

5 years! My opinion on OpenStack

Fu Guangping

In recent years, new technologies have emerged, and technologies such as containers, AI, blockchain, and SDN have been overwhelming. Many OpenStacks have begun to turn to other technologies, and the activity of OpenStack has not been as high as in previous years. Recently, there have been frequent PRs or articles such as the replacement of OpenStack by certain technologies, the abandonment of OpenStack by certain companies, and the middle-age crisis of OpenStack. Many people are worried about the future of OpenStack. It has been more than 5 years since I used OpenStack as both OpenStack Developer and OpenStack User. I want to share my own opinion.

In recent years, everyone is mentioning IOE, and self-controllable. I think the development trend of computing resources in the data center infrastructure will be X86+KVM+ open source cloud such as OpenStack/Kubernetes, which is the most controllable and fastest way many traditional IT companies can take. The strength of OpenStack is in IaaS. I think this one has been done quite well, at least it can be used, enough to meet most of the functional requirements of users' centralized resource management, resource distributed scheduling and distribution. Nova, Neutron, and Cinder are more and more stable and have a lot of features.

From the user's point of view and product perspective, OpenStack's peripheral systems such as management interface, monitoring system, log system, deployment tools, and etc. are not really perfect, but I think this is understandable. Peripheral system companies have their own set of standards. It is impossible for the community to develop a tool that can meet the needs of all enterprises. It is also impossible to use the enterprise integration solution that is out of the box. OpenStack has never planned to do this. The community provides an open API that makes it easy for companies to integrate existing peripheral systems or open source tools such as puppet, logstash, pacemaker, haproxy, and etc.

Those systems which can not set as standards are difficult to be open source, as they might be tied to one or a few homes. Take the example of OpenStack's highly availability. Assuming that the community decides to use pacemaker, many companies are definitely not willing to use it. They must consider the questions, why should we use pacemaker? In case there is a problem with the pacemaker, what should we do? There are a bunch of people who complain OpenStack. We now have a set of OpenStack using a layer of virtual machine to achieve high availability of virtual machines. The problem of memory leaks has not been solved. After eating memory, we have to restart periodically. So the community can't do this, but in the documentation there is an introduction to using OpenPace to make OpenStack high availability, which is enough.

I agree with the UNIX design philosophy, so-called Do one thing, and do it well. There are no commands or tools in Linux that can solve all your problems, but in combination of the commands, it can. And the OpenStack API is like a pipeline. Someone had previously complained about Linux. Why isn't there a standard IDE, set development, compilation, debugging, and packaging? I argued at the time that there was no need for this. I wouldn't say how many tools would be so bloated, and the way I made it would be sure to be all picky. Would users like it? So Linux does not currently have a standard so-called official IDE, but there are many open source and vendor-developed IDEs. Many of these tools are inseparable from vim, gcc, gdb, such as DevC++. OpenStack is like the current Linux kernel. You should never expect OpenStack to be an all-in-integrated-solution. It will help you to make a mirror and install thousands of nodes in a U-disk, and then with it you can monitor all systems and help you handle all kinds of logs. What do you want manufacturers to offer? First of first, the OpenStack community is not a company, don't take OpenStack as Party B, it is not a coolie. As far as people's nature is concerned, everyone wants to do core functions rather than those who are not thankful, why do you have such expectation to a community?

[Regarding OpenStack stability issues, many people do not distinguish whether the root cause is a failure from lower systems or a platform failure. I have encountered](#)

many problems in solving customer problems before, and found that many of them are caused by the underlying network, storage, middleware (message queue, database). In some old versions of OpenStack, the virtual machine may be stuck in a state due to Nova's message loss or the state of Nova and Cinder volume is inconsistent, which may cause the disk to be unloaded. The new version of OpenStack itself has improved a lot. In addition, each component of OpenStack can basically achieve A-A high availability. For example, the API service implements instance redundancy through load balance, and the RPC service itself supports multi-instance deployment. In the new version, cinder-volume also supports A-A high availability through distributed lock. With the shared storage, nova-compute hangs the virtual machine and can also trigger the evacuation migration to implement the virtual machine HA. Coupled with the monitoring system guarantee, the probability that the real OpenStack platform itself will hang up for no reason is not big. Even if some services of OpenStack hang up, the underlying storage and network will not affect the business operation. In other words, to ensure the stability of the OpenStack platform, you first need to ensure the stability of the underlying storage and network. Storage instability can be changed for commercial storage, ovs instability can use Linux Bridge, and even the hardware SDN can be used to replace Neutron. The whole core only retains the API. In short, which one can't feel it, just change it and feel that it can't be fooled. OpenStack is not bound. You are free to combine, and the technical selection in the early planning is very important. Everyone often compares the stability of VMware and OpenStack. Many people say that VMware is very stable, and OpenStack is a problem. For the time being, is it reasonable to compare between OpenStack and VMware? Do you think the underlying hardware is quite comparable when comparing? How much money for VMware, with Netapp, EMC commercial storage, and how much OpenStack with open source free Ceph storage is. When Ceph was not very stable in the past few years, there were often problems with clusters being unavailable, and many people blamed OpenStack.

Of course, there is not a 100% stable system. OpenStack is already a very complex distributed system, and the technologies covered include a huge technical ecosystem of hardware, operating system, network, architecture,

distributed, virtualization, etc., which need a strong level of operation and maintenance. Many people say that OpenStack is too complicated and difficult to operate. I would like to ask which distributed system can be done without complicated and difficult to operate. Hadoop and Spark which system is really operated and not delaminated. The root cause of Operation and Maintenance OpenStack is not because it is complicated, but because of its nature. As an IaaS solution, it requires the skill level of the operation and maintenance personnel to include traditional servers, data centers, distributed, virtualized, and network. Architecture, operating system, automation, database, high availability, message queues, etc., both traditional and modern technologies. High-level services such as Trove also need to understand the operation and maintenance of the database. In fact, many operators and implementers who can't hold OpenStack are really difficult to have these capabilities at the same time. They often don't understand the network of the operating system, do not understand the operating system of the network, understand the network and only understand the underlay tradition. The network does not understand overlay, do not understand SDN, even if it is difficult to find a system that understands the operating system and understands the network, does not understand the underlying virtualization, does not understand the database, relying on a few people can not be able to stand up. Don't think that creating a DevStack or referring to an online tutorial to deploy an OpenStack feels that OpenStack is very simple. I feel confident that one person can get everything in OpenStack. When I have a real accident, I can't find a clue. Is it impossible to determine the OpenStack platform? It is obvious that the underlying network is faulty or the storage is hung. When you see a bunch of OpenStack services ERROR, OpenStack is really bad. Really able to hold OpenStack, must rely on multiple teams to coordinate operation and maintenance and support, big companies are okay, many small companies are difficult to have such a complete operation and maintenance team, often one person or several people are responsible for the communication of the entire project Implement the entire cycle of delivery, testing, operation and maintenance. At least one big company in China has done a good job, and the automated deployment and automated operation and maintenance platform has done a good job. At least our environment has not had any major problems.

Users should treat and accept OpenStack fairly. They should not be prejudiced. I have encountered many customers who started to vomit OpenStack instability, unreliable, and zero tolerance as long as OpenStack has problems. But vcenter out of the question feels taken for granted, after all, there is a large company endorsed by Vmware, and who is endorsing OpenStack?

Also talk about some other OpenStack IaaS+ projects and PaaS related advanced services, such as Trove, Magnum, Sahara, vision and design concepts are very good, very valuable, expand the ecology, but also reflects the community's open mind, mind Patterns and ecological awareness, rather than having to threaten to replace who is fighting for your life, but some projects are not really good, for many reasons. The user thinks that OpenStack is mixed, and the energy is scattered. I think that there are not enough people (many people turn to the container, AI) and the user just needs to be less focused (the deployment of OpenStack will definitely use the virtual machine, but not necessarily the container and big data) There are two important reasons. After all, the higher the level, the more options you can combine, the more dispersed the users, the fewer people you use, the less feedback you have, and the fewer developers you invest. The many OpenStack projects are basically One or two people insist on maintaining.

[OpenStack has never promised to be a public cloud like AWS, so don't blame OpenStack comparing AWS. Think about how many years AWS has taken.](#) OpenStack also has no promise to make a private cloud better than Vmware. What OpenStack can do, in addition to OpenStack itself provides a core, what it can provide depend on the company's skills. I think there are several domestic public clouds based on OpenStack that are not bad, and some companies that do private cloud delivery based on OpenStack are also good. [Don't expect OpenStack to be a public cloud or provide cloud solution out of the box.](#)

[Some people say that the community is not well managed, and the utilitarianism is too strong. The Openstack community engages in the statistics of commits. It is](#)

not the technology itself that is involved in open source. Many of them care about the statistics. I think that "utilitarian" is partially reasonable. Open source also needs companies' support and contribution. Excluding a few of the open-source pure geeks, most companies are considered KPI, which is not charity, and have profit-driven. There are indeed a few people who opportunistically game the statistics, but this is a minority and does not represent the entire community atmosphere. Others say that the community is endlessly squashed. OpenStack members come from major companies all over the world. The inconvenient thing is not to talk about the community. Is it the individual of the company, and often does not endlessly wandering around, endlessly opening meaningful meetings?

Regarding the future of OpenStack, I am still optimistic. Although it is not as active as it used to be, it has been used by a large number of customers on a large scale, including all walks of life, such as the Internet, banks and operators, whether it is a private cloud or a public cloud, and is constantly expanding new clusters. In the new deployment, there will be no possibility of being completely replaced in the short term. Its development trend may be like Linux, even if it is no longer a hot topic on the dinner table, it will certainly exist in IDC everywhere. 10,000 steps back, if the real OpenStack, like the CloudStack of the year, is gradually replaced by another IaaS solution xxxStack1, xxxStack2, OpenStack will also exist in any xxxStack, just like the UNIX system, the form will always exist in Linux, BSD.

Also, for me personally, I sincerely thank OpenStack. From the beginning of the research to toss OpenStack, through OpenStack let me gradually understand and understand cloud computing from the white cloud that does not pass the cloud computing, and understand how cloud computing works and land. In OpenStack development, I also learned a lot of engineering knowledge and skills, how to communicate between distributed system components, how to achieve high availability, software architecture design concepts and so on. At the same time, I also met a lot of ancestors in the field of cloud computing and learned a lot of skills from them.

Originally, the original title was "I am the name of OpenStack". Later, I thought that the word "positive name" is subjectively targeted. After all, there are a thousand Hamlet in the eyes of a thousand people. There must be thousands of different opinions and controversies. I don't think there is anyone who is right or wrong, and there is no [difference between right and wrong](#). [Everyone has different perspectives.](#) , [different demands](#). It is also a good thing for everyone to argue. It is more than a few times to express their opinions. Maybe everyone understands each other better. It also allows everyone to truly understand OpenStack from more dimensions and expose some problems that exist. Tooth decay]. Therefore, this article only represents my personal opinion, there is no crusade, who is the name. If you agree with my point of view, please like me and let me know that someone has the same opinion as me. It is also a good thing to make a debate for everyone. We will fight for a few times and express our opinions. Maybe everyone will understand each other better. It will also enable everyone to truly understand OpenStack from more dimensions and expose some problems that do exist. [呲牙] If you don't agree, I hope that I can point out the disagreement in the comments. If I bump into a collision, I might get a spark:)

about the author:

Fu Guangping, net name int32bit, graduated from Beijing University of Posts and Telecommunications, contributor to OpenStack code, knows the column "OpenStack" author.