Highly Available OpenStack Deployments with NetApp & Red Hat's OpenStack platform

June 26, 2015

Jeff Applewhite
Technical Marketing Engineer, Cloud Solutions Group, NetApp

Introductions

Introduction

Jeff Applewhite

- Technical Marketing Engineer, Cloud Solutions Group
- 5 Years at NetApp, 2.5 Years on OpenStack Engineering team
- ATC on various OpenStack projects

redhat

NetApp

Agenda

1) Introductions
2) NetApp OpenStack Integrations
3) Red Hat and NetApp OpenStack collaborations
4) Deploying RHEL-OSP and NetApp in an HA configuration
5) FlexPod: Cisco, Red Hat, and NetApp unlocking business value
6) Q & A

OpenStack 10,000 Foot View

- **Every release is getting better**
  - New features released on regular 6 month cadence
  - Excellent quality through automated CI/CD DevOps processes
  - Version upgrades are becoming easier
- **Installation toolsets are Enterprise Ready**
  - HA Production Deployments are on the rise
  - The ease of deployment and configuration of NetApp Cinder from RHOS5 → RHOS6 has greatly improved.
  - GUI configuration of NetApp now works “out of the box”.

The NetApp OpenStack Story

Overview

NetApp OpenStack Involvement

- OpenStack Foundation
  - Charter member (Gold)
  - OpenStack Summit sponsors
- 1st Major Storage Provider in Community
  - Upstream Contributions
  - Numerous Production Deployments
  - NetApp is a Deployer of OpenStack

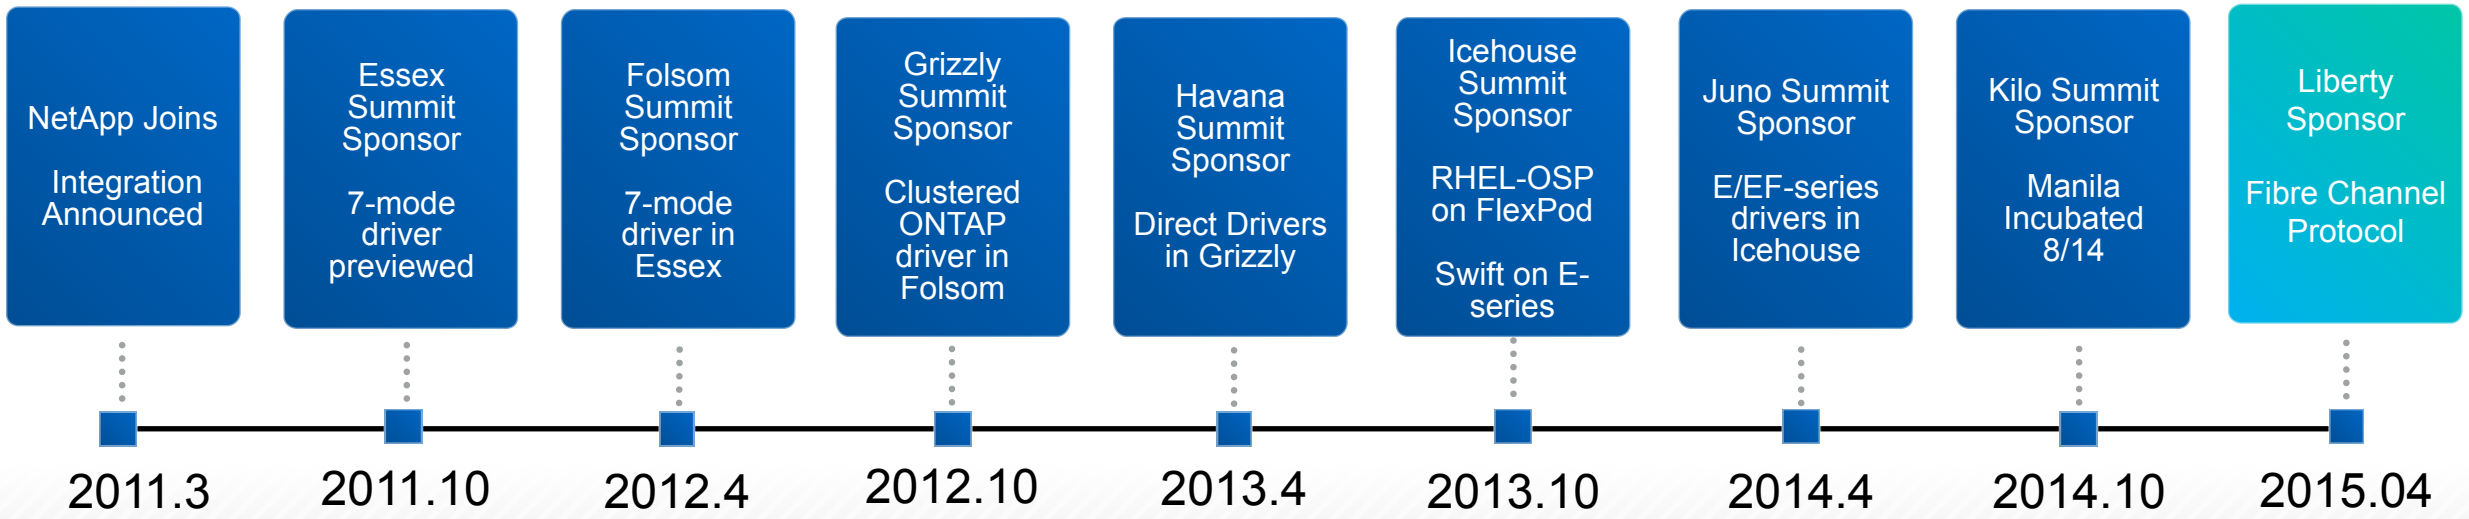

NetApp & OpenStack Deployment

Adoption Accelerating

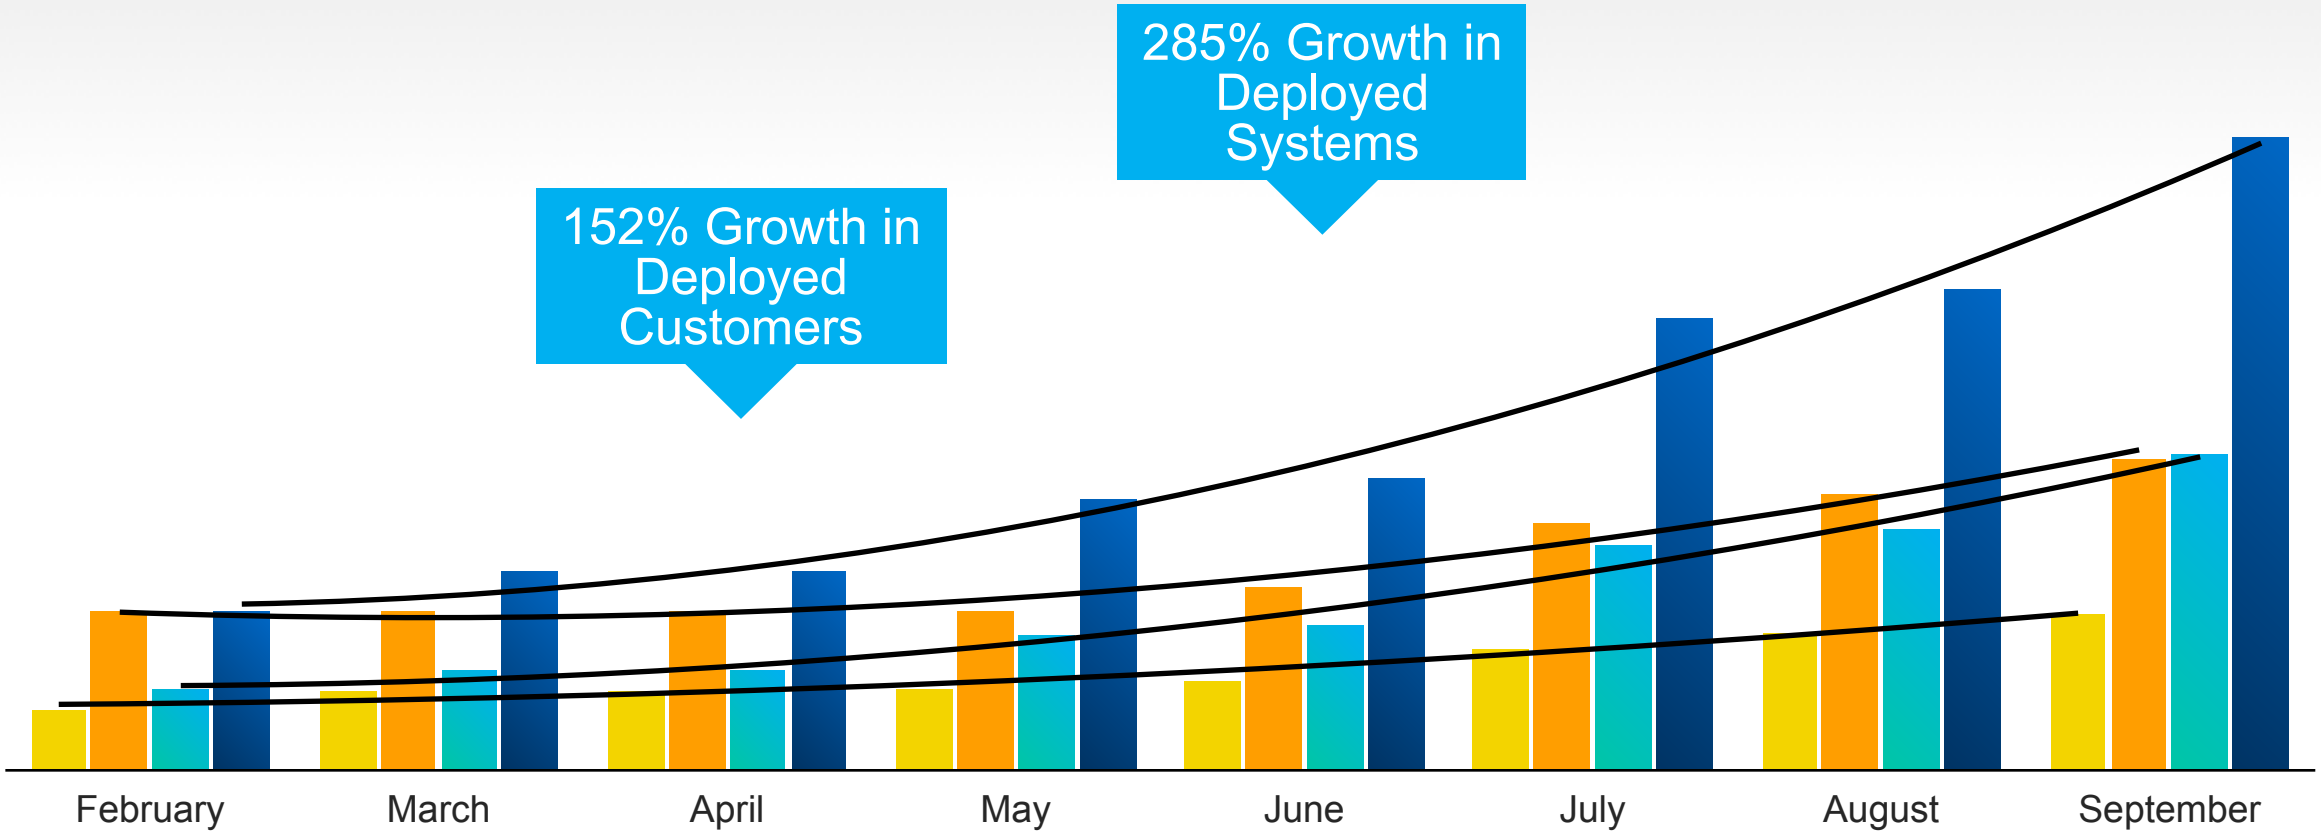

Integration Overview

Manila

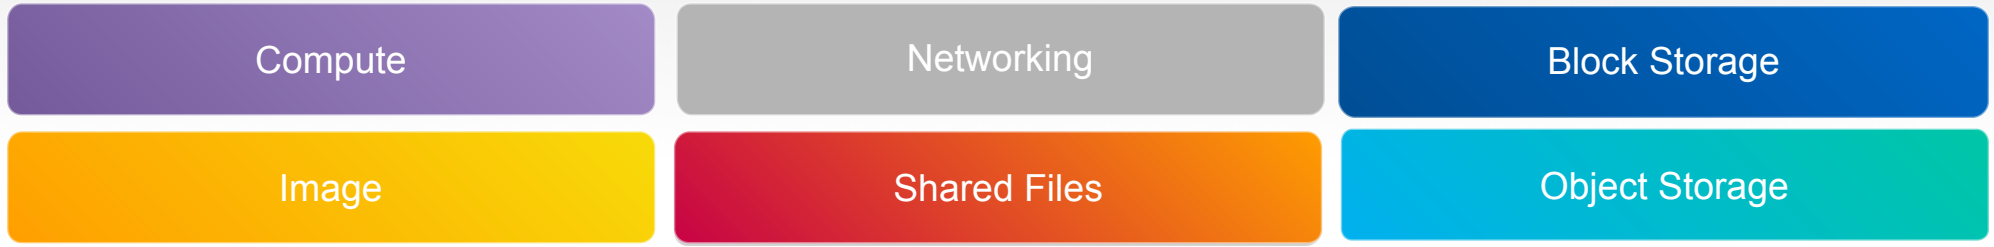

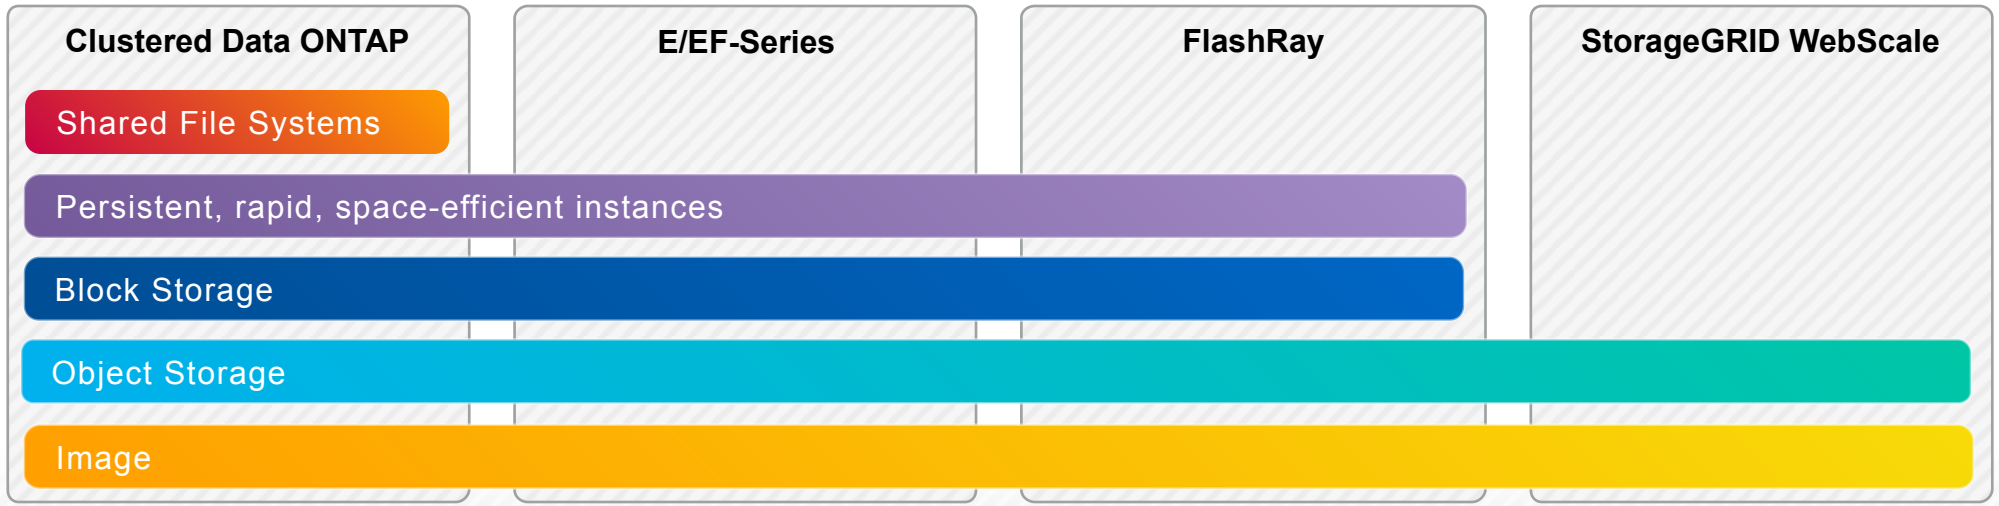

NetApp's Core Competencies

FAS

Continuous Operations

Seamless Scaling

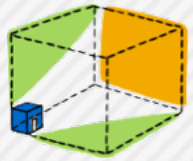

Storage Efficiency

Data Mobility

Data Protection

Unified Architecture

Quality of Service

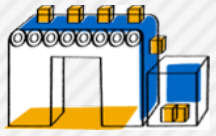

Service Automation

Secure Multi-Tenancy

E-Series

Simplicity

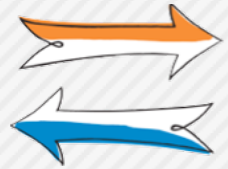

Horizontal Scaling

Consistency

Density

Versatility

Vast Bandwidth

Broad Connectivity

Data Mobility

All-Flash

Glance

- Rapid Cloning
  - **Copy offload** eliminates first network copy from Glance to Cinder host
  - NFS image cache used for subsequent clones.
    - In testing we can FlexClone 20GB in 0.3 second!
- Space Efficiency
  - **Deduplication**: Common 4k blocks are coalesced into a single block
  - When used on Glance image store FlexVol, storage footprint is reduced dramatically
    - Up to 90% disk savings

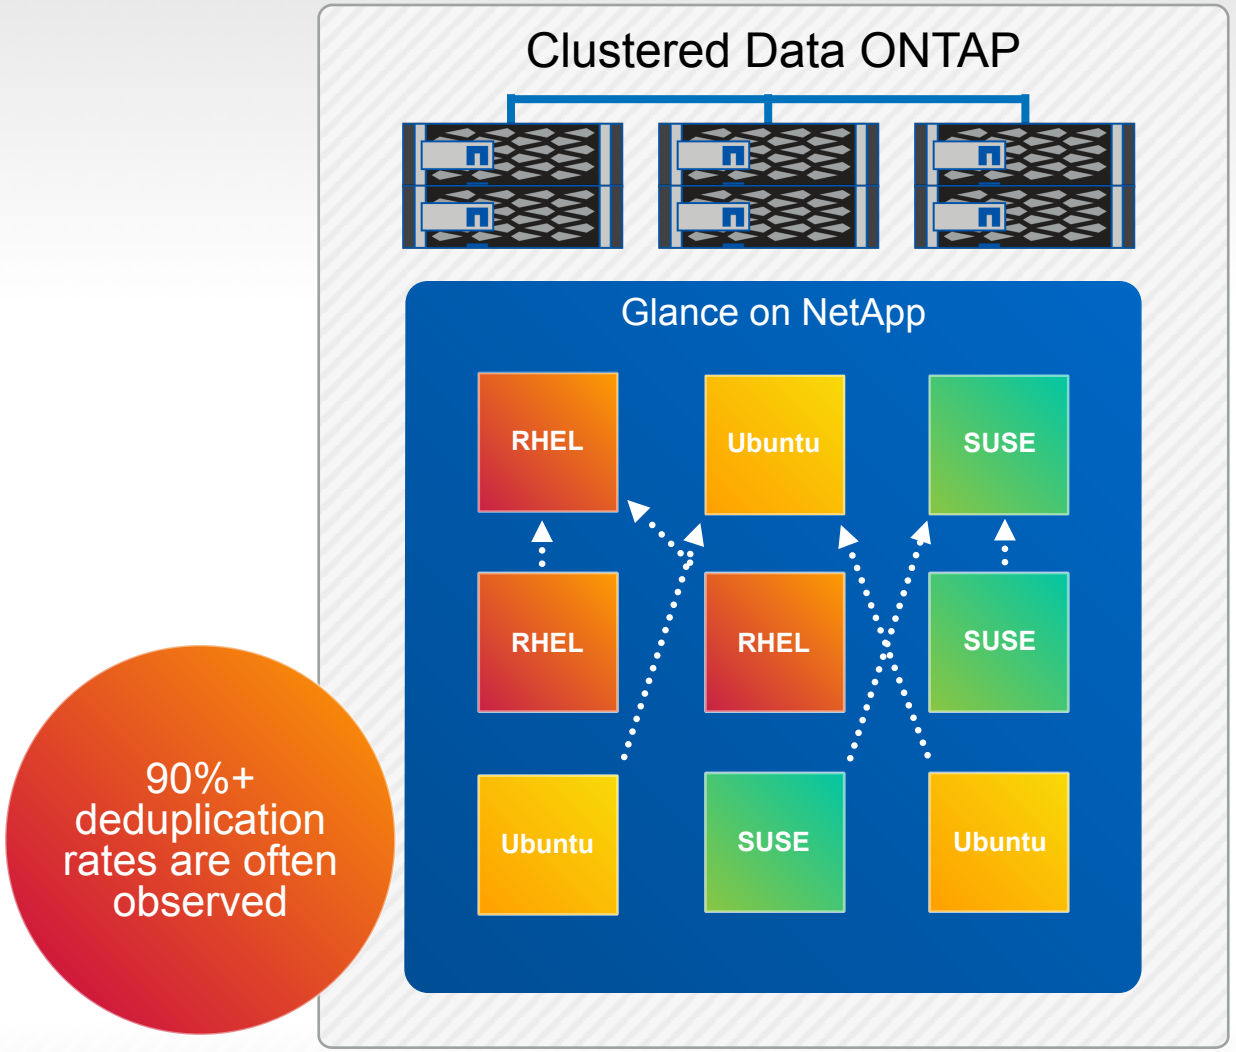

© 2015 NetApp, Inc. All rights reserved.

Cinder deployment with NetApp

Diverse Options for Diverse Use Cases

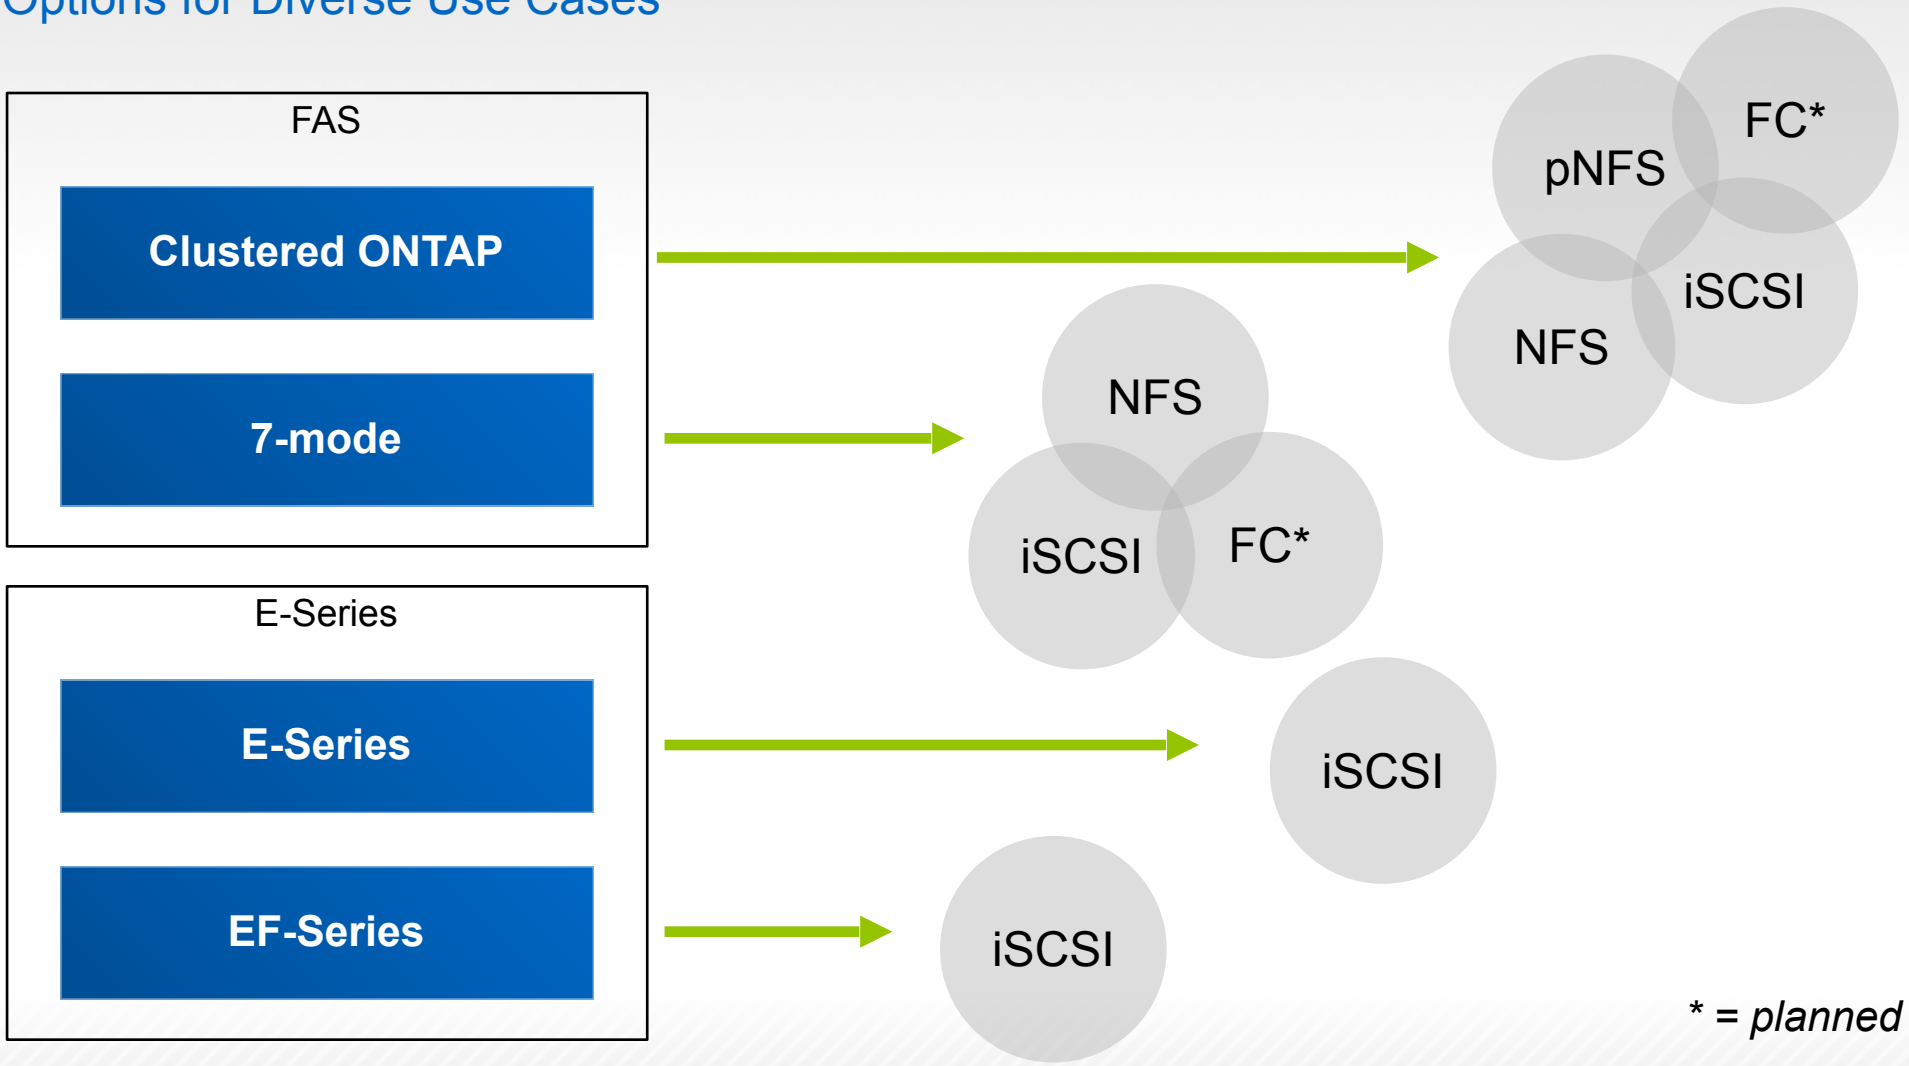

* = *planned for Kilo release*

Deliver workload-aligned block storage offerings

Create a storage service catalog that maps the differentiated features of Data ONTAP to Cinder

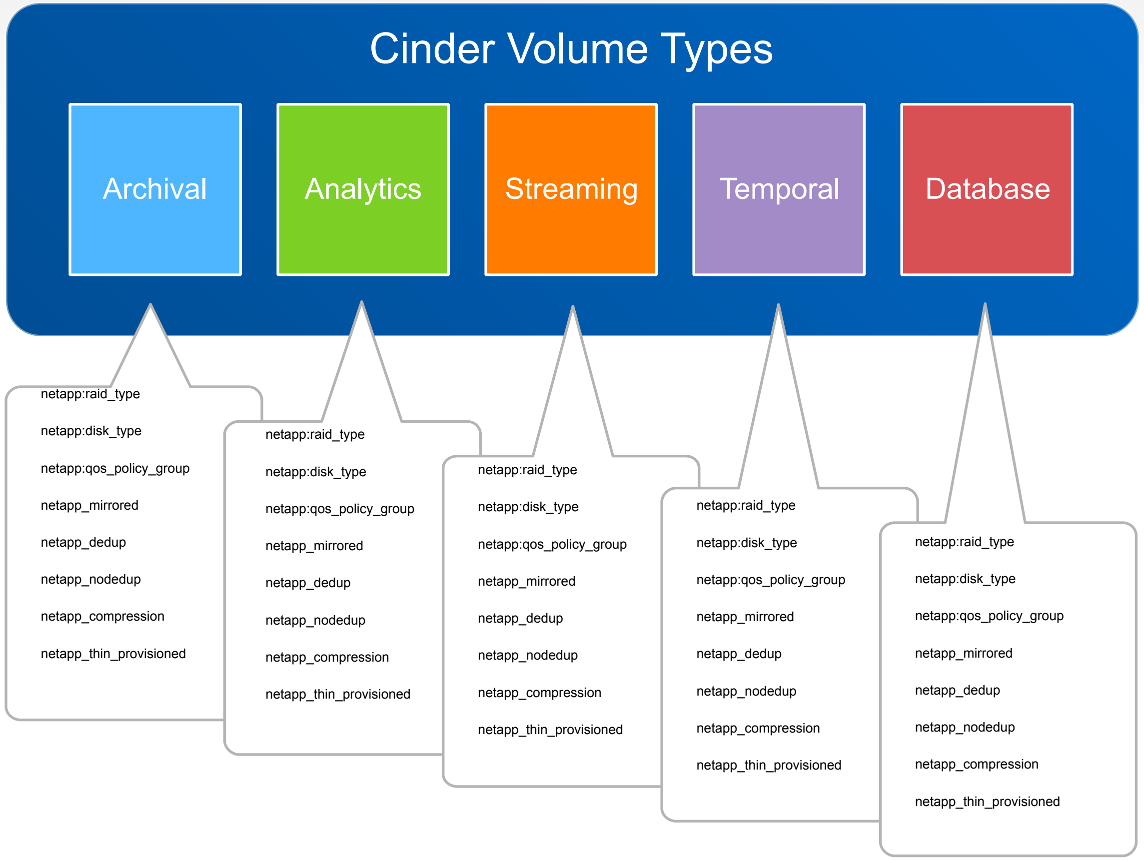

- Map features of underlying NetApp storage to create classes of service
  - Aligned to workloads – for example:
    - Database needs high IOPS with Flash, and data protection
    - Temporal workloads need thin provisioning and deduplication
- Ensure that consumption matches intent
  - Show back, chargeback, etc.

NetApp

Swift

Advantages of NetApp E-Series arrays

- Resiliency
  - As disk sizes increase, so do rebuild times after failure
  - With E-Series DDP, rebuild operations are 8x faster
    - Rebuild traffic offloaded from network to backend storage
- Efficiency of Swift?
  - Swift replicates data 3x times across cluster by default
    - Weight of replication traffic can become limitation to scale
    - As Swift scales, hardware requirements increase linearly
- With E-Series
  - Data replication can be reduced to 1.3x on disk locally
  - Less hardware is needed leading to lower rack space, power, & cooling requirements

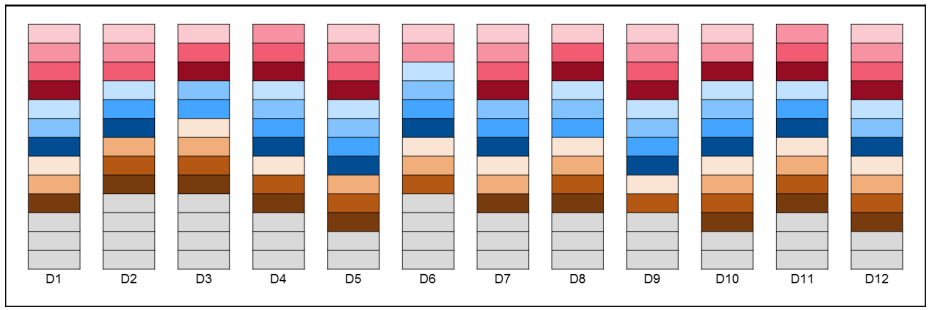

NetApp E-Series DDP

- Dynamic distribution / re-distribution of data“De-clustered” RAID
- Evolution of CRUSH (erasure coding)
- Space and scaling efficiency
- 7 Patents applied for

Swift on NetApp E-Series

Efficient Storage and Scaling with Dynamic Disk Pools

DDP reduces disk rebuild times by

90%

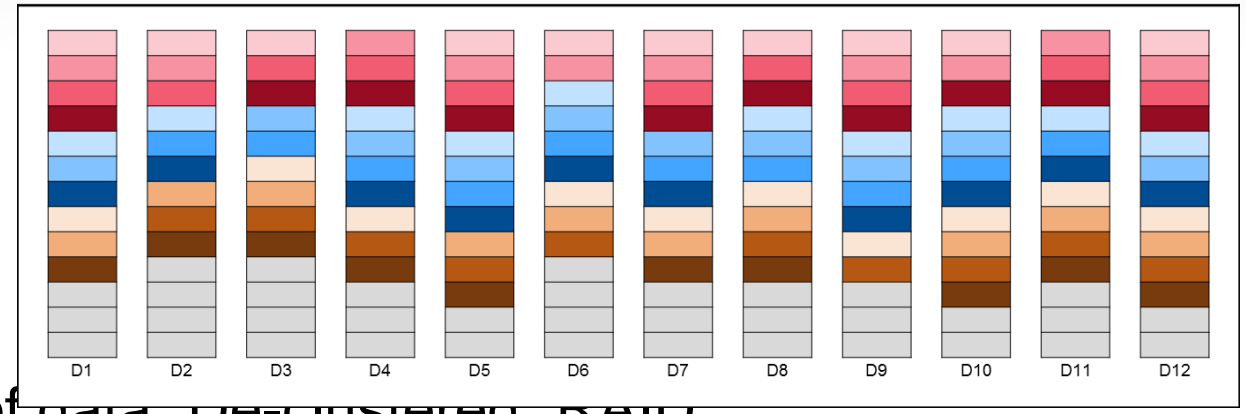

- Dynamic distribution / re-distribution of data "De-clustered" RAID
- Deploying Swift with NetApp E-Series reduces:
  - **Required storage capacity**
  - **Ongoing cost of operations**
  - **Deployment footprint**
  - **Replication traffic between Swift nodes**

Red Hat and NetApp joint efforts in RHEL-OSP6

- Regular syncs to provide for consistency of roadmap and vision
- Puppet modules for managing the Cinder NetApp driver have been integrated in RHEL-OSP since release 5
- The RHEL-OSP installer now has the NetApp Cinder driver exposed in the Cinder GUI deployment pane (RHEL-OSP 6 A1 hotfix or native to RHEL-OSP6 A2 release)
- NetApp internal IT and Engineering OpenStack deployments are live on RHEL-OSP
- Aligning efforts with development teams: Features, bugs, and priorities
- Support – Drivers are certified via Red Hat test suite

Highly Available Deployments of RHEL-OSP 6 and NetApp Storage

Solution Overview

FAS

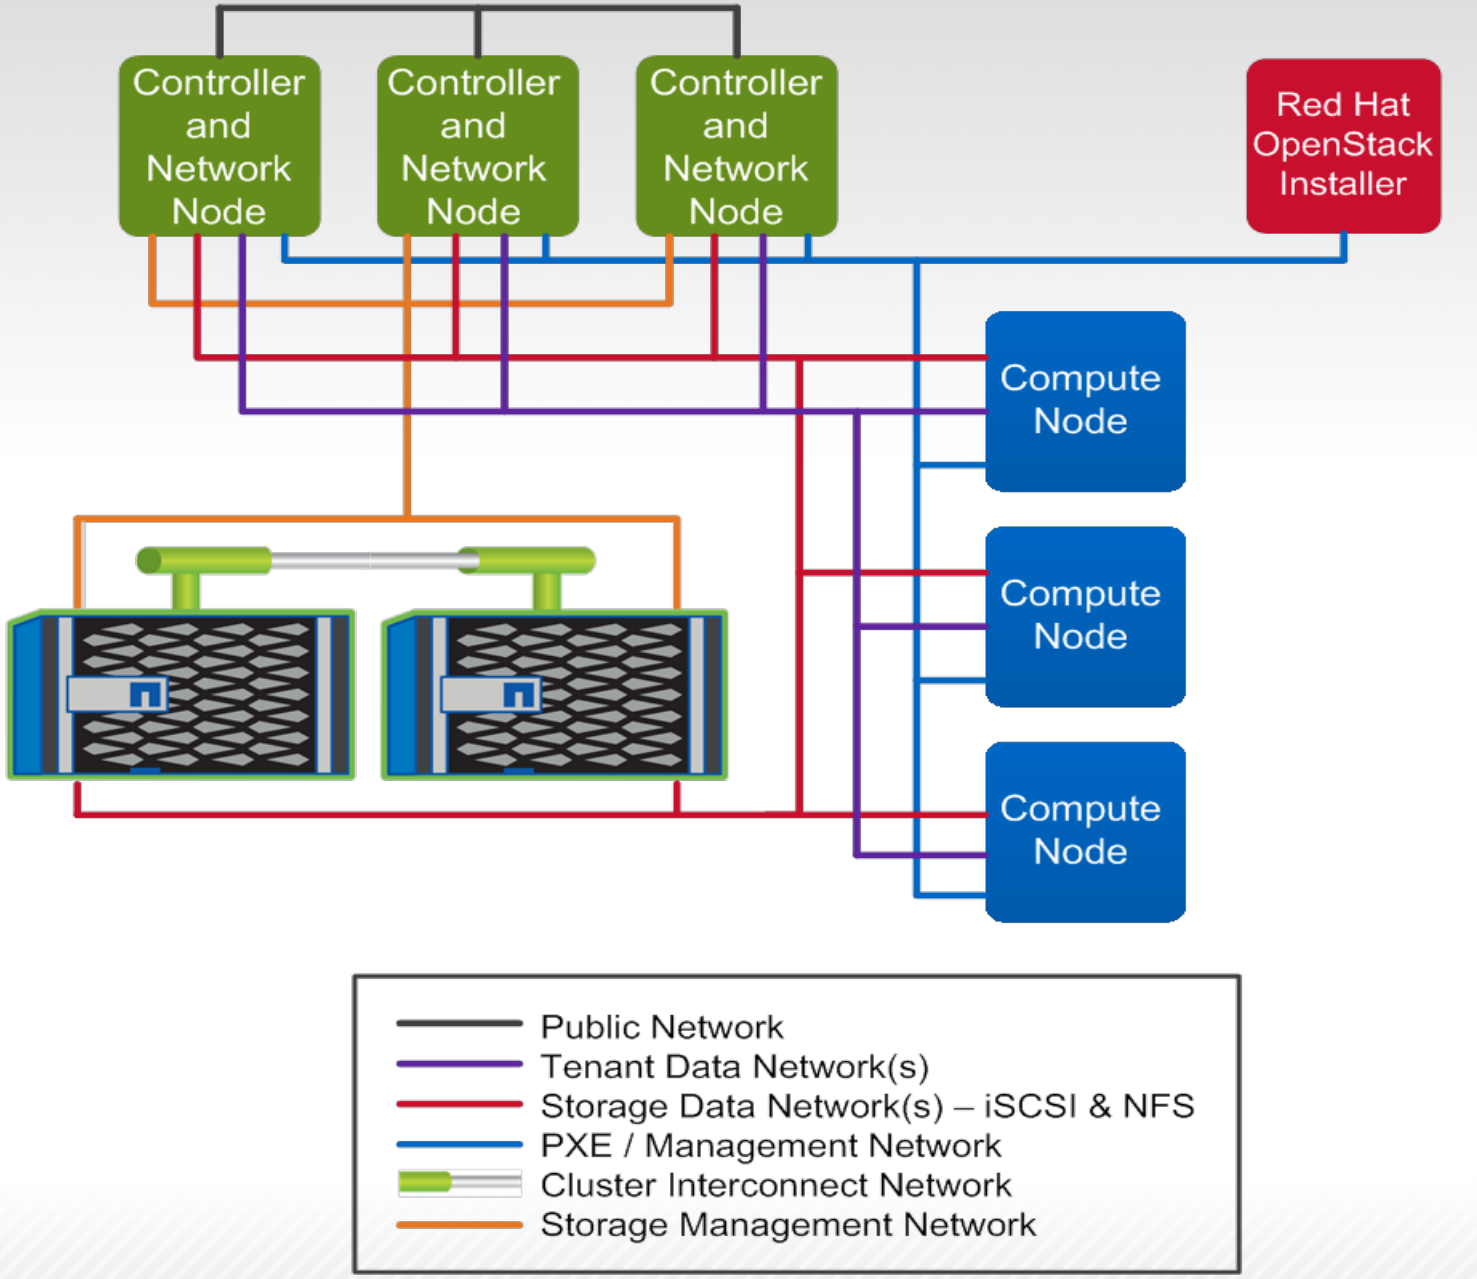

Solution Overview

E-Series

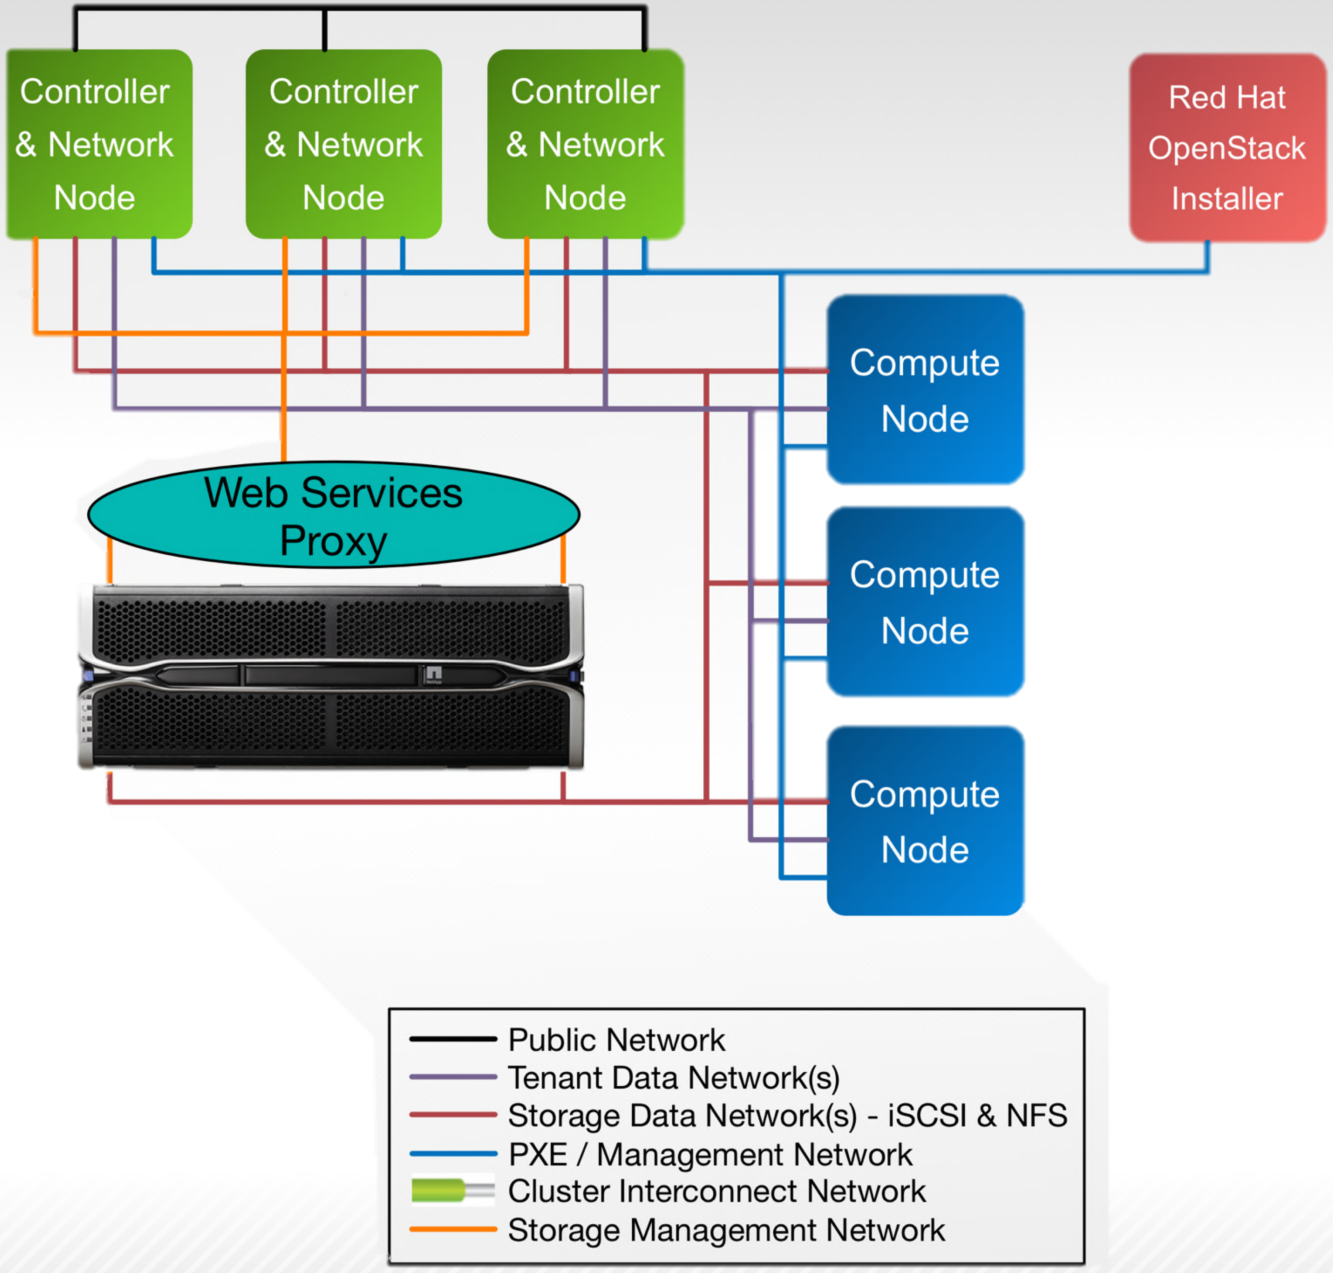

Deployment Choices:

- High Availability
- Networking Subsystem
- Messaging Provider
- Passwords

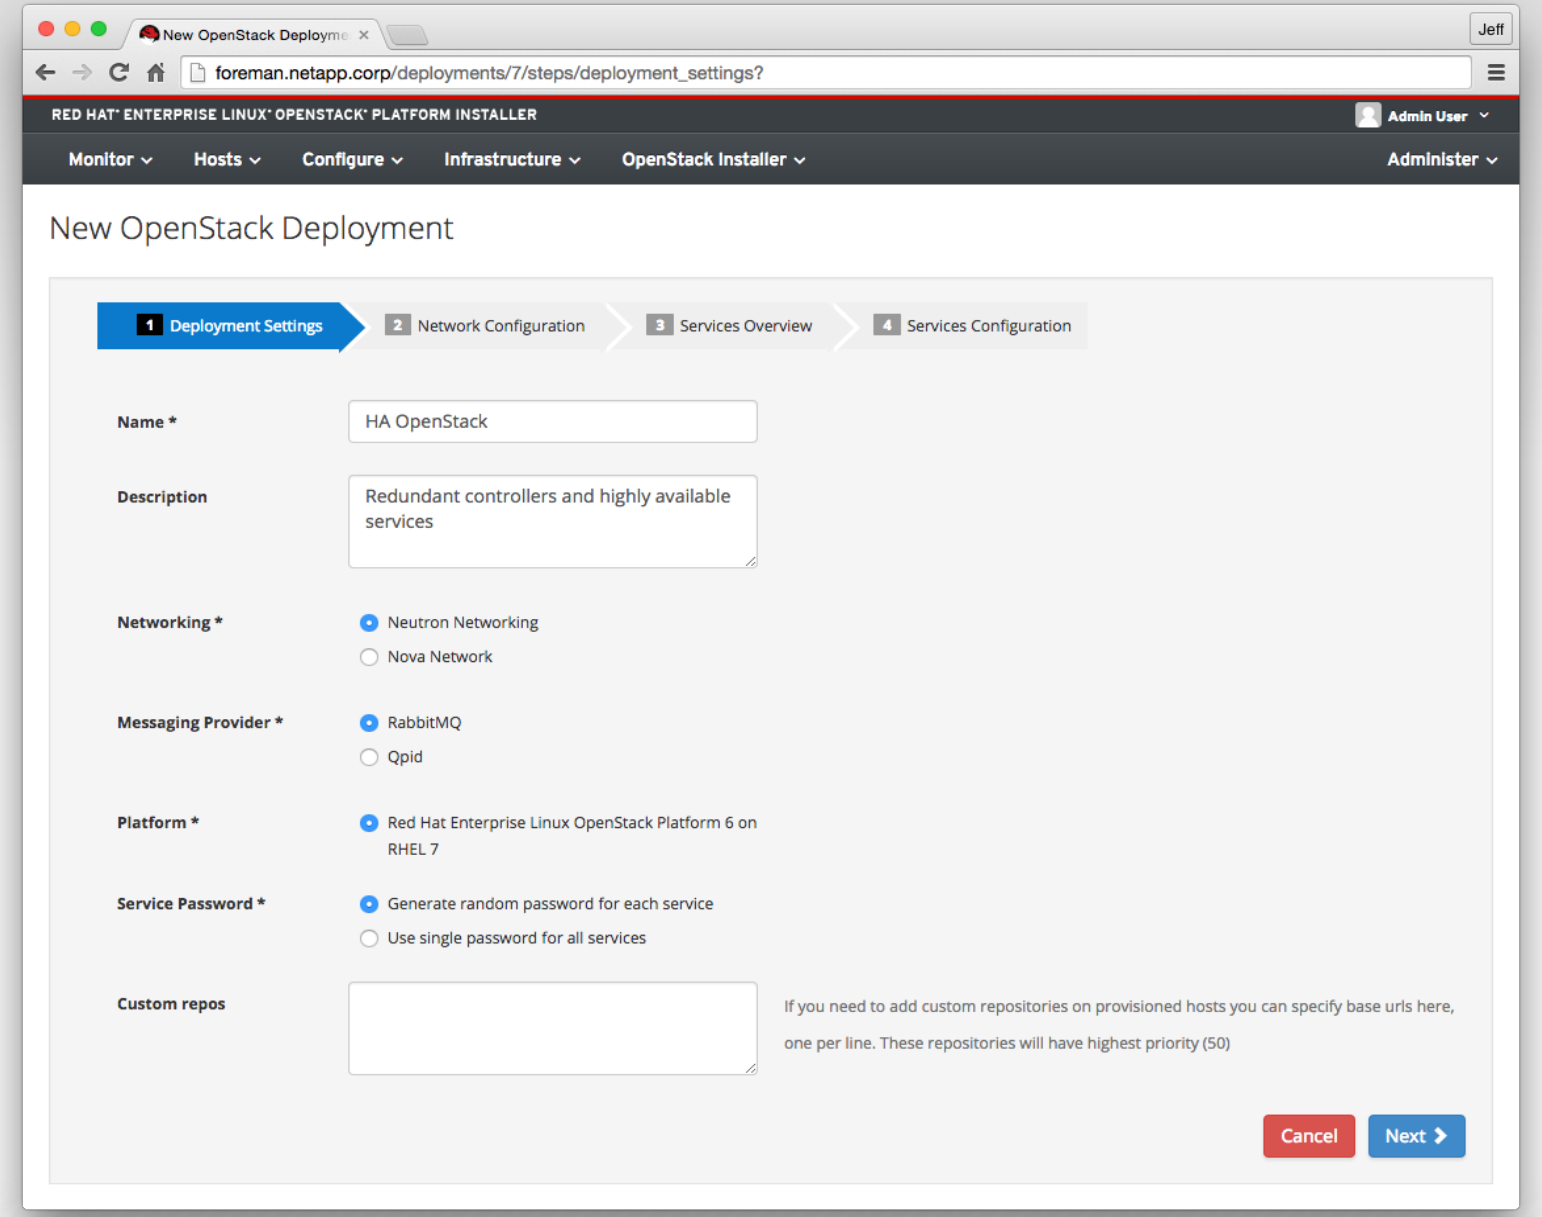

Networking Choices:

- **Click New Subnet to create networks**
- **Drag Networks to their proper location**

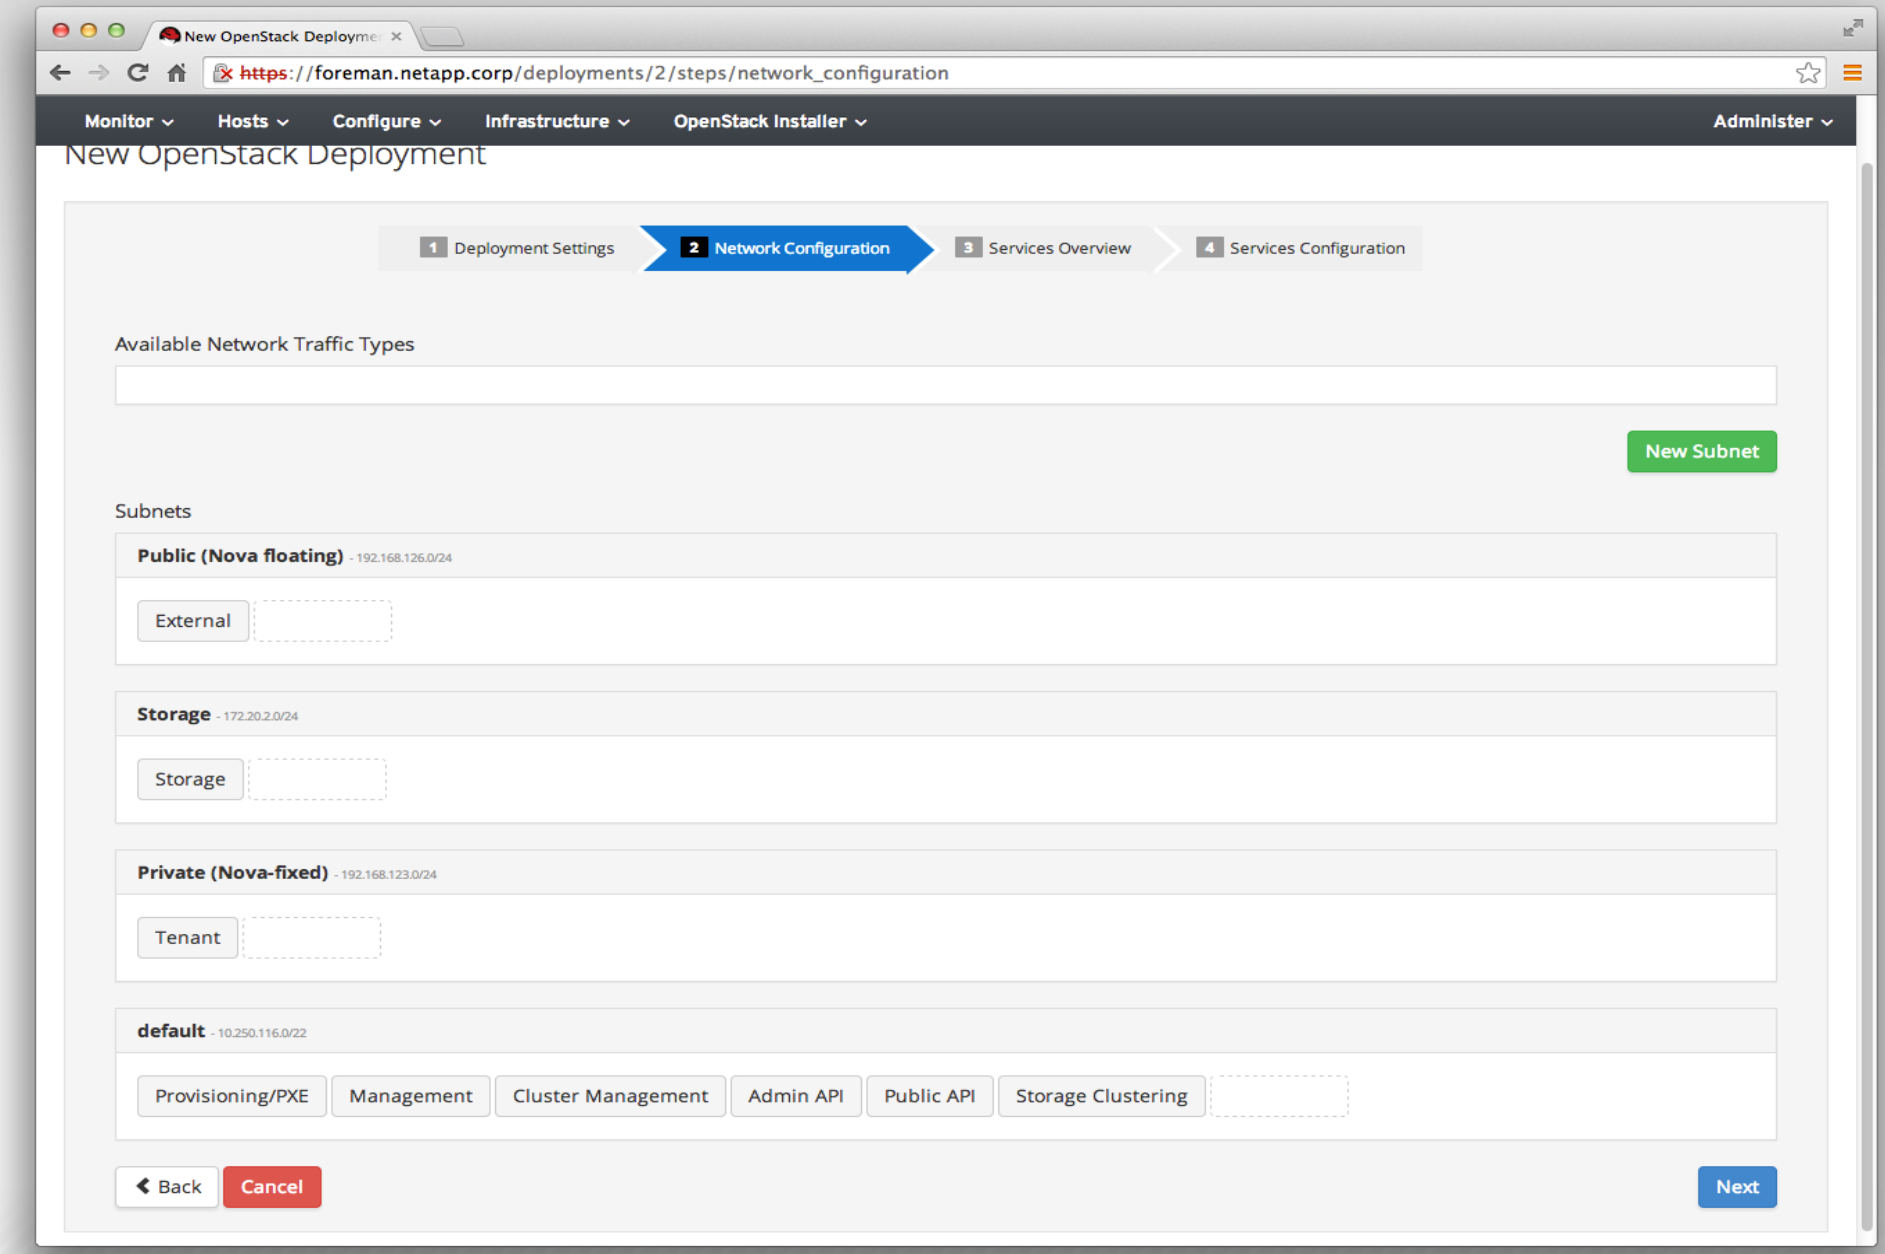

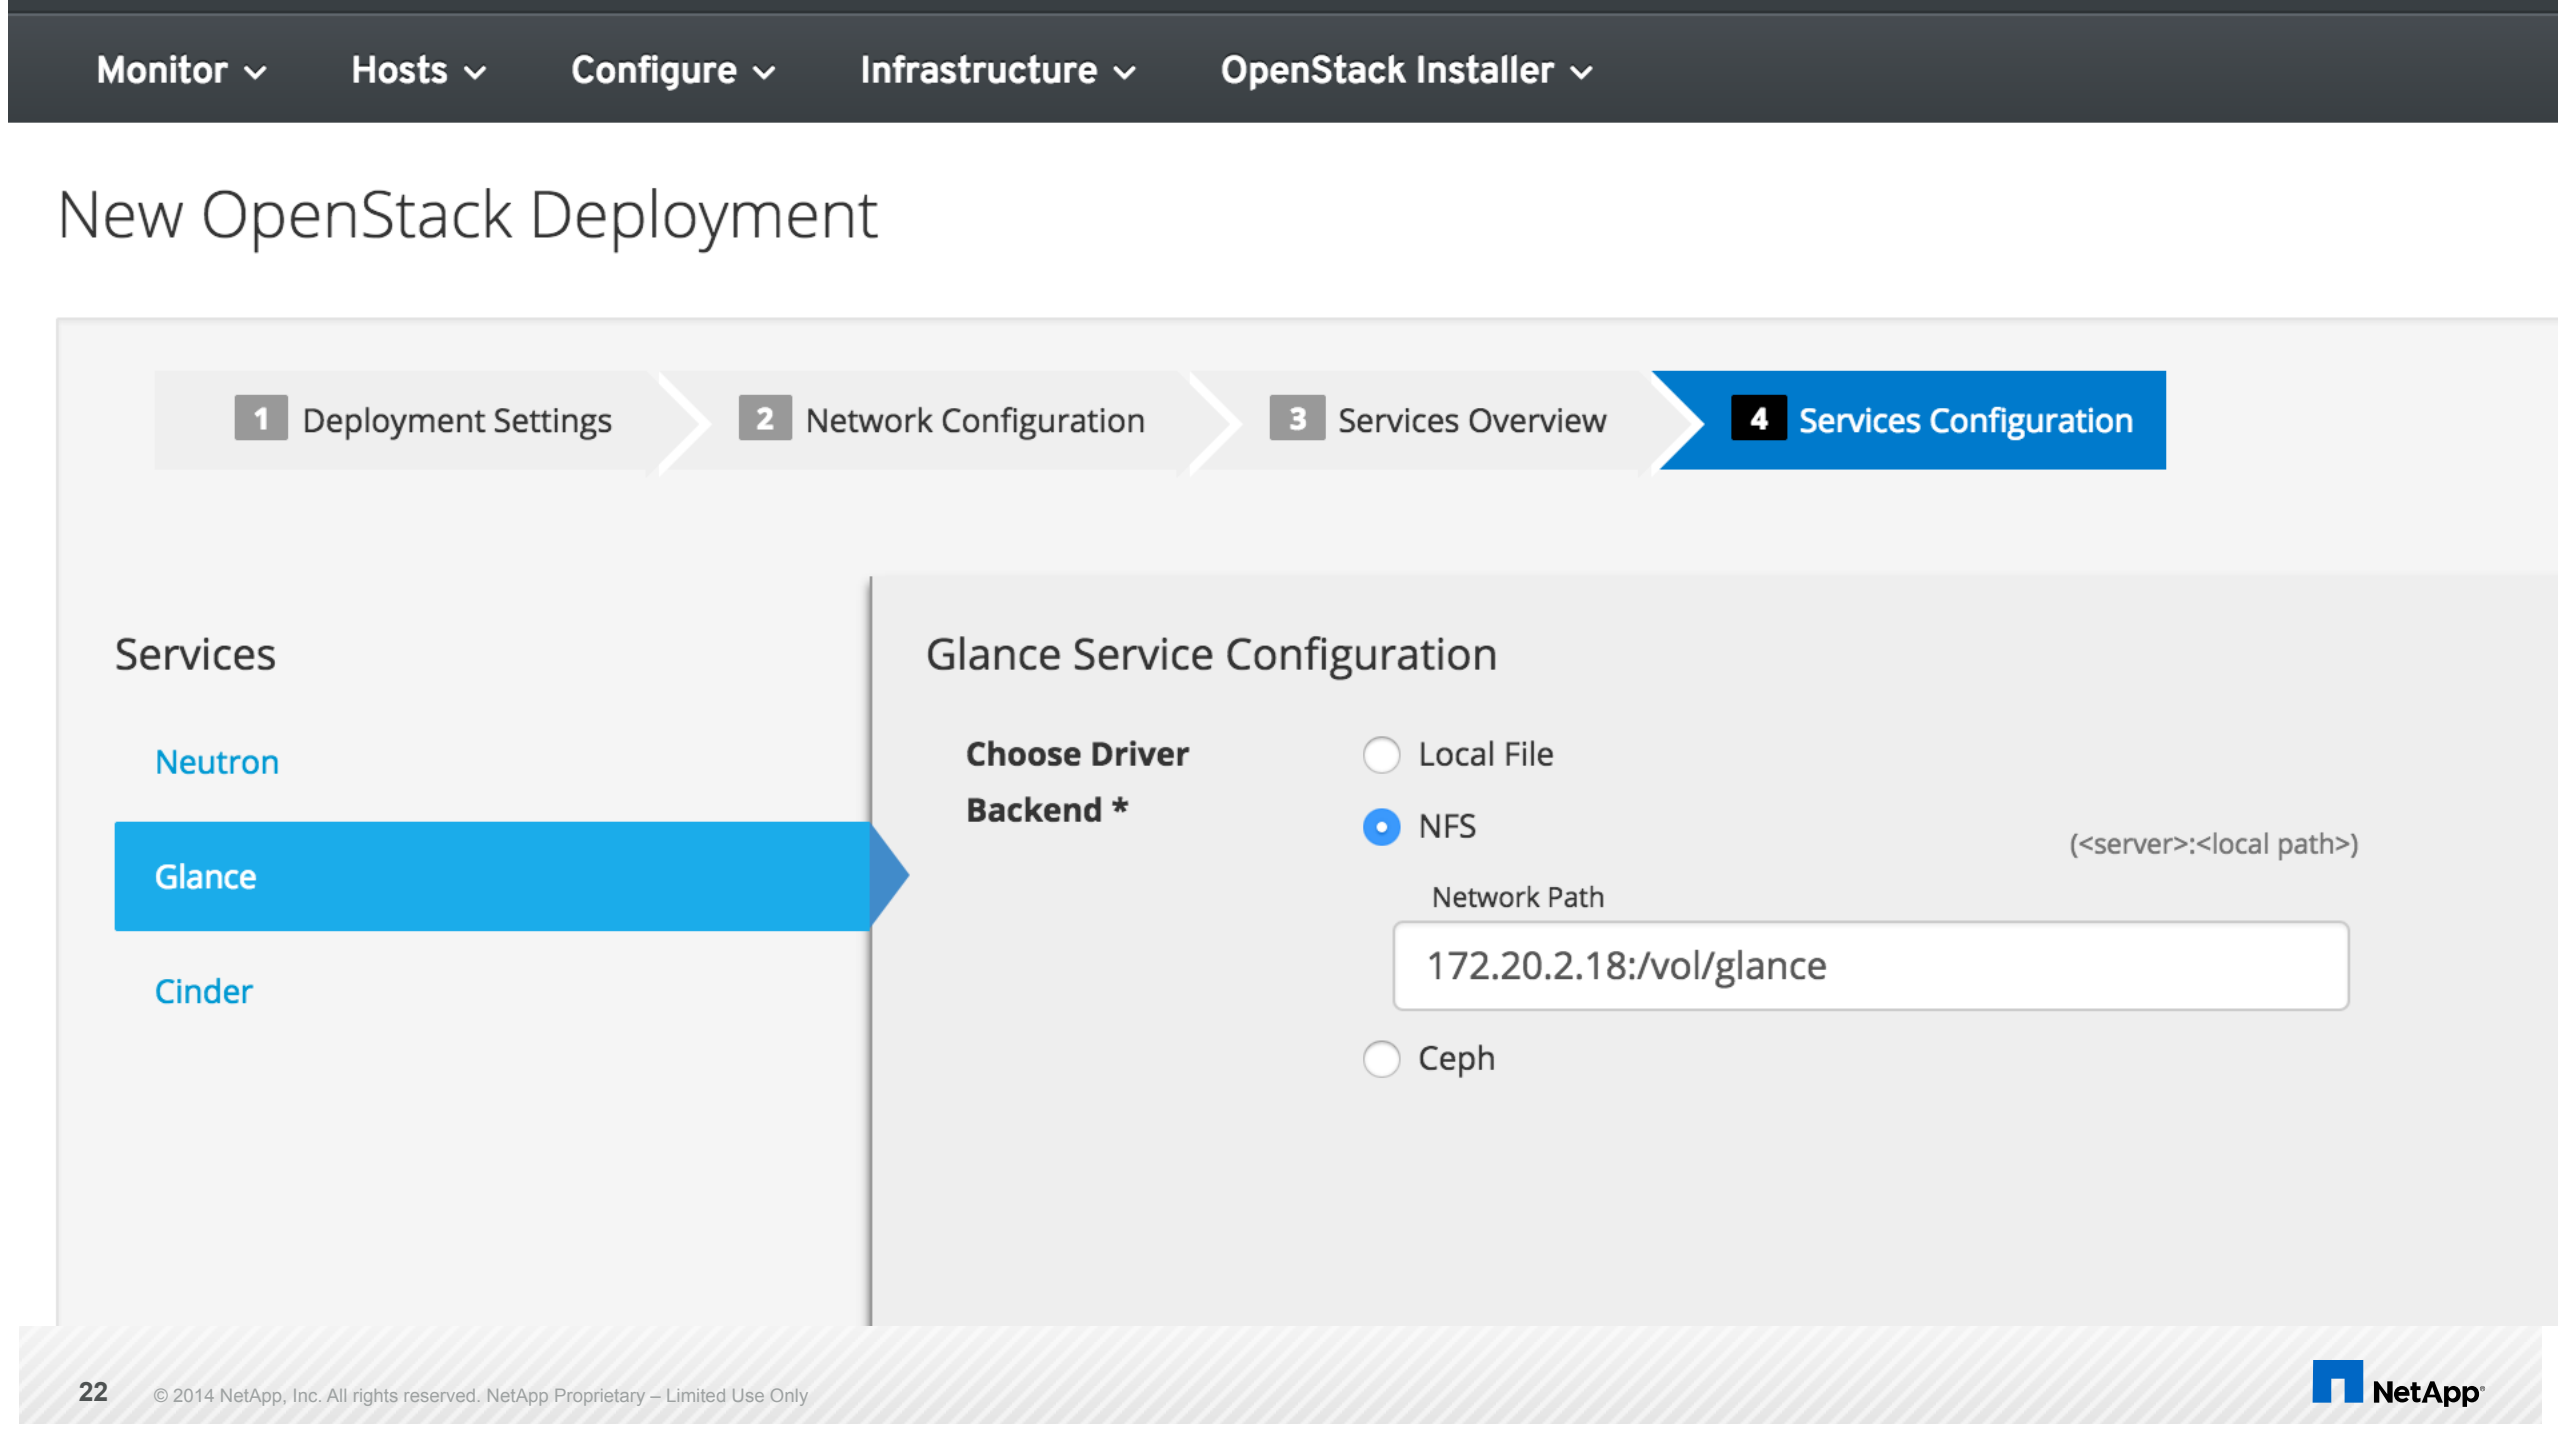

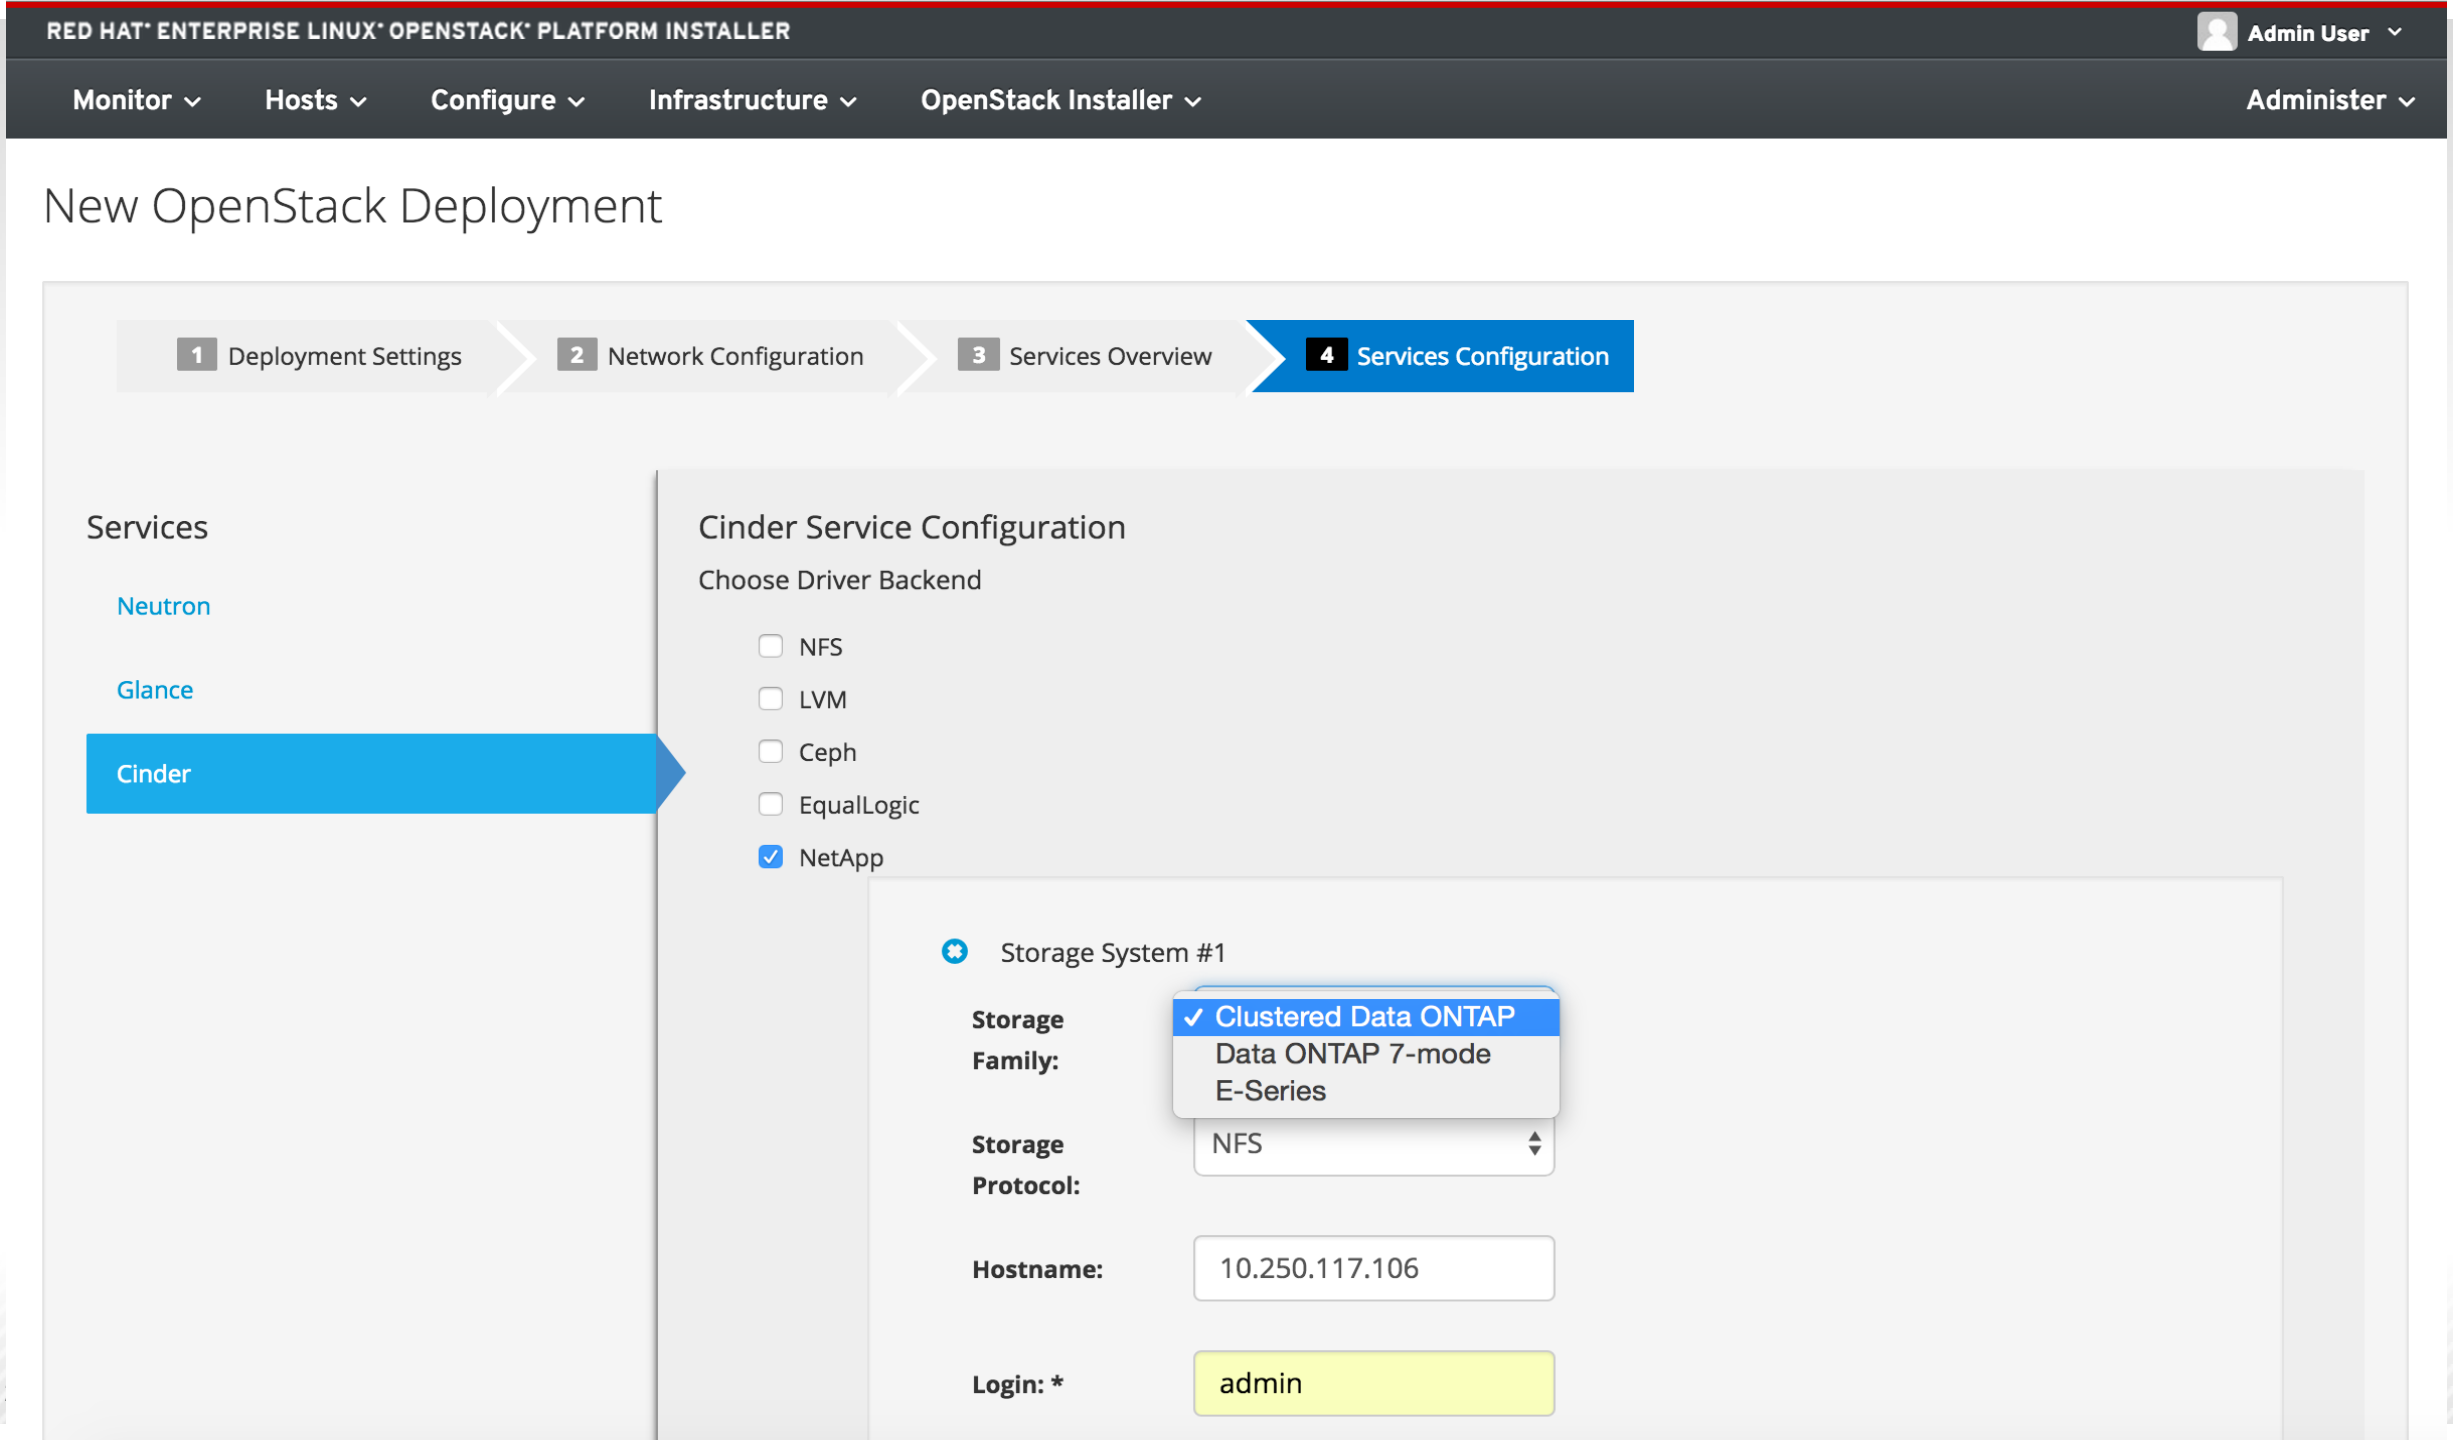
RED HAT® ENTERPRISE LINUX® OPENSTACK® PLATFORM INSTALLER
Admin User
Monitor
Hosts
Configure
Infrastructure
OpenStack Installer
Administer
New OpenStack Deployment
1 Deployment Settings
2 Network Configuration
3 Services Overview
4 Services Configuration
Services
Neutron
Glance
Cinder
Cinder Service Configuration
Choose Driver Backend
NFS
LVM
Ceph
EqualLogic
NetApp
Storage System #1
Storage Family:
Clustered Data ONTAP
Data ONTAP 7-mode
E-Series
Storage Protocol:
NFS
Hostname:
10.250.117.106
Login: *
admin

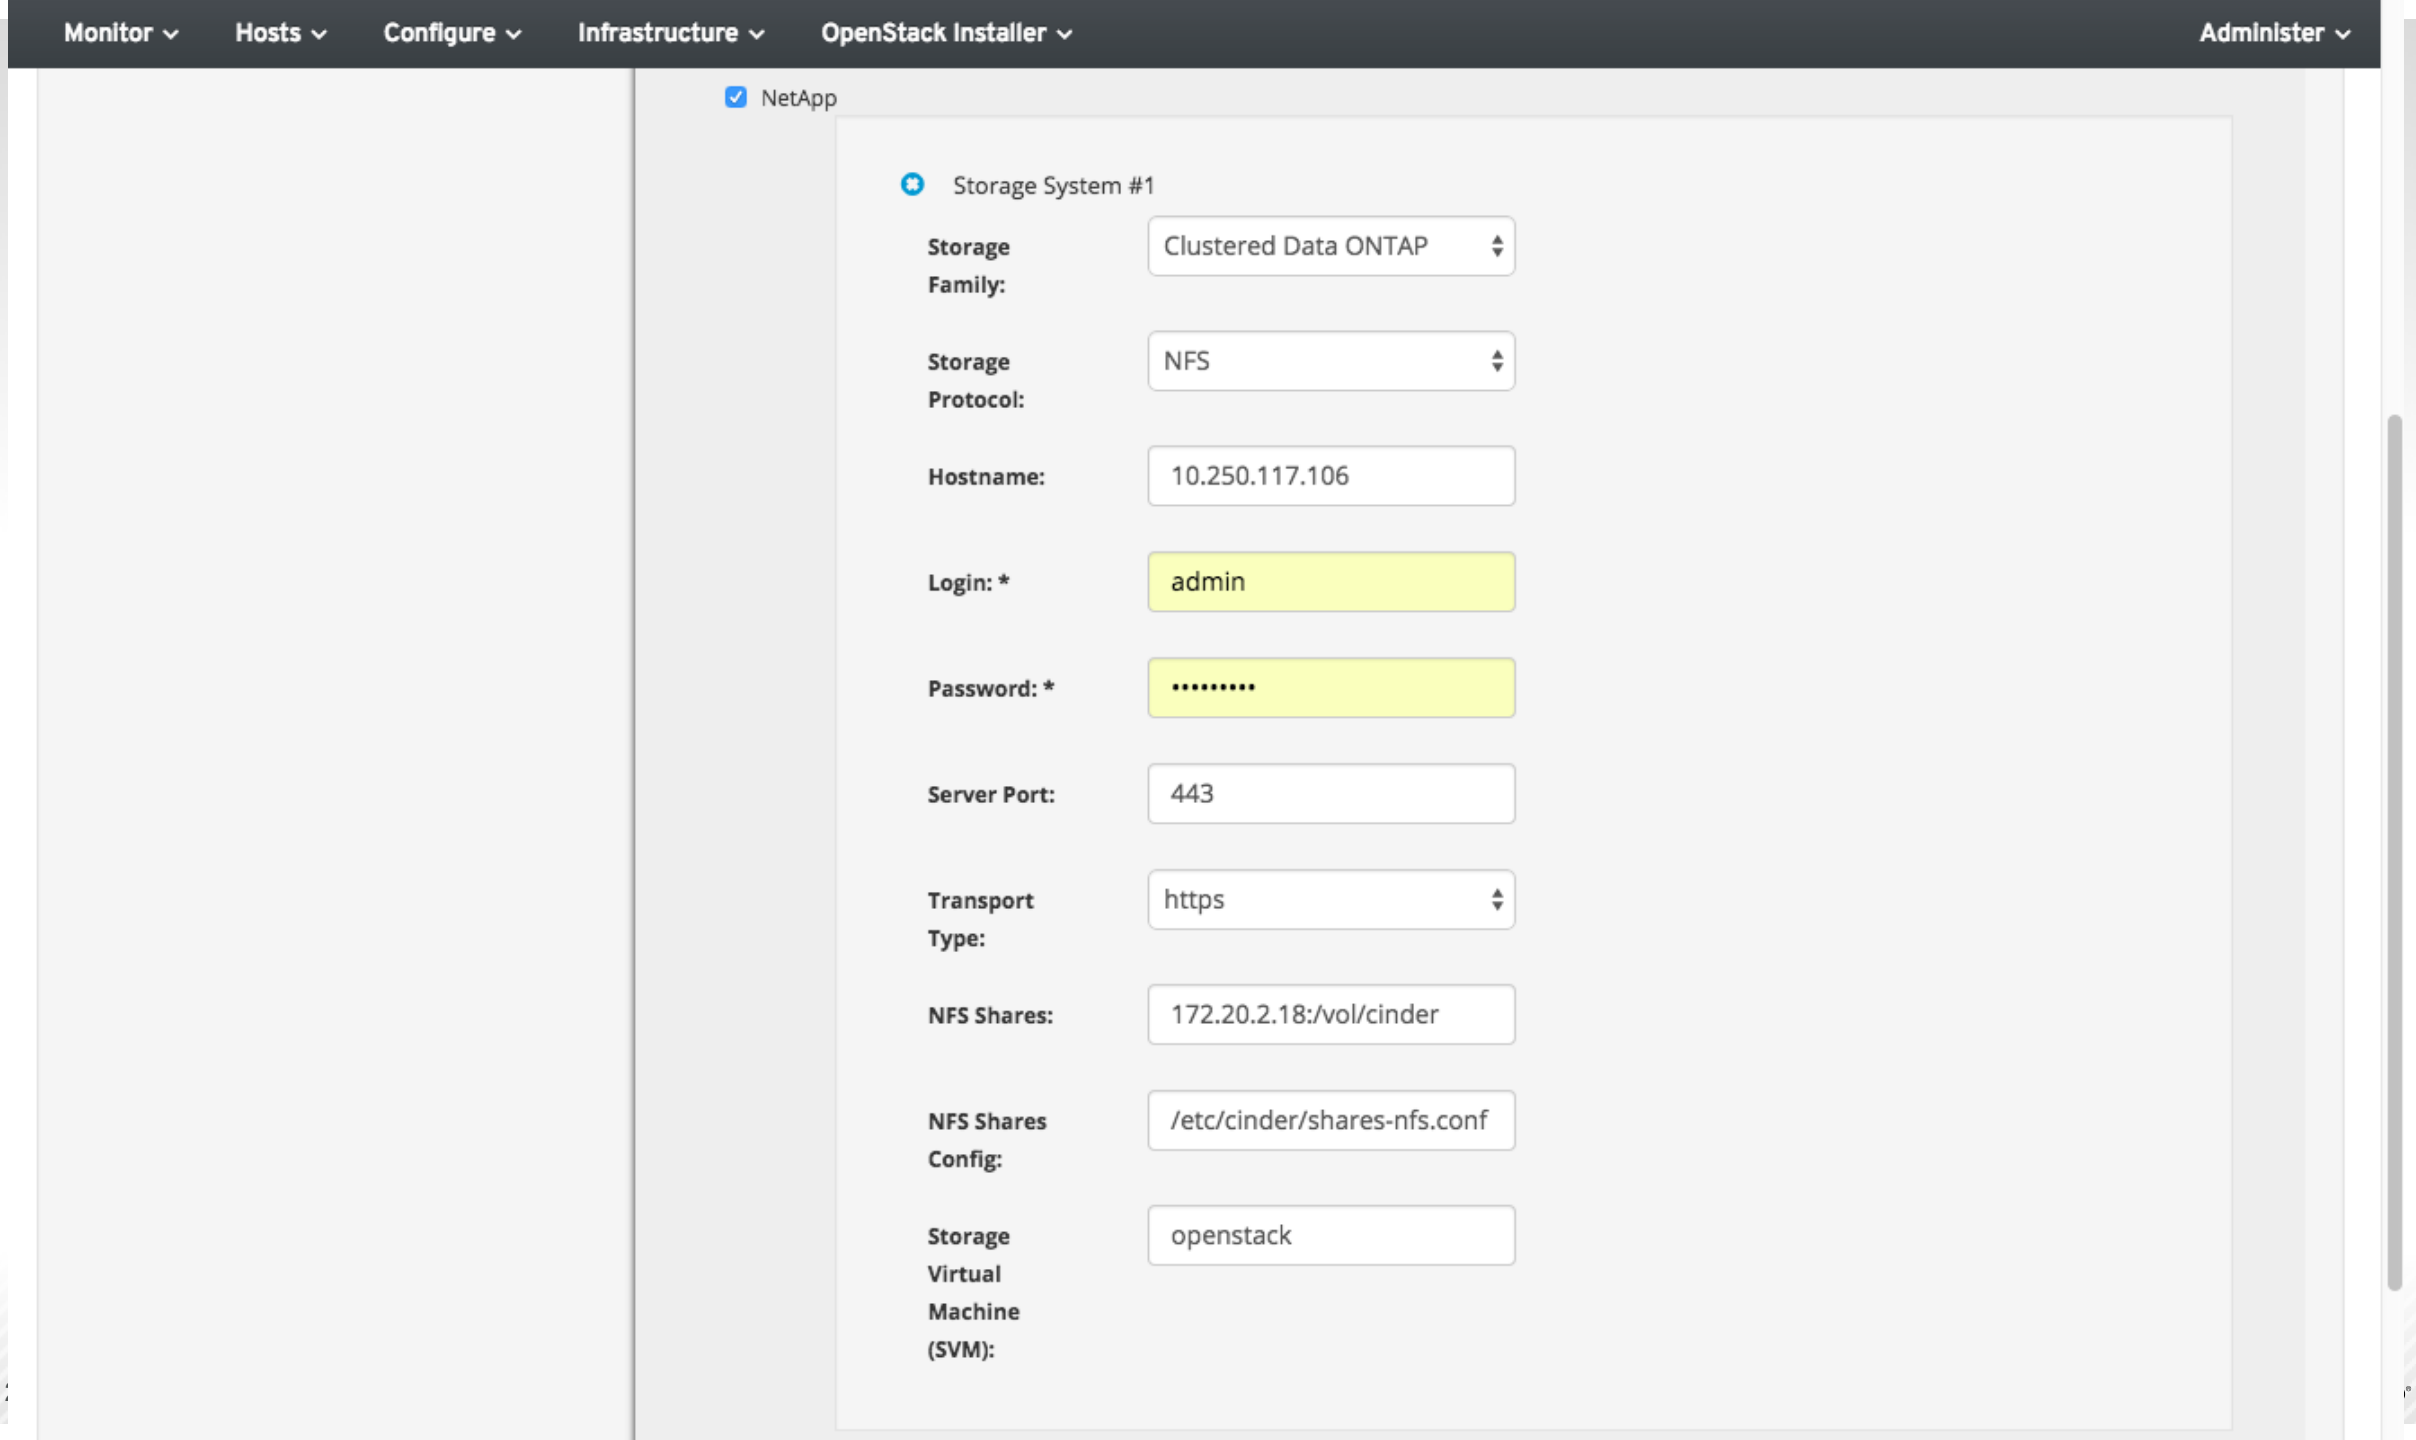
Monitor
Hosts
Configure
Infrastructure
OpenStack Installer
Administer
NetApp
Storage System #1
Storage Family:
Clustered Data ONTAP
Storage Protocol:
NFS
Hostname:
10.250.117.106
Login: *
admin
Password: *
Server Port:
443
Transport Type:
https
NFS Shares:
172.20.2.18:/vol/cinder
NFS Shares Config:
/etc/cinder/shares-nfs.conf
Storage Virtual Machine (SVM):
openstack

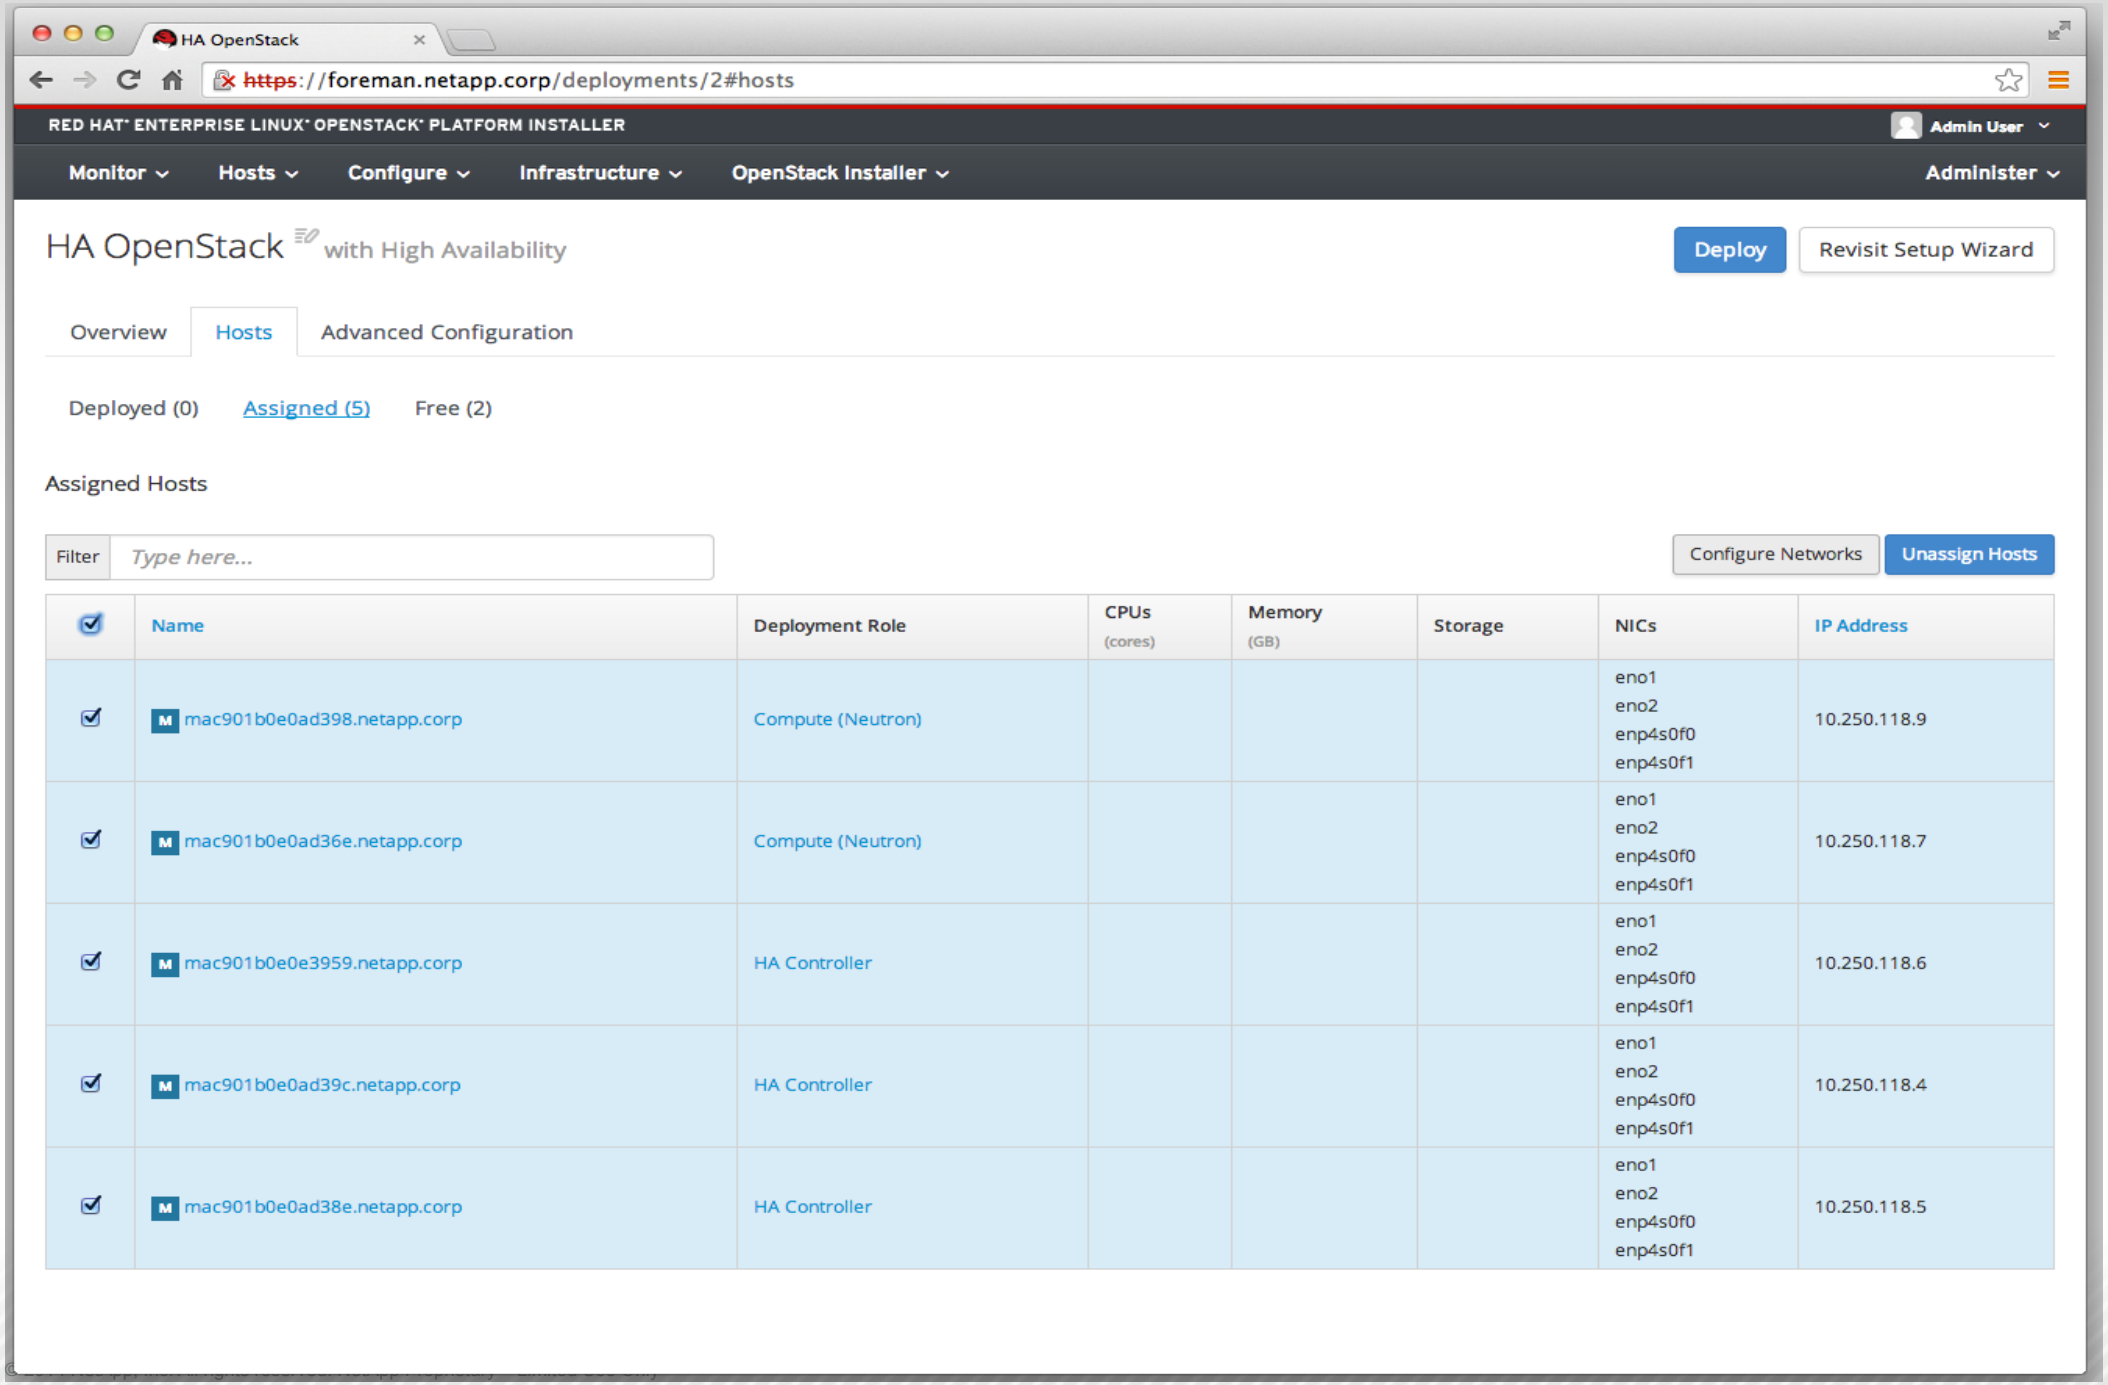
HA OpenStack
https://foreman.netapp.corp/deployments/2#hosts
RED HAT ENTERPRISE LINUX OPENSTACK PLATFORM INSTALLER
Admin User
Monitor
Hosts
Configure
Infrastructure
OpenStack Installer
Administer
HA OpenStack with High Availability
Deploy
Revisit Setup Wizard
Overview
Hosts
Advanced Configuration
Deployed (0)
Assigned (5)
Free (2)
Assigned Hosts
Filter
Type here...
Configure Networks
Unassign Hosts
Name
Deployment Role
CPUs (cores)
Memory (GB)
Storage
NICs
IP Address
mac901b0e0ad398.netapp.corp
Compute (Neutron)
eno1 eno2 enp4s0f0 enp4s0f1
10.250.118.9
mac901b0e0ad36e.netapp.corp
Compute (Neutron)
eno1 eno2 enp4s0f0 enp4s0f1
10.250.118.7
mac901b0e0e3959.netapp.corp
HA Controller
eno1 eno2 enp4s0f0 enp4s0f1
10.250.118.6
mac901b0e0ad39c.netapp.corp
HA Controller
eno1 eno2 enp4s0f0 enp4s0f1
10.250.118.4
mac901b0e0ad38e.netapp.corp
HA Controller
eno1 eno2 enp4s0f0 enp4s0f1
10.250.118.5

Drag networks to the bond0 interface which has 2x 10GbE NICs.

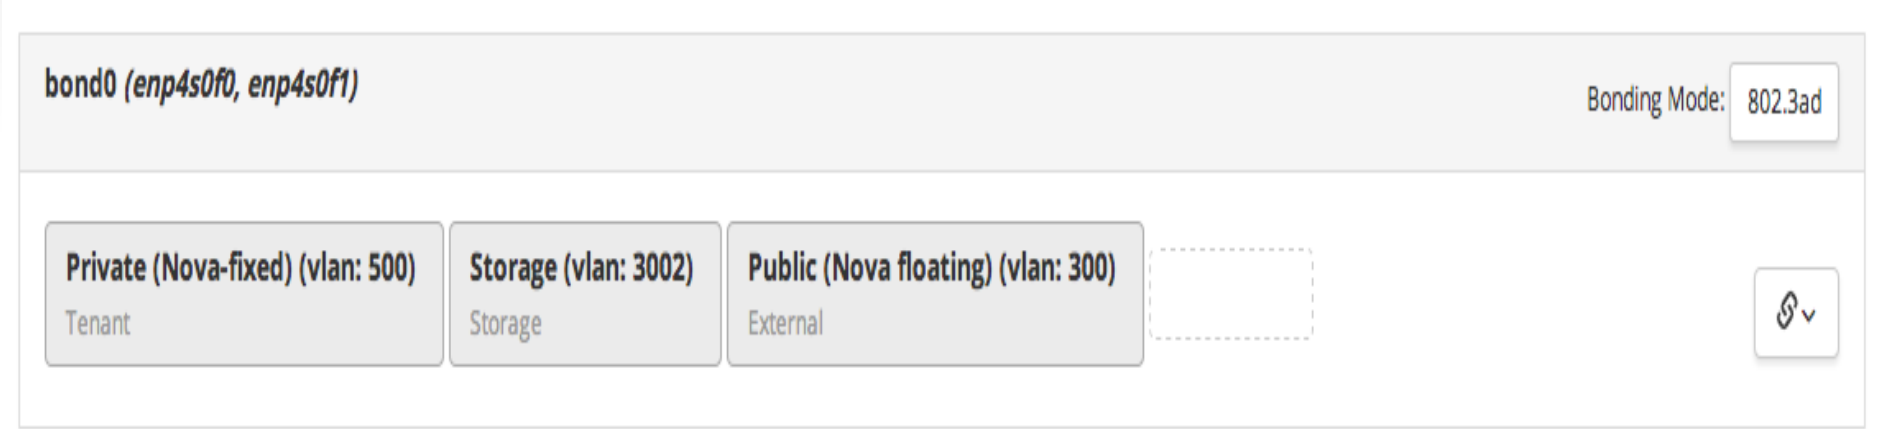

Note: 802.3ad performed best
In our tests with no port errors. See referenced deployment guide for configuration details.

Lessons learned

- **Check Twice, Deploy Once**
  - Verify networking: Verify that the installer and the physical network is correctly setup.
  - Pay close attention to iptables and routing setup on installer.
  - While testing deployments you can prevent Puppet from overriding your changes by setting the immutable flag on a file ex. `# chattr +i /etc/cinder/cinder.conf`
- **Start Simple:** Test with a single controller node and ensure this basic build is successful.
  - Delete deployment, delete discovered hosts and start again.
  - A successful build will stay at 30% quite long. Watch top to see yum, puppet, and then later OpenStack processes running on the node.
- **If Deployment Fails:** Troubleshoot failed nodes on the console of the node
  - `# puppet agent –t --debug`
  - Look at advanced section of deployment to see how variables evaluate. Correct errors.

Enterprise Deployments

Why FlexPod for Red Hat Enterprise Linux OpenStack Platform 6?

FlexPod Platform

- Converged Infrastructure solution developed by NetApp® and Cisco®
- NetApp FAS, Cisco UCS® and Cisco Nexus® switch components

NetApp® FAS + Cisco Nexus® + Cisco UCS®

Verified and Validated Architecture

Cisco® Validated Design (CVD)

NetApp® Verified Architecture (NVA)

- Detailed planning stage
- Collaborative design
- End-to-end validation
- Consistent documentation

OpenStack on FlexPOD

- Speed up Cloud Deployment
- Deliver on Enterprise SLAs
- Increase Cloud Reliability
- Improve Security and Compliance
- Reduce Cloud Implementation Risks
- Take Advantage of Comprehensive Cloud Support
- Create an Open Hybrid Cloud Foundation

Red Hat, NetApp, and Cisco:
Partnering for Proven Excellence

FlexPod: Full-Stack Best of Breed

- Compute
  - Server abstraction with Cisco UCS Service Profiles for easily scalable systems
  - iSCSI SAN boot eliminates local drives in compute nodes for stateless booting
  - Enterprise-class hypervisor with RHEL KVM
- Networking
  - Industry standard and feature-leading Cisco Nexus switching
  - OpenStack Neutron ML2/VXLAN or ML2/Nexus modular drivers in RHEL-OSP
- Storage
  - NetApp Cinder driver configured automatically with RHEL-OSP Installer
  - Unified, scale-out storage: block, NAS, hybrid, all-flash
  - Swift Object Storage on NetApp E-Series array

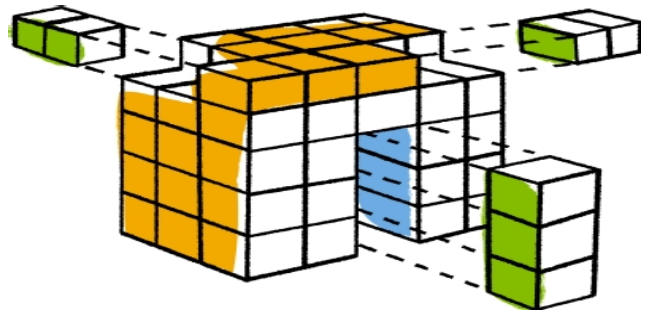

FlexPod: High Availability Out of the Box

- Redundant components
  - Multipath everywhere
  - Dual fabrics
  - Dual storage and network infrastructure devices
- Seamless Upgrades
  - Cisco UCS firmware for compute and network
  - NetApp Data ONTAP OS and firmware for storage
- Nondisruptive Operations
  - Live migration of storage interfaces and volumes across cluster
  - On-line expansion and contraction of compute and storage clusters

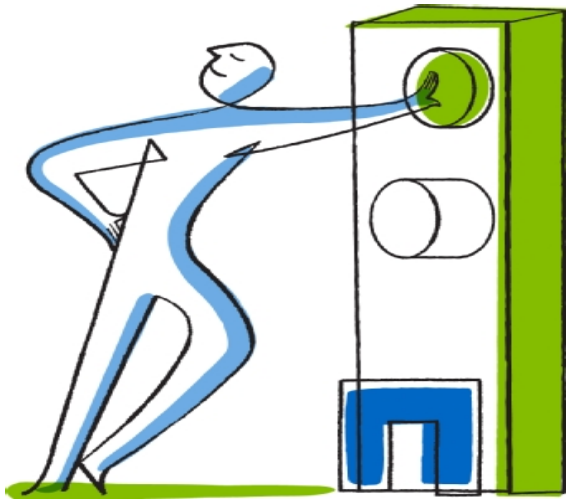

FlexPod: Scaling Up and Scaling Out

- Compute
  - Up to 4 CPUs per server
  - Up to 6TB RAM per server
  - Up to 160 half-width servers in a single UCS domain
  - Multi-UCS domain management with UCS Director
- Storage
  - Up to 8.4PB in a single HA pair
  - Up to 33PB in a SAN or hybrid cluster
  - Up to 101PB across a NAS cluster
  - Up to 250 SVMs in a SAN cluster or 1,000 SVMs in a NAS cluster

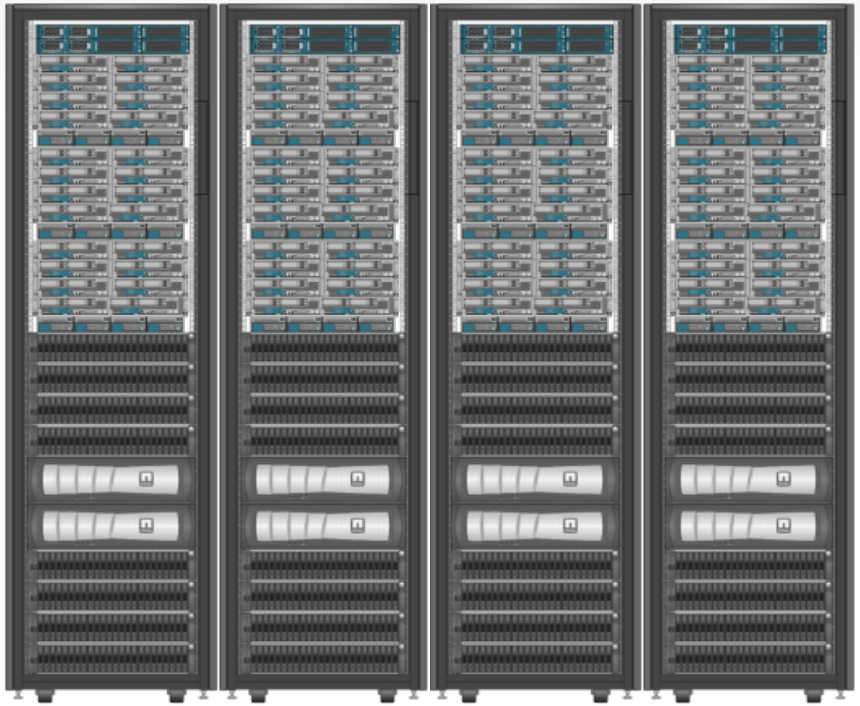

Scale Testing

Comparison with a competitor's published numbers

	Competitor *All-Flash*	Data ONTAP	Percent Decrease
Boot 200 Instances with 100GB Cinder Volumes	37 Minutes With 15 Compute No		89%
Delete 200 Instances and Cinder Volumes	***10 Minutes***		*88%*
Boot 1000 Instances with 100GB Cinder Volumes	150 minutes With 15 Compute Nodes		83%
Delete 1000 Instances and Cinder Volumes	***17 Minutes***	***6.4 Minutes***	*62%*

Clone and Boot 1000 Persistent Instances:

31 minutes!

- Source: NetApp RTP Testing

HA Reference Architecture

Available today!

- TR4323-DESIGN: “Highly Available OpenStack Deployments Built on NetApp Storage Systems”
  - Solution Design document based on Icehouse
  - Includes best practices for networking, storage, high availability
  - Available for download from http://www.netapp.com/openstack/
- TR4378-DEPLOY: “Red Hat Enterprise Linux OpenStack Platform 5 on NetApp Clustered Data ONTAP”
  - Available from http://www.netapp.com/us/media/tr-4378.pdf
- Follow us on Twitter for @openstacknetapp
  - Also NetApp’s OpenStack blog: http://netapp.github.io/openstack/

Questions?