April 2017

OPENSTACK USER SURVEY

A snapshot of OpenStack users' perspectives and deployments

openstack.org

INTRODUCTION

OpenStack's ninth User Survey represents a snapshot of more than 1,400 completed surveys and nearly 600 deployments logged in the first two months of 2017—the largest sample of OpenStack deployments ever.

We believe results are representative of our community as a whole, revealing users' attitudes, technology choices, and uses of OpenStack software. All OpenStack Foundation members, more than 70,000 people across 160 countries, were invited to fill out the survey. Key findings:

MORE CLOUDS: OpenStack's growth is demonstrated by rapid development of new clouds, with 44% more deployments reported on this survey. The average age of a deployment is less than 1.7 years.

LARGER CLOUDS: 37% of clouds have 1,000 or more cores, up from 29% last year, and 16% of clouds are running more than 1 petabyte of object storage, up from 4% one year ago.

DIVERSE USERS: Users of every organizational size and industry find value in OpenStack, 32% have 10,000 employees or more, while 25% have fewer than 100 employees, and users hail from 78 countries.

MATURE TECHNOLOGY: OpenStack's maturity is demonstrated by two-thirds of deployments in production and 89% or greater adoption of key infrastructure services (Nova, Neutron, Keystone, Cinder and Glance).

DEEP ADOPTION: The median user runs 61–80% of their overall cloud infrastructure on OpenStack, while the typical large user (deployment with 1,000+ cores) reports running 81–100% of their total infrastructure on OpenStack.

CONTENTS

SURVEY CREDITS

The OpenStack User Committee includes Edgar Magaña, Jonathan Proulx, Shilla Saebi, Melvin Hillsman and Shamail Tahir. The User Survey working group also includes Arkady Kanevsky, Megan Rossetti, Tim Bell, Mark Smith, Kenny Johnston, Dave Van Everen, Rob Healy, Kostiantyn Volenbovskyi, Yanlu, and Treva N. Williams.

Heidi Joy Tretheway from the OpenStack Foundation staff leads creation of the survey and report, with support from Anne Bertucio, Jonathan Bryce, Mark Collier, Tom Fifield, Jimmy McArthur, Lauren Sell, and Wes Wilson.

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

EXECUTIVE SUMMARY

In the largest survey of OpenStack deployments ever, the OpenStack User Survey revealed widespread adoption, maturity and growth, with increased user interest in innovation and agility.

This User Survey included 44% more deployments and input from 22% more organizations than one year ago, which had been our largest survey to date. This survey's findings focus primarily on year-over-year comparisons and also answer some new questions.

More than 1,300 users, spanning 476 cities and 78 countries, took part in the ninth semiannual survey. 61% of these users and 74% of deployments are physically located outside of the United States.

The research shows that OpenStack has broad appeal and value for organizations of all sizes—32% of users have 10,000 employees or more, while 25% of organizations have fewer than 100 employees.

The median OpenStack user runs 61% to 80% of their infrastructure on OpenStack; among users with clouds of 1,000 cores or more, the median user runs 81–100% of their overall infrastructure on OpenStack.

The business drivers for OpenStack adoption continue to highlight its competitive advantages, including accelerating the organization's ability to innovate and avoiding vendor-lock in. Users say these business drivers are even more important than saving money and increasing operational efficiency, which ranked No. 1 and No. 2 in past surveys.

> *The median OpenStack user runs 61% to 80% of their overall cloud infrastructure on OpenStack.*

OpenStack's measurement of user satisfaction, the Net Promoter Score (NPS), fell somewhat, though we found users with the newest clouds to be the most satisfied users, with a rating nearly double the software industry average.

Containers remain the top emerging technology of interest to users. Among those currently deploying container orchestration or platform services on OpenStack, 47% are using Kubernetes.

The share of OpenStack deployments in production edged up slightly, with two-thirds of clouds in production. The large proportion of clouds in production demonstrates the maturity of OpenStack, while an influx of clouds in proof-of-concept and test stages forecasts healthy growth for the future.

This is also demonstrated by the average age of a deployment—just 1.68 years. 56% of deployments catalogued were launched in 2016 or 2017.

The typical deployment runs nine OpenStack services, with 16% running 12 or more services. All of OpenStack's key infrastructure services are in use by more than 89% of clouds. Heat, Telemetry, Swift,

Rally, Kolla and Barbican also showed significant increases in adoption.

Among the projects of greatest interest to users, and likely to see strong future adoption, Designate, Magnum, Trove and Manila top the list, with Kolla and Barbican interest also up significantly.

The typical size of an OpenStack cloud increased; 37% of clouds have 1,000 or more cores, compared to 29% a year ago, and 3% of clouds have more than 100,000 cores. The number of users running Nova with the scalable "cells" architecture also increased 218%.

Swift object provisioning also saw greater scale. 16% of deployments provisioned more than one petabyte of object storage, compared to 4% last year, and 33% reported storing 100,000 or more objects, vs. 13% last year.

Analyze the results yourself

The community is invited to explore the User Survey data and answer their own questions via a live survey analysis tool. It offers an online dashboard, six global filter categories, and three data sets (2015, 2016 and 2017), and is available at http://www.openstack.org/analytics.

This dashboard captures live data as survey responses are logged. This survey report is based on data collected between Jan. 26, 2017 and March 1, 2017, but the online dashboard includes surveys answered throughout 2017. As a result, online results sometimes vary from this report due to different collection periods.

Although this survey report might reveal different answers compared to the dashboard, we believe both are valuable resources for the community. Please also note that some of the charts presented in this report do not equal 100% because multiple responses were accepted or due to rounding.

Part One

DEMOGRAPHICS

User Survey includes 44% more deployments than ever before

The ninth OpenStack User Survey represents a snapshot of 583 deployments logged between Jan. 26 and March 1, 2017—a 44% increase in deployments compared to the survey conducted one year ago. Additionally, 754 unique organizations contributed to the data in this report, a 22% increase over the April 2016 survey.

The OpenStack User Survey is the largest survey of its kind, spanning more than 2,500 community members' input over the past year. 65% of survey respondents took the survey only once over the past three rounds of surveys (April 2016, October 2016 and April 2017) and 32% of those surveyed in April 2017 also responded to the survey in April 2016.

Consistency from survey to survey suggests that results are representative of our community as a whole. Our community is much broader than the sample included in this survey, so significance testing is done on small variations in the data

2,561 *users surveyed in the last year*

1,533 *in April 2016*

388 *in October 2016**

1,305 *in April 2017*

to determine whether we are seeing a trend or simply an insignificant variance due to different respondents.

One hundred users logged more than one deployment (as many as five deployments from a single respondent). The number of users logging more than one deployment increased 14% over one year ago.

*Only 8% of those surveyed in this cycle also took both the April 2016 and October 2016 surveys. The October survey was a much smaller sample size because it was designed as a "deployments update" only, targeting those with existing deployments.

Figure 1.1

	April 2015	Oct. 2015	April 2016	Oct. 2016*	April 2017
Completed Surveys	1100	1315	1603	388	1406
Deployments	286	267	405	260	583
Multiple Deployments	74	85	88	44	100
Most Deployments from a Single Respondent	5	16	11	11	5

What are the roles of survey respondents?

We ask survey participants about their role in the OpenStack community, and 32% of people indicated they have more than one role.

Cloud architects are more represented in this survey, up 5 points over last year, while cloud operators and sys-admins rose 4 points and CIO/IT infrastructure managers rose 2 points. Representation of upstream developers declined 7 points and app developer-deployers fell by 5 points.

Among the 14% of individuals who noted their involvement in OpenStack in the "other" category, quality assurance, research, security, various developer roles, community manager, project manager, product manager, business development, consultant, and marketing were mentioned.

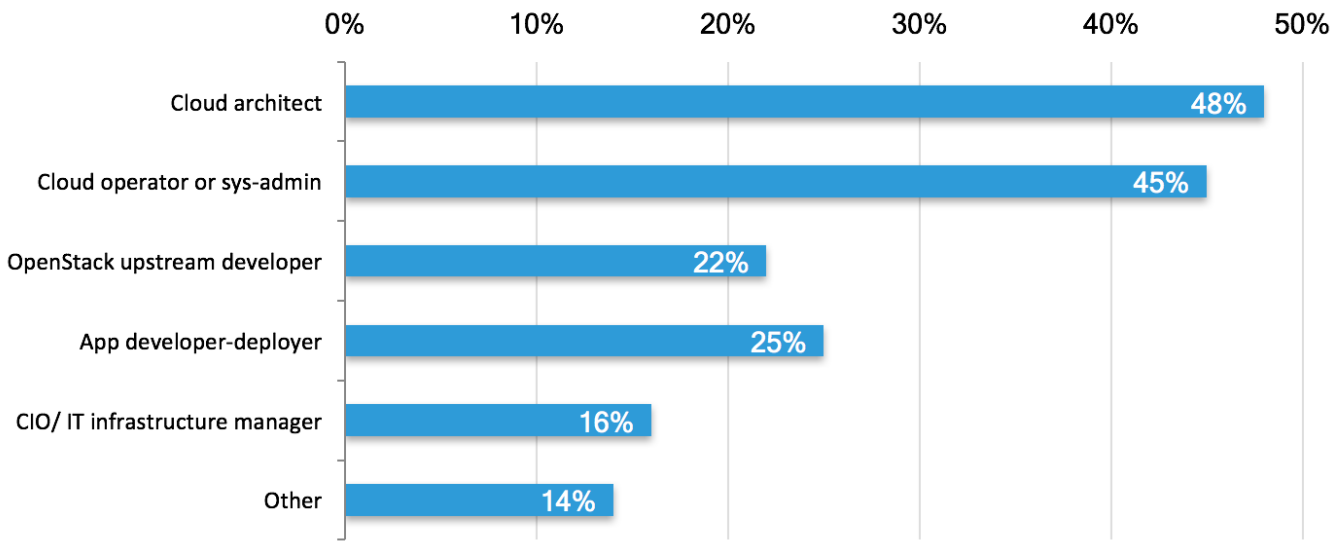

Figure 1.2 n=1275

Which industries use OpenStack?

Top industries represented on the User Survey include information technology, which represents 56% of total respondents; telecommunications at 16%; academic/research at 11%; and finance, retail/e-commerce, manufacturing/industrial, and government/defense, all at 2%.

Among the 9% of organizations in the "other" category, consulting, real estate, web hosting, software and hardware, news, entertainment and advertising were mentioned.

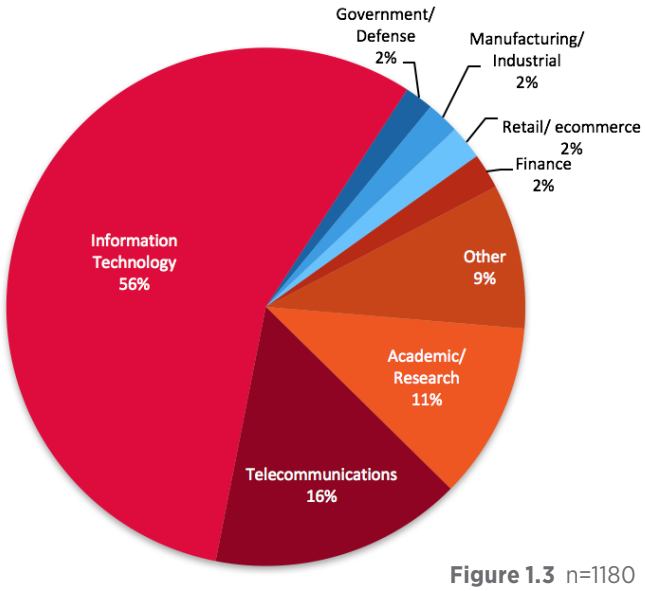

Figure 1.3 n=1180

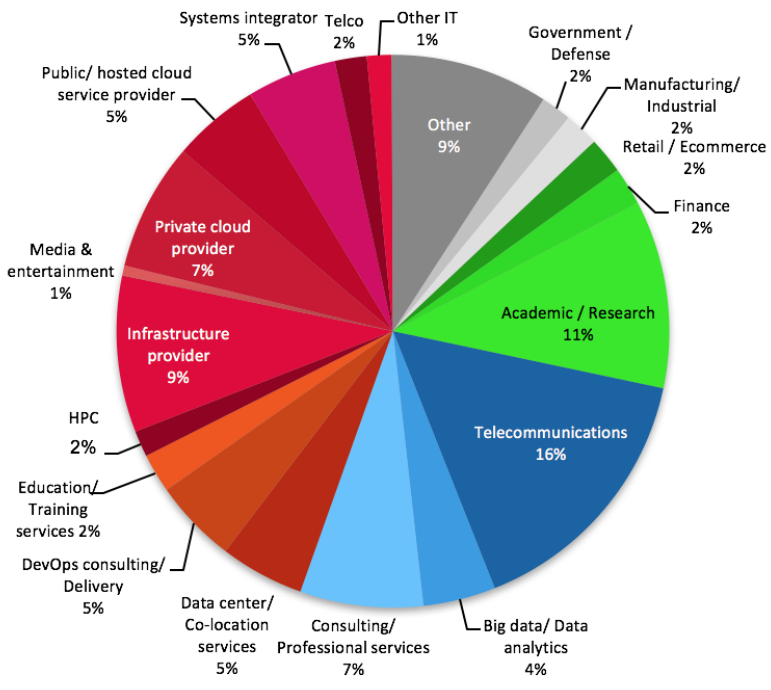

Figure 1.4 n=1180

Looking closer at the IT sector, infrastructure providers made up 16% of the total IT companies, followed by private cloud providers and consulting/professional services, each with 13%. Systems integrators, public/hosted cloud-service providers, DevOps consulting/delivery, and data centers/co-location services rounded out the most prevalent IT categories, each with 9%.

Among the segment of IT users listing their business category as "other," cognitive computing, government contractor, SaaS, SDN, software vendor, support, operating system, and ERP (enterprise resource planning) were noted. This wide variety of use cases and industries validates OpenStack's broad appeal, acceptance and adaptability.

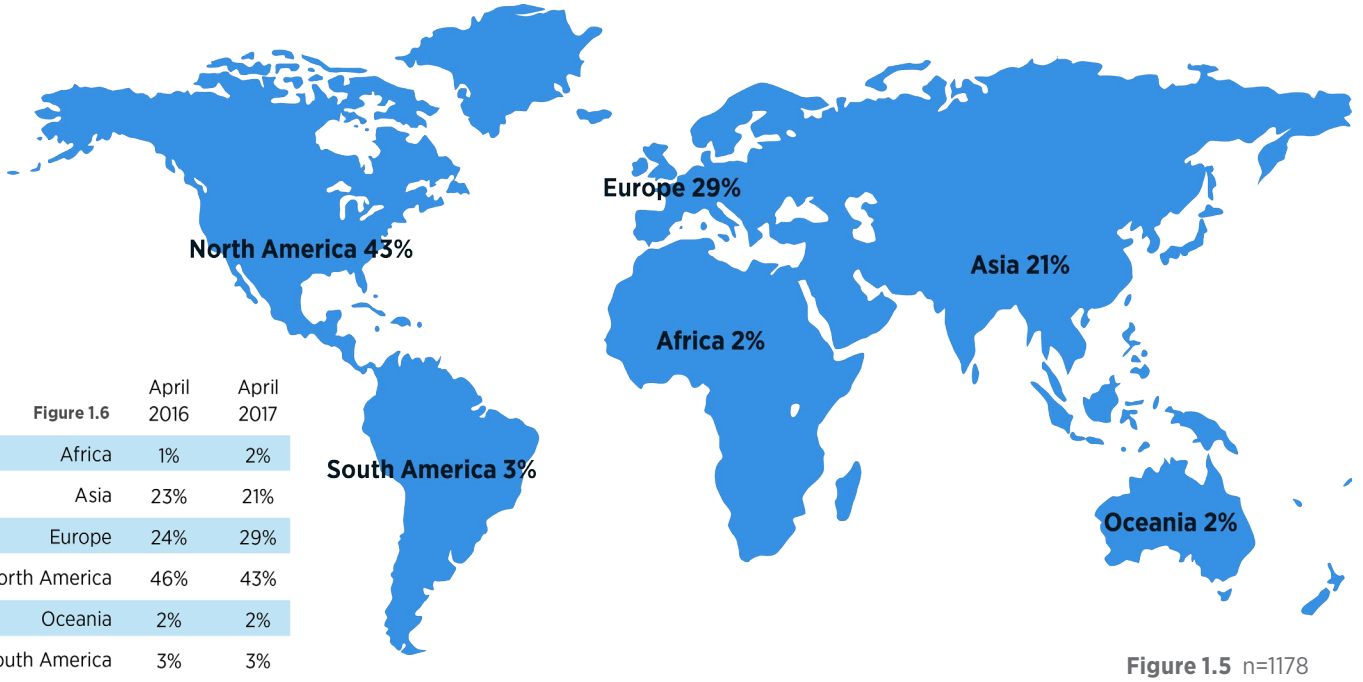

Figure 1.6	April 2016	April 2017
Africa	1%	2%
Asia	23%	21%
Europe	24%	29%
North America	46%	43%
Oceania	2%	2%
South America	3%	3%

Figure 1.5 n=1178

Where in the world are OpenStack users?

This survey showed a significant increase in respondents from Europe, compared to the April 2016 survey, while North America was slightly less represented.

Users are located across 476 cities, with the top locations including:

- Silicon Valley (California, USA) 9%
- San Antonio (Texas, USA) 4%
- Tokyo (Japan) 3%
- Bangalore (India) 2%
- Beijing (China) 2%
- London (UK) 2%
- Raleigh (North Carolina, USA) 2%
- Armonk (New York, USA) 1%
- New York City (New York, USA) 1%

Users span a total of 78 countries, with 61% of respondents outside of the United States, making this our most internationally representative survey to date. Highest concentrations were in:

- United States: 39%
- China: 6%
- Germany: 5%
- India: 5%
- Japan: 4%
- United Kingdom: 4%
- Canada: 3%
- France: 3%
- Australia: 2%
- Brazil: 2%
- Netherlands: 2%
- Spain: 2%

What size organizations use OpenStack?

OpenStack is used by organizations of every size, with one-quarter of users from organizations of fewer than 100 employees, and 32% of organizations with more than 10,000 employees.

This data is consistent with prior reports, demonstrating that OpenStack is used by major enterprises (refer to www.openstack.org/users for case studies and presentations from some of the world's largest brands), mid-size organizations and small companies.

It also demonstrates OpenStack's broad appeal and value for both large and small organizations; indeed, mid-size organizations make up the largest share of users.

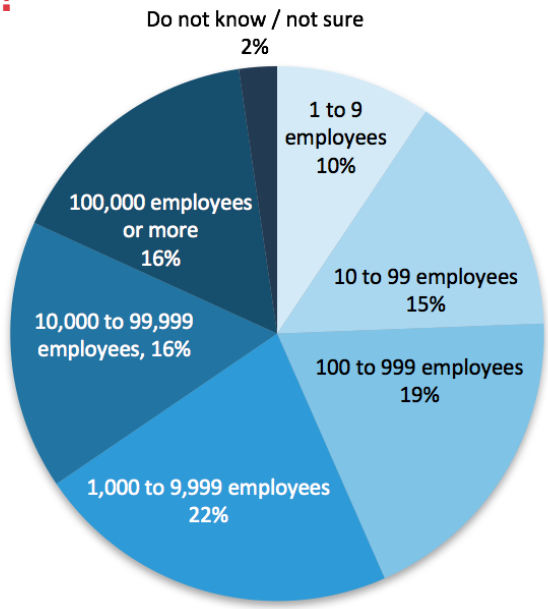

Figure 1.7 n=1176

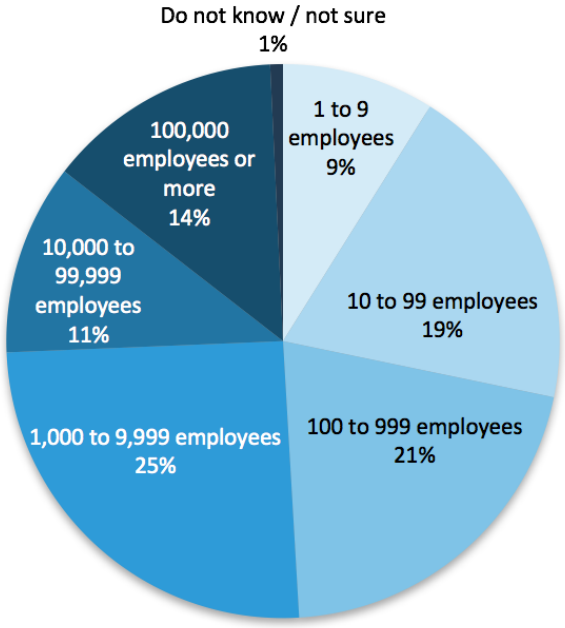

Figure 1.8 n=269

When we analyzed this data through the filter of deployments only, which reduced our sample size from 1,176 people answering the question to 269 unique organizations, the proportion of mid-size organizations using OpenStack increases to nearly half of the pie, while the largest and smallest organizations each take about a quarter.

Additional analysis of trends of all users answering the User Survey over the past two years shows larger organizations increasingly represented (the three largest organization sizes increased 5 to 6 points each), while the smallest organizations' representation declined, down 12 points for organizations of 1–9 employees and down four points for organizations with 10–99 employees. This change could be related to a smaller proportion of upstream developers filling out the survey.

Part Two

USER PERSPECTIVES

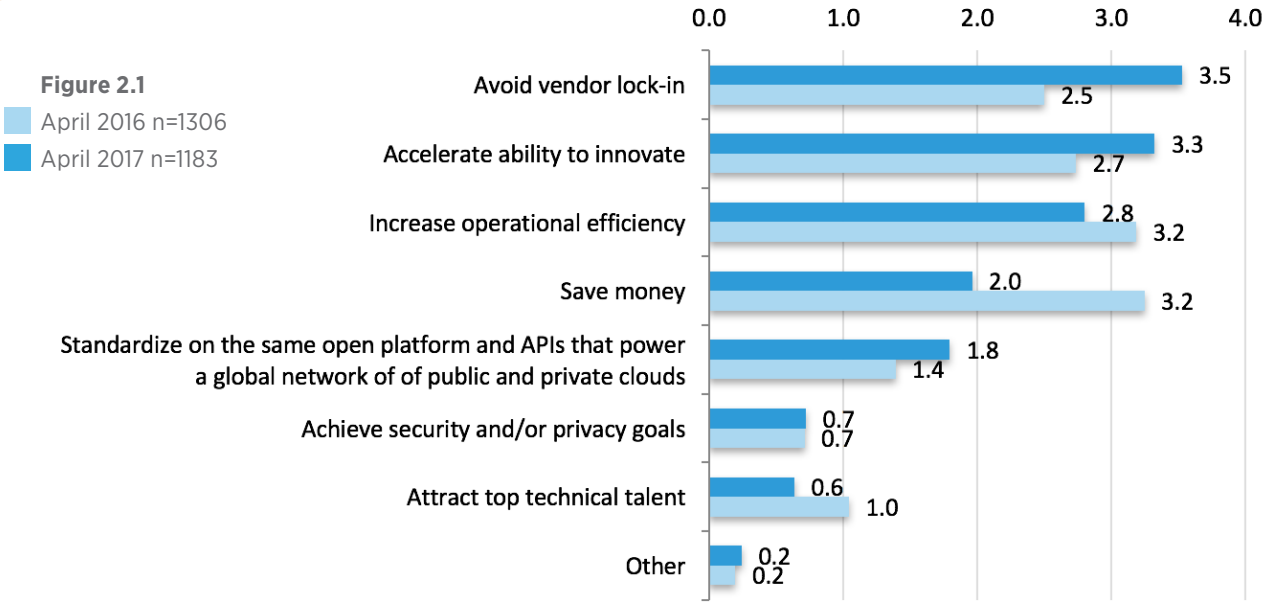

Figure 2.1
April 2016 n=1306
April 2017 n=1183

Why do organizations choose OpenStack?

The business drivers for choosing OpenStack are evolving, with the ability to innovate and avoid vendor lock-in becoming more important to users.

Users were asked to name their top five business drivers and rank these from 1 to 5. To best express this data, each No. 1 rank was assigned a weight of 5 points; each No. 2 rank assigned 4 points, and so on. Point totals were then calculated for both this year's survey and last year's to show changes in business sentiment over time.

We note that of the top three responses one year ago, increasing operational efficiency and accelerating innovation are still among the top three. And although the top five business drivers remain constant, their ranking order has changed.

While "standardizing on the same open platform and APIs that power a global network of public and private clouds" was a top five business driver for 97% of survey-takers one year ago, it was not their No. 1 reason for choosing OpenStack. Its relatively low score here reflects the majority of users ranking it as their fifth reason.

We also looked at business drivers in the context of cloud type, segmenting answers between private and non-private clouds. The variance was not a significant difference, typically just 1 point.

Other business drivers mentioned included research and learning, testing and training, the value of open source and community support, increased stability, NFV, and reducing network complexity.

How likely are users to recommend OpenStack?

OpenStack has long used the Net Promoter Score framework to gauge community sentiment and satisfaction, as well as to identify areas where we can concentrate efforts and resources to significantly improve.

We asked users, "How likely are you to recommend OpenStack to a friend or colleague?" Answers are on a 10-point scale with 10 being the highest. Scores are broken down into three parts: "promoters" answer a 9 or 10, indicating they are very highly satisfied and likely to promote OpenStack; "passives" answer a 7 or 8, which is considered neutral; and "detractors" answer 0-6. A Net Promoter Score, or NPS, is calculated by taking the percentage of promoters and subtracting the percentage of detractors.

Survey responses include more than twice as many promoters as detractors, with detractors being

	April 2015	Oct. 2015	April 2016	Oct. 2016	April 2017
Responses	517	894	1183	278	1044
Promoters	47%	47%	47%	51%	44%
Passives	26%	31%	35%	34%	36%
Detractors	27%	22%	18%	15%	20%
NPS	20	25	29	36	24

Figure 2.2

20% of those surveyed overall. When looking only at those with deployments, nearly half (48%) are promoters, scoring OpenStack a 9 or 10, while the number of detractors is 17%. We also note that the October 2016 survey tends to be an outlier on virtually all questions in this survey due to its much smaller sample size.

How do NPS scores compare to other survey data?

Trends over time showed increased satisfaction until this survey, where the number of promoters fell off by 3 points year-over-year, while detractors increased by 2 points.

We looked at the NPS score data in a variety of ways to determine whether certain segments of users were more or less satisfied. Industries including private cloud providers and telcos rated OpenStack most highly, with an NPS of 39, followed by public/hosted cloud service providers and DevOps consulting/delivery services, both at 37. Those in the "other" industry category rated OpenStack lowest.

Deployments running Mitaka, Newton and trunk releases had an NPS score of 28, compared to all other deployments at 36.

We also looked at deployments by age of the cloud, finding that deployments created in 2016 and 2017 were most satisfied with OpenStack, with NPS scores at 38 and 35, respectively, compared to the oldest deployments (clouds launched between 2010 and 2014), with an NPS of 14. This comparison also yields an interesting insight—deployments made in 2016/2017 are not necessarily using the latest releases (Mitaka and Newton).

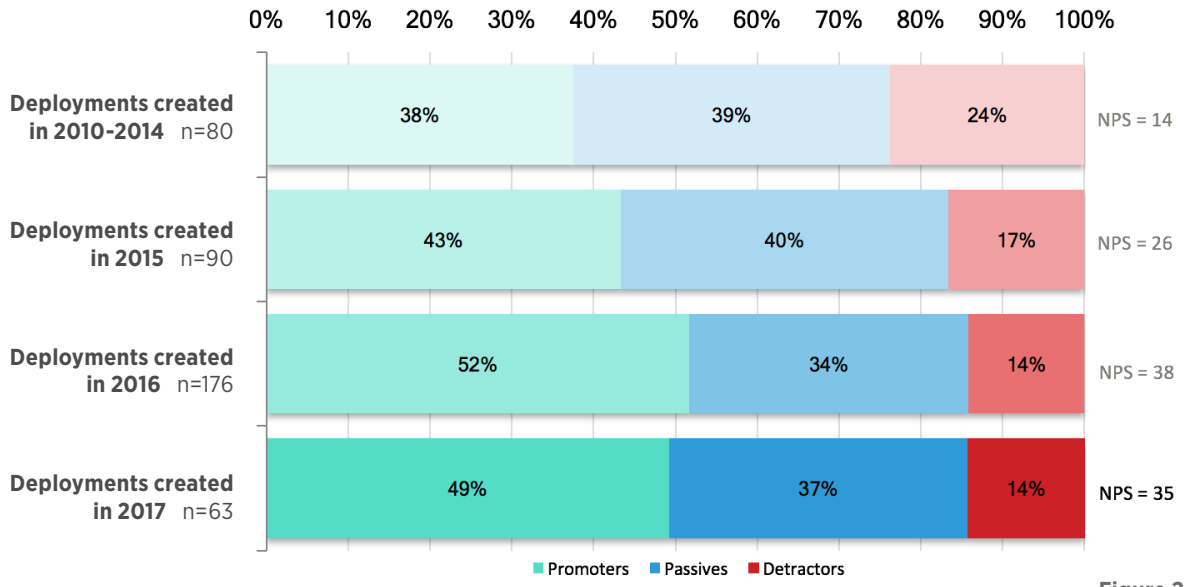

Figure 2.3

Why do users recommend OpenStack?

As a follow-up to the Net Promoter Score, we asked users why they would recommend OpenStack. Of the 682 comments received, a word-frequency analysis showed that platform, community, support, ease of use [easy], deployment, complexity, infrastructure and features were top of mind.

The prevailing sentiment was that OpenStack achieves business and technical goals with flexibility, scale, and agility, while complexity in deployment and operation were most often cited as pain points.

"I appreciate the value of the OpenStack project and its community of developers especially for having designed and engineered a very complex integration of a number of open-source technologies into a single, powerful cloud computing platform," one wrote.

"I believe since the Mitaka release, OpenStack has been a force to be reckoned with—a powerful array of cloud features businesses have come to expect," another said. "With ever stronger global support, I anticipate OpenStack to continue growing in capability and innovation for years to come."

A developer noted, "You can contribute easily if you want a new feature instead of waiting for the vendor to implement it," a comment echoed by many users.

Another noted, "If you're going private cloud, OpenStack is the only viable option that does not lead (sooner or later) to bankruptcy. And, with so many vendors eager to deliver, you stand the best chance to mix and match the best solutions targeting the various functions while still maintaining standard and open interfaces and APIs between components."

> *OpenStack [is] a force to be reckoned with—a powerful array of cloud features businesses have come to expect.*

BUSINESS OBJECTIVES

"OpenStack fully achieved our objectives while having a very short time to install (1-5 weeks), compared to the previous cloud solution," a user commented.

Another said, "OpenStack has had tangible positive effects on the accomplishment of our core mission set. Whereas other technology-centric educational institutions within the [organization] are stuck with legacy—static solutions—we have migrated our entire education platform into an everything-as-code paradigm that is underpinned by OpenStack-Ceph and SaltStack."

The need for cultural change was noted, in that OpenStack "requires a fundamental shift in how the company perceives servers. OpenStack is not a turnkey implementation strategy. Proper assessments need to be done—[before] and after the OpenStack implementation. OpenStack showcases a lot of great principles that, once understood and mastered, bring tremendous value.

Why users recommend, continued

However, it is not something that happens overnight—it is a journey."

The modular nature of OpenStack was appreciated by many users. "Install what you need. Saves a lot of cost. Tons of documentation available to do it yourself."

Users also said they did not want to be beholden to a vendor for cost reasons or the feature roadmap. "Avoid vendor lock-in for freedom of innovations and service extendability," one noted.

Another said, "It's been one of our most important infrastructure improvements, which will lead us to increase our productivity."

CAPABILITIES

"No other platform provides such a diverse set of X-aaS capabilities," noted several users. "OpenStack is a full-featured solution for cloud computing. Documentation is easy to find, leading us to speed up our deployment. We can modify the code and make the solution suitable for our own use."

Users said they enjoyed OpenStack's agility, such as "being able to reduce latency in building environments for users, and developers to execute faster while using existing infrastructure, and also avoid the old ways and methods of going to different departments in executing specific operations that are needed to build an environment. It is about responding faster in minutes and not weeks, while ensuring security and privacy."

A user noted that OpenStack's core services "accomplish at least 80% of our actual needs. And once you learn the quirks, it's rather simple to use and maintain."

"OpenStack enables us to scale out infrastructure management at a fraction of the cost of traditional infrastructure with the security of on-premises asset and data control," a user said.

Another wrote, "OpenStack provides the right level of infrastructure abstraction, automation, and programmability for most of the deployments I work with, regardless of system scale."

"I can integrate any kind of vendor solution and [it is] interoperable with all services. Also, it allows small companies to provide their own solutions to the market."

However, the speed with which OpenStack evolves can be a drawback for some. "OpenStack's development cycle is faster than our internal

Why users recommend, continued

development cycle, which causes a great deal of false starts as the company management re-targets on a new release, causing departments to re-start their efforts and play catch-up while projects move forward with stale releases."

COMMUNITY

OpenStack's community is among the most frequently mentioned positives. "OpenStack has the benefit of thousands of developers all over the world working in tandem to develop the strongest, most robust, and most secure product that they can. Anyone who chooses to can access the source code, make any changes or modifications they need, and freely share these changes back out to the community at large."

"I see that the OpenStack community is very open to new ideas [and has] a strong base of technically strong people," wrote an upstream developer, "resulting in great features implemented well, resulting in a great user experience."

Another upstream developer said, "I've had great activity/communications with the contributors from all over the world [who] are very kind, polite, and honest."

An operator of a mega-scale environment added, "I'm attracted to open standards and avoid vendor lock-in at all costs. One of the biggest selling points for OpenStack is the vibrant community driving such a rich SOI ecosystem."

NFV

Many users said they recommend OpenStack for use in Network Functions Virtualization (NFV). "The platform has achieved a level of maturity in the last couple of releases that makes it a true viable alternative—especially in the NFV space—for cloud infrastructure and management."

Another user praised the contributor community. "This project is enabling acceleration in the development of key technologies such as SDN and NFV. Latin America needs to urgently adopt the new paradigm that has introduced OpenStack to the world."

Why users recommend, continued

"I am a big fan of opensource technologies and I believe OpenStack will be at the forefront in Network Function Virtualization deployment. This is meant to disrupt the usual hardware-centric deployment model, which is not scalable as far as I am concerned. OpenStack will change the ICT/ Telecom landscape."

COMPLEXITY

Comments were mixed regarding OpenStack's flexibility. "Openstack is a great program, but sometimes the complexity of the projects and deployment can be a major problem pushing people into either a distribution and support contract or a SaaS (software-as-a-service)-type offering. Not that those are bad, but would like to see more effort put toward Kolla and Stackenetes-type deployments that lessen the complexity, and ease access to other projects," a user noted.

"OpenStack is a powerful integration engine for creating a cloud platform, however the deployment and operational complexity, as well as poor documentation, are the reason of not getting a score of 10," another said.

On the other hand, many users told us that OpenStack is "very efficient software; easy to deploy, easy to use, and easy to operate." Several praised "the clear documentation and standardization of APIs."

FLEXIBILITY

Whether looking for freedom in technology or vendor choices, users told us flexibility is key. "No other open-source cloud platform even comes close to the feature set and flexibility of OpenStack. Most commercial ones don't either."

> *"OpenStack has ... thousands of developers all over the world working in tandem to develop the strongest, most robust, and most secure product that they can."*

Another user added, "OpenStack as a platform is powerful and flexible. Modern deployment and lifecycle management tools have removed some of the trickier elements of booting up and operating an OpenStack cloud, making it a no-brainer to recommend it for everyone."

EXPERTISE

The need for more talented OpenStack engineers was noted by several users. "Deploying, maintaining and upgrading a production OpenStack-based cloud is no small feat. Luckily the technology itself has vastly improved over the last few years. But finding experienced people is still quite challenging. Without the people—no cloud."

Another user added, "I believe OpenStack is changing the cloud ecosystem the same way Linux changed the OS ecosystem 15–20 years ago. And we are in the 'cloud era,' so I expect high demand for OpenStack expertise."

Besides “free” and “open,” what do users like best about OpenStack?

When we asked users to tell us what they like best about OpenStack—besides the fact that it is free and open—users weren’t shy about reminding us that these two elements are still compelling.

“As if ‘free’ and ‘open’ were not of the most important?” one user said. “How else could I have afforded to build my own cloud installation?”

Word-frequency analysis of the 602 comments showed nearly three times as many people mentioning “community” than any other word (148 instances), as well as many mentions of openness, flexibility, support, standardization and APIs, innovation, stability, features and platform, scalability, documentation, and the ecosystem of contributors.

Here are more excerpts from what users said:

TECHNOLOGY

- It is an integration platform where all of your infrastructure technologies can come together. This increases efficiency to operate and provide infrastructure services. In addition, its strong ecosystem is the best.
- OpenStack is built upon proven technology with scalable underpinnings.

> *It is an integration platform where all of your infrastructure technologies can come together. This increases efficiency to operate and provide infrastructure services.*

- It addresses a very real need. At this moment I have tens of thousands of machines managed by a single OpenStack bare metal cluster. I also have tens of thousands of VMs, and over a thousand hypervisors in a single OpenStack VM cluster. That’s pretty awesome!
- The flexibility and dynamic nature to the platform. As containers become more involved, [there are changes] in how not just workloads, but OpenStack itself, is deployed.
- The stability of the core system. There are issues—even in Newton—but in general it’s really stable.

What users like best, continued

- HTTP APIs are clean enough to be used from whatever clients I choose. I don't have to use special client code.
- [What] I like most is the new projects that 'join' OpenStack, like CEPH in the past, and now we have initiatives with Karbor and others. There is a constant change that is driven by the community that later can help the product become more robust.
- It works! Like any software it has its blindspots, but for most part it works. It needs to get more battle tested for larger scale.

COMMUNITY

- Global collaboration system—everyone benefits from all the achievements of the others.
- Ability to listen to everybody and find common understanding on how certain features can and should be supported.
- Global community and transparency equals the ability for everyone to make a difference.
- I strongly value the community, in particular individuals I have met via OpenStack. Also I appreciate the foundation's commitment to best practices around the software such as open governance, inclusiveness, and being global.

VENDORS & USER CONTROL

- As a customer, I control how I consume OpenStack—DIY, distro, on-premises managed service, off-premises managed service, or public cloud. This is a great thing because it puts the user in control of the consumption pattern.
- The ability to control what is present on my infrastructure without all those crazy dependencies. I can build my cloud exactly how I want due to the dynamic nature of OpenStack.
- Open source avoids vendor lock-in and supports common shared API access to resources.

Which areas of OpenStack require further enhancement?

Word-frequency analysis in this set of 617 responses included 50 or more mentions of deployment, networking and documentation, 37 mentions of Neutron and 31 of Nova. Other hot topics included releases, features, security, management, containers, the project Keystone, and APIs.

These were some key comments, grouped by themes:

FEEDBACK LOOP

- Be more user- and ops-oriented. It's been an issue since the beginning of OpenStack and it still is, even though we've seen huge improvements.
- Better communication from customers to the development teams on what problems they need to be solved and how this can happen.
- Delivery dates and when features will be available are super important.
- Better feature request process, and some way to keep the project teams accountable. I have heard the response, "I don't get paid to work on your feature requests," more than once, which is frustrating because while I understand that is true, it's not benefiting the larger community.

INSTALLATION

- What we have today is varied installation/deployment models.
- Need the installation to become easier—the architecture is still too complex right now.
- Installation, particularly around TripleO and HA deployments, is very complicated.
- A common deployment and lifecycle management tool/framework would really make things easier. Having every distribution use its own tools (Triple-O- Fuel- Crowbar- ...) really doesn't help. And yes, I know that this is not OpenStack's fault but if the community unites behind one tool (or maybe two), we could put some pressure to the vendors.

> *A common deployment and lifecycle management tool/framework would really make things easier. Having every distribution use its own tools really doesn't help.*

- Automate installation. Require consistent installation between projects.
- Standardized automated deployment methods to minimize the risk of splitting the developments in vendor-specific branches.
- Deployment is still a nightmare of complexity and riddled with failure unless you are covered in scars from previous deployments.
- Initial build up needs to be much easier, such as using a simple scripted installer that analyzes the hardware and then can build a working OpenStack. When upgrades become available, it can do a rolling upgrade with 0 down time.

UPGRADES

- The lifecycle could use a lot of attention. Most large customers move slowly and thus are running older versions, which are EOL upstream

Areas for enhancement, continued

sometimes before they even deploy them. Doing in-place upgrades is risky business with just a one or two release jumps, so the prospect of trying to jump 4 or 5 releases to get to a current, non-EOL version is daunting and generally results in either a lot of outage or simply greenfielding new releases and letting the old die on the vine. This causes significant operational overhead as getting tenants to move to a new deploy entirely is a big ask and you end up operating multiple versions.

- Many organizations appear to be moving toward containerizing their OpenStack control plane. Continued work on multi-version interoperability would allow organizations to upgrade a lot more seamlessly and rapidly by deploying newer-versioned containers in parallel with their existing older-versioned containers. And it may have a profoundly positive effect on the upgrade and lifecycle for larger deployments.
- The biggest challenge is to upgrade the production system since there are a lot of dependencies and bugs that we are facing.
- Releases need more feature and bugfix backporting.
- Stop coming out with all of these releases. Only do a release once every two years.

BIG TENT & CORE SERVICES

- Focus on core (i.e fundamental building-block) components.
- Increase emphasis on standards so that there's better consistency across projects, from API interaction through to configuration.
- Stop including every related technology as projects under the big tent; focus on the essential core projects to bring them to excellence.
- Hunting use cases for newborn projects is absolutely disastrous to the overall foundation and goals of OpenStack. Simplification of end users' workflows is all that matters: increasing function while increasing system complexity as little as possible.
- With the advent of the big tent, OpenStack has become quite diffuse. It is hard to tell what is important and there's a sense that development is unfocused and lacks velocity.

DOCUMENTATION

- The official OpenStack docs should better support older releases (not just latest release and maybe the prior release).

Areas for enhancement, continued

- Ask.openstack.org is in dire need of some TLC; it perpetuates the impression OpenStack is hard. It's a big missed opportunity to improve crowd-sourced documentation.
- Ask Openstack should be more structured and should have a better filter to identify if the question is no longer duplicated.
- [We need] better developer documentation in the projects. General dev documentation is fine, but the projects sometimes have next to nothing or seriously outdated documentation about elementary topics.
- Document use cases, best practices, and give architecture advice for new organizations that want to embrace OpenStack.
- Version management of documentation—it's not easy to find information (e.g. operation manuals, APIs) about a specific or latest version of OpenStack.
- Documentation of installation steps. After talking to several people in the OpenStack community at the summits I attended, I realized there was no way—even with an advanced skill set in IT—that I was going to get this to fly on my own.

CROSS-PROJECT

- As with all large projects, everything needs enhancements, but in general I would say the lack of cross-project alignment on feature sets is very problematic. Quite often, new features are introduced in core components (especially Neutron and Keystone) but dependent projects rarely implement or support them properly, which is very frustrating for the user.

> *Projects need to be made to interoperate with each other. An operator should have to know how to deploy 'OpenStack'—not 'Keystone, Glance, Nova, Neutron, etc., etc., etc.,' and each with their own caveats.*

- Better integration of some services, working more together, trying to avoid the "not invented here" syndrome even for other OpenStack projects.
- Projects need to be made to interoperate better with each other. An operator should have to know how to deploy "OpenStack"—not "Keystone, Glance, Nova, Neutron, etc., etc., etc." and each with their own oddities and caveats. Things are still changing too much at this stage in the lifecycle, and it is a massive task for operators to figure out and stay on top of what has changed/broken/deprecated across all the various projects in each release, assuming teams have even documented these things properly.

COST

AWS price reductions have undercut the cost case for our OpenStack deployment in a lot of cases, so my company management directed us to move away from OpenStack except in certain very specific

Areas for enhancement, continued

disk-IO-intense workloads. I have been directed to discontinue my upstream development and community involvement. I am continuing to do upstream development, but only during my lunches, nights and weekends.

STABILITY & SCALABILITY

- Truly large-scale clouds are not well-documented or supported. We're migrating a large number of legacy environment to OpenStack, but often run into scaling issues when attempting to scale beyond a few hundred nodes.
- Stability is my biggest issue. Our deployment switched to Newton in October and API calls have been failing more often since then.
- OpenStack today is much more stable than a few cycles back, and the documentation is more clear. Stability and maturity should always be the primary goal of any project in the community for every cycle.

CONTAINERS

- [We need] a clear and strong container support strategy—there is currently too much overlap with competing initiatives/projects.
- Clarify container support for bare metal, Neutron, SDN, Cinder and Manila.
- Docker project should have a strong OpenStack support, comparable to AWS.

HORIZON

Several users mentioned a desire for a widget-based dashboard, better web front end, and a web UI console integration of different services.

- Integration of some advanced functionality on the Horizon dashboard, such as auto-scaling groups.
- We should have a "Horizon for admins" where more things could be done, since everything with command line is nice but is very difficult when you want to delegate to other people.
- UIs aren't up to commercial offerings such as AWS or vCenter. There are a lot of unanswered questions during installation, particularly with regard to networking. This is largely due to the options available. Prescriptive setups are hard since one size doesn't fit all.

NEUTRON

- Software-defined networking could have better integration. There are many solutions from vendors but their level of integration/automation is far from perfect.
- Neutron needs to be reworked and simpler—we don't need to include every use case under the sun. Kick out the obscure architectures. Quit presenting about containers at each Summit.

Areas for enhancement, continued

LOGS

- Error handling and messages could do with improvement. At the moment when we hit an error in our production OpenStack, we need to trawl through the logs to find the strack trace and then trawl through the source code to determine what caused that crash and fix it, either by cherry picking from upstream or finding a way to configure around it.
- We see great improvement from release to release. Our main is concern is the general cluster stability and the weak reporting (logs).
- While I haven't run the most recent versions, I still think log interpretation may be an issue when troubleshooting, as it's difficult to know what (logged info) is important.

How often do users refer to documentation from docs.openstack.org?

There were no significant changes in the frequency that users indicated referring to OpenStack's documentation. See user comments for additional detail, both praise and critiques of documentation.

DOCUMENTATION
an OpenStack Community Project

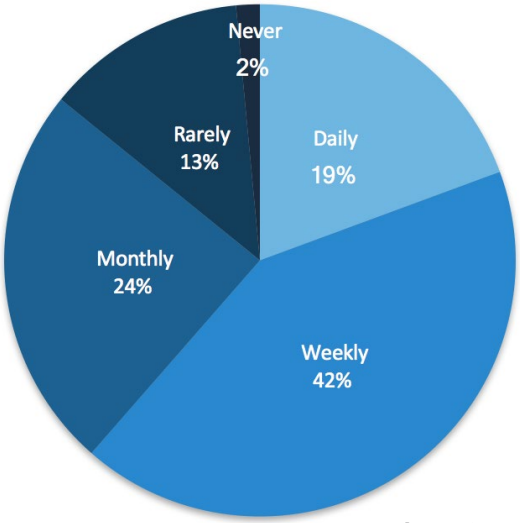

Figure 2.4 n=996

Which emerging technologies interest OpenStack users?

For the fourth cycle in a row, container technology leads the list of new/emerging technologies of greatest interest to OpenStack users. Containers interest was 5 points higher than a year ago, and interest in hardware accelerators was up 6 points.

Those with deployments were more likely to note that SDN/NFV, bare metal, and hardware accelerators were technologies of interest. Those who did not log deployments were particularly interested in hybrid cloud and IoT. There was no significant difference in containers interest between the two groups.

Other technology of note included block chain, cross-region compute and networking, Lambda-like functions (Picasso is a developing project for functions as a service), machine learning and artificial intelligence as a service, and MANO (NFV management and orchestration).

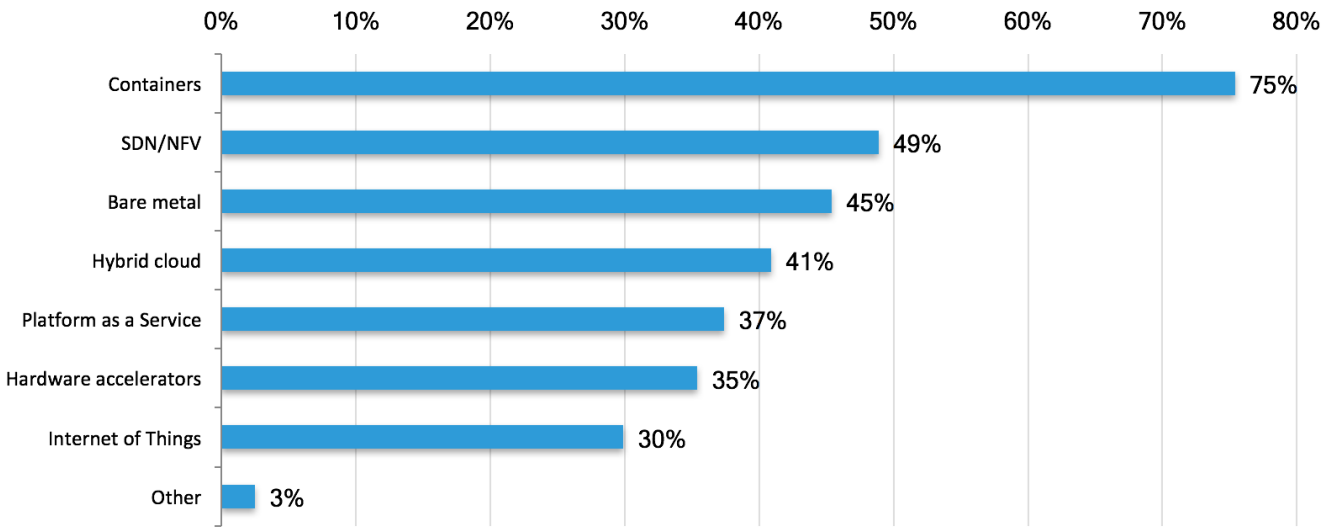

Figure 2.5 n=981

In what stage are OpenStack deployments?

The share of deployments in production increased slightly year-over-year, up one point (not a statistically significant increase). Two-thirds of deployments are in production or full operational use; the remainder are in testing or proof-of-concept phase.

In our October 2016 survey, we showed an even greater number of deployments in production, 71%, however the sample size then (n=233) was less than half of the sample size now (n=482).

Additionally, the October 2016 survey was positioned only as a "deployment update" and sent only to existing deployments, therefore newer clouds (likely to be in PoC or test phase) were less represented. Given these variances, we see clouds in production as relatively constant.

The large proportion of clouds in production demonstrates the maturity of OpenStack, while the substantial group of clouds in development predicts healthy growth for the future.

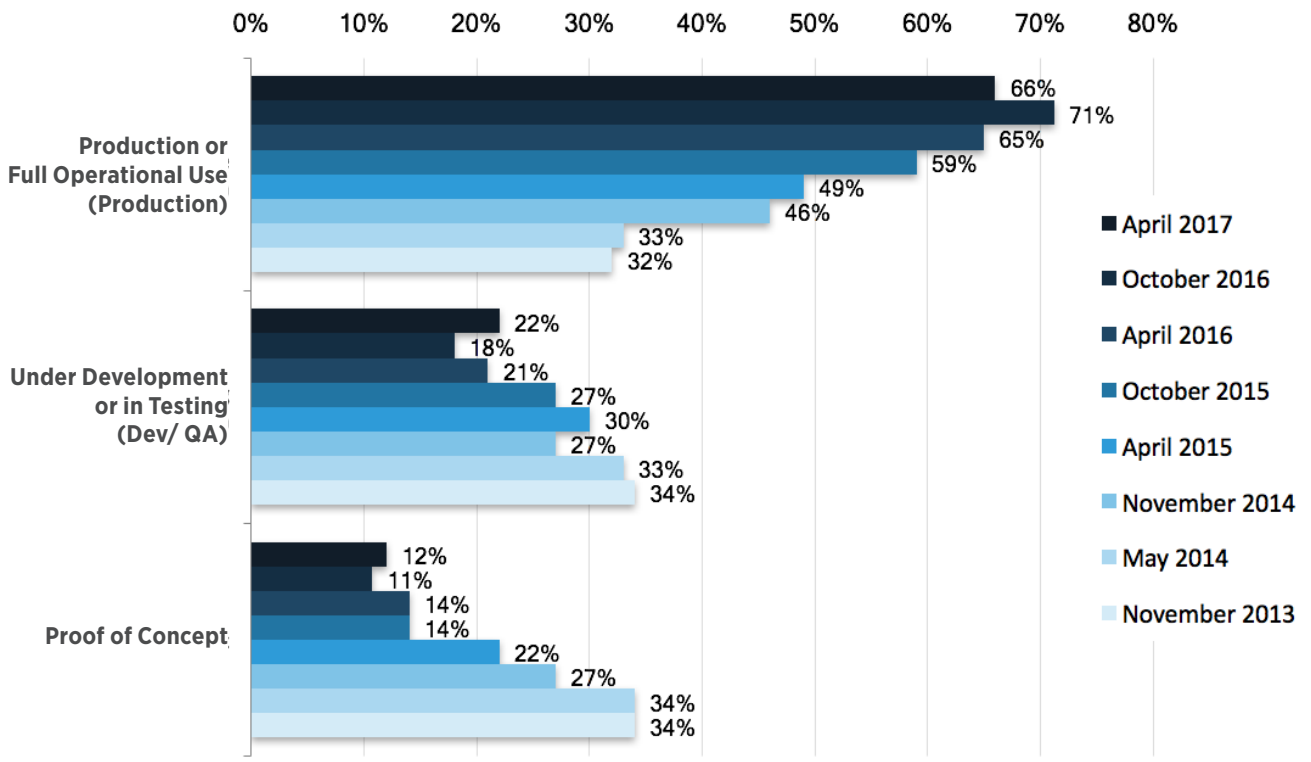

Figure 3.1 n=482 for April 2017

How mature are OpenStack deployments?

For greater detail on the maturity of deployments, we asked users when their deployment was created. Of the 470 deployments that responded to this question, the vast majority were launched in 2016. The average age of a cloud was just 1.68 years.

Breaking down these clouds by stage, we find that 54% of clouds launched in 2016 and 2017 are in production, and among clouds launched in 2015 and earlier, 82% are in full operational use.

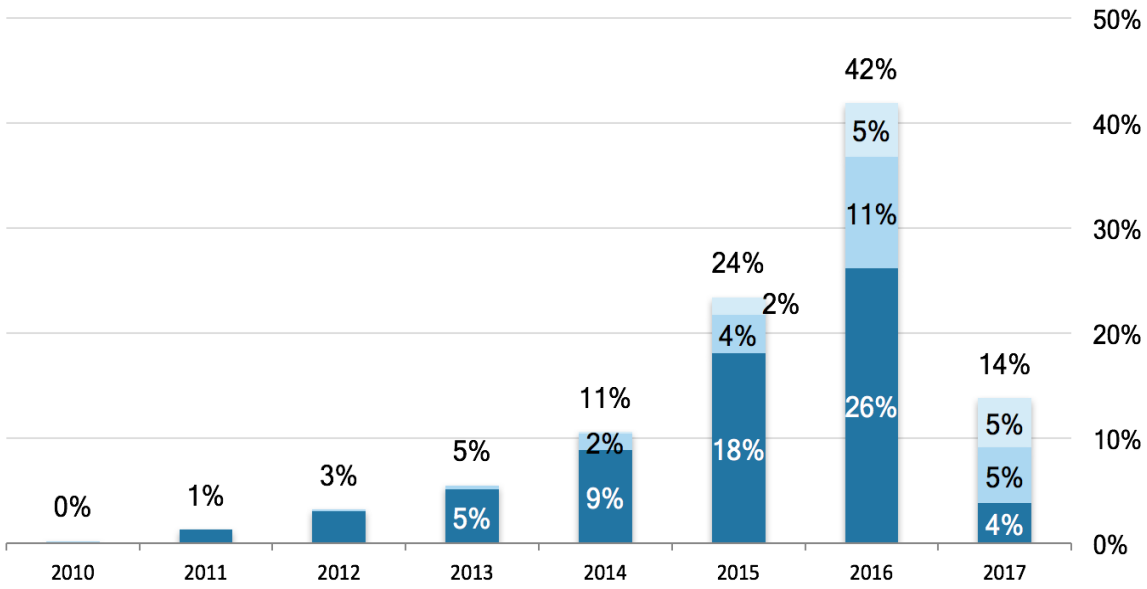

Figure 3.2 n=482 for April 2017

In which countries are OpenStack deployments located?

79% of deployments reported they are physically located in just one country, with 8% reporting they are located in five or more countries. The top countries reported were:

- United States 26%
- Germany 6%
- Canada 4%
- China 4%
- Spain 4%
- United Kingdom 4%
- Australia 3%
- France 3%
- India 3%
- Japan 3%
- Netherlands 3%

What types of clouds are running OpenStack?

On-premises private cloud represented 70% of deployments surveyed, up 5 points from one year ago. While the number of public clouds is growing overall, they made up a slightly smaller share of total deployments surveyed, as did off-premises private clouds.

Community cloud representation increased 2 points.

We compared cloud types both by size and business drivers, finding that private clouds tended to be smaller, although the median size was 100 to 999 cores for both private and non-private clouds.

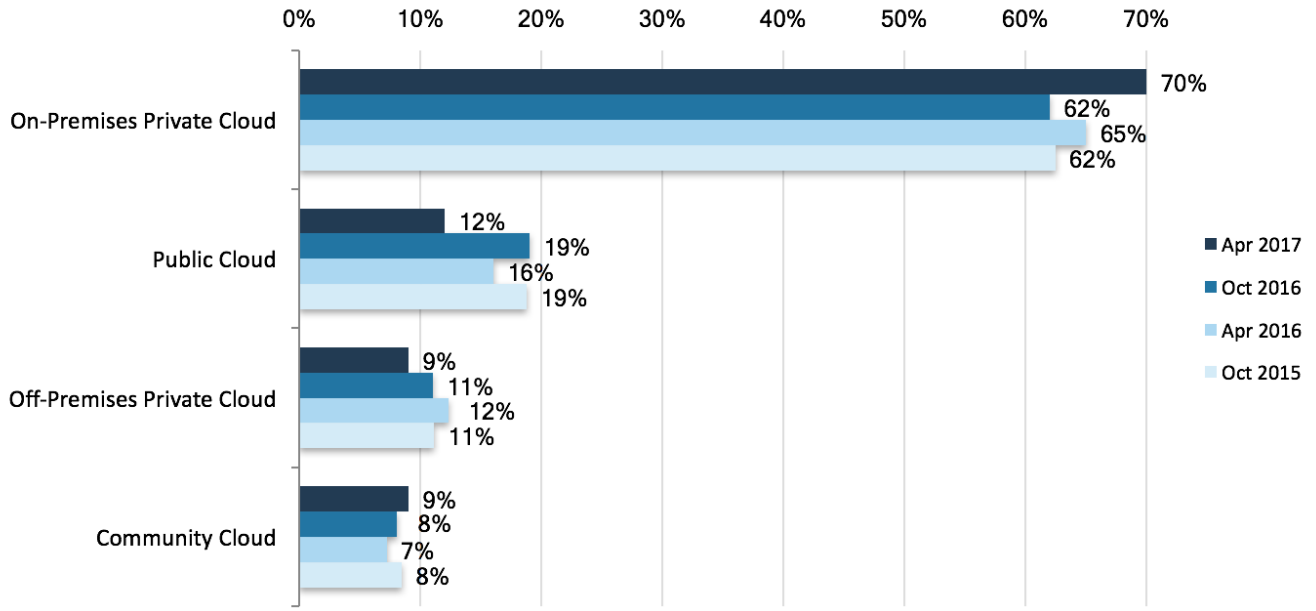

Figure 3.3 n=318

Which releases are deployments at all stages using?

Adoption of OpenStack's releases shows more deployments using Mitaka or Newton than older releases (Liberty and prior). Note that more than one answer was accepted per deployment so the percentage of users on each release adds to more than 100%.

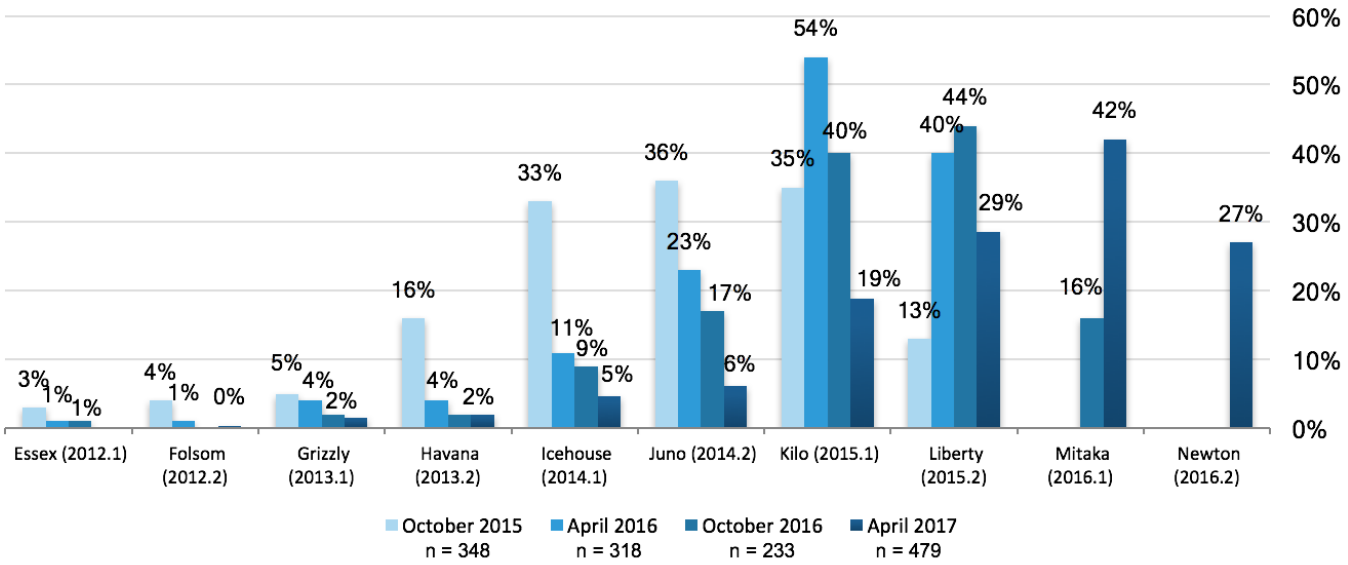

Figure 3.4

ICEHOUSE
THE NINTH OPENSTACK RELEASE

JUNO
THE TENTH OPENSTACK RELEASE

KILO
OPENSTACK GOES TO ELEVEN

LIBERTY
THE TWELFTH RELEASE OF OPENSTACK

OPENSTACK mitaka

OPENSTACK NEWTON

Newton, released October 2016, is the 14th release of OpenStack. It is followed by Ocata (released February 2017, shortly before the survey period ended) and Queens (September 2017).

How are releases being adopted by production deployments?

We also look at release adoption data filtered to deployments in production only, and it is consistent with past survey results showing slightly lagging adoption among these deployments. The Mitaka release is at the midpoint of this curve.

It is important to note that the Liberty release was end-of-life (EOL) and not security supported as of Nov. 17, 2016. Additionally, Mitaka will be EOL and not security supported on April 10, 2017. Combined with the still significant Juno and Kilo userbase, that means the majority of OpenStack users are on an unsupported release.

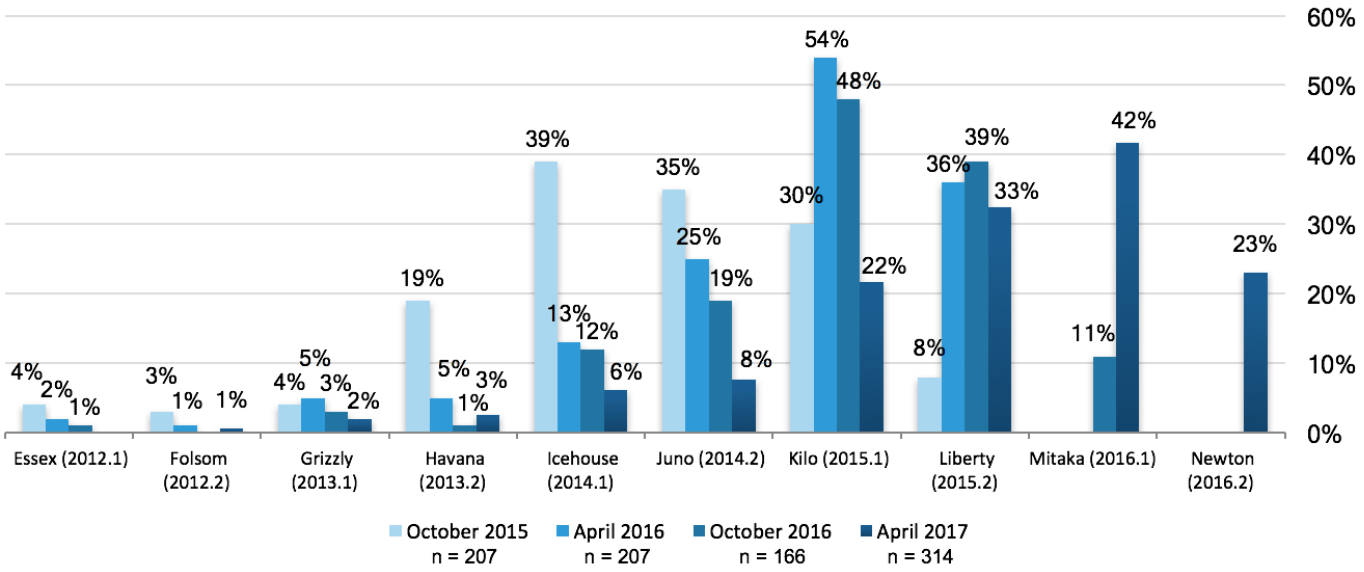

Figure 3.5

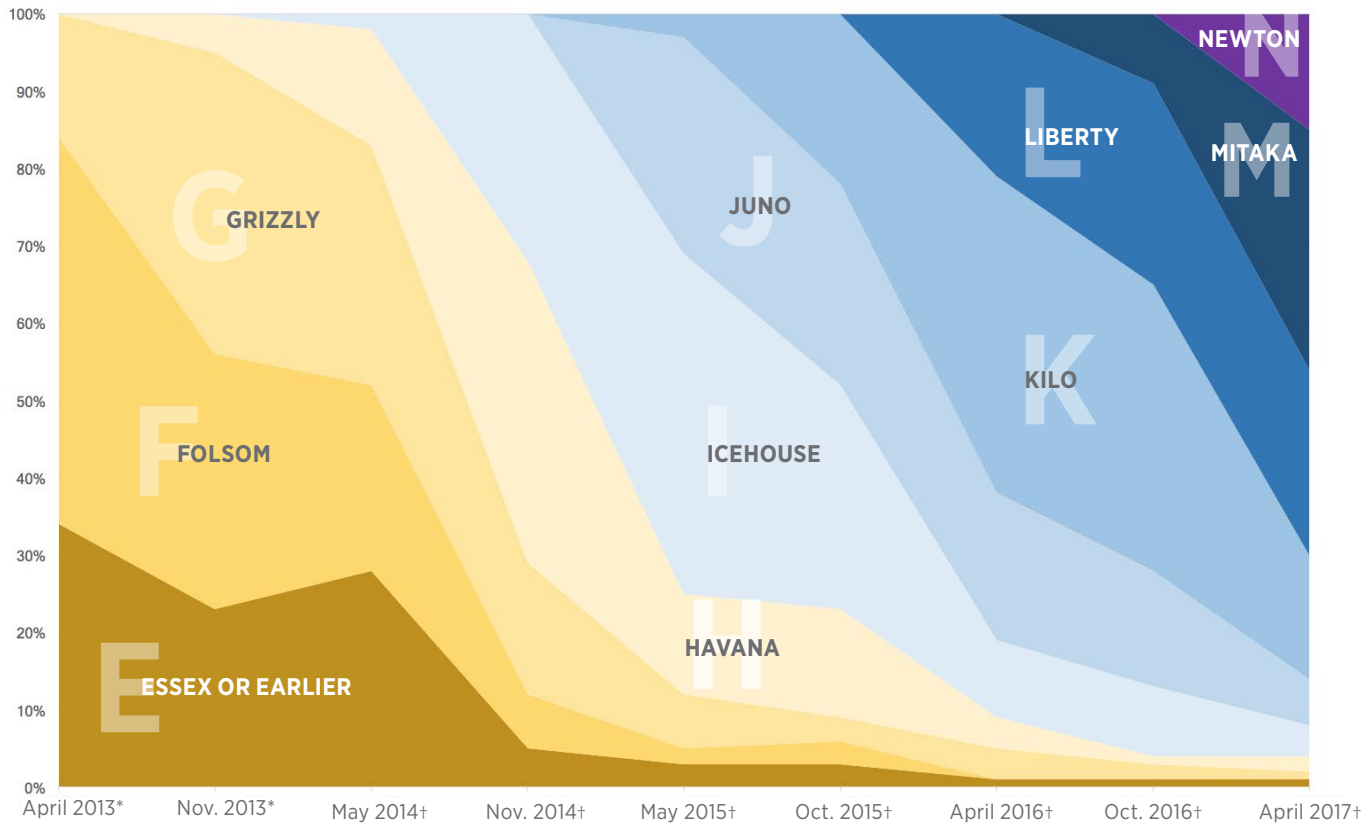

Figure 3.6

How are releases adopted by production deployments over time?

This area chart shows OpenStack deployments gradually adopting new releases, with the majority of those surveyed on the latest three releases. Particularly large areas, such as Icehouse and Kilo, speak to the popularity of these releases.

Note: In this chart, deployments running the Diablo release are wrapped into the Essex release. Diablo represented <7% of deployments in April 2013 to May 2014 and 0% from November 2014 to present.

* April 2013 and November 2013 figures include all deployments.

† May 2014–April 2016 figures represent production deployments only.

Which workloads and frameworks are running on OpenStack?

A wide array of workloads run on OpenStack, with infrastructure services significantly increasing by 11 points in this survey cycle. Software dev/test and QA and CI slipped 8 points, while the rest of the workloads remained relatively constant compared to a year ago.

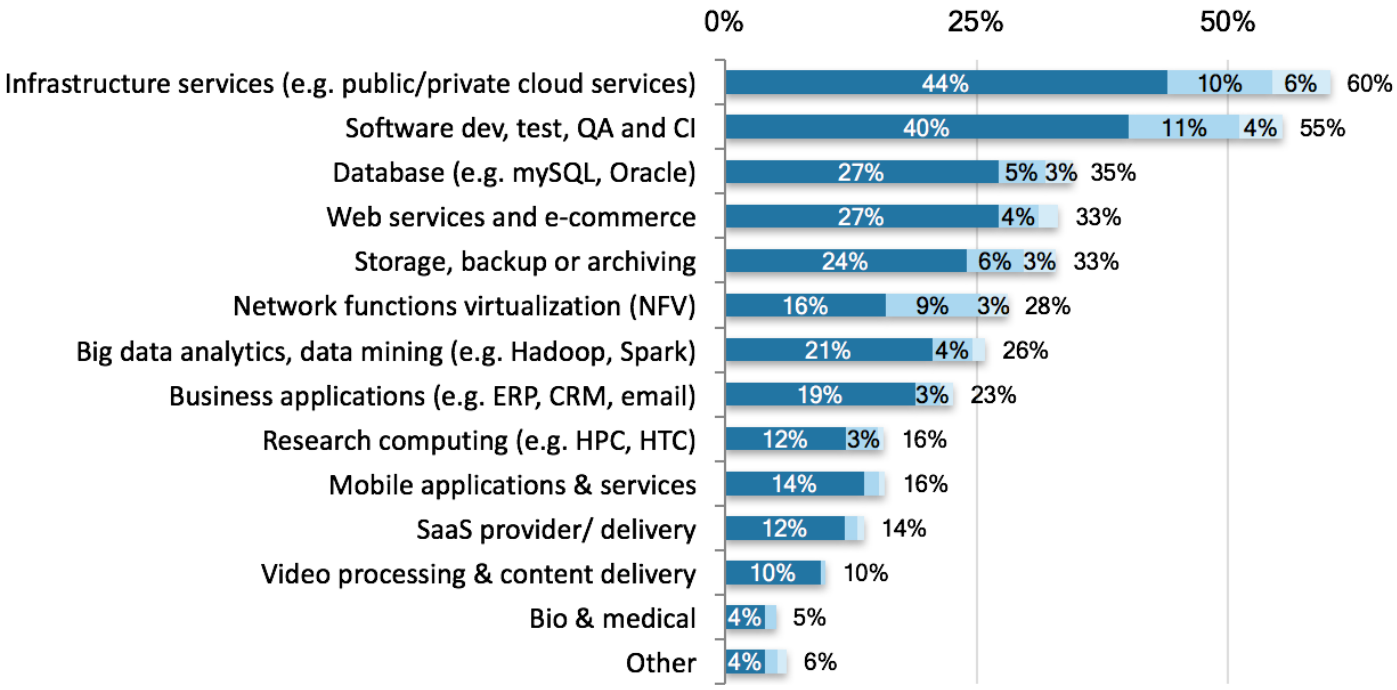

Figure 3.7 n=471

Which projects are used by OpenStack deployments?

OpenStack's core services (Nova, Neutron, Swift, Cinder, Keystone and Glance) have grown to almost total adoption by all clouds in production, while one year ago roughly one-third of them were in test phase. This survey showed significant gains in production use for Keystone, Nova, Glance, Neutron, Horizon, Cinder, Heat, Telemetry, Swift, and Rally.

Learn more about OpenStack projects' features and project maturity at https://www.openstack.org/software/project-navigator.

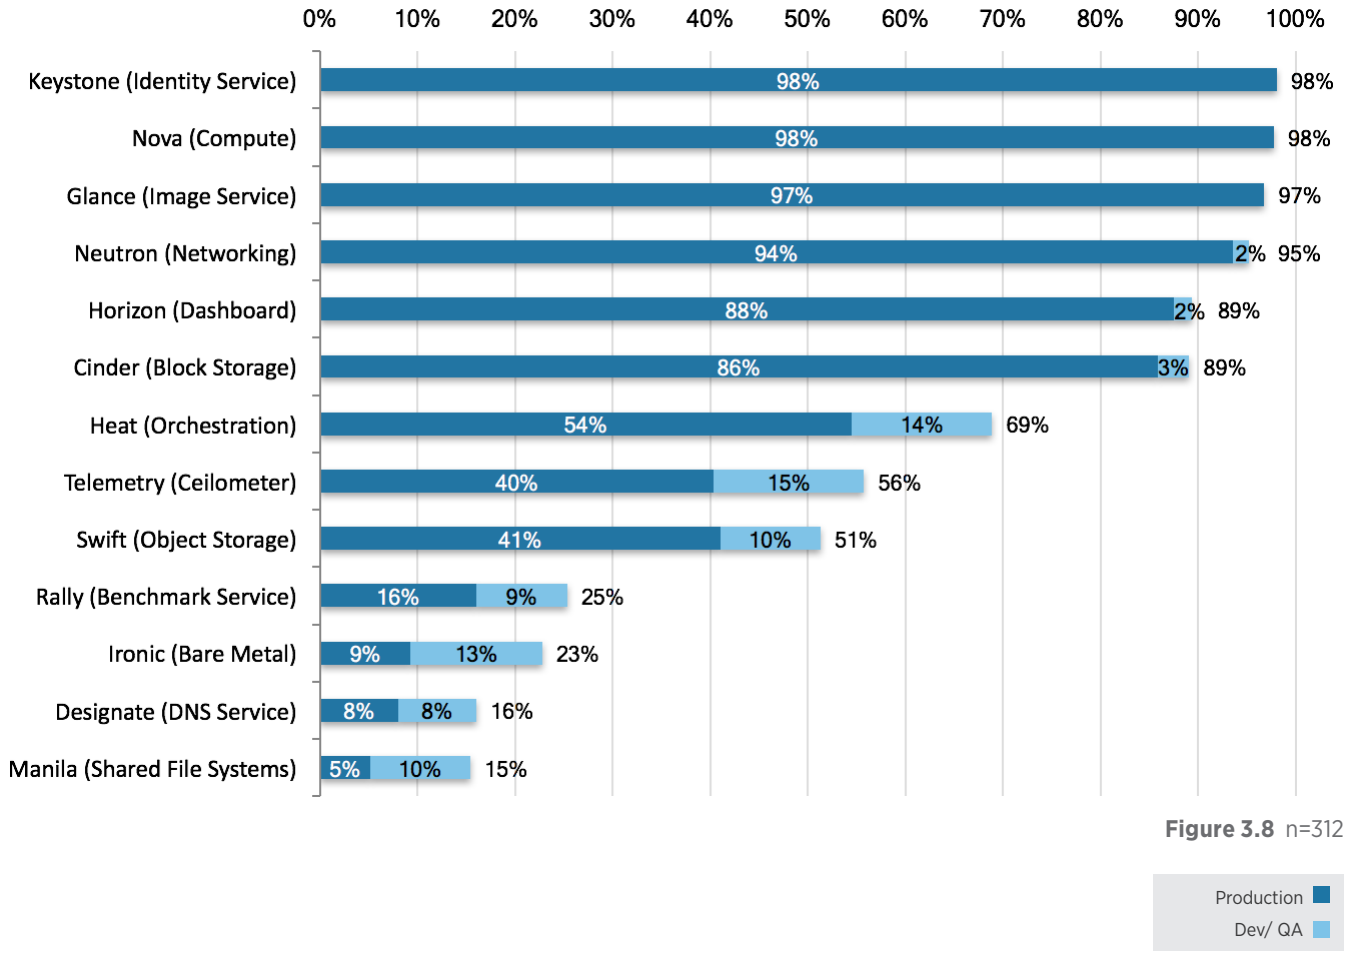

Figure 3.8 n=312

Which projects are used by OpenStack deployments?

Among emerging projects, Kolla and Barbican showed significant increases in adoption compared to April 2016.

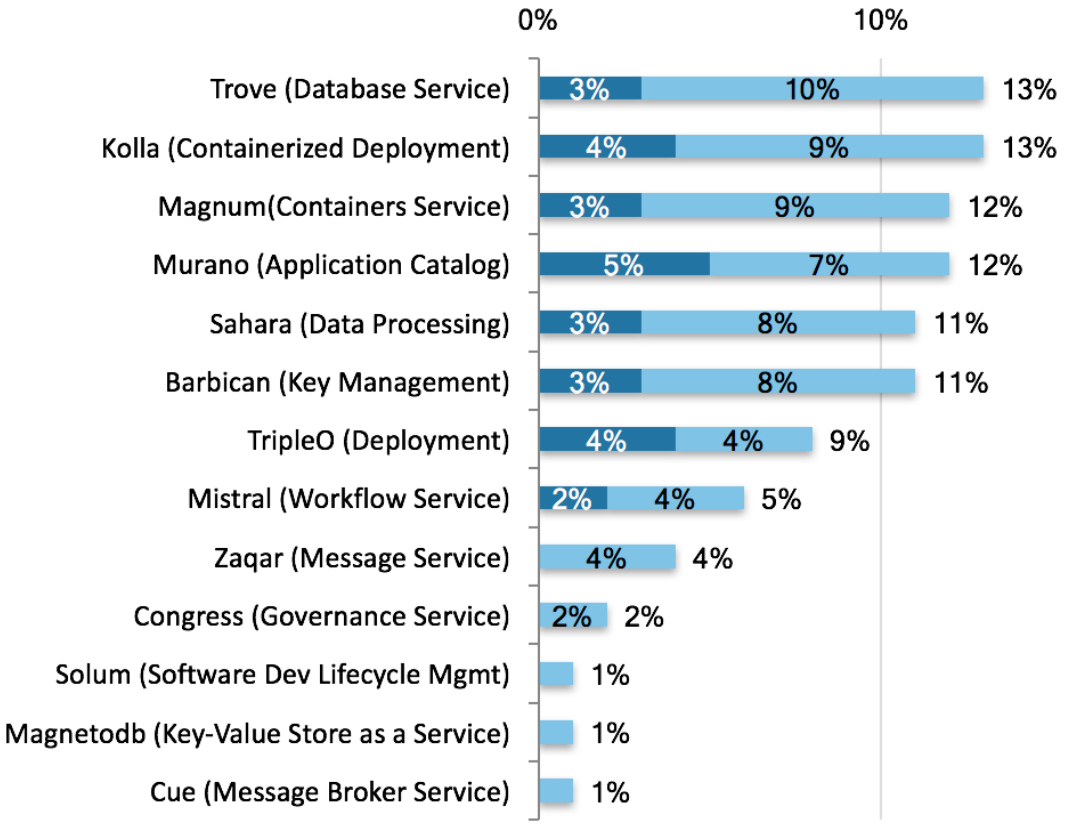

Figure 3.9 n=312

Which projects are users most interested in adopting in the future?

While users are asked about projects they are using in production or test phase, they are also asked to indicate whether they are interested in using other projects in the future.

This survey saw significantly higher interest in Kolla's containerized deployment, Barbican's key management, governance service Congress, workflow service Mistral, message service Zaqar, and Solum, a software development lifecycle management tool.

Interest in Ironic, the bare metal service, remained high at 29% but was significantly (8 points) lower than last year. Likewise, interest in Sahara fell off 9 points. The other projects did not change significantly.

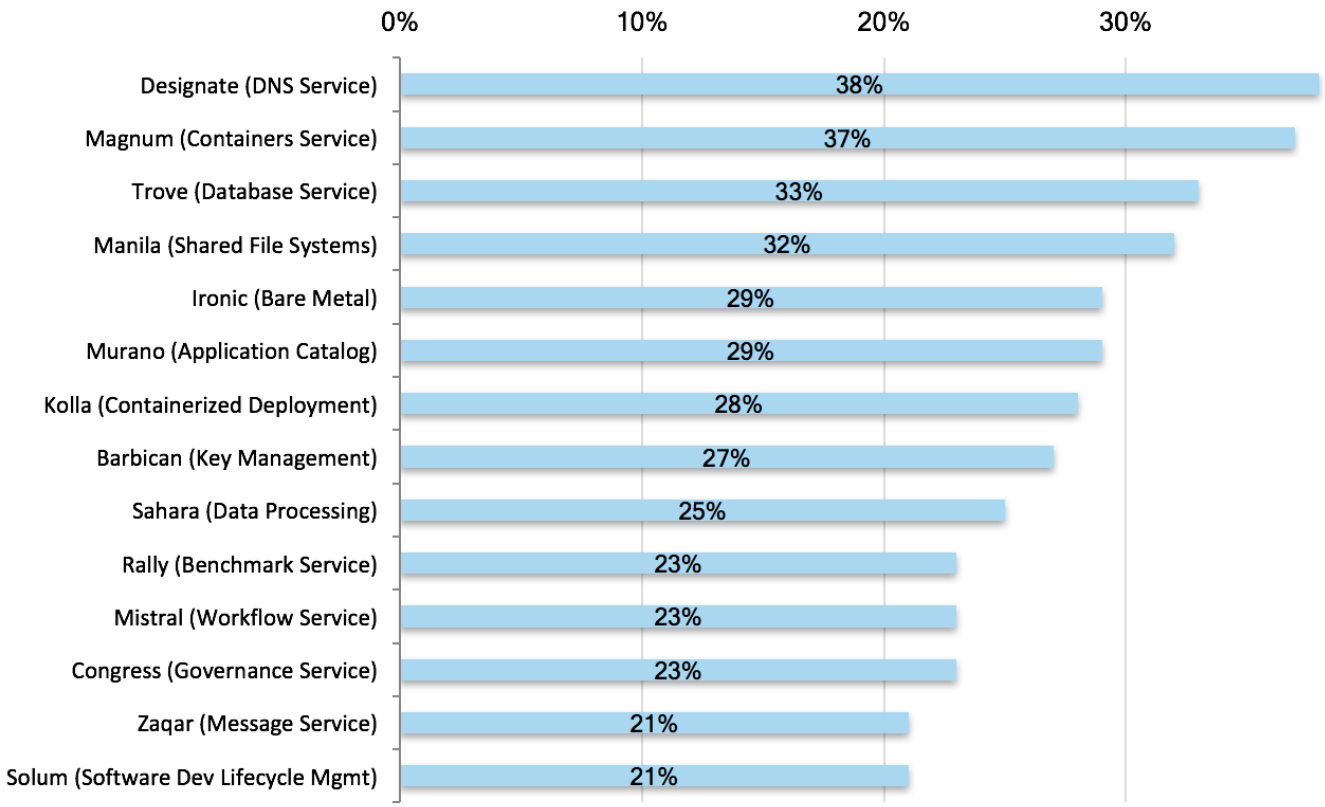

Figure 3.10 n=312

How many projects does a typical deployment use?

The average and median number of OpenStack projects used by all clouds was 9; while for deployments in production, 8 projects on average are used. It is also interesting to note that a significant portion of clouds (16%) use 12 or more projects.

Considering just 25 of OpenStack's 60-plus Big Tent projects are measured for adoption on the User Survey, it is likely that clouds use more, including projects that provide support such as Documentation and I18n (Translation).

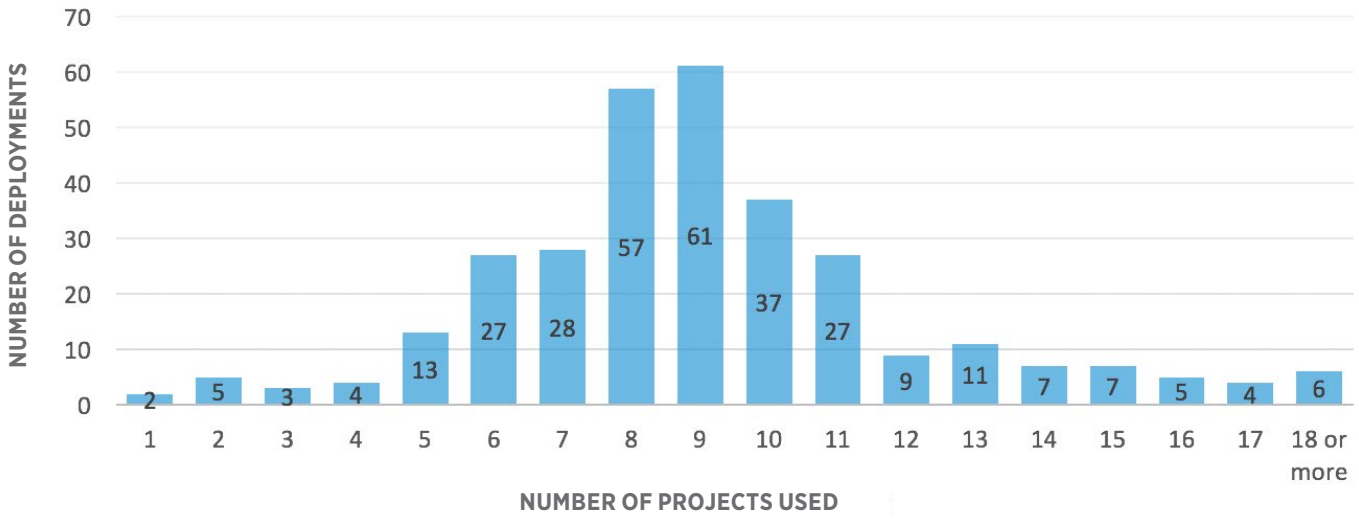

Figure 3.11 n=312

Mascots for some of the most frequently adopted projects:

What toolkits are application developers using with the OpenStack API?

The percentage of app developers using OpenStack clients declined 8 points to 70%; those who wrote their own toolkits also declined significantly by 13 points, and those using libcloud (Python) declined 14 points.

New answer options, which could account for lower answers on other questions, include Shade and gophercloud. Those indicating they do not use or plan to use the OpenStack API increased 8 points.

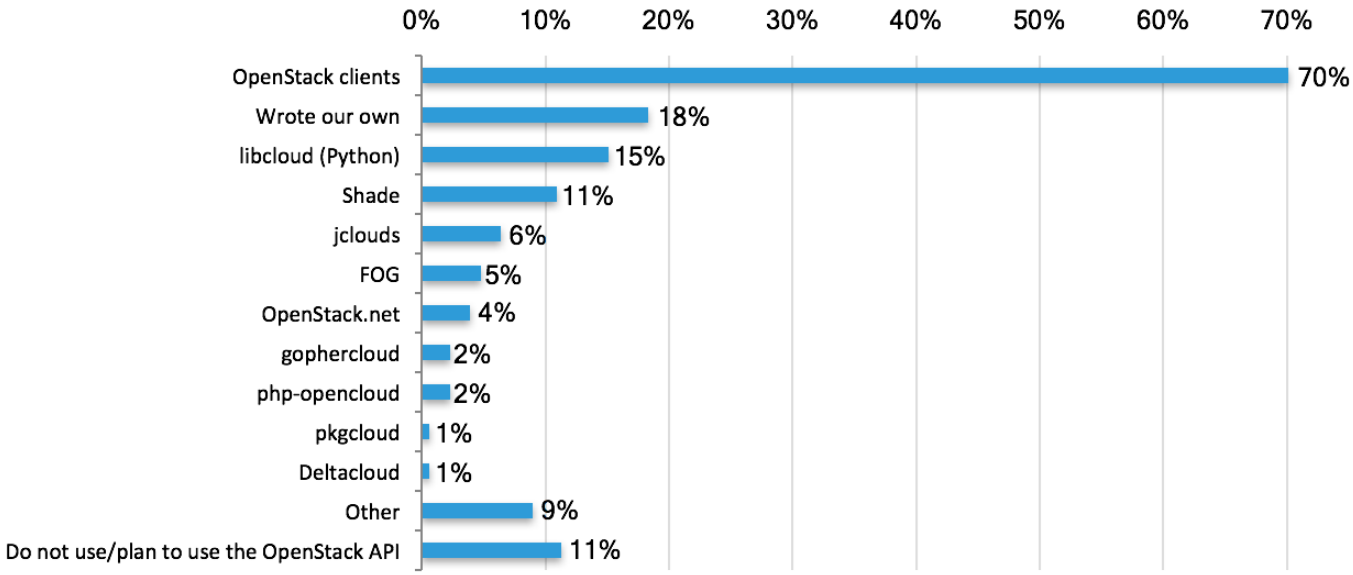

Figure 3.12 n=311

Part Four

CLOUD SIZE

How many users do OpenStack clouds support?

The median cloud supports 100–999 users, with a handful exceeding 100,000 users. There were no significant changes compared to last year's findings.

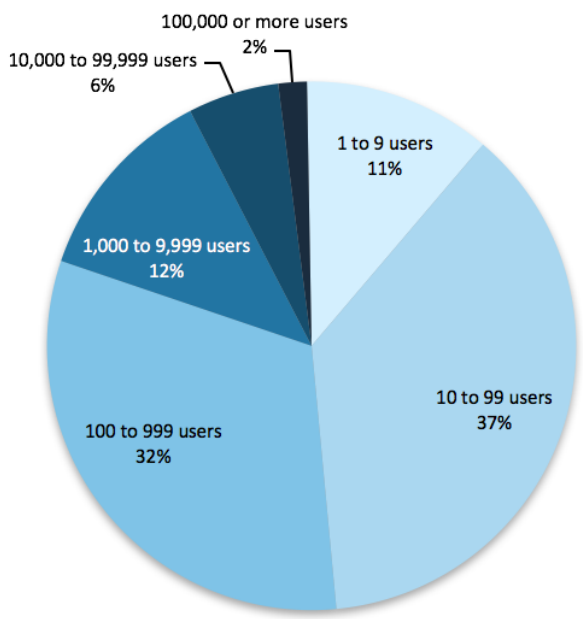

Figure 4.1 n=287

How many physical compute nodes do OpenStack clouds have?

The typical OpenStack cloud has 10–99 nodes, and about 8% of clouds exceed 1,000 nodes. This is 2 points higher than last year. Clouds with 100–999 nodes also increased by 3 points.

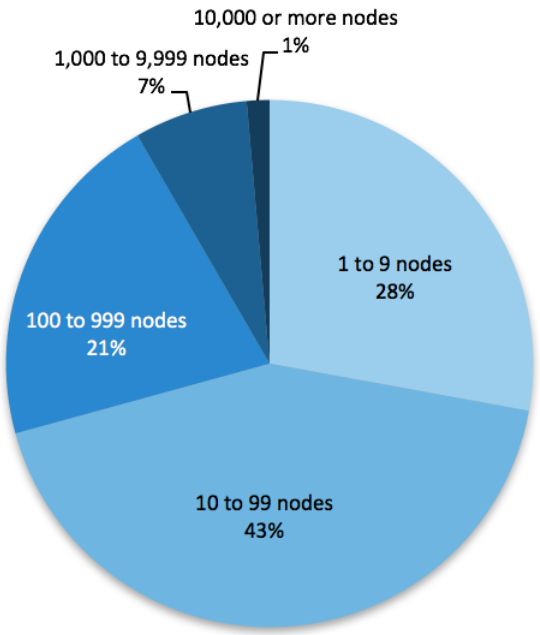

Figure 4.2 n=287

How many processor cores are in an OpenStack cloud?

Just 3% of OpenStack deployments have more than 100,000 processor cores in a single deployment, representing diverse organizations in industries such as research, telecom and ecommerce. Notably, the number of clouds with 10,000–99,999 cores increased significantly, from 4% to 10% of the total pie, while the 10-99 cores slice dropped 7 points.

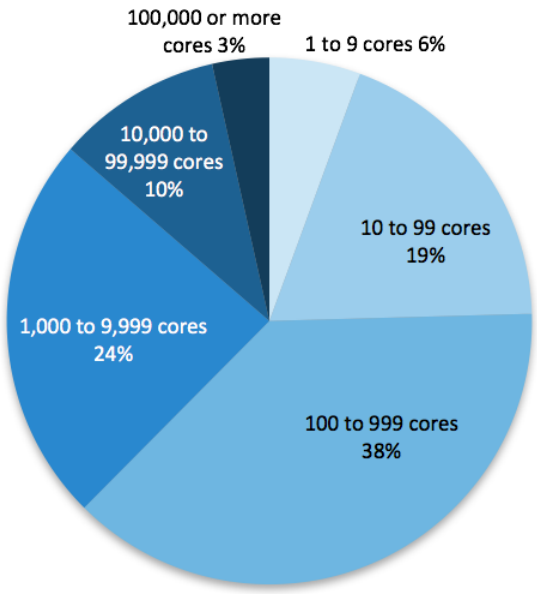

Figure 4.3 n=285

How many instances are in an OpenStack deployment?

The number of instances in a cloud edged up slightly, with a significant 10-point drop in those with 10–99 instances, resulting in small gains across the other segments.

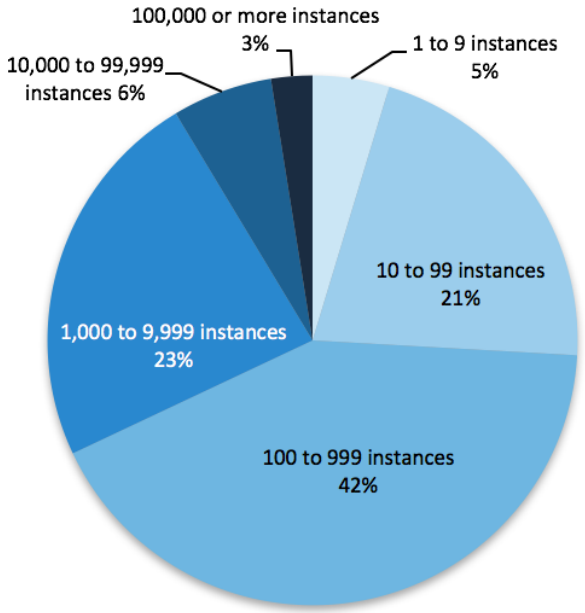

Figure 4.4 n=278

How many usable IPs are in an OpenStack cloud?

The number of usable IPs in a cloud also edged up slightly; the share of those with 1,000 or more IPs was 5 points higher than last year.

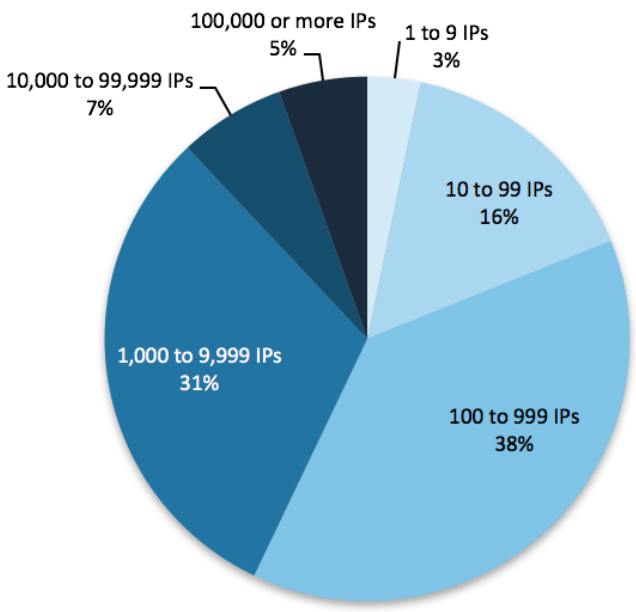

Figure 4.5 n=275

How much usable block storage in a Cinder deployment?

Cinder block storage increased significantly, with the smallest users (less than 9 terabyte) declining from 34% to 21% of the pie. A handful of users noted astronomical storage needs of more than 10 petabytes of storage.

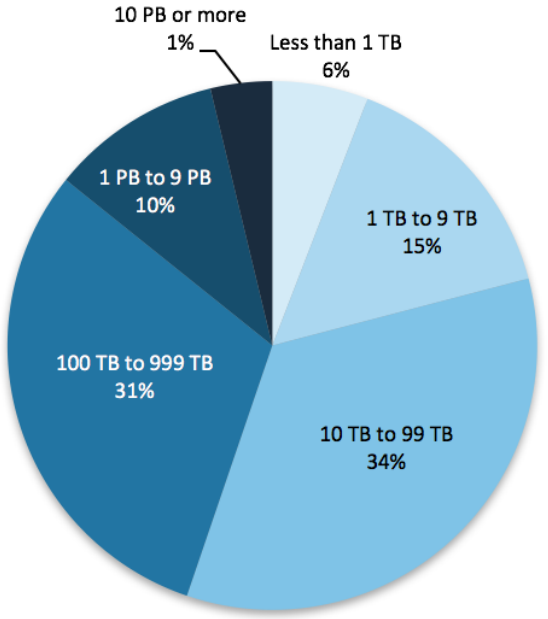

Figure 4.6 n=239

How much Swift object storage is provisioned in a deployment?

As with Cinder's block storage, Swift's object storage showed greater demand in this survey, with 16% of users provisioning more than 1 petabyte of storage, compared to just 4% reporting that amount provisioned one year ago.

Last year, nearly half of users required more than 10 terabytes of Swift object storage. This year, nearly two-thirds of users require that amount.

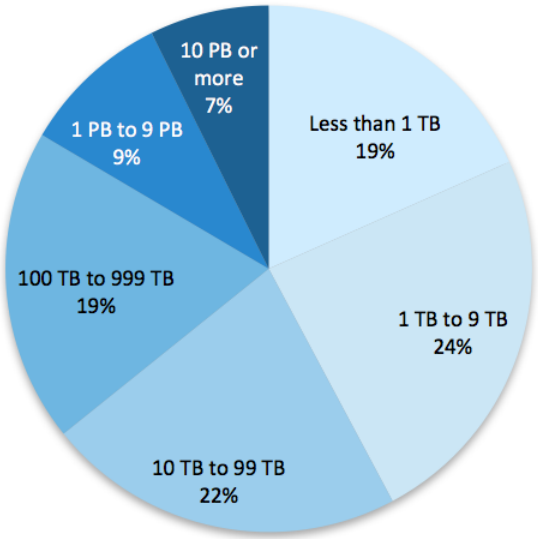

Figure 4.7 n=110

How many Swift objects are stored in a deployment?

The number of Swift objects stored by deployments increased significantly, from 13%, indicating they stored 100,000 objects or more last year, to 33% storing at least 100,000 objects.

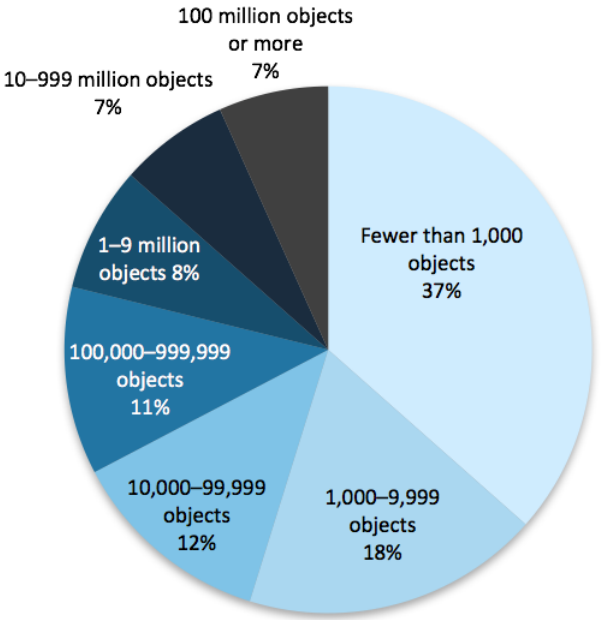

Figure 4.8 n=104

How many Nova cells are used?

While the majority of users aren't currently using Nova cells, about 146 deployments reported using cells, which is substantially more than the 67 users reporting using Nova cells one year ago.

In terms of raw numbers, substantially more users weighed in on this question—355 in this cycle compared to just 94 last year.

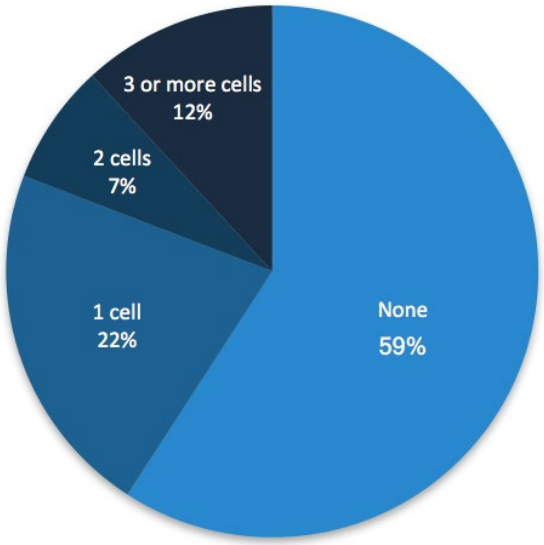

Figure 4.9 n=355

What packages are OpenStack deployments using?

We saw no significant changes from those answering this question.

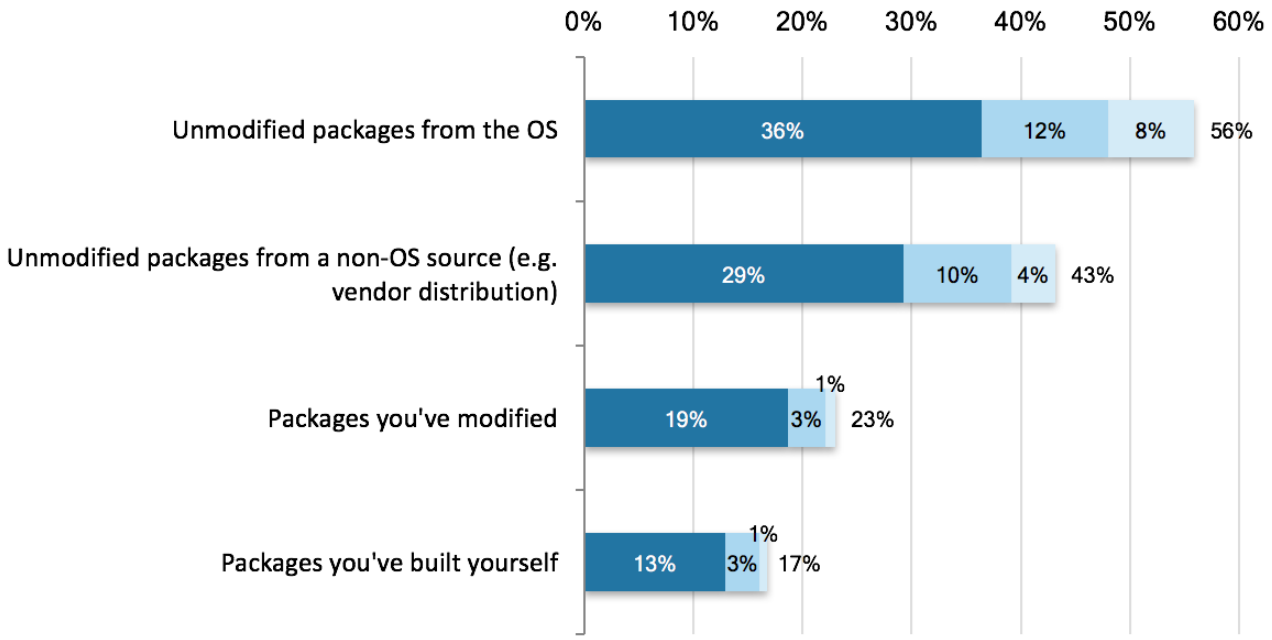

Figure 5.1 n=417

What tools are used to deploy/configure OpenStack clusters?

Puppet declined significantly in this survey, by 14 points, while the rest changed just slightly. Among the "other" answers, we saw several instances of RedHat Director and TripleO.

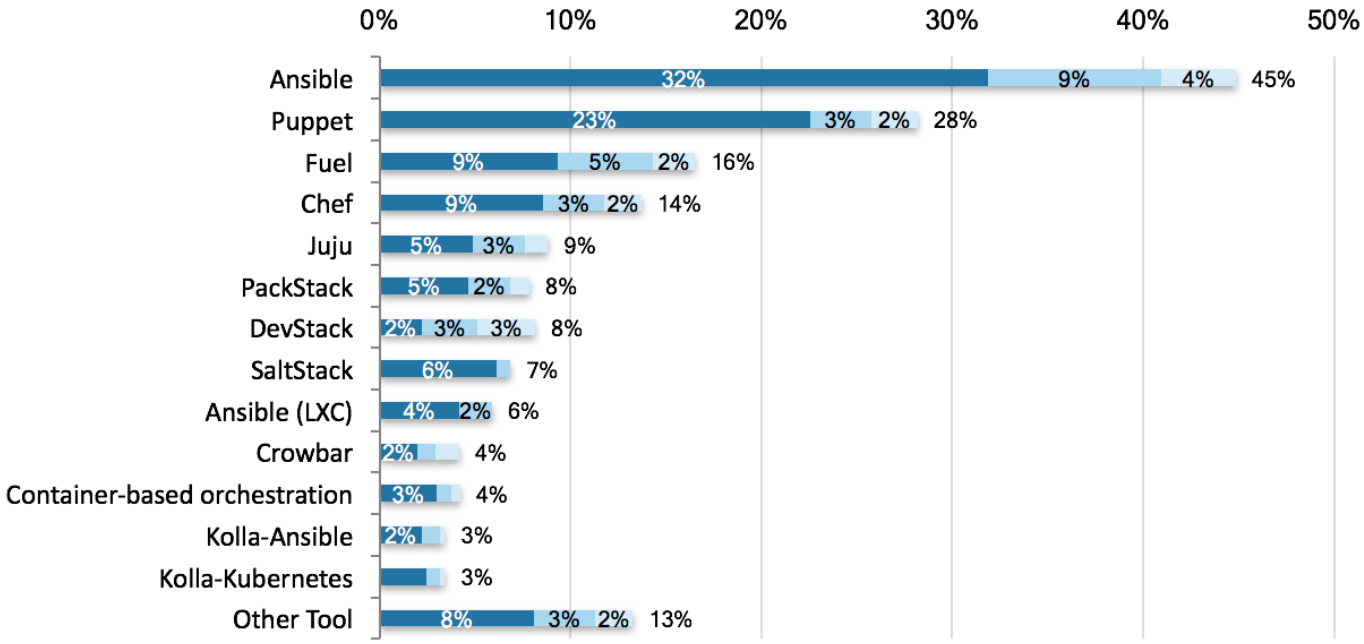

Figure 5.2 n=408

Among those running OpenStack services inside of containers, which container format is used?

With the rise of containers, this is a new question to the survey this year. We found that 20% of the 167 respondents indicated they were running more than one container format.

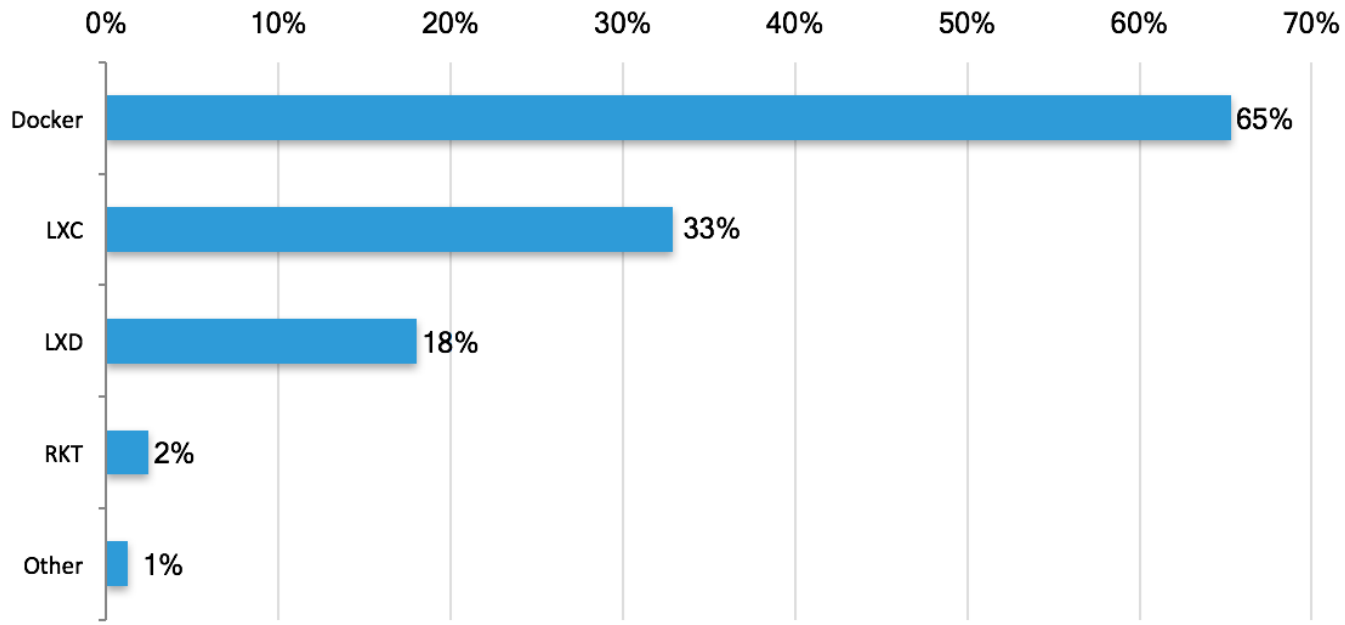

Figure 5.3 n=167

Which container and PaaS tools are used to manage applications on this OpenStack deployment?

While Kubernetes use remained relatively constant (up 5 points, but not a statistically significant difference), those using Docker Swarm increased 7 points and those using OpenShift fell 7 points, while CloudFoundry and Mesos each fell 12 points. Juju was noted among other PaaS tools.

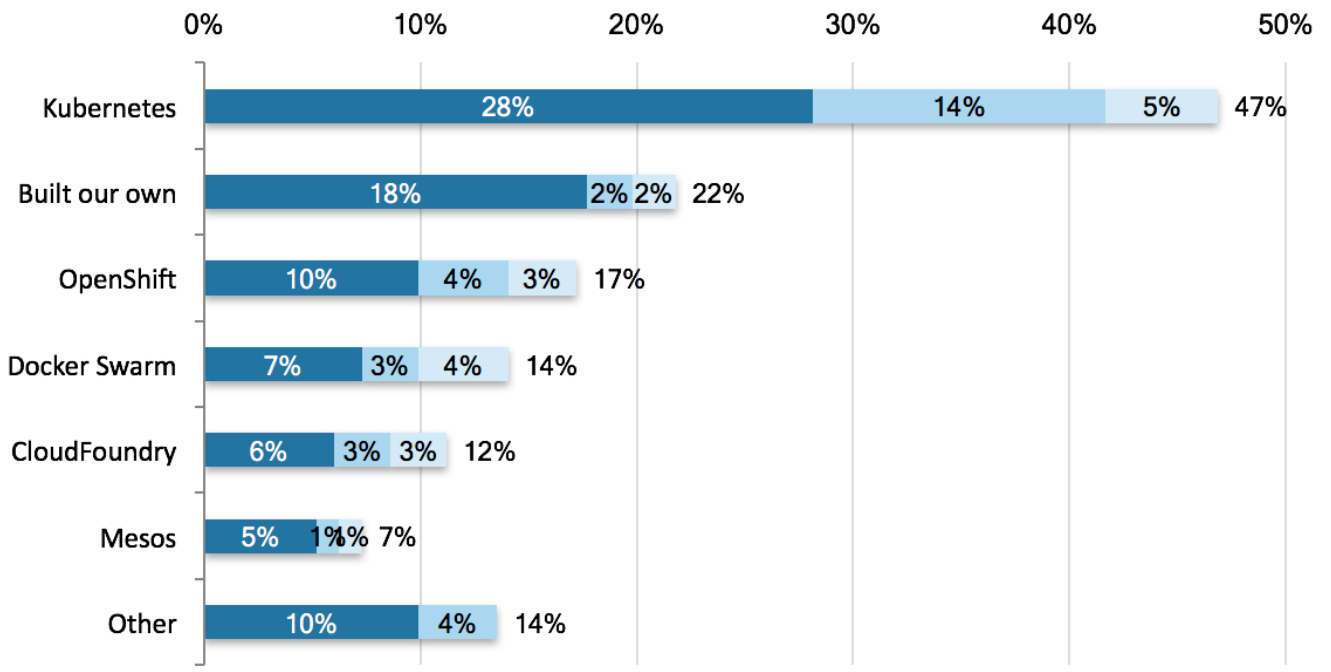

Figure 5.4 n=192

How are container and PaaS tools used in the past 12 months?

Since there are relatively few answers to the previous question, "Which container and PaaS tools are used to manage applications?" (192 in April 2017 and 118 in April 2016), and the amount of crossover among those answering both surveys was relatively low (about a third), we also looked at the combined answers for both April 2016 and April 2017, while de-duplicating those who answere in bost surveys (keeping their most recent answer). Kubernetes leads, followed by OpenShift and CloudFoundry.

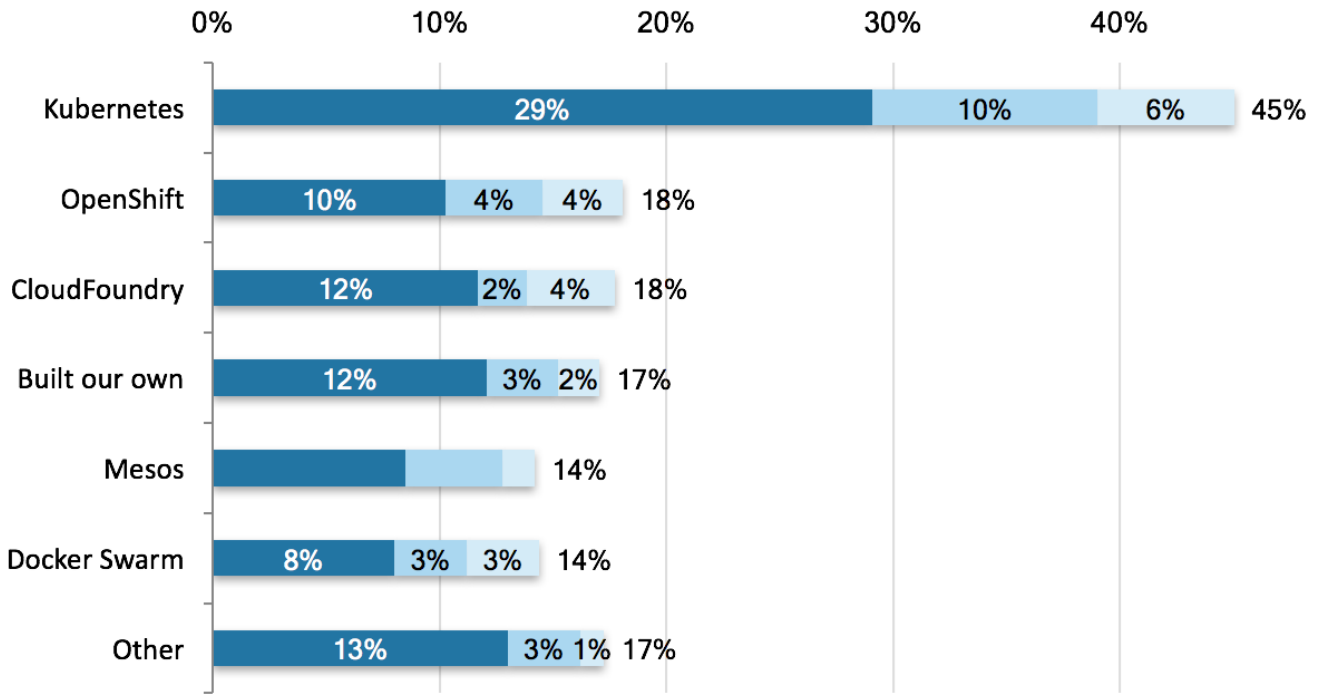

Figure 5.5 n=282

Which OpenStack Compute (Nova) hypervisors are in use?

KVM slipped slightly, down 4 points, and QEMU was up 8 points compared to last year. Despite high number of respondents reporting use of QEMU, its actual use is expected to be much lower, specifically in production environments. These deployments can be counted as using KVM as hypervisor. This is related to fact that most of those deployments are using both KVM and QEMU but should be considered to be using KVM from perspective of this survey.

Bare metal remained constant while ESX moved up 2 points and Docker dropped by 4 points, which is likely related to the deprecation of the Nova-Docker driver. FusionCompute was noted among other hypervisors.

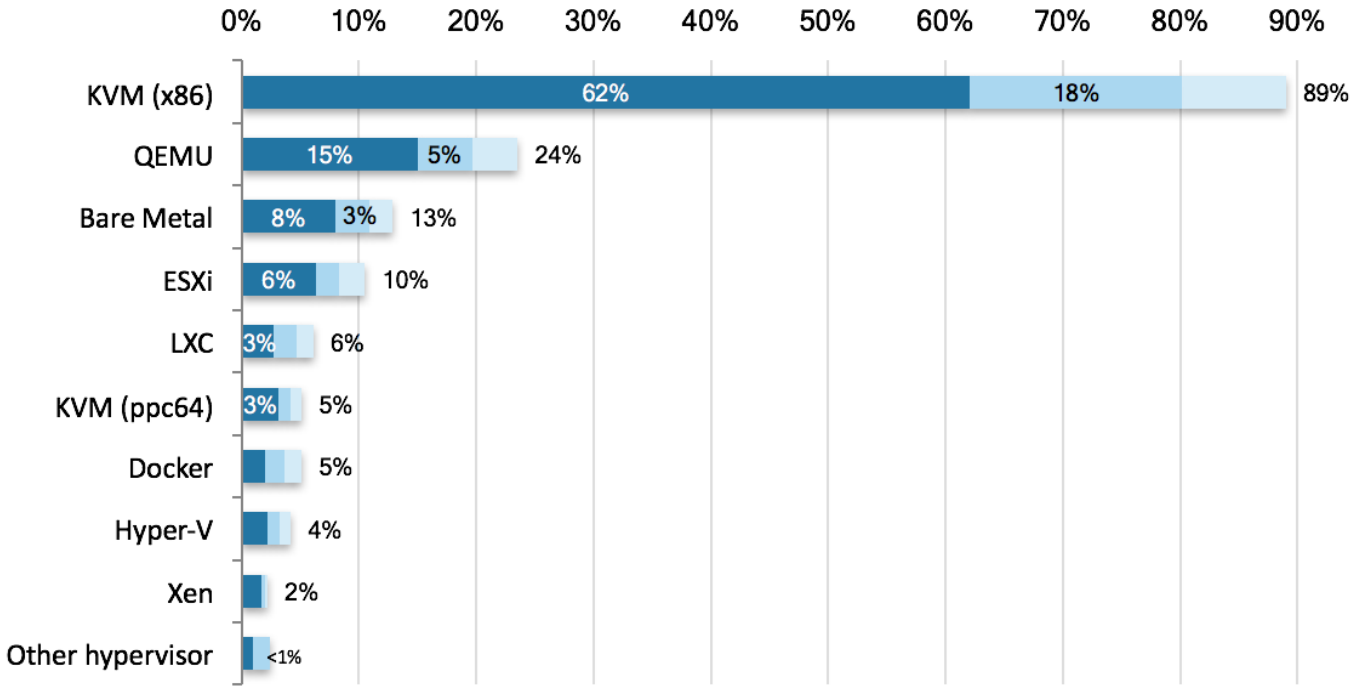

Figure 5.6 n=409

Production
Dev/ QA
Proof of Concept

Which databases are used for OpenStack components?

MariaDB Galera Cluster moved from fourth to first place on this list, up 7 points, while MongoDB slipped 6 points. MySQL with Galera also dropped significantly, down 8 points. This follows a longer-term trend over time where deployments are moving away from MySQL (down 22% over past 18 months) to MariaDB and others. IBM DB2 was noted among other databases used.

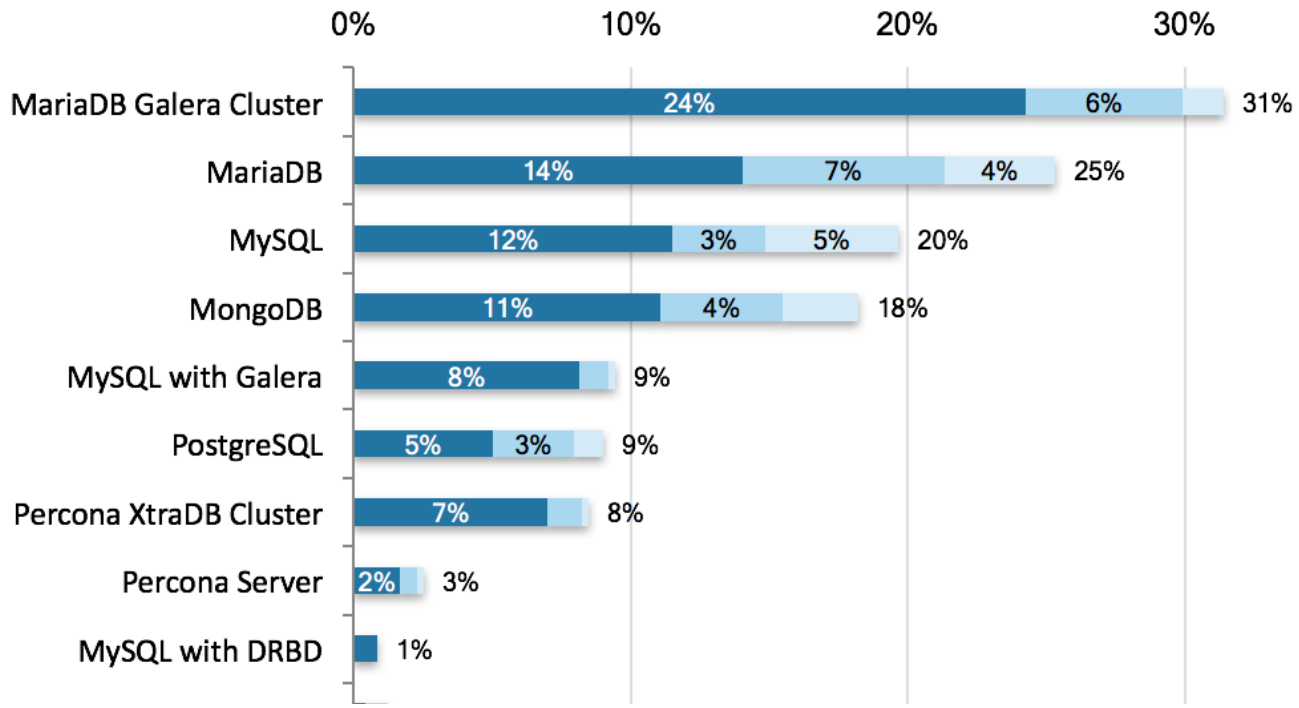

Figure 5.7 n=478

What programmatic interfaces (SDKs, tools, platforms) are exposed to users, customers or developers so they can consume a product or cloud?

Also new to the survey this year, this question found that 70% of programmatic interfaces use OpenStack clients, while "wrote our own" was also popular. Among the "other" answers, we saw use of Terraform.

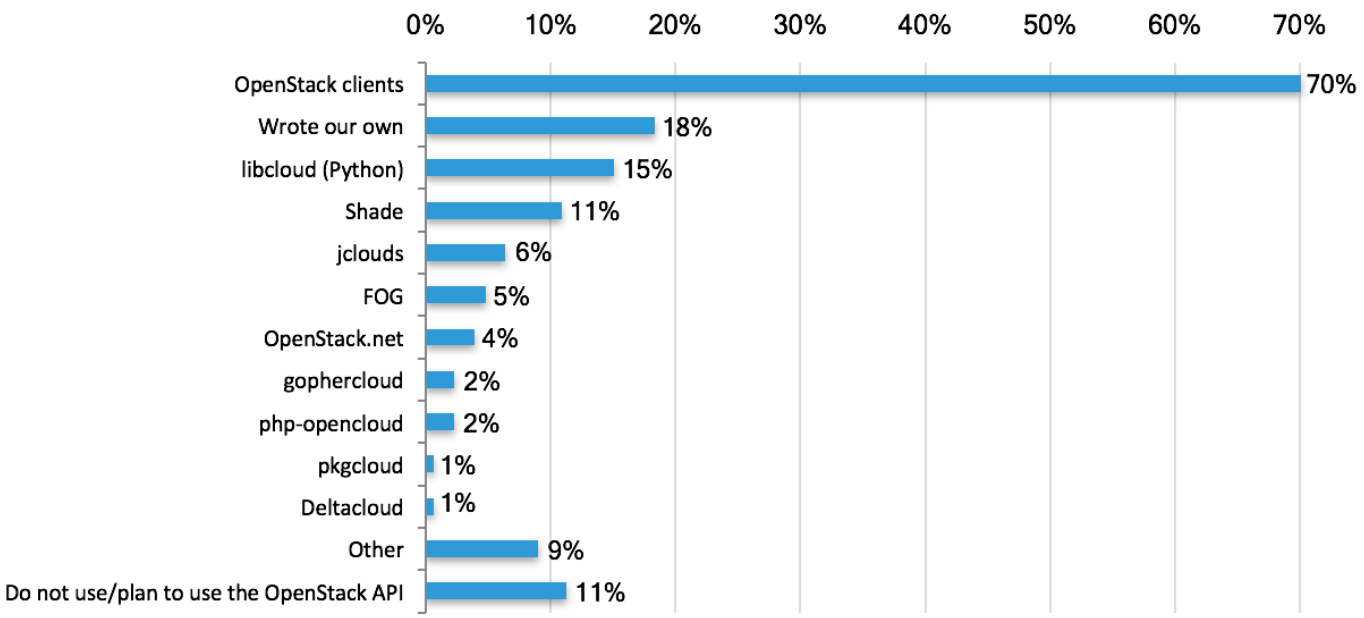

Figure 5.8 n=311

With which other clouds do users interact?

This question was modified in this user survey to be asked of all deployments (not only app developers) and to ask explicitly if users are not interacting with other clouds (the "none" option, whereas in the past, users with no interacting clouds would simply skip this question).

As a result, direct data comparisons to April 2016 are not possible, but generalized comparisons excluding the "none" answer suggest that the portion of users interacting with Amazon Web Services and Google Compute Engine are stable, while there is a large increase in usrs interacting with OpenStack Public Clouds and a drop in those interacting with Azure and OpenStack private clouds.

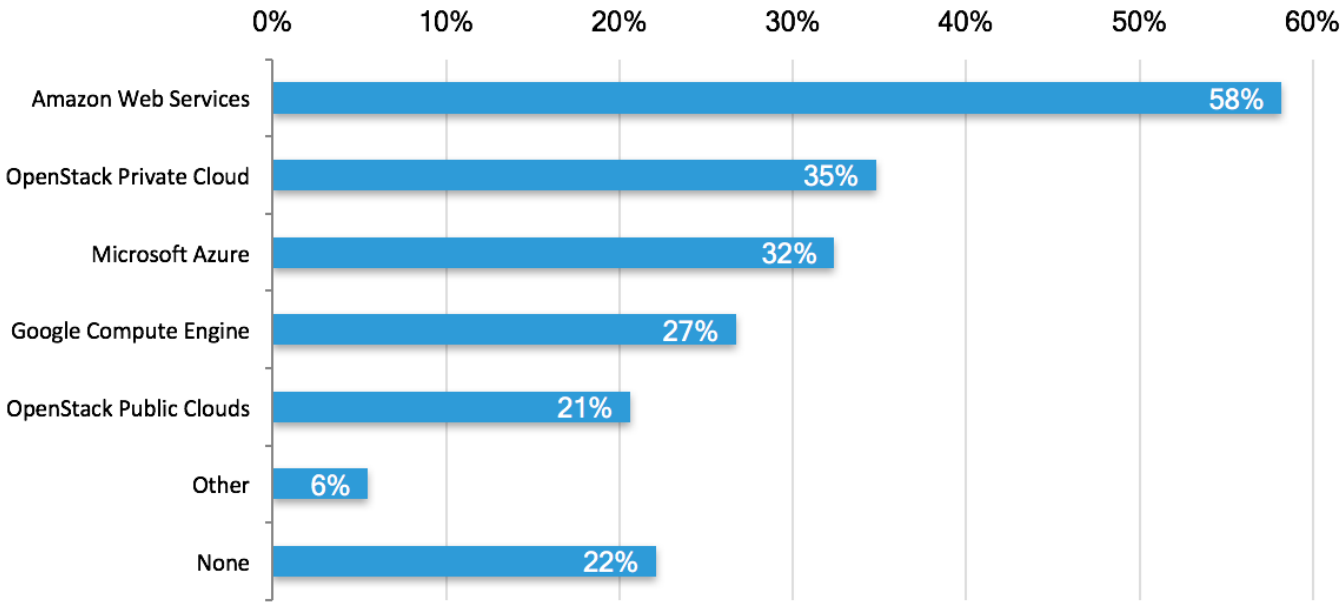

Figure 5.9 n=286

What percentage of users' cloud infrastructure runs on OpenStack?

The median OpenStack user runs 61% to 80% of their infrastructure on OpenStack, with 44% of users running 81% to 100% on OpenStack. This speaks to strong adoption throughout the enterprise, not only on emerging infrastructure and new applications, but also for legacy systems that are being moved to OpenStack.

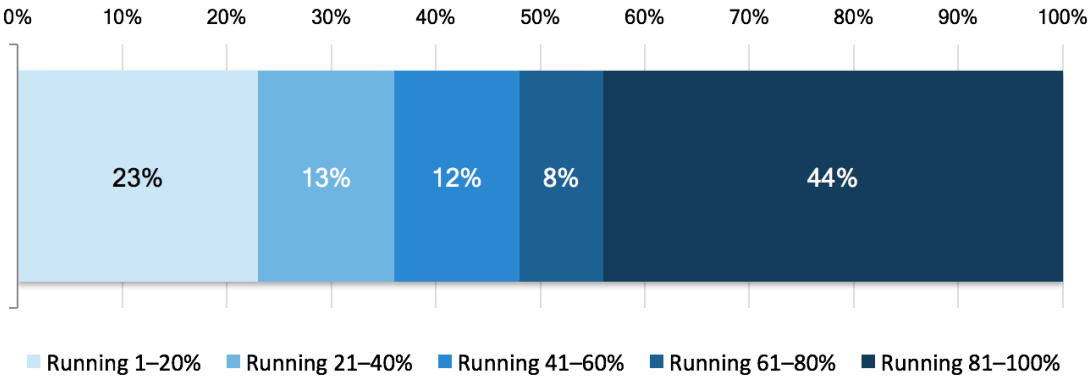

Figure 5.10 n=375

When we dug deeper into this question, which is new in this survey cycle, we looked at the largest clouds—those with 1,000 cores or more—and found that the majority of large users run 81% to 100% of their overall cloud infrastructure on OpenStack.

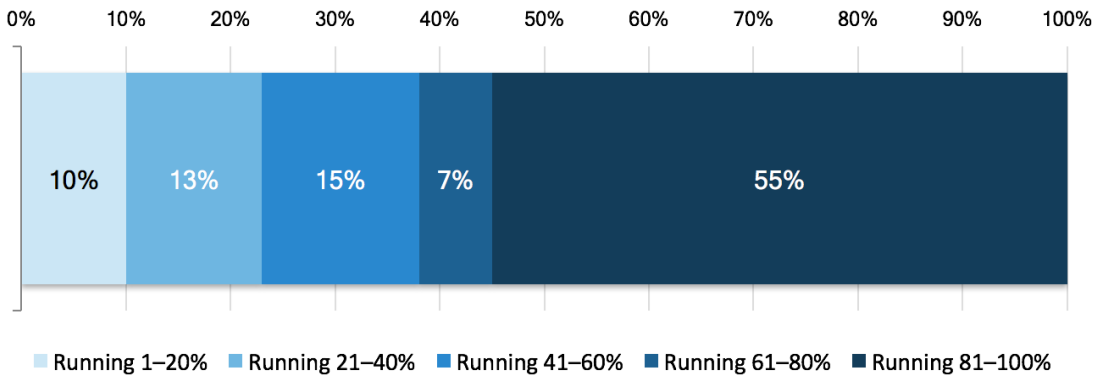

Figure 5.11 n=99

Which OpenStack network (Neutron) drivers are in use?

In this cycle, we substantially revised this question to reflect more than 20 potential answers supplied by the community in past surveys. This doubled the number of possible answers. As a result, it's difficult to provide a direct comparison as some answers from April 2016 have been segmented into multiple options in this survey.

With that in mind, last year Open vSwitch commanded 60% of answers, followed by ML2 at 31%, Linux Bridge at 21%, and Juniper/OpenContrail at 14%.

Neutron drivers not listed here, but which garnered a handful of responses each (1% to 2%), were ML2 VMware DVS, Hyper-V, Arista, ML2 Huawei Agile Controller, Nuage Networks, ML2 Mellanox SR-IOV, ML2 - Big Switch, PLUMgrid and ML2 OpenFlow Agent.

We also note that the number of "other network drivers" dropped from 12% to 5% of total responses, but did not cover all responses.

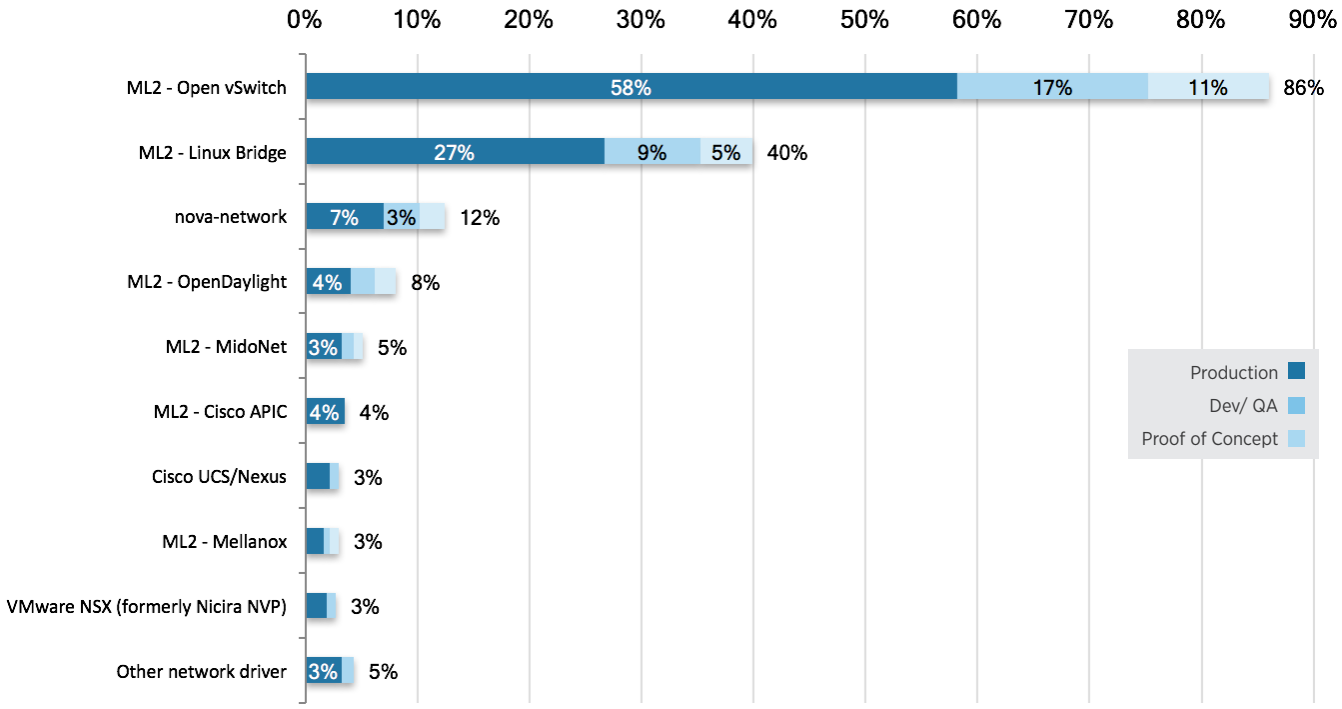

Figure 5.12 n=371

Which OpenStack identity service (Keystone) drivers are in use?

Keystone drivers remained relatively unchanged this cycle, with SQL down 3 points and AD up 6 points, while templated drivers fell off the answer list (they represented 3% of answers last year).

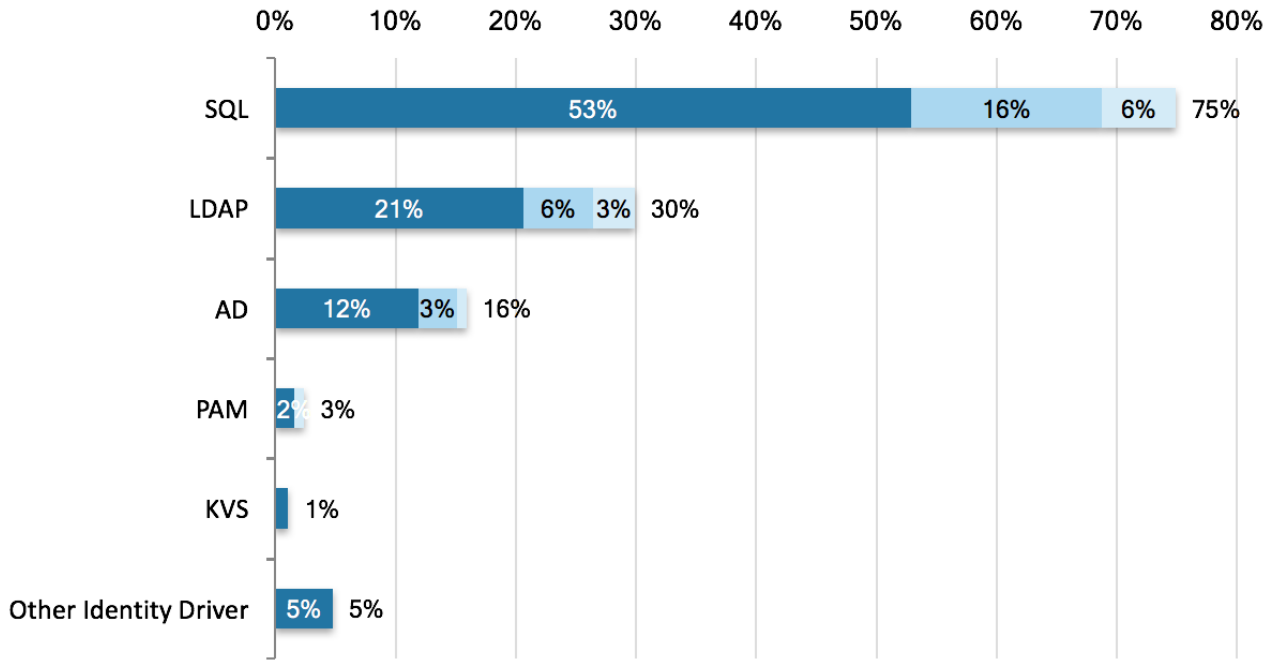

Figure 5.13 n=378

Which OpenStack block storage (Cinder) drivers are in use?

Cinder drivers also remained relatively constant, with Ceph RDB up 8 points and both LVM and NetApp up 3 points.

Just a handful of respondents indicated IBM Storwize, Huawei, HDS, IBM GPFS, Dell EqualLogic, IBM XIV/DS800, Windows Server 2012, Nexenta, SAN/Solaris, HP LeftHand, XenAPI Storage Manager, Sheepdog and IBM NAS.

Among the largest clouds with 1,000 or more cores, Ceph RDB is still dominant, but not used by the majority, while other block storage drivers were also popular.

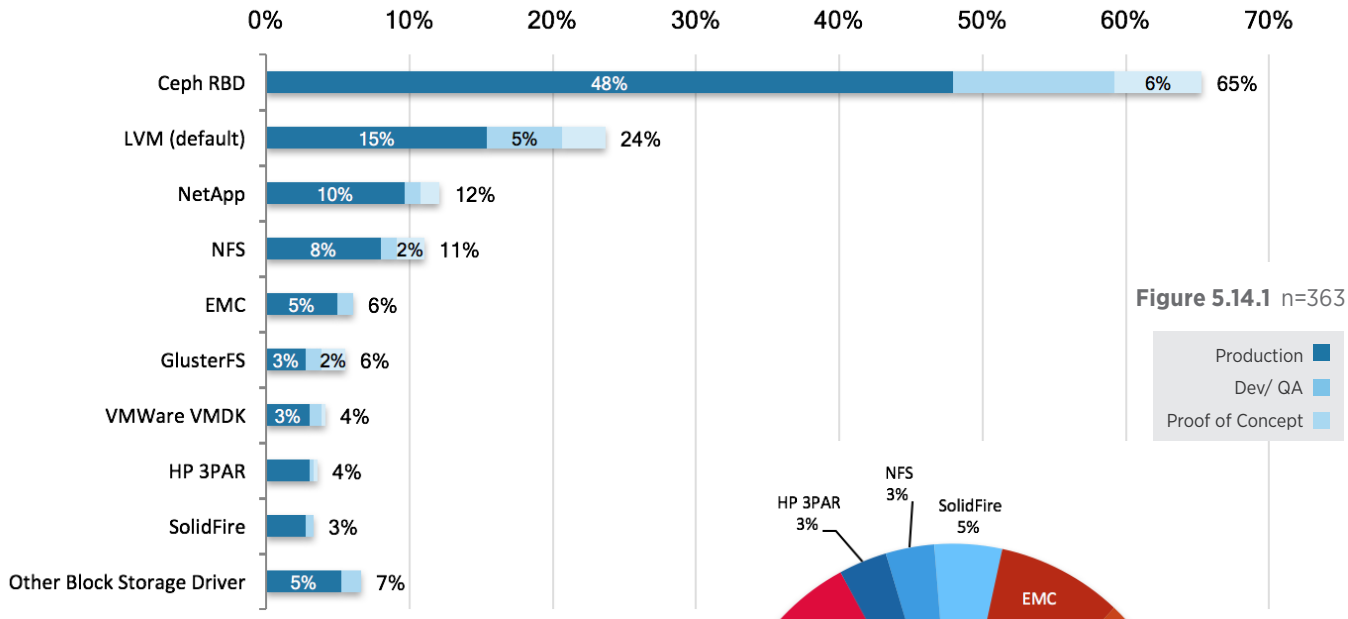

Figure 5.14.1 n=363

Among the largest clouds with 1,000 or more cores, Ceph RDB is still dominant, but not used by the majority, while other block storage drivers were also popular.

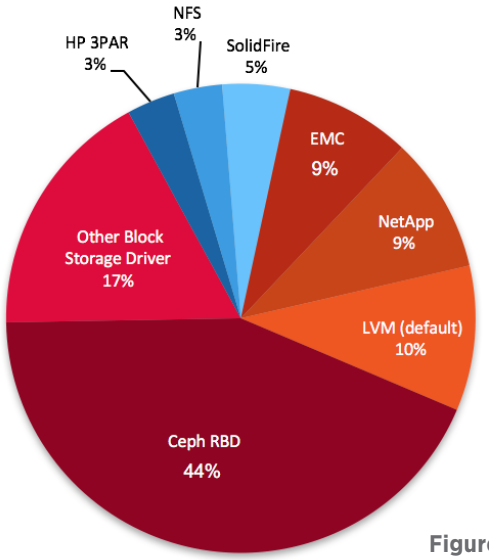

Figure 5.14.2 n=88

Which OpenStack-powered solutions are most used?

A new question on the survey this year, we asked which vendor's products power users' OpenStack clouds. Red Hat, Canonical and Mirantis topped the list.

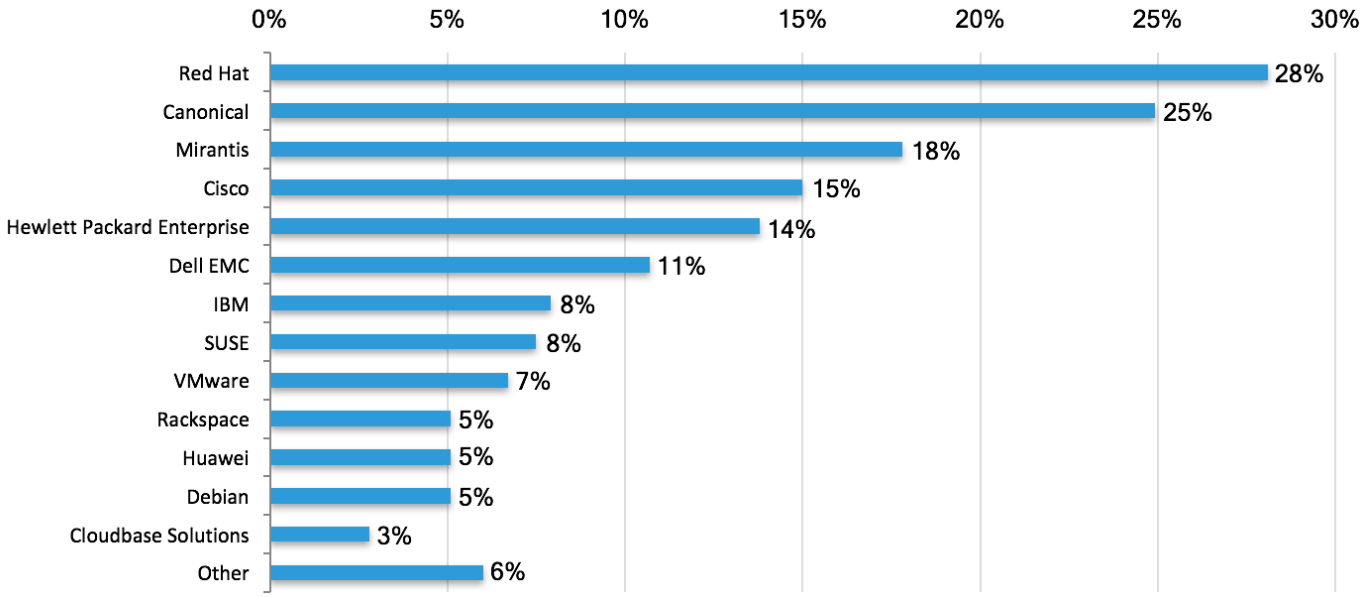

Figure 5.15 n=253

Which operating systems are running OpenStack deployments?

Ubuntu Server continues to lead among operating systems. Its use was down 22 points compared to last year, but still higher than the 45% of users who indicated Ubuntu in the October 2015 survey. Red Hat Enterprise Linux declined 7 points, while others remained fairly stable.

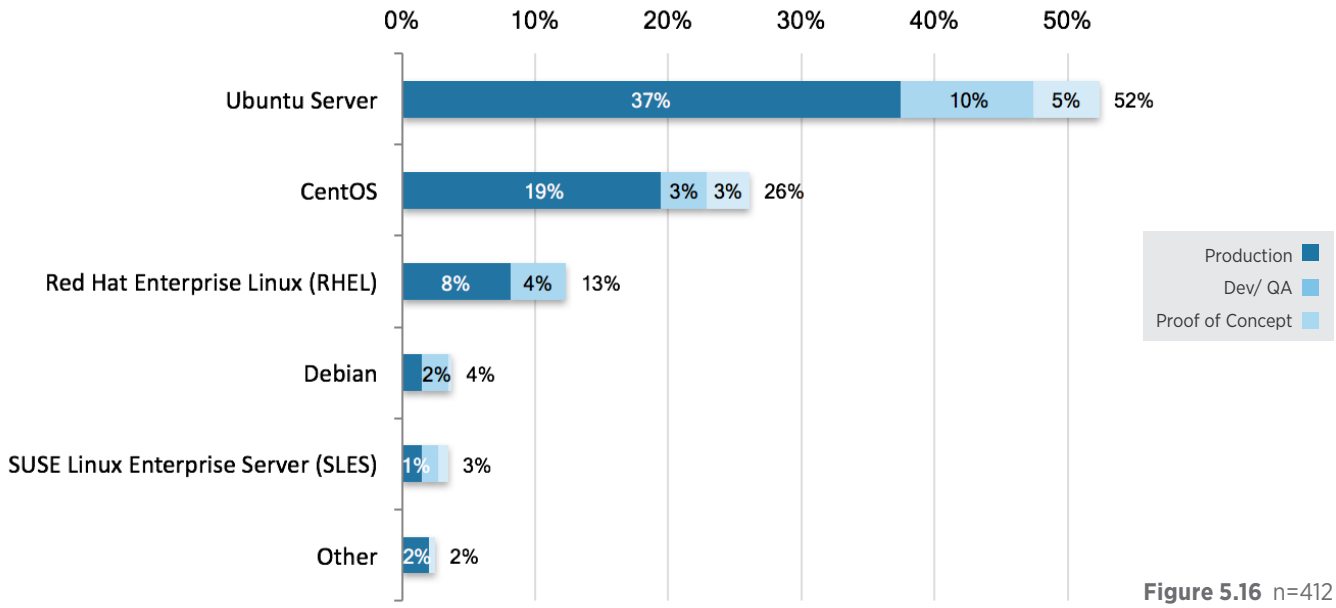

Figure 5.16 n=412

We also looked at the answers from only large users (1,000 cores or more), finding substantially more CentOS use (an increase of 18 points), and less Red Hat Enterprise Linux use (down 11 points) compared to last year.

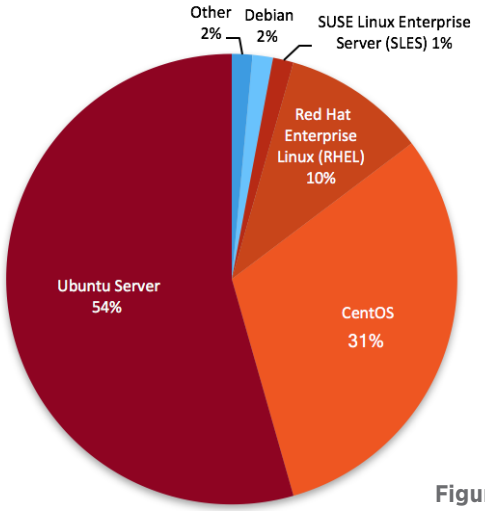

Figure 5.16.2 n=68

Which compatibility APIs does this deployment support?

In the April 2016 survey, "this deployment doesn't use compatibility APIs" was not an answer option, leading many to skip this question (n=139 in April 2016, compared to n=300 in April 2017).

As a result, in this cycle we know that 61% of clouds support compatibility APIs, however we are unable to compare this data set's percentages accurately to last year's. Looking at the data generally, EC2 remains the top API compared to last year, followed by S3, and OCCI and GCE remain just a fraction of the leaders.

In OpenStack's October 2016 Newton release, the EC2 compatibility API was removed from Nova, replaced with an external project that provides the same functionality. This could increase the difficulty of providing EC2, so this statistic could have interesting changes in future surveys.

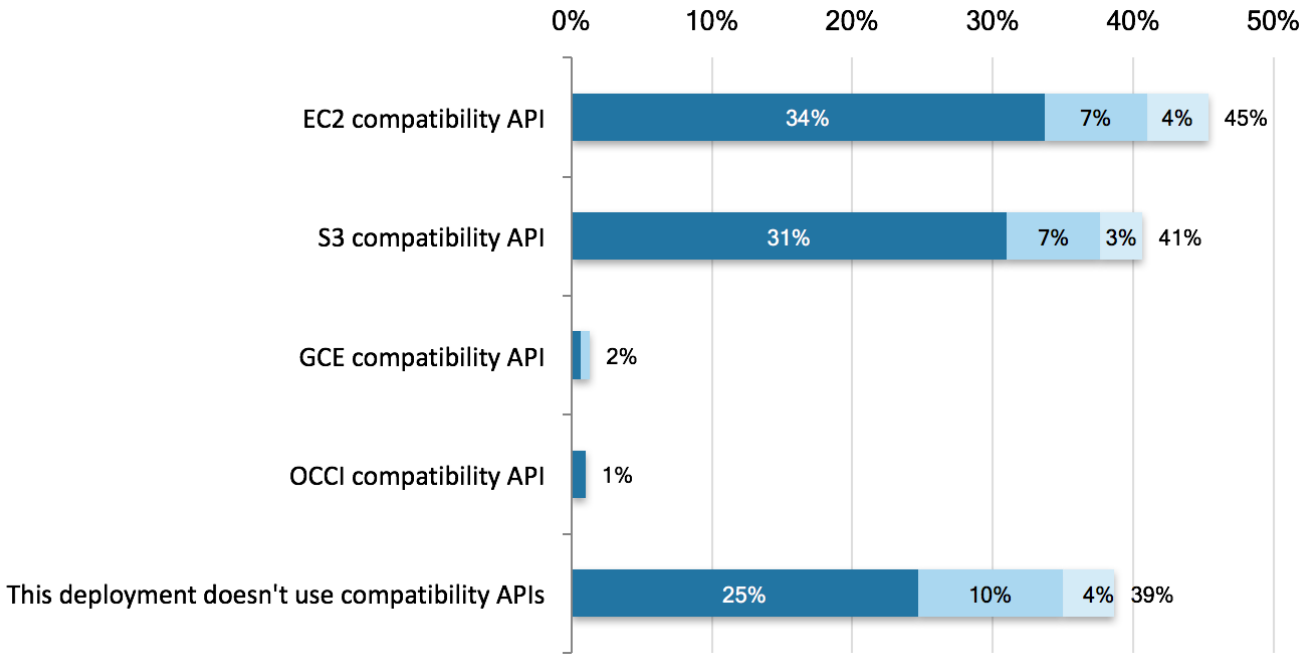

Figure 5.17 n=300

Part Six

PROJECT-SPECIFIC QUESTIONS

This section includes questions supplied by the Project Team Leads, who can submit one question that will be served to users who indicate in their deployments that they are using this project.

The results of these questions were provided to PTLs just prior to the Project Teams Gathering at the end of February 2017 in Atlanta, to give development teams data and context for the decisions they were making.

In some cases, PTLs elected to ask qualitative questions; these are noted at the end of this section and the full answer set was also shared with PTLs.

What do you consider to be the biggest challenges to achieving interoperability between OpenStack clouds?

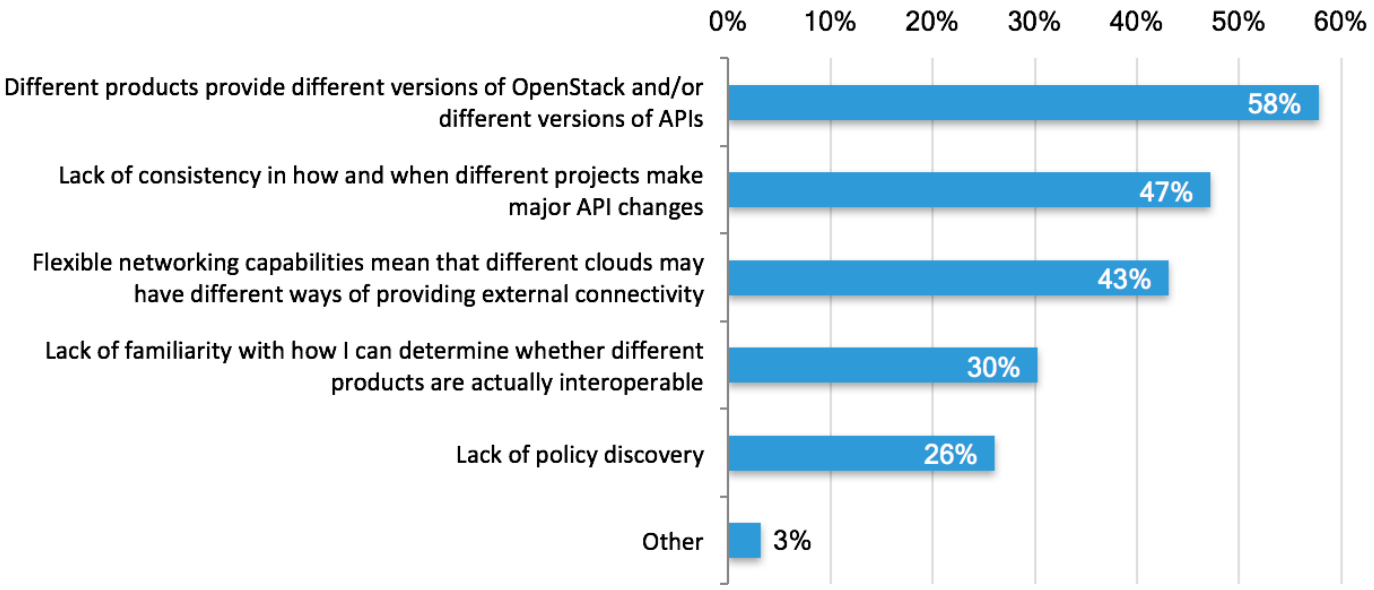

Figure 6.2 n=218

What's the current/expected load on your Heat deployment?

Figure 6.1 n=83

In your opinion, where should the Keystone development team focus their effort(s)?

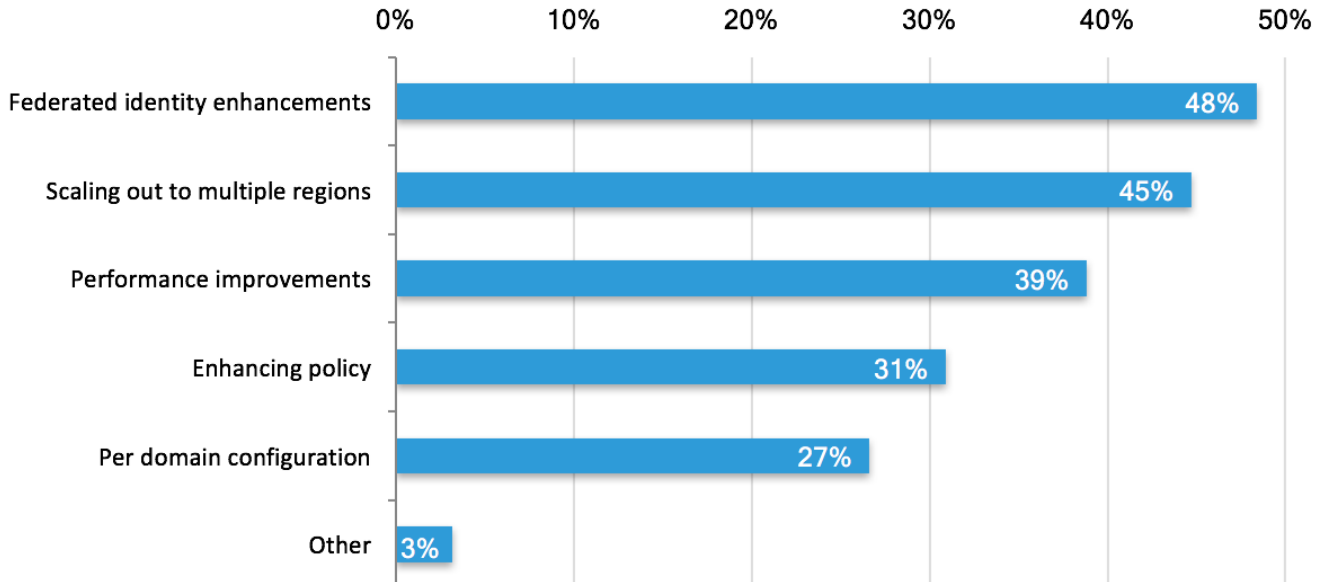

Figure 6.3 n=188

How do you use Kolla?

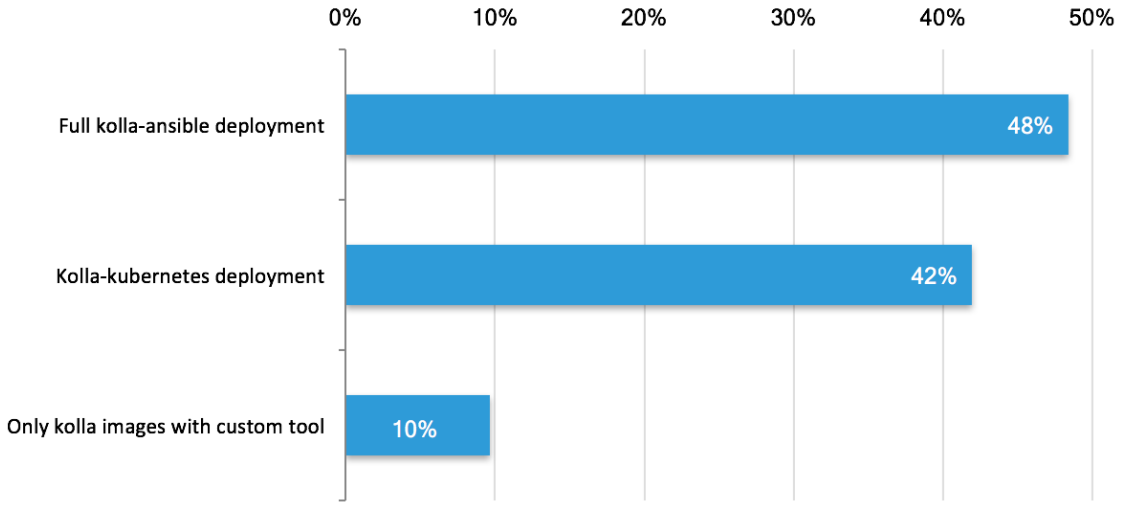

Figure 6.4 n=31

Which OpenStack Shared File Systems (Manila) driver(s) are you using?

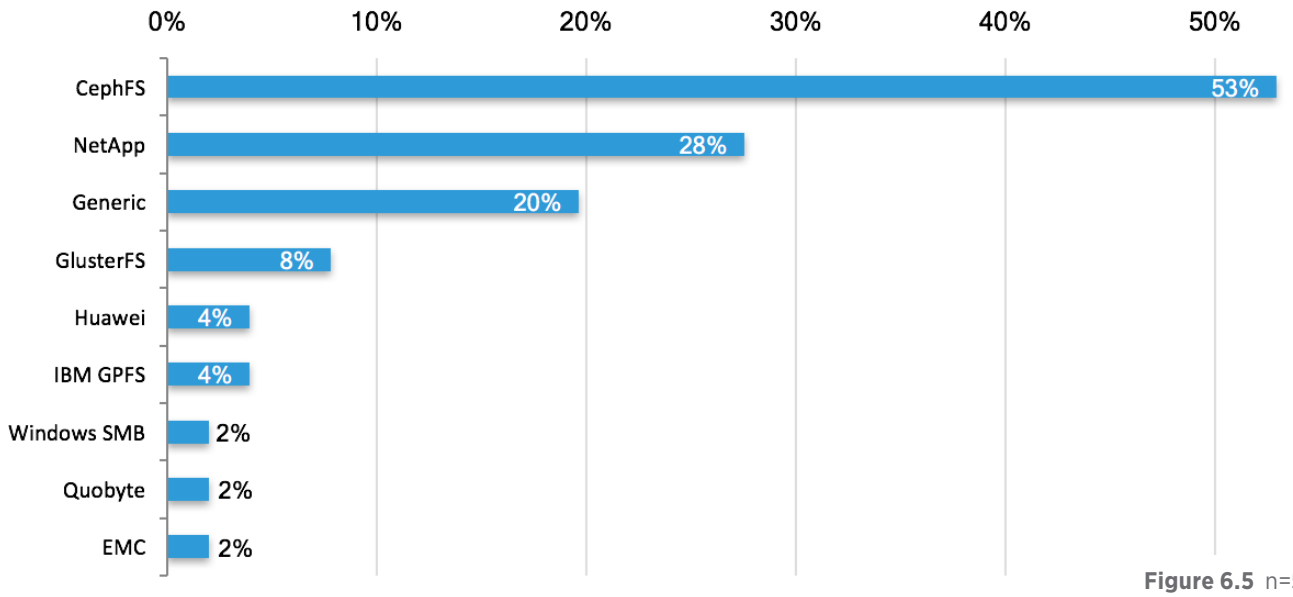

Figure 6.5 n=51

Which features in Neutron are you actively using, interested in using, or looking forward to using?

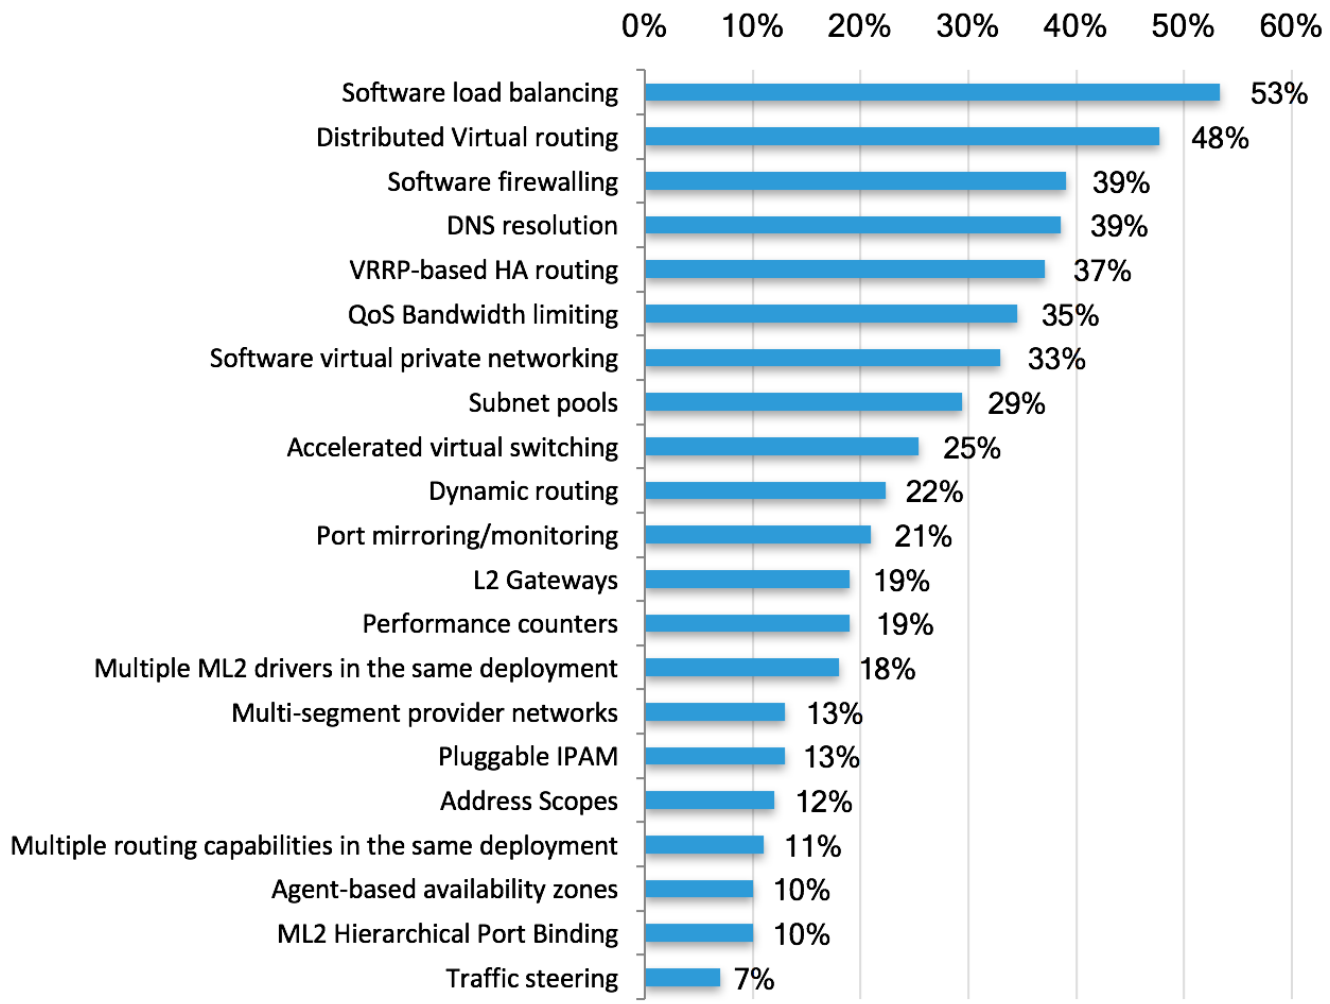

Figure 6.6 n=197

How important is it to be able to customize Nova in your deployment? e.g. classload your own managers/drivers, use hooks, plug in API extensions, etc.

51% Not important: I use pretty much stock Nova with maybe some small patches or bug fixes that aren't upstream.

39% Somewhat important: I have some custom scheduler filters and other small patches but nothing major.

10% Very important: My Nova deployment is heavily customized and hooks/plugins/custom APIs are a major part of my operation.

Figure 6.7 n=221

Which of the following features in RefStack are important to you?

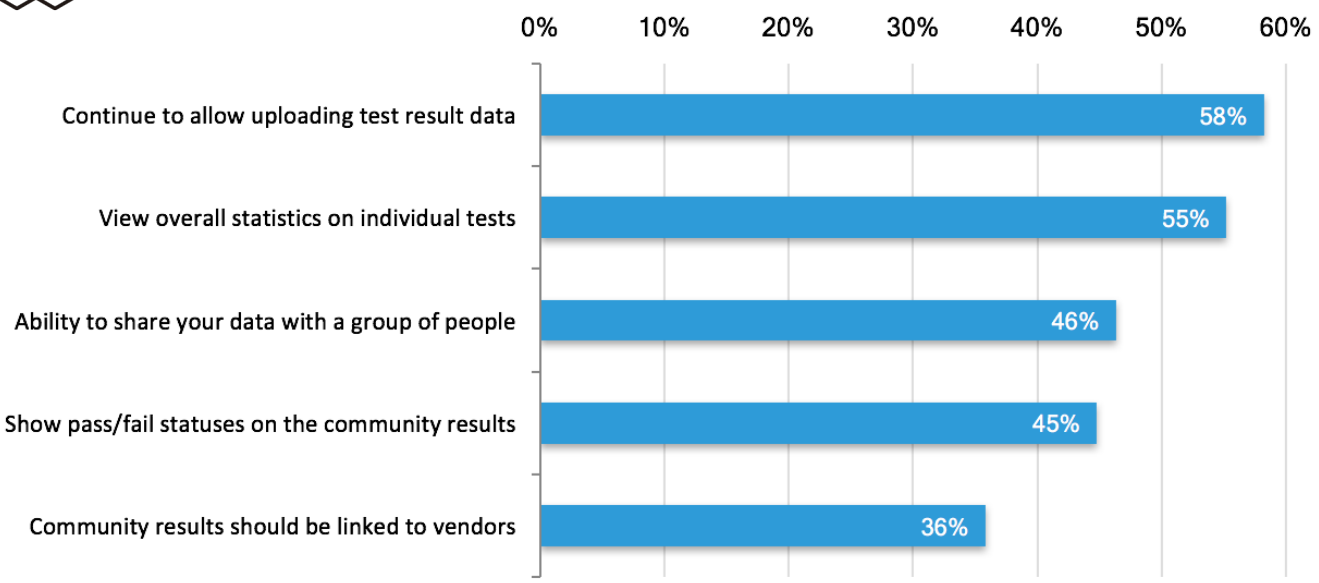

Figure 6.8 n=67

Swift: What kinds of data are you planning to store on object storage in the next 12 months?

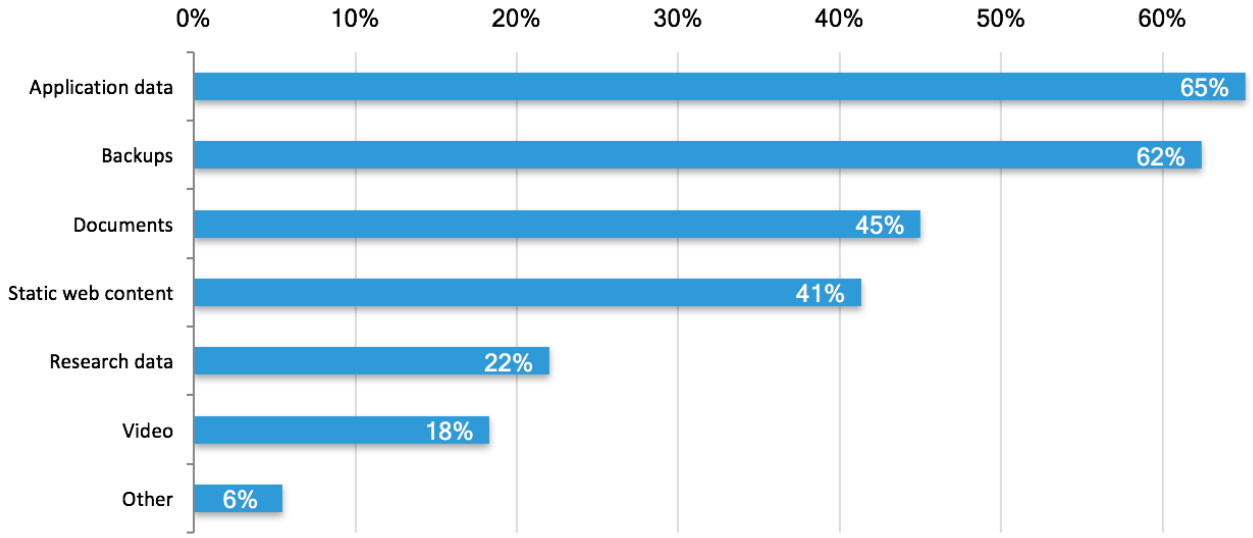

Figure 6.9 n=109

Which dispatcher are you using or planning to use with Ceilometer?

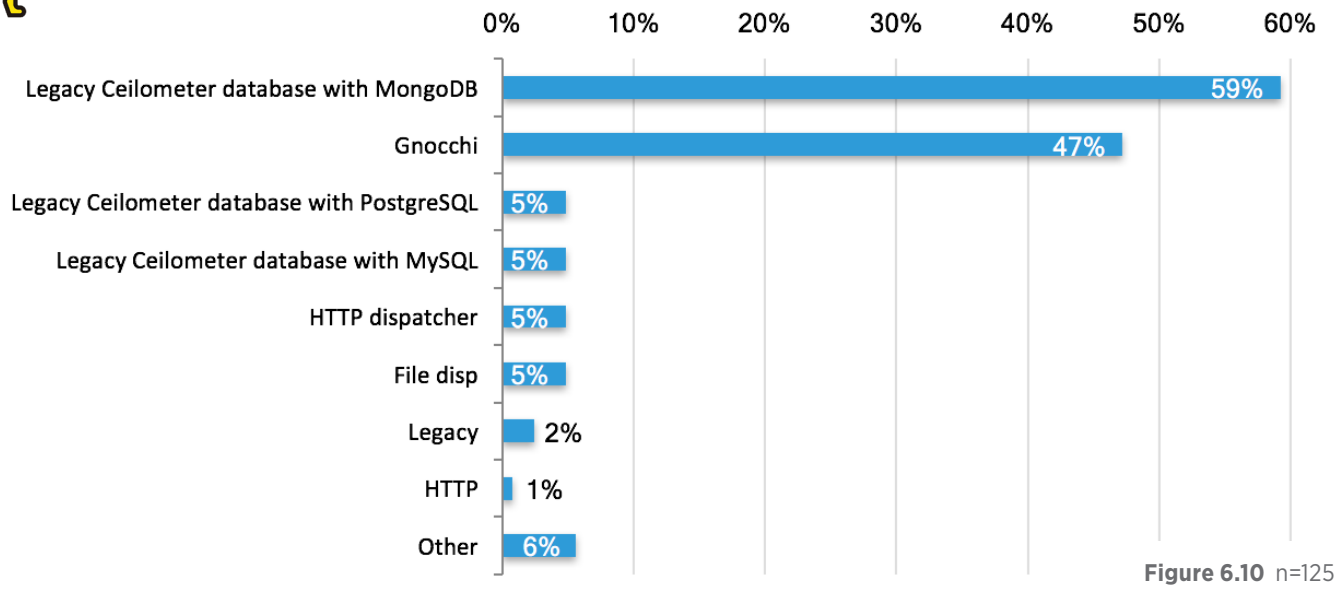

Figure 6.10 n=125

TripleO: As a user/operator, do you value having a common interface for deployment tooling and the deployed cloud? (e.g common APIs and CLI tools)

Figure 6.11 n=51

When do you expect to begin evaluating Database as a Service in your cloud using the OpenStack Trove project?

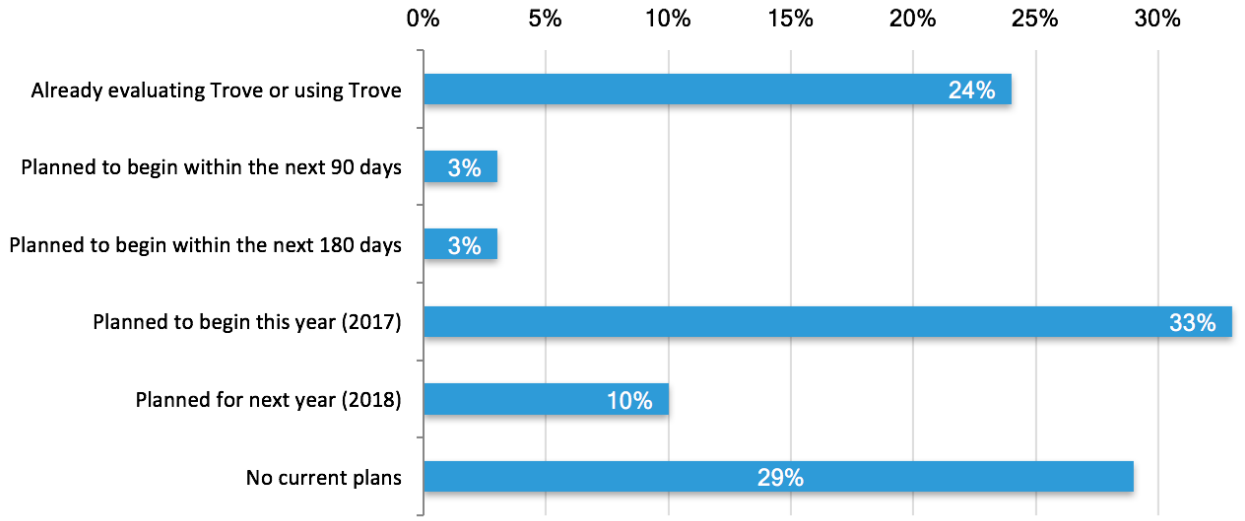

Figure 6.12 n=100

More questions—and answers—for teams to consider...

Some PTLs chose to ask qualitative questions. While the response rates were generally lower, the feedback to the teams offered rich detail. Please contact PTLs directly for more detail. Questions included:

BARBICAN: What is missing from Barbican that might prevent your adoption?

CINDER: If there was one thing you would like to see changed (added, removed, fixed) in Cinder, what would it be?

DESIGNATE: What (if any) features or issues are preventing you from using Designate as a production service?

GLANCE: As you're aware, the Images API v1 supplied by OpenStack Glance is currently deprecated. If you haven't yet moved to the Images API v2, what is preventing you? Please be specific.

INTEROP: Do you use the OpenStack Powered program or Interoperability Guideline compliance as a factor in your OpenStack procurement process? Why or why not?

IRONIC: What's been your most frustrating or difficult experience with Ironic?

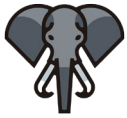

SAHARA: What else can be improved in the Sahara project?

METHODOLOGY

For the fourth survey in a row, the User Committee and OpenStack Foundation partnered with an external independent data scientist to help analyze and report the data.

The User Survey Working Group reviewed data and reporting at each stage of the report creation process, including comment analysis, draft charts/ report, and final report.

This cycle's survey changed some answer sets based on community feedback to reflect our evolving technology choices. Several new questions were added to the survey, and they are noted in the survey text.

The "app developer" section, which last year was served to those who indicated that they create applications on OpenStack, was eliminated. Two of the four questions from that section were moved into the deployments section, and the Foundation is exploring launching an app developer-specific survey in the future.

REPORTING PERIOD

In this report, we primarily looked at trends from answers submitted:

- Jan. 26 - March 1, 2017: survey 2017-01 (Ocata cycle); denoted as the **April 2017** survey
- Jan. 21 – Feb. 29, 2016: survey 2016-01 (Mitaka cycle); denoted as the **April 2016** survey

This 12-month timeframe enables us to see larger trends and compare on an apples-to-apples basis, because the October 2016 survey had a much smaller sample size, limited questions, no qualitative feedback requested, and was targeted toward existing deployments, simply asking users to update their deployment data.

Anyone who completed or updated their deployment in the October 2016 survey cycle had much of their deployment information pre-populated, however, if they had not filled out the user survey since April 2016, they had to re-enter their information.

METHODOLOGY

As in prior surveys, survey logic showed some users certain questions, based on their responses to prior questions. 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. As a result, these questions have a much smaller number of responses. Few questions were required, resulting in variable *n* numbers for each chart.

In addition to quantitative data, we gathered a substantial amount of qualitative data from open-ended questions. The User Survey Working Group includes volunteers from the user committee and community, bound by a confidentiality agreement, who selected representative comments and identified common themes to add more insights to numerical results.

The project matrix of most commonly deployed projects included fewer than half of the 60+ OpenStack Big Tent projects, focusing on core services and function-based projects rather than projects of general use such as Documentation, Release Management, and I18n (translation). We intend to revisit adding more function-based projects to this list on a future survey.

Project Team Leaders (PTLs) were offered an opportunity to submit one question regarding their project to the user survey, in either a select-one, select-all, or short answer field format. PTLs from 17 projects submitted questions. Questions supplied by the PTLs changed substantially compared to the prior survey, and therefore cannot be compared to a prior cycle's data set.

SURVEY INSIGHTS

Learn more about users' technology choices and point of view in the ninth edition of the OpenStack User Survey.

More than 1,300 users and nearly 600 deployments are featured in this report.

Analysis includes demographics, user perspectives, deployments, cloud size, technology choices, drivers, vendors, and project-specific questions.

Hear directly from users with dozens of verbatim comments and critiques.

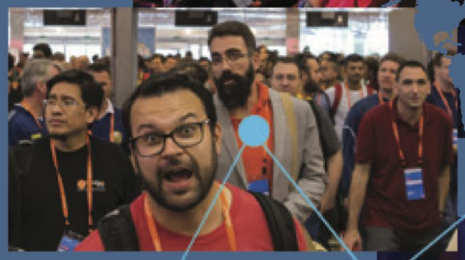

7,845