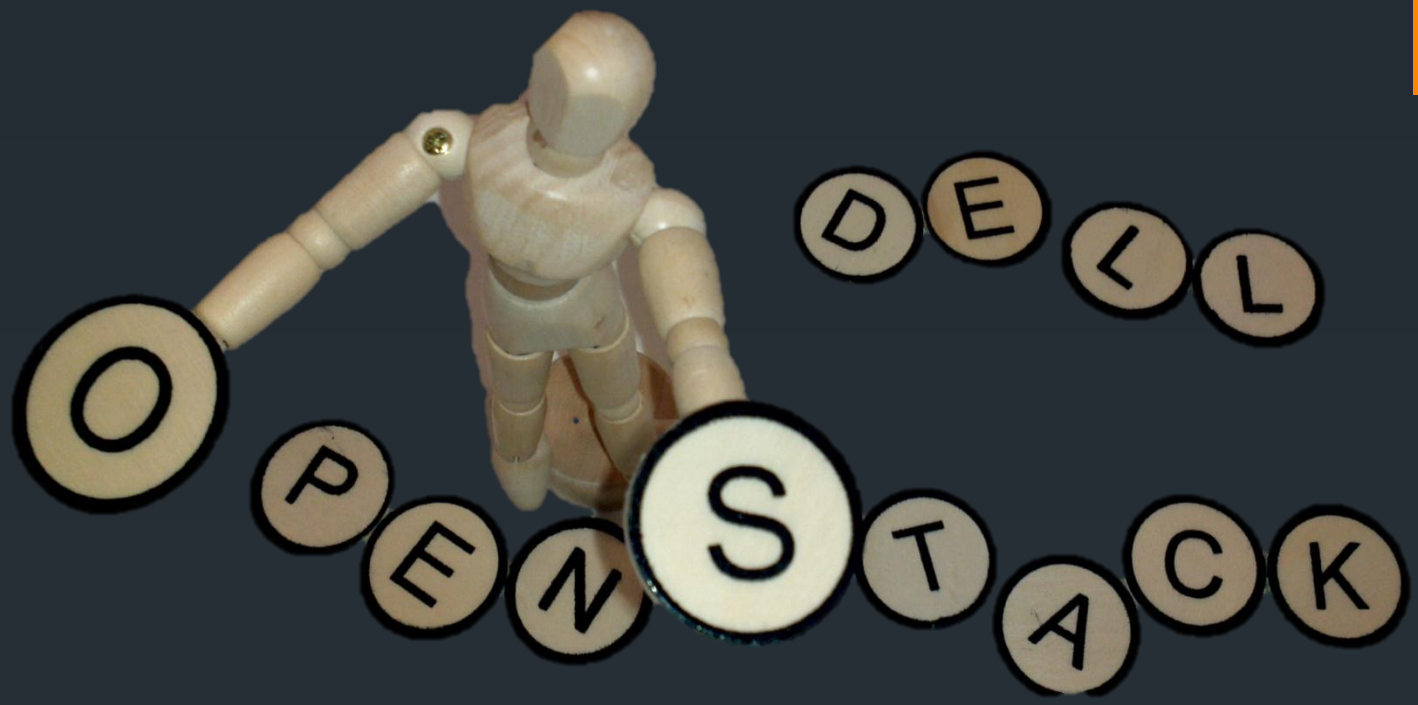

OpenStack & Hardware

Rob Hirschfeld, Dell

Greg Althaus, Dell

Lessons Learned from our Bootstrapping White Paper

- 1 year later
- Customer Tested

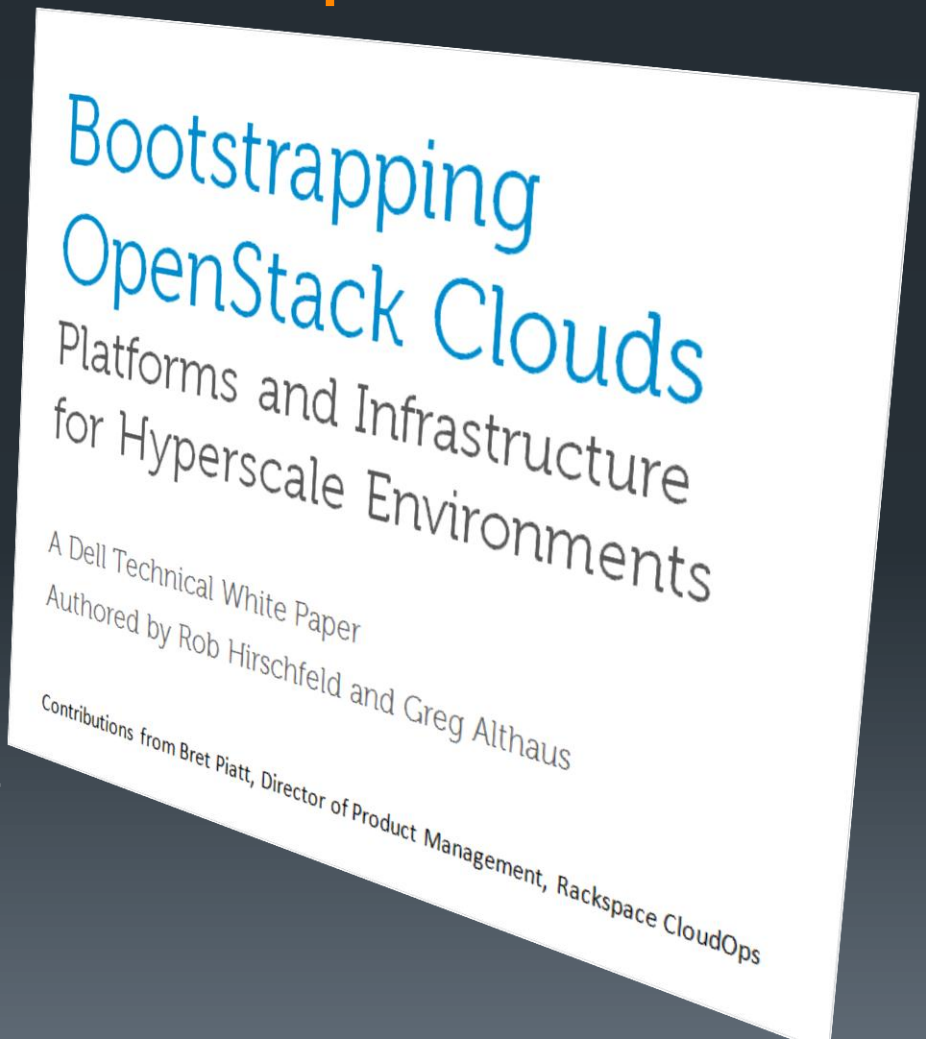

So How did we do?

AnyScale vs. HyperScale

- Need to start small
- Natural Growth Steps
  - 6 for PoC
  - 20 – 100 for Pilot
  - 200+ for Production
- Design patterns replicate to scale out
- Generic hardware configurations

CLOUD IS BIASED TOWARD
LOTS OF SMALL NODES

Fault Zones

Build in expectation for Hardware Failures

- Core to Swift Design
- Nova is not there yet…

- Mitigated by
  - Stacked Switch Design
  - Block Storage backed VMs

Natural Grouping

- Naming
  - “Pod” or “Cluster” or “Zone”

- 150 – 200 Nodes
- Based on Stacked Switches
- Need for Core Routing
- Need to administrative infrastructure
  - <5% of capacity should be control & admin

Node Networking

- Not seeing physical isolation
- Moving aggressively to 10 GbE
- Using Teamed NICs for redundancy

- **Still need innovation – expect changes**

- Dell acquired Force10

Compute Hardware

- **BALANCE, BALANCE & DENSITY**

- VM to Core ratios
  - From 4 or 1 to 1.
- VM to Spindle ratios
- RAM is NOT a major factor,
  - easy to get more than needed
- Networking is a constraint

- 8 to 12 cores are still sweet spot
  - Spindle counts
  - Network capacity

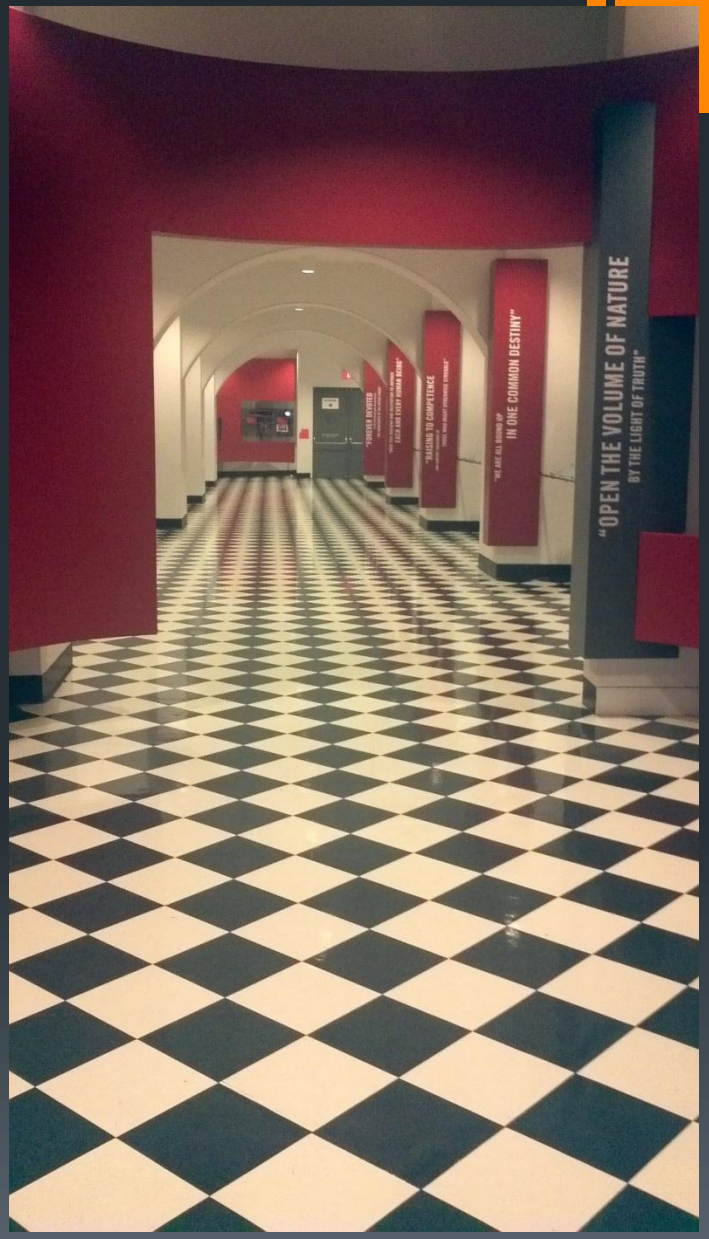

Storage Hardware

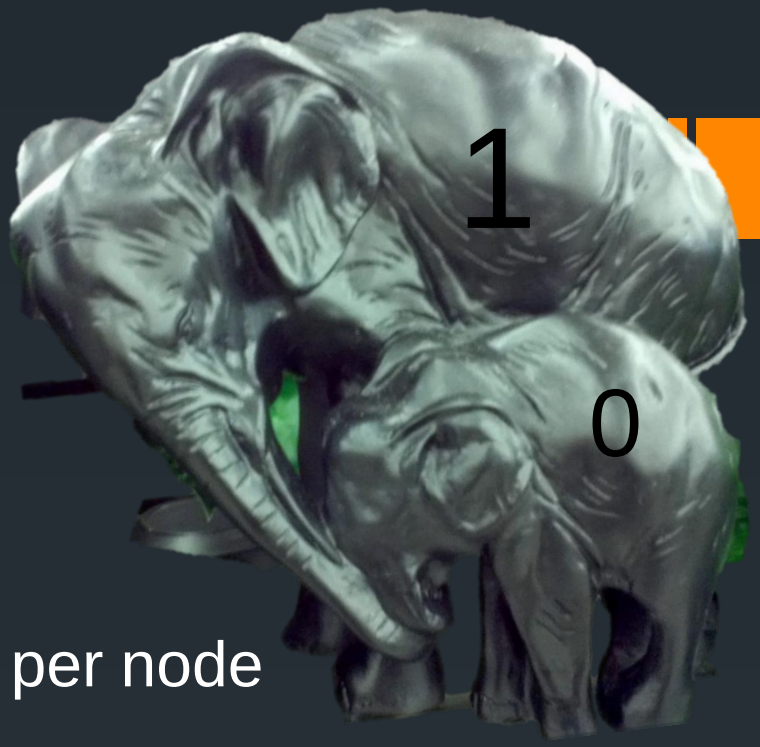

- **ALL ABOUT SPINDLES!**
  - Requests for more spindles per node
  - Dell is listening
- Need Block storage (e.g.: iSCSI)
- Nova Volume flux causing churn
- Drives 10 GbE Discussions

CloudOps

- New white paper!

- Cloud Solution include Operational Model
- **Based on DevOps approaches**

- Critical thinking to manage AnyScale architecture
  - Rolling upgrades
  - Hardware fault tolerance
  - Tip balance toward more small nodes

Questions?

More Information:

http://Dell.com/OpenStack
http://RobHirschfeld.com