

D4Science Infrastructure - Task #10988

Task # 10240 (Closed): Install a modern virtualisation framework to host the 'corporate' services

Install oVirt on dlib14x-19x

Jan 19, 2018 02:34 PM - Andrea Dell'Amico

Status:	Closed	Start date:	Sep 28, 2017
Priority:	Normal	Due date:	
Assignee:	_InfraScience Systems Engineer	% Done:	100%
Category:	System Application	Estimated time:	0.00 hour
Target version:	Migrate from Xen to a new virtualisation system		
Infrastructure:	Development, Pre-Production, Production		
Description			
And also configure gluster as distributed file system.			
Related issues:			
Related to D4Science Infrastructure - Task #11709: Post configuration of the ...		Closed	Apr 30, 2018
Blocks D4Science Infrastructure - Task #11396: Move the VM hosted on the Gane...		Closed	Mar 07, 2018

History

#1 - Jan 19, 2018 02:34 PM - Andrea Dell'Amico

- Status changed from New to In Progress

A link that explains a scenario that we can implement: https://ovirt.org/documentation/gluster-hyperconverged/Gluster_Hyperconverged_Guide/
A link that describes the existing ansible roles that can help manage oVirt: <https://ovirt.org/blog/2017/07/ovirt-ansible-roles-an-introduction/>

#2 - Jan 24, 2018 08:36 PM - Andrea Dell'Amico

- % Done changed from 0 to 20

The oVirt admin engine is running, and the first bits of glusterfs are in place too.

#3 - Mar 07, 2018 06:56 PM - Andrea Dell'Amico

- Blocks Task #11396: Move the VM hosted on the Ganeti cluster to the oVirt cluster added

#4 - Mar 13, 2018 09:02 PM - Andrea Dell'Amico

Still no progress. I've something else to try tomorrow, after than I'll install the engine host outside a VM.

#5 - Mar 18, 2018 05:19 PM - Andrea Dell'Amico

Update:

I was able to complete the installation of the oVirt engine in virtualized/HA mode, and the installation of another host.
But when I tried to reinstall that host to add the HA services needed to migrate the hosted engine, that host went to an inconsistent state.
I guess that all the three hosts involved (dlib17x, dlib18x, dlib19x) had too many inconsistent configuration from my past attempts.

So I decided to remove all the virtualisation packages, configurations, and network setup from the three hosts to start from scratch.
Some problems remain:

- dlib17x has currently no connectivity on the main interface, don't know why.
- dlib14x did not come on line again after a reboot (I don't have the console here)

About gluster: I manually created a new volume after the successful installation. As it's possible to use logical volumes as bricks (this is the way used by the automatic configurator, btw) I created a data volume where each brick is composed by three disks. This simplifies the configuration. disks or partitions can be added to a volume after the volume has been created, anyway.

Last, security: gluster can use ACLs to limit the authorized clients, but as firewalld does not permit to limit the sources, we could face problems with all services hosted on the hypervisors. Maybe we could ask the IIT people to create firewall rules for us, that cover both the hypervisors and SAN public IP addresses? We could block everything incoming but ssh (and perhaps that too).

#6 - Mar 18, 2018 06:29 PM - Andrea Dell'Amico

I've reported the required steps and links to the documentation in the internal wiki:
https://support.d4science.org/projects/aginfraplut/wiki/OVirt_Hypervisors_configuration_and_setup

#7 - Mar 22, 2018 04:22 PM - Andrea Dell'Amico

about the firewall rules: it's possible to use the direct rules functionality of firewalld, to add plain iptables rules. Our firewalld role already supports the scenario.

#8 - Mar 22, 2018 08:06 PM - Andrea Dell'Amico

The hyperconverged host installation went well, but I'm now stuck because there should be a configuration section in the administration dashboard that is not present in our. As it's the one that must be used to configure the storage network it's quite important. So I suspended the operations and asked for help in the ovirt users mailing list.

#9 - Mar 23, 2018 04:18 PM - Andrea Dell'Amico

- % Done changed from 20 to 60

I've found the configuration dashboards that was missing. I never tried to click directly on the network name :-(. Now we have four hosts up, the last two still have a problem with the main network interface that Tommaso is investigating. After that we'll have to add all the available disks to the storage, and then we will be able to start running VMs.

#10 - Mar 26, 2018 07:47 PM - Andrea Dell'Amico

All the hosts are up.

#11 - Mar 28, 2018 09:00 PM - Andrea Dell'Amico

The problem with the interfaces on dlib 14x/17x is bad. There's no way to make them working reliably it seems. I tried various options, following the directions that I've found here: https://helpful.knobs-dials.com/index.php/Forcedeth_notes but without any effect. Some files with the options are present on dlib17x and dlib15x under /etc/modprobe.d/forcedeth.conf.

As the interfaces worked well when they were not used as a bridge, we could try to swap the interfaces: use the intel ones on the public network and the NVIDIA (forcedeth) on the storage one. We could try on one of the servers. If it does not work, the safest solution is to only use the Intel cards without any bonding.

#13 - Apr 06, 2018 05:11 PM - Andrea Dell'Amico

- % Done changed from 60 to 40

Another round: we've found a reliable way to configure the network interfaces on dlib1[4:7]x disabling the bonding on the external interfaces, but the global status is now so messed up that we decided to start from scratch again (for the last time, I hope). We also found a broken disk on one of the servers, so we decided to dismiss one of them and use it as spare parts. So the cluster is going to be composed by the hosts from dlib15x to dlib19x.

The gluster volumes will be destroyed as well.

#14 - Apr 10, 2018 06:51 PM - Andrea Dell'Amico

We abandoned the bond configuration on the dlib1[5:7]x servers and after that the setup went well.

And about the firewall configuration, we've found a way that seems compatible with firewalld. The executed commands are:

```
firewall-cmd --zone=public --add-source=146.48.80.0/21
firewall-cmd --zone=public --add-source=146.48.122.0/23
firewall-cmd --zone=public --add-source=10.0.0.0/8
firewall-cmd --zone=public --add-source=192.168.0.0/16
firewall-cmd --zone=public --add-source=146.48.122.0/23 --permanent
firewall-cmd --zone=public --add-source=146.48.80.0/21 --permanent
firewall-cmd --zone=public --add-source=10.0.0.0/8 --permanent
firewall-cmd --zone=public --add-source=192.168.0.0/16 --permanent
```

```
firewall-cmd --set-default-zone=drop
```

Those commands move the two interfaces on the drop zone so that the traffic is blocked by default, while on the public zone the access is granted to the selected networks only.

#15 - Apr 16, 2018 07:34 PM - Andrea Dell'Amico

- % Done changed from 40 to 80

We have completed the storage setup, i guess we are ready to deploy virtual machines there.

#16 - Apr 18, 2018 02:17 PM - Andrea Dell'Amico

- *Start date set to Sep 28, 2017*

due to changes in a related task

#17 - May 02, 2018 05:16 PM - Andrea Dell'Amico

- *% Done changed from 0 to 50*

The already available templates for the virtual machines need some refinement. I'm going to create a new template for Ubuntu 16.04, with a lot of packages removed, python 2 installed, and ovirt-guest-agent correctly installed and configured. If there are no objections I'll remove the Ubuntu 14.04 template.

#18 - May 02, 2018 05:17 PM - Andrea Dell'Amico

- *Related to Task #11709: Post configuration of the oVirt manager added*

#19 - May 02, 2018 05:20 PM - Andrea Dell'Amico

- *Status changed from In Progress to Closed*

- *% Done changed from 50 to 100*

This can be closed, the remaining tasks pertain to [#11709](#)