# OpenStack is Now Ready for the Enterprise

## Technology Enhancements, Robust Community, & Strong Ecosystem Make Leading Open Source Cloud Operating System Ready for Primetime

## Executive Summary

The OpenStack cloud operating system is becoming a strong alternative to traditional IT service delivery models for large enterprises. In addition to replacing traditional models, OpenStack is establishing itself as an accelerator towards efficient and credible on-premise and hybrid cloud alternatives to remote-only cloud services (Amazon Web Services, Google Cloud Engine, and Microsoft Azure).

Early adopters of OpenStack were focused on research and software development environments, and early releases lacked the stability and feature sets required for broad scale enterprise use. However, OpenStack has come a long way in the last six years with significant technology enhancements, a strong community, and a range of vendor offerings to help fill in the perceived gaps of the upstream code.

Moor Insights & Strategy (MI&S) believes that OpenStack is reaching a tipping point toward mainstream enterprise IT adoption. However, transitioning from a traditional IT deployment model to a cloud operating model is not for the faint of heart, and running a cloud at scale can prove to be a challenge. Enterprise IT organizations looking to adopt OpenStack should take a calculated approach to adoption and consider partnering with OpenStack technology vendors and system integrators to ease the path to migration.

## OpenStack History & Key Tenets

OpenStack is a cloud operating system designed for datacenter-wide control of resources as pools of compute, storage, and networking. It provides management and administration through well-defined, simple web interfaces or with command-line tools. OpenStack began in mid-2010 and was formed by an agreement to converge the efforts of Rackspace (rewriting the infrastructure controlling their servers, the Rackspace Cloud File Platform) and NASA's project Nebula (work from Anso Labs, a NASA contractor). The OpenStack mission statement defines what the project is about and helps shape the goals the community aspires to reach:

> *"To produce a ubiquitous Open Source Cloud Computing platform that is easy to use, simple to implement, interoperable between deployments, works well at all*


Copyright ©2016 Moor Insights & Strategy

*scales, and meets the needs of users and operators of both public and private clouds."*

OpenStack was founded with four guiding principles that have been the cornerstone of the entire project.

1. **Open Source**: A commitment to truly open source software with no editions which cripple or limit feature content.
2. **Open Design Process**: At regular intervals (6 months), the community holds a public design summit to gather requirements and write specifications.
3. **Open Development**: Code is publicly available throughout the development process.
4. **Open Community**: A commitment to maintain a healthy community, necessary to the ongoing development of OpenStack.

Today OpenStack is managed by the OpenStack Foundation, a nonprofit non-stock organization founded in 2012 whose purpose is to develop, support, protect, and promote the OpenStack Project bylaws. OpenStack is published under the Apache 2.0 License.

In addition to being open, the key motivation behind OpenStack is to change the traditional software development and acquisition model in three ways.

1. Traditional comparable software packages are "owned" by the developing companies who naturally derived special value from the authorship even though they may offer some version for sale to the public.
2. Feature content of traditional software is controlled by the developing company who seeks their own best interest.
3. The nature of free / open source eliminates prohibitive cost barriers.

## OpenStack Framework Overview

OpenStack has evolved considerably since the project's inception. The initial Austin release in 2010 included just two core services (nova and swift), while the Mitaka Release in first half of 2016 included six core and 16 optional services. Everything in OpenStack is controlled by the administrator through a well-defined dashboard, with the users empowered to provision their own resources. Figure 1 depicts the OpenStack framework at a high level.


Copyright ©2016 Moor Insights & Strategy

## FIGURE 1: OPENSTACK FRAMEWORK

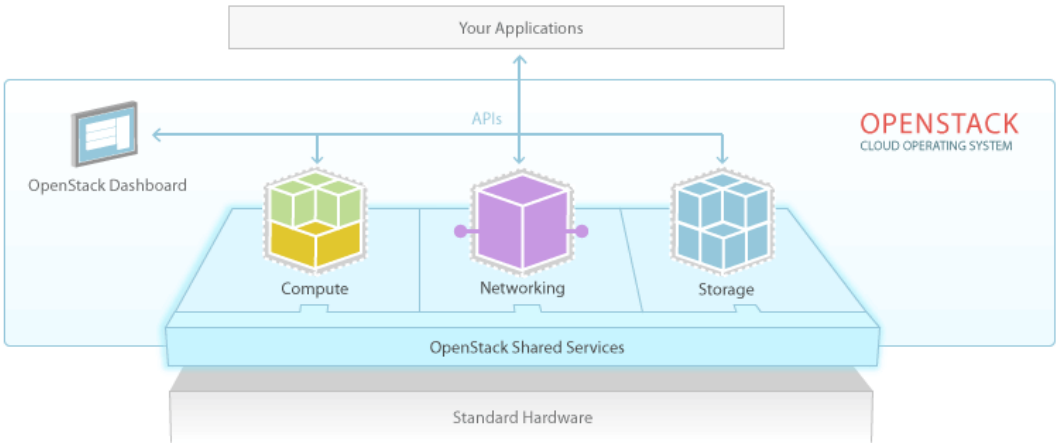


*Source: http://www.openstack.org/software*

OpenStack is made up of services that are the definition and operation of its infrastructure. There are five core services at the heart of an OpenStack deployment, each with a set of well-defined APIs.

- **nova**: Manages the life cycle of **compute** instances including spawning, scheduling, and decommissioning of machines on demand.
- **neutron**: Provides **networking** as a service for other OpenStack services such as nova which is the predominate consumer of neutron's service.
- **cinder**: Provides persistent **block storage** to all running instances including their creation and management.
- **keystone**: Provides an **authentication and authorization** service for all OpenStack services including a catalog of all endpoints.
- **glance**: The **image service** that stores and retrieves virtual machine disk images used by nova during instance provisioning.

One additional core service is available if an object storage service is needed.

- **swift**: An **object storage** service which stores and retrieves arbitrary unstructured data objects in a highly fault tolerant manner (important in a scale out architecture).

Each of the OpenStack core services has been under active development for more than four years. Appendixes A and B provide the release history and a list of optional services that fulfill the functional needs of a cloud architecture. They are optional in that some have chosen different methods to provide these necessary functions.


Copyright ©2016 Moor Insights & Strategy

## OpenStack Target Markets & Use Cases

OpenStack interest and adoption by end users have increased significantly over the last six years. The OpenStack User Survey April 2016 shows the top priorities contributing to adoption are (1st and 3rd) "saving money" and "operational efficiency" and (2nd and 4th) "being more agile / innovative" and "avoiding vendor lock-in". Key target segments that are currently the primary adopters of OpenStack include cloud service providers, telecommunications companies, enterprise IT, and some recent traction within the high-performance computing (HPC) community. This paper focuses on enterprise IT user requirements and adoption trends.

OpenStack allows enterprise IT organizations to act as a cloud service provider for internal business units and project teams. Adoption within the enterprise spans across many vertical markets including financial services, manufacturing, retail, healthcare, and others who are building private or hybrid cloud environments. Figure 2 shows a list of specific workload areas that OpenStack users are running today.

## Figure 2: Workloads & Frameworks Running on OpenStack

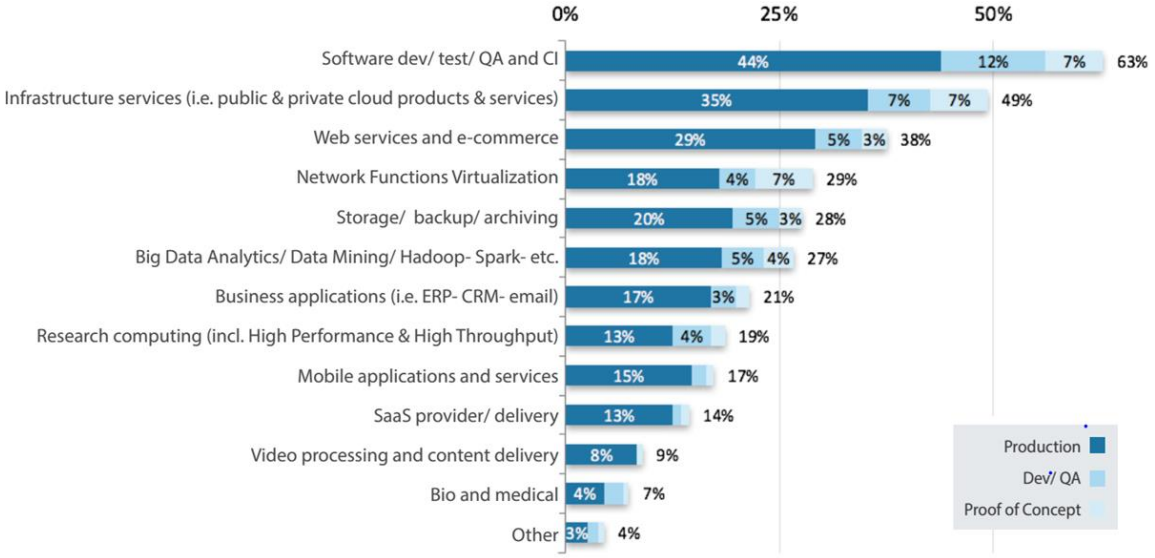


*Source: OpenStack User Survey April 2016 (License Agreement)*

Software development / test / QA has been the predominant early adopter use case, because the OpenStack framework can help developers leverage a continuous improvement / continuous development (CI/CD) model to which many organizations are migrating today. OpenStack has emerged as the *de facto* standard for infrastructure as a service (IaaS) in the private cloud and in some cases a more flexible, cost-effective


Copyright ©2016 Moor Insights & Strategy

alternative to public cloud resources for software development. In addition, OpenStack gives engineers a vital self-service capability to provision (and de-provision) environments autonomously. This capability eliminates the downstream bottleneck caused when waiting for operations staff to do the provisioning.

OpenStack was designed with the original intent of being optimized for "born-in-the-cloud" workloads where application logic is separated from the infrastructure and resiliency is built into the application. As such, workloads with these criteria—like web services, infrastructure services, and ecommerce workloads—are also experiencing success in OpenStack environments. However, the enterprise IT organization's reality includes a long tail of legacy applications like traditional databases, ERP, CRM, and others. To enable broader adoption of OpenStack within the enterprise, enterprise IT will need the capability to include these legacy applications in their OpenStack environments over the long term.

## KEY OPENSTACK REQUIREMENTS FOR ENTERPRISE IT

Enterprise IT interest in OpenStack continues to increase at a rapid rate, as many organizations are deploying or have plans to deploy OpenStack over the next several years. However, by some estimates, a large percentage of enterprise users have failed when attempting to implement OpenStack, and most users have found the implementation experience difficult. Below are some of the top challenges cited by enterprise users who are evaluating and deploying OpenStack.

- **Enterprise-Class Product Maturity for All Workloads**: Many enterprises looking to deploy OpenStack into production see several potential pitfalls that may limit their ability to scale across all of the workloads in their datacenter.
    - **Scalability**: OpenStack has not been broadly proven at large scale by many users (100s to 1,000s+ node clusters). Many users see potential challenges to scale on OpenStack within their environments.
    - **Release Process Maturity**: When compared to the maturity of many commercial software packages, OpenStack is viewed as having shortcomings in terms of bug fixes, test coverage, documentation, and the ability to upgrade to perform rolling upgrades to new releases.
    - **Support for Legacy Applications & Deployment Processes**: OpenStack was originally designed to support born-in-the-cloud workloads where availability is built into the application versus traditional applications that require high availability built in the underlying infrastructure. Lack of


Copyright ©2016 Moor Insights & Strategy

robust features for legacy application support may be viewed as a limitation compared to traditional virtualized environments. In addition, OpenStack is designed to be deployed using a CI/CD deployment model which is a new approach for many enterprise IT organizations.

- **Complexity of Migration, Deployment, & Operations**: Many IT organizations find the task of migrating their applications to deploy and run on OpenStack very challenging. When compared to a packaged hypervisor or other commercial software, enterprise users find OpenStack to be a low-level framework lacking in robust deployment tools. There is also a perception that OpenStack installation and ongoing management processes are designed for consultants and specialists who speak OpenStack 24x7. These challenges are particularly true for users who leverage OpenStack upstream projects and deployment tools versus vendor distributions / tools.

In addition to the above product- and process-related challenges, IT organizations often find a lack of OpenStack skills in the market, adding reluctance to place big bets on OpenStack for broad scale management of private and hybrid cloud environments. To help alleviate these challenges, many enterprise users look to vendor products and services for the capabilities and expertise to adopt OpenStack for their environments.

## THE JOURNEY TO OPENSTACK MATURITY

As enterprise IT seeks to overcome potential challenges in deploying OpenStack, many users want to evaluate the maturity of OpenStack as a whole and of each project of interest. In a traditional enterprise environment, maturity is often interchangeable with the idea of stability and can be measured by the arrival and rate of change requests (usually identifying bugs that need repair). The same metric can be applied to OpenStack. However, OpenStack is as an ever-evolving development program by design and will be unlikely to reach what some in enterprise IT would view as “finished”.

### *MEASURING MATURITY FOR OPENSTACK*

The OpenStack project has done an excellent job publishing information necessary to understand a measure of maturity of each project included in a release. Maturity as defined by OpenStack is a metric (1 through 8) based on a set of distinct stability criteria including the existence of an install guide, support by 7 or more SDKs, an adoption percentage greater than 75%, project achievement of corporate diversity, and whether or not the project has stable branches. Figure 3 is a snapshot of the maturity metric available for the OpenStack core services.


Copyright ©2016 Moor Insights & Strategy

## FIGURE 3: MATURITY DASHBOARD FOR OPENSTACK CORE SERVICES

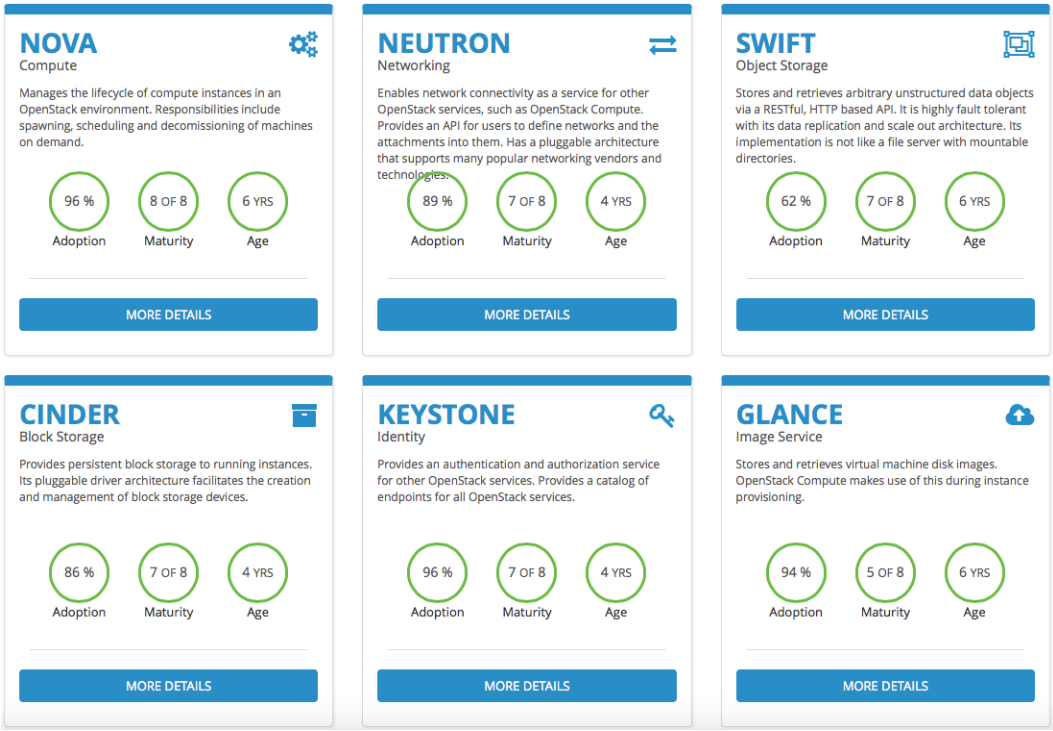


*Source: http://www.openstack.org/software/project-navigator March 2016*

Each of the core OpenStack projects includes very detailed deployment and installation guides with specific guides available for the major Linux distributions.

### *BUG TRACKING*

Bug tracking capability is another important element when looking at stability for enterprise users. Actual bug information for OpenStack projects is tracked on GitHub. Figure 4 shows an example of the nova bug tracker at https://bugs.launchpad.net/nova.

## FIGURE 4: OPENSTACK NOVA BUG TRACKER EXAMPLE

*Source: GitHub 16 April 2016*


Copyright ©2016 Moor Insights & Strategy

While there are no critical bugs noted, the total bug count is far from 0, and this example shows 52 high importance bugs. Many would panic at the thought of high importance bugs reported in the most mature of all the core services, but even in the best enterprise software packages, high importance bugs are discovered from time to time, some of which are shared with users and others managed internally by the vendor. The GitHub community has an additional metric called "Bug Heat" that is an attempt to measure of impact of each outstanding OpenStack bug.

## OpenStack Community & Vendor Support for the Enterprise

Over the last 6 years, the OpenStack community has flourished into a strong group of individuals and companies contributing to the various projects and source code improvement initiatives. For example, the OpenStack Liberty release from October 2015 included contributions from 1,033 individual contributors and 164 organizations.

The proper governance model for any open source project is critical to ensure the project guidelines are understood and followed by community members. OpenStack leverages a community-oriented governance approach. While this approach can be beneficial to create products that align with industry needs, it requires significant support from key industry leaders aligned with the OpenStack mission to help create structure around decision making and governance.

### *Building a Stable Base*

The OpenStack community is heavily focused on developing a stable, robust trunk to help drive maturity, ease of implementation, and standardized capabilities for downstream projects and distributions. A number of groups within the OpenStack community are responsible for driving product requirements for trunk development and enhancements. Figure 5 illustrates the organizational structure for the responsibilities of the user committee, product working group, and other working groups in driving requirements for the multi-release roadmap. In short, the product working group consolidates and integrates feedback from the various working groups and user community to prioritize OpenStack product requirements. In addition, the product working group is responsible for publishing a well-defined three-release community roadmap. The companies and individuals required to deliver on each of the roadmap features commit to provide the development resources needed to execute to the roadmap release schedules.


Copyright ©2016 Moor Insights & Strategy

## FIGURE 5: ORGANIZATIONAL STRUCTURE FOR PRODUCT ROADMAP PROCESS

*Source: OpenStack Foundation*

Enterprise readiness and adoption have been a major focus for the OpenStack community, so the Enterprise Work Group is tasked with ensuring enterprise user requirements are prioritized within the various OpenStack projects and releases. Below are examples of priority requirements the Enterprise Work Group is driving.

- Ease of installation and deployment
- Support for rolling upgrades
- Support for High Availability (HA) OpenStack Services
- Support for HA VMs (including live migration)
- Onboarding "pet" workloads (the OpenStack community refers to legacy applications that require HA infrastructure as "pets", while born-in-the-cloud workloads are considered "cattle"), including…
    - Developing reference architectures to assist developers who are re-architecting applications for the cloud
    - Guidance for enterprise developers to launch and create applications that are OpenStack ready


Copyright ©2016 Moor Insights & Strategy

In addition to the Enterprise Work Group, the Large Scale Deployment Work Group, Application Ecosystem Work Group, and others focus on solving problems relevant to enterprise users like scalability and support for a range of application development environments. Collectively these groups drive requirements for releases, addition of new projects, and support for relevant technologies for the OpenStack user base. For example, the Liberty release included features relevant to enterprise users such as finer-grained management controls, performance enhancements for large deployments, and more powerful tools for managing new technologies such as containers in production environments. The Mitaka release (April 2016) includes enhancements to user experience, scalability, and manageability with a long list of features that include simplified rolling upgrades, easier live migrations, improved APIs, better interoperability, and continued improvements in container integration.

## *Making the Products & Community Better*

The OpenStack Community has been focused on driving improvements to not only the features and capabilities of the solution but also the stability of each project and release. For example, corporate industry leaders sponsored a Global OpenStack Bug Smash event for the Mitaka release. The event was an around-the-clock, around-the-world collaboration of the developer community to fix as many bugs as possible across the wide range of OpenStack projects. The event was also designed to help onboard and grow the number of OpenStack developers, and increase the community's collective knowledge of OpenStack tools and processes.

The OpenStack Foundation and its members are also working to solve the OpenStack skills shortage issue. For example, the community launched a professional certification program and a tool called Project Navigator designed to teach users how to consume OpenStack software. The certification program aims to provide a baseline assessment of OpenStack knowledge and be accessible to OpenStack professionals around the world. The first certification, Certified OpenStack Administrator (COA), will be available in 2016. Project Navigator aims to help users understand the breadth of functionality of the more than 25 cloud-related services or projects available under the OpenStack umbrella, by aggregating important information about each project—such as maturity, release schedule, packaging, and documentation support—into an easy-to-navigate interface. Openstack.org also hosts a training marketplace with information on training classes put on by corporate community members.


Copyright ©2016 Moor Insights & Strategy

## IT TAKES A VILLAGE TO BUILD AN ENTERPRISE-CLASS CLOUD

As is true with many open source communities, OpenStack relies on corporate industry leaders to make investments and provide resources to help grow the community and market acceptance of the OpenStack operating system. OpenStack is heavily dependent on corporate developers for code contributions to new releases, bug fixes, and overall infrastructure health. For example, the Mitaka release is estimated to have 70-80% of its developer resources driven by representatives from 10-12 corporations. Vendors like Intel, Mirantis, Rackspace, RedHat, and others are also focused on the "nuts and bolts" work of code improvements and bug fixes in the underlying infrastructure to enhance stability, manageability, and interoperability. In addition, these organizations are actively engaged with training, documentation, and management of the OpenStack user groups.

Industry leaders are also focused on broadening the OpenStack developer and user base. For example, Intel and Rackspace announced a joint investment to accelerate the enterprise features and adoption of OpenStack. Included in this investment is the development of the OpenStack Innovation Center (OSIC) which includes two 1,000-node OpenStack clusters that are available to the community for upstream development and large-scale testing of OpenStack features and implementations. Users are required to release results of their experiments that were run on the cluster back to the community and also required to upstream applicable code that comes out of their projects.

## THE EXPLOSION OF COMMERCIAL SUPPORT FOR OPENSTACK

Some of the strongest evidence of the evolution and readiness of OpenStack is the number and nature of companies that now offer OpenStack. Many fill the perceived gaps with valuable additional services including other OpenStack tools not in the current release and may include tools of their own making. The OpenStack Marketplace: Distros & Appliances includes offerings ranging from pure distributions of OpenStack, integration with other OpenStack projects and specific vendors tools, to integrated appliances. The vendor community continually provides information to ensure end users are aware of available OpenStack products and services. For example, the Intel Cloud Builders program serves as a go-to-market avenue for industry players to publish reference architectures and solution blueprints.

Pay particular attention to the "TESTED" symbol associated with some items on the OpenStack Marketplace (Figure 6) and how it relates to the offering. The tested symbol


Copyright ©2016 Moor Insights & Strategy

indicates that the listed item contains OpenStack and has been tested / validated for OpenStack API compatibility including date of the release.

## FIGURE 6: OPENSTACK TESTED

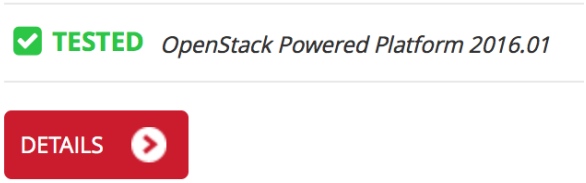


*Source: http://www.openstack.org/marketplace/distros*

The presence of the OpenStack Tested symbol denotes the product is validated as OpenStack Powered Platform, OpenStack Powered Compute, or OpenStack Powered Object Storage.

"OpenStack Powered" is for products that run fully functional instances of the OpenStack software. Figure 7 provides an example of the OpenStack Powered logo. Full detail and meaning are available at http://www.openstack.org/brand/interop.

## FIGURE 7: OPENSTACK POWERED LOGO

*Source: https://www.openstack.org/marketplace*

A major focus area for commercial software services is making OpenStack easier to install and deploy. Figure 8 provides an overview of the open source and commercial deployment tools available today.


Copyright ©2016 Moor Insights & Strategy

## FIGURE 8: DEPLOYMENT / CONFIGURATION TOOLS USED WITH OPENSTACK

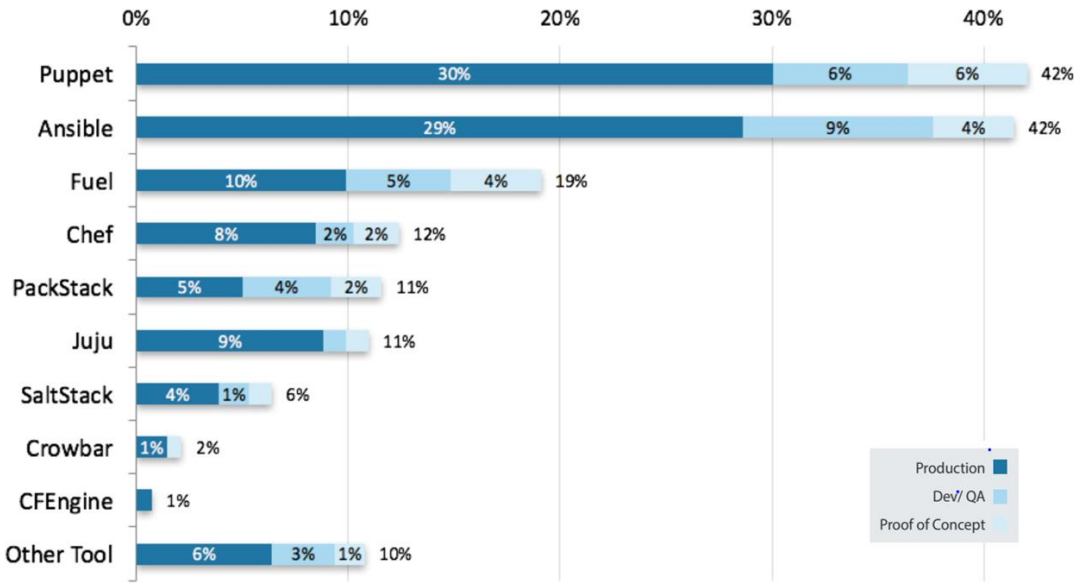


*Source: [OpenStack User Survey April 2016](#) ([License Agreement](#))*

In addition to solving deployment challenges with OpenStack, vendors are creating OpenStack distributions and specialty packages focused on improving the ongoing operational experience. These include support for performance enhancements, rolling upgrades, high availability features, and hardware validation services. Other vendor-focused work includes integration efforts around complementary technologies like containers and traditional virtualization environments. Integrators and service providers are also delivering OpenStack consulting and support expertise to help enterprise users deploy and manage their OpenStack environments. When evaluating vendors, enterprise IT organizations should ensure they choose vendors that complement the standard open source OpenStack code and contribute their enhancements upstream back to the community.

## EARLY ADOPTER EXPERIENCES

One of the best indicators of the adoption and readiness of OpenStack is whether or customers are actively using it. The initial OpenStack users were focused on primarily on research and non-enterprise use cases. However, IT organizations in government, media, retail, finance, consumer goods, manufacturing, and technology are now among


Copyright ©2016 Moor Insights & Strategy

those moving beyond evaluations and pilots to adopt OpenStack-based clouds as a strategic initiative to support the business.

One such company is Symantec, a leading provider of cyber security solutions. Symantec uses OpenStack for both their business IT operations and as the cloud platform which runs all of their applications. Symantec is an “upstream” consumer of OpenStack, choosing to take their selected packages directly from GitHub. Symantec has experienced significant improvements in the operational aspects of OpenStack over the last several release cycles, including the ability to apply patches and upgrades. Based on their experience, Symantec encourages enterprise users looking to use the upstream code deployment approach to remain current on the latest release to ensure they receive the benefit of bug fixes. This requires upgrading to the latest release every six to twelve months. Users looking at less frequent upgrades, more robust bug fixing capabilities, and additional tools or packages to ease in deployment may want to consider using an OpenStack distribution.

## Call to Action

OpenStack has evolved over the last six years into a viable cloud operating system solution for enterprise IT. While early releases lacked the robust features, usability, and deployment model requirements for enterprise use, the upstream capabilities have been enhanced significantly, and ecosystem partners have filled in the remaining gaps with commercial distributions and services.

Enterprise IT organizations should consider a phased rollout approach when introducing OpenStack in their environments.

- **Identify a pilot project** to build familiarity with OpenStack and to test out the viability of partner-enhanced approaches. Ideal use cases include test / development environments or web / cloud based applications. Avoid using legacy applications as the first foray into an OpenStack deployment model.
- **Ensure internal IT staff has the proper OpenStack training and knowledge base** to take on a pilot project. Even using a partner model requires internal teams to have a solid understanding of the OpenStack approach and how it differs from traditional virtualization environments.
- **Determine the best deployment strategy**. OpenStack can be used in many types of cloud environments including dedicated hosting, public, private, and hybrid clouds. Understand your desired adoption model when you begin the


Copyright ©2016 Moor Insights & Strategy

project, as this will have an impact on your rollout approach and partnership strategy from day one.

- **Consider partnering with OpenStack technology vendors and system integrators**. While there are some notable customers who have taken the “roll your own” approach using “upstream” code, MI&S believes many enterprise IT organizations will have a higher likelihood of success by leveraging partners for distributions, software tools, services, and support. The proper partner strategy can translate to faster deployments, access to a richer set of features, avoidance of common pitfalls, and easier staffing. Figure out the right level of vendor support model, given your team and your organization’s strategy.

MI&S believes OpenStack is ready for enterprises whose IT organizations who adopt a pragmatic approach to workload choices, staffing plans, deployment models, and partnership strategies.


Copyright ©2016 Moor Insights & Strategy

## APPENDIX A: RELEASE HISTORY FOR OPENSTACK SERVICES

| | Austin 2010.1 | Bexar 2011.1 | Cactus 2011.2 | Diablo 2011.3 | Essex 2012.1 | Folsom 2012.2 | Grizzly 2013.1 | Havana 2013.2 | Icehouse 2014.1 | Juno 2014.2 | Kilo 2015.1 | Liberty 2015.2 | Mitaka 2016.1 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| **nova** | **x** | **x** | **x** | **x** | **x** | **x** | **x** | **x** | **x** | **x** | **x** | **x** | **x** |
| **swift** | **x** | **x** | **x** | **x** | **x** | **x** | **x** | **x** | **x** | **x** | **x** | **x** | **x** |
| **glance** | | **x** | **x** | **x** | **x** | **x** | **x** | **x** | **x** | **x** | **x** | **x** | **x** |
| *horizon* | | | | | *x* | *x* | *x* | *x* | *x* | *x* | *x* | *x* | *x* |
| **keystone** | | | | | **x** | **x** | **x** | **x** | **x** | **x** | **x** | **x** | **x** |
| **cinder** | | | | | | **x** | **x** | **x** | **x** | **x** | **x** | **x** | **x** |
| **neutron** | | | | | | **x** | **x** | **x** | **x** | **x** | **x** | **x** | **x** |
| *ceilometer* | | | | | | | | *x* | *x* | *x* | *x* | *x* | *x* |
| *heat* | | | | | | | | *x* | *x* | *x* | *x* | *x* | *x* |
| *trove* | | | | | | | | | *x* | *x* | *x* | *x* | *x* |
| *sahara* | | | | | | | | | | *x* | *x* | *x* | *x* |
| *ironic* | | | | | | | | | | | *x* | *x* | *x* |
| *barbican* | | | | | | | | | | | | *x* | *x* |
| *congress* | | | | | | | | | | | | *x* | *x* |
| *designate* | | | | | | | | | | | | *x* | *x* |
| *manila* | | | | | | | | | | | | *x* | *x* |
| *mistral* | | | | | | | | | | | | *x* | *x* |
| *zaqar* | | | | | | | | | | | | *x* | *x* |
| *cloudkitty* | | | | | | | | | | | | | *x* |
| *freezer* | | | | | | | | | | | | | *x* |
| *magnum* | | | | | | | | | | | | | *x* |
| *senlin* | | | | | | | | | | | | | *x* |

**bold = core service**
*italic = optional service*


Copyright ©2016 Moor Insights & Strategy

## APPENDIX B: OPTIONAL SERVICES IN THE OPENSTACK FRAMEWORK

Below is a comprehensive list of optional OpenStack services that fulfill the functional needs of a cloud architecture. They are optional in that some have chosen different methods to provide these necessary functions.

- **barbican**: Securely stores, provisions, and manages of credentials (passwords, keys, X.509 Certificates) and includes a REST API
- **ceilometer**: Collects utilization data (physical and virtual resources) and allows defined criteria to be established to trigger actions
- **cloudkitty**: Provides chargeback and rating
- **congress**: Provides policy supporting governance and compliance
- **designate**: Multi-tenant DNS service with keystone integration
- **freezer**: Automates backup and restore
- **heat**: Orchestration service with an OpenStack-native REST API and a CloudFormation-compatible Query API (allows existing CloudFormation templates to be used)
- **horizon**: Web-based self-service portal used to launch an instance, assign IP addresses, and configure access control
- **ironic**: Provisioning service for bare metal machines instead of virtual machines; includes APIs and plugins and by default will use PXE and IPMI
- **magnum**: Makes container orchestration engines (Docker, Kubernetes) available as first class resources
- **manila**: Shared file service
- **mistral**: Workflow service providing flexible task scheduling to enforce business process and task relationships
- **sahara**: Provides a simple means to provision (and scale already-deployed) Hadoop clusters by specifying several parameters
- **senlin**: Generic clustering service that manages objects exposed by other components (nova, heat, cinder, *etc*.)
- **trove**: Provisions relational and non-relational databases (database as a service)
- **zaqar**: Multi-tenant cloud messaging service for web and mobile developers


Copyright ©2016 Moor Insights & Strategy

## Important Information About This Paper

### *Authors*

Gina Longoria, Senior Analyst at Moor Insights & Strategy
Jimmy Pike, Technologist in Residence at Moor Insights & Strategy

### *Reviewer / Publisher*

Patrick Moorhead, President & Principal Analyst at Moor Insights & Strategy

### *Editor / Design*

Scott McCutcheon, Director of Research at Moor Insights & Strategy

### *Inquiries*

Please contact us to discuss this report, and Moor Insights & Strategy will respond promptly.

### *Citations*

This note or paper can be cited by accredited press and analysts but must be cited in-context, displaying author's name, author's title, and "Moor Insights & Strategy". Non-press and non-analysts must receive prior written permission by Moor Insights & Strategy for any citations.

### *Licensing*

This document, including any supporting materials, is owned by Moor Insights & Strategy. This publication may not be reproduced, distributed, or shared in any form without Moor Insights & Strategy's prior written permission.

### *Disclosures*

Intel is a research client of Moor Insights & Strategy, and this paper was commissioned by Intel. Moor Insights & Strategy provides research, analysis, advising, and consulting to many high-tech companies mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

### *Disclaimer*

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. Moor Insights & Strategy disclaims all warranties as to the accuracy, completeness, or adequacy of such information and shall have no liability for errors, omissions, or inadequacies in such information. This document consists of the opinions of Moor Insights & Strategy and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice.

Moor Insights & Strategy provides forecasts and forward-looking statements as directional indicators and not as precise predictions of future events. While our forecasts and forward-looking statements represent our current judgment on what the future holds, they are subject to risks and uncertainties that could cause actual results to differ materially. You are cautioned not to place undue reliance on these forecasts and forward-looking statements, which reflect our opinions only as of the date of publication for this document. Please keep in mind that we are not obligating ourselves to revise or publicly release the results of any revision to these forecasts and forward-looking statements in light of new information or future events.

©2016 Moor Insights & Strategy. Company and product names are used for informational purposes only and may be trademarks of their respective owners.