openmetal

Leveraging **On-Demand Private Clouds** to

Reduce Costs and Drive Innovation

EXECUTIVE SUMMARY

On-Demand Private Clouds emerged in 2018 from traditional industry names building from their history in on-premise closed source private clouds. Dell's VMWare and HPE's Greenlake are typical examples.

Open source on-demand private clouds emerged in late 2020 from hosting and cloud providers using open source systems for their infrastructure. InMotion Hosting, OpenMetal's parent company, was the first to release.

Currently available open source on-demand private clouds are based on OpenStack and Ceph. These two technologies have been commonly used for on-premise private clouds, have a long history, and are market leading names in private cloud. For example, OpenStack is the 2nd most active open source project in the world, just behind Linux itself. It is used by thousands of private companies, enterprises, and public cloud providers including some of the largest hyperscalers.

The typically cited issue with OpenStack private clouds is the complexity to establish it initially:

- The required skill set of a team to deploy and manage the cloud was a barrier.
- Risk of failure was high.
- Up front investment was necessary.

By eliminating the risks and time to production utilization, on-demand private clouds open a critical door for both smaller IT groups and large companies.

Deploying a private cloud has reached the critical "trivial" marker point. Time to production utilization has moved from being measured in quarters or months in 2019 to being measured in minutes and now just seconds.

Due to this ease of use, open source private clouds will see large growth of hybrid and primary uses.

Executive Summary Continued

Modern private clouds have extensive ways of consuming the resources within the cloud. Containers/Kubernetes, VMs, User Management, Block Storage, Object Storage, Orchestration, and Networking are all common. API first is the norm at the administrative and user levels. User GUIs and Usage Governance is built in.

Common management systems, like Kubernetes and Terraform, can utilize modern private clouds in the same ways they can utilize large public cloud providers. With the advent of on-demand private cloud, expectations are that new management technologies for ML, AI, and other workloads will continue to treat OpenStack as a first class resource provider.

Of note, productization and interoperability of major systems is a top priority for on-demand private cloud providers. Open source often comes with the dual challenge and benefit of "some assembly required". Prescriptive standards were embraced fully in 2020 and will accelerate throughout the years. Specialist providers have also emerged and best-of-breed technologies are being integrated into next generation offerings.

The prescriptive standard varies by provider, but OpenStack and Ceph are common. This allows companies to move between providers. Natively linking on-premises OpenStack and on-premises Ceph with their on-demand counterparts is available now as well.

Numbers vary, but **the cost of on-demand open source private cloud, including management by either company or the provider, is between 10% and 60% less than resources from the public clouds or closed source private clouds**.

Companies that are running their own data centers or in colocation but are not at scale will benefit the most. Companies in competition with Amazon and Google but that are also stuck paying for AWS and Google Cloud will both cut costs and cut having to pay high rates to a competitor.

In closing, the transition of private clouds to being on-demand is transformational both for the companies that can now use the private clouds but also for large public clouds. As new feature parity services become available from the "second tier" providers around the world, the expectation is that pricing pressure on the mega-clouds will drive costs down.

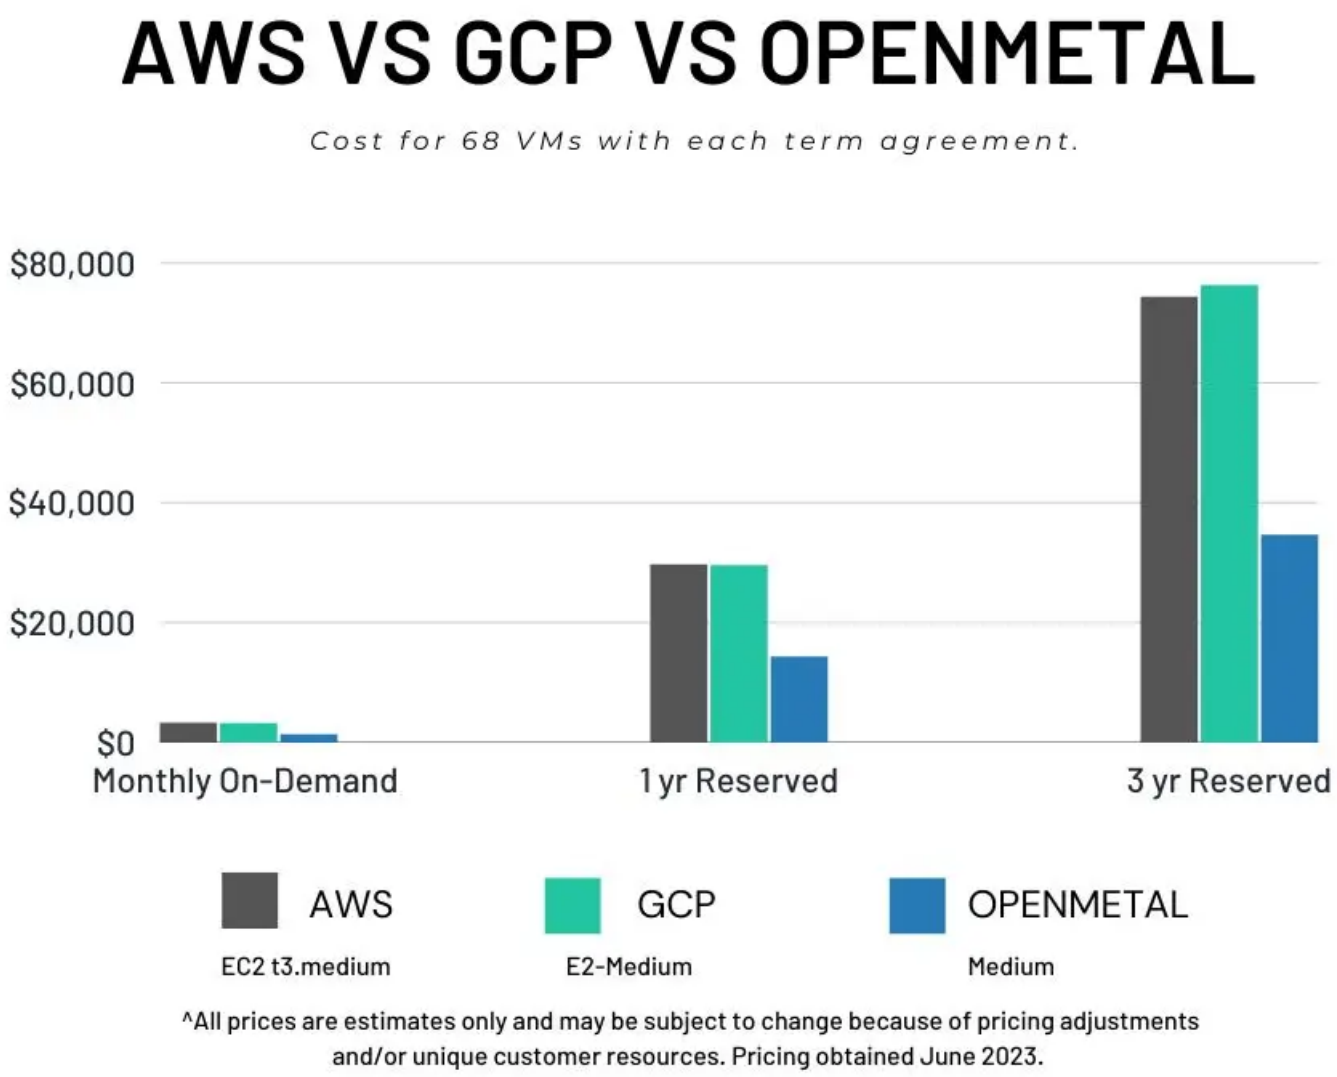

CONTENTS

ON-DEMAND PRIVATE CLOUD DEFINED

FIRST OFFERED BY INMOTION HOSTING

On-demand private clouds emerged in 2018 from traditional industry names building from their history in on-premises private clouds. VMware and Nutanix, for example, accomplished this through partnerships with large public clouds. Open source on-demand private clouds emerged in late 2020, first offered by OpenMetal's parent company, InMotion Hosting, with many mirrors currently under development by other key providers.

As on-demand private clouds are a newer macro service, its definition hasn't been fully settled upon, but application of Agile philosophy around delivery time is typically quoted. Consider the "time to production" as measured by a recognized need for resources to the launch of a VM or container to meet that need.

On a private cloud, time to utilization was **commonly measured in quarters before 2016**, then months in 2016-2017, then weeks in 2018-2019, then in 2020 it fell quickly from weeks to minutes - **about 45 minutes for a 3 server cluster on the OpenMetal platform.**

This "trivial to deploy" checkpoint of under one hour made using OpenStack private clouds even more accessible to teams of all sizes. But in 2022, we pushed the boundaries even further. We dramatically reduced the deployment time **from 45 minutes to a mere 45 seconds**. This speed truly embodies the essence of "on-demand" – provisioning a private cloud environment in near real-time and opening up a world of new possibilities.

FULL PROVISION TIME WENT FROM
QUARTERS TO MINUTES TO SECONDS

TIME TO PRODUCTION

On a private cloud, time to utilization was **commonly measured in quarters before 2016**, then months in 2016-2017, then weeks in 2018-2019, then up through 2022 it fell quickly **from weeks to minutes to seconds**.

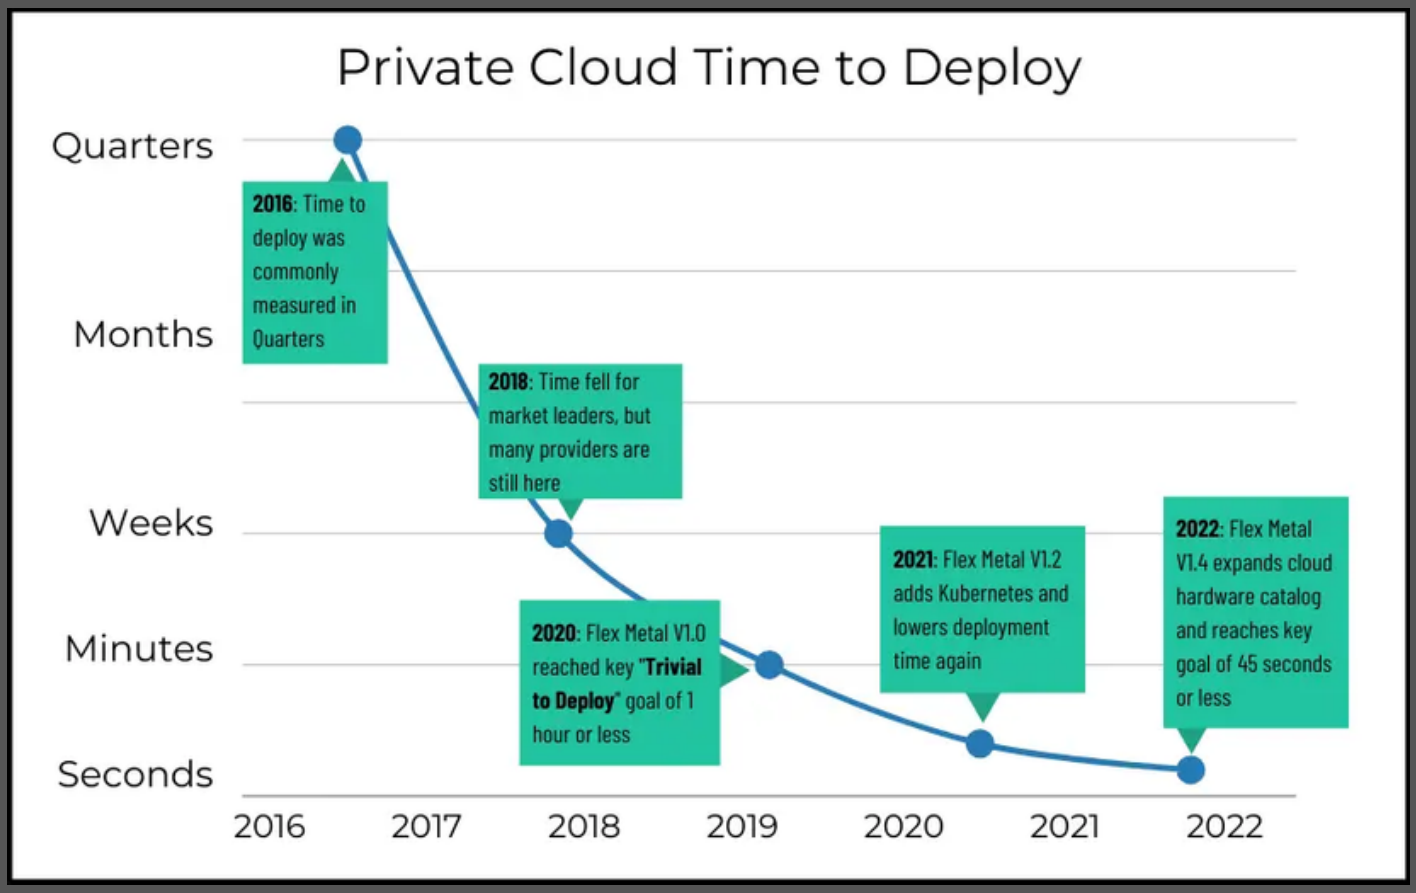

45 seconds

The typical full provision time is now about 45 seconds for a 3 server cluster.

OPENSTACK AND CEPH

The most robust and feature rich open source on-demand private clouds are based on OpenStack and Ceph. These two technologies have been commonly used for private clouds for many years. OpenStack goes back over 10 years originating from NASA and RackSpace.

Due to its passionate and active contributor base, OpenStack continues to be one of the top five most active open source projects. Over the past decade, over 2,200 organizations have voluntarily contributed OpenStack usage data and architecture details across more than 4,000 deployments, as well as feedback for the upstream community to improve the software.

OpenStack offers several advantages over proprietary platforms. First, being open source, there are no licensing fees, leading to significant cost savings. OpenStack boasts high scalability, allowing businesses to easily add or remove resources as their needs evolve. This on-demand nature ensures they only pay for what they use. Additionally, OpenStack's open architecture allows seamless integration with various technologies and platforms a business might already have in place.

Ceph's inherent scalability allows storage capacity to easily grow or shrink alongside virtual machines provisioned within the OpenStack environment. This elasticity ensures businesses aren't limited by rigid storage constraints. Ceph also eliminates vendor lock-in by functioning on commodity hardware, reducing overall storage costs. Ceph and OpenStack integration creates a unified cloud environment where Ceph acts as the high-performance storage engine behind OpenStack's virtual machines and applications - translating to efficient data management, scalability on both compute and storage fronts, and ultimately, a robust and adaptable cloud foundation for any business.

With all of its strengths, two considerable issues are typically quoted when companies explore creating a private cloud for themselves based on OpenStack.

OpenStack and Ceph Continued

First, to create a high-quality OpenStack requires a group of skilled System Engineers. This includes specialists in hardware, networks, security, and Linux. It is a stretch for a small business to have these skills on staff and unlikely a medium business will have more than one of them.

Second, even with the skilled group, most will not have experience with OpenStack. In order to learn to run a private cloud, the IT team has to convince their company to finance a "Pilot Program" of the potential cloud. Prior to on-demand OpenStack, those clouds would cost hundreds of thousands of dollars in server and network gear, plus 3-12 months worth of time. And with that, many, maybe even the majority, of the pilots never turn into a production cloud.

Many enterprise-focused companies, like RedHat, Canonical, and Accenture, successfully help enterprises bridge that gap economically. For smaller IT teams, they simply couldn't access the benefits and cost savings of private cloud.

With the advent of the on-demand private cloud providers, the two most considerable issues have been overcome. This change is expected to lead to extensive adoption by both smaller IT teams and by large companies. Those large companies have always recognized the value but couldn't get corporate backing due to the risks and heavy initial costs.

> "When it comes to choosing cloud, private clouds have always been the gold standard for greater control over security and costs, but complicated to set up. When these setups fail, it weakens the open source community around OpenStack. To drive greater adoption of open source systems, we felt it was imperative to simplify setup and ease the adoption of OpenStack. With OpenMetal, any company, regardless of size or skill, can now spin up on-demand private clouds with no risk or wait."
>
> **Todd Robinson, OpenMetal President**

DEFAULT SERVICES IN A MODERN PRIVATE CLOUD

Modern private clouds are feature packed and ready by default to allow for teams using this IT infrastructure to move quickly.

Containers, VMs, User Management, Block Storage, Object Storage, Orchestration, and Networking are all common.

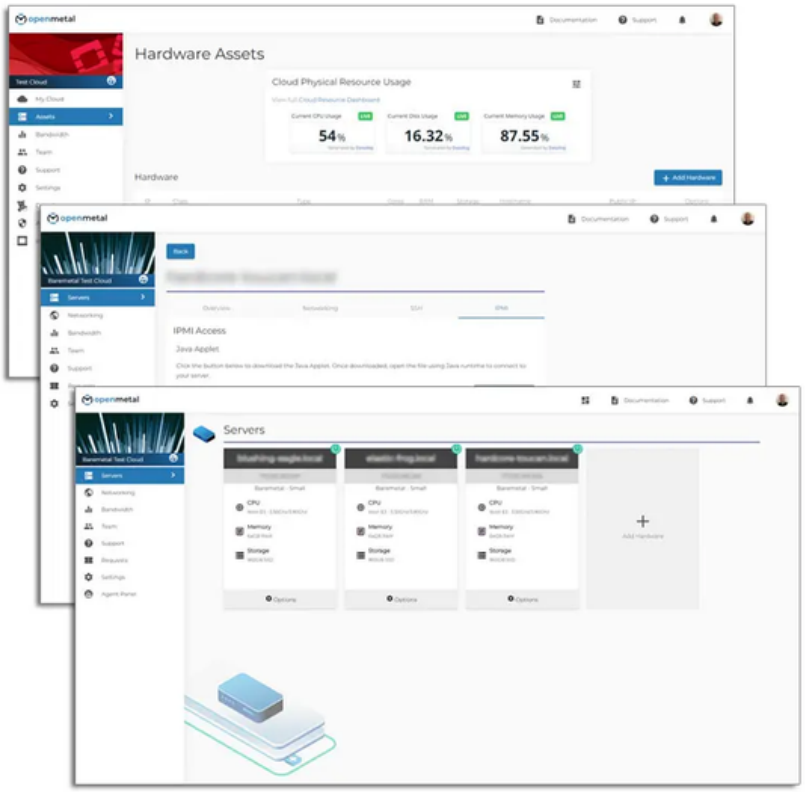

API/Orchestration Ready/GUIs

Modern private clouds are API-first systems. Automation of infrastructure is expected and encouraged. Tools like Ansible and Terraform are in heavy use on private clouds just like they are on public clouds. GUIs are also available for the most common use cases on private clouds.

Empowering Self Service

Private clouds are designed to accelerate access to IT resources. Solid self service processes are standard. In particular, OpenStack has a very robust User and Project management system. Each Project can access their resources through APIs or GUI. Projects have quotas for all resources and the cloud administrator. There is also a robust command line interface that uses the APIs.

Networking

The power of software defined networking can be overlooked, but it is a very important part of the security and easy of use model in private cloud deployments. It is trivial to create switches and routers. Current technologies include VxLAN which allows each Project inside of the cloud to have segregated private networks. Clouds with a sufficient number of hardware nodes can have redundant switching and routing easily.

Virtual Machines (VMs)

The mainstay of cloud is the Virtual Machine. As this is basic table stakes for a cloud, all providers offer VMs. The more advanced providers will offer time savers like including OS images to help improve ease of use. Unique to private clouds is the control of setting the flavors of VMs available for users.

Containers/ Kubernetes

Containers have become standard items also since 2020. Kubernetes is the leading orchestration system for containers and all of the new leading providers already offer support for Kubernetes. This may not be included from all cloud providers with the base cost of your cloud, but leaders like OpenMetal do, so it is expected that others will follow suit.

Containers are supported in two forms:

Containers

Containers have become standard items also since 2020. Kubernetes is common for smaller private clouds, where the containers/pods run within VMs. This method accomplishes two key requirements. First, since containers of one project are exclusive to each VM, the container security issues are addressed. Second, as traditional VM workloads are then safely co-mingled with container workloads, efficiencies of scale are aligned.

Kubernetes

The second form is more typical for large deployments of a single workload. Containers can be run directly on hardware. This is more efficient from a hardware resource standpoint. It is typical to graduate to this approach versus do this by default. Orchestration systems like Kubernetes do not provide safe co-mingling of containers of different users. It is expected the application or administrator planning the workloads will provide the proper segregation.

Block Storage

This is also a mainstay of cloud. Not all providers are created equal with block storage. There are several key items to understand and review when selecting block storage in your on-demand private cloud.

First, hardware considerations. NVMe SSD and spinners must be available options. SATA SSDs are also a typical offering and for many workloads are a solid fit. Spinners are often the only logical choice for large data needs with cost constraints. NVMes are the standard for other data sets. Though more expensive than SATA, they are significantly faster. NVMes are also critical for accelerating spinners.

Second, block storage can be delivered in several ways. Each has positives and negatives.

Block storage can be highly available network storage coming from Ceph or a provider like NetApp. This is an often-recommended solution as it has built-in data protection. It has significant IOPS considerations though as data is traversing the network and making multiple copies.

For the highest performance, block storage can also be a single drive LVM on NVMe. In this case, the application must provide the data redundancy. This is often seen with databases that have mature replication strategies.

Block storage can also be an LVM on top of RAIDed drives on the same server as the VM. This is a solid and familiar alternative when the application can not easily provide redundancy but does need high IOPS.

Key Items to Review

When Selecting Block Storage in Your On-Demand Private Cloud

Hardware Considerations

- **Sata SSDs**
  typical offering
- ***Spinners**
  for large data needs with cost constraints
- ***NVMe SSD**
  standard for any other data sets; more expensive, but faster

Delivery Type

- **Highly Available Network Storage**
  like Ceph or NetApp
- **Single Drive LVM (NVMe)**
- **LVM on Top of RAIDed Drives**

*must be available options

Object Storage

This is often an additional component of most private clouds but is offered by default on the leading providers. Object storage, from an API and usage standpoint, is very standardized so how it is used in private clouds is the same as in public clouds. There are two important factors to consider though with private clouds.

First, it is highly recommended to have object storage included with your initial cloud deployment. This allows for the object storage to be ready for use right out of the box. This is key to empowering teams to move quickly as they can immediately use the object storage without delays.

Second, with most, if not all, initial on-demand private clouds are hyper-converged systems. The object storage will be co-mingled with the systems providing block storage. In these systems, the performance requirements of block storage are the driver for what hardware is in the system. OpenMetal, for example, includes NVMe as the primary storage due to benefits of the extremely high IOPS.

Object storage typically does not need high IOPS and instead cost efficiency is much more important. As usage of object storage grows in your cloud, there will be a tipping point when you will need to add different storage appropriate to that need. The storage used in this situation is large spinning disks with optional NVMe acceleration and CPU compression.

Providers are quite different at this time for how this situation is handled. Look for cost per GB and options like compression on the fly, redundancy and acceleration of the object storage gateway service, and NVMe acceleration.

For OpenMetal, once the cost tipping point is reached in a hyper-converged and converged approach, two options are available based on scale. Both require an addition of three spinner-based servers to your cluster to provide redundancy and data protection. It is likely at that stage several converged units can be returned after migration for a net lower cost. Migration of data to new hardware is native to Ceph and nearly trivial.

AI, ML, and Data Science Applications

As hardware specific to these fields, like the A100 GPU from Nvidia, are added to the on-demand private clouds, we expect a surge in ease of use softwares to emerge. The trend of "MLOps" and "AIOps" systems automating against API first systems will continue. For example, Valohai, a MLOps automation system, offers easy workload processing on the public clouds and on OpenStack. Partnerships between on-demand private cloud providers and industry-specific automation companies will become more common as time goes on.

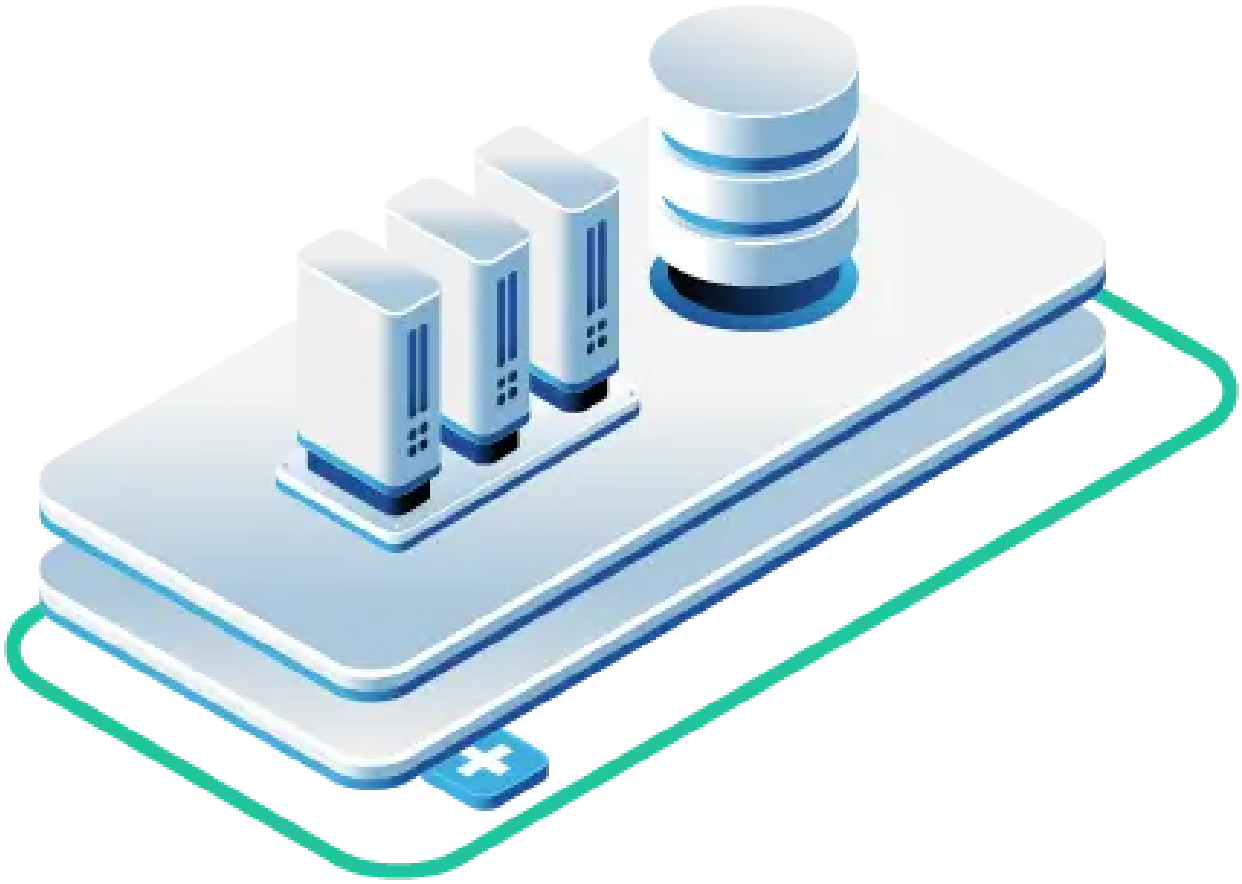

INNOVATION, EMPOWERING AND LEVERAGING STAFF

There are two areas where private cloud has a distinct advantage over public cloud:

Empowering Users

First, as key technologies are included within a modern private cloud, for no additional cost, it empowers users to leverage new technologies on their timeline and without complex red-tape.

If a development team wants to experiment with a Kubernetes-based software deployment, the models and system are already available to them within their current IT infrastructure.

Keeping Talented People

Second, many companies are facing the move from in house data centers or colocation to cloud. The current IT team has significant value and that value is often partially negated by the move to a public cloud.

Moving to a self-managed, on-demand private cloud leverages that existing team. This gives a company the ultimate in cost savings while keeping talented people who contribute much more than just IT services.

The top self-managed offerings will offer onboarding services, ongoing training, and third level support from the provider. They may also offer managed private cloud and starting with that level of service may be logical as well.

SELF-MANAGED PROVIDES THE ULTIMATE COST SAVINGS
AND LEVERAGES THE TALENTED PEOPLE YOU ALREADY HAVE

INTEROPERABILITY AND INTEGRATED SPECIALIST PROVIDERS

It is also important to point out a top priority for on-demand private cloud providers is productization and interoperability of major systems. Open source often comes with the dual challenge and benefit of "some assembly required". This can translate to costs and time delays quickly. This is a clear focus though for providers and prescriptive standards were embraced more and more in recent years.

The prescriptive standard does vary by provider, but OpenStack and Ceph are a very common overall system. This allows companies to move between providers with relatively small adjustments. This includes companies are able to natively link on-premises OpenStack with on-demand OpenStack. The same is true for Ceph and allowing on-premises Ceph to natively mirror to on-demand Ceph.

Equally important, many specialist providers have also emerged. This will continue, and top private cloud providers are working to integrate these best-of-breed technologies. For example, Fleio is an advanced user GUI for OpenStack. In days past, providers would develop their own control panel. Instead, providers are focused on improving the native OpenStack GUI, Horizon, or offering the advanced control panel from Fleio. This is a considerable savings for the providers and yields a higher quality user experience.

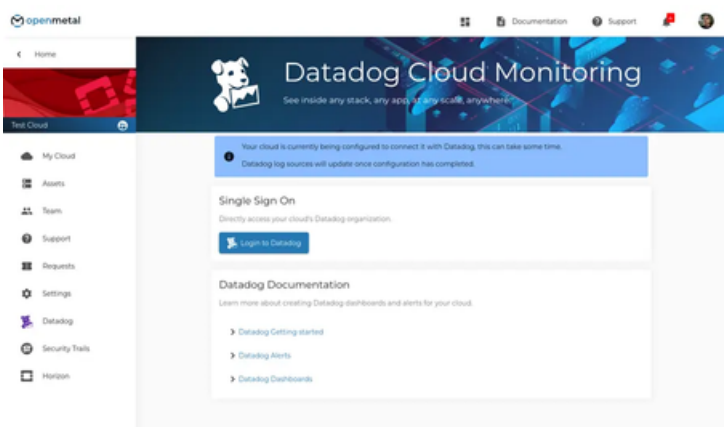

Finally, with OpenStack and Ceph's ubiquitousness, many SaaS providers have quick and easy integration tools available. For example, DataDog, a popular monitoring suite, has an OpenStack plugin that can monitor critical items and place them in a dashboard next to resources being provided from a public cloud. Our OpenMetal portal offers DataDog monitoring, which has been well-received and appreciated by customers.

COST COMPARISONS

Costs and benefits vary by a company's current methods to deliver IT resources. Generally though, the cost of on-demand open source private cloud, including management by either company or the provider, is **between 10% and 60% less** than resources from the public clouds or closed source private clouds.

For companies considering closing their own data centers and reducing headcount, the numbers vary widely. Efficient and at-scale private data centers will only have limited cost savings. Companies in this situation should consider providers that offer managed open source clouds within the customer's location.

Companies that are running their own data centers or are in colocation, but are not at scale or are struggling with bringing in new technologies have both options now available: on-demand in a provider's data center and/or bringing in a cloud provider to their data center.

There is a category of costs that are still being explored. It is common for companies that are using hyperscalers to actually be in competition with that company in some facet of their business. For example, many retailers use AWS, the high margin side of Amazon's business units. That high margin offsets the relatively low margin retail part of Amazon. By utilizing their AWS services, they are funding their competitor!

As on-demand private cloud becomes more common, expectations are that companies caught paying a competitor will quickly move to rectify that painful situation.

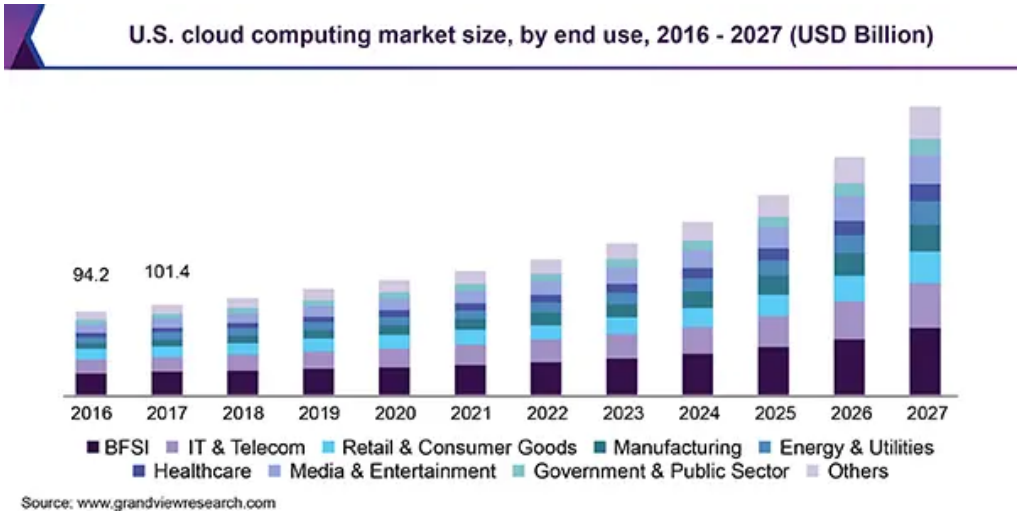

WRAPPING UP

In closing, the transition of private clouds to being on-demand is transformational both for the companies that can now use the private clouds but also for large public clouds. As new feature parity services become available from the "second tier" providers around the world, the expectation is that pricing pressure on the mega-clouds will drive costs down.

HOW DO I GET STARTED?

Speak with our Cloud Solutions Team to get help navigating the world of complex infrastructure decisions.

- **https://openmetal.io/schedule-meeting**
- **sales@openmetal.io**

Get $500 in Free Cloud Credits

Get started with free cloud credits for an OpenMetal On-Demand Hosted Private Cloud. Use the credits to test your Proof of Concept (PoC) before going into production or to get familiar with OpenStack.

CLAIM YOUR FREE CREDITS

openmetal

555 S Independence Blvd
Virginia Beach, VA 23452

EMAIL: sales@openmetal.io
WEB: https://openmetal.io

THE CLOUDX AWARDS 2023
BEST INNOVATION IN Cloud Infrastructure
WINNER

Open Infrastructure FOUNDATION
SILVER MEMBER

FinOps Foundation Member

LINUX FOUNDATION
SILVER MEMBER

CLOUD NATIVE COMPUTING FOUNDATION
SILVER MEMBER