

2018 OpenStack User Survey Report

Executive Summary

The 11th OpenStack User Survey highlights the latest feedback and technology choices of users globally. This year, nearly 900 user deployments offer a snapshot of the project that reflects happier users, the growth of bare metal deployments and Kubernetes, adoption growing beyond IT organizations, and OpenStack as an important component of users' multi-cloud strategies.

Kubernetes remains the No. 1 container orchestration framework running on OpenStack

After taking the lead in 2016, Kubernetes remains the top framework among container and PaaS tools used to manage OpenStack applications. OpenStack and Kubernetes are increasingly deployed as complementary technologies, especially in multi-cloud scenarios.

OpenStack bare metal clouds growing, fueled by Kubernetes adoption

Adoption of Ironic, the bare metal OpenStack cloud service, has grown rapidly with 24% of production deployments now relying on it, up from just 9% in 2016. Among users who run Kubernetes on OpenStack, adoption of Ironic is even stronger with 37% relying on it. Magnum, the CNCF certified installer for Kubernetes on OpenStack, is also widely used among this population, with 16% using it.

Multi-cloud OpenStack users rely on OpenStack

Although OpenStack users often rely on proprietary public clouds like Amazon Web Services (AWS) (44%), Microsoft Azure (28%) or Google Compute Engine (GCP) (24%), the average respondent said that 58% of their infrastructure was powered by OpenStack.

The OpenStack user base continues to expand beyond IT organizations

OpenStack's broad usage among industries beyond information technology (IT) was demonstrated again in 2018, as it was in the previous year, by a significant increase in user responses among non-IT industries, particularly financial services (now comprising 8% of deployments), government/defense (6%) and retail/e-commerce (4%).

OpenStack NPS score reflects happier users

The NPS score of users with deployments rose significantly in 2018 to 41 after being reported as 25 in 2017. After making adjustments in the translated User Survey to clarify confusion around the NPS questions, the score has returned to a level consistent with scores reported in 2016.

Report Methodology

The purpose of the OpenStack User Survey is to gain a better understanding of attitudes, organizational profiles, use cases and technology choices across users' various deployment stages and sizes. 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.

OpenStack's 11th User Survey Report is a snapshot of data collected between August 2017 and August 2018. The data set comprises 1,483 completed surveys representing 858 deployments and 441 unique user organizations.

The 2018 OpenStack User Survey report methodology varies slightly from past reports: Every metric reported represents only those respondents who indicated they have an OpenStack deployment. If a user took the survey but did not report a deployment, their responses are not included in the analysis of this report. Furthermore, not all respondents with deployments answered all survey questions, so the N number (population size) varies by question. The anonymous verbatim comments from all respondents are shared with project team leads.

The OpenStack User Survey is an opt-in survey of OpenStack users around the globe and is available in eight languages: Simplified Chinese, Traditional Chinese, English, French, German, Indonesian, Japanese and Korean. Participants are encouraged through OpenStack community channels to complete the survey. The User Survey is not a market survey and does not include data from 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.

Analyze the User Survey Yourself

Dig deeper into the data at the User Survey Analytics Dashboard available at <http://www.openstack.org/analytics>. The dashboard spans four data sets (2015, 2016, 2017 and 2018).

The User Survey is always open, and the dashboard captures live data as survey responses are logged. Whereas this survey report is based on a snapshot of data collected between August 2017 and August 2018, the online dashboard includes surveys submitted before and after that time frame. As a result, dashboard results sometimes vary from the results in this report due to the difference in data collection periods.

Demographics

Survey Snapshot

- Data set time frame: August 2017 through August 2018
- Deployments: 858
- Unique organizations represented: 441
- Respondents with more than one deployment: 115

Survey Respondents

Nearly 60% of survey respondents with deployments identified their role as either a cloud operator or a cloud architect. However, 46% of respondents reported having more than one role in their organization.

Industries Using Openstack

The users with deployments who reported being a part of the information technology (IT) industry decreased significantly for the second year in a row. The percentage of deployments in the IT industry dropped from 51% in 2017 to 44% in 2018.

Although IT remains the largest industry category overall—representing 44% of deployments—the OpenStack user base continues to expand in non-IT industries. The largest increase was witnessed in financial services (now comprising 8% of deployments), with notable gains seen in government/defense (6%) and retail/e-commerce (4%) as well.

This trend continues to be reflected in organizations publicly sharing their OpenStack use case in 2018 including Banco Multiva, eBay Classifieds Group, Ocado Technology, Progressive Insurance, SBAB Bank, Target and Volkswagen Financial Services.



Where in the World Are OpenStack Deployments?

Respondents to the 2018 User Survey represent 63 countries. The number of survey respondents from Asia witnessed a significant surge for the second year in a row, leading to Asia's ranking as the top region for 2018 survey respondents at 48%. Translation continues to benefit users in Asia, as over 30% of surveys were completed in Chinese, Korean or Japanese. European users ranked 2nd in terms of survey participation, with 26% of total respondents with deployments, while North American respondents claimed the third spot with 20%.



OpenStack Is Used by Organizations of All Sizes

OpenStack is serving as a strong infrastructure solution for both developing companies and established enterprises and is proven at enterprise scale. A diagram of the organizational size of OpenStack users roughly approximates a bell curve and aptly represents the relatively normal distribution of community respondents across organizations of all sizes. Roughly half of survey respondents are employed by organizations with 999 employees or less.



"OpenStack, with its whole cloud ecosystem and community, is a complete solution for IaaS for any size company."

Pawel Stefanski - Mirantis

User Perspectives

Efficiency and Innovation are Top Drivers of OpenStack Adoption

OpenStack users were asked to select their top reasons for choosing OpenStack and to rank these adoption drivers in terms of priority. "Increasing operational efficiency" and "accelerating the ability to innovate" top the list of reasons why organizations choose OpenStack as their cloud infrastructure platform. Other business drivers cited by more than 75% of survey respondents include "avoiding vendor lock-in," "standardizing on the same open platform and APIs that power a global network of public and private clouds," and "saving money." Additional business drivers include "achieving security/privacy goals" and "attracting top technical talent."



"OpenStack is—more than ever—a rich ecosystem of complementary services which together empower you to manage your datacenter in a highly flexible and efficient way."

James Penick - Oath

How Likely Are OpenStack Users to Recommend OpenStack?

The Net Promoter Score (NPS) for the 2018 User Survey indicates OpenStack users are happier with OpenStack than they were last year and more likely to recommend it to others.* The NPS score among deployments in 2018 (41) has increased

significantly from 2017 (25) and remains consistent with 2016 (40).

** In the 2017 report, it was noted that translation may have had an impact on the low 2017 NPS score. In the 2018 User Survey, context was provided for participants to further explain the question. This explanation was subsequently translated and may have had an impact on the NPS score increase.*

I believe in the power and opportunities OpenStack provides to me. In my opinion it is important to have an internal alternative and/or addition to public cloud providers. I assume that a hybrid cloud approach will combine the best of both worlds.

Thomas Lunkwitz - METRONOM GmbH

Which areas require further enhancement?

The No. 1 area that 2018 User Survey respondents cited as requiring further enhancement was documentation. However, documentation also appears in the list of what users like best about OpenStack. Interestingly, users who consult documentation assign a higher NPS satisfaction score to OpenStack than those who do not.

Upgrades and user experience essentially tied for the second most frequently named area in need of improvement. This has created an opportunity for users like OpenStack public cloud provider OVH and Oath (formerly Yahoo!) to share their successful upgrade stories with fellow operators, including sessions and a workshop at the OpenStack Summit Berlin.

More than a dozen respondents also cited networking/Neutron, container integration and ease of deployment as areas the community should focus on improving.

Containers, Bare Metal, Hardware Accelerators Intrigue OpenStack Users

OpenStack adapts and grows with emerging technology, making it an innovation engine for companies as their platforms develop. This year containers and bare metal are emerging technologies of particular interest to OpenStack users, with containers commanding a tremendous lead as the No. 1 emerging technology on the list, cited by 70% of respondents. The closest followers are bare metal at 47% and hardware accelerators at 42%.



“OpenStack is very flexible and helps us to fulfill requirements from multiple use cases within the company. It is very stable and provides a great community that is always willing to support you in any way—from features development to bug fixing.”

Edgar Magana - Workday

Deployment Profiles

Stage of Deployment

OpenStack’s maturity is evident in the stages of deployment reported by users. OpenStack deployments in production saw an increase from 63% in 2017 to 66% in 2018.

“OpenStack is the premiere IaaS environment of our time based on open source fundamentals. OpenStack as a technology has now reached enterprise maturity.”

Bruce Mathews - Mirantis

75% of OpenStack Clouds Are On-Premises Private Cloud

Respondents report that nearly three-fourths of OpenStack deployments are on-premises private clouds. (Note: Because the OpenStack User Survey is opt-in and promoted through OpenStack community channels, it primarily targets users who are directly engaged in the OpenStack community and likely operating OpenStack clouds, which may not reflect public cloud end users.)

While the number of OpenStack public cloud providers who completed the User Survey decreased in 2018, new datacenters were launched by multiple providers including OVH in Strasbourg, Sydney, Singapore, Warsaw and Vint Hill, Virginia; ScaleUp Technologies in Berlin; and VEXXHOST in Santa Clara.



Kubernetes Dominates as Most Popular PaaS/Container Framework Used to Manage Applications on OpenStack

In 2018, Kubernetes claims a substantial lead as the most popular PaaS or container tool used to manage applications on OpenStack deployments, with 61% of respondents naming it, including BBVA, CERN, China UnionPay, Volkswagen AG and Walmart Labs.

Interestingly, 70% of the respondents who said that they are running Kubernetes to manage applications on OpenStack also indicated they are running a multi-cloud environment.

In comparison, the next most popular tool cited (other than “built-our-own”) was OpenShift (which is Kubernetes-based), cited by 14% of respondents.



Multi-Cloud Users Rely on OpenStack

The results of two survey questions when taken in concert underscore the predominance of multi-cloud environments among OpenStack users and their strong reliance on OpenStack to manage the majority of their cloud infrastructure.

Seventy-eight percent (78%) of OpenStack users who answered this question (“With which other clouds do your users interact?”) identified that they are in a multi-cloud environment, with AWS being the leading option at 44%.

Moreover, User Survey participants are also asked what percentage of their overall cloud infrastructure is running on OpenStack. A majority of users responding to this question indicated they are running more than half of their cloud infrastructure on OpenStack.



Project Profiles

Growing Use of Containers on Bare Metal Reflected in Project Adoption

Within the deployments that are using Kubernetes to manage applications on OpenStack, there are notable increases in adoption of certain OpenStack projects, such as Ironic, as users turn to OpenStack to run container orchestration on bare metal. Thirty-seven percent (37%) of OpenStack deployments that indicated they are running Kubernetes to manage applications indicated they are running Ironic in production. In comparison, use of Ironic among all OpenStack deployments in general is 23%. Other projects that saw a significant increase in adoption among Kubernetes users include Magnum, Kolla, Manila, Sahara and Trove.

Among all deployments in production, OpenStack Ironic has been one of the fastest growing projects, going from 9% of deployments running it in production in 2016 to 23% in 2018.

“The many ways that Openstack can be deployed and the multitude of vendor support are great. The ability to integrate dedicated load balancers, containers, databases and more is a huge plus.”

Joss Malasuk - Rackspace

What does the typical OpenStack deployment look like?

This OpenStack deployment profile is based on the most popular response choices to the deployment decisions section of the 2018 User Survey. This data has remained consistent from past surveys, but doesn’t represent the breadth of technology choices that OpenStack users are implementing in their deployment. Find all of the data and response choices at <http://www.openstack.org/analytics>.

Hypervisor: KVM
Cinder Block Storage Driver: Ceph
Identity Service Driver: SQL
Operating System: CentOS
Configuration Management Tool: Ansible

Vendor Snapshot

Diverse Ecosystem Powers OpenStack Clouds

To create a better understanding of the ecosystem supporting OpenStack users, the 2018 User Survey changed the question used in previous surveys (What vendor’s products power your OpenStack clouds?) into two separate questions: one asking which software vendors power their OpenStack clouds and one asking users to identify the hardware vendors that power their cloud. The results below are from respondents with deployments and, as explained in the methodology section above, do not constitute market share data. Furthermore, vendors in the OpenStack ecosystem collaborate with the OpenStack Foundation to promote the user survey among their customers, which may affect the survey results with respect to distribution of users among OpenStack vendors.

The OpenStack User Survey is intended to gather feedback from users about their experience with the platform, including their use cases, scale and technology choices. It is an opt-in questionnaire and should not be interpreted as a market share report. The list of OpenStack vendors is not intended to be a ranking of the market. It is meant to provide transparency about the profile of respondents, which are a subset of the total population of OpenStack users.



“[What I like best about OpenStack is] the support it receives from all the vendors without making it vendor locked.”

Arminster Singh Girgla - Pugmarks InterCloud LLP



Survey Credits

Thank you to community volunteers who helped analyze the response data and review the report. [The OpenStack User Committee](#) includes Melvin Hillsman, Yih Leong Sun, Amy Marrich, Joseph Sandoval and Matt Van Winkle.

Jimmy McArthur and Allison Price from the OpenStack Foundation led creation of the survey report with support from Jonathan Bryce, James Cole, Mark Collier, Lauren Sell and Wes Wilson. The OpenStack Foundation partnered with independent data scientist Kelly Valade to analyze and chart the data as well as Jennifer Fowler and Robert Cathey of Cathey Communications on developing the report. Thank you to Santiago Palenque and Sebastian Marcet of Tipit for their support in creating the custom toolset for both the survey and the analytics dashboard.

Finally, a very special thank you to the entire [OpenStack Internationalization \(I18n\)](#) team. We now present the survey in over seven languages, and this would not be possible without their hard work and diligence.