# Wicked Easy Ceph Block Storage & OpenStack Deployment with Crowbar

Kamesh Pemmaraju, Dell
Neil Levine, Inktank

# Outline

- Introduction
  - Dell OpenStack-Powered Cloud Solution
  - Ceph
  - Crowbar
- Why Crowbar + Ceph in OpenStack?
  - OpenStack Block Storage gap
  - Automation, scale
- What we have done to enable things?
- Customer benefits, drivers – an example
- What's coming next?

# Clouds Require an Operational Focus

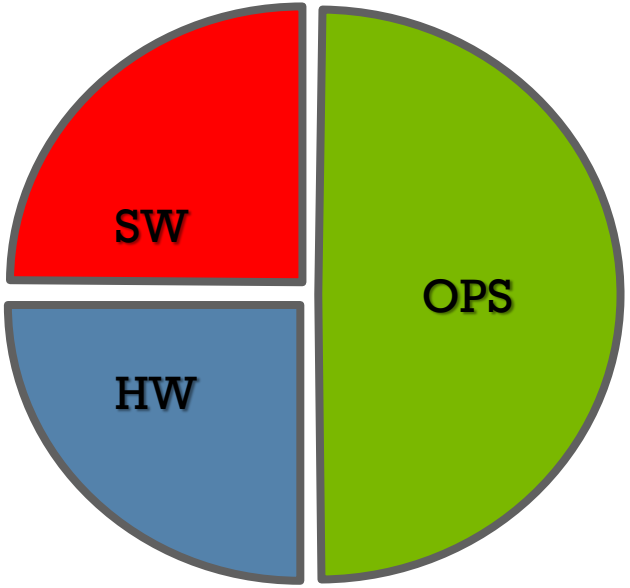


- Clouds demand significant operational and process controls
- Operational decisions drive hardware and software decisions
- We are finding ways to productize operations into best practices

# Dell OpenStack-Powered Cloud Solution



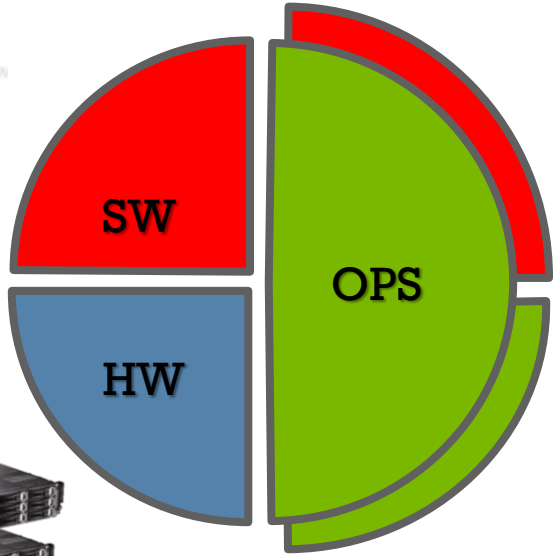


"Crowbar"
CloudOps
Software

Reference
Architecture

Services &
Consulting



# Dell OpenStack-Powered Cloud Solution

| HW + SW + Services | | |
|---|---|---|
| **Hardware** | **HW reference architecture** | • C6220, C6105,R720, R720XD servers<br>• Storage and compute<br>• Force 10 S60 and PowerConnect |
| | **Configuration** | • Min 6 nodes. Max 60 nodes |
| **Software** | **Software** | • OpenStack Installer and continuous integration (Crowbar)<br>• OpenStack cloud SW<br>• Ceph for distributed storage |
| | **Operating system** | • Ubuntu 12.04 (host)<br>• Windows (guests)<br>• Linux (guests) |
| | **Hypervisor** | • KVM |
| **Services** | **Deployment and consulting** | • Hardware integration – onsite or merge center<br>• Software installation – onsite<br>• OpenStack consulting services<br>  – Assessment, design and installation<br>  – Delivered via partnerships with Mirantis and Canonical |
| | **Support** | • HW/Crowbar: Dell ProSupport<br>• OpenStack support via Canonical (coming soon) |

# Ceph – Unified Storage Software

| OBJECTS | VIRTUAL DISKS | FILES & DIRECTORIES |
| --- | --- | --- |
| CEPH OBJECT STORAGE | CEPH BLOCK STORAGE | CEPH FILE SYSTEM |
| A powerful S3- and Swift-compatible gateway that brings the power of the Ceph Object Store to modern applications | A distributed virtual block device that delivers high-performance, cost-effective storage for virtual machines and legacy applications | A distributed, scale-out filesystem with POSIX semantics that provides storage for a legacy and modern applications |

# Ceph – Key Differentiators

- Unified storage platform (Object + Block + File)
- Ceph Intelligent Placement (CRUSH)
- Ceph Intelligent Devices (Self-healing, P2P)
- Ceph Intelligent Objects (Embedded Software Classes)
- Ceph Integration (Linux Kernel, OpenStack, Cloudstack)

# Ceph Architecture

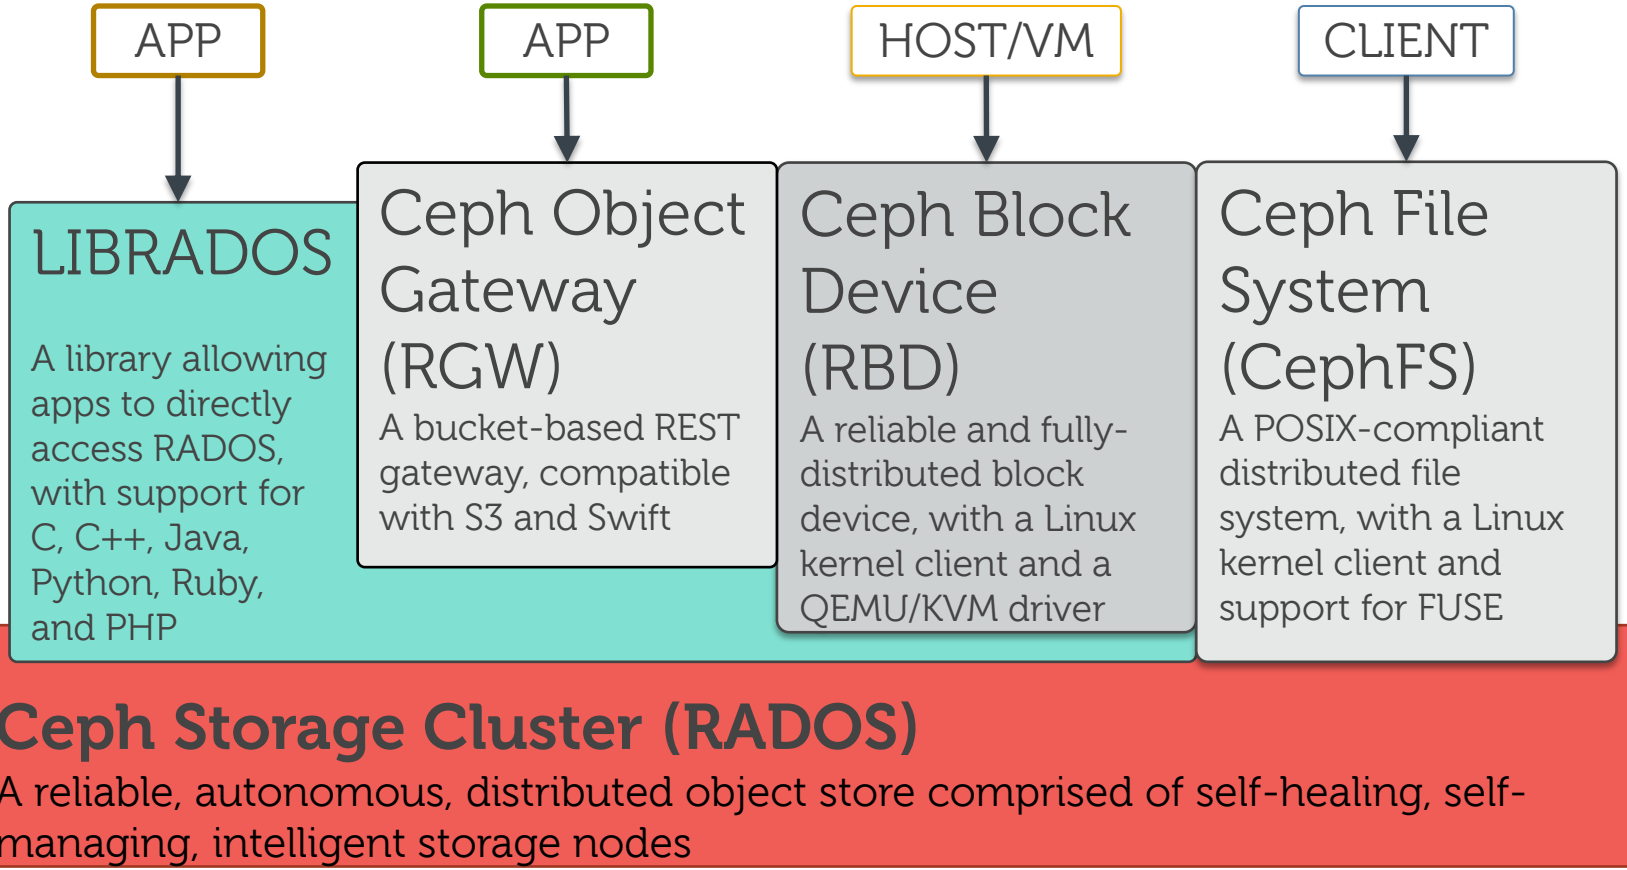

# Ceph RADOS Components


## Ceph Monitor Nodes:

- Maintain cluster map
- Provide consensus for distributed decision-making
- Must have an odd number
- These do not serve stored objects to clients

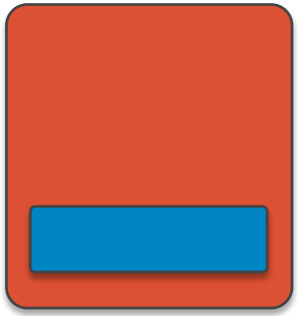

## Ceph OSD Nodes:

- One Ceph OSD daemon per disk (recommended)
- At least three nodes in a cluster
- Serve stored objects to clients
- Intelligently peer to perform replication tasks
- Supports object classes

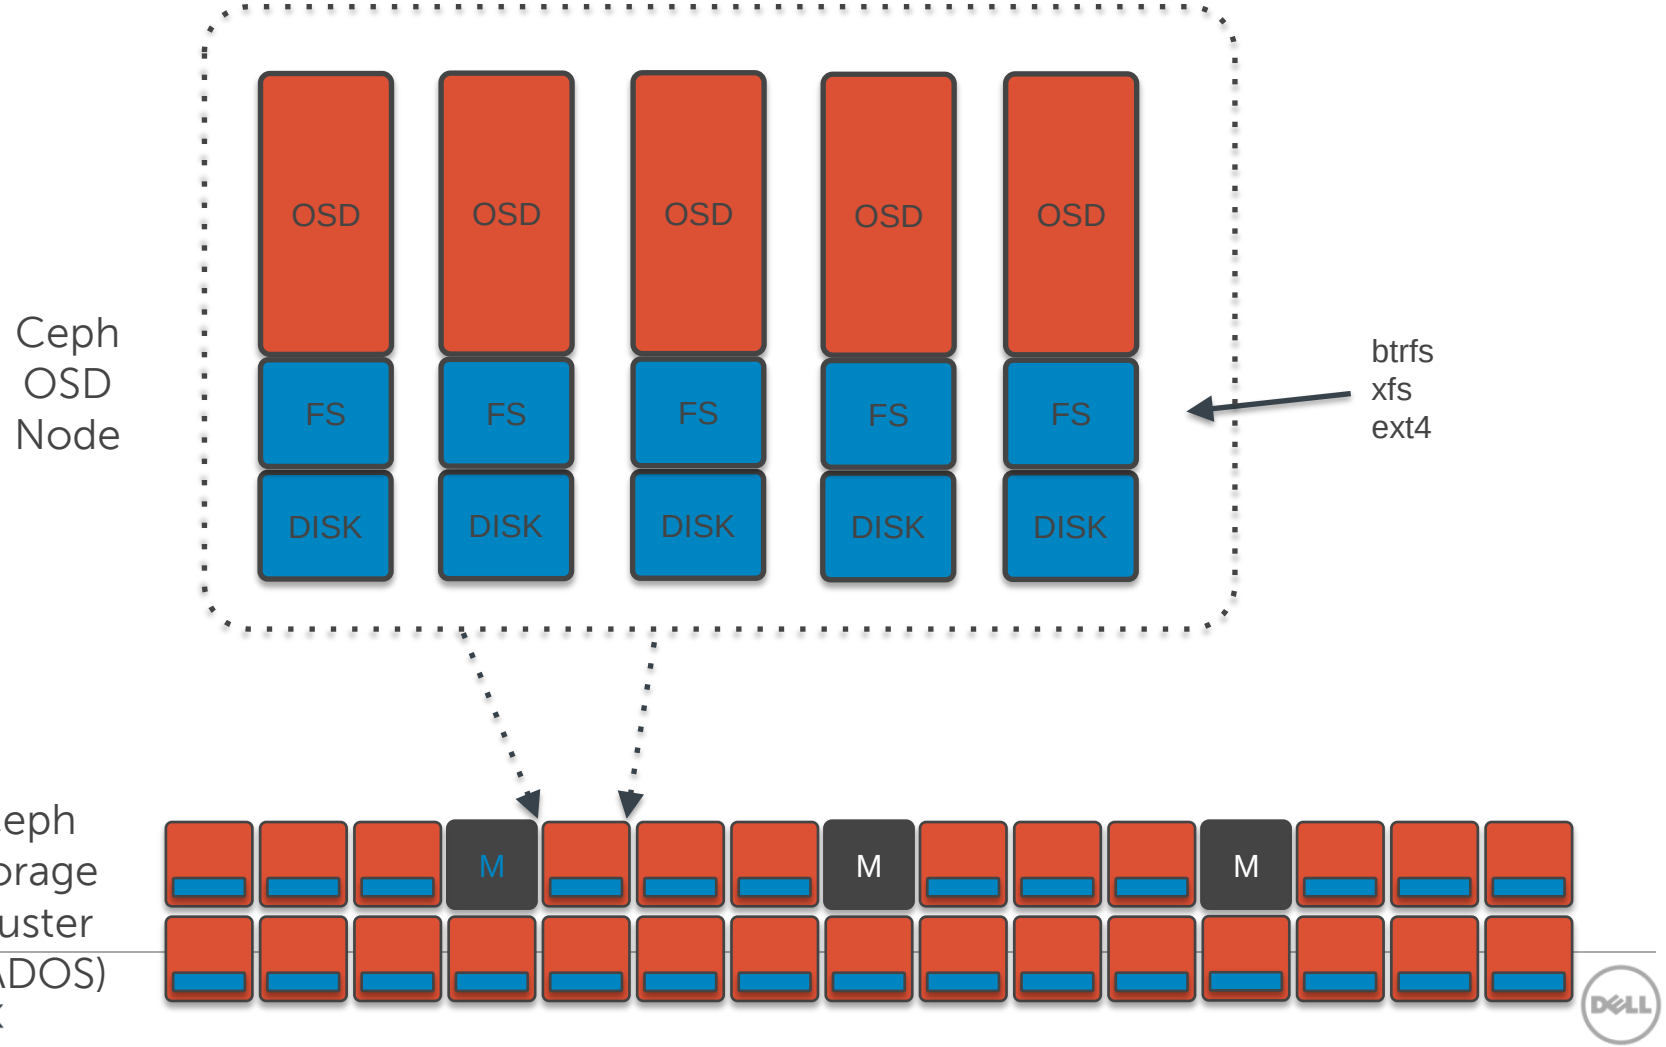
Ceph
OSD
Node
OSD
OSD
OSD
OSD
OSD
FS
FS
FS
FS
FS
DISK
DISK
DISK
DISK
DISK
btrfs
xfs
ext4
Ceph
Storage
Cluster
(RADOS)
M
M
M

# What is Crowbar?

Mission: “A Zero Touch Cloud Installer”

**Servers in boxes to full function cloud in under 2 hours**

- Fast & Flexible
  - Bare metal install including BIOS & RAID config
  - Users can choose how their system is configured (“barclamps”)
- DevOps Embracing
  - Ongoing ***Operations Model*** (DevOps for Clouds)
  - Leverages & Wraps Opscode Chef
- *Open*
  - *Not* specific to OpenStack – Dell using for other Apps
  - *Not* restricted to Dell hardware
  - Apache 2 licensed

# Images vs. Layers: Overview

Images: **Single Unit**

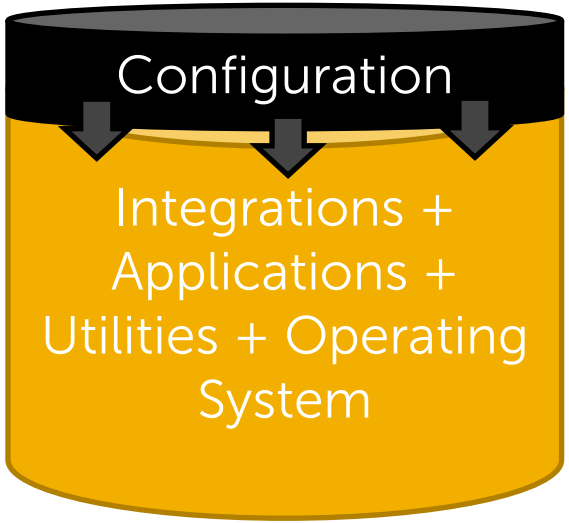


Layers: **Stacked Pieces**

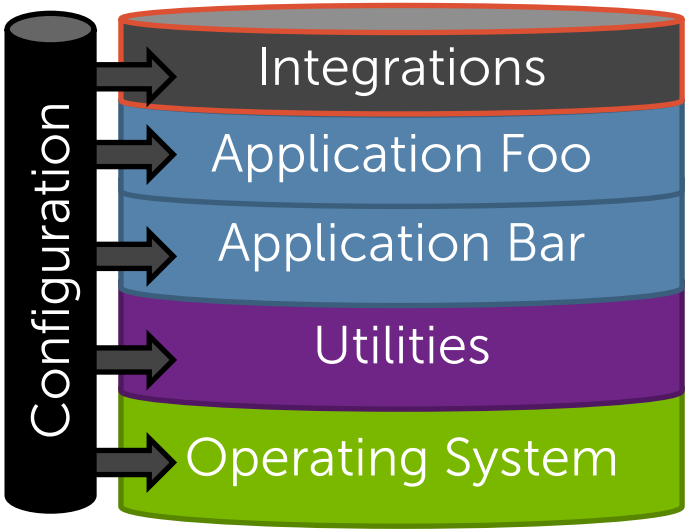

# Images vs. Layers: Lifecycle

## Images: **Replacement**

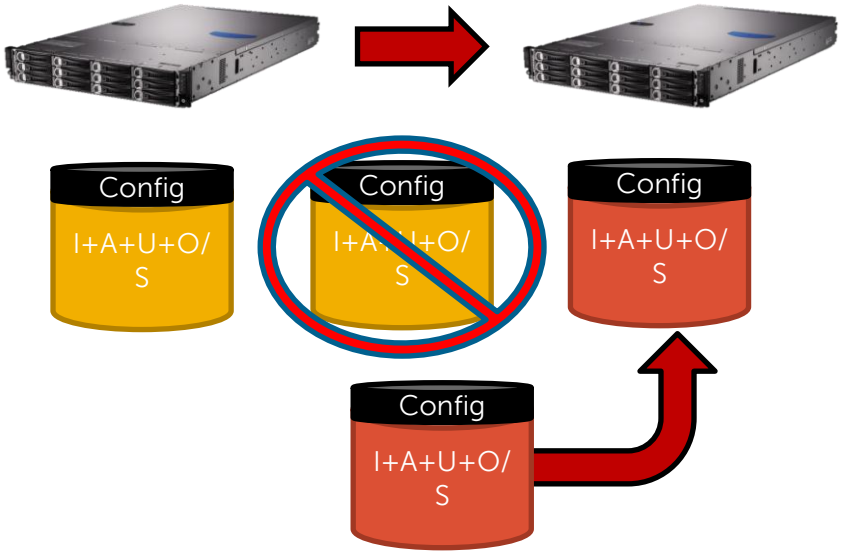


## Layers: **Upgrade**

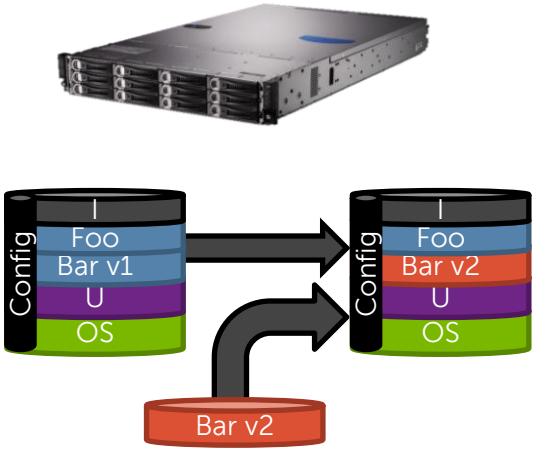

# Crowbar Uses the Layer Model for Deployment

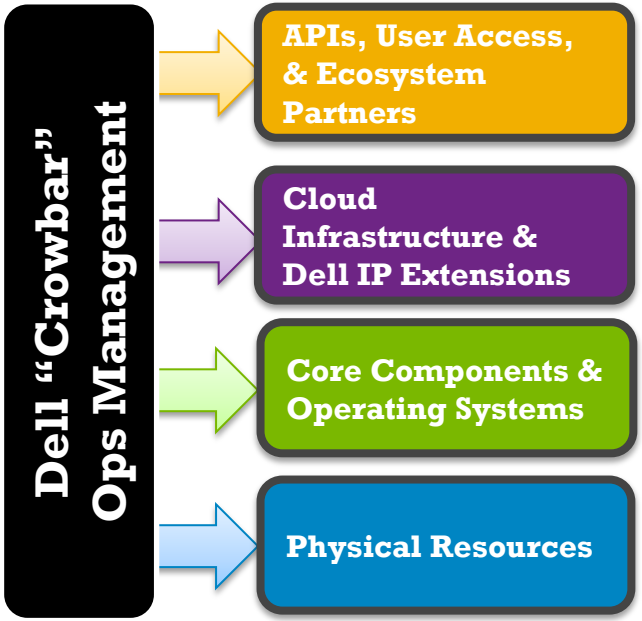


- Includes all the components required to implement an *entire* cloud infrastructure including ecosystems partners
- Pluggable components deploy cloud infrastructure
- Allow for addition of Dell IP and expansion by the community services and customers
- Can integrate with Dell existing products
- Delivers basic data center services and required cloud infrastructure.
- Provision bare-metal servers from box to cloud WITHOUT user intervention (other than racking/cabling and some minimal configuration questions) **NOT EXCLUSIVE TO DELL HARDWARE**

* Crowbar has potential to service other programs beyond OpenStack

# How Does Crowbar Work?

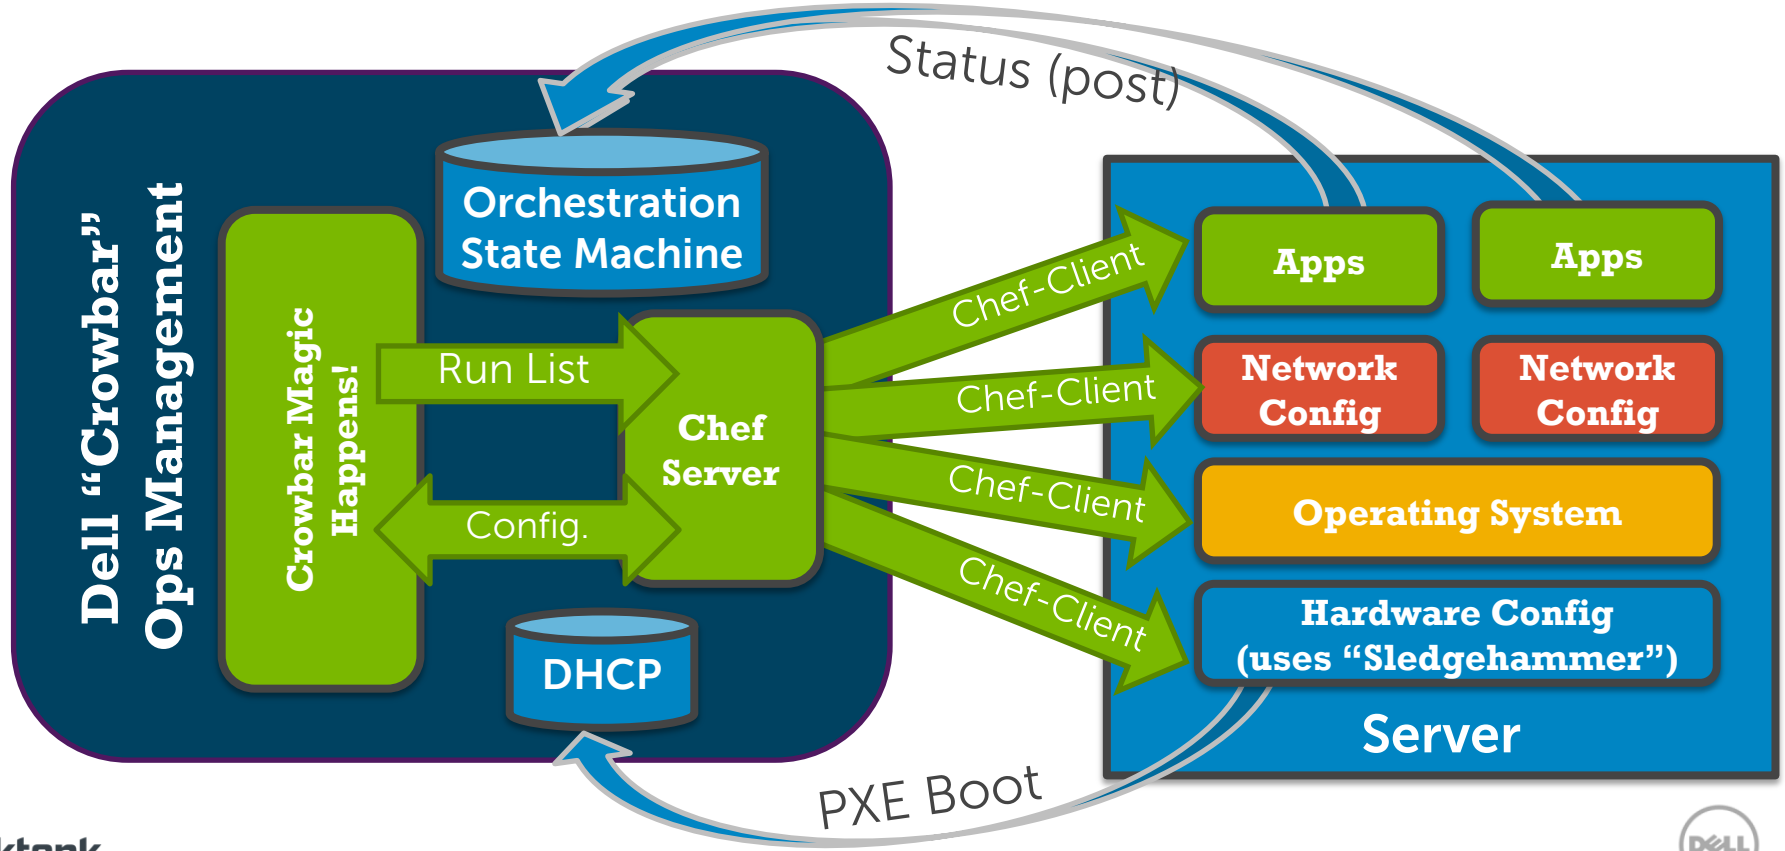

# What is a Barclamp?

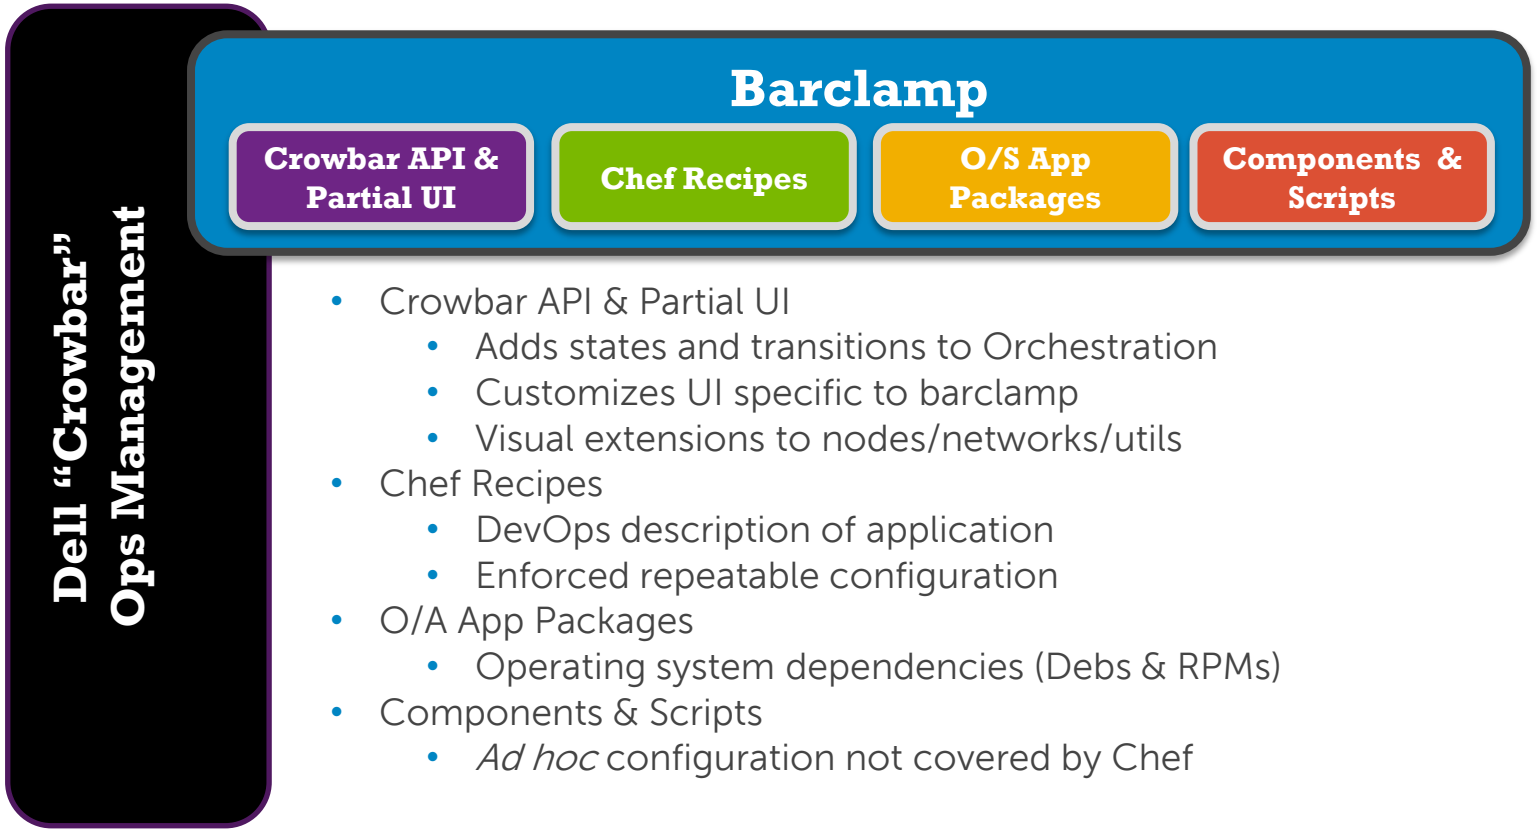


- Crowbar API & Partial UI
  - Adds states and transitions to Orchestration
  - Customizes UI specific to barclamp
  - Visual extensions to nodes/networks/utils
- Chef Recipes
  - DevOps description of application
  - Enforced repeatable configuration
- O/A App Packages
  - Operating system dependencies (Debs & RPMs)
- Components & Scripts
  - *Ad hoc* configuration not covered by Chef

# Modular Design: Barclamps The Configuration You Want to Deploy

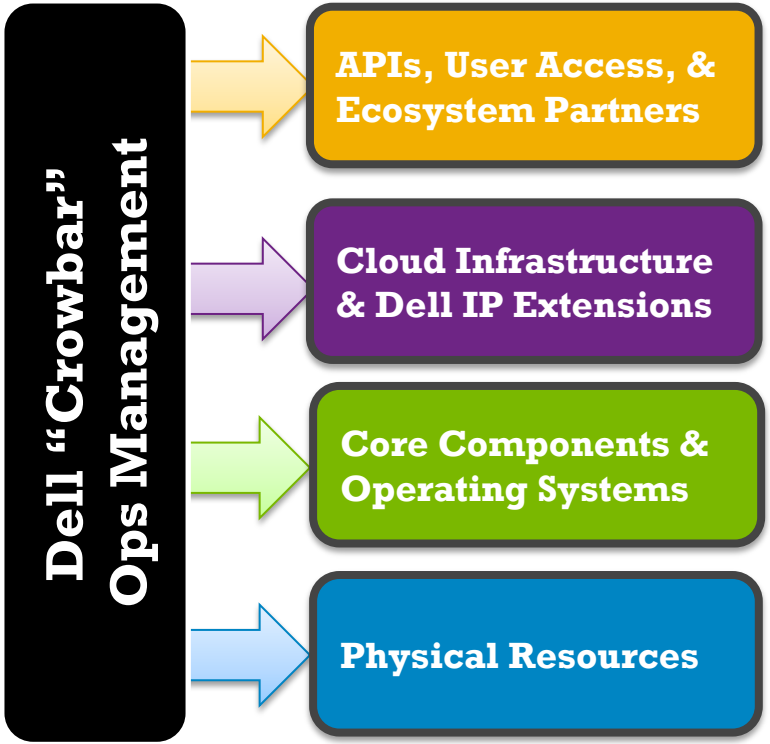


| | | |
|---|---|---|
| Nagios | Ganglia | Dashboard |
| enStratus | RightScale | Horizon |
| Nova | Swift | Glance |
| Ceph | Force10 | Keystone |
| Crowbar | DNS | Logging |
| Deployer | NTP | |
| Provisioner | BIOS | IPMI |
| Network | RAID | |

# The 4 Stages of Cloud Storage Thinking

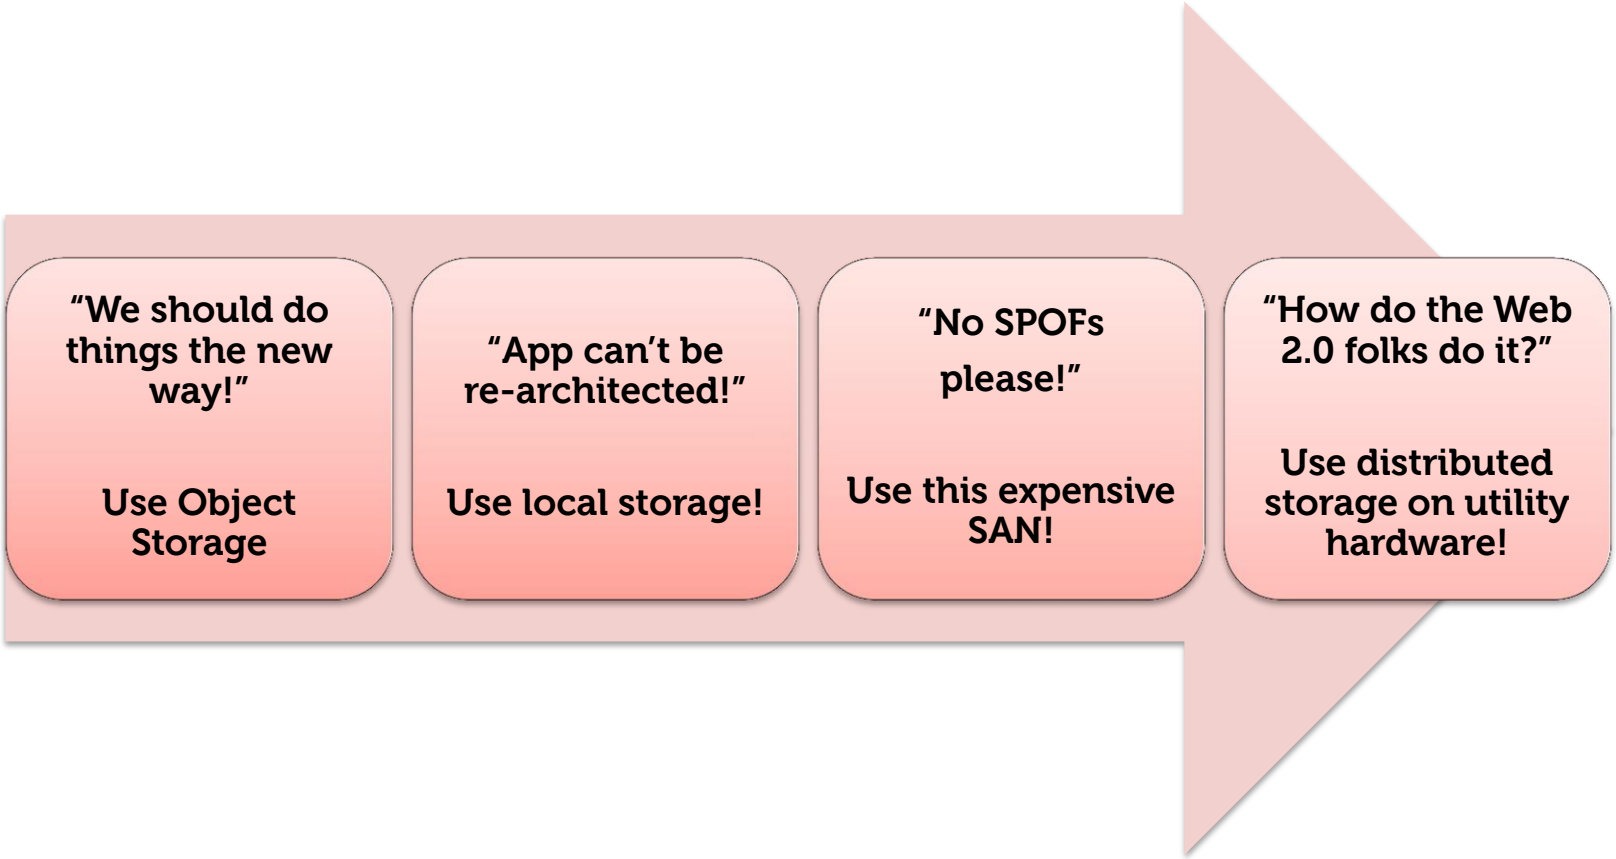

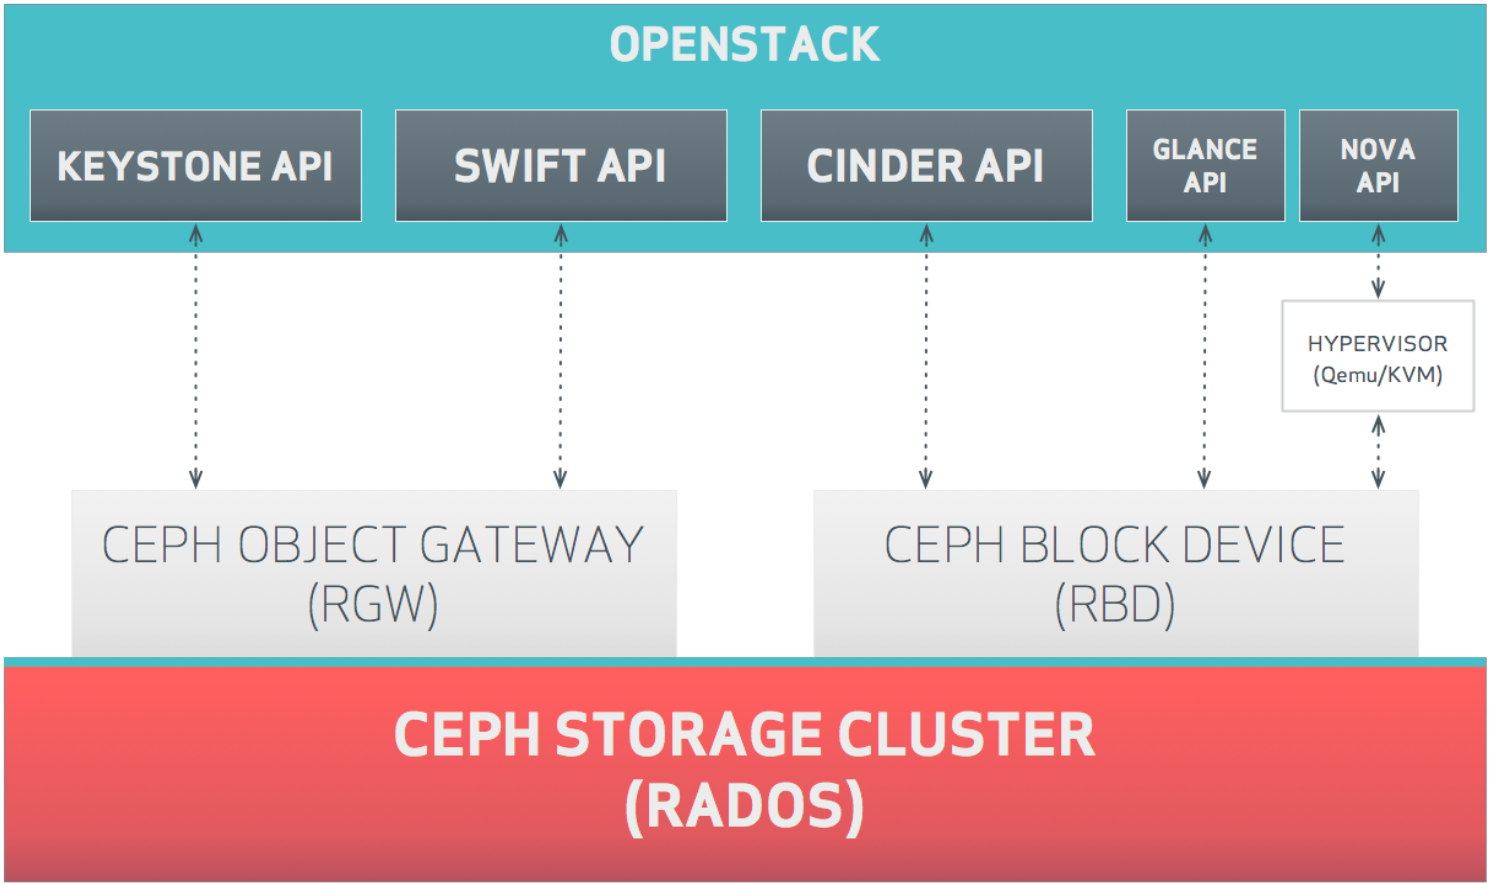
OPENSTACK
KEYSTONE API
SWIFT API
CINDER API
GLANCE API
NOVA API
HYPERVISOR (Qemu/KVM)
CEPH OBJECT GATEWAY (RGW)
CEPH BLOCK DEVICE (RBD)
CEPH STORAGE CLUSTER (RADOS)

# Ceph Block Storage for Dell Openstack

- Integrated: with Cinder and Nova (via qemu-kvm)
- Compatible: with Glance
- Foundations: no SPOFs, self-healing, CRUSH
- Features: Copy-on-Write, Snapshots, Cloning
- Integrated: with Crowbar for automation

# Scale-Out Operations With OpenStack

If you can't automate it, you can't scale it

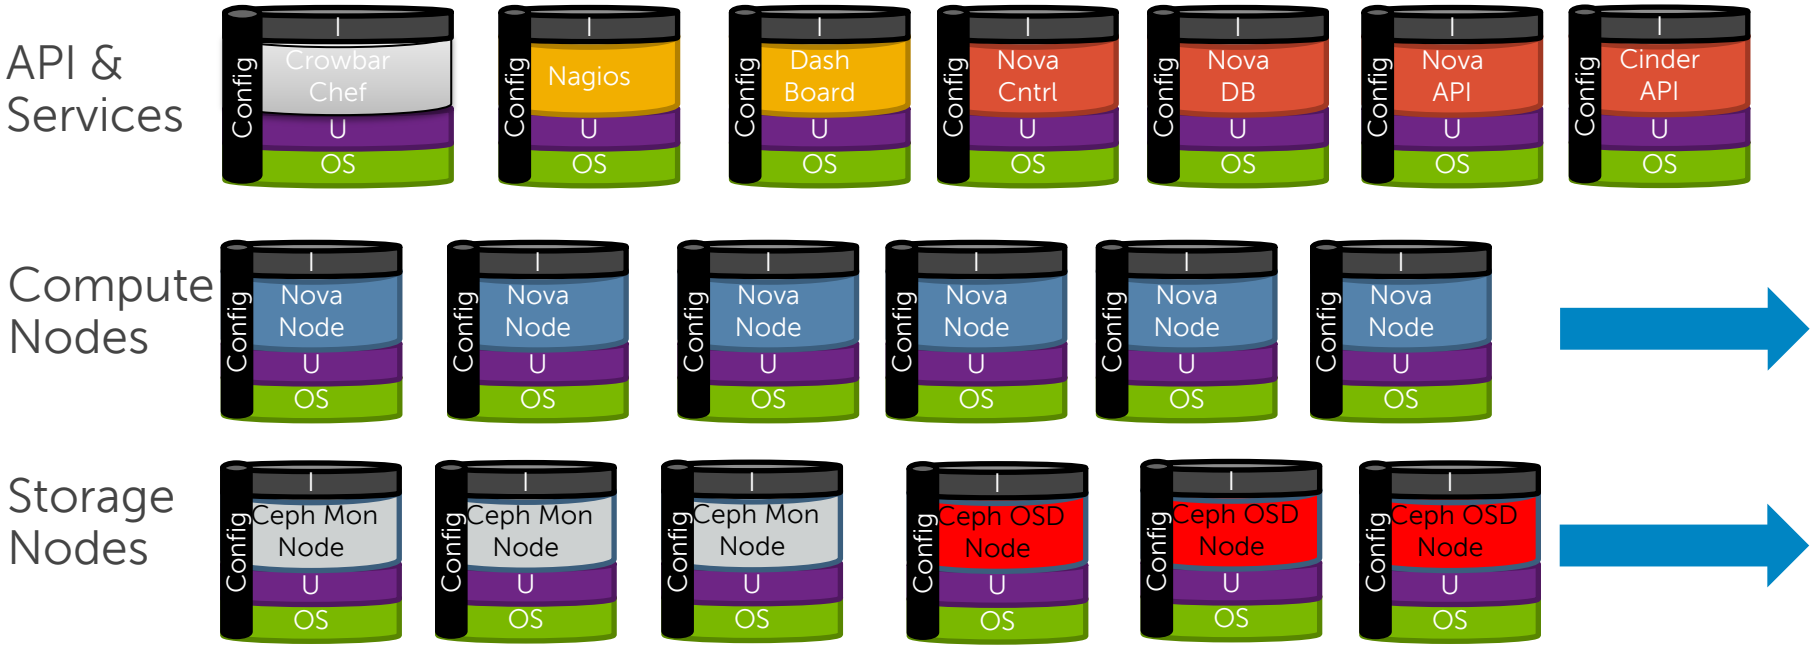

# Ceph Related Barclamps for Crowbar

- https://github.com/ceph/barclamp-ceph
  - ceph.com packages into the Crowbar Admin Node
  - Ceph OSD roles (ceph-store)
  - Ceph Mon roles (ceph-mon)
- https://github.com/crowbar/barclamp-nova
  - Hypervisor with RBD (ceph-client)

# What Have We Done to Enable?

- Dell and Inktank have partnered to bring a complete solution for OpenStack + Ceph + Automated deployment with Crowbar
- The joint solution provides:
  - Crowbar barclamp to
    - Deploy Ceph clusters automatically in a very short time
    - Connect those clusters to OpenStack
  - Professional Services, Support, and Training
    - Collaborative Support for Dell hardware customers
  - Joint Solution
    - Validated against Dell Reference Architectures via the Technology Partner program

# Customer:
## University to Deploy OpenStack + Ceph Using Crowbar

### Situation

- University employs close to 900 researchers and receives hundreds of million dollars in Grants as a Top 10 Research Institution primarily working on Cancer and Genomic Projects. The University has a need to provide a centralized data repository for Researchers in order to insure compliances concerning retention of data.
- The intent of the Data Repository is to provide 2TB of free storage space to each Researcher, with opportunity to purchase from the university more capacity at a very reasonable cost.

### Decision Drivers

- University investigated using an traditional SAN storage solution which was very expenses on a per TB basis. They also investigated using public cloud storage options which also proved to be too expensive. Finally, they looked at Hadoop for the project but found that it was not a good fit for this use case.

- In the end, University chose Dell/Inktank to architect a platform that would be very scalable and provide lost costs per TB and was the best of all worlds that provide compute and storage on the same hardware.

inktank

DELL

# What's Next?

- May
  - Ceph Cuttlefish
- Summer 2013
  - Dell Grizzly support
  - Ceph Dumpling (August)
- https://github.com/crowbar/crowbar/wiki/Crowbar-2.0
  - Crowbar v2.0 work is happening in the open. Check it out and participate!

# Don't Miss the OpenStack Summit Sessions

## Havana Sessions:

– Planning the Ceph Roadmap for Openstack

Wednesday April 17, 2013

1:50pm - 2:30pm in Room - B113

– Features for Ceph with Cinder and Beyond

Thursday, April 18, 2013

9:50 – 10:30am in Rooms – C120+121+122

# Contact Information

## Reach Kamesh and Neil for additional information:


Dell.com/OpenStack

Dell.com/Crowbar

Inktank.com/Dell

Kamesh_Pemmaraju@Dell.com

@kpemmaraju

Neil.Levine@Inktank.com

@neilwlevine


Visit the Dell and Inktank booths in the OpenStack Summit Expo Hall

inktank

DELL