



Contribution ID: 10

Type: Oral

A solution to distribute software packages at Grid sites using Squid

Monday 12 April 2010 15:30 (15 minutes)

In the current model of Grid computing the software needed by virtual organizations (VO) is stored at the sites in a shared area mounted on NFS and jobs running on the worker nodes access the necessary applications through NFS protocol. An alternative model, based on Squid, is proposed in this article. With this approach, the software packages are pulled from the central VO repository to the sites using http protocol. Squid also takes care of managing the transfers of packages from a central repository at the site to the worker nodes, and optimising network throughput over the LAN.

Impact

This project has impact on both the virtual organization (VO) and the site infrastructure. From the VO perspective this solution provides an alternative to the NFS protocol which presents several limitations, mainly due to its not so good reliability and some compatibility issues with other applications, which result in a non negligible source of job failures. For the site infrastructure, this is an interesting alternative to the NFS file system, which can require a quite complicated deployment, whereas Squid relies on the very standard http protocol. Thank to its flexibility, Squid can be configured to optimise the data flow between client and server to improve performance and caches frequently-used content to save bandwidth. Squid also ensures scalability since more servers can be easily configured in load balance.

Keywords

squid cache http computing farm software distribution

URL for further information

<http://www.pic.es/index.gsp>

Detailed analysis

A prototype has been set up to test the possibility of distributing the LHCb specific software packages to the worker nodes (WN) of the site through http protocol using Squid servers. In this framework Squid acts as http proxy and, in addition to that, provides caching functionality reducing the network traffic from the remote software repository of the VO to a central repository at the site. The test setup has a two level hierarchy of squid servers: The first level of the hierarchy consists of one central cache for all the site. The secondary level runs at the worker nodes where a local cache is provided. When a job lands on a WN it issues an http query to download the software tarball from the VO remote repository. If the request is already cached in the local Squid cache of the WN, the package will be immediately available with a disk to disk copy. Otherwise, the Squid server of the WN will escalate the request to its parent cache at the site. The central Squid server receives the request and, if cached, will send the response back through the site LAN, otherwise will forward the http request to the remote VO web server.

Conclusions and Future Work

The feasibility of the proposed model of software delivery has been proved with a prototype setup at PIC Tier1 for the LHCb VO. Future work will aim to put this setup in production, first with a subset of nodes of the local computing farm, and tune Squid parameters in order to optimise the performance. In case the model proves to be more advantageous than the current one, based on NFS, we will consider to extend it to the whole farm.

Authors: BRIA, Arnau (PIC); Dr LANCIOTTI, Elisa (PIC (Spain)); Dr GRACIANI, Ricardo (Universitat Barcelona)

Co-authors: CASAJUS, Adria (Universitat Barcelona); MERINO, Gonzalo (PIC); DELFINO, Manuel (PIC)

Presenter: Dr LANCIOTTI, Elisa (PIC (Spain))

Session Classification: Infrastructure Tools and Services

Track Classification: Experiences from application porting and deployment