

Squid monitoring tools

A common solution for LHC experiments.

Alastair Dewhurst (STFC) and Dave Dykstra (Fermilab)

Introduction

Squid[1] is an HTTP caching proxy for the Web. It reduces bandwidth and improves response times by caching and reusing frequently requested web pages. Squids are becoming an increasingly important service provided by sites to the LHC experiments.

Why monitor them?

The LHC applications that use squids support failover. Therefore if a squid fails, a job running on the grid is still likely to succeed as there should either be other working squids or the request can go direct to the source. However if multiple squids are allowed to fail this could result in a domino affect that could disrupt activities at a large number of sites.

Frontier

The Frontier[2] system reads database queries from central servers and caches the results in site squids all over the world. It was developed for the CDF experiment at Fermilab and adapted for CMS. ATLAS later also adopted Frontier. The squid monitoring services were set up on CMS machines.

CVMFS

CVMFS[3] is a network file system optimized to deliver experiment software in a fast, scalable, and reliable way. It has been adopted by the four LHC experiments. It can use the same site squids as Frontier, although some sites choose to set up separate squids. The monitoring requirements for both Frontier or CVMFS squids are the same. Extra monitoring is needed on central service squids.

WLCG Task Force: Objectives

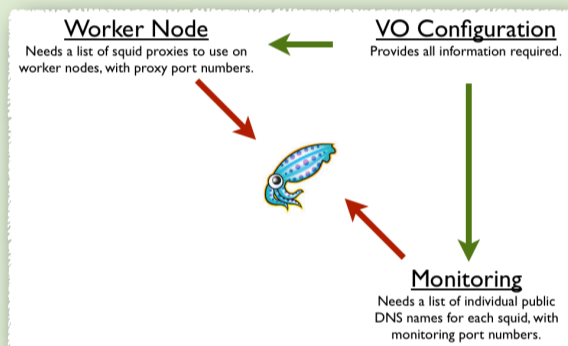
In October 2012 a WLCG task force[4] was set up to look at ways of improving squid monitoring across the WLCG. Its objectives were:

- Make a plan for moving squid monitoring from CMS Frontier management to WLCG management
- Decide how to integrate squid monitoring better with WLCG common operations
- Produce an architecture for a common squid monitoring system that is configured the same way for all VOs

It was beyond the scope of the task force to define how applications discover what proxies to use. The recommendations were published in May 2013.

Squid views

The information required to use a squid is different from the information required to monitor it. The information that can be stored in the standard information systems is different again. The result of this has led to VOs having to provide and maintain all the information needed in order to use and monitor squids.




Information systems

The Grid Configuration Database[5] (GOCDB) is the official repository for storing and presenting EGI topology and resources information. It allows sites to register their services, request they are monitored and allows them to flag when they are in a downtime. The OSG Information Management System [6] (OIM) provides a similar service for OSG sites.


Monitoring Tools

Three main tools have been deployed to monitor squids.


- Site Usability Monitoring[7] (SUM) tests are jobs that are submitted to sites, run on worker nodes and perform a basic task. As the name suggests these tests are primarily concerned with demonstrating that the site is usable so may not inform you if an individual squid is broken. The test is also specific to the type of activity the squid is used for.



- MRTG[8] monitors squids via SNMP and draws graphs showing how much traffic has passed through each interface.
- AWStats[9] monitoring provides much more detailed information about squid usage by analyzing the log files. This has been deployed on squids providing key services.



- Failover Detection[10] monitoring is built on AWStats and currently only covers CMS Frontier. This automatically detects failovers hitting reverse-proxy squids on the central servers and sends emails to shifters/sites about problems.



Recommendations and implementation

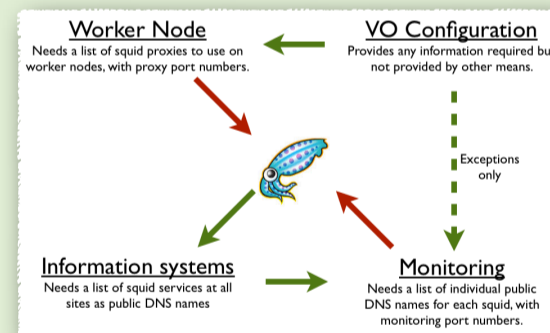
1) It was decided to move the MRTG and AWStats squid monitoring from frontier.cern.ch to a new pair of virtual machines more closely associated with WLCG.

- MRTG was successfully moved and the monitoring can now be found at: wlcg-squid-monitor.cern.ch

2) It was decided that squid services will be registered in the GOCDB & OIM, as publicly available round-robin DNS aliases if there is more than one squid implementing the same service. Sites may have multiple independent squid services registered. GOCDB & OIM will not distinguish between different purposes/applications for squid services.

- This was done and sites can now register their squids in the information services.

3) Additional information needed for squid monitoring beyond that stored in GOCDB & OIM will be maintained by per-VO operations personnel in per-VO files on the squid monitoring servers. These will be combined with the GOCDB & OIM information. The way the new service works is shown below:



4) A SUM test of squid services will be set up, based on non-response to MRTG monitoring probes a few periods in a row. Those squids for which SUM tests are required must be declared in GOCDB or OIM.

5) It was recommended that the Failover Detection monitor be extended to also cover ATLAS Frontier and CVMFS central servers, and an additional SUM test of squid services be added to indicate that failovers are hitting a central server instead of using the site's squid service.

Conclusions

The WLCG Squid monitoring task force was able to fulfill its objectives and made several recommendations for improving the scope and resilience of the squid monitoring service. It will also lead to a reduction in the amount of effort needed by individual VOs to monitor the squids it uses. Some of the recommendations have been successfully implemented while others are waiting to be rolled out at larger scale.

The authors of this poster would like to thank all the members of the WLCG squid monitoring task force, especially Andrea Valassi (CERN), Alessandro Di Girolamo (CERN), Simone Campana (CERN), Luis Emiro Linares Garcia (Universidad de los Andes (CO)) and Barry Jay Blumenfeld (Johns Hopkins University (US)) for their help in coming up with these recommendations. The Fermilab work was supported by the U.S. Department of Energy under contract No. DE-AC02-07CH11359.