

Ganeti playbook

How to use your cluster.

© 2010-2011 Google

Use under GPLv2+ or CC-by-SA

Some images borrowed/modified from Lance Albertson and Iustin Pop

Outline

- Adding nodes
- Organizing nodes
- Repairs

Adding nodes

- Install and configure the node (or let your configuration management do the job)
- Add the node with `gnt-node add <name>` (asks for root passwords if key authentication is not possible)

Organizing nodes

In a big cluster you want to organize nodes into groups. Ganeti will make sure instances' primary and secondary nodes are in the same group.

Rule of thumb: One group per subnet:

```
# gnt-group add group2
# gnt-group rename default group1
# gnt-group assign-nodes group2 node20 node21 node22 ...
# gnt-instance change-group --to group1 instance_name
```

Recovering from master failure

```
# on a master candidate
gnt-cluster master-failover

# use --no-voting on a 2 node cluster
```

(A linux-HA experimental integration is present in 2.7)

Preemptively evacuating a node

We can remove instances from a node when we want to perform some maintenance.

Drain, move instances, check, set off-line:

```
gnt-node modify -D yes node2    # mark as "drained"
gnt-node migrate node2          # migrate instances
gnt-node evacuate node2         # remove DRBD secondaries
gnt-node info node2             # check your work
```

```
gnt-node modify -O yes node2 # mark as "offline"
```


It is now safe to power off node2

Recovering from node failure (1)

```
# set the node offline
gnt-node modify -O yes node3


# use --auto-promote or manually promote a node
# if the node was a master candidate.
```

(This step can also be automated using linux-HA)



Recovering from node failure (2)


```
# failover instances to their secondaries
gnt-node failover --ignore-consistency node3
# or, for each instance:
gnt-instance failover \
  --ignore-consistency web
```



Recovering from node failure (3)

```
# restore redundancy
gnt-node evacuate -I hail node3
# or, for each instance:
gnt-instance replace-disks \
  {-n node1 | -I hail } web
```

(The autorepair tool in Ganeti 2.7 can automate these two steps)



Re-add an off-lined node

After a node comes back:

```
gnt-node add --readd node3
```

Then it's a good idea to rebalance the cluster:

```
hbal -L -X
```

Maintenance

Shutting/Starting down all instances:

```
gnt-instance stop|start --all [--no-remember]
```

Blocking/Unblocking jobs:

```
gnt-cluster queue [un]drain
```

Stopping the watcher:

```
gnt-cluster watcher pause <timespec>|continue
```

Cluster Shutdown

Graceful shutdown before powering off nodes:

```
gnt-cluster verify
gnt-cluster queue drain
gnt-cluster watcher pause 6000
gnt-instance stop --all --no-remember
gnt-job list --running # Check if jobs have completed
```

Emergency shutdown (faster):

```
gnt-instance stop --all --no-remember
```

Cluster Re-start

After a graceful shutdown, return the cluster to service:

```
gnt-cluster queue undrain
gnt-cluster watcher continue
```

watcher will restart all instances in 10-20 minutes:

```
gnt-cluster verify
```

Ganeti upgrades

From the master node:

```
alias gnt-dsh=dsh -cf /var/lib/ganeti/ssconf_online_nodes

# Stop Ganeti
gnt-dsh /etc/init.d/ganeti stop
# Now unpack/upgrade the new version on all nodes. eg
gnt-dsh apt-get install ganeti2=2.7.1-1 ganeti-htools=2.7.1-1
# Now upgrade the config and restart
/usr/lib/ganeti/tools/cfgupgrade
gnt-dsh /etc/init.d/ganeti start
gnt-cluster redist-conf
```

Other resources

The Ganeti administrator's guide

<http://docs.ganeti.org/ganeti/current/html/admin.html>

Conclusion

Questions?

© 2010-2011 Google

Use under GPLv2+ or CC-by-SA

Some images borrowed/modified from Lance Albertson and Iustin Pop

