

Compute Cells for Grizzly

Chris Behrens

chris.behrens@rackspace.com

OpenStack Summit
San Diego, CA USA
October 15th, 2012



Overview

Overview

- Scalable/distributed nova
- Separate DB and AMQP broker per cell
- New nova-cells service
 - Message routing (broadcast and direct-to-cell)
 - Scheduling
 - Separate from host scheduling
 - Periodic broadcasts from children with capacity and capabilities
- Inter-cell communication via RPC
- Cells connected in a hierarchical tree
 - API cell(s) at the top
 - No nova-api in child cells

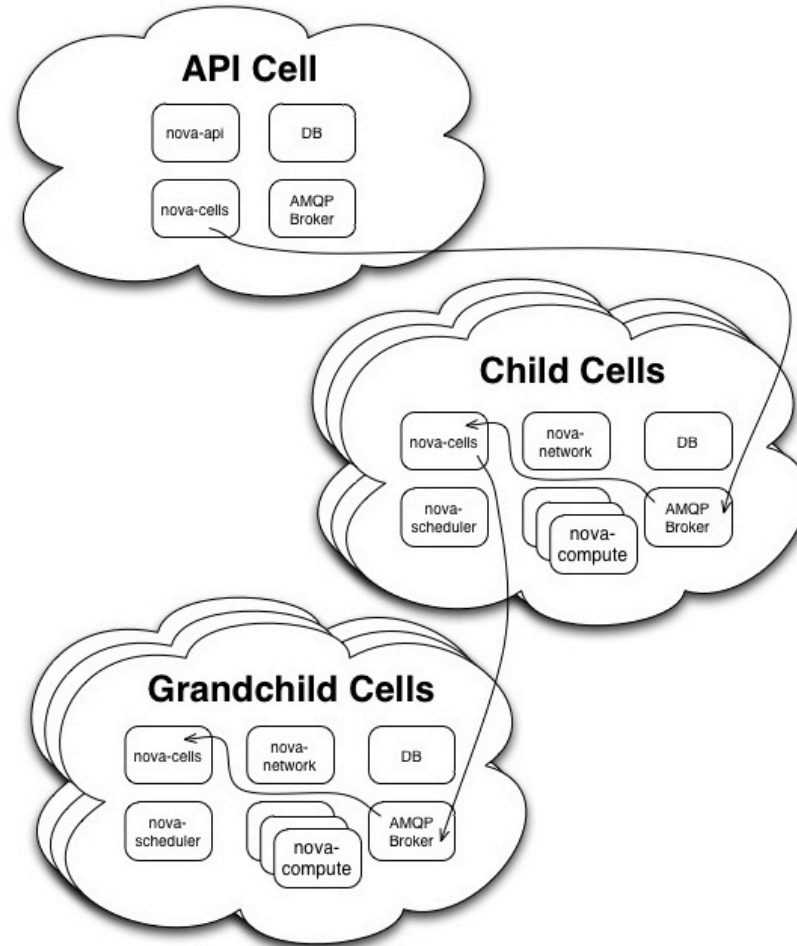
What data is where?

- Each cell has its own DB
- Child cells only have data for its own cell
 - Same data as a normal nova deploy
 - Instances, services, compute_nodes, etc
 - Quotas should be disabled (NoopQuota class)
- API cells have subset of all child cell data
 - All instances
 - Quotas
 - Migrations (due to how Quota code works)
- API cells know nothing about hosts, hypervisors, etc

Minimal services

API cell	Child compute cell
AMQP broker (for RPC)	AMQP broker (for RPC)
DB	DB
nova-cells	nova-cells
nova-api	nova-scheduler
	nova-network
	nova-compute

Routing



Status

Status

- Been running in a large-scale production environment
 - Since 8/1/2012 within Rackspace® Cloud Servers
- Code is available publicly
 - http://github.com/comstud/nova/tree/cells_service
- Didn't make Folsom ☹️
- Needs to merge into trunk ASAP!

What works?

- Building instances
- Actions on instances
- Quotas
- API extensions with cells support
 - Hosts
 - Hypervisors
 - Instance_usage_audit_log

What's not implemented?

- Security groups
- Host aggregates
- Volumes
- Availability zones

Issues

- New ComputeCellsAPI class
 - Overriding of current compute API class not ideal
- Quotas
 - Needs migrations in API cell
 - Commits quotas early in API cell
- Cells manager class
 - Code is confusing
 - Should be broken up into more classes?
- Message ordering between cells
 - Race conditions
 - Amount of messages?

Discussion

Discussion

When can this go to trunk?

- Within next 2 weeks? (I hope)
- Code is proven, as it's running in a production environment
- Getting it into trunk allows us to iterate on it more easily

Discussion

Miscellaneous

- Security group support
- Host aggregates
- ...

THANK YOU

Chris Behrens
chris.behrens@rackspace.com



RACKSPACE® HOSTING | 5000 WALZEM ROAD | SAN ANTONIO, TX 78218
US SALES: 1-800-961-2888 | **US SUPPORT:** 1-800-961-4454 | **WWW.RACKSPACE.COM**