

A Brief Look At

# OpenStack

This eBook discusses topics that anyone considering an OpenStack private cloud will find interesting. It explores OpenStack's market penetration and some of the key drivers for technology adoption. The book also outlines some use cases.



a publication of



# Table of Contents

---

Introduction — Understanding OpenStack

Chapter One — A User's Perspective

Chapter Two — OpenStack's Future

Chapter Three — Use Cases

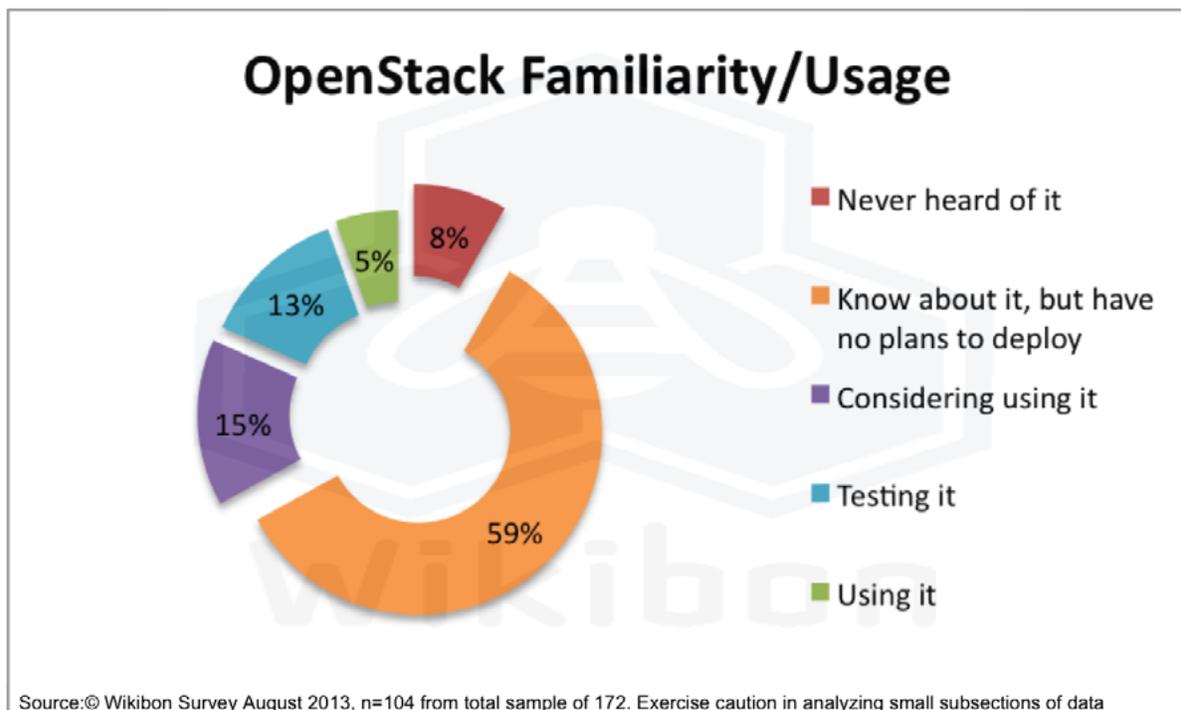


# INTRODUCTION

---

# Understanding OpenStack

The cloud computing revolution continues to change the way organizations view their IT resources. OpenStack cloud software provides a common, open source platform that can be used by corporations, service providers, researchers, and anyone else that wants to deploy a cloud.





# CHAPTER ONE

---



# A User's Perspective

## The Arguments Every OpenStack User Must Make

For some time now, critics have suggested that OpenStack may offer lower total cost of ownership, but fails to deliver on agility and uptime compared to other cloud services.

[A story on ReadWrite](#) quoted industry insiders who suggested OpenStack is merely a way to “fool” senior management into believing that a company has moved to cloud computing. TechRepublic, meanwhile, stated that OpenStack remains too complex to offer real opportunities for innovation.

We asked Sam Charrington, a cloud computing analyst and founder of St. Louis-based CloudPulse Strategies, how he would respond to these criticisms of OpenStack if he were making the business case.

This interview has been edited and condensed for clarity.

## Bright Computing: Why should organizations think about using OpenStack?

Sam Charrington: It is all about achieving greater agility. [OpenStack] is great to address questions such as, "What could we do as a business if it took us minutes to provision resources and not weeks?" Or "What could we do as a business if we could cut the time developers waste on chasing down resources and focus that on more creative activity?"

BC: How can you calculate return on investment (ROI) on an OpenStack implementation more accurately?

SC: *Both private and public clouds should be looked at for transformation, including team, business, and industry transformation. That is the true opportunity. Not everyone, however, sees it as this. If you dramatically shift the way business is done, deliver capabilities faster, or do things that you couldn't do otherwise, it is a softer ROI than being able to replace a bunch of VMware licenses with OpenStack. So you can get too focused on ROI and just look at cost savings, but that could mean you are missing out on more innovation and limit your options for bigger wins down the road.*

BC: What are some of the most common OpenStack implementation challenges?

SC: *The issue is that ease of use continues to get in the way of OpenStack adoption – it is harder to get started than many people would like, and as the platform grows with adding new projects, it is also not getting any easier. Having multiple choices for storage and networking makes it harder to install, but this also creates a lot of flexibility and power too.*

BC: Where is OpenStack weakest when compared to AWS or VMware?

SC: *They are very different beasts. OpenStack is more of a private cloud to date, but as more hosted implementations happen on OpenStack, they could be more comparable. The strength of an OpenStack-based public cloud is portability, and in using its APIs in multiple places. There really are limited on-premises options with AWS, while with OpenStack you can run both on-premises and in the cloud. Also, the AWS API is supported by a lot of tools because of its market presence and it is a lot more mature. VMware hasn't been successful in building its own public clouds and to make things more portable to date.*

BC: Where do you stand in terms of the debate on public vs. private clouds?

SC: *The typical enterprise is experimenting with both technologies. There still is a lot of confusion between virtualized environments and private clouds. And IT managers need to understand the degree of automation that is behind private clouds, such as provisioning, scaling, and management features.*

BC: If you aren't an Internet-native corporation like Facebook or Google, is there still a role that OpenStack can play?

SC: Yes, but it isn't because of your question's implication. I tend to think of OpenStack as a more commercially acceptable version of what Google and Facebook are building for themselves. These companies have their own internal management platforms and put a lot of time and effort into things that aren't really applicable for most enterprises. OpenStack tries to help enterprises that can't build these tools and tries to replicate some of the elements in these platforms. We are starting to see enterprises look more carefully at [continuous development](#), (although) there is a lot of [waterfall](#) development still being done. More of them are trying to transition to DevOps, but it requires a culture change that is significant. The big benefit here is to enable a higher degree of agility, and that makes them useful for enterprises that want to do this.



# CHAPTER TWO

---



# OpenStack's Future

## Why the Internet of Things Will Be Driven by OpenStack

The Internet of Things not only promises a plethora of devices, such as wearables and connected cars, but also guarantees there will be a deluge of data that requires processing. That means cloud infrastructures must keep pace, which experts say provides enterprises [an opportunity to leverage OpenStack](#).

OpenStack is a good foundation for any vendor or supplier of Internet of Things (IoT) platforms and services, according to René Büst, senior analyst and cloud practice lead at Crisp Research.

In Büst's opinion, an IoT platform provider could deliver enablement services, or ways to facilitate online connectivity, by offering OpenStack to a service provider or vendor of IoT products. Whereas large enterprises that have comprehensive experience running complex IT infrastructure could run their own OpenStack environments for IoT applications.

## Still early days

Right now, enterprises adopting OpenStack to support the IoT are looking at it mostly as a means to develop services, said Büst. These include mobile backend services, event hubs for connecting devices such as sensors, streaming analytics services, machine learning services, and edge location services.

“OpenStack is just an infrastructure base,” he said. “It gives you the liberty to basically run every type of owned software and service on it. The more (by default) value-added services that are provided by the OpenStack Project, the easier they are to build and deliver IoT projects.”

Büst said the OpenStack community can help the technology become a viable platform for IoT if they develop more high value-added services that encapsulate specific use cases, such as [OpenStack Sahara](#), which is intended to provide an easier means of provisioning Hadoop clusters to analyze big data.

Forrester Research predicts the IoT trend is about to enter a second stage in its development that will focus on software application platforms that provide pre-built connectivity, security, management, and analytics capabilities.

## An emerging cloud standard

OpenStack APIs are becoming a cloud standard and almost every public, private, and hosted private cloud provider either already has developed or is in the process of developing varying levels of support for the OpenStack APIs

because of its powerful ecosystem. There also has been enterprise adoption from reputable Fortune 500 companies, such as Fidelity, Comcast, and Best Buy.

For instance, Comcast has exhibited a very high profile about embracing OpenStack, which is one of the best examples of an enterprise using the platform to embrace IoT. Comcast has built a private OpenStack cloud to support video streaming applications and communication between its network and customer set-top boxes.

But there are many other IoT scenarios that could be supported by OpenStack, such as transportation systems for tracking buses that also correlate with weather data.

Forrester analyst Lauren Nelson says that enterprises of all sorts are going to find themselves supporting IoT to varying degrees, depending on their business, and already should have started to set up cloud infrastructures to handle the coming influx of data that needs processing.

Building a private cloud for targeted IoT projects is a great way to begin, and some organizations have built a secondary private cloud for that purpose. Organizations are wary of putting valuable data and intellectual property in a public cloud, although one avenue enterprises are pursuing is building out their own OpenStack environment, on premises, and processing some data in a public cloud and bringing it back inside.

## The coming transformation

Adoption of cloud infrastructure such as OpenStack to support IoT, however, has not yet gone mainstream. Many organizations are just beginning the transformation to becoming more software-driven businesses and are doing little software developer enablement.

Nelson continued, “They’re not realizing they have to make changes now or they will be left behind. Even if an organization starts small, it takes at least 10 OpenStack engineers to create an environment—and there is a lack of available OpenStack expertise. They will need people who are dedicated to this and who can learn it.” Remember, even small, targeted IoT projects can escalate and become more expensive.

Enterprises should be aware of everything that needs to be considered before building or adopting a cloud infrastructure based on OpenStack for supporting IoT, including the architecture design and which IoT-ready and IoT-enablement services should be available.

Organizations, especially smaller or midsize ones, should consider if they really want to run an OpenStack infrastructure completely on their own, since OpenStack is complex and the right skills are still hard to come by. Instead, they should leverage OpenStack from a service provider that is managed or OpenStack-ready.

# What Consolidation in the OpenStack Sector Means For Private Cloud Adoption

There is a lot of cash being poured into emerging OpenStack firms, which is a good sign; it proves that there's plenty of interest in enterprises adopting the technology.

In fact, [451 Research](#) said in its OpenStack Market Monitor & Forecast, which focuses on the key public and private vendors that directly provide OpenStack offerings, that the OpenStack market is expected to reach \$3.3 billion by 2018.

"OpenStack has seen tremendous growth over the last four years in terms of investment and community expansion," wrote Al Sadowski, Research Director for 451 Research's Service Provider channel, who authored the report. "The open-source platform is increasingly a consideration for private cloud deployments, and the business models within the ecosystem continue to evolve."

451 Research is not the only one to make such bold forecasts. This year, predicts Ken Rugg, CEO of database-as-a-service vendor Tesora, the number of OpenStack projects in production or being planned will exceed the number of VMware clouds.

Others have suggested that public/private hybrid clouds would be a focus for cloud platform vendors in the first quarter of this year. An increasingly fragmented container market, as competitors rise to challenge Docker's dominance, will nonetheless continue its move to the mainstream.

[Bright Cluster Manager for OpenStack](#) brings a proven cloud / cluster management software solution to the OpenStack community, thereby making it even easier to deploy a complete cloud over bare metal and manage it effectively.



# CHAPTER THREE

---

# Use Cases

## Some Ways OpenStack Can Help Your Business

### Hybrid Public / Private Cloud Solution

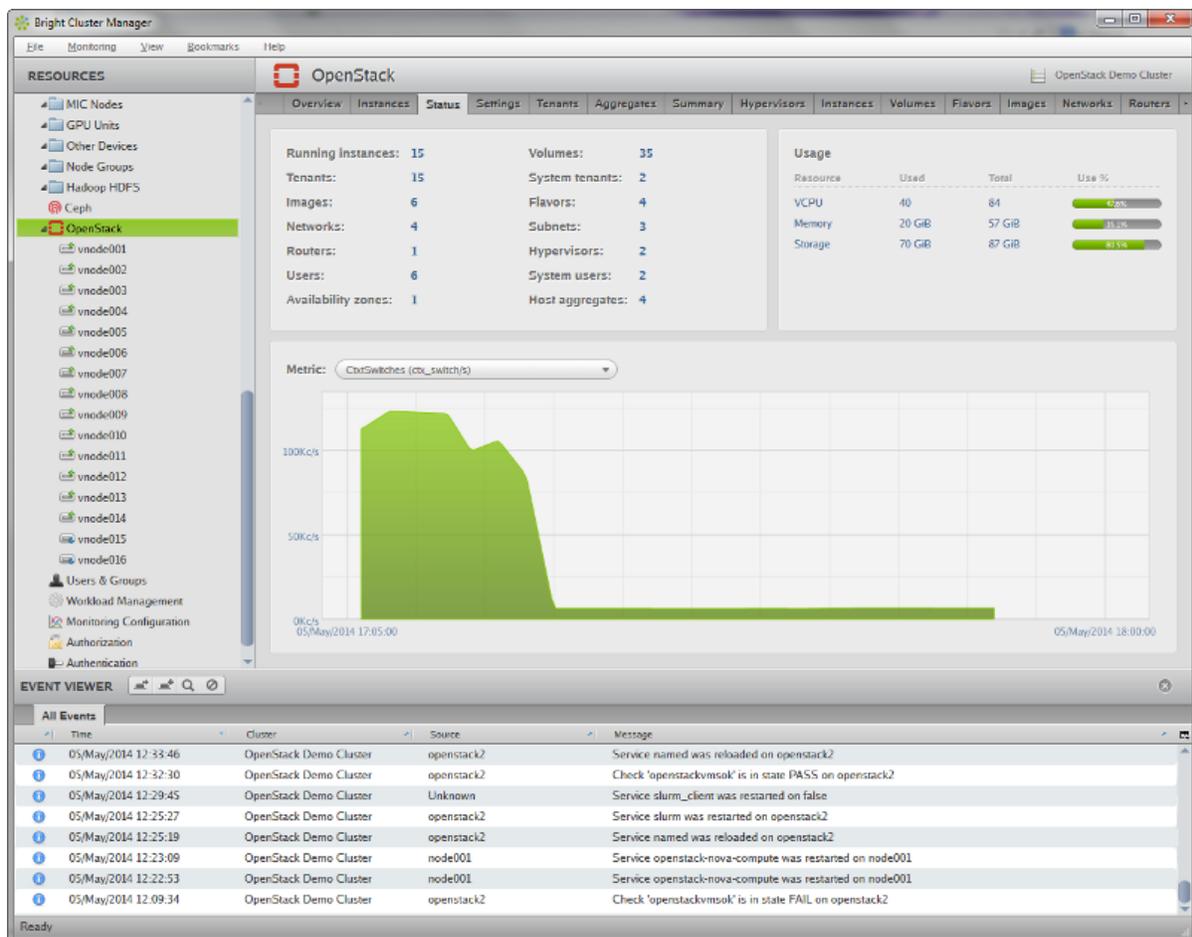
Many organizations are looking to build a private cloud infrastructure in their own data centers in order to reduce their dependence on expensive public cloud services. Using Bright's OpenStack solution puts this capability in the hands of virtually any organization. Bright Cluster Manager for OpenStack can turn a rack of servers into a working cloud in less time than you'd imagine, leaving you free to start deploying services. Now you can choose to run applications in a public OpenStack cloud or your own private cloud.

### Build and Deploy a Private Cloud in your Data Center

Switching your data center from discrete services to a cloud can let you get the most from your equipment. Bright's OpenStack solution makes it easy to build a working cloud, optimizing the utilization of the servers in your data center while maintaining the service levels your customers demand.

# Expandable Cloud Infrastructure to Meet Expected Growth

If your business plan calls for a cloud infrastructure that grows over time, you need a cost effective way to do that. By using Bright Cluster Manager to build and manage your OpenStack cloud right from the start, you can eliminate much of the cost in both resources and time. The task of adding nodes, loading them with the right software, and configuring them properly to join the cloud becomes quick and easy.



Here Bright's Status tab summarizes an OpenStack private cloud in terms of instances, tenants, flavors, software images, networks, users, and more.

---

# SIGN UP FOR A FREE DEMO

Talk to one of our experts and see how OpenStack can help you get the results you really want.

FREE DEMO

---

Bright's solution for OpenStack brings our proven cluster management software to the OpenStack world making it easy to deploy a complete cloud over bare metal and manage it effectively. Our solution leverages Bright Cluster Manager's capabilities to provide single pane-of-glass management for the hardware, the operating system, the OpenStack software, and users.

With Bright Cluster Manager for OpenStack, system administrators can get their own clouds running quickly and keep them running reliably throughout their life cycle — all with the ease and elegance of fully featured, enterprise-grade management software.



a publication of

**Bright** Computing