Making Nova Infrastructure HA/Fault Tolerant

This discussion will focus only on infrastructure and has nothing to do with guests.

http://etherpad.openstack.org/making-nova-components-ha

Goals of discussion

1. Discuss changes that are needed for each of the components.
2. Not solve the issues today.
3. Leave with ideas for blueprints or topics of discussion on the mailing list.

My reference implementation based on where nova is today, Cactus release.

MySQL Active/Passive setup

- Two nodes
- DRDB or shared DAS for file system
- Heartbeat and/or Pacemaker to control resource fail over

Wanted to use stock MySQL included from distributions

- No patches
- No recompiling
- No complex setups

RabbitMQ Active/Passive

- Cluster != HA
- Active/Passive
  - DRDB or DAS for shared filesystem
  - Heartbeat and/or Pacemaker to control resource fail over

RabbitMQ seems limited in scaling options

http://www.rabbitmq.com/pacemaker.html
http://www.rabbitmq.com/clustering.html

Managers Active/Passive

1. Nova-API
    - Should have option to run with real web server
2. Nova-Network
3. Nova-Scheduler
4. Nova-compute
    - Hard to make HA as its only one per node

NTT’s Goal

- High Availability of Nova-Network Service

Use of Open Source Software’s

- Linux-HA (Heartbeat)
- Keepalived

Design Approaches

Active/Standby using Linux-HA

- In this approach, we are trying to switch over from primary network server to standby network server using heartbeat as a HA software.

Detecting failover criteria

Failover will be detected by the peer system in following cases

- Heartbeat service is down
- Network failure(ethernet interface on which the udp packets are sent by the heartbeats is down i.e. eth0, eth1 and so on)P
- Powered off
- nova-network service is not running on primary network server
- syslog has reported any fatal error

Configure Heartbeat

You must configure three files to get heartbeat to work: authkeys, ha.cf and haresources.

- /etc/ha.d/ha.cf - the global cluster configuration
- /etc/ha.d/authkeys - a file containing keys for mutual node authentication
- /etc/ha.d/haresources - This file specifies the services for the cluster and who the default owner is

Cluster Configuration

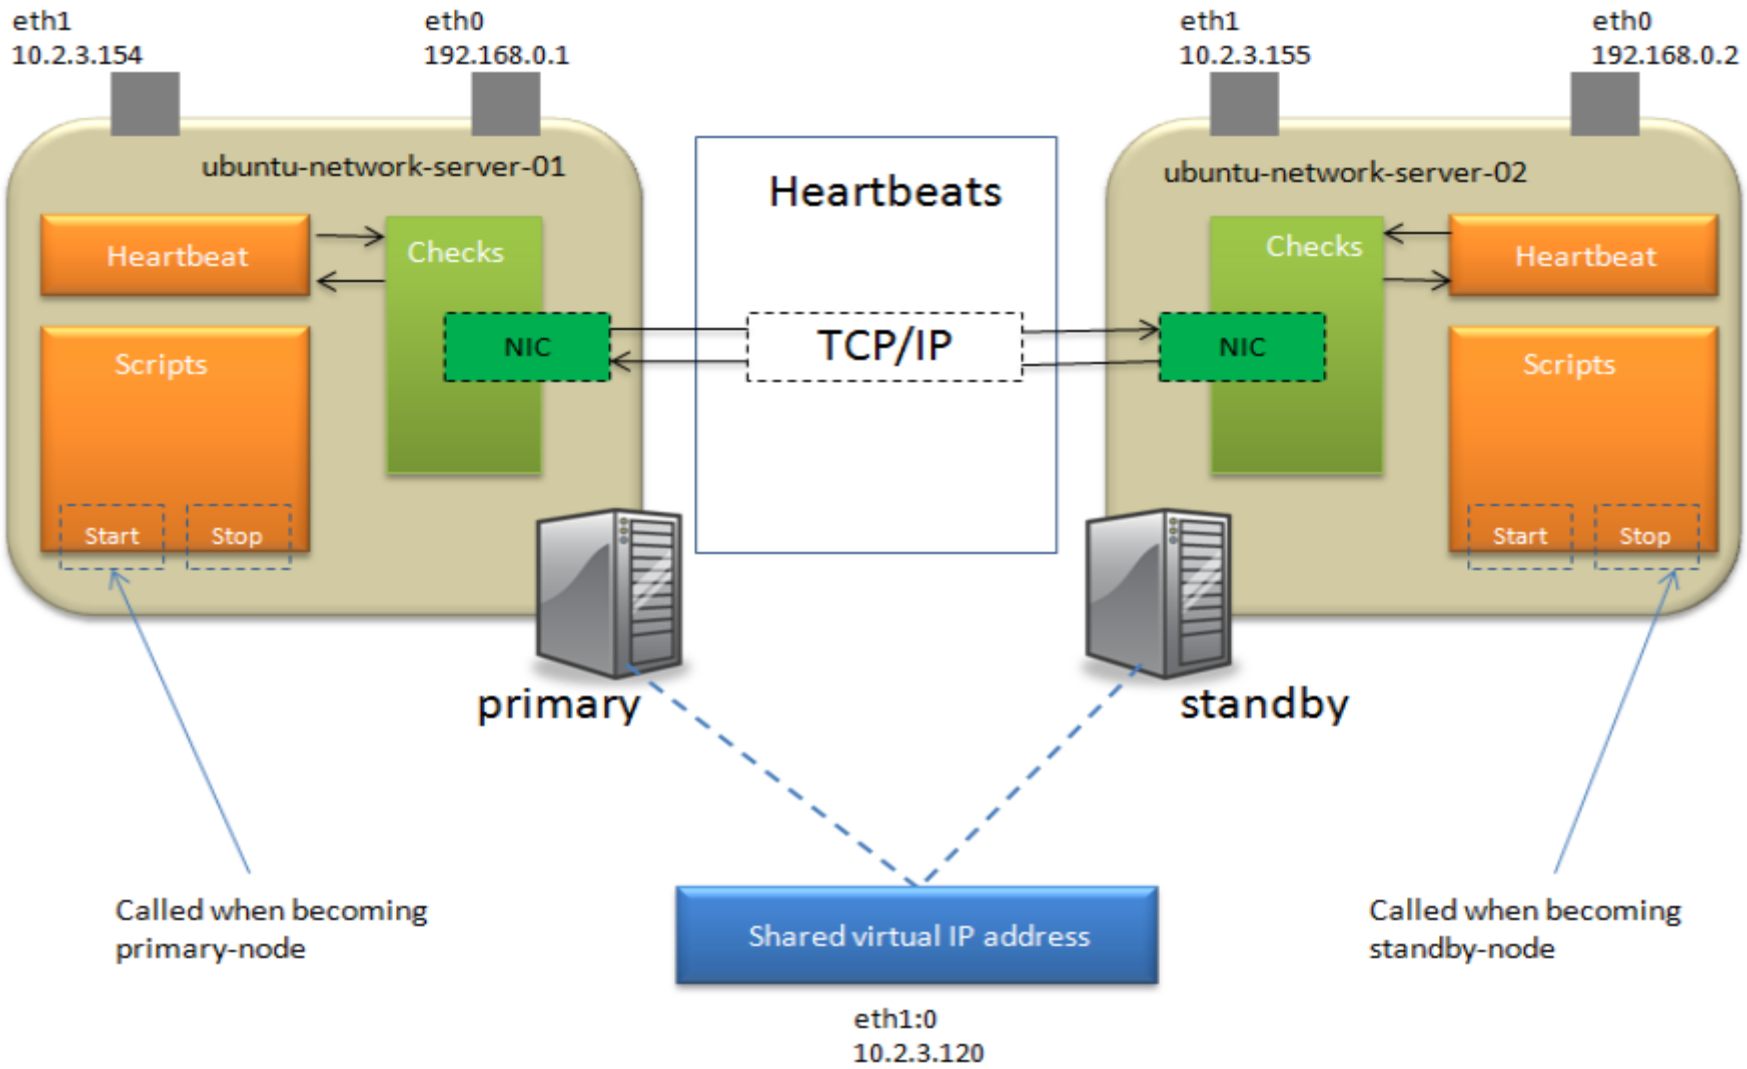

Nova Network Configuration

Nova.conf

- Routing source IP address should be same on both primary and standby network server.

--routing_source_ip=10.2.3.120

- Host flag settings should be same on both primary and standby network server. All network messages communication happens using rabbitmq on topic network.<hostname>. By default hostname is the hostname of the server on which nova-network server is running if host flag is not provided. After failover we don’t want hostname to differ from the primary network server so it should use host flag to get rid of this problem. Most importantly, administrator will have to use the same host to create networks using nova-manage.

--host=somehostname

Test HA

- Start the heartbeat service on the primary and then on the standby network server
- After the heartbeat starts successfully, you should see a new network interface with the IP address that you configured in the haresources file on the primary network server.
- Run new VM instance and associate floating ip address 10.2.3.104 to an instance.
- Ping to 10.2.3.104 from your laptop. Ping should be successful.
- Now simulate failover by simply stopping heartbeat on the primary network server
- You should see nova-network service will be running on the standby network server and nova-network service will be stopped on the primary server.
- Verify that you can still ping to 10.2.3.104. If ping is successful that means you are successful in implementing HA of the network server.

Disconnection Time

- The disconnection time should vary depending upon deadtime parameter you have set in the ha.cf. To find out the actual network disconnection time, I have taken 3 samples using tcpdump (tcpdump -i eth1 -n icmp) with different deadtime settings. But shockingly I got unexpected results.

Dead time (in	Ping Reply Stop	Ping Reply Response	Time difference (in Seconds)
4	15:20:28	15:20:56	28
4	15:21:47	15:22:14	27
4	15:23:46	15:24:26	40
3	16:23:07	16:23:54	47
3	16:25:16	16:26:03	47
3	16:26:59	16:27:46	47
2	16:40:54	16:41:19	25
2	16:41:56	16:42:20	24
2	16:43:02	16:43:26	24

Problems

The failover happens exactly after 4/3/2 seconds but ARP table in the neighboring machines are not updated with the new MAC address of the standby network server. It takes lot of time to update the ARP table if we don't broadcasts a gratuitous ARP to inform the neighboring machines about the change of MAC address for the IP.
I tried broadcasting a gratuitous ARPs to router (in our setup 10.2.3.1 is the router) just before starting the nova network service and recorded the disconnection time again.
ARP command
arping -b -c 1 -s <floating Ip address> -I <public interface> <router ip address>
$arping -b -c 1 -s 10.2.3.104 -I eth1 10.2.3.1

Amazingly, this time the disconnection time recorded was in the range of 4-5 seconds for 4/3/2 deadtime parameter.

Possible Changes to the nova-network

We may need to send gratuitous ARPs for following IP addresses:-
- Routing source IP address
- Floating IP Addresses
  - Gateway address of each vlan(which is assigned to each bridge)

Instead of handling these gratuitous ARPs externally when the failover happens, we think we should make changes to the nova-network source code to address all of the above 3 points. Even if we modify the source code to send the ARP messages, there will be a down time of few seconds. However, this modification sound reasonable than the current implementation in terms of HA.

Note: The down time will depend on many factors like number of VM instances, vlans etc. For POC purpose, we have used 2 VM instances and one Vlan.

Discussion

Database

- Ideas?
- Concerns?
- replacements?

Discussion

RabbitMQ

- Ideas?
- Concerns?
- replacements?

Discussion

Nova-network

- Ideas?
- Concerns?
- Multi-master?

Discussion

Nova-scheduler

- Ideas?
- Concerns?
- Multi-master?

Discussion

Nova-api

- Ideas?
- Concerns?
- run multiple instances with lb in front?