

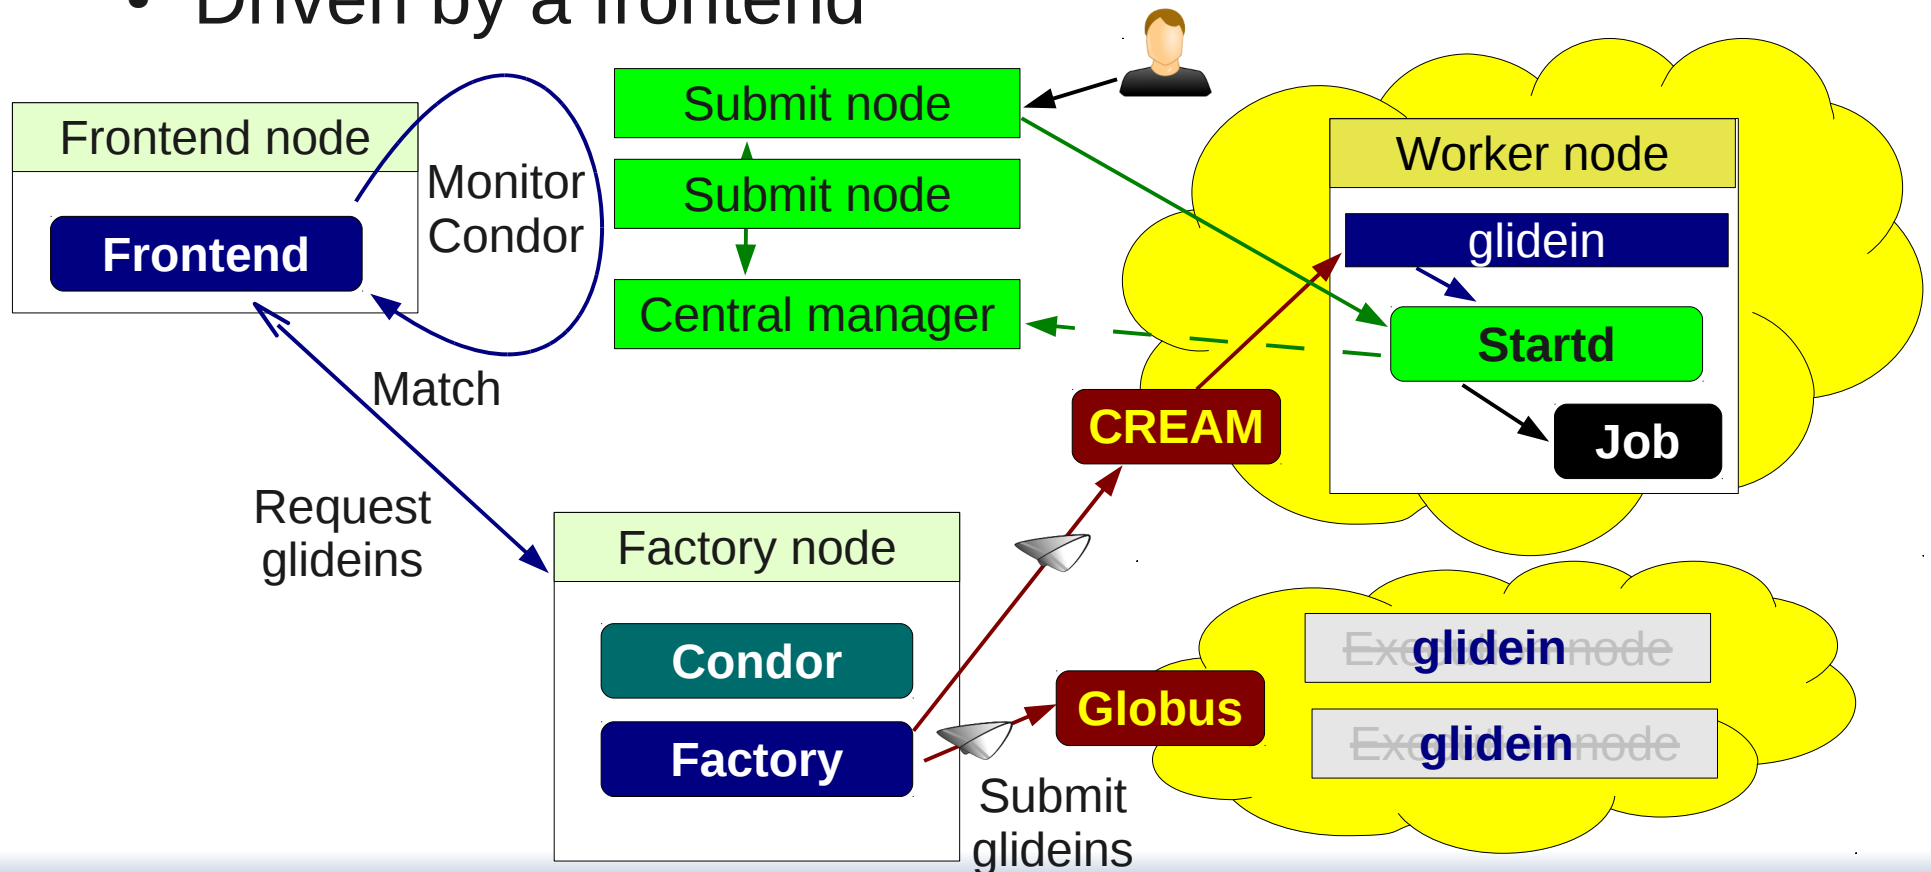
GlideinWMS Training @ UCSD

Configuring a glideinWMS factory

by Igor Sfiligoi (UCSD)

Refresher – Glidein factory

- The glidein factory knows about the sites and does the submission
 - Driven by a frontend

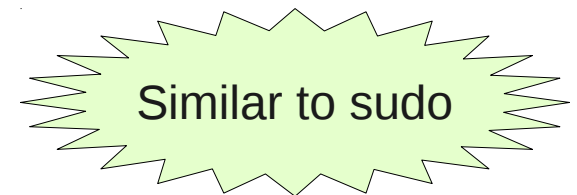


Setup and configuration

- glideinWMS comes with an installer
 - Will help you install all the needed software and do most config
- The factory config you get out of the installer is typically just a rough template
 - You will have to finish it by hand
 - It is in XML format

The Condor part

- The Condor components use standard config
 - Just using multiple schedds for scalability
- The collector uses GSI security for WAN
 - Will need x509 cert and proper security config. to talk to the frontends
 - Local auth. still FS based
- Factory will use **condor_root_switchboard**
 - For UID switching
 - **Must be owned by root and setuid**
 - **Config will need to be maintained by hand**
(More details later)



Privilege separation

- Frontends will be delegating proxies to the factory
 - Each frontend files owned by its own UID
<http://www.cs.wisc.edu/condor/CondorWeek2010/condor-presentations/sfiligoi-condorg.odp>
- The factory uses `condor_root_switchboard` to
 - Create dirs owned by right UID
 - Write files
 - Submit jobs

} UID switching

Must update every time a new frontend added

Auth. source user

Auth. target users

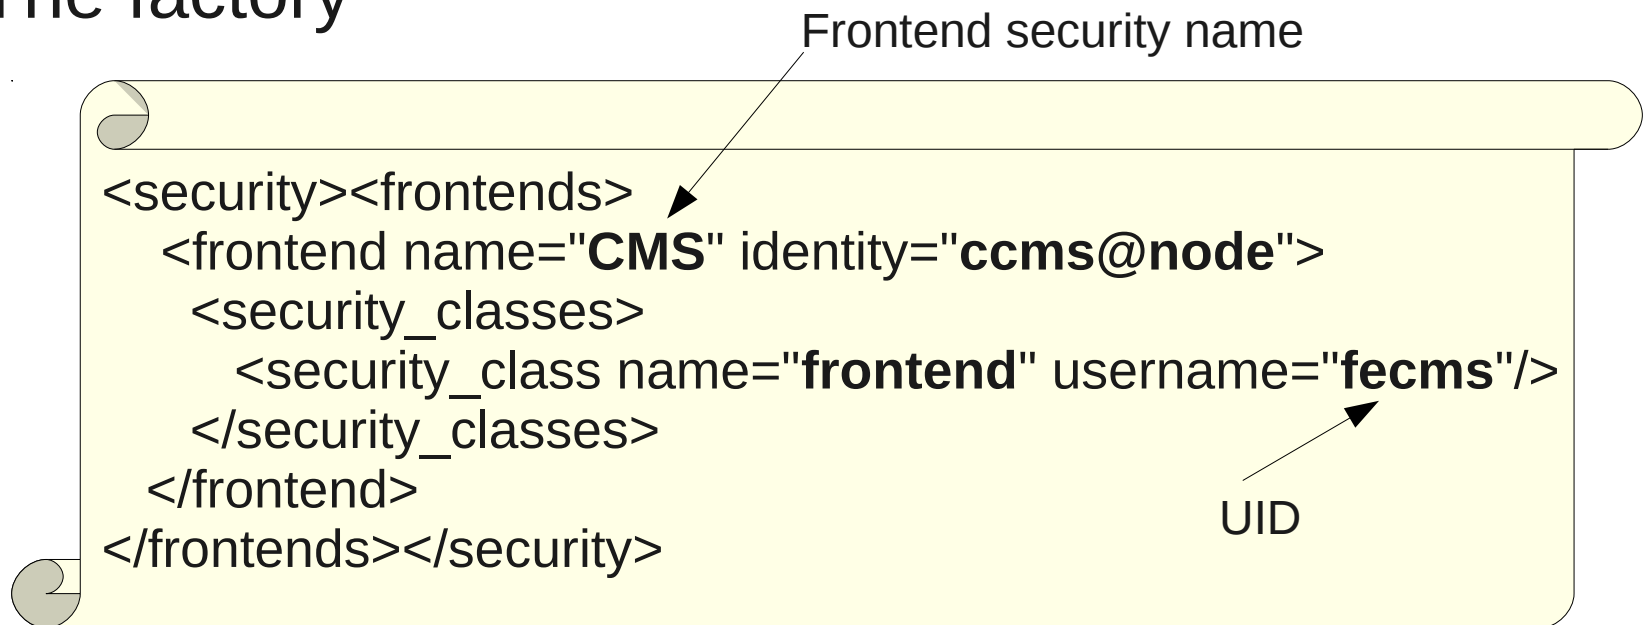
Base dirs for dir creation

`/etc/condor/privsep_config`

```
valid-caller-uids = gfactory  
valid-caller-gids = gfactory  
valid-target-uids = fe1 : fe2 : fecms  
valid-target-gids = fe1 : fe2 : fecms  
valid-dirs = /var/gfactory/clientlogs  
valid-dirs = /var/gfactory/clientproxies  
procd-executable = <path_to_procd>
```


Frontend authorization

- The factory must whitelist all the frontends it supports, both in
 - The collector (where X509 authentication happens)
 - The factory



Adding the sites

- The factory is supposed to know what Grid sites are out there and can be used
- Unfortunately, we don't have good tools for site discovery
 - The installer can query BDII and RESS, but it is very primitive
 - You are better off using other tools (like ldapsearch)
- **Admin needs to manually add sites in the XML file**



Work in progress to improve this

Site attributes

- Each site will have some attributes associated with it
 - Contact info – gatekeeper, jobmanager, RSL
 - Site config – work dir, platform, firewalls, glxec
 - Site limits – mostly wallclock time
 - Site properties – supported VOs, nearby SEs
 - User requested attributes – Can be anything
- Next slides explain the most often used ones

Site contact info

- How to submit to the site

```
<entry name="CMS_T2_DE_DESY_ce4"  
  gridtype="gt2"  
  gatekeeper="grid-ce4.desy.de:2119/jobmanager-lcgpbs"  
  rsl="(queue=cms)(jobtype=single)">
```

RSL often optional

- If you can reference an InfoSys, do it

```
<entry><infosys_refs>  
  <infosys_ref type="BDII" server="exp-bdii.cern.ch"  
    ref="GlueCEUniqueID=grid-ce4.desy.de:2119/jobmanager-lcgpbs-cms,Mds-Vo-name=DESY-HH,Mds-Vo-name=local,o=grid" />  
</infosys_refs></entry>
```

Helps with monitoring

Site config

- Tells the glidein how to behave

```
<entry name="CMS_T1_US_FNAL_ce3"
  work_dir="Condor">
  <attrs>
    <attr name="GLEXEC_BIN" value="OSG"/>
    <attr name="GLIDEIN_Glexec_Use" value="REQUIRED"/>
    <attr name="USE_CCB" value="False"/>
    <attr name="CONDOR_OS" value="rhel5"/>
    <attr name="CONDOR_ARCH" value="x86_64"/>
  </attrs></entry>
```

Special keyword

Note that one can host many condor binaries

```
<condor_tarballs>
  <condor_tarball arch="x86" os="rhel5" version="7.5.x" tar_file="path1"/>
  <condor_tarball arch="x86_64" os="rhel3" version="7.4.x" tar_file="path2"/>
</condor_tarballs>
```

Site limits

- Grid sites typically have wallclock limits

```
<entry><attrs>  
  <attr name="GLIDEIN_Max_Walltime"      value="65200"/>  
  <attr name="GLIDEIN_Retire_Time_Spread" value="3600"/>  
</attrs></entry>
```

- Pressure and sanity limits

```
<entry><config>  
  <max_jobs idle="400" running="5000" held="100"/>  
  <submit cluster_size="10" max_per_cycle="100" sleep="0.2"/>  
  <release max_per_cycle="20" sleep="0.2"/>  
  <remove max_per_cycle="5" sleep="0.2"/>  
</config></entry>
```

Matchmaking attributes

- Proper site attributes

```
<entry><attrs>  
  <attr name="GLIDEIN_Supported_VO"  
    value="CMS, HCC, GLOW, NEBioGrid, GLUEX"/>  
  <attr name="GLIDEIN_SE"  
    value="bsrm-1.t2.ucsd.edu"/>  
</attrs></entry>
```

- Arbitrary other attributes

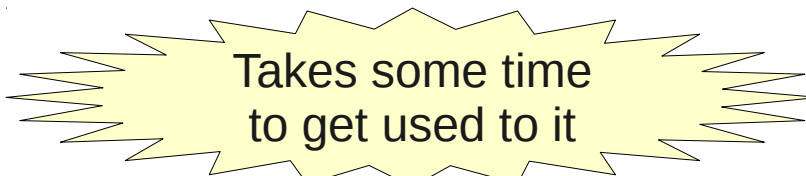
```
<entry><attrs>  
  <attr name="Already_Validated" value="True"/>  
  <attr name="Site_Nickname" value="UCSD_gw2"/>  
</attrs></entry>
```

Final note on sites

- VO Frontend admins trust you to forward their proxies only to trusted parties
 - You should always **think twice before adding** a new site
- The same site may support multiple VOs
 - **If possible, use the same entry** (economy of scale)
 - May not be always possible, though (RSL, attributes, etc.)

A note about reconfigs

- Factory has init.d-like maintenance script
 - `./factory_startup start|stop|reconfig`
- The config file editing unconventional
 - **Cannot** edit the master config file (`glideinWMS.xml`)
 - Must edit a copy of it, typically in `../glidein_bla.cfg/glideinWMS.xml`
 - Then tell reconfig where is the copy
`./factory_startup reconfig ../glidein_bla.cfg/glideinWMS.xml`



Takes some time
to get used to it

Removing sites

- One should never remove an entry from XML file
 - Impossible to add back an entry with the same name (memory effect)
- For long-term removal, just disable it
 - Requires a reconfig (heavy)

```
<entry name="CMS_T2_IT_Pisa_ce1"
        enabled="False"/>
```

- For short-term, put in downtime
 - *./factory_startup down -entry CMS_T2_IT_Rome_ce01*

Pointers

- glideinWMS development team is reachable at glideinwms-support@fnal.gov
- The official project Web page is <http://tinyurl.com/glideinWMS>
- OSG glidein factory at UCSD
<http://hepuser.ucsd.edu/twiki2/bin/view/UCSDTier2/OSGgfactory>
http://glidein-1.t2.ucsd.edu:8319/glidefactory/monitor/glidein_Production_v4_1/factoryStatus.html

Acknowledgments

- The glideinWMS is a CMS-led project developed mostly at FNAL, with contributions from UCSD and ISI
- The glideinWMS factory operations at UCSD is sponsored by OSG
- The funding comes from NSF, DOE and the UC system