

This is Dave's proposal for how to store & maintain squid configuration information for the WLCG. It was used as a basis for discussion and is left here for historical reference.

Terms:

- Squid machine: a computer that runs a squid process
- Squid service: a squid machine or set of squid machines that perform a specific function or functions. Sites may have multiple squid services or a single one. Sites that accept opportunistic grid jobs are encouraged to have a squid service for opportunistic use that is separate from the production service.
- Squid proxy: a squid service used as an http proxy as opposed to a reverse proxy.
- Squid monitoring servers: the pair of machines implementing `wlcg-squid-monitor.cern.ch`

Squid configuration information will take 3 forms:

1. Information System View: a list of squid services at all sites as public internet DNS names. If there is only one squid machine in a squid service, the DNS name can be the primary name or an alias for the machine, but if there are multiple squid machines in the service, the name must be a round-robin alias listing all of the IP addresses of the squid machines (that is, a single address of a hardware load balancer is not allowed). Does not include port numbers.
2. Monitoring View: a list of individual public DNS names for each squid, with monitoring port numbers. The default port is 3401.
3. Worker Node View: list of squid proxies to use on worker nodes, with proxy port numbers. The default port is 3128. May be on a private network. If there are multiple squids, they may be in a round-robin DNS alias or a hardware load balancer. May include backup proxies at other sites, to be tried if all previous proxies have failed. Figuring out how to store this is outside of the scope of this task force, but it is defined here in order to show how it is distinct from the other 2 forms and so the design of the storage for the other two forms does not interfere with this view.

For example: CMS site T2_ES_IFCA has 2 squids on a private network known internally as `squid01prv.ifca.es` and `squid02prv.ifca.es` and externally as `squid01.ifca.es` and `squid02.ifca.es`. So far they have no round-robin names for these, but they would be required to create a private one and a public one, for example `squidprv.efca.es` and `squid.efca.es`. Then the Information System View would be "`squid.efca.es`". The Monitoring view would be `squid01.ifca.es:3401` and `squid02.ifca.es:3401`, and the Worker Node View would be `squidprv.efca.es:3128`.

Details:

1. Store the Information System View in GOCDB and OIM. Maintenance of the information is the responsibility of site administrators.
2. Create the Monitoring View from the Information System View by means of translation files on the squid monitoring servers. The files will be one per VO for each of the two Views, but have the same simple format. Maintenance of the translation files will be the responsibility of operations personnel from each VO. The information system data (that is, site names, squid service names, and VO names) will be read either directly from GOCDB & OIM or via ATP. The CMS VO also will make use of a translation of site names from the Information System View into the CMS site names, either from ATP or directly from the CMS SiteDB. A simple site with only one squid on a public network using the default ports will need no translation entries, but sites with multiple squids, a private worker node network, or non-standard ports will need entries. The translation files will also be able to add whole sites that aren't in the Information System View, but that will be discouraged except for reverse proxies like on Frontier launchpads and CVMFS stratum ones.
3. Whenever Squid-related information is duplicated in more than one source, audits will regularly compare them, and notices will be sent to operations personnel when they don't match. The information can also be stored in different forms (e.g. in AGIS & ATP) but it should come from the above primary sources.

4. The Monitoring View is only needed on the squid monitoring server so it doesn't need to be made available publicly.

Rationale:

1. The responsibilities are very similar to things that are already being done. Storing squid information in GOCDB & OIM is new, but it is very much like other things that site administrators already do and no new functionality is asked of these information systems, just a new field. CMS operations people already maintain a translation file very similar to this (to translate between the Worker Node View and the Monitoring View) on the existing squid monitoring server, and ATLAS operations maintains a python script there that does effectively the same thing.
2. The solution is as simple as possible given the requirements.

Defining how the Worker Node View is stored is outside of the scope of this task force. It is discussed in [HttpProxyDiscoveryProposal](#).

This topic: LCG > SquidMonitoringTFInfoSystem

Topic revision: r6 - 2013-04-22 - DaveDykstra



Copyright &© 2008-2024 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](#)