

Table of Contents

High Availability Implementation for VOMS.....	1
The VOMS Ping script.....	1
The VOMRS Ping script.....	1
Configuration assuming Linux HA.....	1
Resource Switching.....	1
Conclusion.....	2

High Availability Implementation for VOMS

The CERN requirements for VOMS requires a highly available configuration. As discussed in VomsNotes, the high availability functions for the critical VOMS are available as standard in the application using a shared reliable database and an application front end. This does not cover the administration interface which is rated as a Medium criticality service.

A voms-ping function would be required to provide a way for the slave to monitor the master status. voms-ping should take the server name as an argument (like ping) and perform an VOMS application level connection to the server and check that the application can reply. While this does not guarantee the entire application is running, it covers the most common use cases (such as core dump of server or machine motherboard failure).

The VOMS Ping script

This script has been provided here, with the name 'voms-ping' and must be run directly from the server that must be tested, without parameters. Page LCGVomsCernSetup contains the relevant rpm.

Its return value will be 0 if all the server are up and running, and 1 otherwise. In case the result is 1, then the output of the script will list exactly what server had problems, and whether that problem was in the core server or in the admin components.

The VOMRS Ping script

This is in preparation (see table row in VomsServiceMonitor). It should be integrated in LinuxHA and only run on the host which is the master.

Configuration assuming Linux HA

To provide a full high availability function for VOMS,

- Master/Slave set up using Linux-HA and a shared database containing all the state data
- No high availability is provided as part of the VOMSRS interface

Using Linux-HA with a small voms resource script (start/stop/monitor/status) provide this function. The take over time is estimated at around 30 seconds following detection of a failure. There may be a substantial delay between occurrence of failure and detection.

The HA configuration has been implemented as follows

.

In the event of a failure or an operator initiated switch for planned maintenance, the configuration is changed

- Service IP now points to slave server

.

Resource Switching

The VOMS application consists of several components

- VOMS

- VOMS Admin
- VOMRS

The configuration proposed is that VOMS and VOMS Admin should always be running on the master and slave. Only VOMRS should be stopped and re-started when the server switches. The monitoring would also reflect this selection.

https://savannah.cern.ch/bugs/?func=detailitem&item_id=15788#comment2 requests a version-invariant vomrs service. This is needed for leaving file */etc/ha.d/resource.d/gridvoms* untouched across vomrs releases. The LinuxHA set-up is included in rpm CERN-CC-gridvoms-1.1-3 in CDB and on the hosts. The LinuxHA activity is logged in */var/log/messages* both on voms102 and voms103.

Conclusion

For the cost of two machines with small disk space and a highly available database backend, a highly available VOMS implementation can be made which is resilient to network, machine and storage failures.

-- TimBell - 19 Oct 2005

This topic: LCG > VomsWlcgHa
Topic revision: r12 - 2008-08-05 - SteveTraylen



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.
or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback