

Subject: RE: deployment of an install & config management service (Foreman) based on Puppet
From: SCHAER Frederic <frederic.schaer@cea.fr>
Date: 11/5/13 4:48 PM
To: "hepix-config-wg@desy.de" <hepix-config-wg@desy.de>, Cristina Aiftimiei <cristina.aiftimiei@pd.infn.it>

Hi,

It all depends on your environment...

We @IRFU just deployed puppet on our batch cluster, in addition to the few nodes we already had : we now have 240 managed hosts
Our lesson recently learnt is : the puppetmaster is cpu hungry. Very hungry. Even with puppet (3) , ruby and passenger optimizations.

Our config is :

- One node, SL6.2, serving as
 - o Puppetmaster running under passenger (with a recompiled ruby version, see another mail of mine in this thread about that)
 - o Foreman server
 - o Git (gitolite) server
 - o CA server (the puppet CA)
 - o Puppet fileserver
 - o HTTP/Yum mirror for our Oses, epel and any other repositories we might use (and one apt mirror, for our single debian managed host)
 - o Disk is only relevant for mirrors in our case : the mysql DB is 500MB big after a few weeks, because of reports but these can/are cleaned up
- The nodes specs (now):
 - o Xeon E5-2640 0 @ 2.50GHZ (12 phycical cores)
 - o 32GB mem. 27GB free...
 - o 2TB (raid6) data filesystem, 185G used.

Foreman acts as a web frontend, and connects to "foreman proxy" daemons to run his actions, so in case of overload, you can probably split the puppetmaster/tftp /dhcp/foreman server (the proxy must run on the puppetmaster), but these things are a bit sensitive to set up...

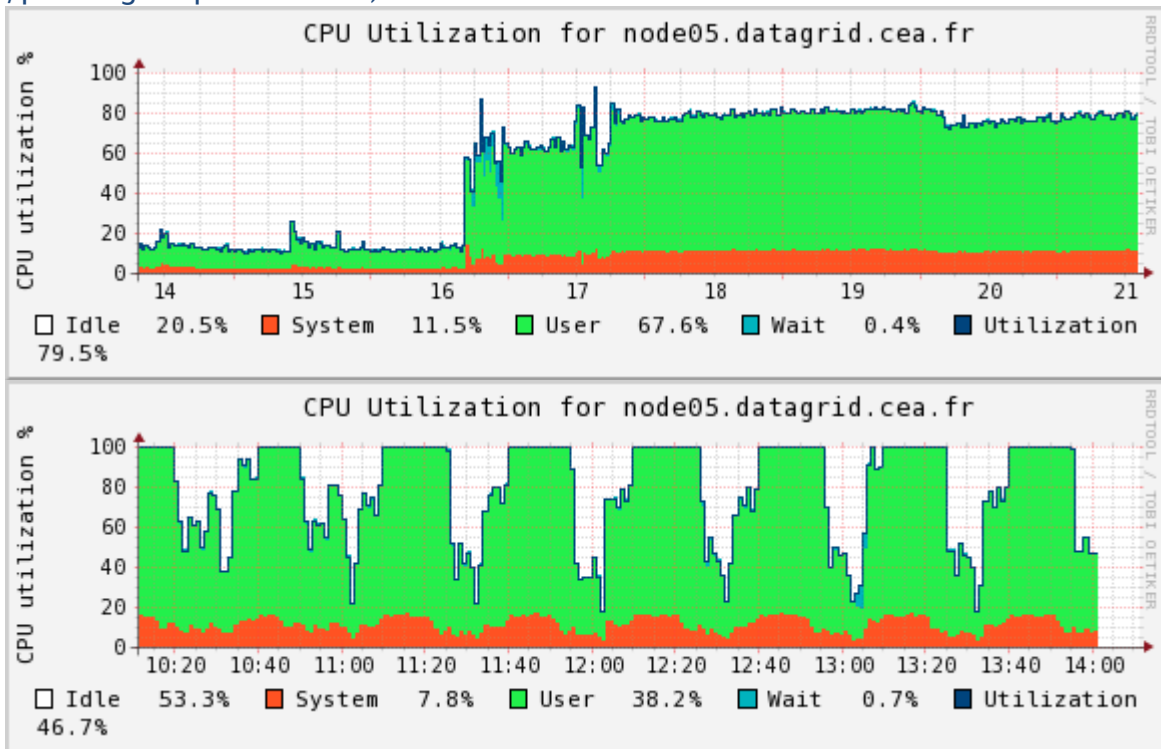
But for a beginning, I would advise to run both/all on same host, as this surely will rule out "same host assumptions" in the code/config.

The node we were using before migrating our batch saw its CPUs melt when we put the 220 additional hosts, it was an old HP DL385G2 (Dual-Core AMD Opteron(tm) Processor 2214 HE).

We did manage to get things running on that host, but despite all my efforts, I could not lower the cpu usage enough so we decided to urgently migrate the puppetmaster to its yet-uninstalled-brand-shiny-new-host we had. My optimizations did cut the compile time from about 60 seconds to about 30 on the HP node, but this was not sufficient. Now, it's about 10 seconds per host.


For the record, we did face huge connection timeouts on foreman/httpd, tcp sendQ/rcvQ issues for same-host tcp connections, and we even went for kernel TCP keepalives changes to improve things.

This is the graph showing the move from 20 to 240 managed hosts (*with* httpd/kernel/passenger optimizations) :



We now have a DELL R620 Xeon E5-2640 0 @ 2.50GHZ (12 physical cores) which seems happy, this is the cpu graph showing the migration effect :

21.10.13 21:36 --- 28.10.13 18:31



Some people prefer not using a puppetmaster at all (they run masterless), this of course removes that issue, but that's another way of working and this implies sharing and securing "your way" the puppet configuration.

You can also implement slave puppetmasters, but we haven't gone much further than making sure this was working on a VM.

Regards

De : Cristina Aftimiei [mailto:cristina.aftimiei@pd.infn.it]
Envoyé : mardi 5 novembre 2013 13:09

À : hepix-config-wg@desy.de

Objet : deployment of an install & config management service (Foreman) based on Puppet

Dear all,

I understood that the topic I'm interested in could be of interest also of this WG, and that probably I would receive helpful advices. Please excuse me if it's not the case.

We are interested in deploying an installation & configuration management system to manage a new cloud infrastructure, based on Openstack, but with possibility to manage more than cloud, probably, in the future. Based on the experiences I saw around we are going in the direction of adopting/deploying Foreman with Puppet.

Regarding the deployment of the Foreman & Puppet Master – Do you have any advice regarding a real (production) deployment? Puppet & Foreman together on the same node, on different nodes, what would be the minimum requirements for the machine/machines where deploying those services (RAM, disk, CPU)

Thank you very much,
Cristina

--

Cristina Aiftimiei

Ist. Naz. di Fisica Nucleare – Padova

Address: via F. Marzolo, 8 – 35131 Padova – ITALY

Phone: +39.049.967.7362

Mobile: +39.3460230488