

Concurrent Adaptative Load Balancing at (@CERN)

Tuesday, 10 July 2018 16:40 (20 minutes)

CERN is using an increasing number of DNS based load balanced aliases (currently over 600). We explain the Go based concurrent implementation of the Load Balancing Daemon (LBD), how it is being progressively deployed using Puppet and how concurrency greatly improves scalability, ultimately allowing a single master-slave couple of Openstack VMs to server all LB aliases. We explain the Lbclient that has been rewritten to be able to use Collectd metrics to determine the status of the node. We also used Go for the Lbclient rewrite as its concurrency features reduce the real time needed for the whole procedure in particular when a node belongs to multiple LB aliases. We explain that the LBD server acts as an arbiter getting feedback on load and health from the backend nodes using SNMP to decide which IP addresses the LB alias will present. While this architecture has been used since long at CERN for DNS based aliases, the LBD code is generic enough to drive other load balancers, so we explain how we investigated using it with HAProxy to provide adaptative responses to load and health monitoring.

Primary authors: REGUERO, Ignacio (CERN); SAIZ, Pablo (CERN); TOTEVA, Zhechka (CERN)

Presenter: TOTEVA, Zhechka (CERN)

Session Classification: Posters

Track Classification: Track 8 –Networks and facilities