



Minimizing the Risks of OpenStack Adoption

White Paper

Introduction

Over the last five years, OpenStack has become a solution of choice for enterprise private cloud deployments, with a lot of mindshare from both developers and IT.

Some of the reasons enterprises prefer OpenStack are:

1. Open APIs, no vendor lock-in
2. Compatibility with the widest sets of server and storage solutions
3. Well-defined APIs, similar to AWS, with rich DevOps tool integrations
4. Community support and training

This has prompted several verticals to adopt OpenStack for their use cases.

OpenStack.org lists 9 different industries:

1. Web / SaaS / e-Commerce
2. Education, research and government
3. Information Technology
4. Healthcare
5. Retail
6. Finance
7. Manufacturing
8. Film / Media / Gaming
9. Cloud hosting / MSP

In each of these segments, companies that are at the forefront of adoption leverage a lot of benefit from having an agile, open-API based, on-premises cloud for their developers, R&D staff and IT. Walmart, AT&T, Comcast, BMW, Time Warner, NTT and Yahoo! are among the companies running tens of thousands of cores in an OpenStack-based cloud. This has drastically reduced their infrastructure expense and improved the agility of their IT processes. However there are many organizations still sitting on the sidelines, who are not sure if OpenStack is ready for them or how should they adopt it.

OpenStack is gaining popularity because of open APIs, rich DevOps tools integration, and wide support for server and storage devices.

Common Use Cases

Like any new technology, there are specific use cases where OpenStack is best. OpenStack is a framework designed for converting a set of compute, storage and network equipment into an API-driven cloud. This allows businesses to achieve the agility, speed and scale they need to stay competitive. Some of the most popular use cases across different verticals include:

1. Application development (CI/CD)
Providing a self-service cloud API for developers, and CI/CD tools. Jenkins and other tools have plugins to directly work with OpenStack APIs to create worker nodes and run a complete CI/CD pipeline. Developers and research staff members can also provision workloads without going through IT. IT teams can create projects with well-defined quotas for their teams and let them play within those limits.
2. IT as a Service
IT teams want to move away from a ticketing-based system to a self-service model. A new workload can be deployed in minutes without going back to storage and networking teams to get the underlying resources. With software-defined networking, one can create private networks, load balancers and firewall rules without involving the networking team to configure physical switches and routers. Similarly with software-defined storage, one can use business policies to carve out some storage pools and use them for workloads.

3. **Big Data Applications**
Running big data applications in a cloud environment has several advantages. Developers can create, expand and destroy big-data clusters on the fly. Some of the supported applications include Hadoop, Cassandra, Spark, and ELK stack. These can be deployed within minutes using predefined templates, adding a lot of agility to the R&D processes that many industries are running to extract value from their data. For example, BioTech companies can provide this cloud to their data scientists for medicine discovery; energy companies can use this for making exploration decisions; or media companies can use this platform for digital transformations.
4. **Network Function Virtualization (NFV)**
Network functions used to be run on appliances that included a combination of hardware and software. This is very time-consuming and hard to scale for most companies. Given that most of these appliances are x86 based, cloud becomes a natural platform to run all of these network functions as VMs. One can even create pipelines where data packets can traverse many such functions for different kinds of testing or transformations. Telcos like AT&T, Verizon and others are leveraging OpenStack across dozens of datacenters for this use case.
5. **SaaS Hosting**
Software is being converted to a service in many sectors. For example, we consume email, productivity software, CRM tools, shared data storage and even tax filing software as a service these days. SaaS provides two fundamental advantages: ease of use and constant updates. This reduces the burden on the consumers as well as vendors. All of the SaaS offerings are hosted on a cloud-based infrastructure, so that they can be scaled based on demand. Many of them start on a public cloud but have since moved to on-premises

clouds for better control, performance and cost efficiency. OpenStack can be used to build this on-premise infrastructure.

6. **Platform for Containers**
Containers and OpenStack are very complementary technologies. One can build a cloud using OpenStack and deploy elastic container clusters on it. This allows enterprises to use both VMs and containers on the same underlying platform based on the application needs.

Complexity cannot be eliminated, someone has to deal with it.

Top Challenges & Solutions

OpenStack was created to enable enterprises to create an AWS-like public cloud on-premises. This requires a very different design compared to legacy virtualization vendors, who are trying to package their existing software to look and feel like cloud.

Challenge #1: High Flexibility

One feature of OpenStack is the very high configurability the framework allows in choosing server, storage and networking solutions. For example, one can choose servers from Dell, HPE, Cisco, SuperMicro, or Lenovo along with some white label solutions. Similarly, one can use storage from a large number of vendors including NetApp, Nimble Storage, EMC, HPE, Hitachi, or Dell Compellent. In terms of networking one can use switches from Cisco, Juniper, Arista, Cumulus and many others. There are close to half a dozen network controllers like VMware NSX, Juniper Contrail,

Midokura, PlumGrid, and Cisco ACI that can be used with OpenStack. This makes it a complex exercise to pick, create and manage the cloud. ***Just because something is flexible, it doesn't mean that you have to exploit it!***

Challenge #2: High Configurability

In building an OpenStack-based cloud, one can choose among at least few hundred ways of configuring various parameters in the configuration files of different services. This requires an expert team of IT professionals in an enterprise, professional support from OpenStack distro vendors or using a pre-configured OpenStack-based solution to get the most efficiency out of your cloud. ***Configurable is an asset as long as you know how to manage it well!***

Challenge #3: Monitoring & Operations

OpenStack allows you to build a cloud, but the monitoring software needs to be installed and managed separately. Many customers use open-source tools like Nagios, Zabbix or vendor solutions to do the monitoring. This adds another software component that you have to install, monitor and manage. Look for vendors that provide built-in monitoring tools with their solution, as that simplifies the installation and keeps the monitoring tool up to date with new features. That way, you will have a single place to call for support. Similarly for operations and troubleshooting, one can either use third party tools like Splunk or Datadog, or see if a vendor has built-in support for these. Also, these tools come in two forms: on-premises and cloud-based. The on-premises version has to be maintained separately and may not updated very frequently. Try to choose the SaaS-based version as much as possible. This will minimize your headache of dealing with the installation, management and running of this tool. Also SaaS vendors will add new features a lot more frequently and make them available to customers. ***Always remember, monitoring and operations are not built into OpenStack; you have to get them separately.***

Challenge #4: Scalability

Building a small cluster with 3 to 5 nodes is quite easy and there are several open source installers available to do it. This is a good way to play with the API and learn

about the cloud. However, many issues come up as one tries to scale the OpenStack deployment to dozens or hundreds of nodes. A larger infrastructure requires better testing, configuration and specific options to work well in practice. It is critical to have support from a vendor, even if you are doing it yourself. Otherwise, you can be stuck in a bad situation and it can take weeks to resolve it based on the expertise of your team. ***Never underestimate challenges that occur with scale!***

Challenge #5: Upgrades

Upgrades are an achilles heel for any software deployment. Upgrading a cloud solution becomes even more complicated with an increased number of services. In the case of OpenStack, you have services on the controller node like nova, neutron, cinder, keystone and others that need to be upgraded. In addition, each compute node has local services that handle compute (nova agents), storage (cinder agents) and networking (neutron agents). Some upgrades also involve a schema change in the persistent database required by OpenStack, and usually each release comes with a mechanism to upgrade the database. All these moving pieces make the upgrade process even more challenging. It is very critical to have support from a vendor in terms of upgrades, so that you have a validated and well-tested upgrade process. Also keep in mind that the more configurable your deployment is, the higher the complexity of the upgrade. It is best to keep your deployment based on some standard configuration and best practices. ***Upgrade risk and time is directly proportional to the number of diverse and multi-vendor components in your system.***

Challenge #6: Multi-vendor support

Given the flexibility and wide support for servers, storage and networking vendors, it is natural to think about building a multi-vendor solution using OpenStack. Some companies even sell a fully integrated stack with components from multiple vendors. Don't be misled by the simplicity of the first day install. What really matters in a solution is the ongoing effort of maintenance, upgrades and management. Look for solutions that minimize that burden on your IT team. Also, it is critical

to have a solution you can scale easily without going through a complex matrix of figuring out how to add different resources. **Pick a solution based on ongoing maintenance, fewer vendors, lower total support costs and flexible scaling.**

Challenge #7: Workload Migrations

Once you have a cloud up and running, one of the biggest hurdles is to migrate the workloads from an existing environment. Very few customers start with a greenfield deployment with no existing workloads. The workload can be running on a VMware-, Microsoft- or KVM-based virtualized environment, or on a public cloud. Most of these environments have a different underlying file format. Either you can do a complex multi-hypervisor deployment with OpenStack or you will have to deal with migration. In the case of a multi-hypervisor deployment, make sure to check for the exact support matrix for supported features on each hypervisor.

Overall, it is better to have a simpler architecture with fewer hypervisors, as each hypervisor other than KVM comes with its own management tools, requiring the management of multiple software stacks. For example,

using OpenStack on top of VMware vCloud is like putting one stack on top of another stack, doubling the complexity of upgrades, software management and troubleshooting. **Pick a solution that offers support for workload migration from private and public clouds.**

How to mitigate these challenges

So far we have looked at various challenges and some of the best practices that you should follow to mitigate each specific challenge. However, that does not make the overall decision easier since different solutions may not all point in the same direction. The figure below distills the decision-making process to a few simple characteristics of your deployment needs and your team's expertise to help you choose the best way to deploy Openstack and be successful!

Choosing the best path

The table below shows the recommended solutions based on few parameters. Follow these and you will avoid most of the hurdles in your journey to a cloud based IT model.

Deployment Scale	Configurability Needed	Internal Team	Operations Budget	Solutions
> 1000 servers	High	> 20 people	Large	Pick a distro, train your team to be experts
> 1000 servers	High	< 20 people	Large	Hire professional services
> 1000 servers	Low	> 20 people	Large	Pick a distro, train your team to be experts
< 1000 servers	High	> 20 people	Medium	Use a vendor with flexible turnkey solution
< 200 servers	Low	< 5 people	Small	Use a vendor with a hyper-converged turnkey solution and software-based management