



## Project Description

[etcd](#) is a key-value store for distributed systems that provides a way for applications of any complexity -- from a simple web app to Kubernetes -- to store and access critical data. Originally created at CoreOS -- now part of Red Hat -- etcd offers consistency guarantees and can tolerate machine failures even in the leader node. Applications can read data from and write data into etcd. It can watch specific keys or directories for changes and react, enabling applications to reconfigure themselves accordingly. It is the primary datastore for Kubernetes.

etcd joined the CNCF in 2018 and remained in incubation for two years. Last November, it hit graduated status within CNCF, joining the ranks of Jaeger, Vitess, Helm, and others.

This report assesses the state of the etcd project and how CNCF has impacted its progress and growth. Without access to a multiverse to play out alternative scenarios, it is impossible to sort out causation. However, we can document correlations. This report is part of a series of project journey reports published by CNCF focused on graduated projects.

## Project Snapshot

etcd made its [first commit to GitHub](#) on June 6, 2013. Between [joining](#) CNCF on December 11, 2018, and today, etcd has added:



\*Note: These statistics were collected with the [DevStats](#) tool, which CNCF built in collaboration with CNCF project communities. "Contributor" is defined as somebody who made a review, comment, commit, or created a PR or issue.

## CNCF Premise of Open Source Software Development

A basic premise behind CNCF conferences, including [KubeCon + CloudNativeCon](#), and open source in general, is that most interactions are positive-sum. There is no fixed amount of investment, mindshare, or development contributions allocated to specific projects. Just as open source development is based on the idea that, collectively, we are smarter together than any one of us alone, open source foundations work to make the entire community better. Equally important, a neutral home for a project and community fosters this type of positive-sum thinking, and drives growth and diversity that we believe are core elements of a successful open source project.

## CODE DIVERSITY

### Diversity Across Company Size and Type (End User, Vendor, Foundation)


From its roots at CoreOS, the etcd project has grown to incorporate meaningful code contributions from more than 600 organizations. High-velocity open source projects like etcd garner wide adoption and contribution from both vendor and end-user communities. As such, etcd code contributions come from a wide range of companies, fostering user-driven innovation.

The diversity of vendor contributors is expanding; in 2018, project creator CoreOS was acquired by Red Hat, now by far the largest contributor to etcd. Contributors to etcd include many of the world's largest tech companies, such as Google, Red Hat, Amazon, Alibaba, IBM, and Facebook, as well as fast-growing mid-size companies, like Indeed and Cockroach Labs. Contributions also come from dozens of smaller businesses and startups, such as Heptio and Weaveworks. Contributing organizations to etcd are well distributed between vendors and end users, demonstrating that end-user innovation can foster and sustain fast-growing, successful projects.


## By the Numbers

The top two contributing companies to etcd as of the end of the December 2020 reporting period were Red Hat and Google, with 69% and 5% of contributions, respectively. Red Hat and CoreOS provided the majority of code contributions to the project over its first two years. Since then, the etcd project has diversified to include many additional companies. The total number of companies contributing code has increased by 33% since etcd joined CNCF, from 450 to over 600. Red Hat's percentage of all contributions has decreased somewhat, but the company has continued to contribute a high volume of code while Google, Alibaba and Amazon have expanded their contributions. This indicates a healthy dynamic in which the project originators and early joiners continue to contribute high volumes of code, but encourage other organizations to contribute a greater percentage of code over time, sharing stewardship and growing the community. Another key project health indicator is the number of contributors. etcd has enjoyed 66.8% expansion of individual contributors over the two years since the project joined CNCF. During the five years before joining CNCF, etcd accumulated 2,335 contributors; in the two years since, it has added 1,741 contributors.

### etcd contributions by company (cumulative)




Cumulative growth of etcd contributions by company from Q1 2014-Q1 2021



Percentage breakdown of etcd contributions by company from Q1 2014-Q1 2021


### Percentage of etcd code contributions by company

## Total Number of Companies Contributing etcd code



Cumulative number of companies contributing by quarter (Q1 2014-Q1 2021)

## Number of Contributors to etcd Project



Cumulative growth in contributors by quarter (Q1 2014-Q1 2021)

# GEOGRAPHIC DIVERSITY OF CONTRIBUTORS


Contributors to etcd have come from more than four dozen countries spread across six continents. The geographic diversity of contributions expanded quickly from 12 countries in the first year of the project to 26 during the second year. The chart shows the percentage of contributors over time, broken down by country (based on self-reported location on GitHub).



Map of project contributors by country from Q1 2014-Q1 2021


## Percentage contributors to etcd by country

Percentage Contributions to etcd by Country



Change in number of monthly contributors by country as a percentage of total (Q1 2014-Q1 2021)

etcd Contributions by Country (cumulative)




Change in number of contributors by country (Q1 2014-Q1 2021)


## DEVELOPMENT VELOCITY

Having gained across several key velocity metrics since joining CNCF, etcd is flourishing. One way we track developer velocity is with the following formula:  $\text{velocity} = \text{commits} + \text{PRs} + \text{issues} + \text{authors}$ . We also look at the growth of PRs, code commits, and issues filed as separate line charts. A third way to examine velocity is by looking at the cumulative number of contributors over time. The charts below illustrate rising velocity for etcd.

### Monthly velocity of etcd



— commits — prs — issues — authors



### etcd commits, PRs, and issues

Growth of etcd pull requests, code commits, issues, and authors over time (Q1 2014-Q1 2021)

## EDUCATION, EVENTS, AND SPONSORSHIP

Growth of community participation in education, events, and sponsorship is a reliable proxy for the health of a project.

The etcd project actively participates in KubeCon + CloudNativeCon North America, China, and Europe through a variety of presentations and talks from community leaders. Since joining CNCF, the project has been the topic of seven talks and presentations across all three CNCF flagship events.




# MARKETING GROWTH AND PROGRAMS

## Website and Social Media

When etcd joined CNCF in December 2018, the foundation started promotional efforts to help sustain, nurture, and expand the etcd community. These have included blog posts, email newsletter mentions, and social media support. Thanks in part to these efforts, public awareness of and interest in etcd have grown quickly. Google Analytics data for etcd shows an increase in pageviews since the project was contributed to CNCF, totaling more than 1.9M to date.

etcd has over 1k followers on Twitter, having increased its following by more than tenfold since joining CNCF.



Growth in monthly pageviews from December 2018 to February 2021


# PROJECT DOCUMENTATION\*

Continuous additions to and improvements of project documentation are essential for the growth of any open source project. Robust documentation is critical to educating new users and helping existing users resolve problems and understand a project’s capabilities. etcd documentation has significantly increased over the years. Since joining CNCF, the number of authors and companies committing documentation to etcd has grown by 17% and 12%, respectively. As of this report, 249 authors have committed, and 93 companies are involved in committing, documentation. The number of documentation commits has increased by 25% to reach 1,594 since etcd joined CNCF.

In December 2020, etcd embarked on a [documentation improvement program](#) with support from the CNCF. The goal is to improve the organization and SEO performance of documentation as well as contributor and user experiences with it.


*\*Documentation for etcd is authored in .md files. CNCF uses the DevStats tool to automatically collect and count statistics of all relevant .md files in the etcd repositories in GitHub.*

### etcd Documentation Committers



Growth in participation in etcd project documentation (Q1 2014-Q1 2021)

### Documentation Commits



Cumulative growth of etcd project documentation commits (Q1 2014-Q1 2021)

# CONCLUSION

CNCF is committed to **fostering and sustaining an ecosystem** of open source, vendor-neutral projects by democratizing state-of-the-art software development and deployment patterns to make technology accessible for everyone. We hope this report provides a useful portrait of how CNCF is fostering and sustaining the growth of etcd.

“The IBM Cloud team has been deploying etcd into production for four years now. With hundreds of thousands of etcd instances created for Kubernetes and OpenShift over that time period, it has proven to be a reliable technology that we are happy to work with. In addition, etcd’s watch mechanism has proven critical to our microservices architecture, enabling a very resilient declarative model for our cloud service. etcd, along with Kubernetes and OpenShift, is the technology our business is built on.”

- **JAKE KITCHENER, SENIOR TECHNICAL STAFF MEMBER, IBM**

“Having etcd as our meta-store in Placement Driver and our inspiration for Raft implementation in production has proven to be a great choice for TiKV and TiDB, ensuring data consistency and high availability across TiDB clusters. We are proud and glad to be part of its graduation journey, and we’d love to be involved in its ecosystem development more in the future.”

- **ED HUANG, CO-FOUNDER AND CTO, PINGCAP**

“After seven years of development, etcd has reached maturity and become the cornerstone of many distributed systems. The most important decision for its success was joining the CNCF community and growing its maintainers across many organizations. We are excited to see its graduation at CNCF. etcd is the centerpiece powering the container service and many other critical services at Alibaba Cloud. We are looking forward to improving its stability, reliability, and performance with the community in the future.”

- **XIANG LI, ENGINEERING DIRECTOR, ALIBABA CLOUD**

“Open source software powers our lives in so many ways. From Linux to Kubernetes, open communities of builders from all sizes of organizations and walks of life spend considerable time creating and maintaining projects that underpin much of the internet, telecommunications, finance, transportation, gaming, retail, and healthcare systems we use every day. etcd is one of these critical projects, and we’re proud to have etcd as a core part of Amazon EKS and to be involved in helping the project grow and thrive. We are fervent supporters of etcd’s graduation and look forward to collaborating with etcd and other CNCF projects to build secure, reliable, powerful, and scalable open source software.”

- **BOB WISE, GENERAL MANAGER OF KUBERNETES, AWS**