

# Compute Cells

Chris Behrens

[cbehrens@codestud.com](mailto:cbehrens@codestud.com)

comstud@IRC

<http://comstud.com/FolsomCells.pdf>

[http://etherpad.openstack.org/  
FolsomComputeCells](http://etherpad.openstack.org/FolsomComputeCells)

# Overview

- Scalable/distributed nova
  - No need for complicated DB/rabbit clustering technologies
- Separate DB and AMQP broker per cell
- New nova-cells service
  - Message routing (broadcast and direct-to-cell)
  - Scheduling
- Inter-cell communication via RPC
- Cells connected in a tree
  - Supports multiple parent cells

# Cells, not zones

- Cells replace zones (zones removed in Essex)
- No API service in child cells
- API cells at the top know about all instances 'under' it.
- Cell scheduling separate from host scheduling
- Cell scheduling based on periodic broadcasts from child cells
  - Capabilities
  - Capacity

# 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

# API cell differences

- Override default compute API class
  - `--compute_api_class=nova.compute.cells_api.ComputeCellsAPI`
- `rpc.cast/calls` to scheduler/compute hosts go to nova-cells service
- Separate DB server
  - Separate view of instance state
- State checking/setting in API cell as well as in child cell

# Compute cell differences

- Instance info pushed up on DB updates
  - instance\_update
  - instance\_destroy
  - instance\_fault\_create
- Periodic pushing of instance info to parent cells
  - In case of lost messages, etc

# Status

- WIP review up
  - <https://review.openstack.org/#/c/4378/>
  - Review is a bit behind (will get latest code up this week)
- Basic functionality works
  - build, rebuild, reboot, resize, snapshot, etc.
  - State changes in child cells make it to API cell
- Cell scheduling with filter/weighing capabilities
  - Basic RAM-based capacity weigher
- New cells API extension
  - Add/remove cells
  - Force syncing of instance information in children

# Demo

- Build an instance
- Reboot the instance
- Build another instance



# To Do

- Get latest code up for review
- Get initial work merged into master
- Re-factorings needed in compute API
- API should have its own DB/view of state even for non-cells
  - New tables specific to API
- Refactor host and cells scheduler to share code
  - Filtering and weighing

# Questions/Discussion