIN_Is_Monitor =!= TRUE)&&(JOB_Is_Monitor =!= TRUE)">
      <match_attrs />
      <schedds>
        <schedd DN="/DC=org/DC=doegrids/OU=Services/CN=userpool.fnal.gov" fullname="userpool.fnal.gov" />
        <schedd DN="/DC=org/DC=doegrids/OU=Services/CN=userpool.fnal.gov" fullname="schedd_jobs1@userpool.fnal.gov" />
        <schedd DN="/DC=org/DC=doegrids/OU=Services/CN=userpool.fnal.gov" fullname="schedd_jobs2@userpool.fnal.gov" />
      </schedds>
    </job>
  </match>

  <monitor base_dir="/var/www/html/vofrontend/monitor" flot_dir="/opt/javascriptrrd-0.6.3/flot" javascriptRRD_dir="/opt/javascriptrrd-0.6.3/src/lib"
  jquery_dir="/opt/javascriptrrd-0.6.3/flot" />
  <monitor_footer display_txt="Legal Disclaimer" href_link="/site/disclaimer.html" />

  <security classad_proxy="/etc/grid-security/vocert.pem" proxy_DN="/DC=org/DC=doegrids/OU=Services/CN=frontend-server.fnal.gov"
  proxy_selection_plugin="ProxyAll" security_name="frontenduser" sym_key="aes_256_cbc">
    <credentials>
      <credential absfname="/tmp/x509up_u" security_class="frontend" trust_domain="OSG" type="grid_proxy" vm_id="123" vm_type="type1"
      pool_idx_len="5" pool_idx_list="2,4-6,10" />
    </credentials>
    </security>
  <stage base_dir="/var/www/html/vofrontend/stage" use_symlink="True" web_base_url="http://frontend-server.fnal.gov:9000/vofrontend/stage" />
  <work base_dir="/opt/vofrontend" base_log_dir="/opt/vofrontend/logs" />
  <attrs>
    <attr name="GLIDECLIENT_Rank" glidein_publish="False" job_publish="False " parameter="True" type="string" value="1" />
    <attr name="GLIDECLIENT_Start" glidein_publish="False" job_publish="False" parameter="True" type="string" value="True" />
    <attr name="GLIDEIN_Expose_Grid_Env" glidein_publish="True" job_publish="True" parameter="False" type="string" value="True" />
    <attr name="GLIDEIN_Glexec_Use" glidein_publish="True" job_publish="True" parameter="False" type="string" value="OPTIONAL" />
    <attr name="USE_MATCH_AUTH" glidein_publish="False" job_publish="False" parameter="True" type="string" value="True" />
  </attrs>
  <groups>
    <group name="main" enabled="True">
      <config>
```

WMS Collector

Available schedds

Security credentials

Factory configuration example - Entry

Entry configuration

Limits (in glideins)

```
<entry name="EXAMPLE_ENTRY" enabled="True" auth_method="grid_proxy" trust_domain="OSG"
gatekeeper="gatewayname.fnal.gov/jobmanager-condor" gridtype="gt2" rsl="(queue=default)
(jobtype=single)" schedd_name="wmscollector.fnal.gov" verbosity="std" work_dir="OSG">
   <config>
      <max_jobs>
         <per_entry held="1000" idle="2000" glideins="10000"/>
         <default_per_frontend held="50" idle="100" glideins="5000"/>
         <per_frontends>
            <per_frontend name="FRONTEND:SECURITY_CLASS" held="50" idle="100" glideins="5000"/>
         </per_frontends>
      </max_jobs>
      <release max_per_cycle="20" sleep="0.2"/>
      <remove max_per_cycle="5" sleep="0.2"/>
      <submit cluster_size="10" max_per_cycle="100" sleep="0.2" slots_layout="partitionable">
         <submit_attrs>
            <submit_attr name="RequestMemory" value="2000">
         <submit_attrs/>
      </submit>
   </config>
   <allow_frontends />
   <attrs>
      <attr name="CONDOR_ARCH" const="True" glidein_publish="False" job_publish="False"
      parameter="True" publish="False" type="string" value="default"/>
      <attr name="CONDOR_OS" const="True" glidein_publish="False" job_publish="False"
      parameter="True" publish="False" type="string" value="default" />
      <attr name="GLEXEC_BIN" const="True" glidein_publish="False" job_publish="False"
      parameter="True" publish="True" type="string" value="NONE"/>
      <attr name="GLIDEIN_Site" const="True" glidein_publish="True" job_publish="True"
      parameter="True" publish="True" type="string" value="FNAL EXAMPLE SITE"/>
      <attr name="GLIDEIN_CPUS" const="True" glidein_publish="True" job_publish="True"
      parameter="True" publish="True" type="string" value="1"/>
      <attr />
   </attrs>
   <monitorgroups/>
   <files/>
   <infosys_refs />
</entry>
```

Partitionable

Slot description (assuming/imposing 1 core)

GlideinWMS

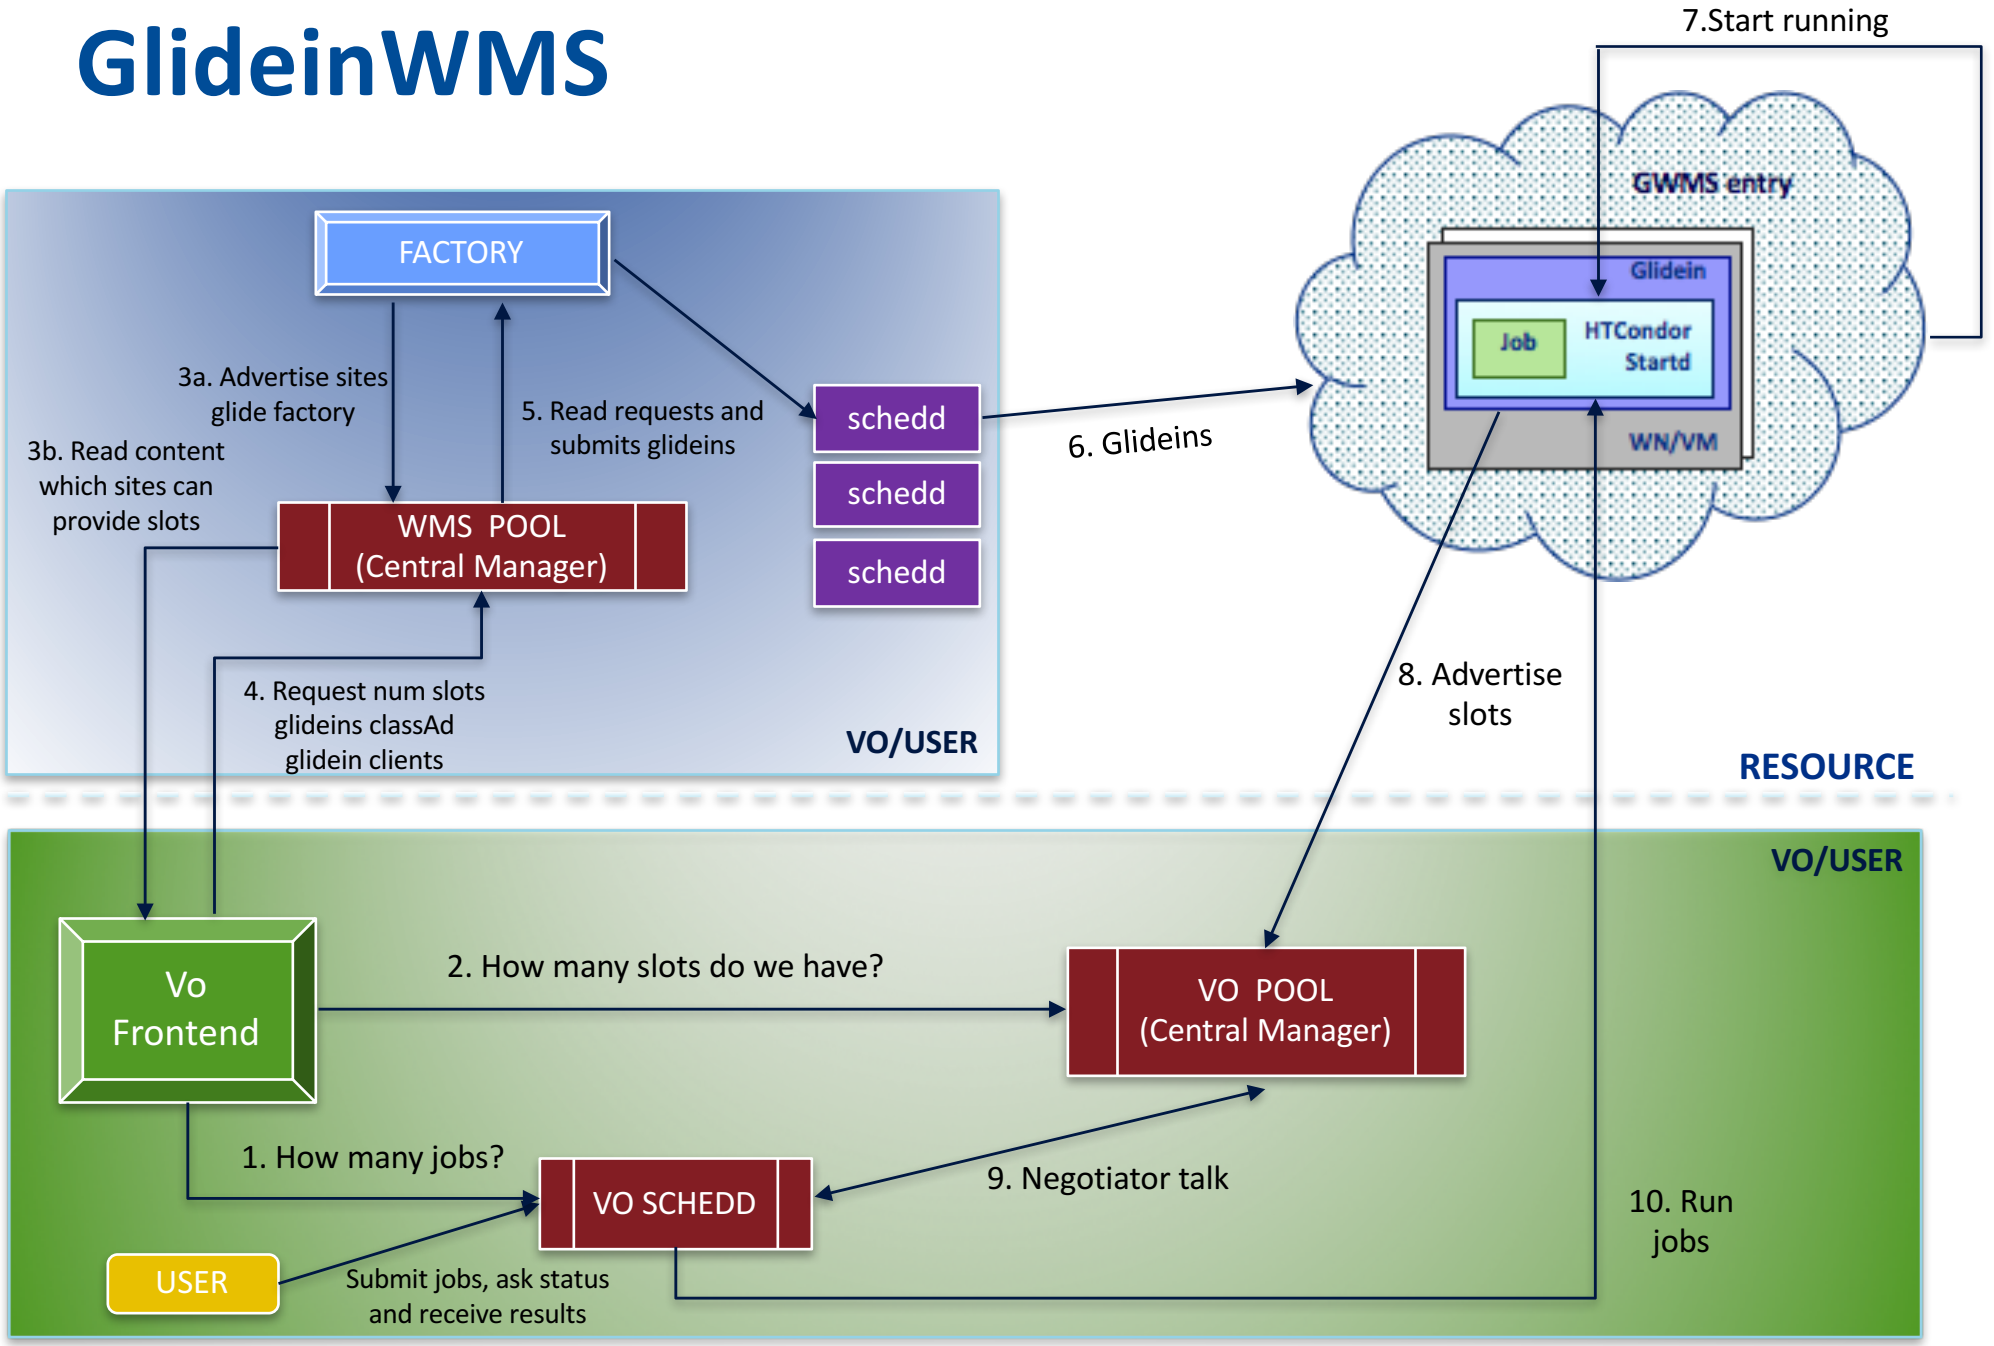

Fermilab

Major GlideinWMS Deployments

- Beta version was called “GlideCAF” in CDF
  - Began testing in 2005
- CMS Global Pool—regularly 200000+ cores
  - Redundant master nodes at CERN and Fermilab
  - Combines production and analysis jobs
- FIFEBATCH / FermiGrid
  - Integrates 18000 on-site cores of FermiGrid with up to 12000 offsite cores.
  - DUNE is using it for standard production and analysis
  - Demonstrated a pool with 2.01 million cores (NOVA 2018)
- Open Science Grid
  - Multi-VO structure shares the same Factory at UCSD
- HEPCloud
  - Portal integrating multiple resources

How it is used?

- Can be used directly
  - HTCondor
- Integrates well in hybrid systems
  - OSG-Connect
  - FermiGrid
- Used by workload/workflow managers:
  - JobSub
  - ProdAgent, CRAB
  - POMS
  - Pegasus
  - HEPCloud

- Scientific computing
- GlideinWMS
- **HTCondor**
- Resources
- Monitoring
- Links
- Demo

HTCondor

- HTCondor is a Workload Management System
  - i.e.: batch system or Local Resource Manager

- Open-source batch system implementation
  - Fault tolerant
  - Robust feature set
  - Flexible
  - Local Center for High Throughput Computing (UW Madison)

HTCondor ClassAds

- HTCondor principles: two parts of the equation
  - Jobs: quanta of work
  - Machines: available resources

- ClassAds is a language for objects (jobs and machines) to
  - Express attributes about themselves
  - Express what they require/desire in a match (similar to personal classified ads)
  - Structure
    - Set of attribute name/value pairs
    - Value : Literals (string, bool, int, float or an expression)

Example Match

Pet Ad

MyType = “Pet”

TargetType = “Buyer”

Requirements =

DogLover =?= True

Rank = 0

PetType = “Dog”

Color = “Brown”

Price = 75

Breed = "Saint Bernard"

Size = "Very Large"

...

Buyer Ad

MyType = “Buyer”

TargetType = “Pet”

Requirements =

(PetType == “Dog”) &&

(TARGET.Price <= MY.AcctBalance)

&&

(Size == "Large"||Size == "Very Large")

Rank = (Breed == "Saint Bernard")

AcctBalance = 100

DogLover = True

. . .

Dog == Resource ~= Machine

Buyer ~= Job

HTCondor Workflow

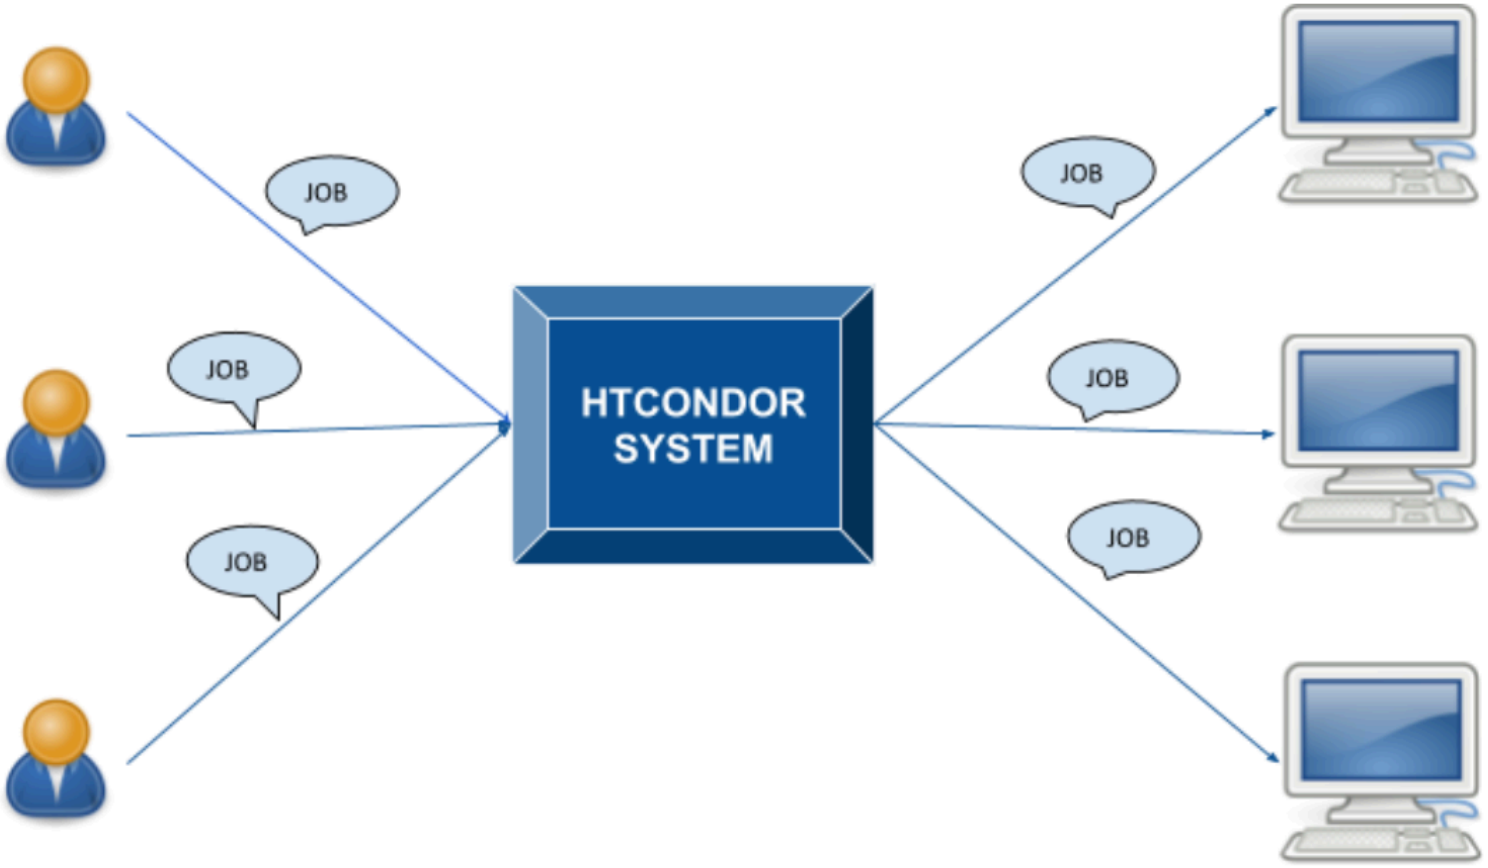

HTCondor components

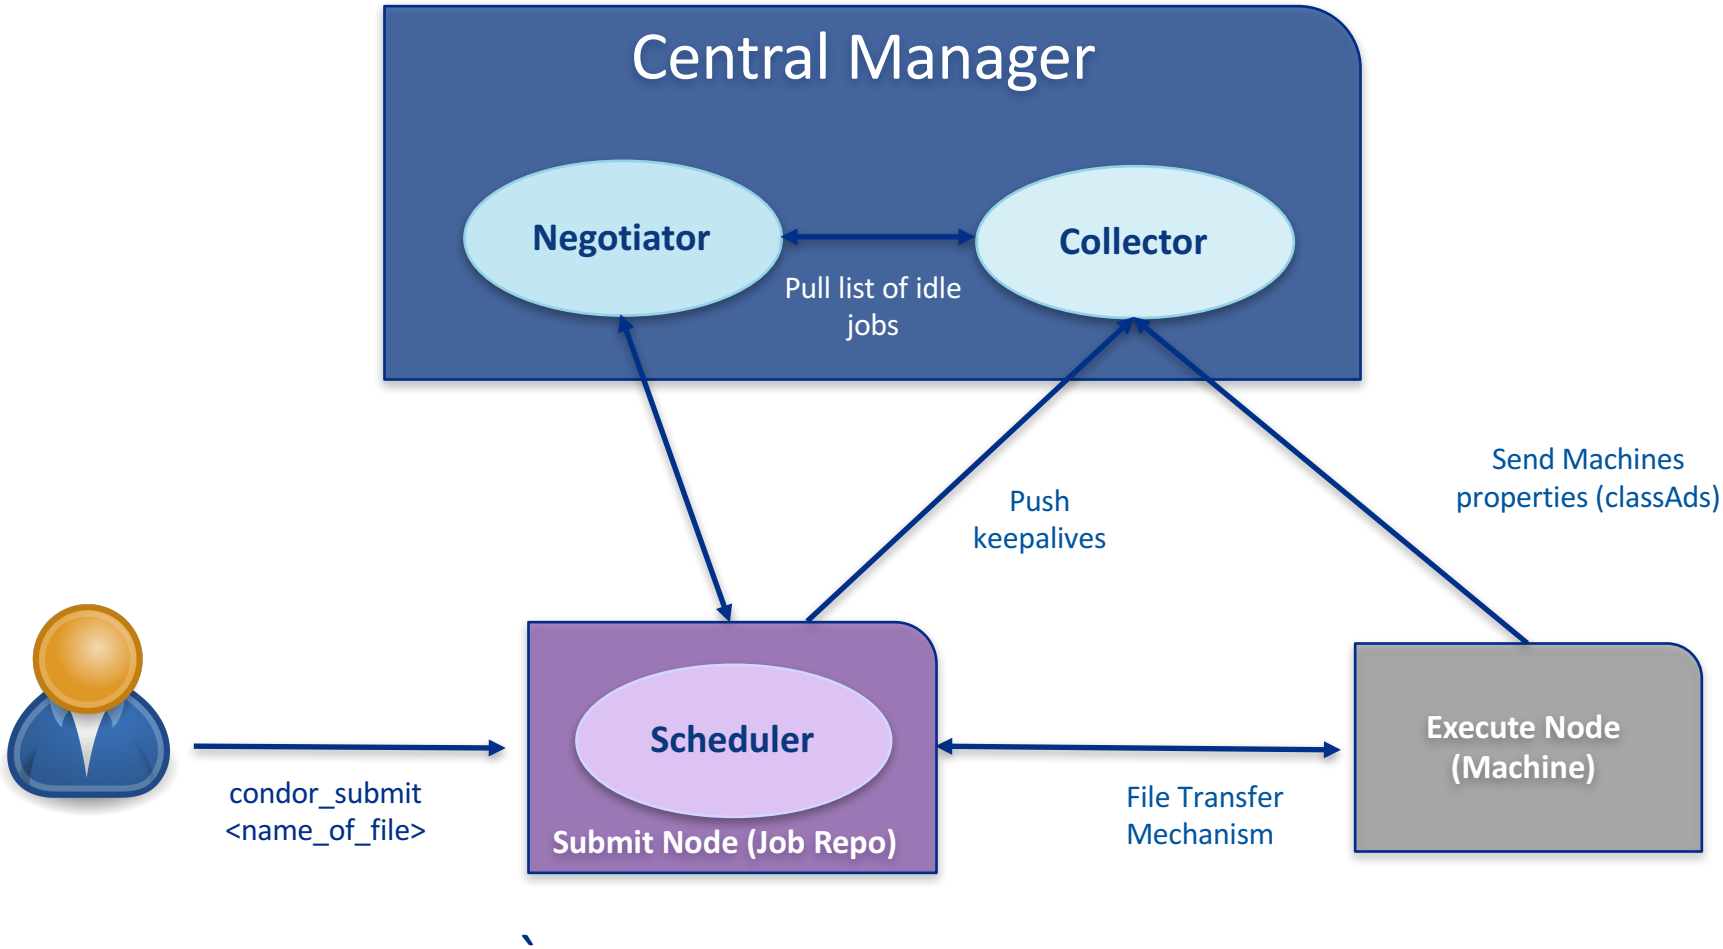

HTCondor components (daemons)

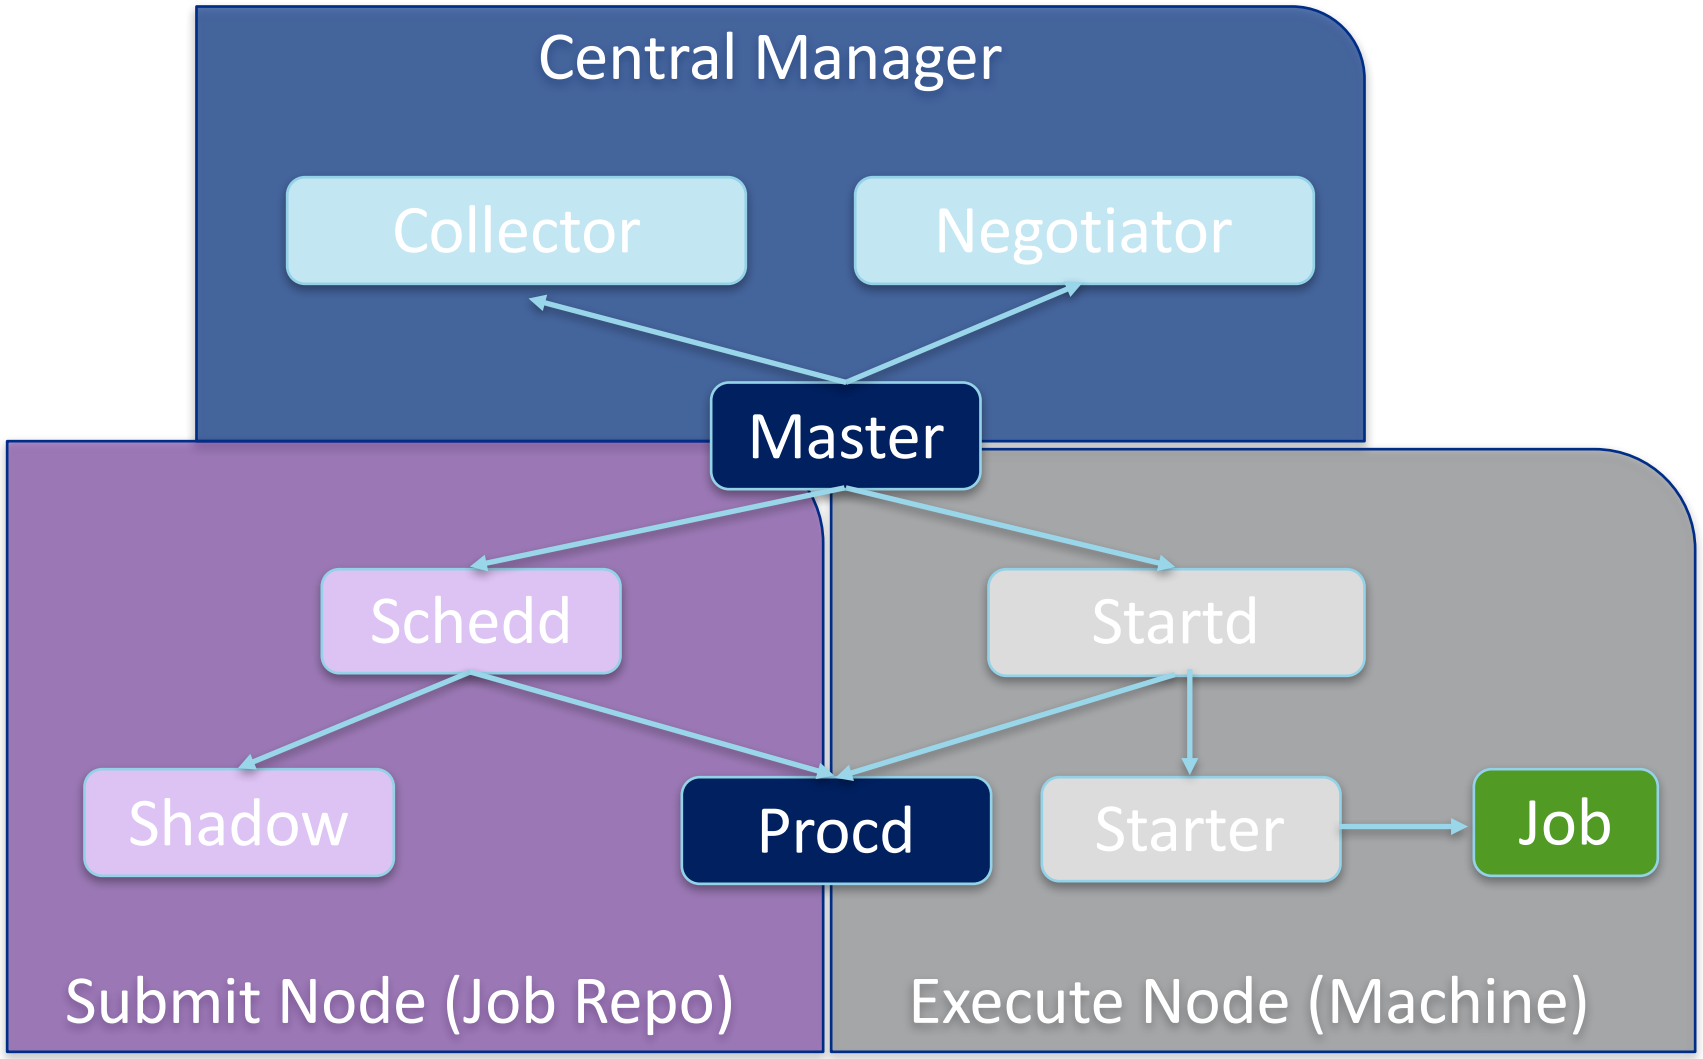

- Scientific computing
- GlideinWMS
- HTCondor
- **Resources**
- Monitoring
- Links
- Demo

Glideins run on Execute Nodes

- This is a machine (worker node, host, node, resource), managed by a (Local) Resource Manager
- More frequently virtual than not
- Characterized by its resources (dimensions):
  - CPUs (or total number of cores)
  - RAM (memory)
  - Disk
- There can be other special resources that the node provides: GPUs, access to devices, software, …
- The Glidein will receive all the node or part of it
- Sometime is not easy to identify everything used by a job

Simple (confusing) scenario

Historically 1 job was running on 1 glidein on 1 worker node using 1 CPU with 1 core.

- Terms got mixed up
- Systems were handling these interchangeably

This is no more

- Systems are more flexible
- We need to know what we want to count or control

Units of work and resources

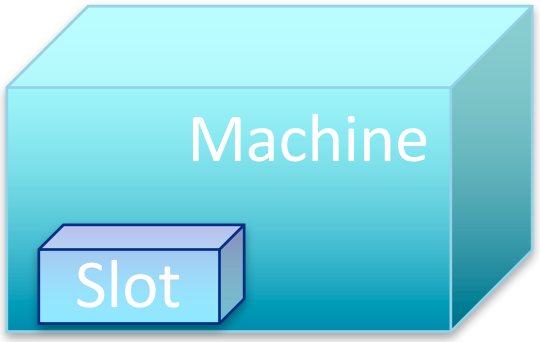

Terms used by HTCondor

- Job
- Machine (Startd)
- Slot (vm, Starter): multidimensional partition of a machine
  - Static
  - Partitionable
  - Dynamic

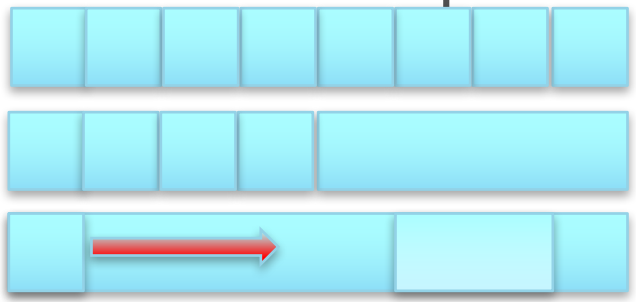

Glidein

- Pilot sent on a Machine (or more)
- Allows partitioning policies
- Job for the Factory

Job and Machine ‘dimensions’

- Job request
  - request_cpus: number of cores, integer, default 1.
  - request_disk: amount of disk space in Kbytes, default to sum of sizes of the job's executable and all input files (or image size)
  - request_memory: amount of memory space in Mbytes, default to executable size
- Machine
  - Cpus: number of cores, integer, by default the available cores
  - Disk: amount of disk space on this machine available for the job in KiB, by default the available space
  - Memory: amount of RAM in MiB in this slot

- Over and Under provision are possible

Partitioning in an overlay system

- Dimensions: Cores, Memory, Disk, Lifetime
- The resource (e.g. GPGrid) partitions its Execute nodes
- GlideinWMS further partitions the resources it receives

- E.g. 64 Cores machine split in 16 or 32 cores cluster slots; 16 or 12 cores Glideins in 4 or 2 cores partitionable slots; 2 or 1 core jobs

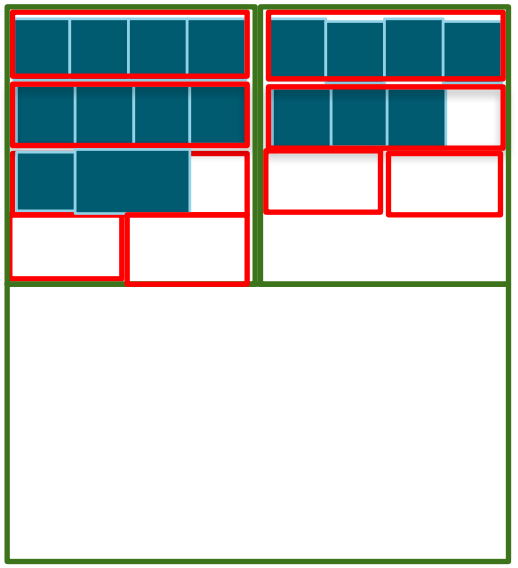

- Issues
  - Fragmentation (unused)
  - Flexibility (vs Complexity)
  - Under or over provisioning (overbooking or be prudent)
  - Scaling (big slots, fewer slices)

Units of work and resources

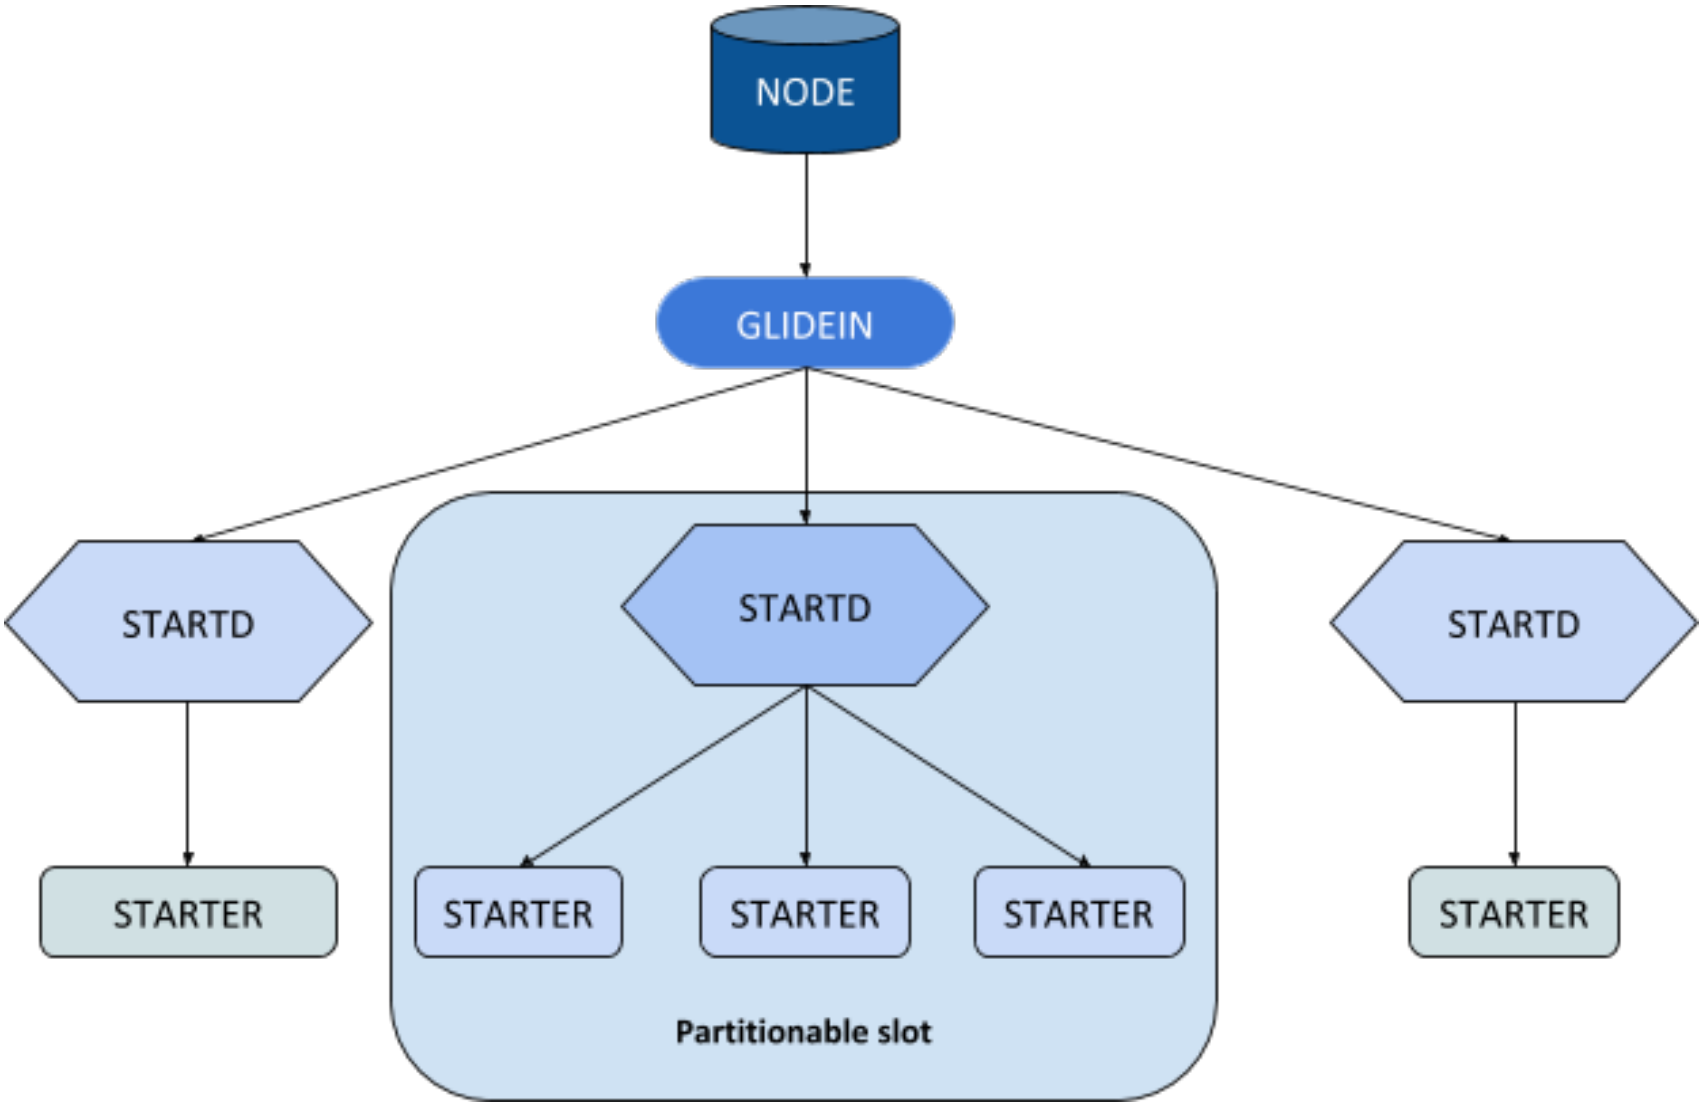

- Scientific computing
- GlideinWMS
- HTCondor
- Resources
- **Monitoring**
- Links
- Demo

Why monitoring is so important?

- Something may break, what?
- CE
  - May refuse to accept glideins
  - May not start glideins
  - Fail to tell us what the status of the job is
- The worker node may be broken/misconfigured (validation fails)
- Networking may not work properly
- Central Manager never hears from the Startd
- Schedd cannot talk to Startd
- Security infrastructure could be broken (CAs missing)
- Jobs not matching

Monitoring resources

- Logs
  - Frontend
  - Factory
  - HTCondor
  - CE
  - Glidein
- ClassAds
- GlideinWMS monitoring
- Monitoring FIFEMON

GlideinWMS monitoring

Glidein Factory Status - gfactory_instance@SDSC

Choose a Glidein destination...

•View has: 20 Total Entries - (CMS_T0-Frontend_cmspilot)

Select a view: Active only Troubleshoot Period

CMS_T0-Frontend_cmspilc

Autoupdate (30 mins)

Update Table Reset All Selections

Link

XML last update: Tue May 28 17:02:05 2019

		Status:									Requested:		Client Monitor:						
Entry Name		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
CMSHTPC_T2_CH_CERN_ce503	↑	5	20	0	20	0	0	0	0	50	1283	21	12	9	79921	38	304436	50	213
CMSHTPC_T2_CH_CERN_ce504	↑	3	20	0	20	0	0	0	0	30	1283	21	0	0	0	30	304436	30	209
CMSHTPC_T2_CH_CERN_ce505	↑	5	28	11	5	0	0	12	32	50	1286	21	31	17	79921	19	304436	50	212
CMSHTPC_T2_CH_CERN_ce506	↑	4	25	0	25	0	0	0	19	40	1285	21	24	3	79921	16	304436	40	224
CMSHTPC_T2_CH_CERN_ce507	↑	7	21	13	0	0	0	8	44	70	1286	21	53	41	79921	25	304436	80	211
CMSHTPC_T2_CH_CERN_ce508	↑	14	20	20	0	0	0	0	32	140	1286	21	106	61	79921	52	304436	160	644
CMSHTPC_T2_CH_CERN_ce509	↑	57	20	0	20	0	0	0	0	570	1286	20	395	188	79921	163	304436	570	224
CMSHTPC_T2_CH_CERN_ce510	↑	113	20	0	20	0	0	0	0	1130	1293	21	835	461	79921	271	304436	1130	328
CMSHTPC_T2_CH_CERN_ce511	↑	120	20	0	20	0	0	0	0	1200	1305	21	881	445	79921	274	304436	1200	324
CMSHTPC_T2_CH_CERN_ce512	↑	24	23	0	23	0	0	0	37	240	1290	21	165	106	79921	65	304436	240	324
CMSHTPC_T2_CH_CERN_ce513	↑	107	23	0	20	0	0	3	0	1070	1292	21	806	395	79921	241	304436	1070	322
CMSHTPC_T2_CH_CERN_ce514	↑	40	21	18	3	0	0	0	39	400	1298	21	284	207	79921	78	304436	390	321
CMSHTPC_T2_CH_CERN_cet01_10	↑	222	20	0	20	0	0	0	0	2664	1070	20	2053	1168	79938	603	304436	2664	180

GlideinWMS monitoring - Kibana

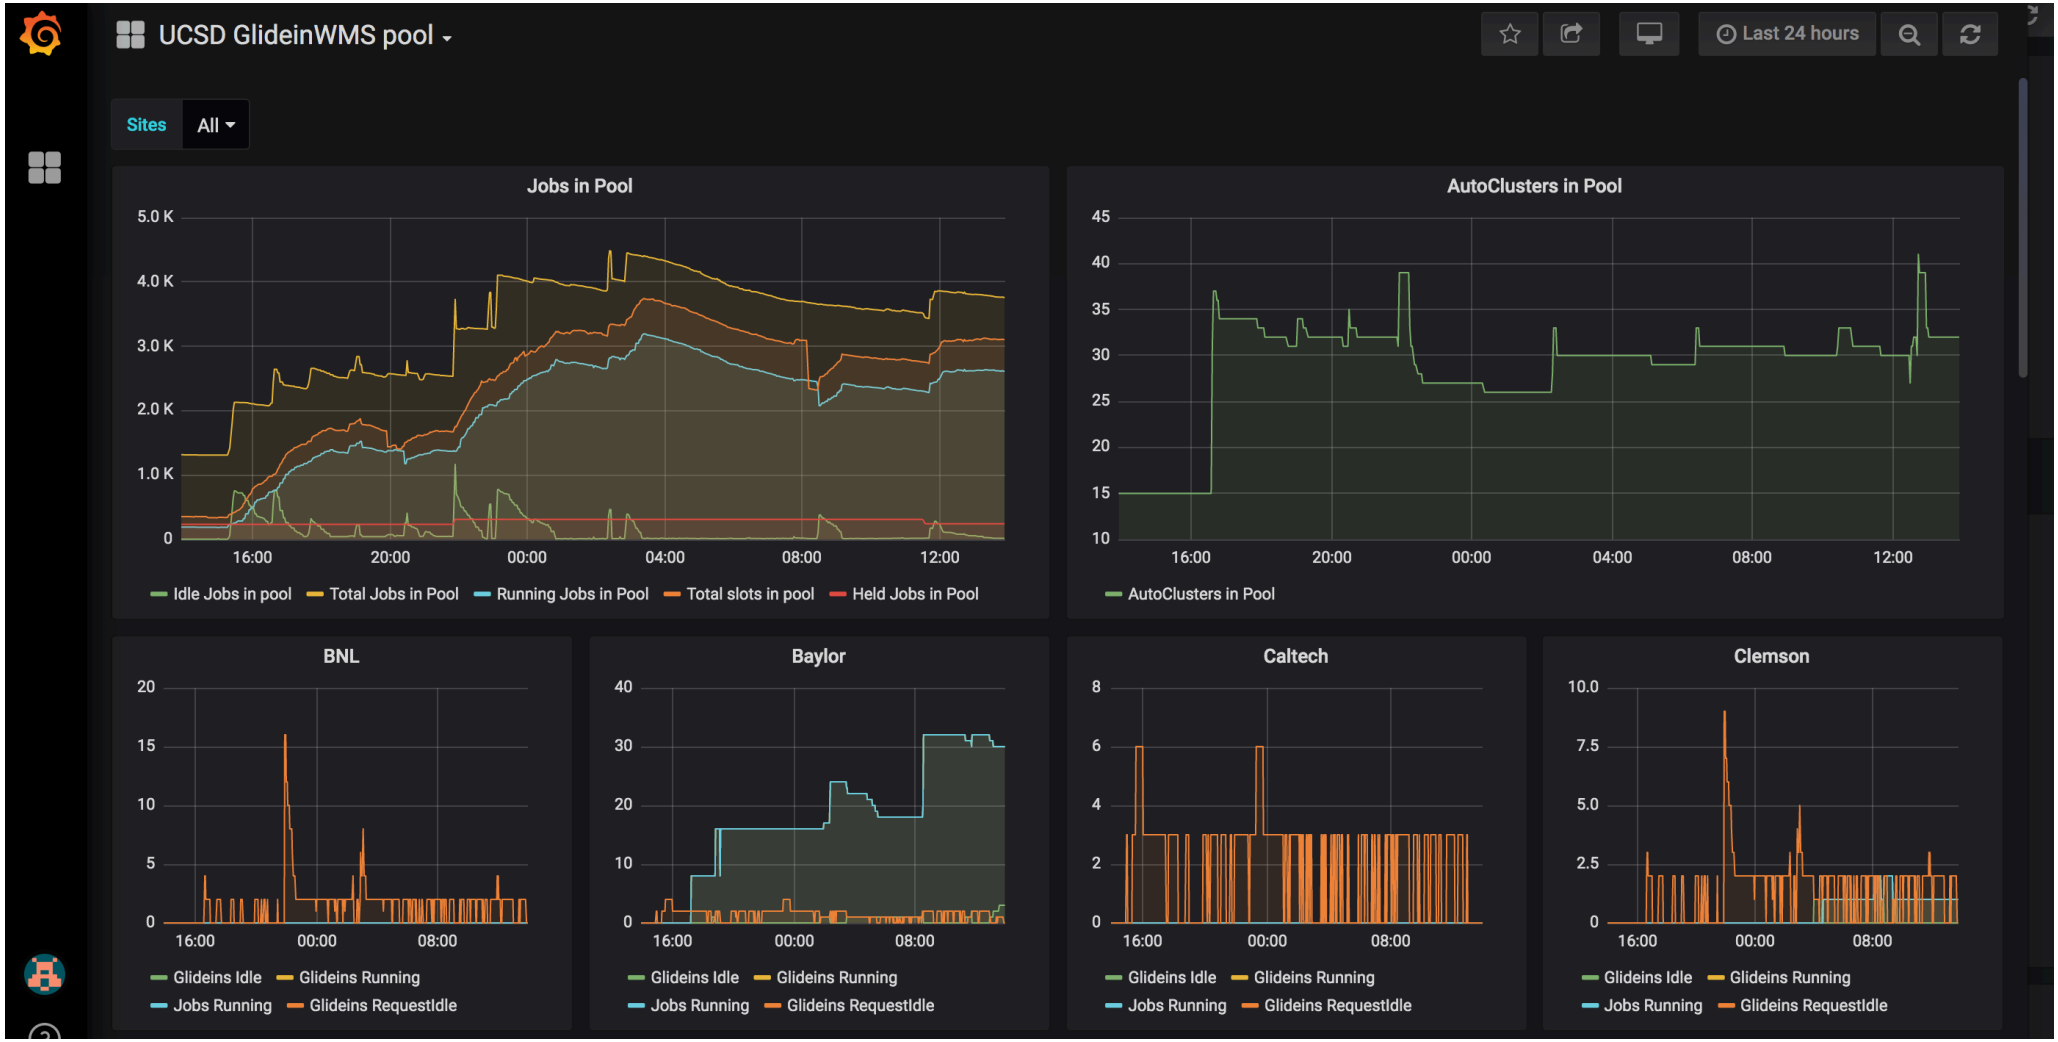

Fermilab

GlideinWMS monitoring - GRACC and Graphana

GlideinMonitor

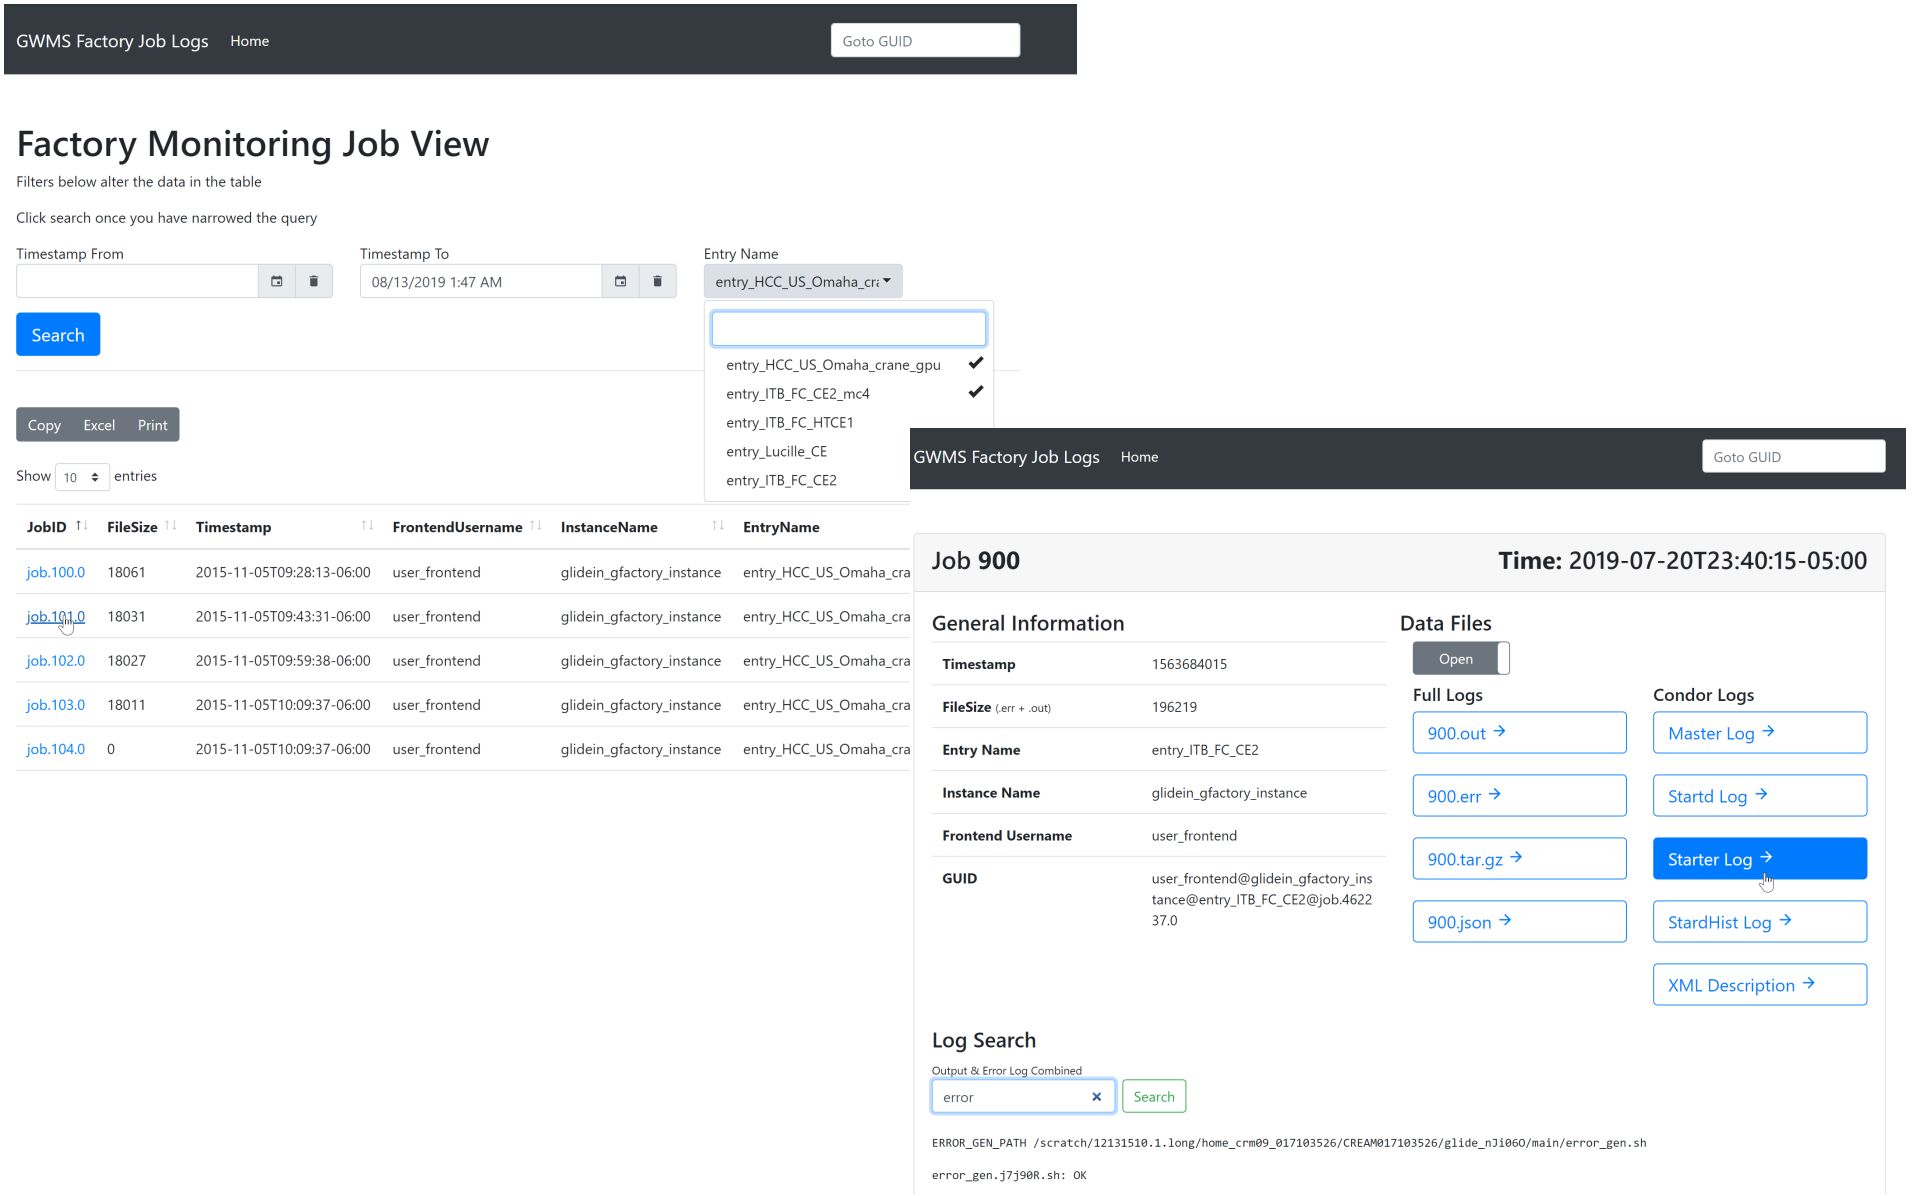

- Scientific computing
- GlideinWMS
- HTCondor
- Resources
- Monitoring
- **Links**
- Demo

Links – GlideinWMS Website

glideinwms.fnal.gov

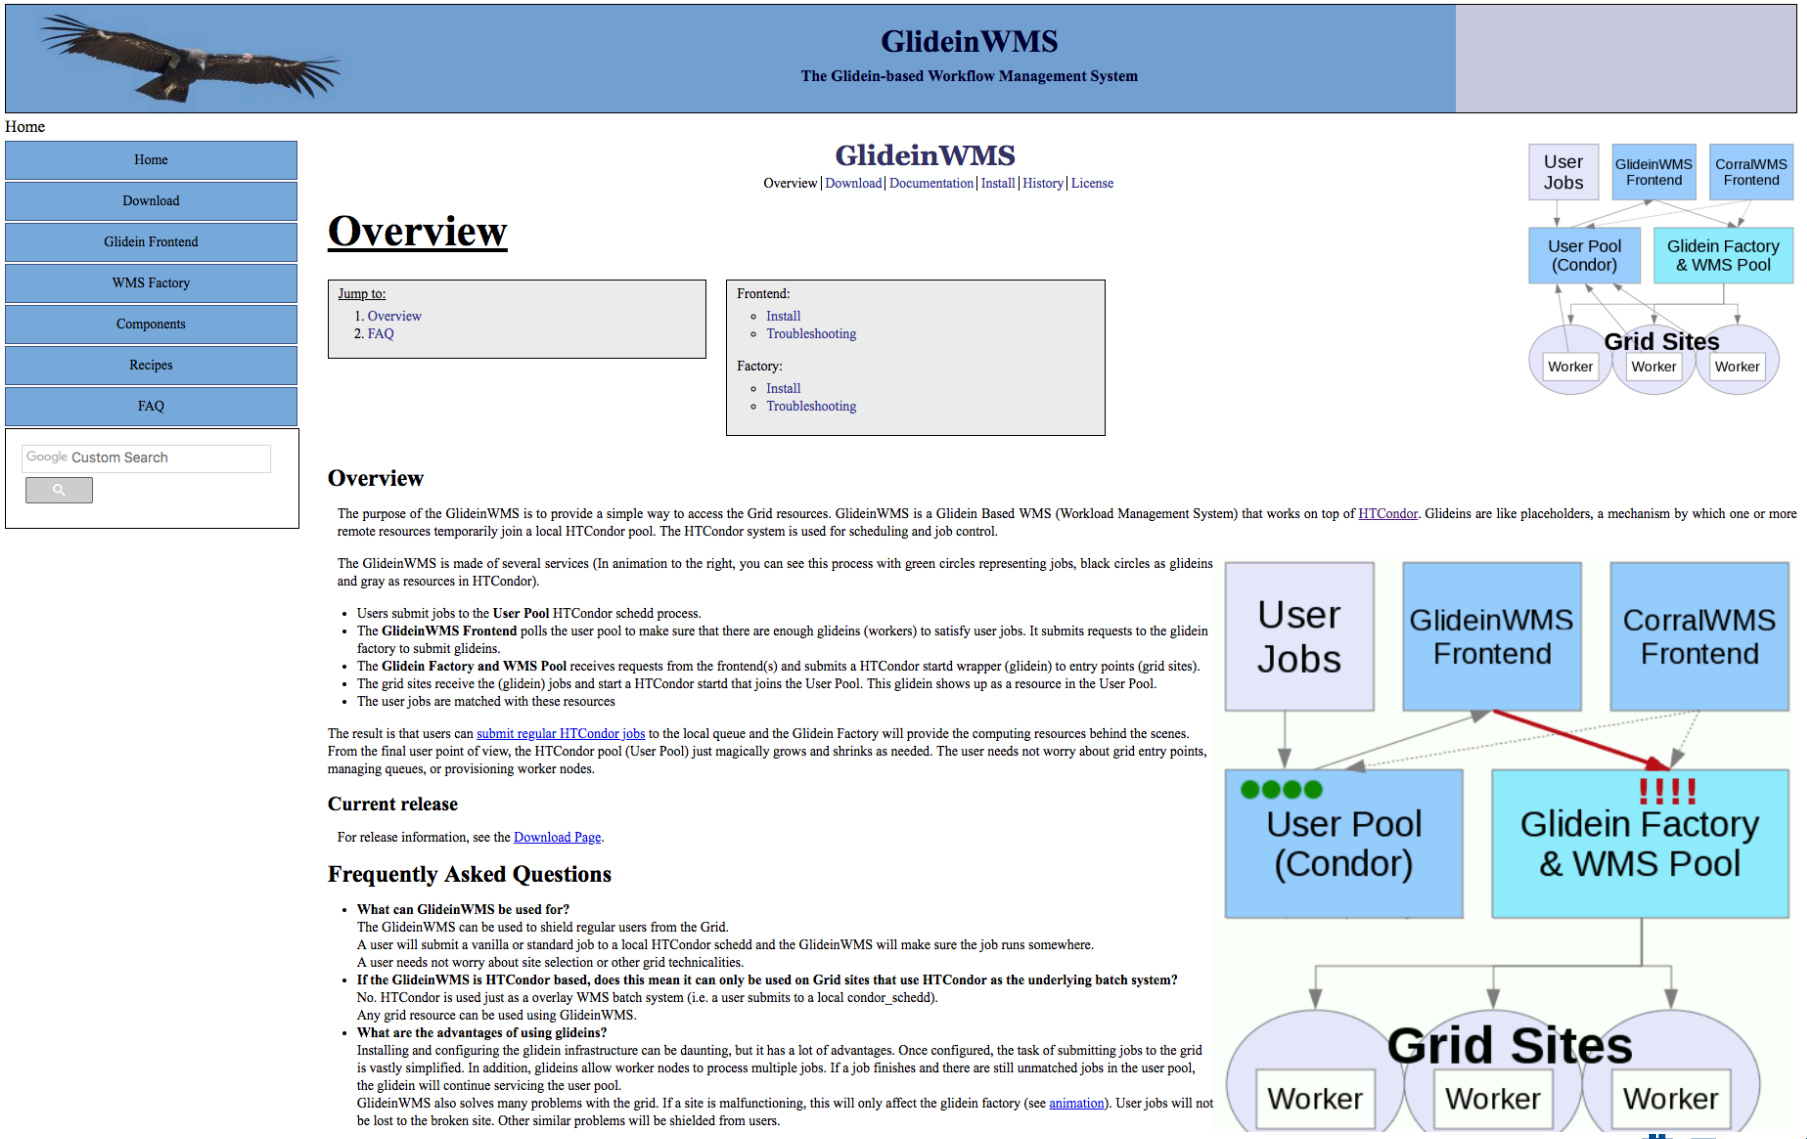

Fermilab

Links – GlideinWMS API

glideinwms.fnal.gov/api/

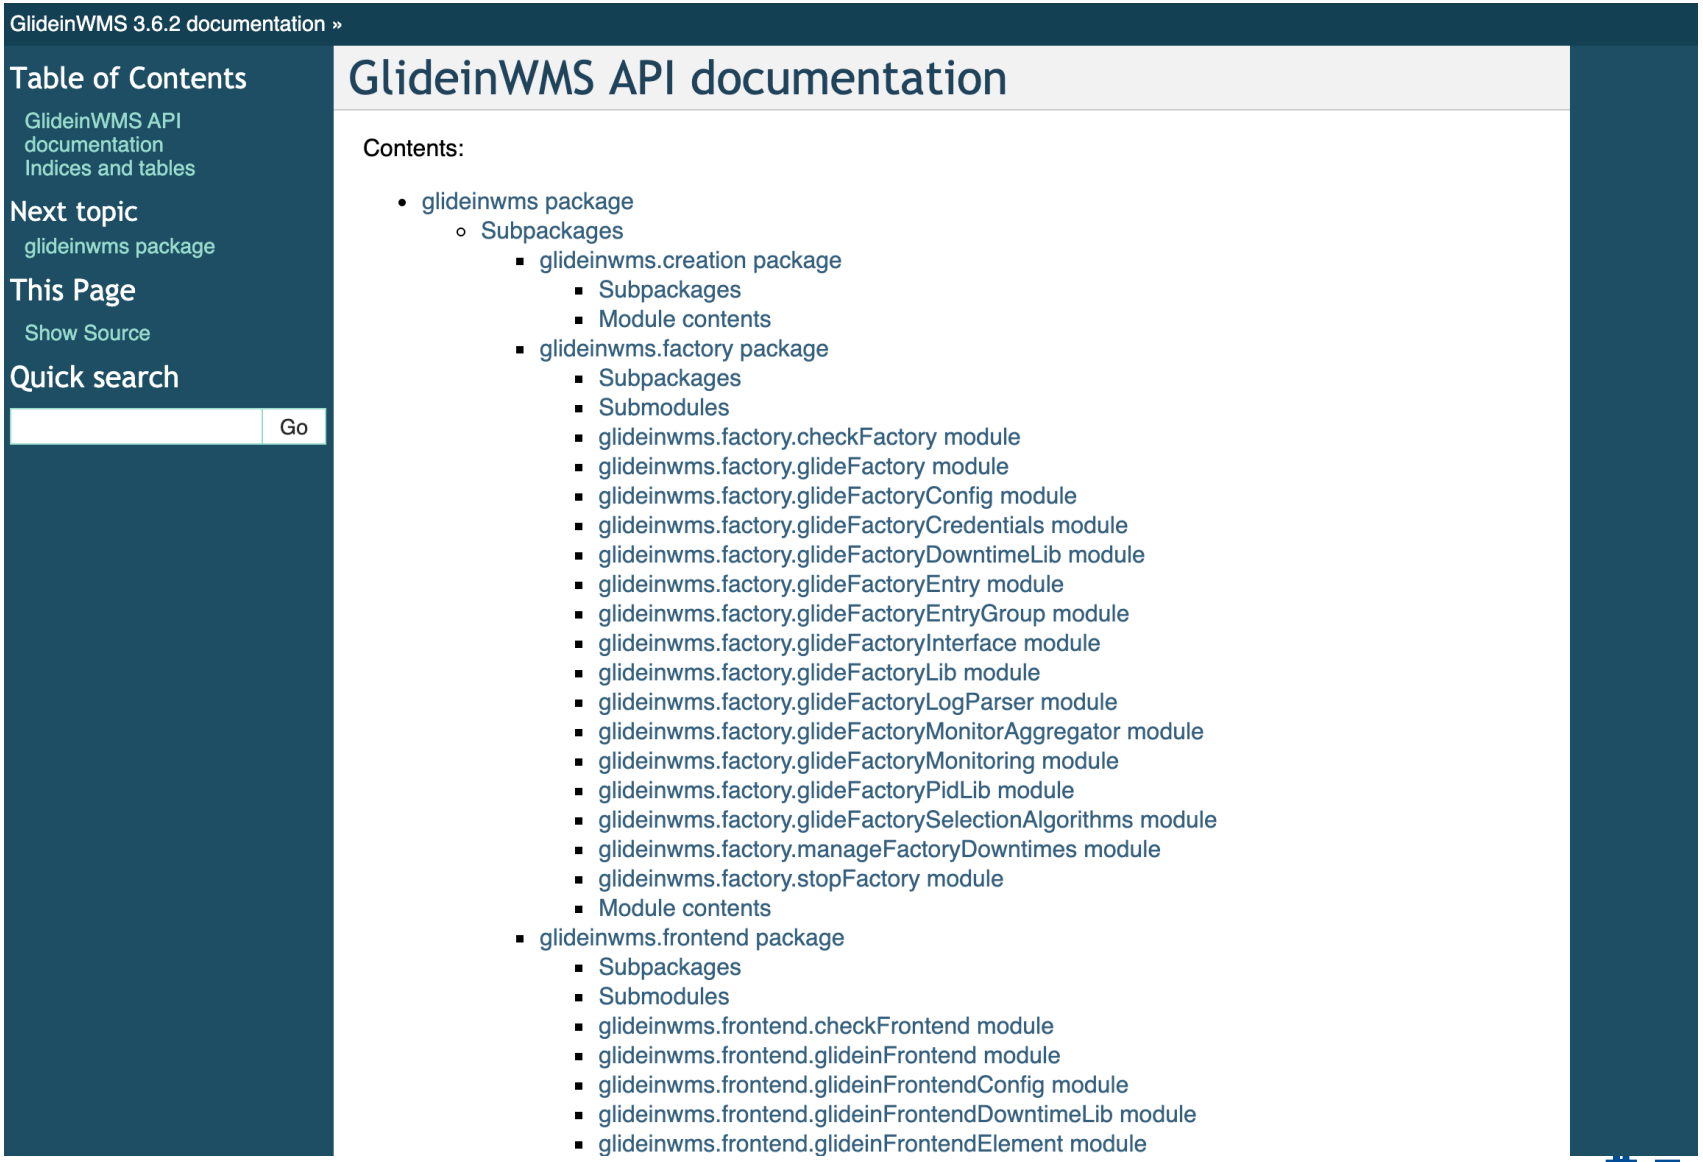

GlideinWMS 3.6.2 documentation »

Table of Contents

GlideinWMS API documentation
Indices and tables

Next topic

glideinwms package

This Page

Show Source

Quick search

Go

GlideinWMS API documentation

Contents:

- glideinwms package
  - Subpackages
    - glideinwms.creation package
      - Subpackages
      - Module contents
    - glideinwms.factory package
      - Subpackages
      - Submodules
      - glideinwms.factory.checkFactory module
      - glideinwms.factory.glideFactory module
      - glideinwms.factory.glideFactoryConfig module
      - glideinwms.factory.glideFactoryCredentials module
      - glideinwms.factory.glideFactoryDowntimeLib module
      - glideinwms.factory.glideFactoryEntry module
      - glideinwms.factory.glideFactoryEntryGroup module
      - glideinwms.factory.glideFactoryInterface module
      - glideinwms.factory.glideFactoryLib module
      - glideinwms.factory.glideFactoryLogParser module
      - glideinwms.factory.glideFactoryMonitorAggregator module
      - glideinwms.factory.glideFactoryMonitoring module
      - glideinwms.factory.glideFactoryPidLib module
      - glideinwms.factory.glideFactorySelectionAlgorithms module
      - glideinwms.factory.manageFactoryDowntimes module
      - glideinwms.factory.stopFactory module
      - Module contents
    - glideinwms.frontend package
      - Subpackages
      - Submodules
      - glideinwms.frontend.checkFrontend module
      - glideinwms.frontend.glideinFrontend module
      - glideinwms.frontend.glideinFrontendConfig module
      - glideinwms.frontend.glideinFrontendDowntimeLib module
      - glideinwms.frontend.glideinFrontendElement module

Links – OSG Docs: install

https://opensciencegrid.org/docs/other/install-gwms-frontend/

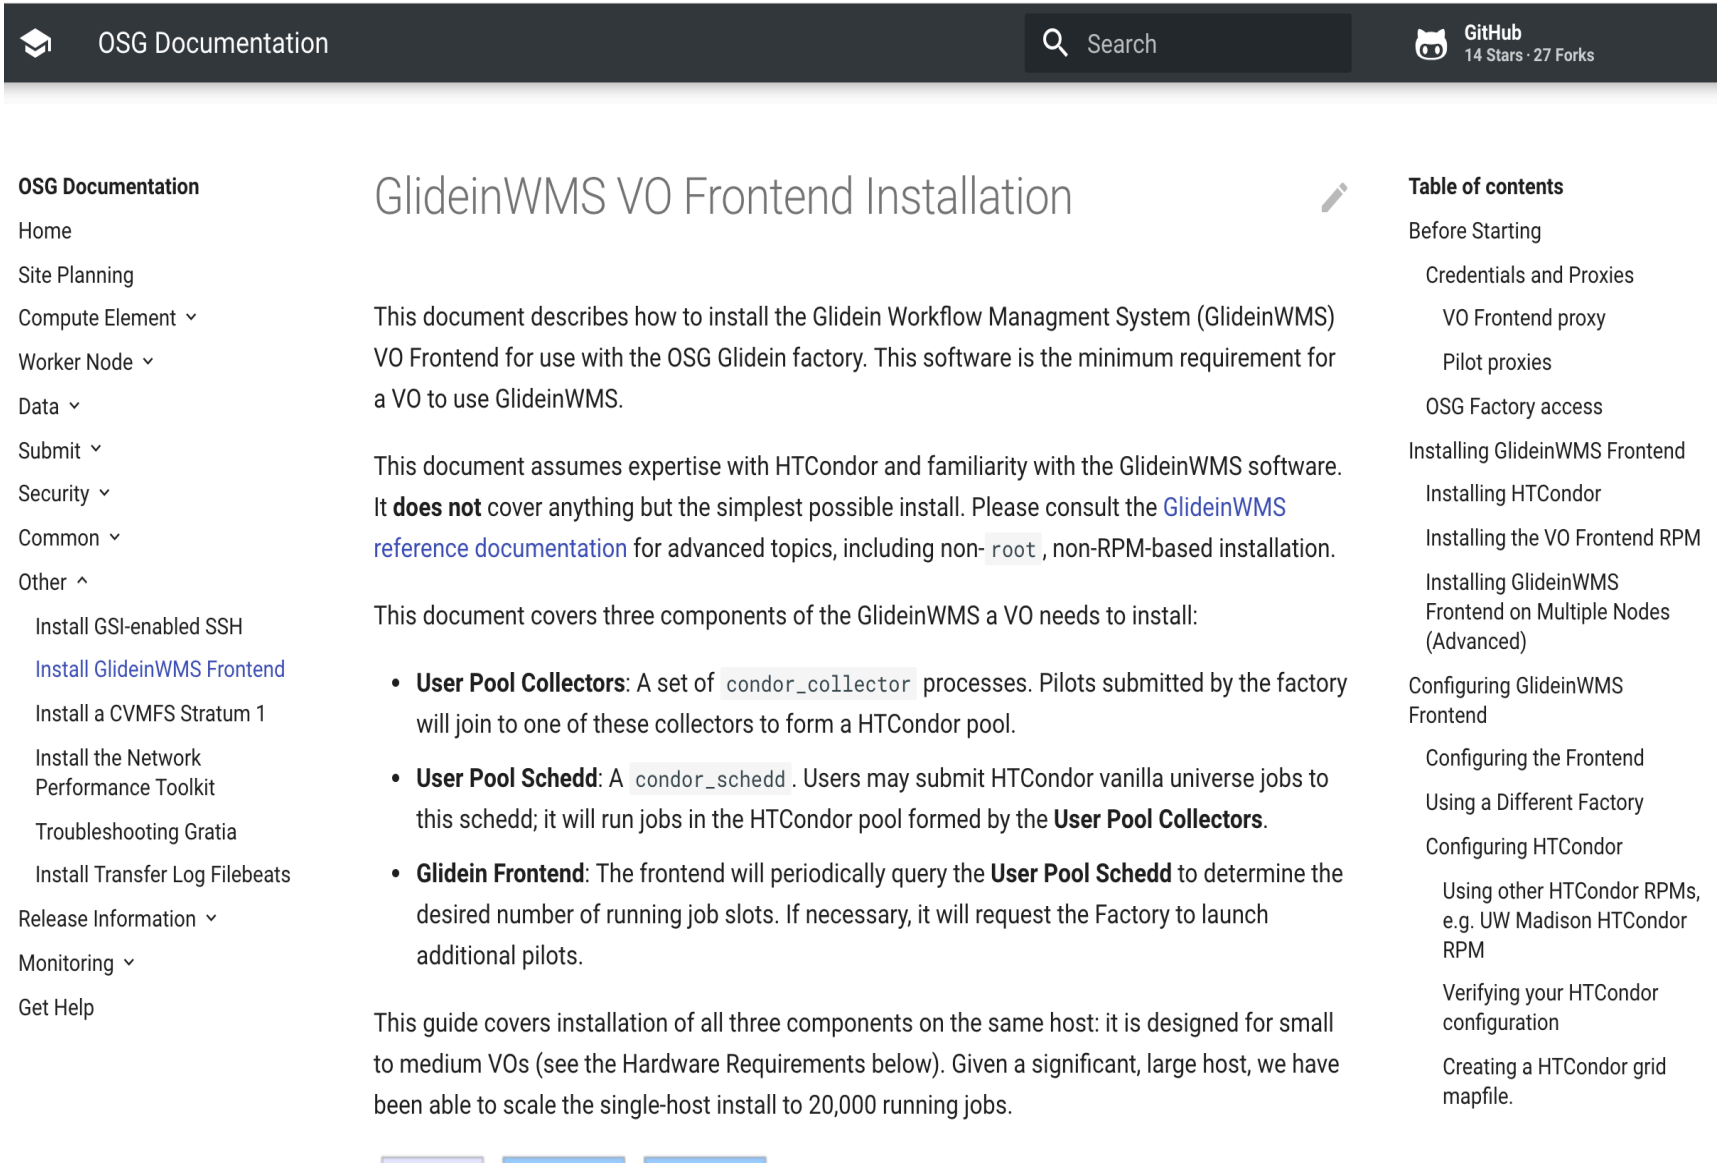

Links – OSG Docs: install

https://opensciencegrid.org/operations/services/install-gwms-factory/

OSG Operations | Search | GitHub 2 Stars · 17 Forks

GlideinWMS Factory Installation

This document describes how to install a Glidein Workflow Managment System (GlideinWMS) Factory instance.

This document assumes expertise with HTCondor and familiarity with the GlideinWMS software. It **does not** cover anything but the simplest possible install. Please consult the GlideinWMS reference documentation for advanced topics, including non-root, non-RPM-based installation. In this document the terms glidein and pilot (job) will be used interchangeably.

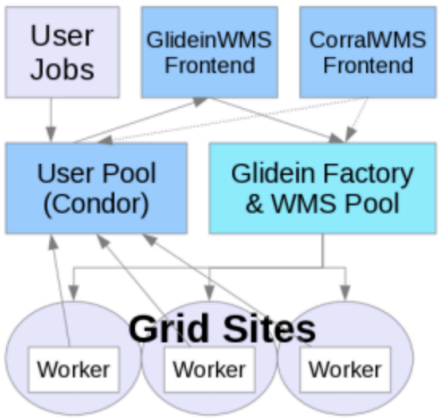

This parts covers these primary components of the GlideinWMS system:

Links – Weekly CI email

CI build of GlideinWMS_CI workflow for slf7 Succeeded

owner-glideinwms@listserv.fnal.gov <owner-glideinwms@listserv.fnal.gov> on behalf of **cireports_jenkins@fnal.gov <cireports_jenkins@fnal.gov>**
listserv
Monday, 3 June 2019 at 1:29 AM
Show Details

CI build of GlideinWMS_CI workflow for slf7 Succeeded

Build number: 917
GlideinwmsCI Web App: here
Jenkins build: here

HOSTNAME: buildservice008.fnal.gov
LINUX DISTRO: Description: Scientific Linux release 7.6 (Nitrogen)
PYTHON LOCATION: /scratch/workspace/glideinwms_ci/label_exp/RHEL7/label_exp2/swarm/venv-2.7/bin/python
PYLINT: pylint 1.8.4, astroid 1.6.0Python 2.7.5 (default, Apr 8 2019, 14:44:40) [GCC 4.8.5 20150623 (Red Hat 4.8.5-36)]
PEP8: 2.5.0

GIT BRANCHES	PYLINT				UNIT TESTS			FUTURIZE
	FILES CHECKED	FILES WITH ERRORS	TOTAL ERRORS	PEP8 ERRORS	#TESTS	#ERRORS	#FAILURES	FILES TO BE REFACTORED
branch_v3_4	203	0	0	3528	510	0	0	2
master	203	0	0	3435	503	0	0	2
master_ci	199	0	0	3610	465	0	0	2
v35/20215_rpm	203	0	0	3435	503	0	0	2
v35/21940	206	0	0	3446	577	0	0	3

branches to refactorize:
branch_v3_4
master
master_ci
v35_20215_rpm
v35_21940

Git commits in the last day:
^^
./glideinwms/.git:
origin http://cdcvs.fnal.gov/projects/glideinwms
./glideinwms_master_ci/.git:
origin http://cdcvs.fnal.gov/projects/glideinwms
^^

SVN: last 10 commits in the trunk/tag:
^^
^^

CI report sent to:
ci_build_reports@fnal.gov,thein@fnal.gov,marcom@fnal.gov,glideinwms@fnal.gov

Cheers,
CI workflow script

Links – Weekly CI email (cont)

https://buildmaster.fnal.gov/buildmaster/view/GlideinWMS/

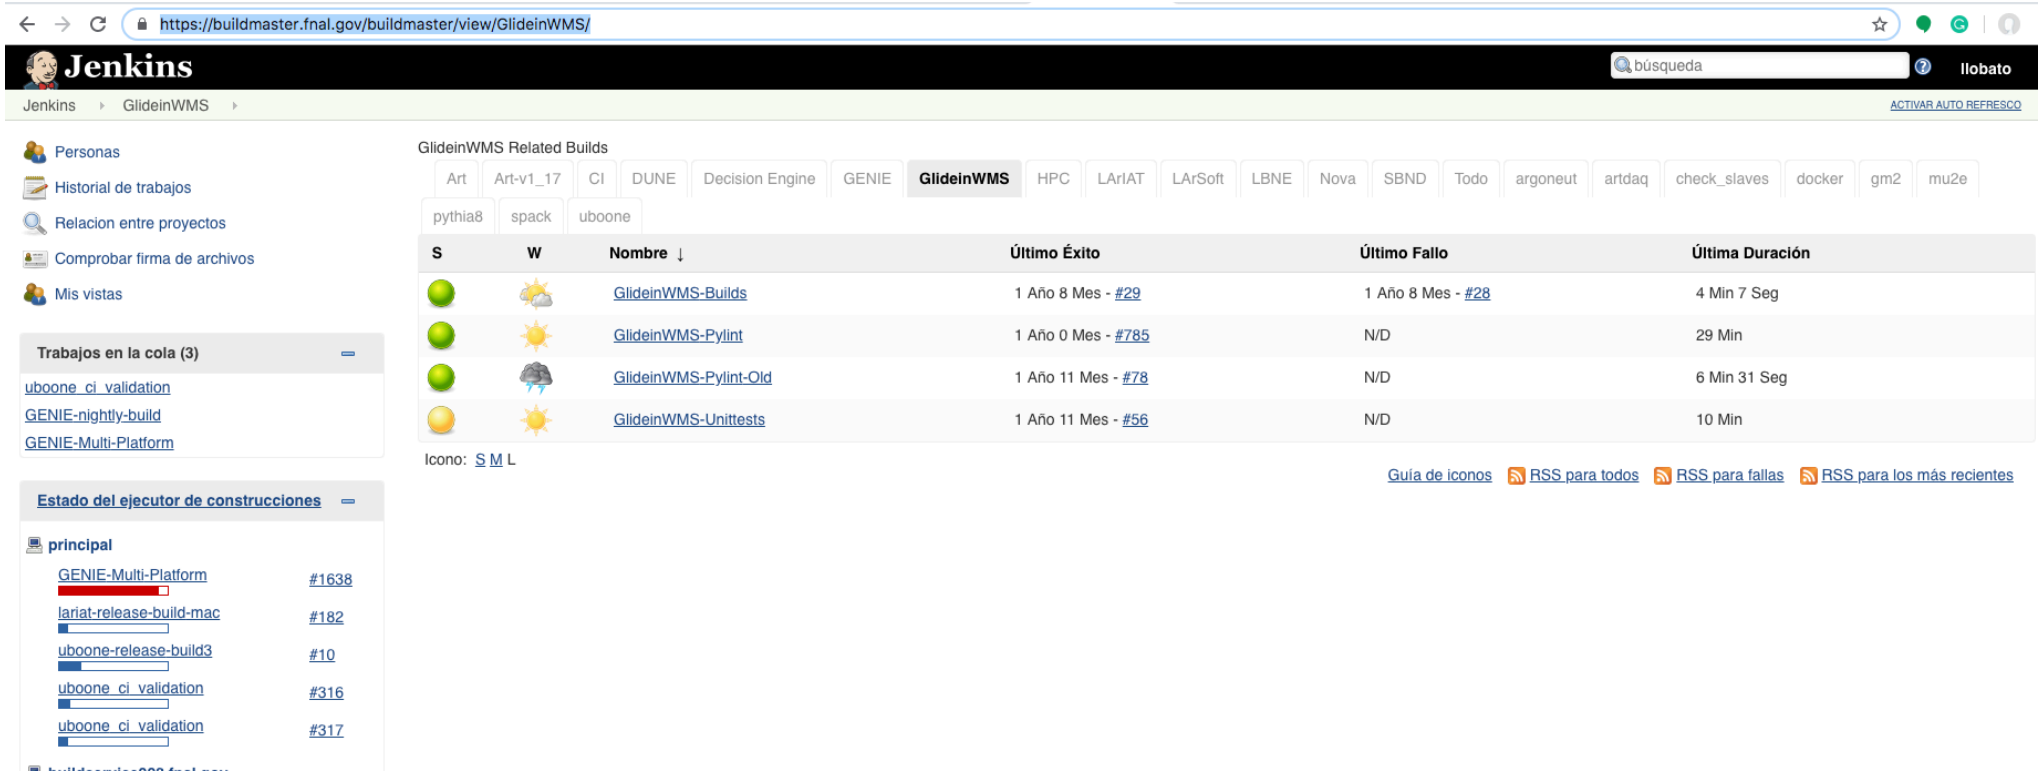

Links - Redmine

https://cdcvs.fnal.gov/redmine/projects/glideinwms/wiki#For-New-Member-Orientation-see-NewMemberOrientation

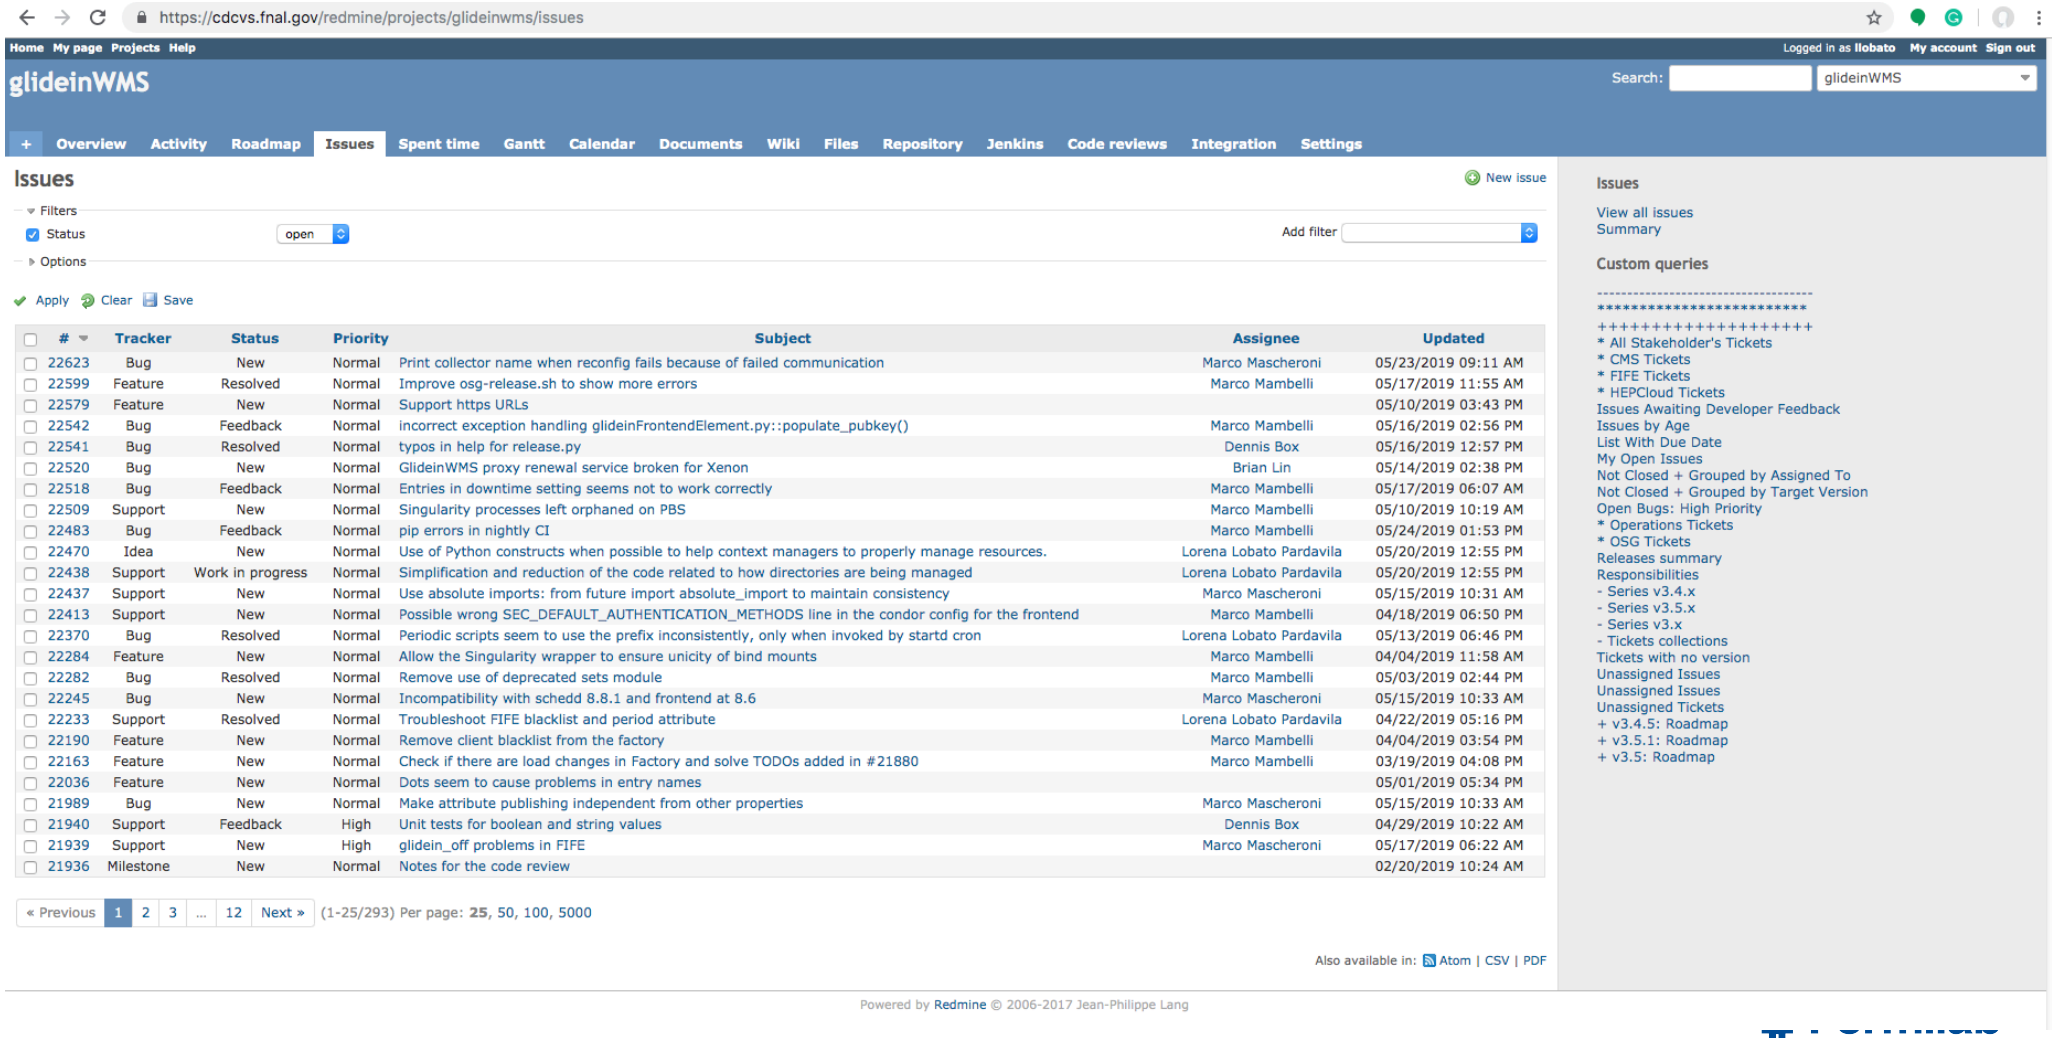

Links – Git repo

git clone ssh://p-glideinwms@cdcvs.fnal.gov/cvs/projects/glideinwms

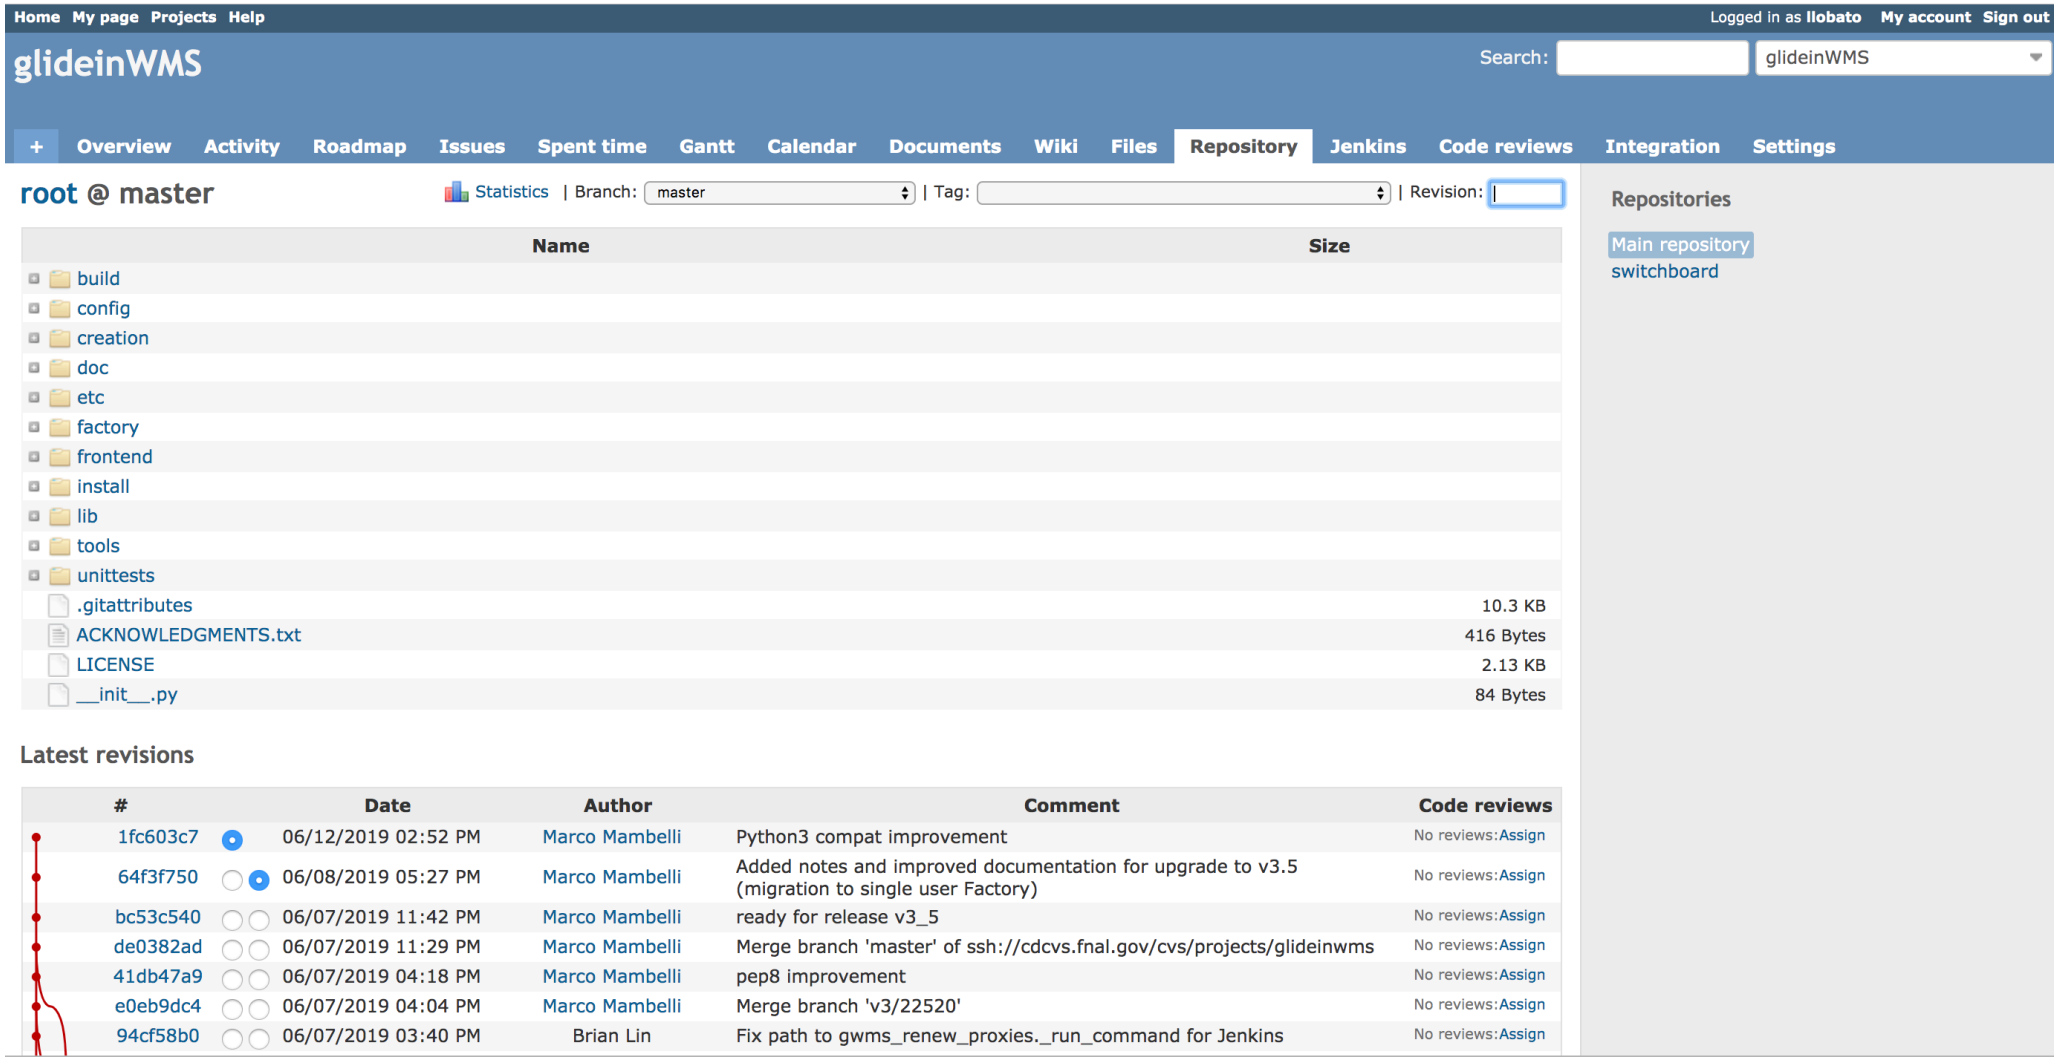

Links – GitHub

https://github.com/glideinWMS

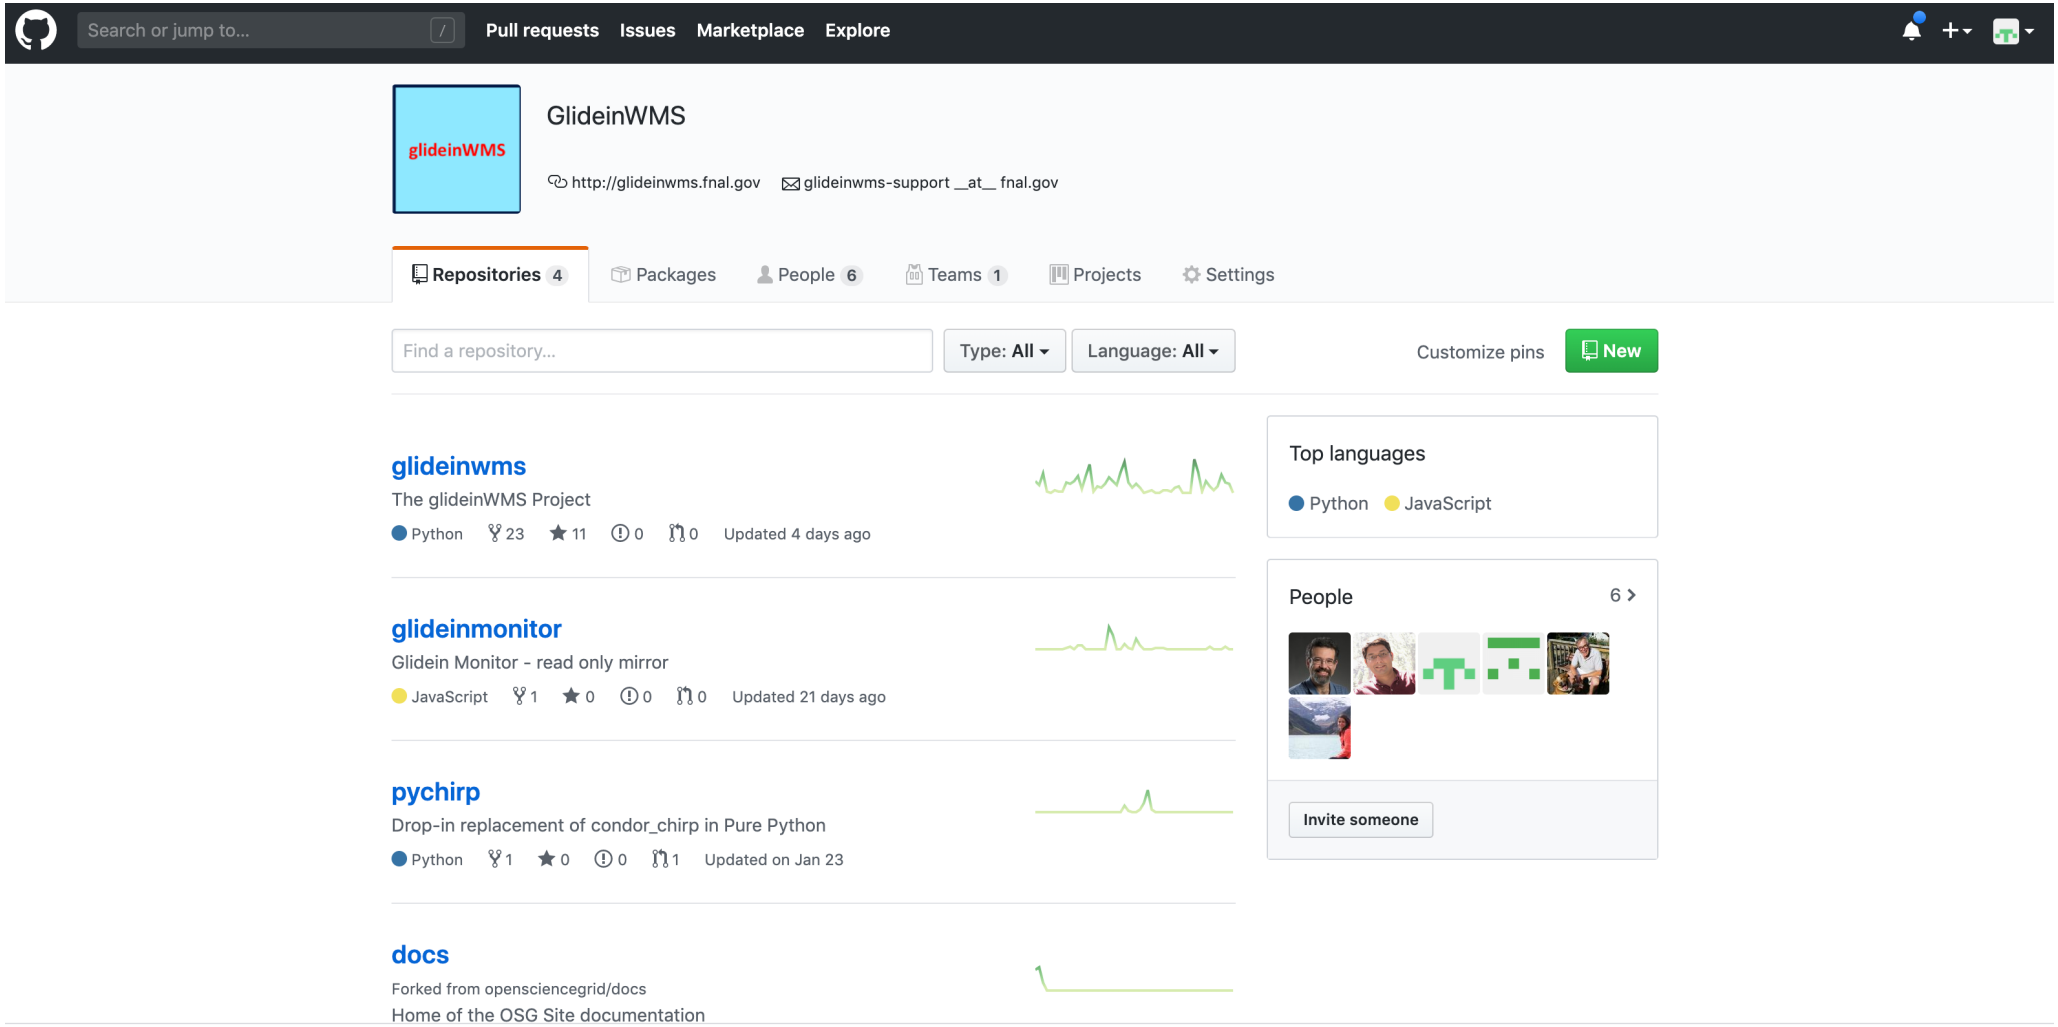

Fermilab

- Scientific computing
- GlideinWMS
- HTCondor
- Resources
- Monitoring
- Links
- **Demo**

Summary

- Scientific computing requires a lot of computing
- GlideinWMS is a pilot based resource provisioning tool for distributed High Throughput Computing
  - Provides reliable and uniform virtual clusters
  - Submits Glideins to unreliable heterogeneous resources
- HTCondor provides the tools
  - ClassAds, schedd, startd, collector
- Worker nodes on Resources
  - cores, memory and disk
  - Can run multiple jobs
- Good monitoring helps troubleshooting
- http://glideinwms.fnal.gov/

References

- HTCondor slides are based in part on a presentation by Todd Tannenbaum and the HTCondor team http://www.cs.wisc.edu/condor/
- GlideinWMS slides are based in part on previous presentations by the GlideinWMS project developers

Extra slides

Glidein based Workload Management System

- The Grid is a sum of thousands of independent Grid sites
- Choosing where to try to run the jobs is not a trivial task -> Extra operations
  - To keep in mind: reliability, scalability, priorities, location, quotas…
- Resources are located in independent pools
  - How minimizing the waste?
- Need abstraction layer for submission to any site
- GlideinWMS makes it easier and transparent to the users
  - Division of operations
- Advantages of a local batch system (HTCondor based overlay system)

Glidein based Workload Management System

GWMS is a workload manager for scientific jobs: it provisions resources from different sources and makes them available to users jobs in a single virtual cluster

- The cluster provided to you is virtual because the hosts are in different places and sometime not even there until you need them
- Uses Glideins or pilots to acquire and prepare the resources, or worker nodes, where the jobs run
- Based on HTCondor

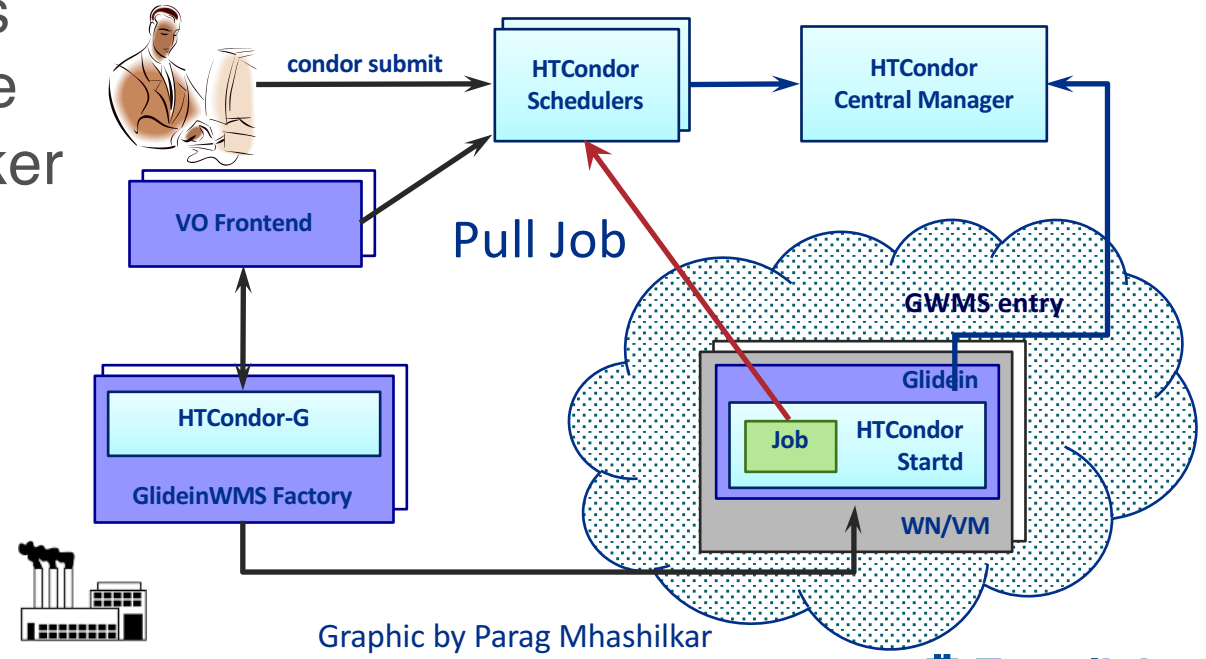

Fermilab

HTCondor building blocks in Glidein WMS

- The Factory works with an HTCondor pool, WMS pool, to submit Glideins to different resources
- The HTCondor Glideins are pilots that launch a startd that registers on a second HTCondor pool, User pool
- User jobs are matched and execute on the resources
- The Frontend monitors the user schedds and notifies the Factory about the need for more Glideins

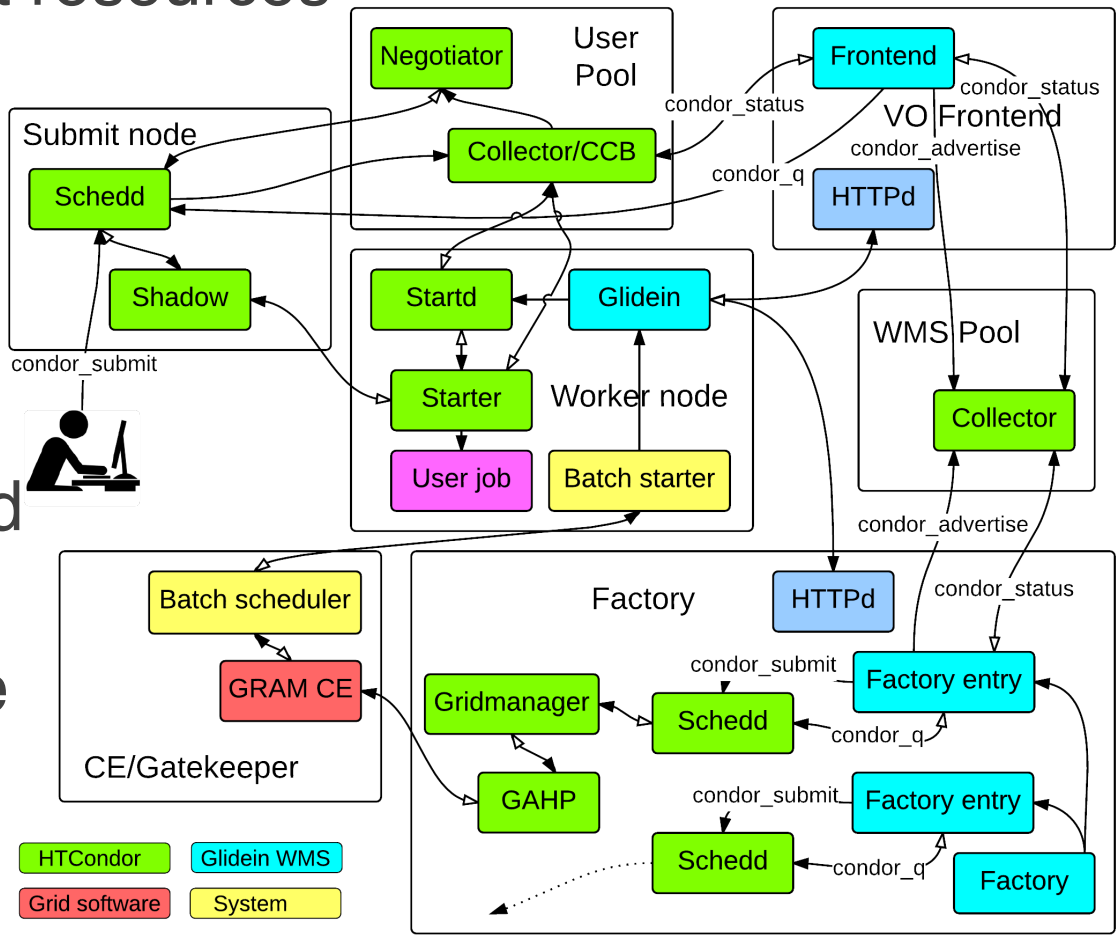

Fermilab

New resources added over time

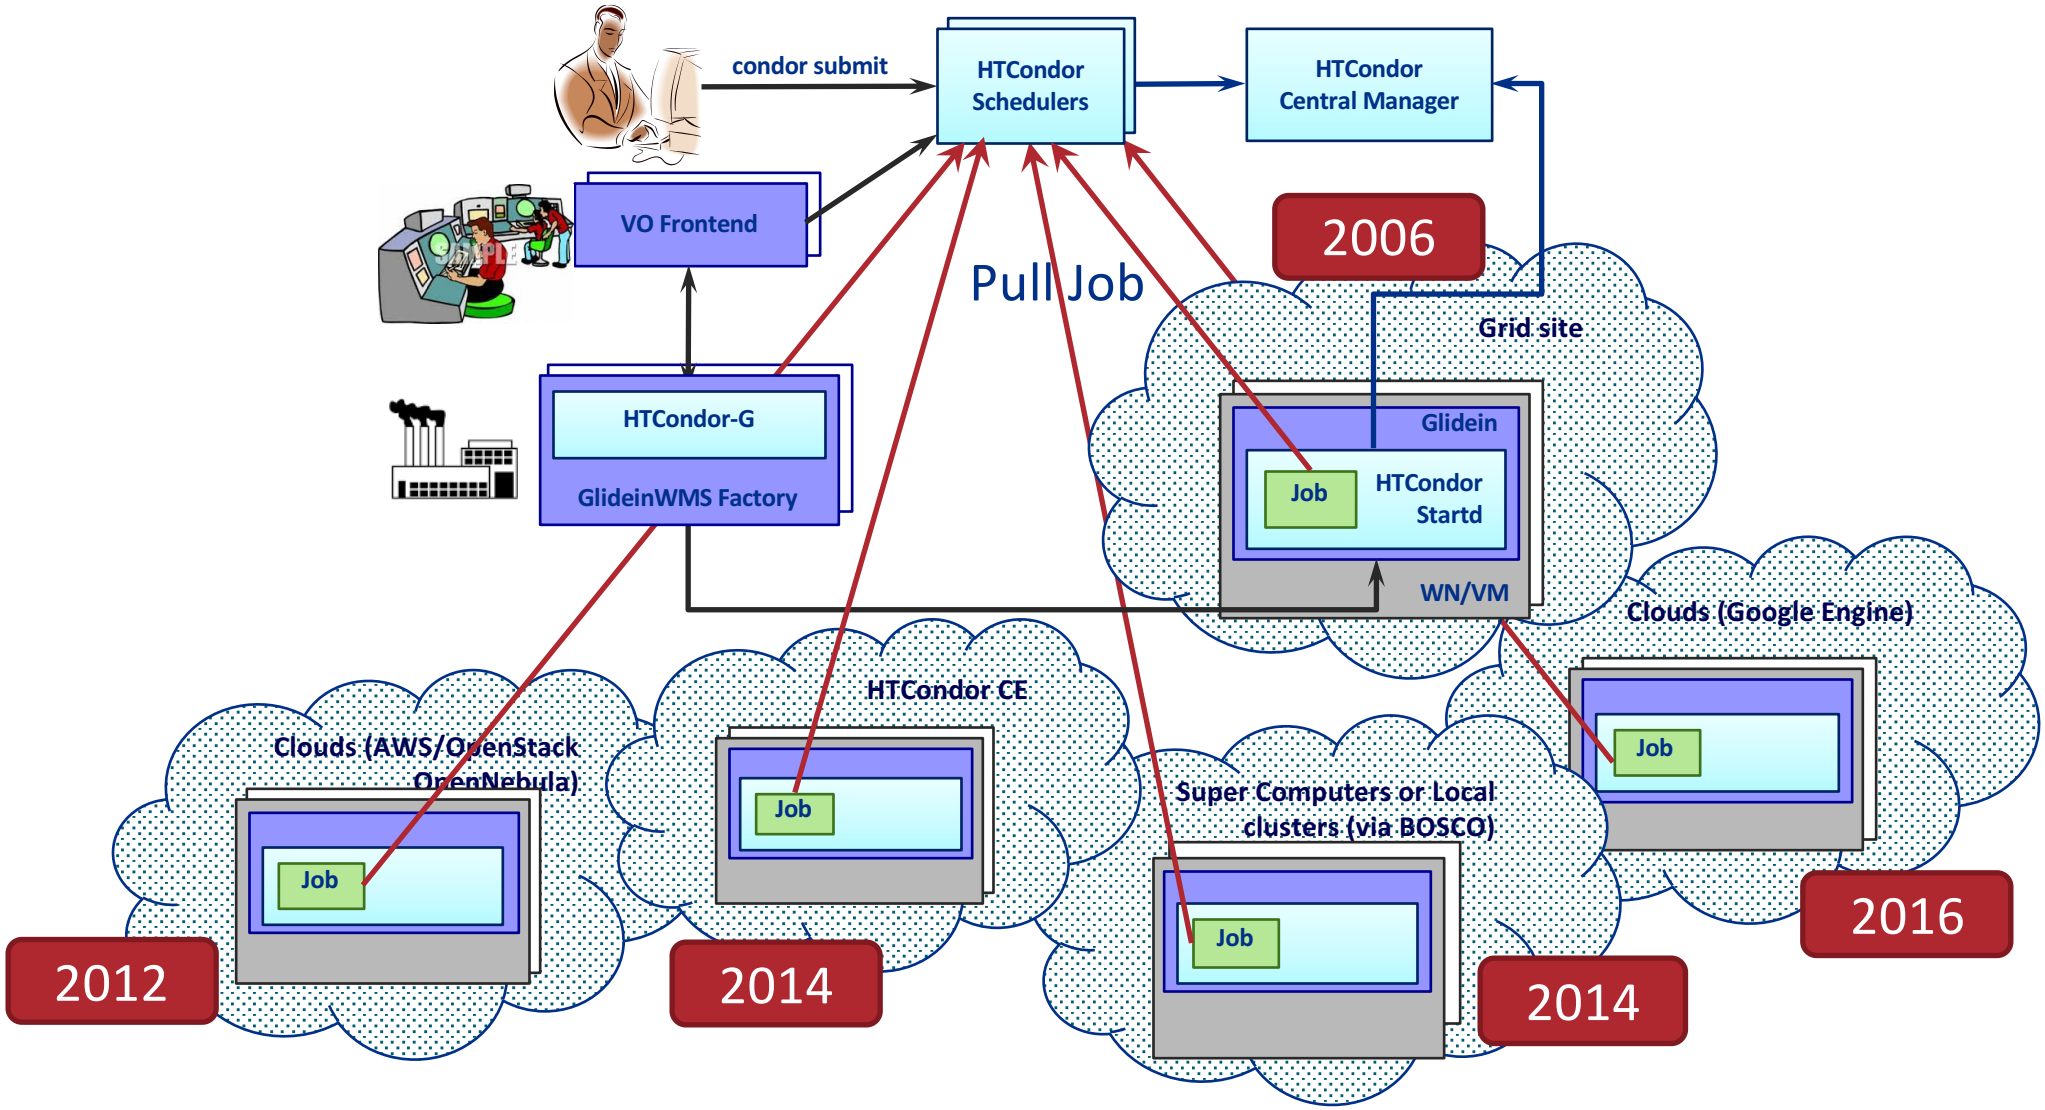

Fermilab

Job and Machine ‘dimensions’

- Job request
  - request_cpus: number of cores, integer, default 1.
  - request_disk: amount of disk space in Kbytes, default to sum of sizes of the job's executable and all input files (or image size)
  - request_memory: amount of memory space in Mbytes, default to executable size
- Machine
  - Cpus: number of cores, integer, by default the available cores
  - Disk: amount of disk space on this machine available for the job in KiB, by default the available space
  - Memory: amount of RAM in MiB in this slot
- Over and Under provision