

October 2015

OPENSTACK USER SURVEY

A snapshot of OpenStack users' attitudes and deployments



openstack.org



INTRODUCTION

This is the sixth survey of OpenStack users since April 2013, with a goal of better understanding attitudes, organizational profiles, use cases, and technology choices across the community's various deployment stages and sizes.

This survey report analyzes respondents who completed or updated the survey during a two-week window in September 2015, and questions represent some modifications from prior surveys in keeping with the evolution of the OpenStack platform.

The survey questions and report are guided by the OpenStack User Committee with the support of Foundation staff and an independent data scientist.

This survey represents a snapshot of 1,315 users and 352 deployments, provided voluntarily. Insights are intended to provide feedback to the broader community, and to arm technical leaders and contributors with better data to make decisions regarding the roadmap and feature enhancements.

The User Survey is not a market survey and does not express all OpenStack deployments worldwide.

When reading the report and statistics, remember that the typical respondent is an informed user or someone actively engaged in the OpenStack community.

We have not attempted to normalise the survey data to be more widely representative, since we believe respondents are more representative of decision makers.

Contents

3	Demographics
7	User Perspectives
15	App Developer Perspectives
18	Deployments
25	Deployment Decisions
34	Cloud Size
38	Current Issues
42	Methodology

Survey Credits

The OpenStack User Committee includes Subbu Allamaraju, Tim Bell and Jon Proulx.

OpenStack Foundation staff Jonathan Bryce, Lauren Sell, Tom Fifield and Heidi Joy Tretheway supported the creation of the survey, presentation and narrative report.

The OpenStack Foundation partnered with independent data scientist Kelly Valade to analyze and chart the data.



EXECUTIVE SUMMARY

In OpenStack's sixth and most comprehensive user survey, OpenStack shows its **increasing maturity**, with 60% of deployments in production compared to just 32% reported in our survey less than two years ago.

There are extremely high rates of adoption of OpenStack's core services, with significant interest in more than a dozen additional projects, indicating **future growth and adoption**.

More than twice as many OpenStack users participated in this survey cycle. Their No. 1 business driver for choosing OpenStack being the ability to **accelerate innovation** and compete by deploying applications faster.

Most users indicated strong interest in **emerging technologies** such as containers, Network Functions Virtualization (NFV), and Platform as a Service (PaaS).

The **flexibility** of the platform is readily demonstrated, exemplified by the top five block storage drivers eking a few percentage adoption points to a veritable plethora of others, the proliferation of container systems, and the continued rise of non-hypervisor drivers in use with Nova such as Ironic bare metal.

As always, there are **areas for improvement**. Taking the biggest hits this round were Neutron, which was criticized for its **complexity**; Ceilometer, for its **scalability**; and **documentation** for a variety of issues.



OpenStack is loved by its users
for its technology
and its community.
Users with deployments in
production are OpenStack's
greatest advocates.



OpenStack's overall complexity remains a barrier to entry, with the "big tent" governance change creating confusion among many users and concerns about a lack of **focus** and coordination among services.

On the whole, OpenStack continues to be loved by its users for its technology and its community, with dozens of comments lauding **inclusiveness, access, and support**.

OpenStack earned a Net Promoter Score of 25, which is substantially higher than the software industry average of 19, and among users with deployments in production that score rockets to 43, putting OpenStack in the top tier for excellence.





DEMOGRAPHICS

SURVEY PARTICIPATION DOUBLES

OpenStack’s sixth user survey of the community saw strong participation in a compressed timeframe.

While the prior survey was open for six weeks, this survey period focused on a two-week window to gather a focused snapshot of users’ attitudes and applications of OpenStack software.

	MAY 2015	SEPT 2015	CHANGE
Completed surveys	634	1,315	+107%
Deployments	270	352	+30%
Surveys with multiple deployments	41	85	+107%
Most deployments from one person	5	16	+220%
Application development responses	230	392	+70%

How does the community work with OpenStack?

The OpenStack community showed its varied interests and abilities in answering the question, “What are the ways you work with OpenStack?” Nearly 29% of those surveyed—more than 350 people—indicated they work in multiple disciplines. Among those who listed “other” as their role were architects, end

users, marketers, students and developers learning OpenStack, translators, UX designers, technical writers, testers, salespeople, researchers and vendors. Based on the number of “other” responses, it may make sense to clarify or add response options next cycle.




Figure 1.1 n=1220

DEMOGRAPHICS

Which industries use OpenStack?

The Information Technology industry continues to dominate OpenStack users who took this survey, making up nearly two-thirds of all responses (64% in this survey cycle, compared to 54% in the prior survey). Specifically, the IT services provided include:

- Big data
- Call center
- Cable TV and ISP
- Data center / co-location
- Government identity & security
- HPC
- Online video games
- Professional services / IT consulting
- Public cloud / hosting / private cloud services
- Software development / DevOps / delivery
- Storage solutions
- System Integrator
- Telco and networking
- Web hosting / email hosting

The top industry segments after IT were telecommunications (12%), academic/research (10%), and 2% each of film/media, finance, and manufacturing/industrial.

Compared to the prior survey, there are only slight variations of representation by industry, except for academic/research, where the user survey responses made up 19% of our total last cycle but just 10% this cycle.

This should be seen as indicative of a difference in survey population, not necessarily in OpenStack’s population overall.




Figure 1.2 n=1130

Breaking out the “other” segment, which represents 8% of responses are government/defense, retail, healthcare, energy, consumer goods and professional services. Additional industries of note include advertising, automotive, ecommerce, gaming, and online travel agency.

We also looked at which industries are using OpenStack in production deployments. More than half of most industries have OpenStack in production, which is consistent with the overall 60% of deployments recorded in this survey at a production stage.

INDUSTRIES BY DEPLOYMENT STAGE

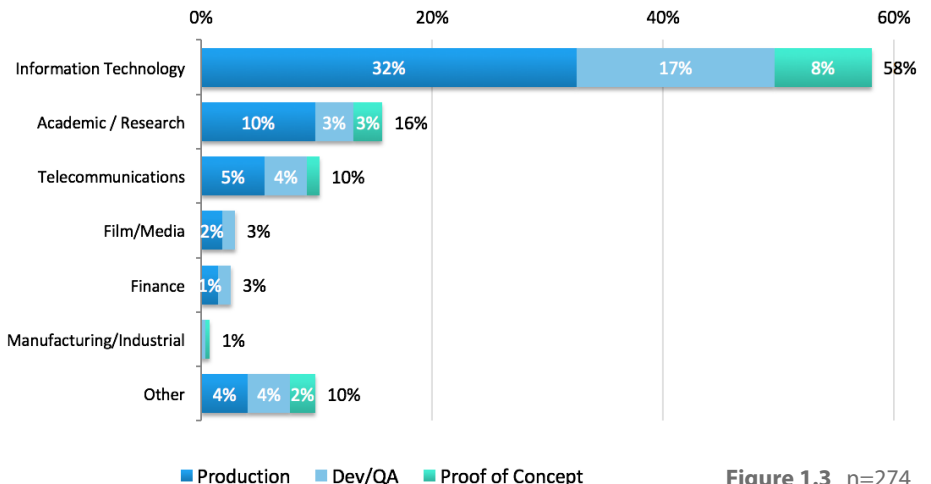


Figure 1.3 n=274

Where in the world are OpenStack users?




Figure 1.4 n=1183

OpenStack boasts a thriving global community of users. Among the 76 countries represented in this survey cycle, the most frequently represented are:

- United States: 39%
- China: 8%
- India: 7%
- Japan: 6%
- France: 4%
- Germany: 4%
- United Kingdom: 3%
- Canada: 3%
- Brazil: 2%

Growth in Asia has accelerated, now representing 28% of responses, up from 23% just six months ago. South America was also represented slightly more in this survey.

Drilling down further, we looked at the cities in which OpenStack users are located—nearly 500 in all. Top cities with high concentrations of OpenStack users included Silicon Valley area cities, plus Raleigh, San Antonio, Tokyo, London, Bangalore and Beijing.

DEMOGRAPHICS

What size organizations use OpenStack?

The median size of a business using OpenStack falls into the 1,000 to 9,999 employee range, and 22% of OpenStack users surveyed classified their organization as such. The even distribution of users across all sizes suggests that OpenStack is a strong solution for both enterprise and developing companies.

We also compared organizational size to the prior survey. In this cycle, larger organizations were more significantly represented than smaller organizations, which shows that the survey population is different, but not necessarily that OpenStack users as a whole are changing.




Figure 1.5 n=1186

CHANGE IN OPENSTACK USER SIZE SINCE LAST SURVEY

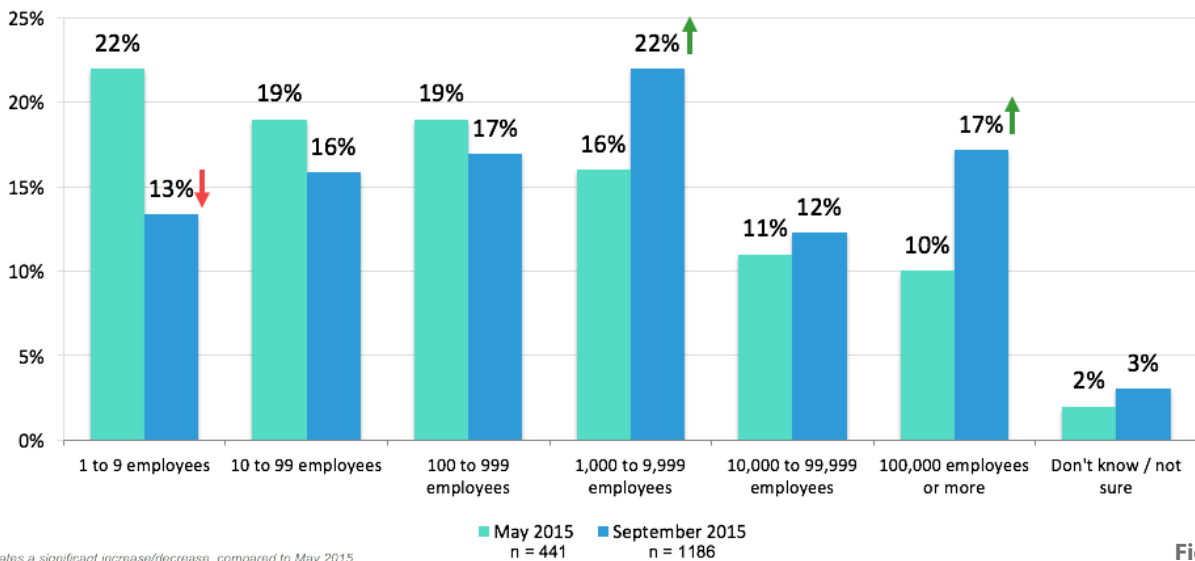


Figure 1.6



Why do organizations choose OpenStack?

According to those surveyed, three factors were most often cited in influencing their decision to use OpenStack. About three-quarters of respondents ranked these among their top five business drivers:

- Accelerating my organization’s ability to innovate and compete by deploying applications faster (77%);
- Avoiding vendor lock-in with an open platform and ecosystem, including flexibility of underlying technology choices (76%); and
- Increasing operational efficiency (75%).

Another perspective on this data considers the relative rank (from first to fifth most important) of the seven drivers presented in this study. Of the 67% of users surveyed who listed “Save money over alterna-

tive infrastructure choices” as a motivating factor, all of them listed it as their No. 1 reason.

Operational efficiency earned 23% of first-place ranks and the majority of third-place ranks, which is consistent with results from the previous survey.

Many respondents indicated additional business drivers beyond the seven specified in our list. Some of these reasons included:

- Productizing OpenStack (cloud / solutions provider)
- Support use of open source software
- Be on the cutting edge / participate in the hype
- Orchestrate or offer NFV platform
- Fast prototyping, automating testing
- Third party compatibility

TOP BUSINESS DRIVERS

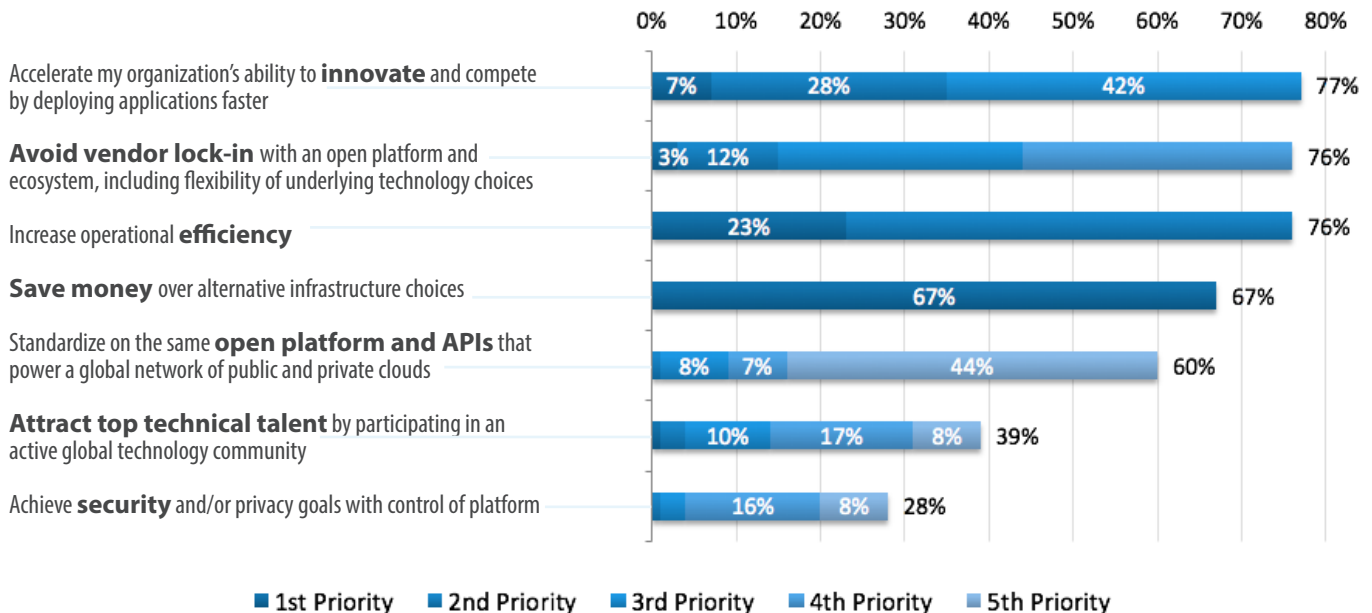


Figure 2.1 n=953

How likely are users to recommend OpenStack?

We asked the classic Net Promoter Score (NPS) question, “How likely are you to recommend OpenStack to a friend or colleague?” Responses ranged from 0–10, with 10 being the highest.

Answers fell in these three categories:

Score 9 or 10: “Promoters”

Score 7 or 8: “Passives”

Score 0 through 6: “Detractors”

In keeping with the NPS methodology advised by Satmetrix, OpenStack’s score is calculated by taking the percentage of promoters (those rating OpenStack a 9 or 10) and subtracting the percentage of detractors (ratings of 0-6). The result can range from -100 to +100. OpenStack’s Net Promoter Score was 25.

Context is critical when viewing this number. Consider the NPS scores released by Satmetrix in its 2015 US Consumer Study, published in May 2015. Satmetrix surveyed 30,000 respondents in 22 categories. The highest NPS scores were achieved by department/specialty stores (industry average NPS ~57) brokerage/investments (industry average ~45), tablet computers (industry average ~44) and smartphones (industry average ~40).

Software and apps earn typically lower NPS scores (industry average ~19). Satmetrix released a chart detailing NPS high/low scores earned by industry, as well as the 25th, 50th and 75th percentile scores, and the industry average.




Figure 2.2 n=893

In the Software and Apps category, the scores were approximately:

- Lowest: -10
- 25th percentile: 7
- 50th percentile: 22
- 75th percentile: 28
- Top score: 53

In context with the prior survey, we found a marked drop in the NPS score, however this could be attributed to several factors including a major change in the population sample (270 responses in May, and 893 in October), and the fact that the rating scale might have been somewhat ambiguous.

USER PERSPECTIVES

For this cycle’s NPS score, we also wanted to know if specific categories of users had strongly varied NPS responses. Respondents self-identifying as app developer gave OpenStack an NPS of 30; cloud operators registered an NPS of 27, and upstream developers ranked it 26. The NPS score is lower due to individuals identifying as “other” types of contributors.

Additionally, we looked at NPS scores through the lense of business size. Large users in the 10,000–99,999 category rated OpenStack most strongly, with a 33 NPS, and the 1,000–9,999 category registered a 28 NPS. However, the data did not indicate that OpenStack was less well-regarded by small organizations; those with just 1-9 employees returned an NPS of 27.

Finally, we looked at NPS scores based on deployment stage (looking only at scores from those who submitted deployment information), and found NPS scores highest among those with deployments in production, suggesting that those actively work-

NPS COMPARISON TO PRIOR CYCLE DATA

	MAY 2015	SEPT 2015
Responses	270	893
Promoters	57%	47%
Passives	32%	31%
Detractors	11%	22%
NPS	46	25

ing with OpenStack deployments are most satisfied with this technology.

In comparison to the prior survey of May 2015, NPS scores were slightly lower overall, and again they were strongest among users with deployments in production. Given that this is only the second time this data has been gathered, trends with data significance are not yet available.




Figure 2.4

Why do users recommend OpenStack—and why don't they?

We followed up the Net Promoter Score rating question with, “What is the primary reason for your rating?” Among OpenStack promoters, flexibility, functionality, community, an open and standardized API, and ease of use were reasons they chose to recommend it.

“I have been a member of the OpenStack community for over two years and during that time I’ve seen small startups and large enterprises immediately see a return on their OpenStack investment and **reduced their expenditures compared to AWS**,” one user wrote.

“Having worked on OpenStack and worked with several other virtualisation management platforms, OpenStack is the **most flexible and gives the most control**,” another added. Echoing this, a different user said, “OpenStack is a flexible set of projects whose **benefits can be realized by a variety of organizations, large or small**.”

In a similar vein: “OpenStack is currently the best choice of available options for operating a public or private cloud. It has broad community and industry support and a consistent set of APIs (both for cloud admins and users).”

“The **pace at which innovation is happening** at OpenStack is truly amazing and will be very hard for any other technology to catch up,” said another. “The host of features available and the possibilities excite me.”

Among detractors, OpenStack is seen as complex and fragmented. “Development is happening at a rapid pace, as is needed, but **‘productization’ is lagging**,” said one. “It will come, but it’s still sometimes ‘a box of nuts and bolts, assembly required.’

Documentation is lacking and some things simply don’t work.”

Another commenter said, “There’s a lot of great things in OpenStack, but it’s a Homercar. Its features sound great, but the resulting platform/product is impractical. It’s **designed for consultants and implementation specialists** who speak OpenStack 24x7. It’s an awful product with a lot of cognitive overhead for a business that wants to just get work done, but also get that work done on OpenStack.”

“It **requires certain skills and knowledge** to use and adjust a business. So it might not be recommended to some of colleagues,” added one user in the “passive” category. Another in that category said, “It’s **relatively stable and far too complex**. Plugins are still difficult and it requires a team instead of an individual to deploy and maintain.”

In general, passives said their recommendation would be context-sensitive, based on whether the tool was right for the job, and whether the operator was capable of using it well.

Speaking to complexity, a user commented, “OpenStack is complex, but it provides a path toward offering cloud computing solutions, be it virtualization, containerization, PaaS or SaaS infrastructure that is **easy to integrate** with vendor and open source solutions to build a public or private cloud that **matches business objectives**.”

Another added, “OpenStack is now a **reliable technology** to make software development and deployment on a cloud platform **easier, faster and cheaper**. **No vendor lock-in**; plain vanilla functionality. It will be the de facto standard for generic cloud platforms.”



There is no stasis in OpenStack. Problems are presented to the community and then they are solved. There always lies something on the horizon—a goal to be reached—and masses of individuals working in concert toward common change.



What do users like most about OpenStack?

Community was cited frequently as the reason users appreciate OpenStack. “The community is so large, but everyone in it seems to value and respect one another. This creates a **safe place for anyone**, client or vendor, to **grow and thrive**,” said one user.

“Community, community, community. This is one of OpenStack’s biggest strengths,” another user added. “Not just the personal interactions (ability to talk directly to devs when things go pear shaped) but also the vendor ecosystem. There are **many channels** for me as an operator to be able to **help shape OpenStack’s future**—and that’s important.”

On a similar note, a user wrote, “I really enjoy working with a **community that extends beyond the company I work for**—it opens up conversations about goals and means to them that wouldn’t otherwise happen.” Another added, “I like being able to influence and direct.”

As a result of this positive, connected community, there is substantial talent to work on hard problems, many users said, and an “active community inno-

vating and **producing features FAR quicker than commercial vendors**.”

That rapid pace of development was a significant theme among commenters. “[There is] **constant evolution**,” wrote one user. “There is no stasis in OpenStack. Problems are presented to the community and then they are solved. ... You can always feel the march of progress. There always lies something on the horizon—a goal to be reached—and masses of individuals **working in concert towards common change**. That overwhelming, constant drive provides both technological and psycho-social benefits in our organization. OpenStack pushes our technology further as we are driven to be more than what we currently are.”

OpenStack’s open framework and APIs, enabling interoperability and preventing vendor lock-in, were another top reason users gravitate toward OpenStack. “**‘Open’ is everything in software**. Without ‘open,’ there is only proprietary software, and history has shown us the ineffective nature of proprietary software.”

USER PERSPECTIVES

“OpenStack provides a **neutral and agnostic** mechanism to manage an agile infrastructure in the Data Center,” wrote another survey-taker. Many users cited OpenStack’s flexibility, and the benefits of “abstraction from vendor choices—agility of being able to replace underlying technologies without having to change dev access.”

OpenStack’s commitment to being open source enhances speed and accessibility, said many users. “We can deploy OpenStack—fix it—**see what it is doing without being slowed down by licensing issues or lack of access**,” a user wrote. “You can learn things in your basement. It’s becoming a standard way to access cloud/vendor API. [It is] actively developed—bugs get fixed, features get added, [and] there is room for innovation.”

Another user appreciated the “ability to build complex environments with automation out of relatively

simple to understand and flexible services.” Scalability when handling complex needs was critical to many users.

“Modularity makes dealing with with OpenStack easier, as most things are divided into small, **logical services that focus** on one thing,” another user told us, adding that it is “a safe choice for long-term system development.”

“My career path is shaped around the success of OpenStack,” said a user, who appreciates the vibrancy of the community. “People talk about OpenStack being the default go-to private cloud. I like Horizon being a first-class citizen, rather than a pretty add-on. The **maturity of the product** is making the dismissers of OpenStack in the early days take note and change their minds.”

“

My career path is shaped around the success of OpenStack.
The maturity of the product is making
the dismissers of OpenStack in the early days
take note and change their minds.

”

Which areas of OpenStack require further enhancement?

The key themes that emerged when we asked users about where OpenStack needs to grow include the complexity of Neutron, scalability of Ceilometer, documentation, projects that don't keep up with other components of releases, and the overall complexity of OpenStack in terms of deployment, configuration, and upgrades.

Additionally, there remains some confusion about OpenStack's new "Big Tent" approach to releases, which focuses on a core set of six services and includes a "big tent" of many more projects than were included in integrated releases for the May 2015 and prior cycles. Some users were concerned that a big tent model produced a lack of focus or coordination among projects.

One user said the most important improvement to OpenStack must be "**Cohesive architecture**. Every single component seems to reinvent the wheel when it comes to providing an API server and a client app. These should be integrated completely. If I want to spin up a new service I should be able to take the OSLO API server and configure it to my needs. Ideally with little coding. Neutron has an interesting model of plugins, which is just different from the other subsystems. Until some single person or team nails this down and enforces a single worldview in terms of how this stuff should work. Today, **all of the client apps work in slightly different ways.**"

Another commenter highlighted, "**documentation** and ease of building docs. I would like to contribute to them but lack time and interest in building a system to build docs. I pulled down the repo once and it's stagnant on my machine with a single edit because of the process of adding to them and committing them upstream." In a similar vein, another commenter said, "I would most like to see effort sunk in



The most important improvement must be cohesive architecture.



fixing bugs, increasing test coverage, and paying down technical debt."

Interestingly, while many users wanted greater interoperability and integration among projects, some asked for more loosely coupled components to enable them to run different components together. Additionally, several users and operators felt that projects didn't take their needs into account adequately. Several called for **operations and development to collaborate more** closely, and some users noted that upgrades require too many hours of work for small teams.

"In my mind, the mission statement of OpenStack as an IaaS is to provide a single consistent API to wrap all of the underlying physical infrastructure. But the **user experience is fragmented, and there is no clarity or focus on the mission statement of each [project]** within the OpenStack ecosystem," wrote one user. "I'll pick on compute. ... Nova is very, very much entrenched in VM and has not embraced the other variants of compute now supported by OpenStack (containers and bare metal)... This is not a graceful, simple or beautiful experience for users. **Nova should be agnostic of the underlying compute technology**—be it VM, BM, containers, or whatever is going to come next."

USER PERSPECTIVES

Which new/emerging technologies are OpenStack users interested in?

The User Survey often provides a forward-looking sense of companies that are considering or testing emerging technologies. Of the 841 survey responses answering this question, more than three-quarters were interested in containers, followed by 59% interested in Network Functions Virtualization (NFV).

OpenStack users have a wide variety of interests in emerging technologies, as evidenced by several dozen additional responses to this question. Other technologies of note that are on our users' radar include:

- ARM-based virtualization
- Bare metal
- CI systems
- Cloud federation
- DBaaS (Trove)
- Desktop as a service
- Integrated HA / improved VM migration
- Real-time data processing & Apache Zeppelin
- SDN (OpenDaylight)
- Standard/proprietary APIs
- VPC
- Windows Nano-Server




Figure 2.5 n=841

How often do users refer to documentation?

OpenStack's documentation at docs.openstack.org continues to be a well-used resource, with almost a quarter of users checking it daily and the vast majority using it at least weekly.

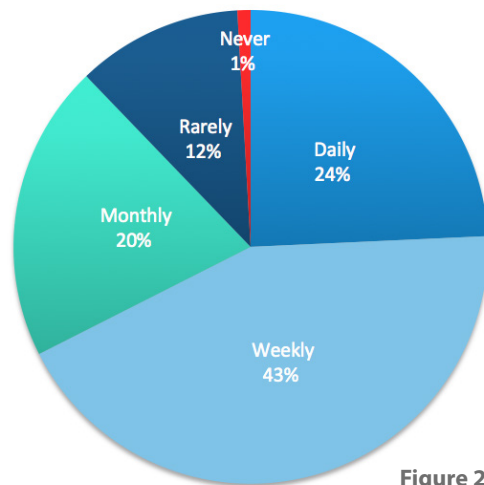


Figure 2.6 n=577



What toolkits are application developers using with the OpenStack API?

As in the last cycle, we used survey logic to supply this section's questions only to those survey-takers who self-identified as an OpenStack application developer.

More than 350 app developers weighed in on this question. Note that the total number of responses exceeds 100% because users had the option to select more than one response in all three of these application developer questions.

Comparing the rankings of SDKs to six months prior, the libcloud remains most popular, while fog and jclouds continue the battle for second and third place. The javascript library pkgcloud has leapfrogged PHP-opencloud and Deltacloud, as did the SDK for the Microsoft platforms, OpenStack.Net.



Libcloud remains most popular, while fog and jclouds battle for second and third place.



Other toolkits noted by our users included Gophercloud (1%), shade (1%), Juju and OpenStack4j (Java). The number of users who were not using a toolkit remained stable, potentially hinting at issues with the uptake of SDKs.




Figure 3.1 n=361

What improvements to OpenStack would enhance developing applications?

We asked app developers—those users who indicated in the demographic portion of this survey that they were involved in OpenStack development—to share their experiences. Of greatest concern was documentation, and 22% of respondents mentioned this issue, including struggling with documentation that was frequently nonexistent, inaccurate for the version of the API they were using, or lacking code samples to make it useful.

Following this, 10% of commenters noted that interoperability among clouds was an issue, with one responding, “Half of the time involved in developing new applications is spent troubleshooting behav-

ior ‘quirks’ and reporting them via provider tickets/upstream bugs, then coming up with suitable workarounds until proper fixes are eventually deployed.”

Closely related was consistency of APIs design and implementation between projects, with several users expressing frustration at the different way resources are managed differently in the various client tools. Quality of other SDKs also got a bad rap.

Closing out the list were many project- or feature-specific issues, with fewer comments about complexity and reliability.

With which other clouds do app users interact?

Nearly 250 of the 392 app developers weighed in on this question. Use of an OpenStack cloud with Amazon remains most popular, though compared with the previous survey the use of multiple OpenStack clouds for an application has increased substantially.

In the “other” category, VMware was most often noted, followed by Digital Ocean, CloudStack, customer private clouds, Fujitsu public cloud, Softlayer and Joyent.




Figure 3.2 n=248

APP DEVELOPER PERSPECTIVES

What stack do app users choose?

Nearly 250 app developers weighed in on this question. Among "other" stacks, these were most often noted:

- Linux- Apache- MySQL- Python
- Linux- Apache/Nginx- MySQL- Python
- Linux- Django- Arakoon- Javascript
- Linux- Nginx- MySQL- Python
- Linux- Nginx- Postgres- Go
- Linux- Nginx- Python- uWSGI
- Linux+docker- nginx- postgresql- rails+node
- Linux- Apache- Python WSGI- HTML/JavaScript
- Linux- Apache- Nginx- Ruby on Rails- node.js- Erlang- Postgres




Figure 3.3 n=245



DEPLOYMENTS

In what stage are OpenStack deployments?

In this User Survey cycle, we see that OpenStack is increasingly mature, with more deployments moving into production stage as compared with prior survey data. Production deployments are nearly double the rate reported in the survey conducted two years prior.

- September 2015
- May 2015
- November 2014
- May 2014
- November 2013




Figure 4.1 n=351

Where are OpenStack deployments located?

In North America, the United States had 81% of surveyed deployments. In Asia, Japan has 36% of surveyed deployments, followed closely by China, with 30%.

In Europe, deployments were dispersed among 21 countries, with Poland deployments representing

20% of these, followed by 16% in the United Kingdom and 12% in France.

When comparing this geographic data on deployments to the locations of survey-takers, there are slightly more deployments represented in North America and Europe.

DEPLOYMENTS

What type of clouds are running OpenStack?

We saw the majority of deployments reported as on-premise private clouds, and well over half of these are in production.





Figure 4.2 n=261



DEPLOYMENTS

Which releases are deployments using?

When the survey was conducted, OpenStack’s twelfth software release, Liberty, was two to four weeks away from its release date of Oct. 15, 2015.

We asked which releases all deployments were using, and found them fairly evenly distributed over the most recent three releases—Kilo, Juno and Icehouse. The total number of responses adds to great-

er than 100% because some deployments indicated more than one release.

Additionally, we segmented this data to focus on production deployments only, which represent 60% of the total reported deployments. This chart shifts slightly back, but remains concentrated on the three most current releases at the time of this survey.

ALL DEPLOYMENTS, BY RELEASE

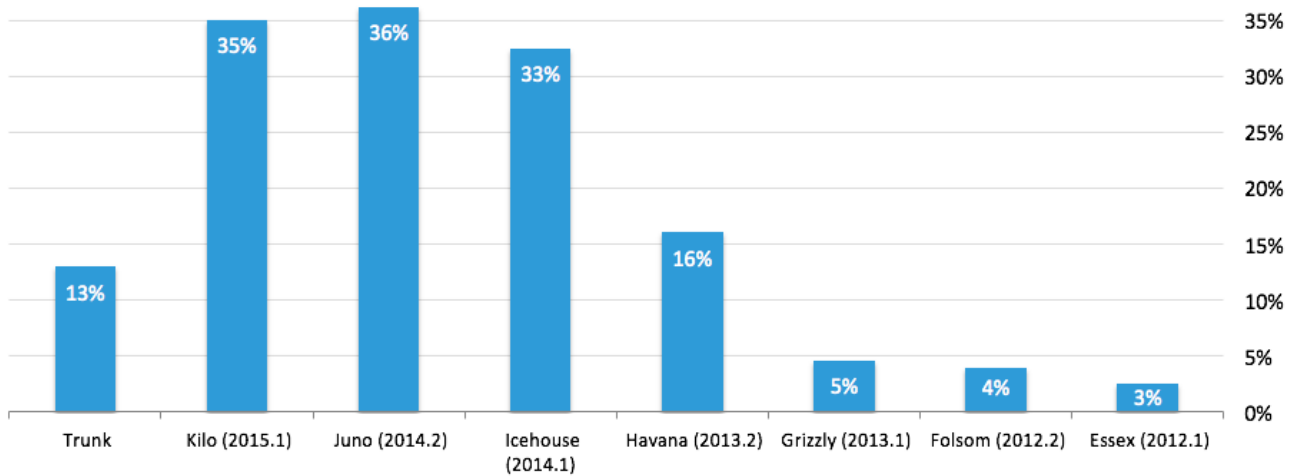


Figure 4.3 n=348

PRODUCTION DEPLOYMENTS, BY RELEASE

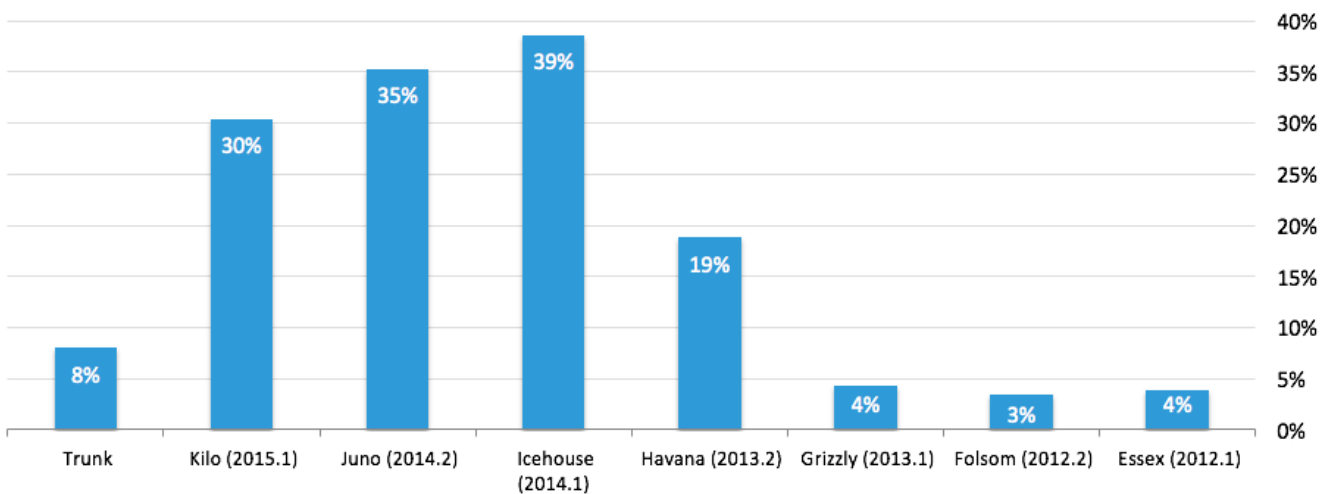


Figure 4.4 n=207

DEPLOYMENTS

PRODUCTION DEPLOYMENTS, BY RELEASE, OVER TIME

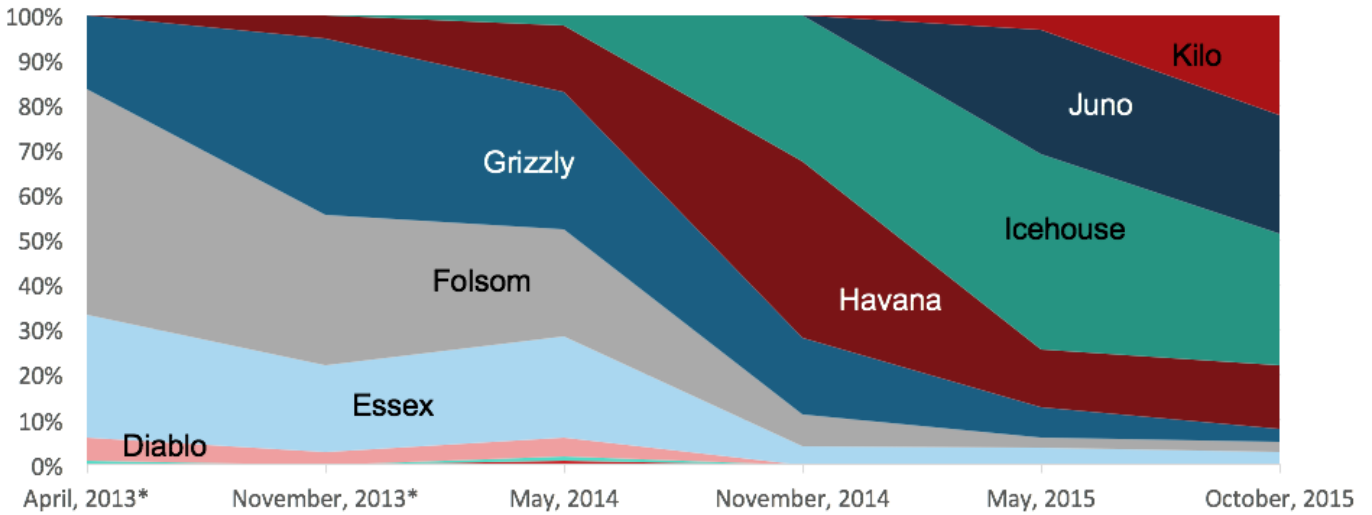


Figure 4.5

We also looked at the OpenStack releases in production deployments over time. In the first two instances (both 2013), the asterisk indicates it is based on all deployment data, but in later surveys we refined this count to show production deployments only. It shows a clear transition to more recent releases over time.

DEPLOYMENTS

Which projects do OpenStack deployments use?

OpenStack’s core services—Keystone, Nova, Glance, Neutron, Cinder and Swift—are among the most commonly used OpenStack projects. This chart shows overall use of the most popular projects in production or testing phase.

Neutron showed the highest degree of increased adoption among projects, from 84% to 88%, compared to the prior survey.

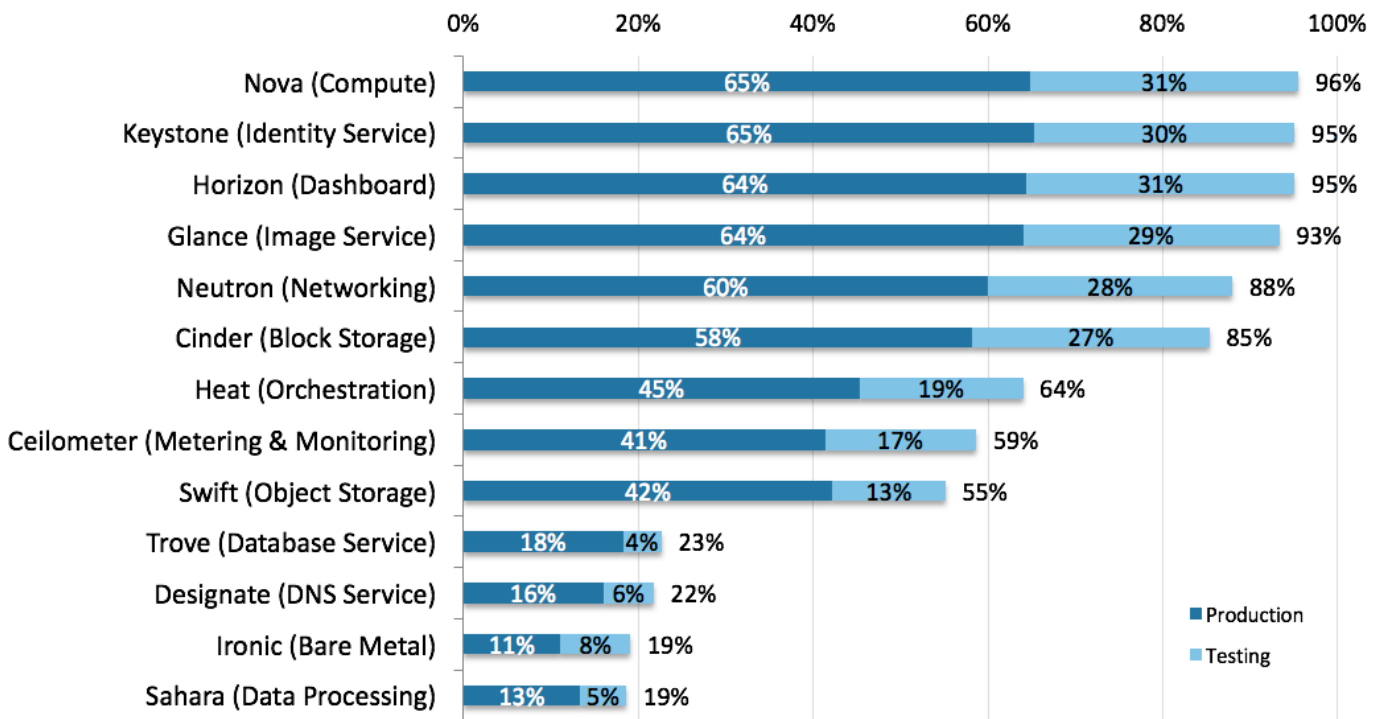


Figure 4.6 n=225

DEPLOYMENTS




Figure 4.7 n=225

Thirteen additional projects are emerging and showing increasing adoption. Figure 4.7 shows the total percentage of deployments in production or testing that also indicated these projects are in production or testing stage.

So, what does the future of OpenStack look like? In addition to looking at the projects in testing phase for indications on where full production will happen next, we asked users who submitted deployments at any stage if they have interest in using these OpenStack projects in the future (figure 4.8).

Additional projects showed interest at less than 2% response rates. Note that Murano and Magnum, as well as more emerging projects on the “other projects” list, might actually have greater degrees of interest but due to survey design this data was suppressed.

TOP PROJECTS OF INTEREST

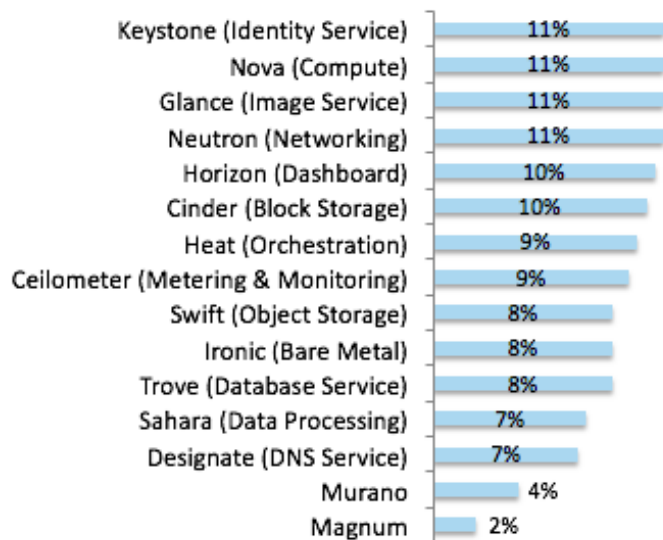


Figure 4.8 n=253

DEPLOYMENTS

Which workloads and frameworks are running on Openstack?

We also asked which workloads and frameworks are running on OpenStack and segmented the data based on the deployment stage, with production deployments again in dark blue.

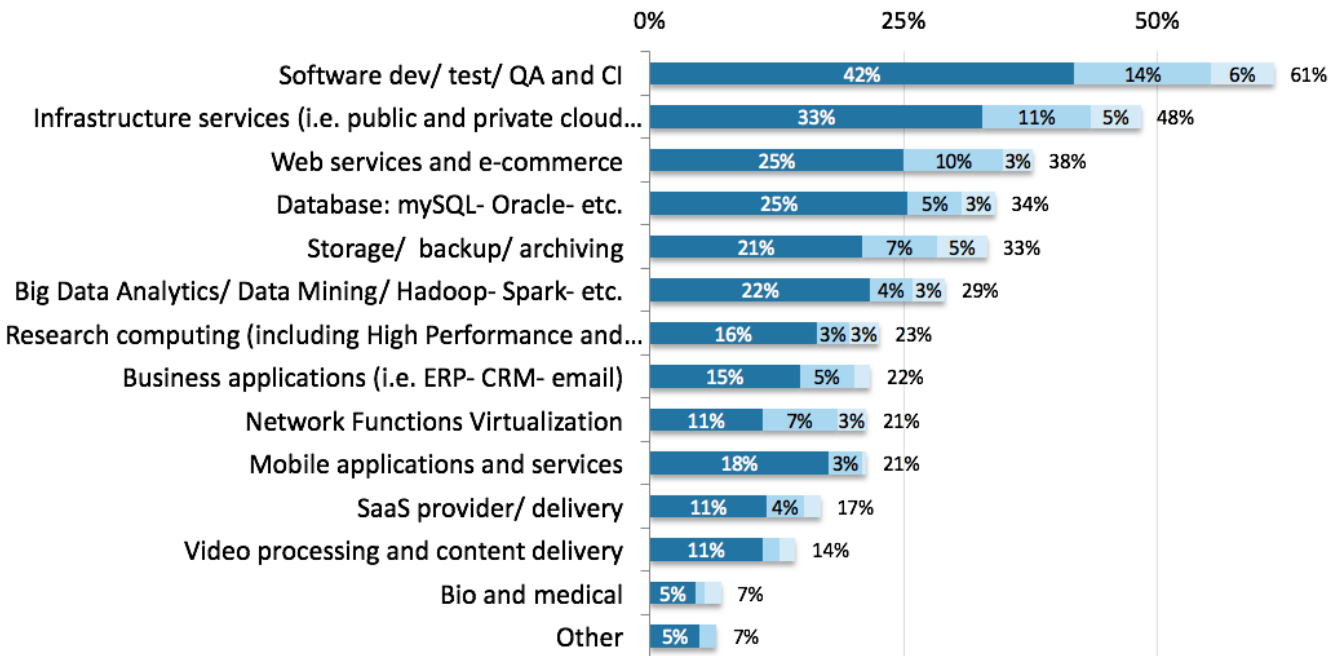



Figure 4.8 n=244





DEPLOYMENT DECISIONS

What packages are OpenStack deployments using?

Most OpenStack deployments use unmodified packages, either OpenStack or through a non-OpenStack source such as vendor distribution.

However, the number of respondents indicating they are having to create their own packages or modify packages in some way has increased by about 10% from last survey.

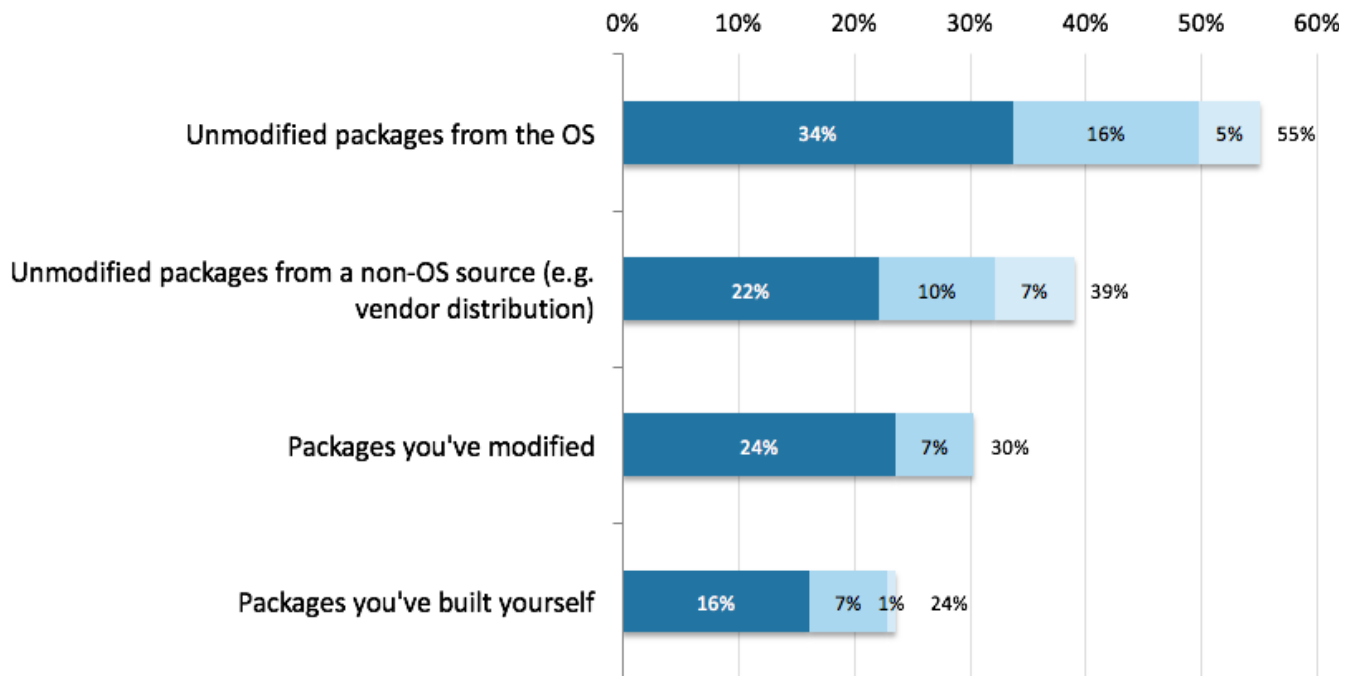



Figure 5.1 n=267



DEPLOYMENT DECISIONS

What tools are used to deploy/configure OpenStack clusters?

Among 261 deployments that indicated which tools they are using to deploy or configure OpenStack clusters, Puppet was the top choice.

From the “other” category, we saw several instances of Foreman, RedHat, Cobbler, Docker, and custom scripts.





Figure 5.2 n=296



DEPLOYMENT DECISIONS

What PaaS tools are used to manage OpenStack applications?

CloudFoundry was the most popular PaaS tool, followed closely by Kubernetes. OpenShift has a large proportion of its deployments in the development/testing stage.

Other PaaS tools were primarily custom, with a few mentioning Apache Aurora, Juju, Murano, and Scalr.





Figure 5.3 n=82



DEPLOYMENT DECISIONS

Which OpenStack Compute (Nova) hypervisors are in use?

KVM provides the majority of Nova hypervisor support. Use of non-hypervisor drivers for nova, such as the Ironic Bare Metal driver, and container systems docker and lxc has increased.




Figure 5.4 n=328

Which databases are used for OpenStack components?




Figure 5.5 n=293

DEPLOYMENT DECISIONS

Which OpenStack Network (Neutron) drivers are in use?

Open vSwitch is the most common choice for Neutron drivers, though the bulk of deployments selecting it are in proof of concept stage.

Among more than 21 network drivers selected by

deployments in this survey, Brocade, Nuage Networks, Big Switch, A10Networks, NEC OpenFlow, Meta Plugin, Ruijie Networks, and IBM SDN-VE were selected in less than 3% of total responses and are therefore not represented on the chart below.




Figure 5.6 n=297

We drilled down in this data to find out which Neutron drivers are used by production clouds with more than 1,000 cores.

There are 43 responses meeting this criteria, so this should be considered directional data only and not necessarily representative.




Figure 5.6.1 n=43

DEPLOYMENT DECISIONS

Which OpenStack Identity Service (Keystone) drivers are in use?

SQL is most often used for the OpenStack Identity Service (Keystone) driver. Other Keystone drivers were primarily custom integrations or layers, or hybrids.





Figure 5.7 n=279



DEPLOYMENT DECISIONS

Which OpenStack Block Storage (Cinder) drivers are in use?




Figure 5.8 n=258

Ceph RBD dominates the choice for Cinder drivers, but more than 20 different drivers were listed by OpenStack deployments. All of the top 5 drivers decreased in relative usage compared to the past survey, as the number of drivers increases overall. Not shown in the chart below are Sheepdog, HDS, Windows Server 2012, Nexenta, Scality, Mellanox, and Huawei, all of which had fewer than 2% of the total deployments using their drivers.

In the “other” category were several custom-built drivers in addition to Brocade, Dell, Hitachi, NEC, Nimble, Quobyte, StorPool and zfs.

We also looked at the largest production clouds (more than 1,000 cores) to find out which Cinder drivers they select.

Among those meeting this criteria, Ceph RBD was most often selected, followed by NetApp, LVM and SolidFire.

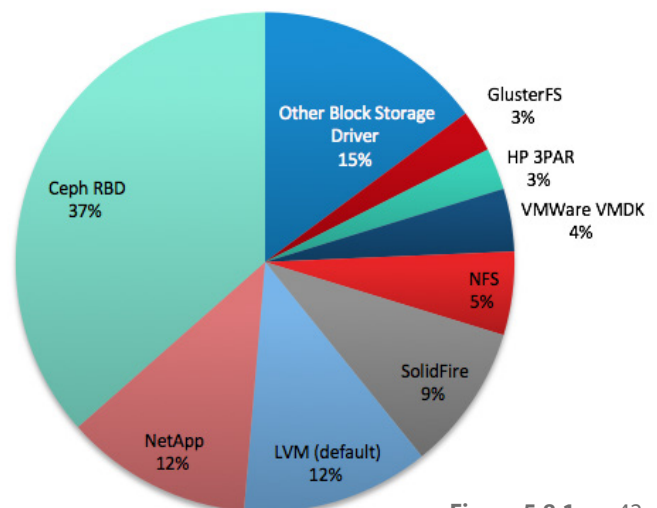


Figure 5.8.1 n=42

DEPLOYMENT DECISIONS

Which operating systems are running OpenStack deployments?

Ubuntu Server provides the majority of operating systems for running OpenStack deployments, increasing its lead by 10% over the past six months. CentOS remains steady in second place with about

a quarter of deployments, while its sister RHEL increases by 2%. SLES increased by 1%, as did Fedora and the Scientific Linux distribution.




Figure 5.9 n=281




Figure 5.9.1 n=35

Going one level deeper, we found that among OpenStack deployments with more than 1,000 users, of which there were a nonrepresentative set of 35 responses, 65% reported using the Ubuntu server, followed by CentOS selected by 23%.

DEPLOYMENT DECISIONS

Among clouds that support compatibility APIs, which APIs are supported?

Roughly one-third of deployments responded to this question.





Figure 5.10 n=111





Throughout this section, we focused on data from deployments in production or testing, excluding those deployments in early proof of concept stage, to get a more realistic view of cloud size.

How many users do OpenStack clouds support?

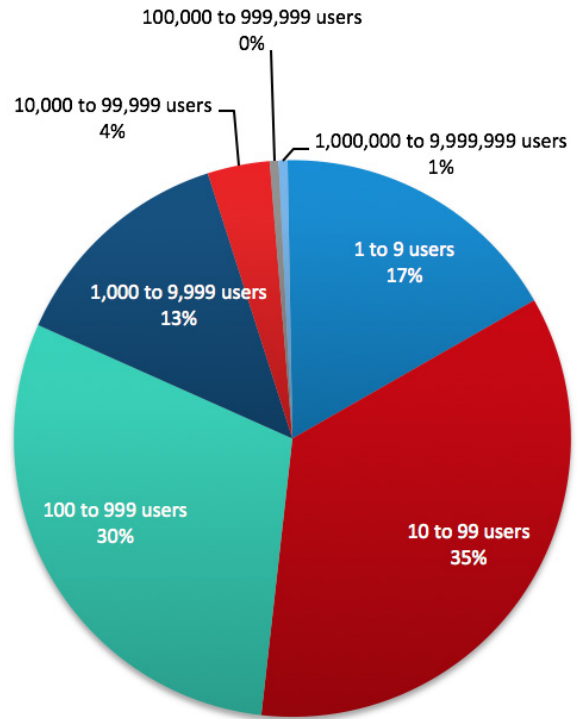


Figure 6.1 n=154

How many physical compute nodes to OpenStack clouds have?

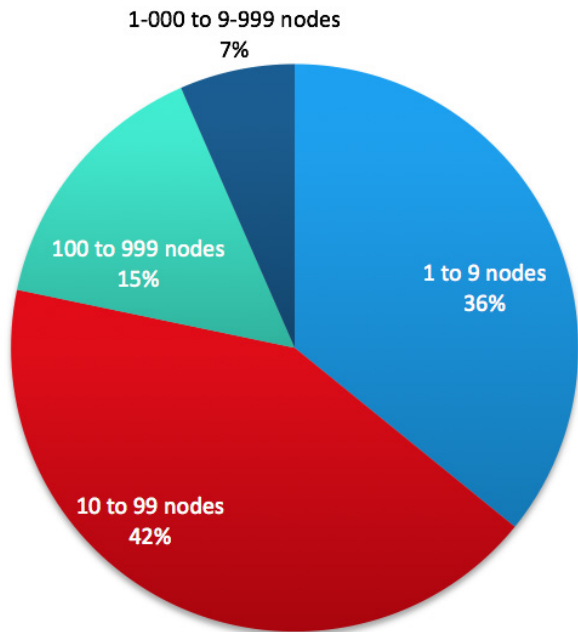


Figure 6.2 n=198

CLOUD SIZE

How many processor cores in an OpenStack cloud?

In addition to our standard data cut, we took a look at the trending growth in OpenStack cores over time, comparing this survey to the prior survey's user data.

The results of this are only directional, as the scale changed slightly and there is a significant difference in the total amount of deployments surveyed.




Figure 6.3 n=195




Figure 6.4

How many instances in an OpenStack deployment?




Figure 6.5 n=195

How many usable IPs?




Figure 6.6 n=188

CLOUD SIZE

What is the size of the deployment's Cinder block storage?




Figure 6.7 n=165




Figure 6.8 n=94

How much Swift object storage is provisioned in a deployment?

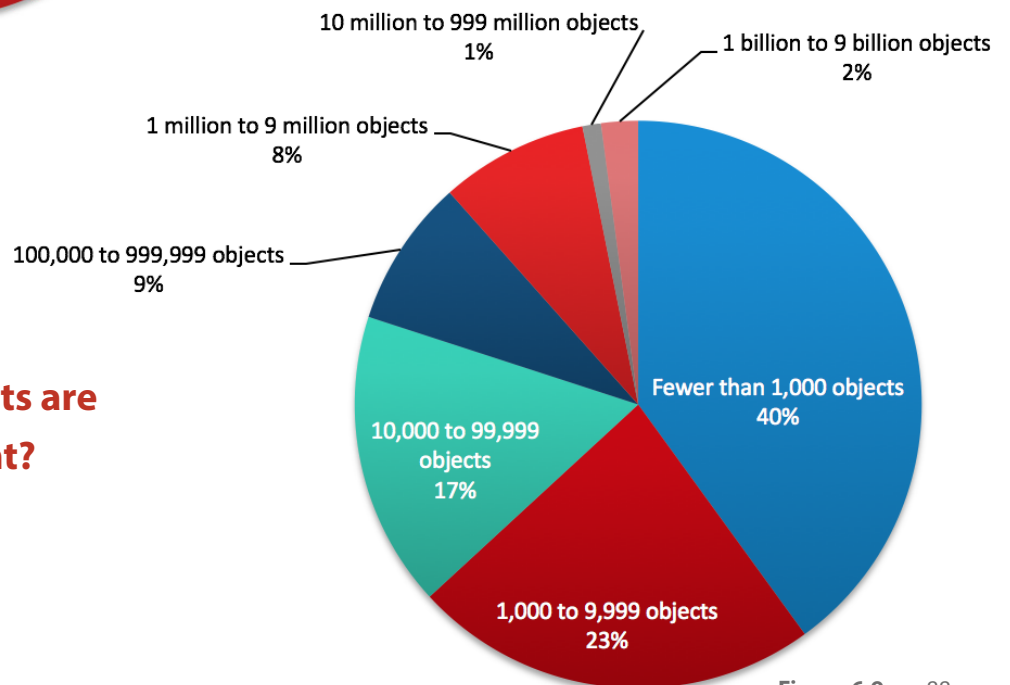


Figure 6.9 n=88



This portion of the User Survey report is conducted at the request of Project Technical Leaders (PTLs), who had an opportunity to submit their own question for consideration.

In most cases, survey logic prompted survey-takers for a response to this question only when they indicated in the deployments section that they were using a given project.

What would allow a deployment using nova-network to migrate to Neutron?

While Neutron is already adopted by 88% of OpenStack users, this question addressed the rest, earning 62 responses to inform developers of barriers to adoption.

In the follow-up question, “Please provide feedback (positive or negative) about Neutron, which will help plan future releases,” simplification was a key theme.

The biggest complaint was its complexity—with 14% of commenters addressing this, such as, “Please look at building a simplified network module that provides a basic but scalable SDN capability.”

A further 14% of respondents referred to related topics such as Linux bridge, difficulty troubleshoot-

ing, flat networks and documentation. There is still a clear desire for high availability throughout the project, which should be achieved in a simple way.

Another 26% of commenters said that fixing bugs, improving reliability, performance and scalability should be areas developers should be working on, compared to 9% who had a feature request and 4% who mentioned the higher-level services.

Few users had positive remarks. “We like Neutron, however, it’s a bit complicated,” said one user. “Compared to nova-network, it’s very complicated.”

Others stated that replacing Neutron with a vendor product had solved some of their problems.





Figure 6.10 n=62

How likely are deployments to use Swift's storage policies or erasure codes in the next year?



Are deployments using Nova's new live upgrade process?

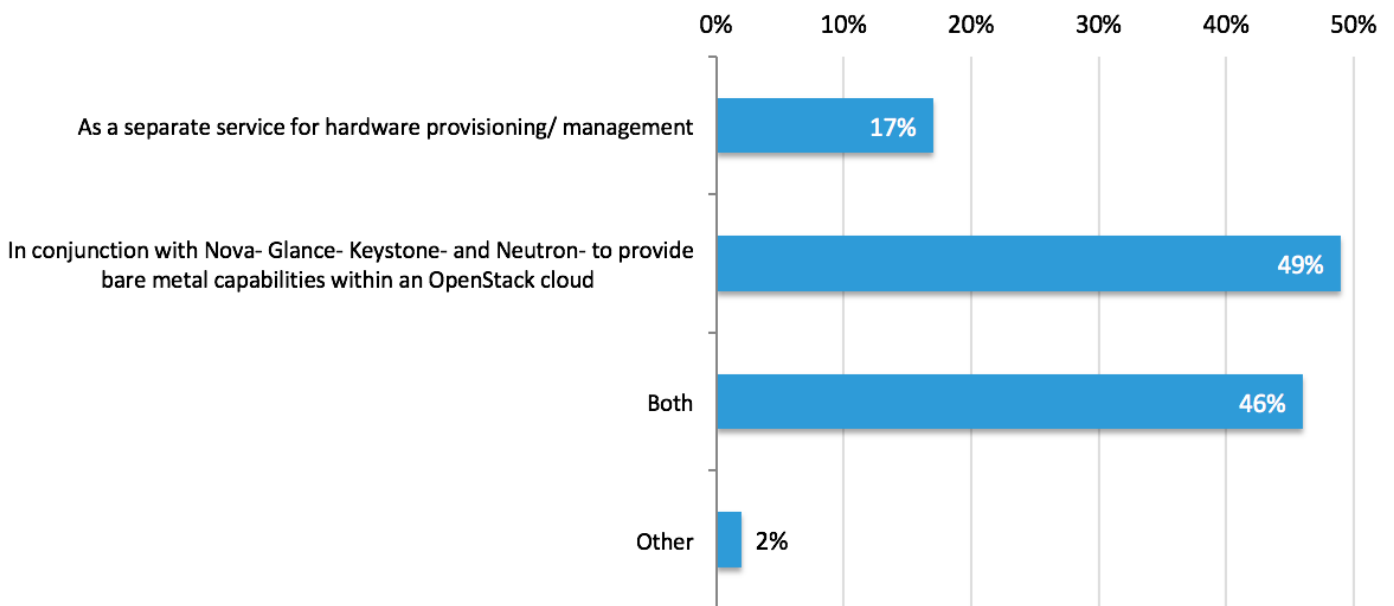
Most (78%) of the 200 respondents indicated they are not using this process, hinting that it should be better communicated to users.

How is the community supporting interoperability?

In the deployment section, users were asked if they were willing to test with RefStack. Of the total eligible for testing, nearly 40% volunteered to take part in this, a total of 89 community members.


How is ironic being used, or planned for use?

Out of 121 responses, nearly half indicated a combination of both options. Due to a survey design error, this question yielded greater than 100% total responses because “select all” was offered instead of “select one.”



What are the primary uses of Ceilometer data?

Among 154 deployments that indicated they are using Ceilometer and answered this question, more than two-thirds use its data for statistic reports as a resource metric. In addition, we evaluated the size of Ceilometer users and found no significant difference in size, based on cores.





The User Committee and OpenStack Foundation partnered with an external independent data scientist to help analyze and report the data. This cycle's survey streamlined some questions, added questions based on emerging technology, and clarified some questions for accuracy. As a result, not all results can be compared exactly to previous cycles' answers.

Additionally, this is the second survey with isolated data outputs, meaning that the first four survey's results were mixed, but since the April 2015 survey, each survey's data is isolated to allow for comparisons and, eventually, trend analysis.

In this report, we looked at answers submitted Sept. 14 – Oct. 5, 2015. We compared related questions to the data gathered between March 9, 2015 – April 16, 2015. This provides us with a smaller, but more representative sample of current state of OpenStack deployments.

As in the last survey, survey logic showed some users certain questions, based on their responses to prior questions. For example, if a deployment is using Swift, we asked additional questions about Swift. This helped us keep the survey as short as possible, and generally improves the quality of results by limiting the response size to just those who have a valid answer.

In addition to quantitative data, we gathered a substantial amount of qualitative data from open-ended questions. The User Committee and foundation staff selected representative comments to add more insights to numerical results. All comments are checked to ensure anonymity and relevant responses are forwarded to project technical leaders.

Changes to the sections About You/Your Organization/Your Thoughts since the Kilo cycle survey:

- Removed request for job title
- Revamped OpenStackInvolvement classification
- Changed survey logic so that sections relevant to respondent's classification would be asked (For example, if a respondent did not indicate that they are an application developer, we did not show them that that section)
- Removed the InformationSources question
- Improved the NetPromoter question to be more easily understood, and provided the opportunity for direct feedback related to this question
- Replaced the question regarding Container Technologies with a more general question regarding Emerging Technologies
- Revamped the general text box questions to reduce their number and improve their instruction
- Application Development
- Significantly reduced the number of questions
- Moved the question about Other Clouds to this section

Deployments changes:

- Deployment location question was simplified
- Workloads question was significantly revamped and simplified
- Added new questions regarding RefStack, PaaS Tools, Stacks, and questions for feedback on Neutron, Swift, Sahara, Nova, Ironic and Ceilometer
- Improved survey logic regarding projects: For example, if the respondent indicated they were running Ceilometer, we asked additional questions regarding Ceilometer

Errata

Some questions intended to be asked of users of a certain project were asked of all projects. Where respondents who did not indicate running a project responded to the project-specific question, we chose to exclude the this data from quantitative results.

