

NetApp Cloud: OpenStack



Rob Esker

Senior Product Manager



Jeff Whitaker

Cloud Solutions Manager

This article is the second in a series on the NetApp® cloud that explores how we help you to take full advantage of private, public, and hyperscale cloud resources.

In the last few years, [OpenStack](#) has emerged as the leading open infrastructure-as-a-service platform for private and public clouds. As NetApp CTO Jay Kidd said in his [2014 predictions](#), "This is the year that OpenStack distributions become more product than project," so if you haven't already done so, this the perfect time to start learning about OpenStack and thinking about how it might fit in your organization.

As a charter member of the OpenStack Foundation, NetApp has been involved for a number of years in adding to OpenStack storage functionality, including diverse deployment options and access to unique features for Cinder block storage and optimized Swift object storage. We're also driving new innovation with the Manila file share service.

In this article I explain in more detail what OpenStack is, provide a few details on how NetApp contributes to OpenStack, and share some thoughts on what your organization should do now.

What Is OpenStack?

One of the best ways to understand OpenStack is through a comparison to Linux®. It's often said that OpenStack is to the cloud what Linux is to servers. Where Linux provides an open—and extensible—operating environment for individual servers, OpenStack provides an open and extensible operating environment for cloud infrastructure.

The OpenStack project started in 2010 as a joint initiative between NASA and [Rackspace](#). Since then, more than 300 companies have joined and 188—including NetApp—actively contribute to the project.

One of the key things about OpenStack is that the various services that OpenStack provides—compute, storage, networking, and so on—are API-compatible with their equivalent [Amazon Web Services](#) (AWS) capabilities. This means that if you already have or you create an application that runs on AWS, you can run the application in any OpenStack environment, including an on-premises data center. The potential to be able to move applications between AWS and OpenStack is proving very attractive. NetApp has worked with a number of early OpenStack adopters to "repatriate" AWS applications that were more economical to run in house. We've also helped with

Explore

OpenStack Summit in Atlanta

The OpenStack Summit is a five-day conference for developers, users, and administrators of OpenStack cloud software. The [next OpenStack Summit](#) will take place May 12-16 in Atlanta, Georgia. NetApp will present multiple technical sessions including solution overviews, deep dives into Manila, detailed OpenStack deployments, and more. Come visit NetApp and enter to win an iPad® mini™.

[▶ More](#)

applications that default to on-premises but that periodically need to "burst" to use cloud resources.

OpenStack Design

OpenStack is modular in design, consisting of a collection of separate development projects all under the same umbrella. The resulting OpenStack services can be used to create large pools of processing, storage, and networking resources, all managed through a dashboard that gives administrators control while empowering users to provision resources through a web interface. Tenants of an OpenStack cloud can manually select and configure services. Applications running in an OpenStack cloud can select and configure services programmatically using OpenStack application programming interfaces (APIs).

Although the component modules are all designed to work together, you are also free to choose only the components you need. For example, a service provider might add one or two modules to complement capabilities it already offers.

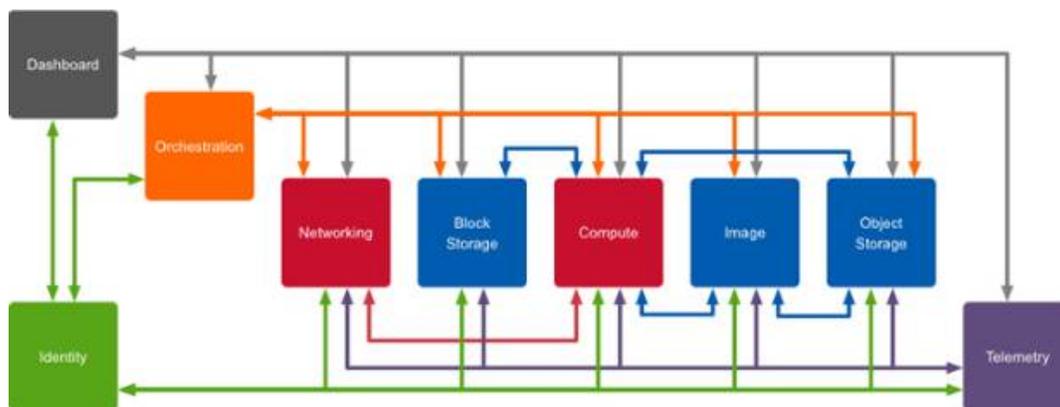


Figure 1) OpenStack logical architecture showing how the various modules fit together.

OpenStack [Havana](#), which was released in October 2013, consists of the following major modules:

- **OpenStack Compute (Nova):** Provides on-demand computing resources by provisioning and managing virtual machines using any hypervisor
- **OpenStack Image Service (Glance):** Provides a registry of bootable disk images and a service for storing and retrieving those images
- **OpenStack Block Storage (Cinder):** Provides a "block-storage-as-a-service" capability with persistent block devices mapped to OpenStack compute instances with support for a variety of storage solutions
- **OpenStack Object Storage (Swift):** Provides a fully distributed scale-out, API-accessible storage platform that can be integrated directly into applications or used for backup, archiving, and data retention
- **OpenStack Identity (Keystone):** Provides a central directory of users mapped to the OpenStack services they can access that acts as a common authentication system across the cloud operating system and can integrate with existing back-end directory services
- **OpenStack Network Service (Neutron):** Provides a pluggable, scalable, and API-driven system for managing networks and IP addresses; the pluggable back-end architecture lets users take advantage of basic commodity gear or advanced networking services from supported vendors
- **OpenStack Dashboard (Horizon):** Provides administrators and users with a graphical interface to access, provision, and automate cloud-based resources; the extensible design makes it easy to plug in and expose third-party products and services, such as billing, monitoring, and additional management tools
- **OpenStack Telemetry (Ceilometer):** Provides common infrastructure to collect usage and performance measurements within an OpenStack cloud; initial targets are monitoring, metering, and input to billing
- **OpenStack Orchestration (Heat):** Implements a service to orchestrate multiple composite cloud applications using the [Amazon Web Services \(AWS\) CloudFormation](#) template format through both an OpenStack-native and Cloud Formation-compatible API; it is intended, in part, to facilitate movement of

workloads from AWS to OpenStack deployments

OpenStack Releases

OpenStack has an aggressive release schedule, with a new release every six months. (Note that the releases are named alphabetically: Austin was the first, followed by Bexar, Cactus, and so on.) Each release typically includes new features and new modules, and may include projects in "incubation" for future releases.

The [OpenStack Icehouse release](#) came out on April 17, 2014. It includes significant enhancements to all of the modules described previously, as well as one new project module:

- **OpenStack Database as a Service (Trove):** Provides the ability to deploy database services in an OpenStack environment

Projects in incubation in Icehouse include:

- **OpenStack Bare Metal (Ironic)**
- **OpenStack Messaging (Marconi)**
- **Data Processing (Sahara)**

Every new release is followed by a design summit and conference where the features for the next release are decided on. (See the sidebar accompanying this article for information on the upcoming summit.)

OpenStack Distributions

OpenStack is defined under an Apache 2.0 license that makes it more commercially friendly than other large open-source projects, enabling broad early involvement and adoption by a wide range of companies.

The analogy between OpenStack and Linux extends to include a very similar distribution model. Just as companies like Red Hat and SUSE created packaged Linux distributions based on the public Linux source code, a dozen or more OpenStack providers, including Red Hat, SUSE, Mirantis, Rackspace, and Metacloud, are creating packaged versions of OpenStack. The OpenStack foundation has done a good job of defining the core features of OpenStack such that all distributions deliver the key capabilities, avoiding fragmentation.

NetApp actively partners with many of the OpenStack distributors. Check out [TR-4284: Reference Architecture: Deploying Red Hat Enterprise Linux OpenStack Platform 4 on NetApp Clustered Data ONTAP](#) for one recent example.

What Does NetApp Contribute to OpenStack?

NetApp is a charter member of the OpenStack Foundation and, since early 2011, a contributing member of OpenStack. In addition to providing integrations with NetApp products, we have a leadership role in the community with the addition of a major new service and have sponsored the last seven OpenStack Summits.

NetApp has made significant contributions in storage-related functionality to recent releases of OpenStack, including Essex, Folsom, Grizzly, Havana, and, now, Icehouse. One of our goals is to provide integrations within OpenStack that allow you to utilize all the capabilities of NetApp FAS and E-Series storage and to make these integrations part of the upstream source so they are available in every packaged distribution.

There are many reasons why enterprises choose NetApp FAS storage, and those reasons apply equally to the cloud. These reasons include advanced storage efficiency with data deduplication, compression, cloning and thin provisioning, Integrated Data Protection with a complete suite of products developed over many years, continuous availability, performance assurance with QoS, and the scalability inherent in the clustered Data ONTAP®

operating system. NetApp addresses critical cloud needs with feature sets hardened against enterprise requirements. As you make the transition to an "as-a-service" IT model, we want to enable you to continue to take advantage of the capabilities that deliver the best value from your storage.

NetApp OpenStack integrations include:

- NetApp specific drivers for Data ONTAP, E-Series, and EF-Series systems
- Storage Service Catalog (SSC) capabilities
- Enhanced and persistent instance creation with copy offload
- pNFS support
- Optimized deployment of OpenStack Object Storage (Swift)

In addition, NetApp has been actively engaged in creating a shared file system service code-named Manila.

Drivers to Enable NetApp Storage in OpenStack Environments

NetApp has contributed a series of drivers that enable NetApp storage and data management capabilities to be provisioned through OpenStack Block Storage (Cinder). In addition to iSCSI options, NetApp enables the use of NFS. Using a file-oriented protocol like NFS as a back end for the block-oriented Cinder service provides a means to achieve greater scale. These drivers provide NFS mounts to individual hypervisors and files are in turn provided as virtual block devices to the hypervisor.

The OpenStack drivers for NetApp clustered Data ONTAP are available in both a "direct-to-storage-virtual-machine" mode and optionally one capable of providing additional service deployment automation and adaptability. Drivers are also available for NetApp Data ONTAP 7-Mode.

An exciting new addition to the Icehouse release introduces NetApp E-Series and EF-Series Cinder drivers to provide additional options for requirements demanding highly dense, high-performance characteristics.

Storage Service Catalog Capabilities

The NetApp Cinder drivers allow you to construct a catalog of differing storage capabilities to meet a diverse base of application and tenant needs. The Cinder Storage Service Catalog can be defined with various efficiency, performance, availability, and protection attributes. The catalog entries themselves (referred to as Cinder Volume Types) can be defined in very granular form or more commonly represent a collection of capabilities most appropriate for a particular tenant use case, as shown in Figure 2.

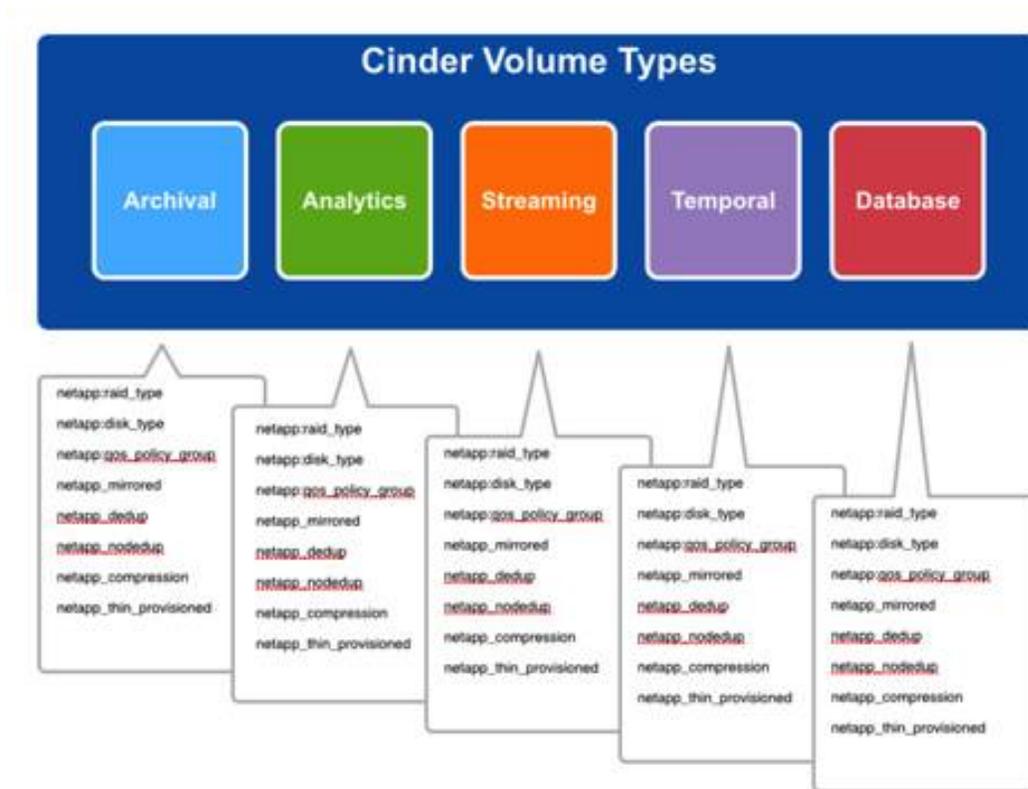


Figure 2) NetApp provides storage service catalog capabilities for OpenStack Cinder.

Enhanced and Persistent Instance Creation and Copy Offload

The NetApp Cinder drivers for clustered Data ONTAP make use of NetApp cloning technology to quickly and efficiently create multiple instances (guest virtual machines) from Glance images. The effect is very similar to that achieved with the rapid cloning capabilities NetApp has made available in the various incarnations of the [Virtual Storage Console](#).

Multiple layers of optimization are employed. The first volume created from an image requires a full copy from Glance by default; however, with Glance configured to store images on a NetApp NFS share, even the first copy can be an efficient clone operation or hardware-assisted copy.

This feature allows the use of NetApp Copy Offload when creating new Cinder volumes from Glance images or existing Cinder volumes. A NetApp specific NFS client binary is used to communicate between the Cinder driver and NetApp storage.

pNFS Support

[Parallel NFS](#) (pNFS) is an extension of NFS 4.1 that provides improved performance, especially with clustered Data ONTAP. No special configuration in Nova or Cinder is required to enable pNFS; mount options can be set in the cinder.conf file.

OpenStack Object Storage (Swift) Optimized Deployment

The [NetApp OpenStack Deployment and Operations Guide](#) contains a reference architecture for the deployment of OpenStack Object Storage atop NetApp E-Series systems that lowers the cost of operations and improves scalability. Swift, by default, makes three copies within a single site for any single item stored, whereas E-Series Dynamic Disk Pools can reduce that to just under a 1.3 multiple. E-Series represents a preferred back end for Swift to improve scalability by reducing the number of replicas (data copies) required for each object in the store.

Manila

Although much of the total storage shipped worldwide underpins shared file systems, up to this point OpenStack has not included a shared file system service. This is something that struck NetApp as a pretty big omission. NetApp designed, prototyped, and built a new service—code-named Manila—that is the equivalent of Cinder for shared or distributed file systems. Like Cinder, Manila is a provisioning control plane for orchestrating, providing capacity, and so on.

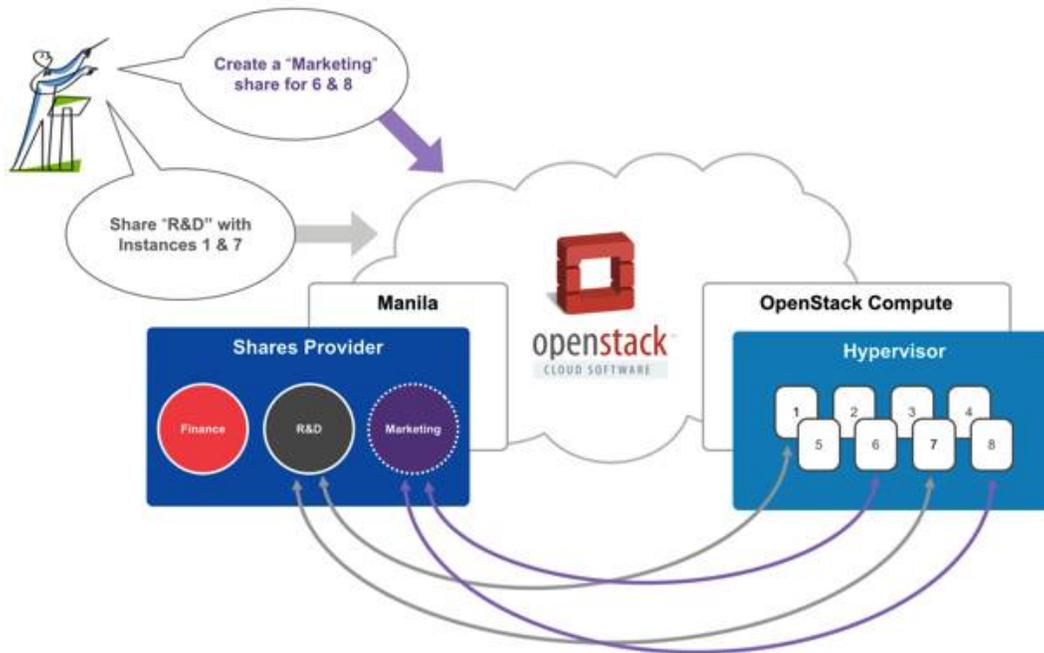


Figure 3) NetApp is developing Manila as a shared file system service to be incubated in an upcoming OpenStack release.

So, Manila is not the actual shared file system itself; it's the control plane that can, for instance, provide access to an existing CIFS share or create a new NFS export and map it between VM instances X, Y, and Z. NetApp is working within the OpenStack community to get Manila into incubation for the Juno release cycle due in October 2014.

What Should You Do Now?

Companies that are gravitating toward OpenStack are doing so for a variety of reasons:

- OpenStack isn't proprietary so it's harder to get painted into a corner.
- OpenStack is hypervisor agnostic: It works with a wide range of hypervisors.
- OpenStack is well supported by many independent software vendors and hardware providers.
- Traditional applications can be supported as well as those built in a more "cloud-native" form.
- Existing AWS applications can be brought on the premises.
- The rich and growing set of development tools for AWS is attractive.
- OpenStack is a good fit for companies moving or desiring to move to a more cloud-friendly (as-a-service) infrastructure

If some or all of these reasons apply to your company, you may want to consider an OpenStack deployment if you aren't already doing so. NetApp is working with more than three-dozen customers to power production OpenStack deployments and a similar number of proof-of-concept implementations.

At a minimum, you should monitor what's happening and track the new developments with each OpenStack release. Keep an eye on OpenStack.org and netapp.com/openstack for the latest information. You may also want to move to a proof of concept so that you can begin to experiment with the technology and build knowledge and expertise.

NetApp is committed to continued OpenStack development and support, and we partner with best-in-class integrators and OpenStack distributors. At the same time, we are committed to provide you with a wide variety of cloud options in addition to OpenStack so that you can choose the best options—software and hardware—for your needs. Find out more about all available NetApp cloud options [here](#).



By Rob Esker, Senior Product Manager and Jeff Whitaker, Cloud Solutions Manager

Rob's involvement with OpenStack dates to the origins of the community in 2010. He founded and continues to lead NetApp's efforts in this area. He's been engaged as an architect of deployments and an ATC (technical contributor to the community), serves as NetApp's OpenStack Foundation representative, and is a regular speaker at OpenStack Summits, conferences, and user groups. Rob is responsible for strategy, integration, and development of NetApp portfolio technologies in and around OpenStack as well as business development, organizational, and process alignment for going to market. A former U.S. Marine, he resides in Austin, Texas, with his wife and two daughters.

As a cloud solutions manager, Jeff focuses on OpenStack and cloud software management. For over five years he has been with NetApp, where he has managed the NetApp virtual storage appliance product family, including Data ONTAP Edge, the first virtual machine version of Data ONTAP. In the past couple of years Jeff has been heavily involved in cloud product direction and strategy, working directly with customers on their public and private cloud needs. He has been in the industry for nearly 20 years and has been involved in both marketing and selling storage and networking products for the service provider space to high-performance computing markets.

Quick Links

- › [Tech OnTap Community](#)
 - › [Archive](#)
 - › [User Groups](#)
 - › [PDF](#)
-