Alibaba Cloud

Alibaba Open-source Fast Facts
2019 November

Overview

Since the first batch of open-source projects launched in 2011, developers at Alibaba have been actively contributing to open-source communities. Alibaba now boasts over 180 open-source projects, contributing codes to all aspects of software solutions, including cloud infrastructure and machine learning, container, enterprise-class system, database and network.

Alibaba and its computing arm Alibaba Cloud have been taking prominent role in a number of worlds' leading open-source communities, such as Linux Foundation, Apache Software Foundation, MariaDB Foundation, Cloud Foundry Foundation, Cloud Native Computing Foundation

- In the star ranking of world's largest developer community Github, Alibaba has gained over 690,000 stars, with about 20,000 contributors, as one of the top ten organizations.
- Since 2011, Alibaba has contributed to open source communities including Cloud Native Computing Foundation, Aliance for Open Media, Cloud Foundry, Hyperledger, Open Container Initiative, Continuous Delivery Foundation, The Apache Software Foundation, MariaDB Foundation, The Linux Foundation.
- Linux Community, Alibaba has contributed over 290 patches
- Alibaba's open-source project "Dubbo" "RocketMQ" "Weex" "JStorm" were donated to Apache Foundation, "Dubbo" and RockeMQ became the top-level project at Apache
- Alibaba Cloud's OSS has become the third official recognized cloud storage by Hadoop
- Alibaba Cloud has helped MySQL to identify nearly 300 bugs, contributed all the patches to MariaDB

Milestones

2012 Alibaba Joined Free Software Foundation
Aug 2015 Alibaba became Advisory Board of Xen Project
Nov 2015 Alibaba became Silver Member of Apache Software Foundation
Nov 2015 Alibaba donated Jstorm to Apache Foundation for further incubation
Oct 2016 Alibaba Cloud Engineer Lixun Peng become the first Chinese staff member of MariaDB Foundation
Nov 2016 Alibaba donated RocketMQ to Apache Foundation for further incubation
Dec 2016 Alibaba Cloud became founding member of Containerd.IO
Dec 2016 Alibaba donated Weex to Apache Foundation for further incubation
Feb 2017 Alibaba Cloud became Gold Member of Linux Foundation
Feb 2017 Alibaba Cloud became Platinum Member of MariaDB Foundation
Feb 2017 Alibaba Cloud joined Network Time Foundation
Feb 2017 Alibaba Group joined Ceph to contribute to the distributed storage system
May 2017 Alibaba Cloud became Gold Member of Cloud Native Computing Foundation
June 2017 Alibaba Group joined Open Compute Project as platinum member
Sept 2017 RocketMQ became top-level project at Apache Foundation

Oct 2017 Alibaba announced new open source projects “OpenMessaging” and “ApasaraCache” that supported by international open source community

Dec 2017 Alibaba Cloud became Platinum Member of Cloud Native Foundation

April 2018 Alibaba Cloud won the 2018 MySQL Community Awards: Corporate Contributor

May 2019. The Apache Software Foundation (ASF) announced Apache® Dubbo™ as a Top-Level Project (TLP). Apache® Dubbo™ was originally developed at Alibaba and open-sourced in 2011. The project joined ASF in February 2018.

Major open source projects

1. **AliSQL** is a MySQL branch originated from Alibaba Group. It is based on the MySQL official release and has many feature and performance enhancements. AliSQL has proven to be very stable and efficient in production environment. It can be used as a free, fully compatible, enhanced and open source drop-in replacement for MySQL. AliSQL has been an open source project since August 2016. It is being actively developed by engineers from Alibaba Group. Moreover, it includes patches from Percona, WebScaleSQL, and MariaDB. AliSQL is a fruit of community effort. Everyone is welcomed to get involved.

2. **RocketMQ** is a third-generation distributed message-oriented middleware launched in 2012, with an enterprise version used by over a hundred companies and institutes. Host to all of Alibaba’s circulation messages for the Singles’ Day sales, RocketMQ delivered 1.2 trillion messages for the 2017 event, peaking at 170,000 per second. In September 2017, the Apache Software Foundation announced that Alibaba’s RocketMQ had achieved the status of a top-level project (TLP).

3. **ApsaraCache** has been open-source since Redis. It is a Redis branch of Apsara DB used by the Alibaba Cloud, Alibaba open-sourced ApsaraCache 2.8 in Oct 2107 with certain v3.0 functionalities to provide a stable, smooth performance to users. ApsaraCache is commonly used across various web categories, including live broadcast, gaming, and lifestyle. Open-sourcing ApsaraCache aims at facilitating long-term growth and improvements, and reflects Alibaba’s development capabilities. The upcoming version will remove host stability bottlenecks caused by frequent AOF rewriting and include full-volume synchronization to mitigate replication interruptions from weak network connections.

4. **Pouch** is a rich container technology developed in-house by Alibaba for various system services offering safety-isolation protection and isolation dimensions. Currently, over hundreds of thousands of Pouch containers support large-scale co-location of internal offline and online operations. Traditional download methods often cause download failures, where large file sizes make retries inefficient, and mass downloads result in file collapses on the client side.

5. **Egg** is an enterprise-class, basic web framework developed to overcome deficiencies found with Node.js. When micro-service architecture is adopted, specific tasks such as service granularity, API interface and development maintenance are involved, and the front and back-end face the conflict of flexible user experience vs. back-end service general utility. Node.js

has managed to overcome this widely-criticized weakness, and npm package is now the most popular peer, with vast improvement in single-threaded problems, avoidance of burdensome callbacks with the help of ES norm, a type system based on mature TypeScript, and better overall performance.

6. **Alibaba's development of Java specifications** and corresponding automatic testing tools (DE test plug-in components [IDEA, Eclipse]), is the result of a wide-spread collaboration between a number of senior Alibaba technical staff, has enabled developers to implement specifications more conveniently and quickly. The plug-in component scans the codes and shows ones which are not compliant with the specified blocker/critical/major levels in the lower part or even the IDEA. Real-time testing functions are also provided based on the inspection mechanism, and quickly spot issues when coding. For the former codes, one-key repairing functions are employed for certain rules.

7. **AliOS Things**, the integrated platform used by all Alibaba Cloud IoT business devices, has been transplanted to 21 types of chips from 17 domestic and foreign chip makers, including STMicroelectronics, Espressif, and Beken. After being made open-source this October, AliOS Things has released 3 major versions and regular updates, rich documents, and supports an active developer community.

8. Further to AliOS Things, Alibaba open-sourced **AliOS Lite** in February 2018, a lightweight version of aliOS based on JavaScript that supports all intelligent processes for IoT devices with lower CPU performance capacities, such as smart watches and cameras. When the industry is struggling to find IoT-based operating systems that are highly compatible and pose lower barriers to development, AliOS Lite provides processing solutions for devices with memory as little as 256 MB, and guarantees uniform action for various APIs, all while integrating Alibaba's proprietary machine intelligence capabilities in face recognition and image classification.

9. **Atlas**, a flexible Android development framework, dockerizes to solve large-scale team cooperation issues, which is suitable for the development of large or small-sized apps on system versions higher than Android 4.0 and supports dynamic publishing. The industry usually uses plug-in component frameworks. However, Atlas adopts modularization, which divides businesses and reuses common parts in consideration of reusability. The project adopts the approach of OSGI (Open Grid Services Infrastructure (OGSI)) and separates business to independently form a bundle for isolation and decoupling operations, so that concurrent development, fast iteration and dynamic deployment can be realized.

10. **JStorm** is an enhanced version of Apache Storm, built from scratch in Java instead of in the Clojure programming language that forms the code base for Apache Storm. JStorm provides a distributed programming framework very similar to Hadoop MapReduce, but unlike Hadoop MapReduce, runs a 24x7 topology which is ideal for streaming data and real-time computation. JStorm also guarantees fault-tolerance, with a new worker process spawned once an existing worker process crashes. JStorm is also 20% faster, on average, than Apache Storm under any scenario.

11. **Dubbo** is a high-performance, java based RPC framework open-sourced by Alibaba. As in many RPC systems, dubbo is based around the idea of defining a service, specifying the methods that can be called remotely with their parameters and return types. On the server side, the server implements this interface and runs a dubbo server to handle client calls. On the client side, the client has a stub that provides the same methods as the server.