




# Block Storage in the Open Source Cloud called OpenStack

June 3, 2015



# Webcast Presenters



**Alex McDonald, Vice Chair SNIA-ESF  
NetApp**



**Walter Boring,  
Software Engineer, HP**

- The material contained in this tutorial is copyrighted by the SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - ◆ Any slide or slides used must be reproduced in their entirety without modification
  - ◆ The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

# Agenda

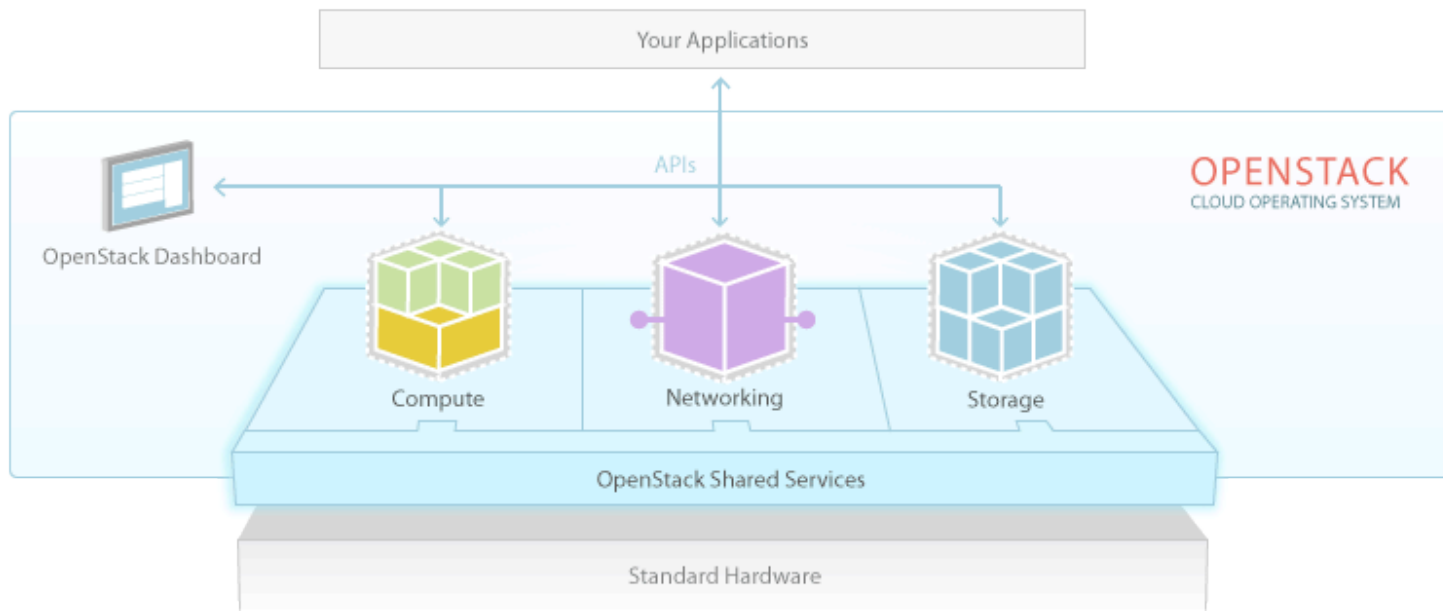
What are we going to cover in this presentation?

- What is OpenStack and it's history
- What are the key components to Block Storage
- How do volumes get attached to virtual machines
- What transport protocols are supported
- What are some of the future efforts

# What is OpenStack?

According to <http://www.openstack.org/software>

“OpenStack is a cloud operating system that controls large pools of compute, storage, and networking resources throughout a datacenter, all managed through a dashboard that gives administrators control while empowering their users to provision resources through a web interface.”



Code available under Apache 2.0 license. Design tenets – scale & elasticity, share nothing & distribute everything

# Open Development Process

Cloud computing fabric controller, the main part of an IaaS system



- **Time based release Cycles**
  - New software release every six months, with interim milestones
- **Twice Yearly Design Summits**
  - Immediately following software release to plan next version
  - Sessions led by developers and Project Technical Leads
- **Broad Contributions**
  - 1000+ developers, from over 50 companies worldwide
- **Elected Leadership**
  - Developers elect their own Project Technical Leaders

# OpenStack Goals

- **Open Platform**
  - Community driven
  - Technology accessible in many ways
    - Hourly via source, linux distributions, appliances, DIY
- **Empower users and developers**
  - Deep engagement from users and developers
  - Users have more control of their own destiny
- **Broad, global support from companies**
  - Not driven by a single Company; no single source
  - Aggregate investment is huge

# OpenStack History

Where did it come from?



- **Started in 2010**
  - Joint project between Rackspace Hosting and NASA.
- **First release October 2010**
  - Austin was the first release name and only included Nova and Swift
  - Every 6 months since then has been a release.
  - Sessions led by developers and Project Technical Leads
- **Cinder created in Folsom release**
  - Cinder came out of Nova Volume codebase
  - Released on September 27, 2012
- **Latest release is named Kilo**
  - Officially released on April 30, 2015



# OpenStack Capabilities

<b>Compute</b> (Nova)	Provision and manage large pools of on-demand computing resources
<b>Object Storage</b> (Swift)	Petabytes of reliable storage on standard gear
<b>Block Storage</b> (Cinder)	Volumes on commodity storage gear, and drivers for more advanced systems like HP, IBM, EMC, Red Hat/Gluster, Ceph/RBD, NetApp, SolidFire, and Nexenta
<b>Networking</b> (Neutron)	Software defined networking automation with pluggable backends
<b>Dashboard</b> (Horizon)	Self-service, role-based web interface for users and administrators
<b>Shared Services</b> (keystone, glance)	Multi-tenant authentication system that ties to existing roles (e.g. LDAP), Image service

# OpenStack Compute (Nova)

Cloud computing fabric controller, the main part of an IaaS system



## ➤ Relevant Components of Nova

- **REST API**
  - Exposes the provisioning capabilities of Nova
- **Scheduler**
  - Determines which compute node to use for requests
- **Compute Manager**
  - Handles the provisioning requests from the scheduler and sends requests to libvirt
- **Libvirt Volume**
  - Manages the discovery and removal of volumes

# OpenStack Block Storage (Cinder)

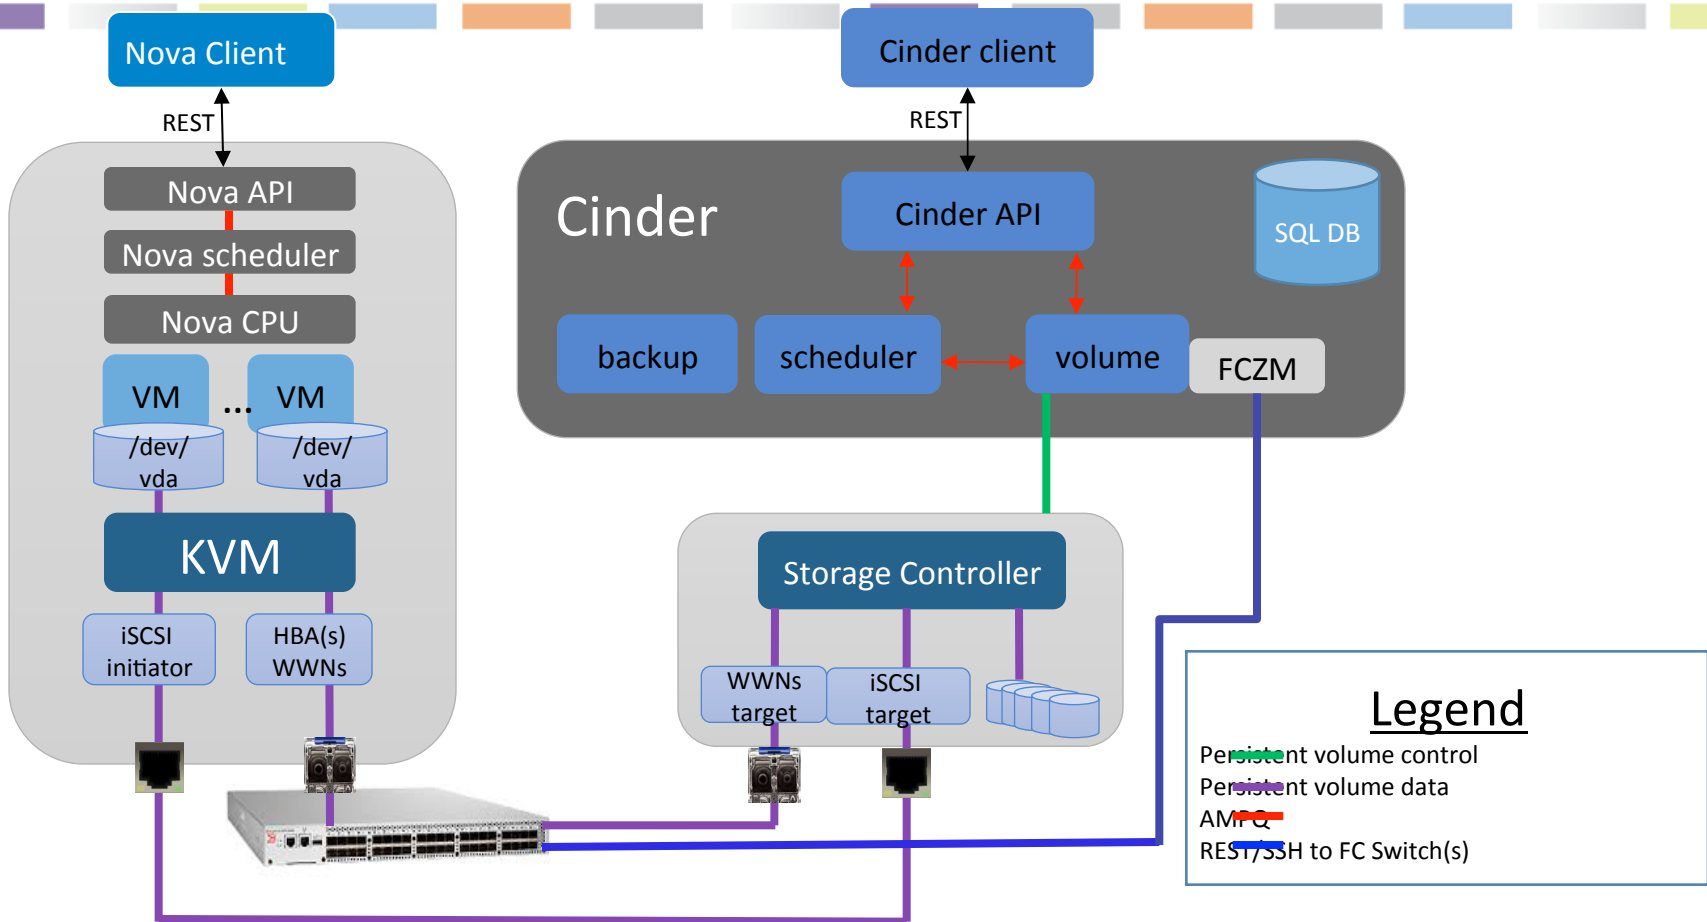
Provides software defined block storage via abstraction and automation on top of various storage systems.



## ➤ Relevant Components of Cinder

- **REST API**
  - Exposes the provisioning capabilities of Cinder
- **Scheduler**
  - Determines which Cinder storage system to send provision requests
- **Volume Manager**
  - Handles the provisioning requests from the scheduler and sends requests to storage system driver
- **Backup**
  - Provides volume backup and restore cinder volumes to external services (Swift, Glance)

# High Level Architecture



# OpenStack Block Storage

## The volume attachment process



- Nova manages the volume attachment process
  - Attachment is initiated by nova client
    - Nova volume-attach <nova id> <volume id>
  - Nova collects initiator information
    - Connector object is passed to Cinder which contains transport specific initiator information.
  - Nova calls Cinder to export volume
    - Cinder exports the volume from the specific cinder backend using the initiator connector information passed from Nova
    - Cinder replies with target information
  - Nova discovers host volume using target info
    - Nova's libvirt volume drivers are used to discover the volume in the host OS, depending on the transport protocol information passed back from Cinder

# OpenStack Block Storage

## Supported transports



### ➤ Supported attachment transports/protocols

- iSCSI
  - TCP (iscsiadm default)
  - Hardware transports
    - Be2iscsi, bnx2i, cxgb3i, cxgb4i, qla4xxx, ocs
  - iSER
- Fibre Channel
- Network
  - NFS, SMBFS, RBD (Ceph), GlusterFS
- GPFS (General Parallel file system)
- AoE (ATA over Ethernet)
- Vendor specific
  - Scality
  - Quobyte

# OpenStack Block Storage

Future work



## ➤ Future efforts

- Consolidate attachment code into python library, os-brick
  - Cinder has already migrated to os-brick
  - Nova migration in progress
- Add new additional transports to os-brick
  - Several vendors are working on new custom transports
- FCoE
- Break out Fibre Channel Zone Manager into standalone library
- Complete the addition of volume multi-attach

# OpenStack Block Storage

## Useful Links



- <http://www.openstack.org>
- <http://docs.openstack.org>
- <http://www.openstack.org/software/kilo>
- <http://bugs.Launchpad.net/cinder>
- <http://bugs.Launchpad.net/nova>
- <http://wiki.openstack.org/wiki/CinderSupportMatrix>
- <http://review.openstack.org>
- <http://status.openstack.org>
- <http://status.openstack.org/zuul>
- <http://en.wikipedia.org/wiki/OpenStack>
- [http://en.Wikipedia.org/wiki/OpenStack#Release\\_history](http://en.Wikipedia.org/wiki/OpenStack#Release_history)



# After This Webcast

- This webcast and a PDF of the slides will be posted to the SNIA Ethernet Storage Forum (ESF) website and available on-demand
  - ◆ <http://www.snia.org/forums/esf/knowledge/webcasts>
  
- A full Q&A from this webcast, including answers to questions we couldn't get to today, will be posted to the SNIA-ESF blog
  - ◆ <http://sniaesfblog.org/>
  
- Follow us on Twitter @ [SNIAESF](https://twitter.com/SNIAESF)



**Thank You**

