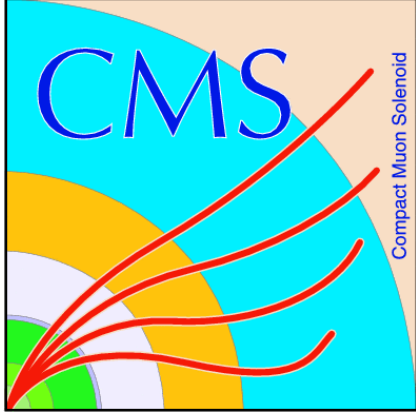

glideinWMS setup at CERN and how we manage our machines

Ignas Butėnas (CERN), Marian Zvada (KIT), Jorge Amando Molina-Perez (CERN)

Current setup

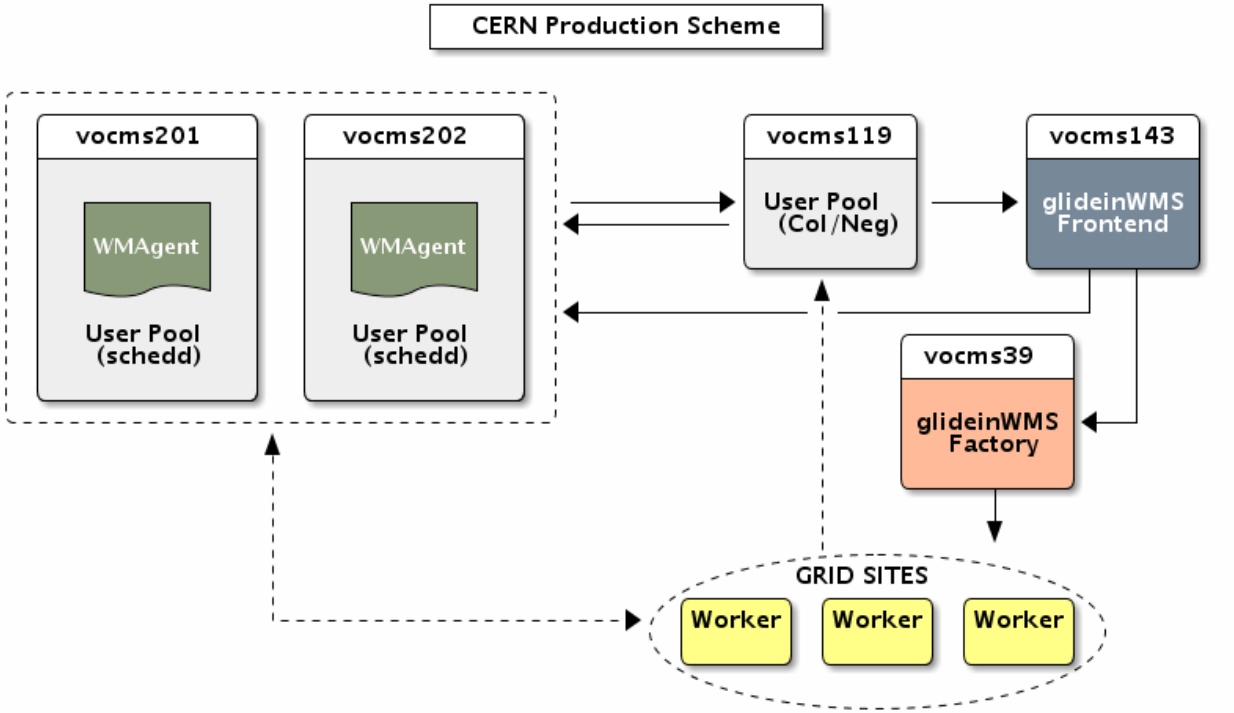

- **vocms201/202**: 8 cores, 48GB RAM
- **vocms119**: 8 cores, 24GB RAM
- **vocms143**: VM, 2cores, 8GB RAM
- **vocms39**: 8 cores, 16GB RAM

- Equivalent to PROD **ITB environment** (but only VMs used) for testing purposes

- All machines (as most of the machines at CERN) are managed by the service which is provided by the CERN IT – **Quattor**.
- Also all the setup is behind the **CERN central firewall**, which is strict requirement at CERN.
- We trying **to follow same procedures and structure** of deployment for all the services we deploy.

Intro

- **VOC:** the person who is the “gates” between experiment (CMS) and CERN IT. The main person who manages all VOBoxes (providing new machines, doing initial setup, tuning current setup, communicates to CERN IT if needed help from them).

- **Service responsible or operator**: person(s) who are responsible for the service and operates it. Has the access (with AFS account, sudo access) to the machines and can modify machine’s settings and profiles.

Account settings

- While we deploying services, we **do not** use AFS accounts.
- But **AFS accounts** (whitelisted ones) **can login to the machines**. For that user needs to be added to the service called E-GROUP (provided by CERN IT).
- **ROOT account usage** at CERN is also very strict, but it is possible to run service as a root if needed. Not logging as root directly of course, but adding users to the machine's profile, who will have sudo access to it.
- **We USE local, shared accounts** to deploy and run services. Local accounts and groups were created and service responsible controls who can login to the machine (managing SSH keys).
- All the settings are managed through the **Quattor**.

Quattor (1/2)

- **Quattor is** a system administration toolkit providing a powerful, portable and modular tool suite for the automated installation, configuration and management of clusters and farms running UNIX derivates like Linux and Solaris.

- **Every machine has it's own quattor profile** (templates) defined in PAN language and stored in CDB.

- All the software need to be installed in to the machine, OS version, filesystem settings, system files and settings, account settings, firewall rules and more, **must be defined in those templates**. Otherwise quattor can just simply **reset** them.

- **NO MANUAL changes on the system is allowed**.

- Only IT, VOC and SR / operators with the permissions **can** change and update templates.

Quattor (2/2)

- All templates are stored in CDB and can be managed through CDBOP session (operator logs in to lxvoadm and starts cdbop session). CDBOP is like subshell and gives you environment and commands to work with templates.

- After applying changes / creating new templates operator submits the changes and they are saved to CBD.

- It is always possible to have separate templates for PROD and ITB for example. So firstly you can commit changes to ITB, test them and then apply to PROD.

- Then IT / VOC / operator uses CCM component to fetch templates to the machine, SPMA component to install new software needed and NCM-NCD component to (re)configure the system, according to the changes have been made in quattor.

Firewall

- Machines are running **local firewall which is a requirement** at CERN. We can't just simply shut down and run without it.
  - We are trying to convince that they can trust our setup(and also that local firewall really affects scalability of our setup), so we want they allow us run without local firewall.

- All CERN machines (almost all) **must run under the CERN central firewall**, which is controlled by CERN IT. Every machine should pass security test before going to production and get ports open for the "outside world".

- Local firewall rules are defined in quattor. As happened already few times – all manual changes are restored by quattor if it is not specified in template. MUST be careful with it.

Alarms and recovery

- **All machines have basic set** of Lemon (service provided by IT) and SLS (another service to monitor the systems) **alarms** defined. If something goes wrong with the machine itself – IT reacts.

- Basic **set of alarms could be extended** and even service alarms could be added. Then it will raise alarms in IT systems and service responsible / IT operator or sysadmins will react. For that service recovery procedures must be defined. *We still miss it in our setup*.

- **Communication** with CERN IT is through ticketing systems, emails and phone if really needed.

Summary

- IT / VOC / Service responsible (operator) controls machine thought the quattor.
- No manual changes on the system are allowed.
- Local firewall controlled by SR (operator), central firewall controlled by CERN IT. Both are requirement.
- Services are deployed by following (as much as possible) the standard deployment procedure defined by CMS.
- Machines and services have defined alarms, which helps to react on system failures and must have recovery procedures defined for each of the alarms.