

OpenMetal Cloud

# IT Leadership Buyer's Guide


---



# CONTENTS

---

<b>OpenMetal Cloud Background</b> .....	<b>3</b>
<b>Time to Production</b> .....	<b>4</b>
<b>Private Cloud Cores</b> .....	<b>5</b>
<b>Support Levels</b> .....	<b>12</b>
<b>Wrapping Up</b> .....	<b>14</b>
<b>Company Information</b> .....	<b>15</b>



This guide is intended for IT leadership and corporate executives. The advent of on-demand private cloud has blurred the line between public and private clouds. This guide provides a framework for this new technology that can increase competitiveness and lower costs.

# OPENMETAL CLOUD BACKGROUND

Use of open source private cloud has a long history with various benefits, costs, and risks. The risks and costs often outweighed the benefits for all but the largest companies or for hosting providers.

This changed significantly with the advent of on-demand private cloud. 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 three server cluster on the OpenMetal platform.

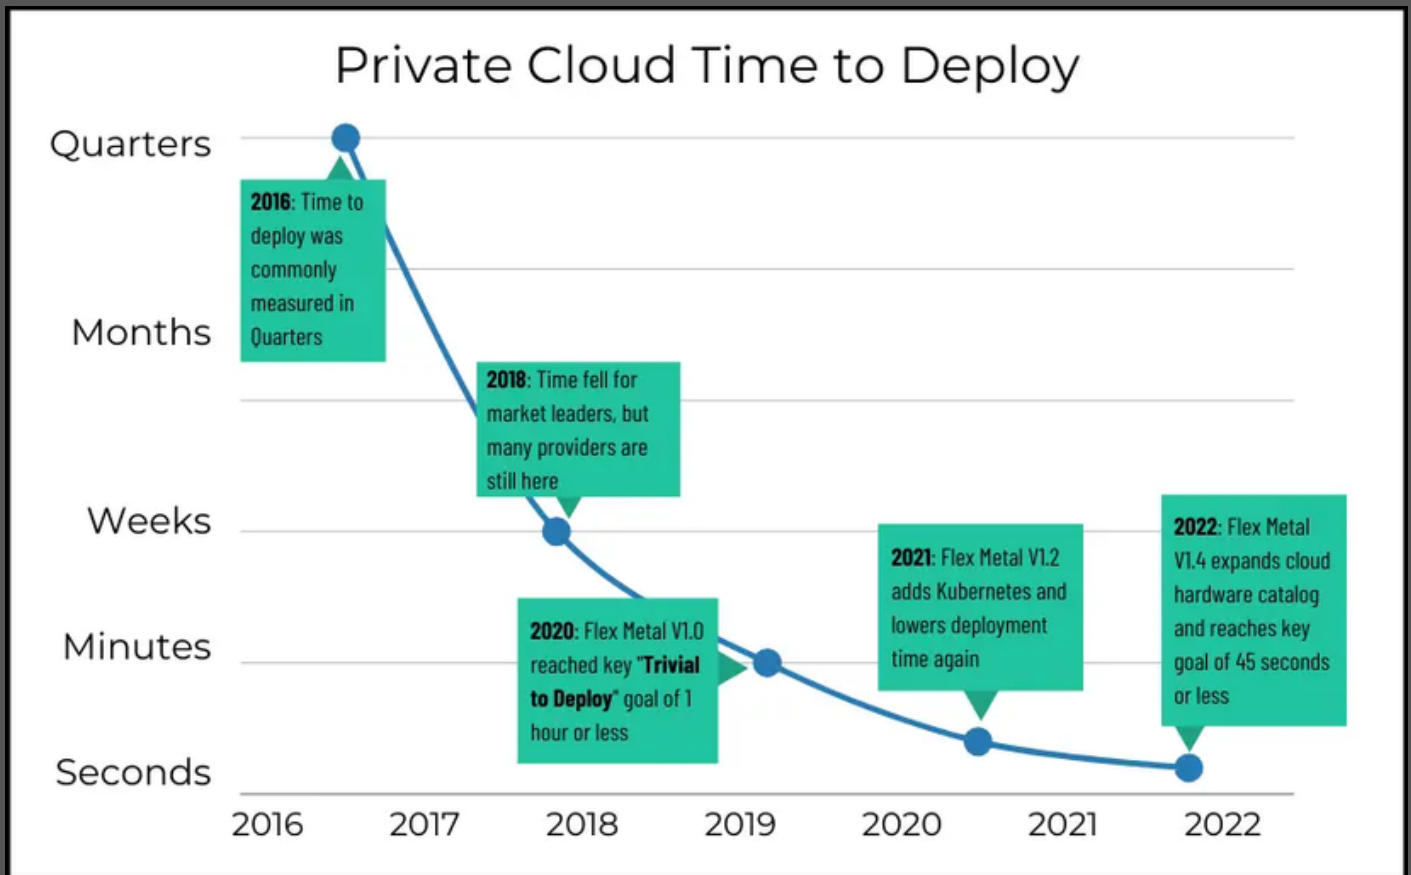
This "trivial to deploy" checkpoint of under one hour made using OpenStack private clouds even more accessible to teams of all sizes. But in 2022, we pushed the boundaries even further. We dramatically reduced the deployment time **from 45 minutes to a mere 45 seconds**. This speed truly embodies the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

This yields a cloud that includes support for containers/ Kubernetes, VMs, block storage, object storage, orchestration, and networking. User management and governance models are also standard.



# 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 three server cluster.

# PRIVATE CLOUD CORES

## How do OpenMetal Clouds work?

Since it is trivial to spin up and down full clouds, you now have the option to separate workloads into separate clouds based on usage patterns, security models, or geographic considerations. Larger clouds are more efficient, but smaller clouds have many unique benefits to consider.

At the heart of each of your clouds is the Private Cloud Core (PCC). The PCC is a hyper-converged set of three servers that provide all resources like VMs and Block Storage. It also includes the Control Plane which manages the cloud itself and provides Networking.

**There are four flavors of Private Cloud Cores across three generations of hardware: Small, Standard, Large, and XL.**

### SMALL

Smalls are ideal for testing, training, and smaller production implementations. A good example of a small deployment is a secure environment for critical data that is separated from your main IT infrastructure.

### LARGE

Larges are great for more demanding and resource-intensive needs as they include over double the RAM, more storage, and an extra CPU with higher specs than in the Standard.

### STANDARD

Standards are the typical building block as they provide significant resources as a hyper-converged cloud and are also ideal as a control plane for very large clouds. It is trivial to lower the workloads on a PCC if it needs to handle more control plane work.

### XL

Our XL option can support even the most demanding workloads. It offers even better CPU specs, along with double the RAM, storage, and included egress of Large.

# Private Cloud Cores Continued

## How do I control my clouds?

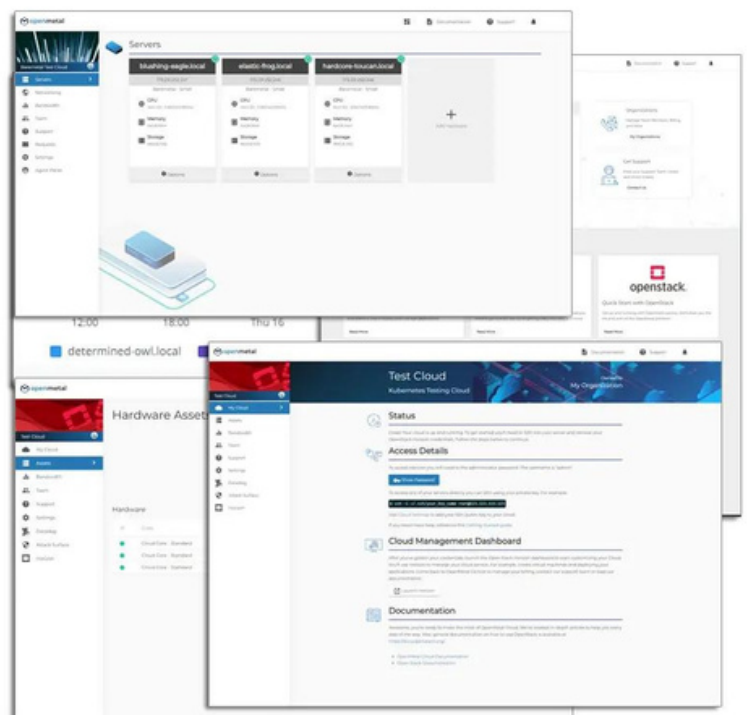
Your top-level control panel is known as OpenMetal Central. From within Central you can create clouds with a few clicks. **The clouds are completely autonomous from each other and have their own hardware, user interfaces, APIs, and public IP space.**

### From OpenMetal Central you can, per cloud:

- Set up billing restrictions
- Add and remove cloud administrators
- Add and remove hardware
- Access Horizon, the cloud's autonomous user interface
- Access billing history
- Contact support and account managers
- Delete the cloud

### From within each cloud you will have CLI, API, and GUI access to those cloud's resources. From within the cloud you can control:

- Networks including routers, switches, and access to public IP space
- Projects - this provides self service access for groups/departments/customers, etc.
- Project users
- Resource limitations for projects
- Object storage
- VM images and flavors available to users
- Orchestration



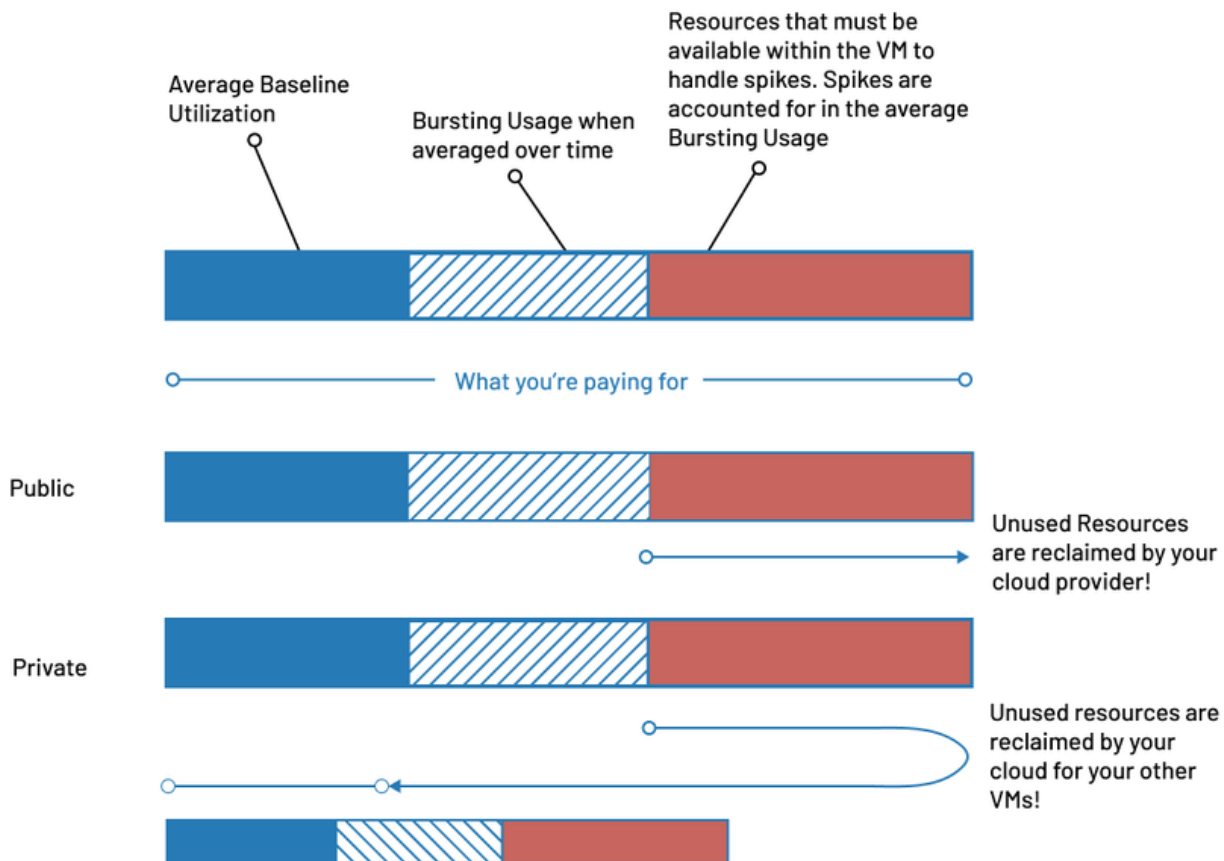
# Private Cloud Cores Continued

## How many VMs will I get in my cloud?

You may be surprised that public clouds “oversubscribe” their CPU’s threads by as much as 16x. RAM “oversubscription” ratios of 1.5x also appear to be in use. There is very good logic and real world data backing up this decision for them. For your cloud, you will control the ratios based on your anticipated workloads. Your Account Manager can help with this. We recommend 12x to start for CPUs and 1.5x for RAM.

To see an estimate of VMs, use our [Pricing Calculator](#). It takes into account if the cluster is hyper-converged and thus the control plane needs to have resources reserved before RAM and CPU can be used for VMs.

Of note, all servers have small boot drives that are separate from the main drives. To calculate VMs on your own, you take each server’s RAM minus 7GB of RAM for the control plane and 2GB of RAM per Ceph OSD (drive or drive partition) then multiply it by 1.5. Next you take each server’s CPU threads and minus 3.5 threads for the control plane and then multiply by 12.



# Private Cloud Cores Continued

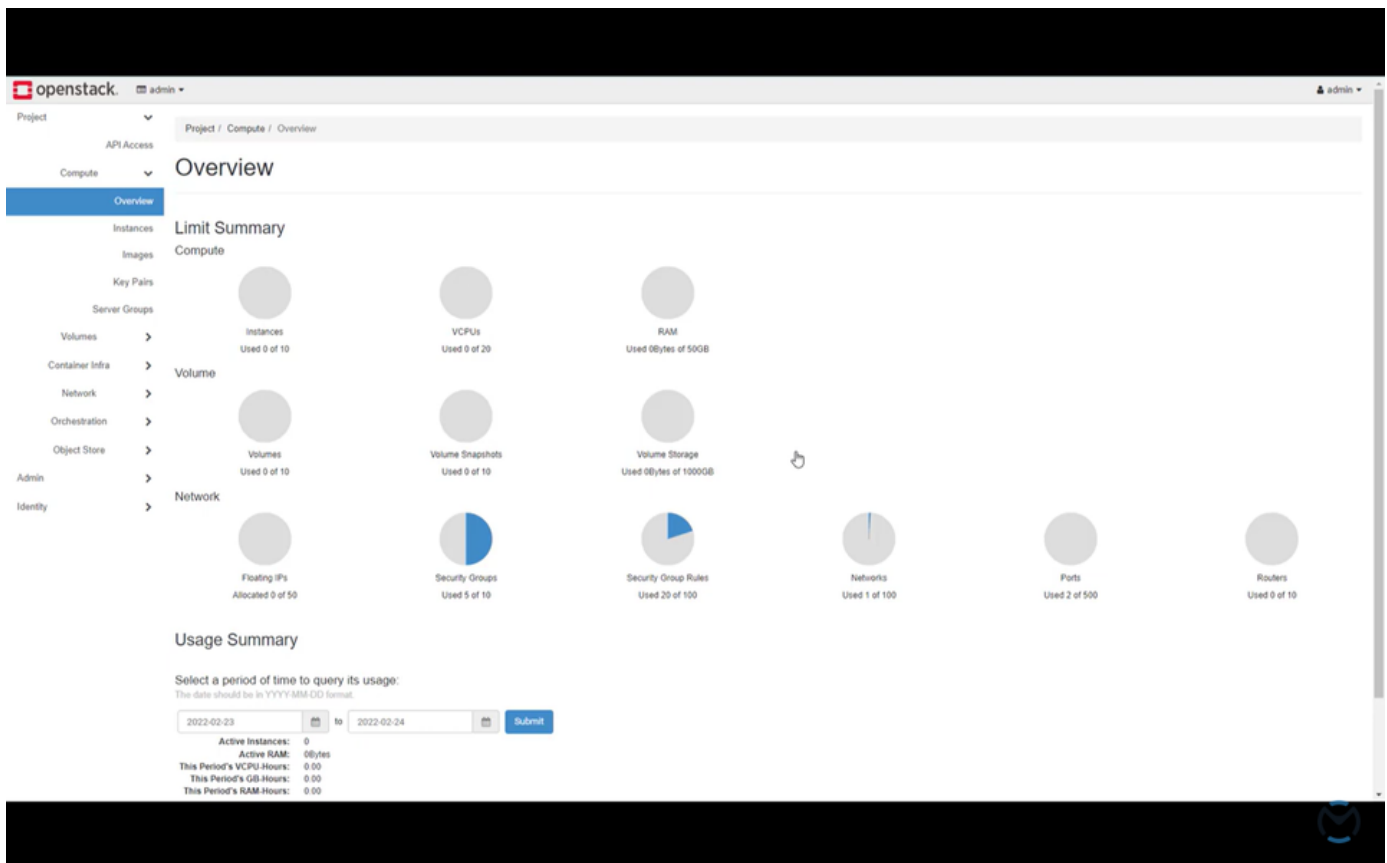
## How much storage will I get in my cloud?

You will see RAW disk space and usable disk space within our documentation. RAW is what actually drives supply, usable is what can be allocated after data redundancy.

We recommend a minimum of “2x replication” for redundancy. This is akin to RAID 1 and means your data exists in two places at one time. Other common settings are 3x replication and erasure coded 3+2, 4+2, 5+2, etc.

Erasur coding is the most efficient for RAW versus usable, but requires more resources and is slower than 2x replication. 3x replication is common as well and enabled by default unless your storage costs override the data protection benefit of 3x replication.

To calculate storage on your own, add together your RAW and divide by your replication level. For erasure coding it is more complex but starts at about a 40% drop with small clusters but goes to as efficient as only a 12% drop for large clusters.





# Private Cloud Cores Continued

## How do I grow my cloud?

It is always important to note that it is easy to create clouds so sometimes the need for more resources can best be met with another cloud. For an existing cloud, it can be grown one of three ways:

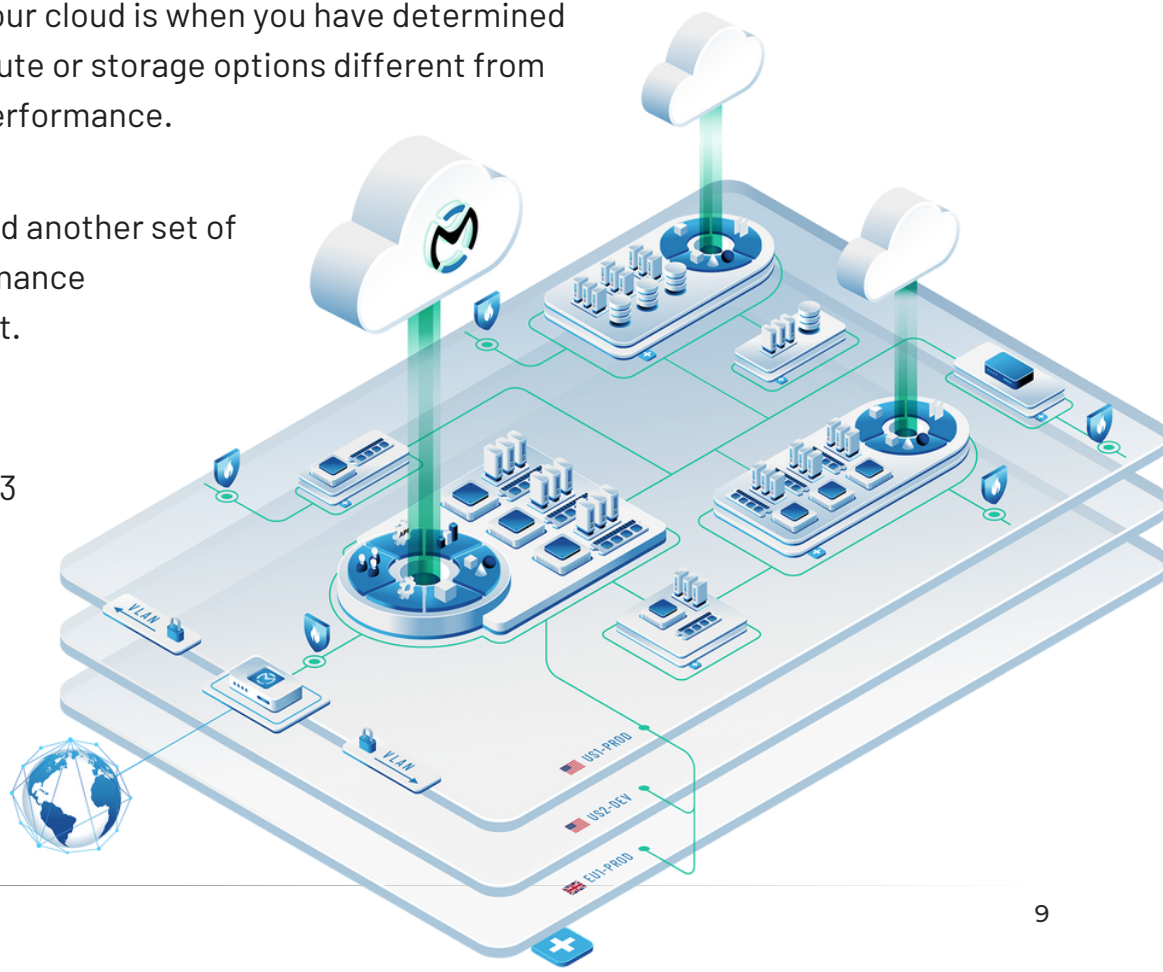
First, you can grow your Private Cloud Core by adding additional converged nodes. These converged nodes are spec matches for your PCC and allow Ceph to use these additional drives within the existing storage cluster. Thus, you will grow your compute and storage at the same time.

Second, you can select compute nodes. These nodes will come with various CPUs and storage but will not be joined to your existing Ceph by default. The drives on the compute nodes are not part of an HA cluster and can either be local LVMs or ephemeral (it will not survive a VM being deleted). When using storage on a compute node the application must provide for data protection as both LVM and ephemeral drives are single points of failure. The upside for storage is NVMe and it is incredibly fast.

The third way to grow your cloud is when you have determined you want specific compute or storage options different from your current cluster's performance.

In this case, you will need another set of servers with the performance characteristics you want.

For storage, clusters have a minimum size of 3 servers for data replication and to be highly available.



## Private Cloud Cores Continued

---

For compute, you can add one server at a time but it is recommended to add three or more servers at one time to allow proper performance if one server must be removed from the cluster for any reason. For example, if you are growing from three Small compute servers, you can add one Large to it, but it is better to add three Standard compute servers, then migrate your VMs to the Standards and return the three Smalls.

Adding or removing nodes can be done within the OpenMetal Central interface and via an API call. Once that server is within your network, you will decide what current VMs will be migrated to the new compute.

---

### **HOW FAST WILL MY STORAGE BE?** BLAZING! OR AS FAST AS YOU NEED.

---

We have three types of storage - NVMe for ultra high performance, SATA SSD for high performance, and spinning HDD for large data stores at relatively slow performance. Large spinner based servers, like our 12 x 12TB option, also come with a set of NVMe drives to accelerate the storage pool to a medium performance level. Reference architectures with FIO performance stats are available from your Account Manager.

When using highly available storage, like a Ceph cluster, there are tradeoffs. More replicas take longer to complete writes. A 2 replica will be faster than a 3 replica which will be faster than erasure coding. Traversing the network lowers throughput. Most deployments will mix how storage is accessed based on the applications.

For another example, drives (excluding boot drives) on a hardware node that are not part of your storage cluster can be made available directly to a VM running on that node. For example, high performance MySQL services are often configured to use a locally mounted NVMe drive with MySQL asynchronously replicating across the network to another VM with local NVMe. Thus HA is accomplished via the software.

This is more performant than using the Ceph's highly available block device mount by 3-4x. It does require additional knowledge but that configuration is common.

# Private Cloud Cores Continued

## How do I set spending limits?

Spending limits are set directly in OpenMetal Central per cloud. Once set, if a server is requested that will put the daily bill over the maximum, the requests will fail and indicate a budget has been reached.

## How are networks or IPs handled?

All clouds support private networks including on the fly creation of routers, switches, and subnets. VxLAN technology is used for separation of projects that are VM based. From an application level VLANs created with VxLAN are indistinguishable from a VLAN. Hundreds or even thousands of VLANs are supported in this way. For bare metal, you are supplied a total of 6 VLANs with each cloud and can access additional if required.

You can grow or shrink from within each cloud. For public IP addresses, we supply you with a /28 to start. Additional IPs are available for monthly per IP fee. We also support assigning our routers to broadcast your IP blocks (/24 and larger via SWIP) and route them to your deployment.

US-East

Mock Deploy Test

My Cloud

Assets

Bandwidth

Team

Billing

Support

Settings

Datadog

Horizon

Documentation

Support

### Networking Assets

#### Inventory IP Address Blocks

CIDR Range	Start	End	Hosts
173.231.252.0/28	173.231.252.1	173.231.252.14	14

#### Provider IP Address Blocks

CIDR Range	Start	End	Hosts
173.231.255.192/28	173.231.255.193	173.231.255.206	14

#### Network Information

Node specific details about your cloud's physical ethernet interfaces, bonded interfaces, and Virtual LAN configurations.

Cloud Core - Standard      Hostname: color-home.local      IP Address: 173.231.252.5

##### Physical Ports

Physical ethernet network interface and bonded interface configurations.

Interface	Mac Address	MTU	Bond Master
eth0	02:00:00:af:67:45	9000	bond0

# SUPPORT LEVELS

## What support levels do you offer?

Many IT departments and IT leadership are looking for flexible support levels as they navigate this transition to off-premises hybrid cloud solutions. There are two standard offerings, plus we can provide tailored offerings.

- **Hardware Level Ticket Based Break/Fix** - This is included with all clouds. Your team works with us using standard response tickets and emergency response tickets with appropriate SLAs for each. This is the lowest cost level and is intended for teams that either already manage clouds or will be training to take over the clouds.
- **Managed Private Cloud** - This includes the hardware and software level break/fix plus management of the OpenStack and Ceph. You'll have a lead IT professional responsible for your organization's needs and the health of your deployment. Access to your IT and subject matter experts are directly within a shared Slack channel, private Slack, or other communications channel if necessary. Planning and execution of migrations, team training, and more can be included. Starts at \$800/month plus a per box fee.

The screenshot shows the OpenMetal support portal. At the top left is the OpenMetal logo. The main header area features the text "Help Guides and Support" and a sub-header "We work 24/7/365 to make sure your service is fast, safe, and secure. Our mission is to make sure you are 100% satisfied with our products and customer service." Below this, a "Start Here" section contains three buttons: "View Documentation", "New Request", and "View Requests". At the bottom, a "Reach the Team" section includes a chat icon. A left sidebar menu lists various account and support options, with "Support" highlighted.

### What skills will my team need to run the clouds ourselves?

The complexity in private clouds comes primarily from the architecting stage, not the operational stage. As your team will have access to already running clouds, much of the complexity has been eliminated. Though running a cloud is much easier than ever before, the following factors will help you navigate this decision.

- 1.** It is important to remember that individual clouds can be spun up in a matter of minutes. Testing, training, and dry runs become simple and fast in a safe environment. This is very empowering for your team and is one of the most important attributes of on-demand private cloud.
- 2.** OpenStack and Ceph are mature technologies. There are vibrant communities around both with extensive documentation and training. Both certainly are complex but configurations have been supplied that work for the majority of workloads.
- 3.** We provide a [Cloud Operator's Training Manual](#) that can be used as your base training for all cloud operations staff. For System Administrators that have been running large deployments, the tools and processes will be familiar. Previous knowledge of Ansible will speed up training. For experienced System Admins, the Cloud Operator's Manual is designed to take between 30 and 60 hours to complete.
- 4.** We can provide Tier 3 support in case your team has situations they can not solve.
- 5.** You can elect to use managed private cloud to get started and transition after your team reaches sufficient expertise.

# 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](mailto: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\*\*](#)



555 S Independence Blvd  
Virginia Beach, VA 23452

EMAIL: [sales@openmetal.io](mailto:sales@openmetal.io)

WEB: <https://openmetal.io>

