

Monitoring

- What to monitor?
- What facilities does Ganeti offer to help monitor

© 2010-2011 Google

Use under GPLv2+ or CC-by-SA

Some images borrowed/modified from Lance Albertson and Iustin Pop

What to monitor

- "Right Now" things (pagable)
 - node down
 - instance down
 - cluster status
 - DRBD issues
- "Historical" things
 - Capacity
 - Utilization

Monitoring Clusters

The master IP should:

- ping
- answer to SSH
- respond to RAPI calls ("version" is a good 'no op')
- "gnt-cluster verify" output should not contain the word "ERROR"
 - `grep -q ERROR: /srv/ganeti/info/$CLUSTERNAME/current`
 - FIXME is the above file generated by an open source package?

Monitoring Nodes

- nodes should be pingable, ssh'able
- load average:
 - Xen: shouldn't be above 2.0 for very long
 - KVM: ...more complicated...
- DRBD issues
 - Nagios monitoring for DRBD:
 - <http://code.google.com/p/ganeti/wiki/DrbdDevicesMonitoring>

Monitoring Instances

- Do you provide monitoring for your users or are they responsible for it?

- Instance owner does it:
 - Do they have the skill?
 - Can you give them access to your monitoring system?
- Monitoring on behalf of the owner:
 - Best way to know if you are providing good service
 - "Fixed before the user notices"
 - Owner should be able to temporarily disable paging if instance will be down intentionally
 - Who gets paged? you or instance owner or both?

Historical Monitoring

Keep long history of utilization for capacity planning, budgeting, and troubleshooting.

- What to collect:
 - Hourly (or more) history of disk I/O, RAM and network utilization
 - Daily/weekly history of # nodes, # instances, # instance create/delete
 - Uptime statistics
- Troubleshooting:
 - Did problem appear when disk I/O reached a certain level?
 - Is the current network utilization heavy or light?
- Capacity and budget planning:
 - When will resources be exhausted? (don't be surprised)
 - Based on current growth, what do we need to buy?
 - How much money have we saved by using virtualization?
 - How many users/machines do we have?

When to start monitoring

"It isn't a service if it isn't monitored. If there is no monitoring then you're just running software"

- **Start now.**
- Monitor before service goes "live" to users
- Until it is monitored, its "beta"
- You can't collect history in the future
 - ...even if this means recording number of instances & nodes in a spreadsheet once a month

Conclusion

Questions?

© 2010-2011 Google

Use under GPLv2+ or CC-by-SA

Some images borrowed/modified from Lance Albertson and Justin Pop

