

# DNS load balancing and failover mechanism at CERN

*Monday 13 February 2006 14:00 (20 minutes)*

Availability approaching 100% and response time converging to 0 are two factors that users expect of any system they interact with. Even if the real importance of these factors is a function of the size and nature of the project, today's users are rarely tolerant of performance issues with system of any size.

Commercial solutions for load balancing and failover are plentiful. Citrix NetScaler, Foundry ServerIron series, Coyote Point Systems Equalizer and Cisco Catalyst SLB switches, to name just a few, all offer industry standard approaches to these problems. Their solutions are optimized for standard protocol services such as HTTP, FTP or SSH but it remains difficult to extend them for other kinds of application. In addition to this, the granularity of their failover mechanisms are per node and not per application daemon, as is often required. Moreover, the pricing of these devices for small projects is uneconomical.

This paper describes the design and implementation of the DNS load balancing and failover mechanism currently used at CERN. Our system is based around SNMP, which is used as the transport layer for state information about the server nodes. A central decision making service collates this information and selects the best candidate(s) for the service. IP addresses of the chosen nodes are updated in DNS using the DynDNS mechanism.

The load balancing feature of our system is used for variety of standard protocols (including HTTP, SSH, (Grid)FTP, SRM) while the (easily extendable) failover mechanism adds support for applications like CVS and databases. The scale, in terms of the number of nodes, of the supported services ranges from a couple (2-4), up to around 100. The best known services using this mechanism at CERN are LXPLUS and CASTORGRID.

This paper also explains the advantages and disadvantages of our system, and advice is given about when it is appropriate to be used.

Last, but not least, given the fact that all components of our system are build around freely available open source products, our solution should be especially interesting in low resource locations.

**Primary author:** Mr BAHYL, Vladimir (CERN IT-FIO)

**Co-author:** Mr GARFIELD, Nicholas (CERN IT-CS)

**Presenter:** Mr BAHYL, Vladimir (CERN IT-FIO)

**Session Classification:** Computing Facilities and Networking

**Track Classification:** Computing Facilities and Networking