

Ganeti 2.7 and beyond



Evolutions of a cluster virtualization manager.

Guido Trotter <ultrotter@google.com>

- Google, Ganeti, Debian


© 2010-2013 Google

Use under GPLv2+ or CC-by-SA

Some images borrowed/modified from Lance Albertson and Justin Pop

Outline

- Community news
- New features
- Medium age features
- Old features



New logo



Community News

- Time Based release process, one freeze every 3 months.
- Synnefo: a private cloud based on Ganeti. Tomorrow in the Cloud devroom at 2pm.
- RPMs for centos and RHEL, community maintained.
- In general, bigger community.

New os definitions

- snf-image: an os definition that installs from inside a VM
<http://www.synnefo.org/docs/synnefo/latest/snf-image.html>

- `ganeti-os-defs`: two improved os definitions for linux and non-linux OSes
<http://sourceforge.net/p/ganeti-os-defs/home/Home/>
- `ganeti-instance-image`: first image based instance installer
<https://code.osuosl.org/projects/ganeti-image>

What's new in Ganeti

- Network management
- External storage
- Exclusive mode node parameter
- Hroller
- Better KVM support
- Linux HA
- Openvswitch

Network management (2.7)

`gnt-network` allows to define networks and their parameters (`netmask`, `gateway`, ...) and to associate them to nodegroups (which defines the physical ways in which the nodes have the network available).

A new `nic` parameters available through `gnt-instance` makes it possible to connect an instance nic to a network, thus inheriting the parameters, and doing IP allocation if necessary.

(contributed by `grnet.gr`)

External storage (2.7)

External storage providers can be configured to seamlessly attach and detach instance disks from nodes, and to create them, thus allowing the use of any storage system in conjunction with Ganeti.

(contributed by `grnet.gr`)

Exclusive mode (2.7)

A new node parameter (that should be changed at cluster or nodegroup level) allows changing the node allocation policy so that resources are reserved for instances exclusively, instead of shared. Currently this mode acts on how disks are allocated on lvm storage (each physical volume is assigned to only one instance, rather than shared by striping), but in the future it will be also used for different resources.

This functionality is meant for more predictable workloads, to minimize contention between instances for node resources.

Hroller (2.7)

This new htool allows calculating how to best perform parallel maintenances on a cluster. It currently support offline maintenances, and will be expanded soon for online ones.

Better KVM support (2.7)

New parameters have been added to better tweak an instance's SMP CPU configuration, support extra hardware, and even pass arbitrary parameters to the `kvm` command line.

Linux-HA resource agents (2.7)

Experimental OCF agents are available, to configure ganeti services as a resource in Linux-HA.

- Master resource: perform master failovers should the master node fail
- Node resources: offline a failed node

This can be used in conjunction with our upcoming autorepair system to fully automate your cluster.

Openvswitch support (2.7)

Very basic support for connecting to an openvswitch network, under KVM.

Just connects the instance interface to the openvswitch port described in the link parameter. More work is needed to support extra parameters such as rate limiting or explicit vlan tagging.

What's coming next

- Monitoring agent
- Exclusive mode for CPUs, spindles, memory
- Rolling reboot support
- Better support for non-lvm storage
- Better hotplug support
- Autorepair (!)

Ganeti 2.6

Stable since mid-2012:

- Instance policies
- Memory ballooning
- CPU Pinning
- RBD support (ceph)
- Disk parameters
- OVF converter

Instance policies (2.6)

Allow limits at cluster and nodegroup level on instance sizes and hardware sharing rates. This permits better forecasting for cluster free space, makes sure users with resize-instance powers don't exaggerate, and permits distinctions by underlying hardware type.

Memory ballooning (2.6)

Allow dynamic resizing of instance memory.

- `MAXMEM` and `MINMEM` instead of `MEMORY`
- Manual resizing of memory without instance shutdown
- Automatic resizing of memory by Ganeti (for now at startup and migrate)

In the future automatic resizing to make space to other instances.

CPU Pinning (2.6)

Allow specifying which CPUs an instance should run on, for full control of your hardware, and better instance-specific guarantees.

Rados Block Device support (2.6)

Allow instances based on the block device behind the ceph file system.

(contributed by grnet.gr)

Disk parameters (2.6)

Allow configuration of some storage parameter that were constant or just build-time. Examples are:

- DRBD resync algorithm and parameters (e.g. speed)
- the default VG for meta-data volumes for DRBD
- number of stripes for LVM (plain disk template)
- the RBD pool

OVF converter (2.6)


Exporting of Ganeti instances to the OVF format, and importing of OVF images into Ganeti.

Customizable master ip setup (2.6)

Rather than forcing the master ip to be "just" in ip added to the master-netdev interface, allow a script to do any arbitrary action to set the node as a master. This allows turning up a new interface, using policy routing, or even a routing daemon to direct traffic to the master node.

What's "old" in Ganeti

- DRBD and LVM storage backends
- KVM and Xen hypervisors support, live migration
- instance allocation and cluster balancing
- remote api and cli control
- IPv6 support
- SPICE support



Ideas

- Making the virtualization entry level as low as possible
 - Easy to install/manage
 - No specialized hardware needed (eg. SANs)
 - Lightweight (no "expensive" dependencies)
- Scale to enterprise ecosystems
 - Manage simultaneously from 1 to ~200 host machines
 - Access to advanced features (DRBD, live migration)
- Be a good open source citizen
 - Design and code discussions are open
 - External contributions are welcome
 - Cooperate with other "big scale" Ganeti users

Conclusion

- Check us out at <http://code.google.com/p/ganeti>.
- Or just search for "Ganeti".
- Try it. Love it. Improve it. Contribute back (CLA required).

Questions? Feedback? Ideas? Flames?

© 2010-2011 Google

Use under GPLv2+ or CC-by-SA

Some images borrowed/modified from Lance Albertson and Justin Pop

