



WMAgent @ CERN

Ignas Butenas (CERN), Jose Hernandez
(CIEMAT), Marian Zvada (KIT)

What will be covered?

- Current instances we operate
- Deployment
- Operations

Instances we operate

- WMAgent (for glideinsWMS) is currently installed on those machines:
 - **Instances connected to CERN PRODUCTION factory:**
 - vocms201. With 48 GB of RAM. PRODUCTION instance.
 - vocms120. With 24 GB of RAM. Will be used for CRAB.
 - Vocms121. With 24 GB of RAM, used for integration tests for WMAgent.
 - **Instances connected to FNAL factory:**
 - Vocms202. With 48 GB of RAM. PRODUCTION instance.
 - **Pre-Production instance connected to CERN ITB factory:**
 - Vocms164. VM.

Additional services

- To work with WMAgent you also need Request Manager and WorkQueue services to be running.
- Those two used by DataOps are now deployed on a separate machine (vocms144). This is WMCore 0.7.x.
- 0.8.x series of those services are deployed in CMSWEB. DataOps running scale tests on it. The services in vocms144 will be phased out soon.

Machine preparation

- While submitter nodes are memory hungry, we always request for the machines with **as much RAM we can get**.
- All the machines are installed and managed with the help of **quattor profiles**.
- **Firewall settings:** only SHARED_PORT (4080) used for condor is opened for WAN access. WMAgent itself do not require ports to be opened for WAN access.
- All the other fine tuning of the system is done by following recommendations given by Condor team.
http://research.cs.wisc.edu/condor/condorg/linux_scalability.html

Deployment

- For the deployments we try to **follow standard deployment procedures** as much as we can.
- We use shared account, local shared accounts and sometimes root to run and deploy services.
- Condor on those machines are always deployed with the help of **glideinWMS installer**.
- WMAgent deployment itself is very easy, just by following **deployment and management scripts** provided by developers. Those scripts also follow standard procedures we have defined.

Operations

- Still in early stage, so no very strict rules defined. No schedule for deployment windows.
- **Developers, FacOps, DataOps and Integration team** are working together to deploy, upgrade and keep instances running.
- Integration team deploys and tests the new versions released. Production instances are deployed by FacOps.
- If something goes wrong – FacOps and DataOps reacting and fixing the problems. If more help needed – developers and integration team are contacted.
- Urgent fixes are applied on the fly. New version are deployed to production only after it was tested on pre-prod.