

Creating Load-Balanced Services on top of Cloud Infrastructure and Puppet

Vítor Gouveia,
vitor.gouveia@cern.ch

IT-PES-PS

- OpenStack Images
- Availability Zones
- Load Balancing
- Foreman parameters

- There are new images available...

The images that should be used with puppet and **ai-bs-vm** should **not** contain the “CERN” word on their name:

- SLC6 Server - x86_64
- SLC6 CERN ~~Server~~ - x86_64
- SLC6 Server - i386
- SLC6 CERN ~~Server~~ - i386
- SLC5 Server - ...

- All images have a timestamp associated with them
 - E.g. [130920] - 20/09/2013
 - use the latest one (unless you use osrepos_date and want to avoid downgrades... improvements under discussion)

- The ai-bs-vm environment variable to specify availability zones:

AIBS_VMAVAILZONE_NAME

– Ex:

```
AIBS_VMAVAILZONE_NAME="cern-geneva-b"  
AIBS_VMIMAGE_NAME="SLC6 Server - x86_64"  
AIBS_HOSTGROUP_NAME="foo/spare" ai-bs-vm  
higgsbox
```

- An user can specify a particular “location” in which a host should be booted.
 - group together servers in terms of availability, other features could be external network connectivity or redundant power

```
$ nova availability-zone-list
+-----+-----+
| Name           | Status   |
+-----+-----+
| cern-geneva-a | available|
| cern-geneva-b | available|
| cern-geneva-c | available|
+-----+-----+
```

Each of these zones have different local network switch and power inputs. A failure in one hardware component of the cern-geneva-a should not affect VMs running in cern-geneva-b.

- The new CC Wigner: Batch starts using it, availability for other services to be announced.

- To share the load among several VMs behind the same service is usual to use a load balancing Service.
- The DNS Load Balancer (lbd) is the recommended one.
- There are two instances
 - One for Quattor managed machines
 - One for Puppet managed machines
- If you want to create a new load balancer for your service:
 - Do a DNS Load balancer aliases request using the web form in <https://lbweb.web.cern.ch/>
 - The form opens a ticket to CS group that creates the DNS zone and passes the case to PES group to define the alias in the servers.
 - We are working to automate all this and integrate with Openstack (LBaaS).
 - Keys: defines which of the Puppet LB server or the Quattor LB server will be able to update the DNS alias:
 - IT/PES (Puppet managed machine)
 - IT/FIO (Quattor managed machine)

```
class xxxx::loadbalancing {
  class {'lbclient': }
  lbclient::config { 'xxxx config':
    nologin => 'on',
    tmpfull => 'on',
    sshdaemon => 'on',
    xsessions => 'on',
    afs => 'off'
  }
  # LB Alias to which I belong
  lbd::client{ 'xxxx LB alias':
    lbalias => 'xxxx.cern.ch',
  }
}
```

- The elements are set **'on'** will be used for health monitoring of the load balanced alias.
 - Ex:
 - nologin => 'on'
the existence of either files /etc/iss.login or /etc/nologin will be checked so the machine will be removed from the load balanced alias when they exists
 - afs => 'on'
the machine will be removed from the load balanced alias when afs is not running
- More details:
<https://twiki.cern.ch/twiki/bin/viewauth/AgileInfrastructure/DNSIbdAlias>

- No state management in Puppet yet (no SMS)
- For the moment, make sure *nologin* is on (default) and touch `/etc/iss.nologin` to remove a machine, delete the file to add it back
- When the sms replacement (roger) is available, it will handle removal of machines in maintenance from DNSLB

- In principle you cannot mix Puppet and Quattor machines on an alias.
 - Because we generate the config files very differently in the two environements:
 - Using CDBsql in Quattor
 - Using PuppetDB in Puppet
- However, as a migration aid, we have implemented the possibility of adding “hardcoded” Quattor alias members to an alias managed by the Puppet lbd.
 - Using static data in hiera for the LBD service.
 - You have to request the Quattor aliases to be added
 - to it-puppet-hostgroup-ailbd/data/hostgroup/ailbd.yaml
 - For any change, send a request to [DNS Load Balancing](#)

- This requires 3 steps:
 1. Ask the same alias to be created in the Puppet environment
 - General Request to DNS Load Balancing in the service-portal
 - Don't forget to specify a list of Quattor machines you want to keep in the alias, if you want to mix both Puppet and Quattor machines
 - When this is done, the DNS alias is still managed by the Quattor LBD and includes only Quattor machines
 2. Ask netops to change the Key for the DNS alias to the Puppet one
 - I.e. ticket to Network Operations asking to change the key for DNS alias xxxx from IT/FIO to IT/PES
 - When this is done, the DNS alias becomes managed by the Puppet LBD and includes Puppet (and static) Quattor machines
 3. Ask that the alias is removed from Quattor
 - General Request to DNS Load Balancing in the service-portal

- Load Balancing as-a-service
 - We are working to automate all this and integrate with Openstack (LBaaS).
 - No need for tickets anymore, no specific web interface, much easier to create/manage DNS aliases
- Initially, only DNSLB
- With LBaaS we are also considering/evaluating other technologies to complement DNSLB:
 - HAProxy for services that require session affinity/stickyness
 - HAProxy is much more expensive, so is not intended to replace DNSLB
- But options limited by the absence of support for floating IP on the CERN network.

- writefirewall (true by default)
 - firewall rules are managed by puppet
- alarmed (false by default)
 - set to true to enable interactions with the Lemon Alarm System
 - enable sysadmins access
- All the other parameters should be configured through hiera

