

VOMS background, options and plans
Compiled by Jeremy w/c 18th June 2012

1) What are the VO Management Server components?

- MySQL database which is a persistent repository for VO membership information,
- VOMS admin which provides the Web UI/services to maintain the VO membership. This requires Apache/Tomcat.
- VOMS server, a daemon process which services the voms-proxy-init requests.

2) What are the associated machine loads?

- The VOMS daemons run 1 process per VO and the resident memory is typically 2.5MB/process. For 23 VOs this means 1.35GB of virtual memory. The daemons run comfortably in a virtual machine.
- Network traffic is low. At peak perhaps 1 connection per second.
- VOMS-admin consists of a single tomcat java process with each VO having its own web application. Requires of order 6GB VSZ and CPU usage is generally low. Background tasks (such as checking membership status) lead to shortlived peaks of activity. The associated mail server load is negligible.
- The MySQL instance requires of order 400MB VSZ and low CPU load.

Currently there are no machine load issues with any of the VOMS components. Replication is generally done to remove a single point of failure.

3) How many VOs per sever?

- The gridpp VOMS instance hosts 23 VOs. Roughly half are active.
- The Glasgow VOMS hosts 6 local VOs. There is low activity.
- The NGS VOMS hosts 33 VOs but only a couple of the VOs are active.

We currently serve about 20 active VOs. The growth rate is now relatively low with a few additional VOs per year.

4) What configurations do other countries/institutes use?

- Most NGIs are considering backup options
- IT, CZ, IBERGRID, GRNET, NL, Ukraine have database and server replication and it has proved reliable for a number of years
- LHC VOs have servers at CERN, BNL and FNAL. The CERN VOMS setup

uses 6 machines (for 14 VOs) but only because Oracle crashes often! The 2 CPU VMs only see use of 1 CPU.

Of the countries already actively replicating, most use mySQL replication for the data and have a single master VOMS-admin.

5) What is replicated and how?

- It is possible to replicate all components but VOMS-admin in particular is best maintained as a single master. Duplication can lead to multiple (confusing) renewal warnings from different nodes (as we saw in May).
- VOMS servers (daemons) are trivial to replicate but need to access a single backend database.

There are 3 main approaches being used to replicate the databases:

- a) MySQL replication (can be secure). Requires single VOMS-admin master and end-to-end authorization. Likely more resource intensive than other methods. It is asynchronous one-way replication. Requires VOMS versions to be in sync.
- b) VOMS-admin replication script (Used by WLCG VOMS). Allows for independent VOMS-admin instances. Very easy to configure pull mechanism but not secure. May not support multiple user certificates. VOMS versions can be different.
- c) PHP based VOMS-admin replication (in EPEL). Supports multi-master replication with low resource usage. No DB sync needed. Can use stunnel. Good in principal solution but least tested approach. Needs more investigation.

6) Implications

- If the VOMS-admin is not replicated it becomes a single point of failure, but other than registration it is only needed for grid-mapfile creation and these are cached indefinitely so impacts are minimal.
- If the VOMS-admin is replicated we potentially encounter complications with registration services if the VOMSRS component is not turned off.
- VOMS client tools connect to the VOMS daemons so the VOMS server hosts need to be configured in UIs, CEs etc. Setting up several additional servers is little more work than one additional server. The basic VOMS installation is relatively straightforward and can run as part of a VM.

7) Conclusions

- There are no machine load issues currently driving the need to replicate VOMS, but there are clear redundancy advantages.
- There are several ways to replicate the VOMS data. Doing so via unencrypted routes is the way several existing implementations do it. It comes down to an issue of privacy. Since there are ACLs on the data implemented by VOMS-admin my inclination is to try and pursue a secure replication route if the loads are not significantly higher.
- We need direct experience of the replication options to make a firm decision on how to replicate the MySQL data.
- We need to understand how many VOs GridPP may support in the future to decide whether there is a need for more than one VOMS-admin instance. If there is then as a minimum we need 2 sites per implementation. If GridPP will only support of order 25 VOs then we need just one VOMS-admin instance and can setup 3-4 servers with the VO data replicated.
- It is possible to run a VOMS-admin instance in each Tier-2 but it is not recommended due to potential complications with machine and user communications.

8) Outcomes

Following discussion Manchester will host the master VOMS instance and additional servers (with replicated data) will be held at Oxford and Imperial.

9) Next steps

- Check VOMS rationalization options
- Plan Manchester transition from Computer Science to Physics
- Understand work to be done to introduce new backup VOMS servers
- Confirm secondary sites (Oxford & Imperial)
- Agree timetable for implementations
- Discuss best option for data replication
- Setup VOMS at additional sites
- Start replication
- Announcements to community
- Configure VOMS clients
- Switchover
- Follow up to include review of VO membership dates and configurations