

Scaling the Cloud

Journey from Cells-v1 to Cells-v2

Surya Seetharaman

17/05/2018


Openstack at CERN

Statistics from Scaling Point of View



Nova Without Cells

Without Scaling - Original Design



How do we Scale... ?

Compute Service is scaled using Nova Cells


→ What are Nova Cells ?

1. A functionality to scale the OpenStack compute cloud.
2. Hosts in the compute cloud are divided into groups i.e cells.
3. Each cell has its own DB and MQ.
4. Cells are in a tree-like structure having a top 'api' cell and child cells.

→ Why do we use Cells ?

1. A strategy for scaling.
2. Failure Isolation.
3. Provides Elasticity/Flexibility.

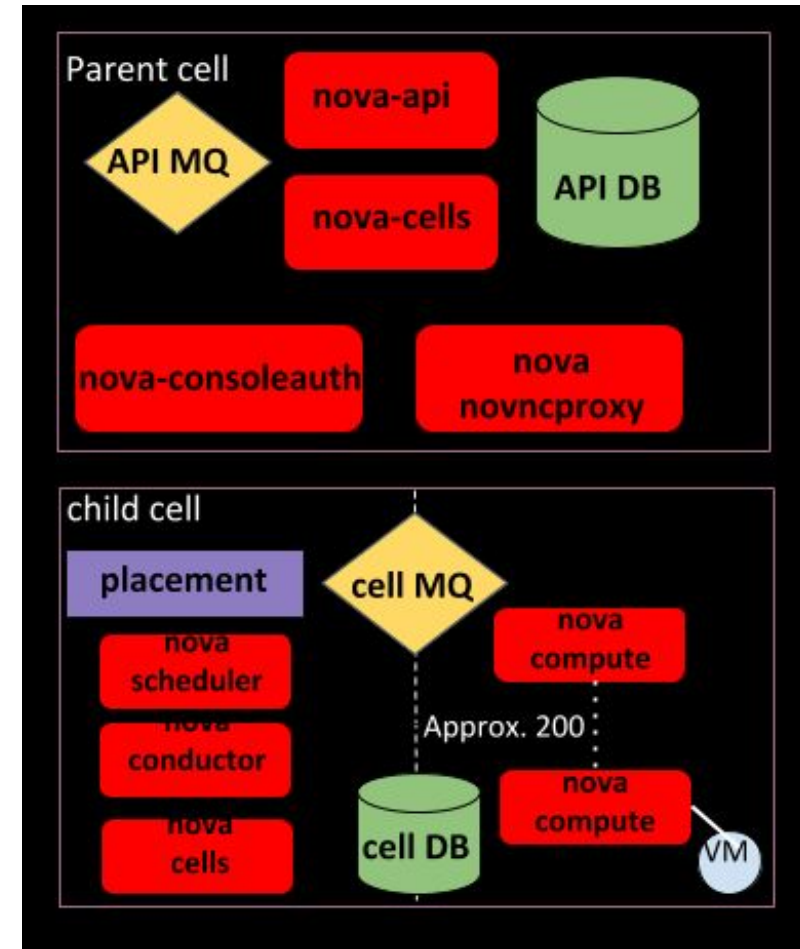
Nova with Cells-v1



Limitations of Cells-v1

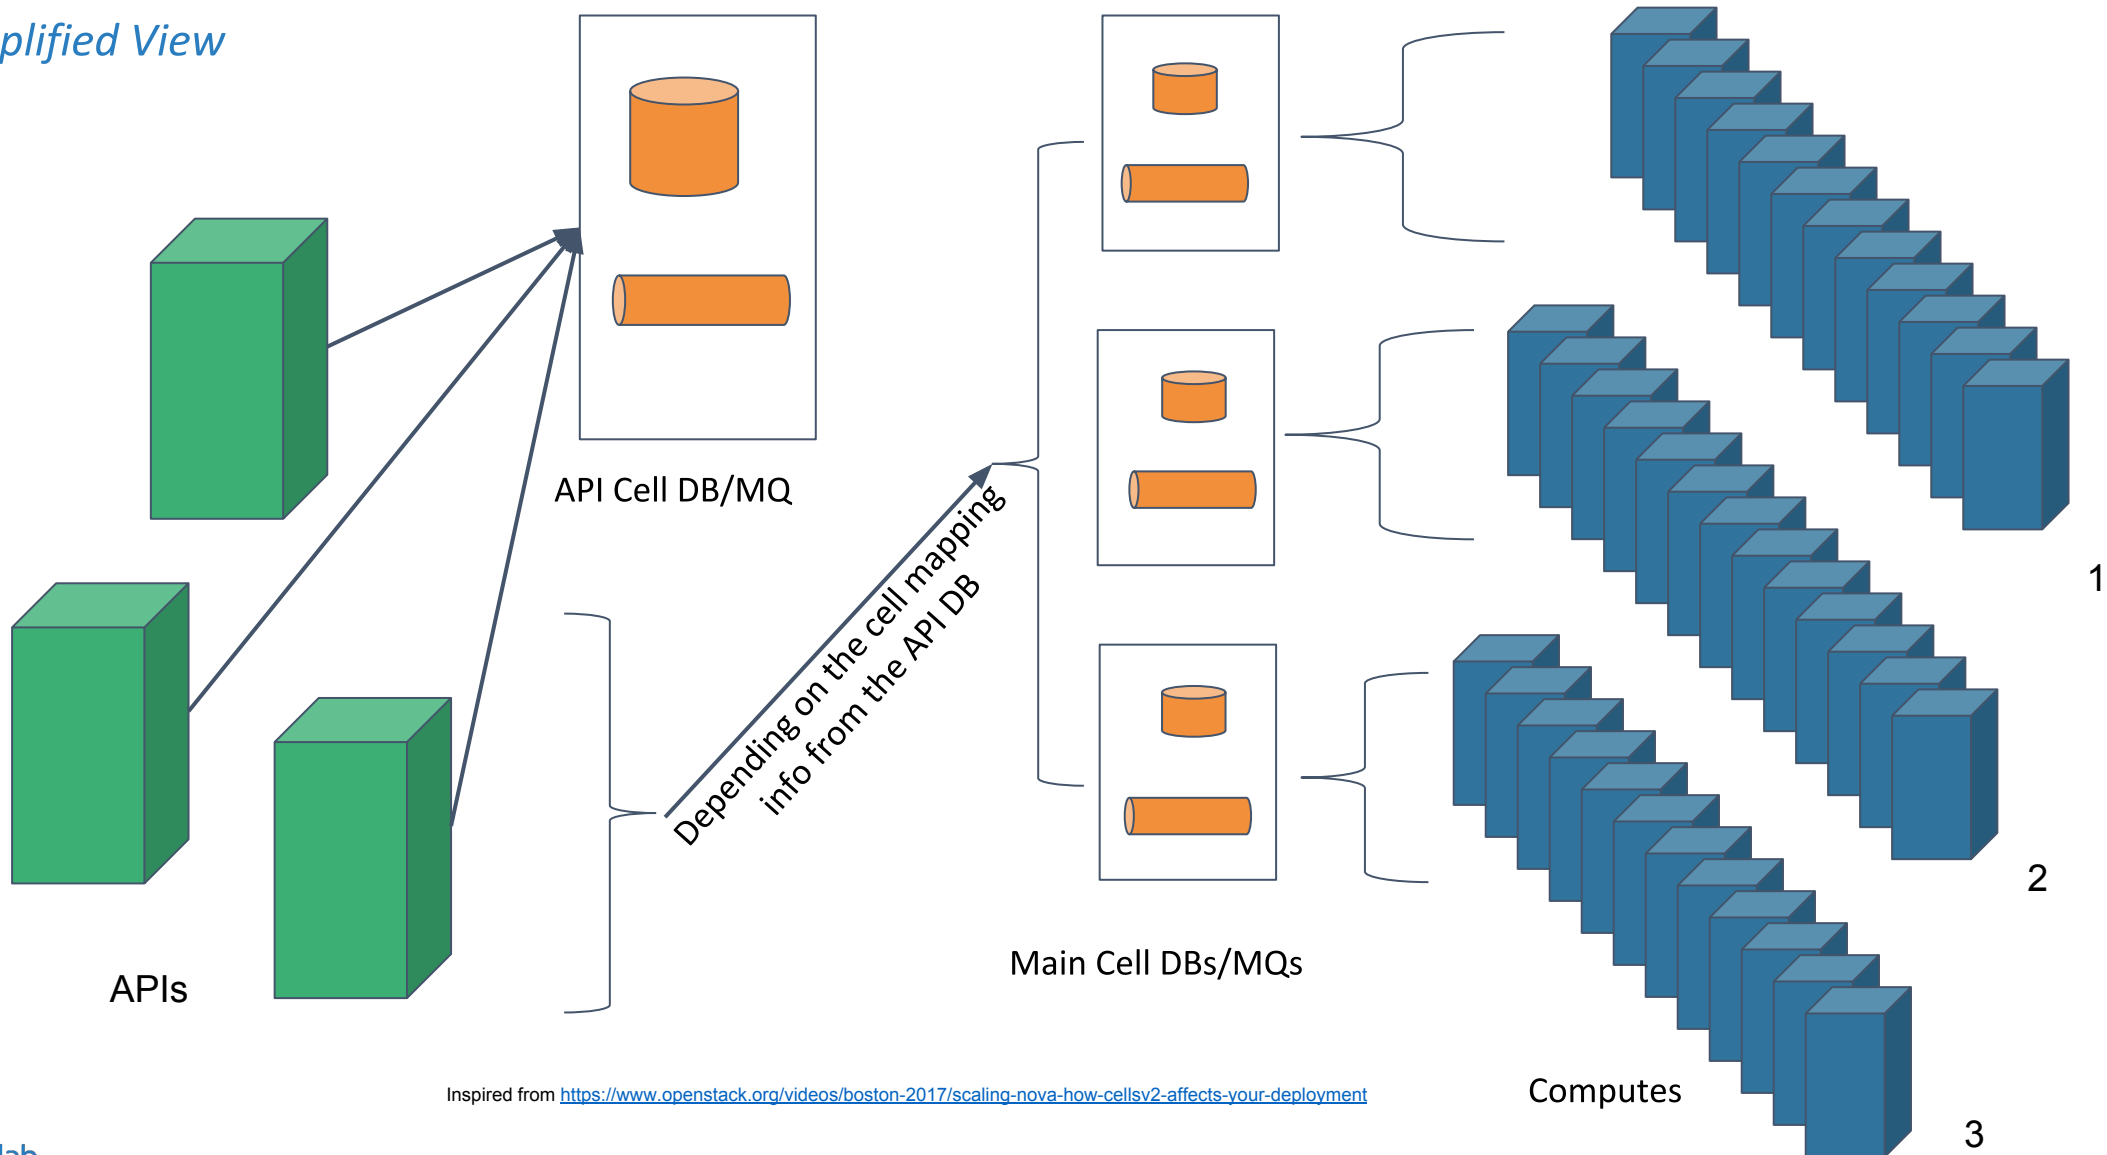
Upstream development has stopped for v1.

- Syncing/Python Replication is annoying.
- Top level cell scheduler has limitations.
 - ◆ It's the story of grass on the opposite side seeming more green.
- Race Conditions since its not built-in in a core manner.
- Two different code paths for nova and nova with cells.
- No path for upgrade from no cells to cells-v1.
- Data duplication in the parent cell.



Nova with Cells-v2

Simplified View



Inspired from <https://www.openstack.org/videos/boston-2017/scaling-nova-how-cellsv2-affects-your-deployment>

Migration to Cells-v2

Completely integrated with the whole nova code - it's in the main code path

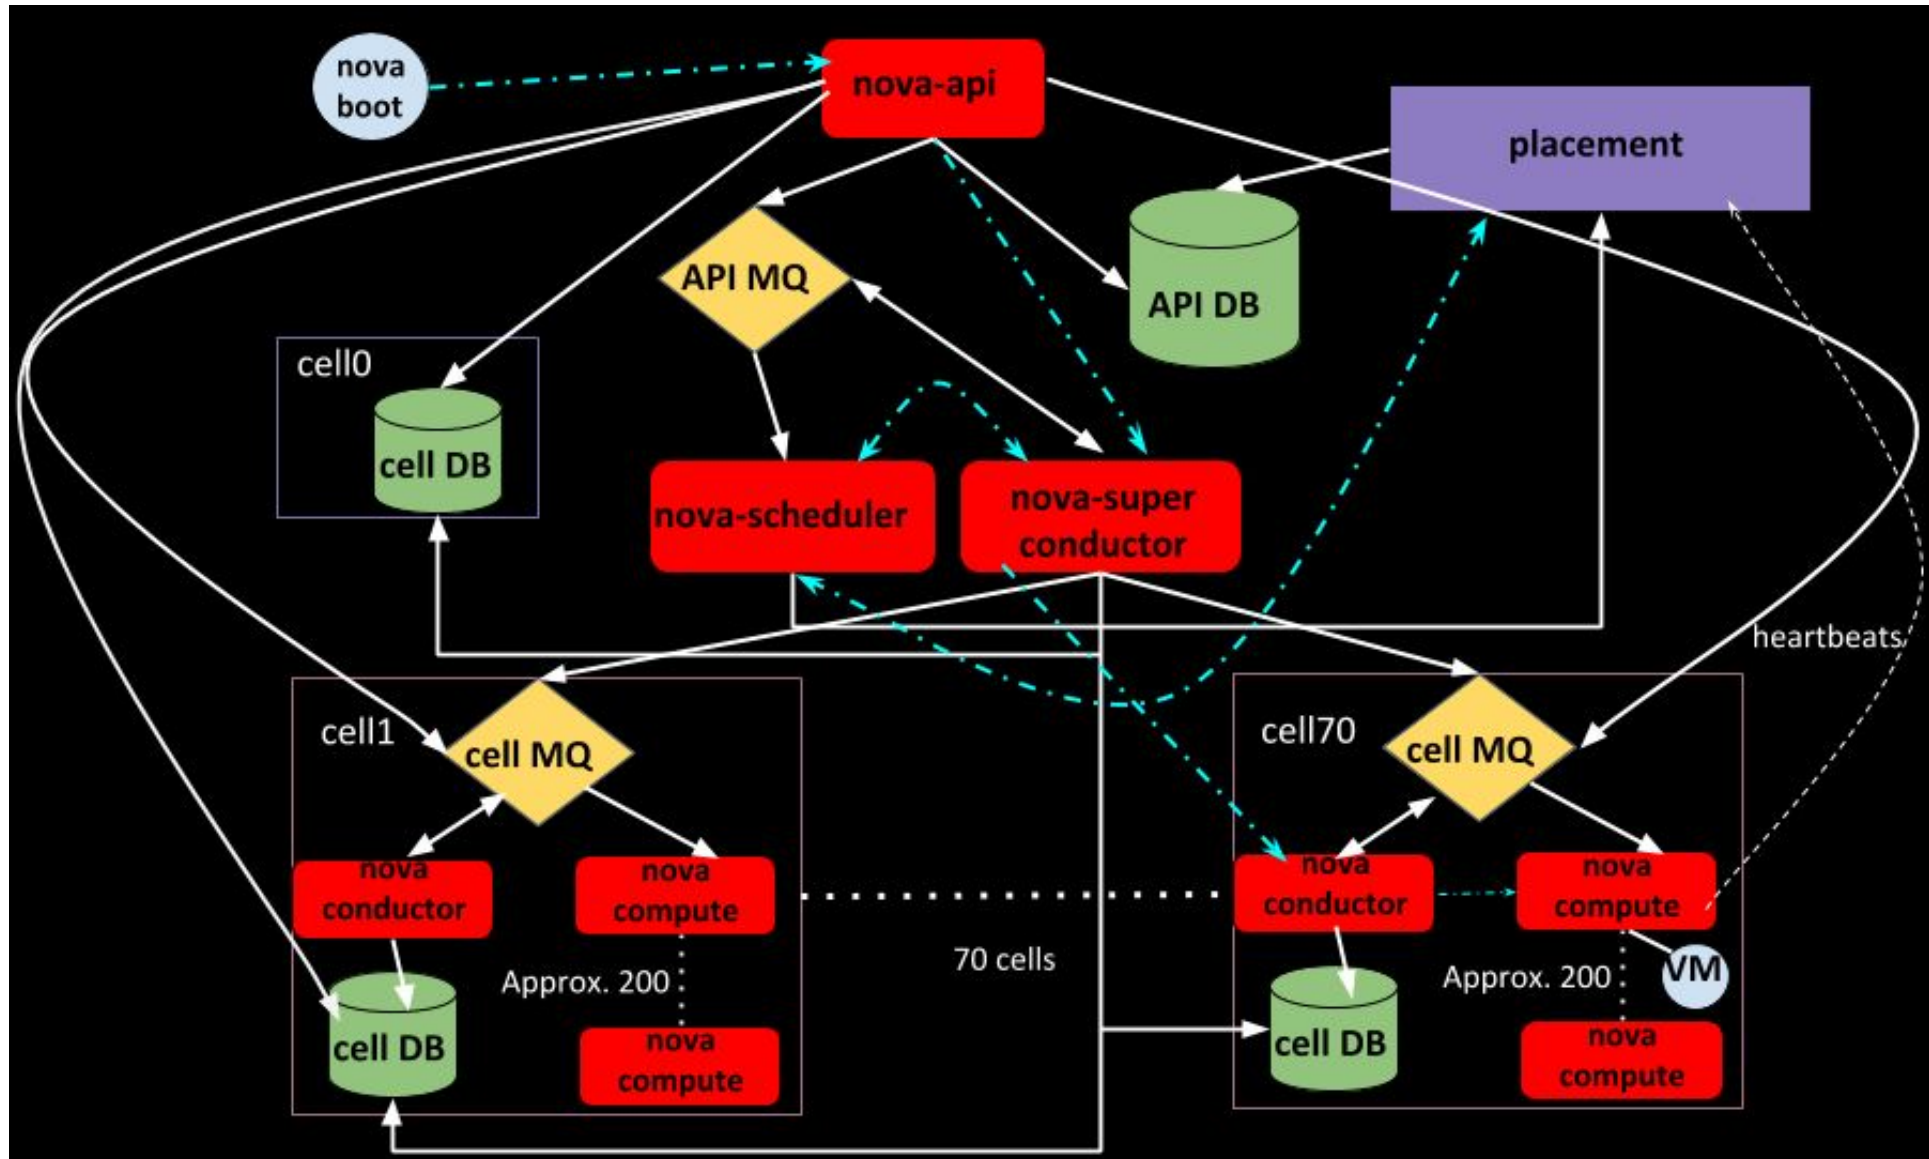
- Basically overcomes all the issues with Cells-v1.
- Global data moves to the API DB (all data in one place, no duplication).
- No need of syncing, no more upcalls.
- Single scheduler - knows about all nodes - pre-claims resources.
- DB/MQ info of each cell in API DB - no more “nova-cells” service.
- Most importantly, this is the maintained cells version upstream.
- All deployments are single cell-v2 or multi-cell.

On 25th April 2018 CERN moved to Cells-v2

One top level API cell and ~70 child cells; each having ~200 computes.

Nova with Cells-v2

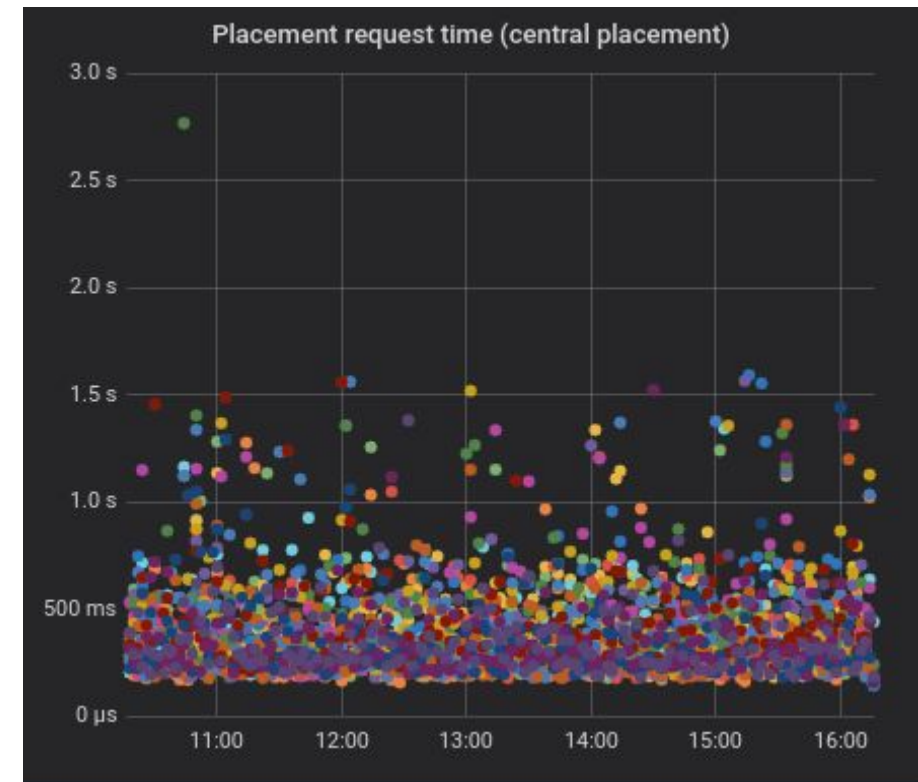
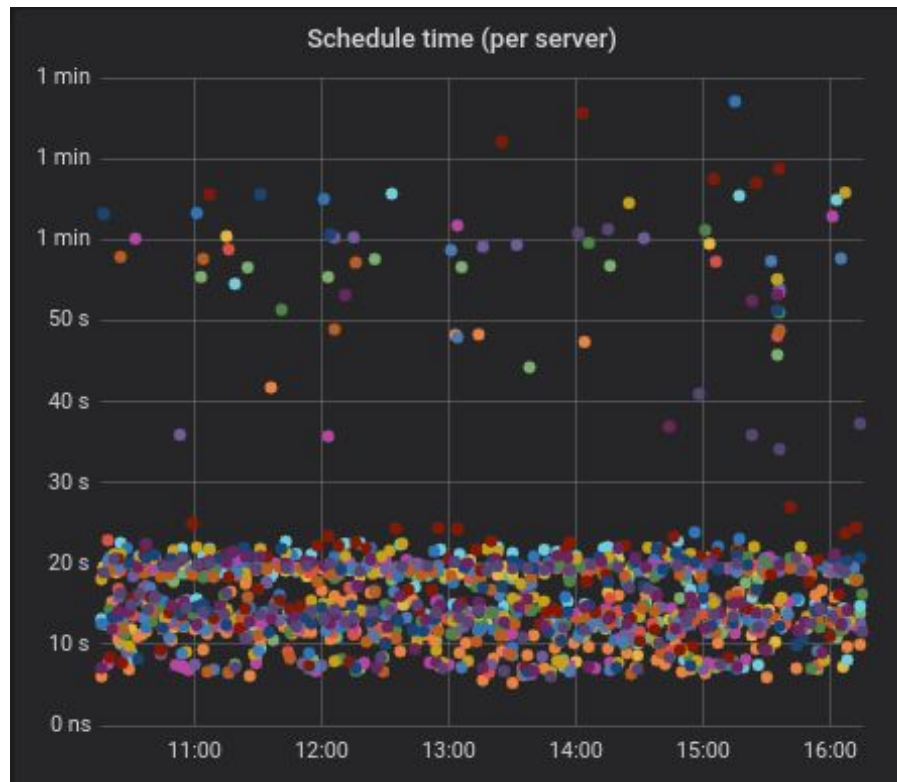
With the main components



CERN Cells - Current Status

Main Caveats

- Global Scheduler
- Compute to Placement heartbeats
- Performance Optimization with Multi-cells-v2



Project Update

Past

- Integration with upstream cells-v2 subteam, contributions / involved with development and pushing features needed for CERN / bug fixes related to preparing for the upgrade.

Present

- Overcoming the issues we face currently with cells-v2 / one of the first to move to multi-cells-v2 - (super exciting :)).

Next couple of months

- Continue developing multi-cells-v2 / optimizations / resilience
- Scaling placement / optimizing scheduling time.

References:

1. http://stackalytics.com/?company=cern&user_id=tssurya&release=all&metric=commits



Questions?

surya.seetharaman@cern.ch

tssurya @ openstack-nova (IRC)

suseetha @ mattermost