OPENSTACK

2020 Open Infrastructure Foundation Annual Report

Contents

1. Chairman’s Note
2. A letter from the Executive Director
3. Introducing the Open Infrastructure Foundation
4. OpenInfra Foundation Project Updates
5. Airship
6. Kata Containers
7. OpenInfra Labs
8. Section 8
9. Section 9
10. OpenStack
11. Section 11
12. Section 12
13. Section 13
14. Section 14
15. Section 15
16. Section 16
17. Section 17
18. Section 18
19. Section 19
20. Section 20
21. Section 21
22. Section 22
23. Section 23
24. 10 Years of OpenStack
25. Section 25
26. Superuser
27. 10 Years of OpenStack
28. COA
29. Ironic Whitepaper
30. StarlingX
31. Zuul
32. OpenInfra Foundation Collaboration
33. Magma
34. OpenDev

35. OpenInfra Foundation Working Groups
36. Diversity & Inclusion Working Group
37. Interop Working Group
38. Open Infrastructure Community Highlights
39. OpenInfra User Groups Growth
40. Outreachy Program
41. Students in OpenStack
42. Section 42
43. Section 43
44. Section 44
45. OpenStack Upstream Institute
46. Press & Analyst Highlights
47. Superuser
48. Event Highlights
49. Open Infrastructure Summit
50. OpenInfra TV: Bringing the Summit to Asia
51. OpenDev
52. PTGs / Forums
53. OpenStack Days / OpenInfra Days / Cloud Operator Days

Chairman's Note

We feel a universal spirit of hope as 2021 begins. It is of no doubt to anyone that 2020 will be entered in history books as a tumultuous year. A year that has impacted us as individuals, the open infrastructure Foundation, our open source efforts and community. Though we will all recall the pandemic impacts, for me, 2020 will always be remembered as a pivotal beginning. It will be recorded in my journals as the year that launched the Open Infrastructure Foundation.

An author once wrote, "Sometimes when you're in a dark place you think you've been buried, but actually you've been planted". Though at times, 2020 felt pretty dark, together we planted the seeds for the Open Infrastructure Foundation. Those seeds have now sprouted. And with good community roots, those sprouts will be prosperous in 2021 and the decade. I am bullish for the future of open Infrastructure. I am excited to see how the technology grows and evolves. And I am excited to see the fruits of that growth being adopted, deployed and used throughout the world.

One of the foundational strengths of our community is from the committed leadership, not just with the OpenInfra Board, but also with the technical steering, project and group leads. That's the strength of our roots. That's the root source of our vision. That's the promise of our future.

It's with that thought, I wish to express gratitude to the hard working 2020 members of the OpenInfra board. Thank you, Allison Randal, Amy Marrich, Arkady Kanevsky, Bill Renxudong , Chris Price, Clemens Hardewig, Daniel Becker, David Krawczynski , Fred Li , Jeff Moyer, Johan Christenson, Julia Kreger, JunWei Liu, Kurt Garloff, Li Kai, Mark Skarpness, Mohammed Naser, Prakash Ramchandran, Ryan Beisner, Ryan Van Wyk, Ruan HE, Shane Wang , Shannon McFarland ,Tim Bell, Tytus Kurek , Vijoy Pandey and Xin Zhong for your tireless efforts, continued commitment to the success of the Foundation and for your visionary leadership.

Let's cheer on the new year and to the successful growth of the seeds planted over these past years.

With all success, Alan Clark OpenInfra Board Chair, 2020

A letter from the Executive Director

Over the last few years, we've been steering our Foundation through a series of transitions to align with the evolution of our community and the changing landscape of open source and cloud. This started in 2017 with an evaluation of strategic focus areas where we wanted to play a part in developing open source solutions. In 2018, we developed our pilot project program with Kata Containers, Zuul, Airship and StarlingX. These projects went through the confirmation process with the Board in 2019 and demonstrated their success in building communities that were delivering software consistently through open development.

Then, in 2020 it was encouraging to see the Board, staff and community come together to deliver a successful and pivotal year for the Foundation despite the challenges of a global pandemic. We came together and persevered to continue the transition, officially becoming the Open Infrastructure Foundation in October 2020.

While the global community was not able to meet in person, everyone still convened to discuss how we collaborate to deliver the OpenInfra mission: building communities who write software that runs in production. We saw users who have been around for years talk about their growing OpenStack production deployments like Workday, LINE, and CERN, as well as new users talk about how they're running other open source projects in production like Ant Group using Kata Containers, Verizon running StarlingX, and Volvo running Zuul for their autonomous driving software.

We also celebrated multiple OpenInfra software releases that continued to develop through online events like the Project Teams Gathering (PTG) and endless asynchronous meetings. Our upstream communities continue to prioritize features enhancing integration with other open source projects, delivering the requirements from operators running production deployments. With OpenStack Ussuri and Victoria, StarlingX 4.0, Kata Containers 2.0, and seven releases from the Zuul developer community, users are continuing to see integration improvements with technologies like Kubernetes, Prometheus, QEMU, and Ansible.

The work is not done. Together with the 60 organizations who joined the launch of the OpenInfra Foundation and over 110,000 community members, we are starting the next decade of open infrastructure tackling challenges including hardware diversification, deployment diversification, government regulations like data sovereignty, and the availability of more open source than ever.

I am confident with the learnings of the past ten years and the momentum behind the OpenInfra mission, the next decade will see more innovation, collaboration, and open source software running in production.

Want to be a part of the next decade of open infrastructure? Join us!

Jonathan Bryce, Executive Director, OpenInfra Foundation

Introducing the Open Infrastructure Foundation

A few months ago at our virtual summit, we announced that the OpenStack Foundation was finally becoming the Open Infrastructure Foundation. What's behind the name change? What is Open Infrastructure?

Infrastructure is the underlying resources required for developing an activity. For computing infrastructure, it used to be just physical servers onto which applications would be deployed. But in the last 20 years we added more infrastructure abstractions (virtualization, infrastructure as a service, application deployment APIs...) to facilitate the development and deployment work. And with those abstractions, the role of the infrastructure provider became more important. Open Infrastructure is the idea of providing that infrastructure using collaboratively-developed open source technologies.

Computing infrastructure was traditionally the domain of proprietary software. OpenStack was started in 2010 to break that tradition: collaboratively build an open source infrastructure-as-a-service solution that could be used by private infrastructure providers and public infrastructure providers alike. That was a bold bet, against the largest and most successful companies in the world. A lot of people doubted it would be successful.

And yet, ten years later, OpenStack is undoubtedly a success. Millions of CPU cores of computing infrastructure are provided through OpenStack. More than 70 regions of public clouds powered by OpenStack are available for users worldwide. Last year saw our 22nd (on-time) release, with more than 20,000 changes merged in only 5 months. A level of activity that placed OpenStack among the 3 most active open source development communities in the world.

Beyond OpenStack's success, these last ten years validated the idea of open infrastructure, which so many doubted. Several new open source projects appeared in this space, including Kubernetes which provides a clean interface at the top of the infrastructure for developers to deploy their applications. Cloud-native applications and frameworks, which are designed to run on top of programmable infrastructure, are highly complementary to the open infrastructure idea. Today, open source is the default for infrastructure software.

The OpenStack Foundation was created to develop, support, protect, and promote OpenStack. But by doing so, we crystallized a community of open infrastructure providers, operating this software in production and ready to share their experiences between peers. At our global summits, at our local OpenStack Days, this community came together to share their experiences between peers, operating OpenStack and other open source solutions.

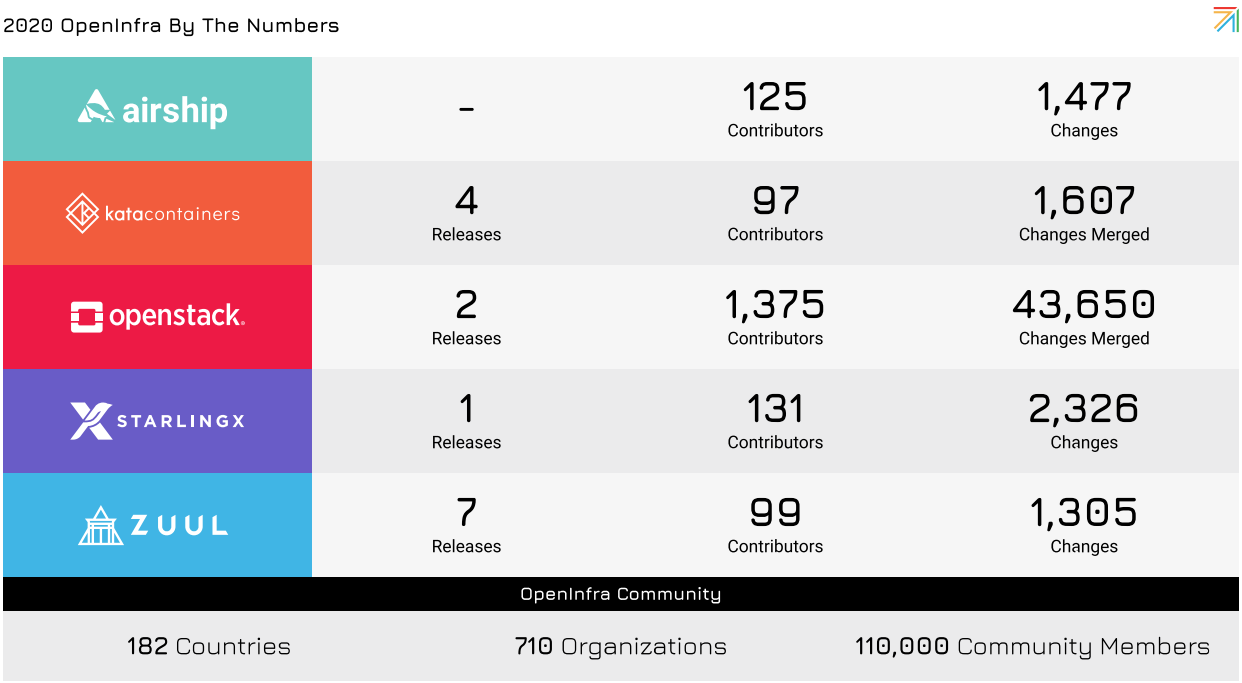
In the next decade, we expect more and more open source usage, especially in the infrastructure area. Edge computing and AI workloads place requirements on infrastructure that are best addressed by open source solutions. New requirements emerge all the time. Climate change raises questions on the sustainability of the centralized datacenter model and drives the emergence of distributed clouds inside cities that can directly heat water or buildings. Data sovereignty pushes for local, interoperable infrastructure. Basic access to connectivity raises new challenges that proprietary solutions cannot address.

In all those new requirements, open infrastructure can and will help. The next decade will see the emergence of several new open infrastructure projects, calibrated to answer tomorrow's challenges. Openly-developed solutions, collaboration between communities, integration between software components, information sharing between infrastructure operators will be more necessary than ever.

And this is why we evolved into the Open Infrastructure Foundation. For the last 10 years, the OpenStack Foundation has been supporting open infrastructure projects, encouraging collaboration without boundaries between communities, organizing events and federating a community of open infrastructure operators. The next decade will see more open source in infrastructure, more need for collaboration and integration… more need for an open source Foundation focused on open infrastructure.

So this is the new journey we are embarking on. More than just a name change, it is our chance to answer the challenges of tomorrow, to support new projects, to welcome new members: a new Foundation to drive the next decade of open infrastructure. Join us !

Thierry Carrez, VP Engineering, OpenInfra Foundation



OpenInfra Foundation Project Updates

Airship 2020 Project Update

Since Airship was announced as an OpenStack Foundation pilot project at the 2018 OpenStack Summit in Vancouver, it has provided operators with a unified, fully-declarative, versatile platform that transforms bare-metal infrastructure into a resilient Kubernetes cluster supporting user-defined workloads. At the Denver Open Infrastructure Summit in 2019, the Airship project committed to overhauling the platform by improving document management, re-imagining the upgrade workflow, and utilizing additional Cloud Native projects to provide Airship functionality with the next generation of Airship: Airship 2.

This year, the Airship community is delighted to announce that it has delivered on its commitment. While Airship 2 remains in beta, development is quickly progressing towards a first-quarter general availability release in 2021. Additionally, Airship 2 has been designated as a Certified Kubernetes Distribution through the Cloud Native Computing Foundation's Software Conformance Program, guaranteeing that Airship provides a consistent installation of Kubernetes, supports the latest Kubernetes versions, and provides a portable cloud-native environment with other Certified Platforms.

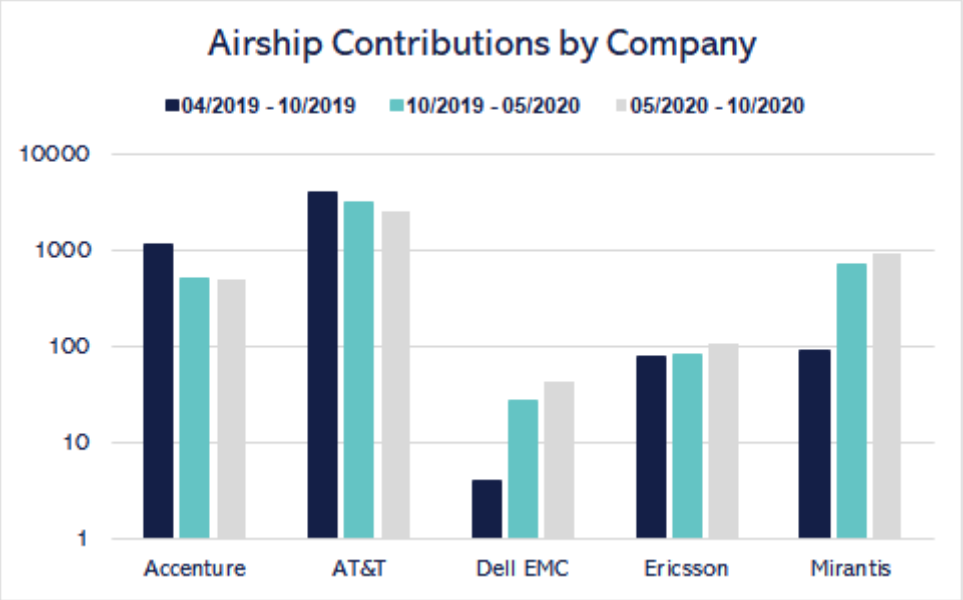
The first change that operators will notice when previewing Airship 2 is its enhanced document management capabilities. Using the Airship 2 command-line-interface, airshipctl, operators can organize and deliver YAML documents that describe an Airship 2 region with phases: logical groups of functionality that are the building blocks of a site. Airshipctl renders phases with

Kustomize, a tool that has been widely adopted by the Kubernetes community. Using Kustomize with airshipctl, operators can keep their YAML footprint small with advanced manipulation tools that reduce data duplication.

Operators familiar with Airship 1 will also notice improvements to the Airship upgrade process. Cloud Native tools such as the Baremetal Operator with Metal3 and Ironic, Kubeadm, and Kustomize have replaced the Airship 1 control plane. In Airship 2, operators can drive upgrades using Airshipctl and let Kubernetes handle the rest—the core Airship components are ephemeral and live outside the control plane.

Airship 2 supports public cloud providers, enabling operators to use the same workflow to manage their workloads on bare-metal, Microsoft Azure, Google Cloud Platform, AWS, and Openstack. Companies with growing requirements can rely on Airship for consistent deployment and management of workloads on Kubernetes, knowing that OpenDev and third-party continuous-integration provided by Ericsson validate these integration points.

2020 also brought greater diversification to the Airship community contributor base. The community elected its second Working and Technical Committees in 2020, adding additional representation from Dell EMC and Insight Global to the existing share of committee members from Accenture, AT&T, Ericsson, and Mirantis. Several companies more than doubled their contributions to Airship. Ericsson's contributions, or commits and reviews, were up 32.5% in the Victoria release from their total in the Train release, Dell EMC 91%, and Mirantis increased their contributions by a whopping 889%.



In February 2020, the Cloud Infrastructure Telco Taskforce (CNTT) chose Airship for their OpenStack reference implementation, providing the global telecom community with a model for deploying NFV workloads using Airship. CNTT seeks to accelerate deployment and management of NFV applications by providing models, or reference implementations, for various telecom use-cases. CNTT includes representatives from over thirty operators, including AT&T, Verizon Wireless, China Mobile, and Deutsche Telekom.

As Airship finalizes the general availability release of Airship 2, operators will see additional NFVi use-cases in Treasuremap. The Airship community revealed two new projects, the Support Infrastructure Provider (SIP) and the Virtual Node Operator (ViNO), at the end of 2020 to provide Kubernetes multi-tenancy in Airship for Containerized Networking Function (CNF) use-cases. In conjunction with Airship, the OpenStack-Helm project will also release OpenStack on Airship 2 using airshipctl phases.

Last year, the Airship community also saw the introduction of a new Vulnerability Management Process where operators can disclose security vulnerabilities in private, the adoption of GitHub issues for release tracking and planning alongside CNCF communities, and the first Airship User Survey. With an upcoming general availability release of Airship 2, the explosive growth of the community, and additional NFVi features, 2021 is shaping up to be another prosperous year for the Airship community.

Get involved:

- Github: https://github.com/airshipit
- IRC: #airshipit
- Mailing Lists: lists.airshipit.org
- Slack: airshipit.slack.com

Kata Containers 2020 Annual Report

Kata Containers seamlessly delivers the speed of containers with the security of virtual machines. Kata Containers became a pilot project in December 2017, in conjunction with the Open Infrastructure Foundation's evolution from being the home for OpenStack to becoming the home of open infrastructure collaboration. In April of 2019, Kata Containers was the first Open Infrastructure Foundation pilot project to graduate, becoming an official open infrastructure project. In 2020, Kata Containers achieved another milestone by releasing the 2.0 release during the Open Infrastructure Summit in October.

Kata Containers 2.0 delivers improved performance and observability enhancements as the community continues to address the challenge of providing secure, lightweight, fast and agile container management technology across stacks and platforms. It is a continuous improvement for Kata Containers to keep moving into the cloud native infrastructure fabric invisibly by reducing overhead and improving security, operability and debuggability.

In 2020, Kata Containers community continues to present its strategic relevance, well-defined governance procedures, commitment to technical best practices and open collaboration, and, most importantly, an actively engaged ecosystem of developers and users.

The project continues to cultivate a global, engaged and growing community as evidenced by the 2020 stats: 4 major releases with 3,016 commits made by 128 authors, representing more than 20 companies.

Kata Containers production deployments continued to expand in 2020, for example:

- In Ant Group, Kata Containers is running on thousands of nodes and over 10,000 CPU cores, and part of the deployment has been upgraded to a 2.0 pre-release version even before Kata 2.0 was officially released;

- IBM Cloud uses Kata Containers to provide a secure environment for their CD Pipeline Service running millions of build containers each month.
- Red Hat intends to provide Kata Containers as a secondary runtime for OpenShift Clusters. This will be enabled through the use of an operator that knows how to install and configure all of the required dependencies needed to run Kata isolated workloads.

The Kata Containers Architecture Committee also went through two election cycles: one in February and another one in September, 2020. Now, the project is led by Archana Shinde (Intel), Eric Ernst (Apple), Fabiano Fidêncio (Red Hat), Samuel Ortiz (Intel), and Xu Wang (Ant Group), . The Architecture Committee members ensure that Kata Containers is kept aligned with its goal of open collaboration and innovation around container speed and security.

The Kata Containers community members collaborated with several communities in 2020. Contributors presented seven related topics at the October Open Infrastructure Virtual Summit. Kata Containers also joined, as a featured project, in the Open Source Hackathon in China. The event attracted developers from several companies, such as Ant Group, Alibaba, Intel and Inspur. At the event, participants openly shared their deployment experiences and adoption plans. Kata Containers community continues to work closely with the OCI (Open Container Initiative) and Kubernetes communities to ensure the project's compatibility.

Looking ahead to 2021, Kata Containers community will focus on improving its integration to the cloud native ecosystem, and continuing to innovate with open collaboration.

Kata Containers' project code is hosted on Github under the Apache 2 license. Learn about the Kata Containers, how to contribute and support the community at katacontainers.io. Join these channels to get involved:

- GitHub: github.com/kata-containers
- Slack: bit.ly/katacontainersslack
- IRC: #kata-dev on Freenode
- Mailing list lists.katacontainers.io

OpenInfra Labs

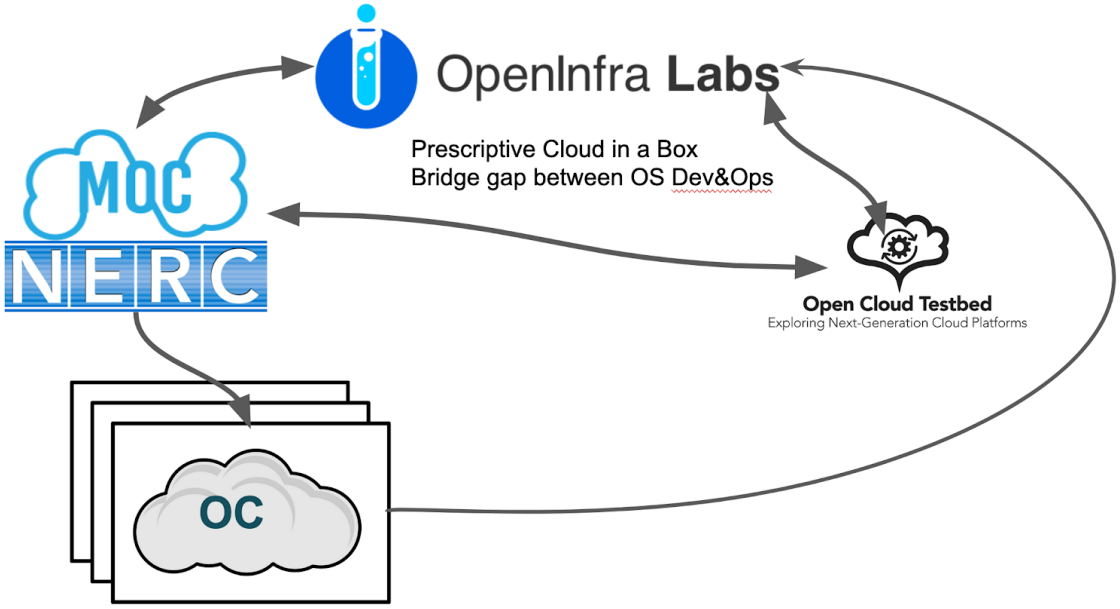
Operators face several challenges with Open Source clouds. Without a standard set of installation steps and services for monitoring, onboarding, offboarding, and billing capabilities, there is no way to encapsulate best practices into configurations or compare information between clouds. Operators frequently require solutions that involve multiple open source projects. Project team developers often have no visibility into how their software is combined with other projects, deployed, operated, or used. As a result, when the operators turn to the community for help it can be difficult for a project team to reproduce or debug problems.

Launched in 2020, the goal of OpenInfra Labs (OI Labs)is to address these challenges by bridging the gap between operators and developers, enabling a community for operators to share their knowledge, providing open source developers real visibility into how their software is used, and creating a path for the open source community to take full advantage of the agility that public clouds have achieved with DevOps. OpenInfra Labs builds on the initial use case of the Mass Open Cloud (MOC), a real-world platform delivering open source projects as services. The MOC is a collaboration between Boston University, Harvard University, MIT, Northeastern University, UMass, industry partners, and the Commonwealth of Massachusetts.

Over the past year, OI Labs has grown from an idea on a whiteboard to a community collaboration featuring a series of connected projects focused on the same goal - to develop tools and spaces to produce results, increase the visibility of requirements within and across open source communities, and accelerate progress in closing the gaps to enable open source projects to be integrated into reproducible clouds that can be federated together.

The efforts associated with OI Labs frequently overlap multiple open source projects - for example:

- Cloud in a Box: Prescriptive Cloud Installations which encompass monitoring, onboarding, offboarding, and billing/reporting.



- Operate First: Closely coupled to OpenInfra Labs is an initiative called Operate First, which adds deployment in a realistic production cloud environment to the workflow of software engineering. By experiencing how projects function, scale, and impact cloud operations, Operate First creates a feedback loop that improves the code and ensures that projects work well with each other. Projects striving to operate first use OpenInfra Labs to develop and refine the techniques for operating first in an open environment. Ensuring that code is cloud-ready is yet another way of improving the state of the art of open cloud operations.
- Project Caerus: Typically, distributed data-parallel compute engines (e.g., Spark and TensorFlow) and distributed storage systems (e.g., S3 and Ceph) are developed and managed independently. The de-coupling of compute and storage has made it difficult to perform end-to-end optimization or to fully exploit hardware resource capacity available in the overall system. Furthermore, compute engines interact with the storage substrate only through standardized APIs that constrain the information exchanged between compute and storage and render their coordination inefficient and expensive, relying on application developers to work out their own optimization strategy.Project Caerus is an initiative focused on bridging the gap between distributed compute and distributed

storage platforms commonly used for big data and AI applications. Caerus aims to create a new open ecosystem that allows compute and storage platforms from different sources to operate in a concerted fashion to substantially improve application performance, resource utilization, and application developer productivity.
- Elastic Secure Infrastructure (ESI): A set of services/systems to permit multiple tenants to flexibly allocate baremetal machines from a pool of available hardware, create networks, attach baremetal nodes and networks, and optionally provision an operating system on those systems through the use of an associated provisioning service. The code development features a mix of upstream OpenStack work (ironic and networking-ansible) and custom ESI code.
- Project Wenju: Addressing the "last mile" challenges of moving AI projects into production.

Other projects focus on improving the process the community uses for creating and maintaining software.

Key Highlights from 2020

- March 2020: Announcement of New England Research Cloud (NERC)project by Boston University and Harvard University at the 2020 Open Cloud Workshop; Introduction of Operate First
- May 2020: Monitoring Working Groupbegins regular meetings; Project Caerusannounced
- July 2020: Project Wenjuannounced
- August 2020: Project Caerus Kickoff
- September 2020: OpenInfra Labs Governance Meetings Begin
- November 2020: Telemetry Working Group planning meetings begin with representation from researchers and industry.
- December 2020: Telemetry Working Group kickoff meeting and Charter Review
- December 2020/January 2021: OpenShift/Kubernetes CNV cloud in a box pilot replicated.

Highlights in the first half of 2021

- Operate First building pipelines for initial open source projects in use
- Pilot ESI at Mass Open Cloud
- Telemetry working group defines a stack that is sufficient for meaningful, realistic experiments.

OpenStack Project Update 2020

Ussuri

On May 13th, with the help of over 1,000 contributors spanning 50 countries, the OpenStack community released its 21st version, Ussuri. The focus areas of the release were reliability of the core infrastructure layer, enhanced security and encryption, and versatility for emerging

use cases like edge and container environments. As a community, our goals were to make OpenStack's codebase be python3-only (dropped Python 2.7 support) and standardize our approaches to project specific contributor documentation.

- https://www.openstack.org/software/ussuri/
- https://governance.openstack.org/tc/goals/selected/ussuri/index.html

Victoria

Later in 2020, the OpenStack community released Victoria, OpenStack's 22nd version. Native Kubernetes integration was one of the primary focus points; Kuryr, for example, implemented support for custom resource definitions so that it won't have to use annotations to store data about the OpenStack object in the k8s API. More generally, there were also features added to support more diverse architectures and standards, such as direct programming of FPGA's and additional TLS support. Other community goals for the release were migrating upstream CI/CD testing to the new Ubuntu LTS Focal and switching the last of the legacy Zuul jobs that were automatically derived from their Jenkins job to native Zuul v3 jobs. 160 orgnanizations, 45 countries, and almost 800 community members worked together to make the VIctoria release a success! We thank every one of them for their hard work and dedication to OpenStack.

- https://www.openstack.org/software/victoria/
- https://governance.openstack.org/tc/goals/selected/victoria/index.html

New SIGs

TaCT

The Testing and Collaboration Tools (TaCT) SIG was created to serve the role previously occupied by the OpenStack Infrastructure team and support OpenStack’s project-specific testing and collaboration tooling and services. The OpenStack Infrastructure team formerly existed to care for the continuous integration and collaboration infrastructure on which the OpenStack community relies. With the rise of the OpenDev Collaboratory, the majority of the Infrastructure team’s former systems, administration activities, and configuration management repositories were no longer occurring under the authority of OpenStack. The TaCT SIG maintains, in cooperation with the OpenDev project, the tooling, and infrastructure needed to support the development process and testing of the OpenStack project.

Large Scale

This group was formed to facilitate running OpenStack at a large scale, answer questions that OpenStack operators have as they need to scale-up and scale-out, and help address some of the limitations operators encounter in large OpenStack clusters. The work of the group is organized along the various stages in the scaling journey for someone growing an OpenStack deployment.

It focuses from the starting configuration stages and goes through monitoring, scaling up, scaling out and upgrading and maintaining. That path was successfully traveled by many before. The job of the SIG is to extract that knowledge and provide answers for those who come next.

Hardware Vendor

The goals of this SIG is to provide a place where vendors, and those interested in vendor-specific things like drivers and supporting libs, can work together and collaborate openly to enable OpenStack services to integrate and work well on all hardware platforms.The Hardware Vendor SIG is still forming and growing and it currently owns and manages the Python client for Dell's DRAC. The SIG is currently welcoming for other vendors to host their projects too.

Technical Committee Changes

Merging of TC/UC

For a long time, the OpenStack community has had two committees helping to direct their efforts. While it was great to have two perspectives giving guidance, unfortunately this approach recently lead to some siloing within our community. In 2020, the Technical Committee and User Committee meld into one group, which also resulted in some changes of the election process for the TC. Starting August 1st of 2020, when Technical Committee elections are held, Active User Contributors are also included in the the electorate so that they have equal say in their representation. Having the Technical Committee as a united group removes the distinction between developers and operators even more and makes them all equal contributors. This means that operators can run for seats in the committee, where they only need to make contributions to be eligible, that can be reporting bugs, making code or documentation changes, etc. At the same time the developers in the community and on the Technical Committee are encouraged to take on active user contributor tasks and ensure equal representation of operators.

- https://opendev.org/openstack/governance/commit/6a4953f71f3f7a237add862f4aad594896d81634

TC Size Change

Throughout 2020, at each of the technical elections, we reduced the size of the Technical Committee by two people down to our current size of 9 members. The size of the TC is a trade-off between getting enough community representation and keeping enough members engaged and active. Before this change, the size (13 members) was defined in 2013, as we moved from 5 directly-elected seats + all PTLs (which would have been 14 people) to a model that could better cope with our explosive growth. Since then, 13 had worked well to ensure that new members could come in at every cycle untill recently. To avoid burning people out, and keep infusions of new contributors being cycled into the committee, we decided to reduce the size. As a result, the committee has been joined by long term developers and operators like Dan Smith and Belmiro Moreira.

- https://governance.openstack.org/tc/reference/charter.html#number-of-seats-to-elect

Project Changes

As a continuously evolving project OpenStack went through a few governance and process related changes over 2020 to ensure maintainable and efficient operation of the comunity and the project teams.

The concept of distributed project leadership was announced during the second half of the year to help the teams to share responsibilities among themselves better. If a project team opts in to this model that means they will not have a dedicated Project Team Lead (PTL). The necessary tasks to guide the project are taken on by liaisons; the three required roles are release, tact-sig and security liaison. There is no guideline if one or more people fill in these roles. There are also some additional recommended roles to take on to perform tasks such as preparing the team for events or ensure a smooth process to onboard new contributors to the team. The distributed leadership model doesn't affect existing models, such as PTL with liaisons.

In order to make the Technical Committee more efficient the process of making updates to a project became faster. Changes, like adding a new repository to a project or adding a tag to its repositories, required a one week waiting period even if enough votes from the TC were added to the review faster. In the new process two votes from TC members in favor to the change is enough for the chair to approve the request without a waiting period. In case there is a disagreement once the change is merged it can be reverted which than goes through the 'formal vote' process to ensure that enough discussion happens before making a decision.

- Distributed leadership
  - https://governance.openstack.org/tc/resolutions/20200803-distributed-project-leadership.html
  - Oslo
- Sped up process of adding projects
  - More efficient TC
  - https://review.opendev.org/c/openstack/governance/+/726932
- Project Retirements?
  - Streamlined?
  - Make space for new projects?

Tags

Supports Standalone

While OpenStack services work well together, there are users that might not want to run all of them and might prefer other technologies over some of the core components of the project. As a result, some services have been modified so that they can be operated independent from the rest of OpenStack (e.g. Cinder Storage with a Kubernetes cluster) without losing their core functionality. In order to easily identify which services are able to be run standalone without other OpenStack services they are marked with the 'Supports Standalone' tag.

- https://governance.openstack.org/tc/reference/tags/assert_supports-standalone.html

Kubernetes Starterkit

Kubernetes has become the go-to container orchestration system to run containerized applications, most commonly on top a cloud platform. As one of these platforms OpenStack can supply multitenant isolation between different Kubernetes clusters. As OpenStack has a number of services to build infrastructure with, it can be challenging to decide the minimum

set to use as a base for Kubernetes. The Kubernetes starter kit tag recommends a minimal set of OpenStack services to provide the necessary resources to Kubernetes and the workloads to operate.

- https://governance.openstack.org/tc/reference/tags/starter-kit_kubernetes-in-virt.html

10 Years of OpenStack

Open Infrastructure Newsletter

The Open Infrastructure Community Newsletter regularly shares the latest developments and activities across open infrastructure projects, events, and users supported by the Open Infrastructure Foundation. The newsletters are sent out to the community in English through email, and they are also published to the OpenStack official WeChat account in Chinese.

In 2020, 11 newsletters were distributed to the open infrastructure community. There have been more than 45,000 people who have opened at least one newsletter for the past year and more than 74,782 unique views on all the newsletters.

Looking forward to 2021, we will continue to deliver the most updated quality content to the open infrastructure community and create awareness among the broader audience in the community. Check out past newsletters, subscribe to the newsletter, and if you would like to contribute content, please email community@openstack.org.

Superuser

The Open Infrastructure Foundation continued to support the popular Superuser publication, covering the open infrastructure ecosystem news, case studies, event recaps, product updates and announcements, project release and more. In 2020, superuser has reached more than 250,000 unique visitors from 193 countries. Editors and contributors are members of the Open Infrastructure Foundation, the OpenStack and adjacent open source communities, and professional journalists.

10 years of OpenStack

Another milestone that happened in the global community was that OpenStack, one of the top three most active open source projects, marked its 10th anniversary in 2020. On the official birthday, the community has hosted a 10 years of OpenStack virtual birthday celebrationwith more than 130 community members from over 30 companies tuning in and participating around the world. Community members from OpenStack User Group Vietnam also organized an OpenStack 10th birthday celebrationon July 4th.

Starting in July, we have published 28 blogs and interviewed many community members about their experiences in the past 10 years with the OpenStack community. With over 3,200 views on 10 years of OpenStack blog seriesand more than 10 articles published by the ecosystem companies, we have successfully concluded the 10 years of OpenStack celebrations with the global community. Thank you for the 10 incredible years of contribution and support. We absolutely couldn’t have done it without you!

COA

In 2020, 185 community members took the Certified OpenStack Administrator examwith a passing rate of 82 percent. Another example of the magnitude of the global OpenStack community, community members represented 17 different countries—United States, United Kingdom, Switzerland, Spain, Slovakia, Romania, Poland, Netherlands, Italy, Indonesia, India, Greece, France, Egypt, Colombia, China, and Chile.

Ironic Whitepaper

This year the Bare Metal Special Interest Group came together to complete a whitepapercovering the topic in which they focus. This group consists of participants of the OpenStack Bare Metal Logo Program, OpenStack Ironic developers, Bare Metal infrastructure providers, along with other Open Source community contributions from adjacent communities such as Metal3.io. The whitepaper conveys use cases and articulates how participants have improved their Bare Metal management through the use of Ironic. Coinciding with the release of the whitepaper, the Foundation helped launch ironicbaremetal.orgwhich is a website to serve as a basic entry point into learning about Ironic from a context outside of the OpenStack community's core focus.

StarlingX

StarlingX is a fully integrated, open source platform that is supported by the Open Infrastructure Foundation. The project integrates together well-known open source projects such as Ceph, Kubernetes, the Linux kernel, OpenStack, and more to create a cloud platform that is optimized for edge and IoT use cases. During the first half of 2020 the OpenInfra Foundation Board of Directors confirmed the project as top-level Open Infrastructure project.

StarlingX 4.0 was released in August 2020. The platform is architected as a fusion between OpenStack and Kubernetes to create a robust and flexible environment for edge workloads, let them be bare metal, containerized, or virtualized. In the 4.0 release the majority of the infrastructure services are containerized to ensure easier and more flexible management and operation. The community has also continued to have increased focus on security related functionality during the release cycles in 2020.

The 4.0 release is integrating the Ussuri release of OpenStack and it was also certified as a conformant Kubernetes distribution of the v1.18 version. The platform now has support for Kata Containers, as a container runtime, with optional support for Time Sensitive Networking (TSN) functionality. The contributors also integrated Active Directory for Kubernetes API authentication and Certification Manager for more automation in that area.

The community had continued focus on making it easier for users to deploy and operate the platform as well as to help new contributors to start participating in the project. As part of this effort documentation was a high priority item during the year where new contributors added a lot of new content in the area of operation and user guides.

Following the already established processes to build and maintain an open governance model the community held two elections to fully re-elect the Technical Steering Committee (TSC) and one election to elect Project and Technical leads. This means that the community's whole leadership was refreshed through elections during 2020.

The community was particiapting in industry events to spread the word about the progress the conributors are making as well as to provide the opportunity for the contributors and new comers to come together and discuss project related matters, such as design, development and testing related topics. There were presentations about the project at virtual events such as the Open Infrastructure Summit, OpenDev and Edge Computing World and contributors discussed technical details about the project at the two Project Teams Gathering (PTG) events. The community organized one very successful virtual hackathon at the beginning of 2020 to increase unit test coverage throughout all the StarlingX services that are developed by the contributors of the project.

The project had 2237 changes committed by 111 authors from 7 organizations during the year.

2020 was also the year of commercial adoption. T-Systems and Verizon are now deploying StarlingX as part of their infrastructure to utilize the platform's capabilities that make it ideal to deploy it as a distributed system to power use cases in the telecom industry.

The StarlingX code is hosted on Github under the Apache 2 license. Learn about the project, how to contribute and support the community at starlingx.io. Join these channels to get involved:

- Source code: opendev.org/starlingx
- IRC: #starlingx on Freenode
- Mailing list lists.starlingx.io

Zuul

Zuul is an open source CI/CD platform designed for test-driven open source projects and software development organizations who need to gate against multiple projects and systems before landing a single patch. Since 2012, Zuul has been proven at scale as a critical part of the OpenStack development process. In 2018, Zuul version 3.0 was released with the intention of making Zuul useableeasier to use outside of the OpenStack project. Since then we've seen the number of users and use cases grow tremendously. You will find Zuul behind cloud, ecommerce, and automotive software ensuring that every commit passes tests and is ready to be merged.

In 2020 the Zuul community landed 1,305 changes to the project. These changes found their way into seven Zuul releases and five Nodepool releases. By releasing early and often the Zuul project has ensured that bug fixes and features are available to user as quickly as possible. Over the last year major new features include integration with Gitlab as a code review platform, the addition of a Nodepool Azure driver to provision test resources, and an authenticated admin REST API. Zuul and Nodepool now require Python 3.6 or newer and jobs can be run with Ansible 2.7, 2.8, and 2.9.

The new authenticated admin REST API provides tenant scoped admin access to Zuul users. This allows tenant specific users to enqueue/dequeue changes for testing, promote changes in dependent pipelines, and create autohold requests for failing jobs.

Zuul has found its way into automotive development. Modern cars contain millions of lines of code [0] and that software needs to be tested. Volvo says, "Premium cars need premium tools, and Zuul is a premium tool." In fact, the One Pedal Drive found in Volvo's Polestar 2 was built

with Zuul. Volvo specifically mentions that cross project dependencies and project gating are incredibly important to their workflows now. BMW also points out that gating is a key feature that enables their teams to work effectively on the ever more complex problem of building automotive software.

Zuul's integration with the Gerrit project continues to get better. Gerrit is now running a Zuul instance that is used to test Gerrit plugins. Plugins can opt into this testing and a number have including the replication, code-owners, and oauth plugins.

Looking ahead to 2021, users should be ready to upgrade to Zuul version 4.0 and eventually 5.0. These new version numbers have been selected to properly communicate the operational impacts for the planned changed. The jobs and user facing configuration for Zuul is not changing. Zuul version 4.0 will require a SQL database connection and the Zookeeper connection must be made with SSL/TLS. These changes prepare the way for Zuul version 5.0 which will enable scale out event handling and scheduling of jobs. Zuul version 5.0 will remove the last remaining single point of failure in a Zuul deployment.

Zuul will also continue to make incremental improvements. We expect Ansible 2.10 support, Digital Ocean provisioning in Nodepool, and Bitbucket code review integrations to land in the next year. The web UI will be updated to expose the new admin REST API functionality via web browsers as well. The Zuul project welcomes contributions to continue to expand the number of code review and resource provisioning integrations supported by the project. If you have an integration you would like to add please get in touch with us.

If you would like to get started on your Zuul journey definitely say hello to the community. We do our best to help those with questions. There are also a couple of companies working to help you get jump started now. Vexxhost runs a hosted and managed Zuul service [1], and Acme Gating will help you run your own Zuul or develop features that you need [2].

We are grateful to everyone that has contributed to the Zuul project in 2020. The community has done really great work to improve Zuul this year. The ground work for large architectural changes like those necessary for Zuul version 5.0 is not an easy task to accomplish. This community has managed to make those preparations while still doing maintenance and upkeep. Thank you to all of our 2020 contributors!

Zuul is free and open source software licensed under the Apache 2 license. If you would like to get involved, join us:

- On the mailing list: zuul-discuss@lists.zuul-ci.org
- On IRC: #zuul on Freenode
- Code review: https://review.opendev.org/q/status:open+-is:wip+project:%2522%255Ezuul/.*%2522

[0] https://www.nytimes.com/2017/06/07/technology/why-car-companies-are-hiring-computer-security-experts.html
[1] https://vexxhost.com/solutions/managed-zuul/
[2] https://acmegating.com/

OpenInfra Collaboration

Magma

Magma is an open-source software platform that delivers an open, flexible and extendable mobile network core solution. Magma's mission is to connect the world to a faster network by enabling service providers to build cost-effective and extensible carrier-grade networks. It is designed to be access network agnostic (4G, 5G or WiFi), flexibly supporting a radio access network with minimal development and deployment effort.

Magma was developed as an open-source project as part of Facebook Connectivity's incubator program, with significant contributions from the OpenAirInterface Software Alliance. In July of 2020 Magma's code and documentation was contributed to a new, independent software project magmacore.org.

Magma's architecture consists of three major subsystems:

- The Access Gateway (AGW) provides network services and policy enforcement. In an LTE network, the AGW implements an evolved packet core (EPC), and a combination of an AAA and a PGW. It works with existing, unmodified commercial radio hardware. Support of 5G radios is under active development.
- Orchestrator is a service that provides a simple and consistent way to configure and monitor the wireless network securely. The orchestrator can be is normally hosted on a public or private cloud. The metrics acquired through the platform allow you to see the analytics and traffic flows of the wireless users through the Magma web UI. Orchestrator also provides REST APIs allowing simple integration with monitoring and management systems.
- The federation gateway allows Magma to integrate with other Mobile Network Operator core network functions using standard 3GPP interfaces. It acts as a proxy between the Magma AGW and the operator's network and facilitates core functions, such as authentication, data plans, policy enforcement, and charging to stay uniform between an existing MNO network and the expanded network with Magma.

Community

At the Open Infrastructure Summit in October 2020, the Open Infrastructure Foundation announced its support of the Magma project. With help from the Foundation, magmacore.org has been structured to make it easy to contribute to the project. It is centered around a developer-led governance with Maintainers providing code review and approval for contributions and a small Technical Committee selected from among the project maintainers providing overall architecture and roadmap oversight. In addition, the Open Infrastructure Foundation also helped launch the Magma Slack workspace back where many community and technical conversations take place. After the Summit announcement, the Slack workspace grew by over 100 people.

The Open Infrastructure Foundation also assists with the monthly Magma community meetings where project updates and use cases are presented. Community members also get the opportunity to ask questions directly to the Magma engineers. The first community meeting took place in November 2020 and had over 20 attendees. View the meeting recordings here.

OpenDev

OpenDev believes that free software needs free tools. OpenDev is a collaboratory for open source software development at a huge scale. Its focus is on code review, continuous integration, and project hosting provided exclusively through open source solutions like Git, Gerrit, Zuul, and Gitea.

As the world shifted to remote work and online collaboration in 2020, the need for the services provided by OpenDev has become more apparent. OpenDev has stepped to the plate and made a number of changes to help make 2020 as productive as possible all while minimizing disruption to existing services and users. In particular we would like to call out the new formalized governance for the project, new services like meetpad, and service upgrades (like Gerrit).

In 2020 OpenDev formalized its governance with two main bodies: the Service Coordinator and the Advisory Board. The Service Coordinator is an individual elected from OpenDev's contributors whose primary role is to communicate planned changes to users and incorporate their feedback. The Advisory Board is made up of liasons from the various projects using OpenDev and organizations that contribute compute resources. They serve as the primary communcation point between users and the service. If you represent an OpenDev using project or a contributing organizationand would like to get more involved feel free to join the Advisory Board.

One of the big changes we have had to deal with as a community is the transition to virtual events. The OpenDev team has deployed a Jitsi Meet instance which is paired with our Etherpad service to act as a video conferencing with built in note taking tool. This service has been used in multiple Project Team Gatherings at this point and helped us keep an event format that is familiar to those who have attended in person events. Outside of events we have seen this service used for team meetings and other places where video conferencing is helpful. We expect this service to replace the existing Asterisk based phone conference tools that we previously deployed.

Several services have also seen major upgrades during 2020. The most notable (and likely noticeable) is Gerrit, but Gitea and Etherpad have also been upgraded. For Gerrit this upgrade was long overdue and gets us much closer to a modern upstream Gerrit experience. We are hoping that this change will help us work with upstream Gerrit more closely going forward, bringing our needs into upstream development. A number of features come with this upgrade including use of NoteDB for data storage and single page diffs. We continuously upgrade our Zuul deployment, but over the course of 2020 it got some major new features and a new look to its web dashboard as well.

Our long running effort to migrate away from Puppet for service deployments and configuration management continues. Several more services have been migrated onto Ansible and Docker including Hound, Grafana, and our IRC accessbot. This effort enables us to use tools that more developers are comfortable with to manage our services and gives us better control over what is being deployed.

One of the major tools that has enabled our transition to well tested Docker based deployments is Zuul's image build pipeline job library. Zuul provides a set of building blocks to run Docker image builds for changes using local image repositories to share the resulting

images between test jobs and changes. This allows any OpenDev user to test sequences of changes to different services and know that each change in the sequence is a landable production state.

On the continuous integration side of the house our CI jobs helped 1,800 users merge more than 54,000 changes. We continue to keep our test environments up to date, pulling in new distro releases and removing old EOL versions. Keep this in mind as CentOS's announcement that version 8 will be EOL at the end of 2021 means that we will be removing it at the end of this year. CentOS Stream 8 test nodes are available as are alternative distros with long term support and stable releases.

OpenDev continues to be a community run project. These services exist because we all help work together to make it happen. 2020 was a difficult year in many ways, but for OpenDev one of the most noticeable difficulties has been losing individuals who have long been a part of the community providing tremendous amounts of help to OpenDev. We would like to thank Andreas Jaeger and Monty Taylor for all the work they have done over the years. This also serves as a reminder that we can always use more help. If you are interested in getting involved do not hesitate to reach out.

Find us at:

- On the mailing list: service-discuss@lists.opendev.org
- On IRC: #opendev on Freenode
- Code review

OpenInfra Foundation Working Groups

Diversity & Inclusion Working Group

The Diversity and Inclusion Working Group continued reaching out to the other OIF projects and inviting them to join our efforts as we are not OpenStack specific and hope to continue offering opportunities for diversity at events and in the community.

Our main focus for the year was the creation of the Open Infrastructure Foundation's stance on divisive language. All OpenInfra projects were contacted and over the course of the work all projects gave input except for Kata. The stance was approved by the Board and the community with several recommendations being relayed to the Board for continuing efforts from them as well as the creation of a Wiki page with suggestions to be used when addressing the replacement of divisive words.

The D&I WG also helped to lead OpenStack's presence during Grace Hopper's Open Source Day. We had about 13 mentors from the community and 45 attendees throughout the day.

In 2021, we plan on continuing to aid the OpenInfra projects in their efforts to remove divisive language and to use more inclusive words. We also hope to be able to resume efforts with the community during virtual and in person events.

Interop Working Group

The OpenInfra Interop Working Group has been issuing guidelines for OpenStack Logos and Branding programs as part of Marketplace development every release as listed below:

- OpenStack Powered Platform
- OpenStack Powered Compute
- OpenStack Powered Object Storage

During Ussuri we had two add-ons:

- OpenStack Powered DNS
- OpenStack Powered Orchestration

In Victoria we added one more add-on: the OpenStack Powered Shared File System.

The Interop repos were moved under the OpenInfra Foundation in Ussuri and are consolidated in Victoria to enable QA and test team to plan for ongoing efforts in Open Infrastructure projects Kata, Airship and StarlingX for container conformance.

The forward looking statement from the OpenInfra Interop Working Group is to start the ground up efforts to set new guidelines for Integrated Wallaby release and define the Open Infrastructure Container Conformance Program. The current work in the technical team has been efforts around containerizing Refstack to move to Ansible and Go besides python 3.7 support for test tools and libraries used by all Open Infrastructure Projects. The Baremetal logo program was deferred to future releases after new guidelines are in place.

Open Infrastructure Community Highlights

OpenInfra User Groups Growth

In 2019, OpenInfra User Groups migrated to the Meetup Pro platform, which allowed the groups more visibility on the web, as well as a streamlined way to organize and communicate.

In 2020, we saw the number of groups grow past 40, adding communities in places like the Congo and Tijuana.

Amidst challenging circumstances, User Groups showed their commitment to community, organizing in new and unique ways. When in-person events were put on hold, several groups took on the task of hosting virtual meetups covering a variety of topics including hybrid cloud, storage, containers, and edge. Some groups even organized watch parties for the first virtual Open Infrastructure Summit that took place in October.

While we know User Groups are looking forward to meeting in person again, we’re sure that virtual meetups are still going to be around because of the valuable accessibility they provide so many community members around the world.

Outreachy Program

For the current round, December 2020 to March 2021, 5 interns were selected:

- Kuryr: T.I. Fasoyin from Nigeria
- Manila: Dina Saparbaeva from Karakol, Kyrgyzstan
- Ironic / OpenStack SDK: Anuradha from India
- Ironic: Oluwatosin Farai from Nigeria
- Manila: Paulali from Kenya

Students in OpenStack

Students are the future- the future of the workforce, the future of the tech industry and the future of open source. The OpenStack community has been slowly building relationships with more and more universities, getting them involved in work upstream in the community. These relationships are mutually beneficial; students get practical experience and build skills in communication and working collaboratively and the OpenStack community gets fresh perspective and the excitement that comes with new contributors.

Boston University Senior Design Program

Three students- Nicole Chen, Ashley Rodriguez, and Mark Tony- at Boston University selected OpenStack Manila as their senior design project. With the help of community mentors- Goutham Pacha Ravi, Victoria Martinez De La Cruz, Jeremy Freudberg, Maari Tamm and Kendall Nelson- the students have begun working on implementing the Manila APIs in the OpenStackSDK. There is a community effort to reach feature parity between the individual python clients for OpenStack services and the OpenStackSDK. The students were onboarded into the community, they have begun attending community meetings, and they also have been working on the specification outlining the work they are beginning and the details of implementation. If they are able to complete all the work in the OpenStackSDK, they can then begin implementing those features in the OpenStack Client. They will continue their work for the remainder of their school year, wrapping up in May of 2021.

North Dakota State University Capstone Program

For several years now, we have participated in North Dakota State University's Capstone program that gives students a chance to select a project- either open source or industry- to work on throughout the spring semester. At the start of 2020, the OpenStack community had four students- Dawson Coleman, Luke Tollefson, Noah Mickus, and Steven Glasford- helping to add support for configuring TLS ciphers and protocols on Octavia load balancers. Michael Johnson, Adam Harwell, and myself mentored students through their semester helping them learn not only about Octavia, but about OpenStack, and open source. With their hard work, they landed TLS cipher and protocol support before the end of the Ussuri release and it became a large part of the promotion and marketing around that release. The rest of their work landed shorty after feature freeze and was included in the Victoria release that happened in October of 2020.

Oregon State University Internship

This year the Foundation started a partnership with Oregon State University funding one student to work in the OpenStack community on the Glance project to give them real world experience working on an open source project. Kendall Nelson and Erno Kuvaja mentored Khuong Luu. He worked on a variety of tasks in Glance. Khuong made several documentation improvments including adding missing information and reording topics to make them more readable and the information more discoverable. He cleaned up config options that were causing orphan processes in Glance when a user sent a SIGHUP signal. Khuong also cleaned up a bug being caused my a missing image property. While he didn't spend time on one large feature, his efforts cleaned up a lot ot technical debt. He finished up his internship when the school year concluded and we thank him for all his hard work!

We have a new student intern getting started with his onboarding process. Ryan Zimmerman, once up to speed on OpenStack and our community, wil begin work on the OpenStackSDK and helping OpenStack services reach feature pairity with their own project specific python clients.

OpenStack Upstream Institute

The OpenStack Upstream Institute (OUI) training is an interactive course that teaches attendees about how the OpenStack community operates including the community structure and processes. The training also includes hands-on practice to teach participants how to use the tools that the community depends on. These tools -such as Git, Gerrit, and more - are maintained by the OpenDev (https://opendev.org) community. As in person events were not possible during 2020, we had to re-invent OUI to make it suitable for being hosted in a virtual format.

The virtual training had a slightly different structure compared to the in person classes. The training was split into two 3-hour long blocks: OpenStack specifics and OpenDev tools. The OpenStack module covered items such as governance, OpenStack release cycles, ways to communicate within the community, and more. The attendees learned how the community operates as we went through the topics during the first block. The class utilized active synchronous conversation between students and mentors, where attendees are able to ask questions from our mentors, who are all active community members.

The OpenDev tools module was also a hands-on class that focused on teaching attendees to setup and use tools that are used for software and documentation development processes as well as task tracking. Many of these tools are used by multiple Open Infrastructure communities and adjacent projects. By organizing this way, we were able to welcome and educate students who might not have an interest in OpenStack specifically which gave us a broader audience. During the class the attendees worked through exercises with guidance from the mentors to give them practice using the tools. During the training we also had breakout rooms students and mentors could hop into to ensure any technical issues students had were solved before the end of the class.

OUI, with the new structure, was held during the Open Infrastructure Summit in October. The class material was also utilized at other events to hold the training in different formats. Ghanshyam Mann hosted a virtual OUI session before the Summit and Amy Marrich led a group of over 15 mentors to educate participants of the Grace Hopper Conference's Open Source Day on how to get started in OpenStack.

Press & Analyst Highlights

Public relations efforts are led by a distributed team of professionals in the United States, Europe, and Asia Pacific. This team engaged analysts and journalists worldwide in dialogue, proactively delivering news, commentary, and contributed byline content for both the Foundation and for its individual projects. Community news highlights are available at https://openstack.org/news.

Key media relations initiatives included the following:

- Launch of the Open Infrastructure Foundation
- Project software releases including OpenStack Ussuri and Victoria, Kata Containers 2.0, and StarlingX 4.0
- Virtual Open Infrastructure Summit
- Virtual OpenDev event
- Celebration of the OpenStack project’s 10th anniversary
- Momentum of OpenInfra projects in the broad open infrastructure and cloud IT landscape

Analysis shows that OpenStack/Open Infrastructure Foundation news coverage comprised more than 24,100 total placements in 2020. Media coverage was distributed globally (31% APAC, 30% North America, 22% Europe, by volume), and extended into 94 countries.

As evidenced by the highlights shared below, media attention has shifted from questions about OpenStack’s future to an active interest in use cases and the progress of the other projects supported by the OpenInfra Foundation.

In 2020, the Open Infrastructure Foundation adopted a virtual briefing format, shifting from in-person analyst day sessions at the physical events. During the course of the year, Foundation staff leadership conducted briefing sessions with more than two dozen analysts in Asia, North America, and Europe, providing updates on both the evolution of the Foundation’s business model and the status of individual projects in the organization’s portfolio. Firms with strong ongoing interest in our community’s work include IDC, Forrester, Redmonk, 451 Research, and others.

Analyst interest in the community’s work has shifted steadily since the early days of the Foundation. While early coverage necessarily focused on the OpenStack project, the focus has broadened substantially to include other projects under the Open Infrastructure Foundation. Additionally, coverage has shifted to include Open Infrastructure Foundation projects as components of broad-based solutions that leverage open infrastructure technologies such as Kubernetes.

Superuser

The Open Infrastructure Foundation continued to support the popular Superuser publication, covering the open infrastructure ecosystem news, case studies, event recaps, product updates and announcements, project release and more. In 2020, superuser has reached more than 150,000 unique visitors from 193 countries. Editors and contributors are members of the Open Infrastructure Foundation, the OpenStack and adjacent open source communities, and professional journalists.

Event Highlights

As many organizations experienced in the first few months of 2020, the Open Infrastructure Foundation had several challenging decisions to make as the pandemic began to spread rapidly around the world. Having the safety of the Community as the top priority, the decision was made to move the entire 2020 program, including the PTG (Project Teams Gathering), OpenDev series, and Summit to a virtual format. With the help of Foundation and event sponsors, partners, and the support of the entire Community, we were able to continue our mission to collaborate, share knowledge, network, and continue improving the software with the broadest global reach ever despite the inherent challenges of the year.

Open Infrastructure Summit

The first virtual Open Infrastructure Summitwas held in October 2020, gathering 10,000 attendees (largest Summit attendance to date) from 120+ different countries to discuss 30+ open source projects. In addition to use cases from users like Volvo Cars Corporation, Ant Group, GE Digital, Société Générale, China Mobile, European Weather Cloud, Workday, and China Tower- plus a big announcement about the future of the Foundation:

- The Summit kicked off with the Keynote announcementthat the OpenStack Foundation was officially becoming the Open Infrastructure Foundation
- **Volvo Cars Corporation**discussed how it successfully uses Zuulfor continuous integration in a range of different software components, and how Zuul is a game changer that helps them avoid merging broken code.
- **Ant Group**, the leading peer-to-peer payments processor in China, presented why isolation methods such as scheduling isolation and LLC isolation are used in combination with Kata Containers.
- **GE Digital**presented the tools, migration procedure, and highlighted the biggest challenges in upgrading from OpenStack Newton to Queens with minimal downtime.
- **SK Telecom**5GX Labs won the 12th Superuser Award for developing a containerized OpenStack on Kubernetes solution called SKT All Container Orchestrator (TACO), based on OpenStack-helm and Airship, all while contributing upstream to both the OpenStack and Airship projects.

Summit videos are available on the Summit videos page. Thank you to our Summit sponsorsfor supporting the event!

OpenInfra TV: Bringing the Summit to Asia

At the annual OpenInfra events, the global community convenes to discuss the challenges and success of integrating different open source projects. Averaging attendees from over 60 countries at face-to-face events and attracting attendees from over 120 countries for the virtual OpenInfra Summit, the Foundation events team concentrates on making sure all events cater to the diversity of the audience.

For the virtual event, time zones were one of the biggest challenges as one block of time is never perfect for every single person in the world. The team selected four communities to pilot OpenInfra TV—China, India, Japan, and Korea—to provide local coverage of Summit highlights presented by community leaders in their local language.

OpenDev

The OpenDev event, held in-person in previous years, was a three part virtual series that focused on three topics: Large-scale Usage of Open Infrastructure Software, Hardware Automation, and Containers in Production.

OpenDev events bring together the developers and users of the open source software powering today's infrastructure, to share best practices, identify gaps, and advance the state of the art in open infrastructure. For each topic, participants joined discussion oriented, collaborative sessions where they explored challenges, shared common architectures, and collaborated around potential solutions. Each topic's agenda was set by a programming committee made up of community members who met for weeks leading up to the event to make sure it was tailored to the broader community.

The three events gathered over 2,000 registrants from over 70 countries.

Here are some of the highlights:

Large-scale Usage of Open Infrastructure Software - June 2020

The event kicked off with user stories from Blizzard Entertainment, OpenInfra Labs, and Verizon who shared their own scaling challenges in short presentations followed by Q&A. After that, the discussion focused on a variety of topics including upgrades, high availability (HA) networking, bare metal, and more. At the conclusion of the event, participants looking to get further involved were encouraged to join the OpenStack Large Scale SIGto continue sharing challenges and solutions around scaling.

Thank you to the programming committee members: Beth Cohen, Roman Gorshunov, Belmiro Moreira, Masahito Muroi, and Allison Price!

Read the full recapand catch up on the discussion recordings.

Hardware Automation - July 2020

This OpenDev event focused on topics around hardware provisioning lifecycle for bare metal, bare metal infrastructure, networking and network security. Participants representing over 200 companies spent three days sharing their knowledge and discussing their experience of building and operating software that automates their hardware in cloud environments. One of

the ways mentioned to keep the momentum going after this event was to join the OpenStack Multi Arch SIGand OpenStack Bare Metal SIGto continue sharing challenges and solutions around Hardware Automation. Julia Kreger also highlighted the latest white paper from the OpenStack Bare Metal SIG: "Building the Future on Bare Metal, How Ironic Delivers Abstraction and Automation using Open Source Infrastructure."

Thank you to the programming committee members: Keith Berger, Mark Collier, Julia Kreger, Mohammed Naser, and James Penick!

PTGs / Forums

In June of 2020, we hosted our very first virtual Project Teams Gathering. We were pleased to have over 700 people register for the event - more than we've had for any in person PTG. By taking it virtual instead of in-person, we had new people participate that have otherwise never been able to attend the PTG. Over the course of the week, 44 teams across Airship, Kata Containers, OpenStack, StarlingX, and Zuul met to discuss their upcoming releases and make progress on standing items. The format of the PTG was different than any event we've hosted before because it was all virtual and we needed to accomodate for all timezones. We settled on 4 hour blocks of working hours with 4 hours inbetween and allowed teams to schedule themselves to better fit their day to day schedules and time zone restrictions. The PTG again made use of the community developed PTGbot with some tweaks to help keep people aware of active discussions and find where and when teams were meeting.

The community had two options to use for their video meetings: the Open Infrastructure Foundation subsidized Zoom rooms, or the OpenDev team hosted meetpad rooms. Many teams wrote up summaries of their discussions and accomplisments throughout the week:

- Airship: https://www.airshipit.org/blog/airship-update-june-2020/
- Kata Containers: https://medium.com/kata-containers/kata-updates-from-the-project-teams-gathering-72cc565e5675
- OpenStack: https://www.openstack.org/blog/victoria-vptg-summaries/
- StarlingX: https://www.starlingx.io/blog/starlingx-vptg-june-2020-recap/

We continued to meet virtually throughout the year and the week after our first Virtual Open Infrastructure summit, we decided to again, keep people safe, and run it like we did in June. In October of 2020 the teams met on Zoom and in Meetpad to talk to each other and collaborate. Over 500 people registered for the event spanning 45 countries. The event hosted communitiy members across 46 teams and had their usual technical discussions which you can find summarized here:

- Airship: http://lists.airshipit.org/pipermail/airship-discuss/2020-November/001135.html
- OpenStack: https://www.openstack.org/blog/wallaby-vptg-summaries/
- StarlingX: http://lists.starlingx.io/pipermail/starlingx-discuss/2020-November/009958.html

OpenStack Days / OpenInfra Days / Cloud Operator Days

Overcoming the local challenges of the pandemic, the global open infrastructure community organized four OpenStack/OpenInfra Day events (a combination of virtual and face-to-face in Turkey, China, Indonesia and Korea - educating over 10,000 attendees around the world.

The typical attendee reached at these local events are cloud and IT architects, software developers, platform and solution engineers and product management.

The Tokyo community introduced a new local event, Cloud Operator Days, a two day online event that gathered 1,400 attendees to discuss challenges of operating cloud infrastructure.

These one or two day regional events are organized by the local community and supported by the OpenInfra Foundation to include project workshops, upstream training, and booths from the local ecosystem. These events continue to be an excellent touchstone for the community to engage with our ecosystem, and to gather local open source developers and users to collaborate, share use cases, and support for the OpenInfra projects.