openmetal

OpenMetal.io

CHOOSING THE RIGHT CLOUD DEPLOYMENT MODEL FOR YOUR BUSINESS

CLOUD IS NOT A DESTINATION, IT IS A PHILOSOPHY

Cloud has matured dramatically in the last 10 years. In the beginning there were only a few cloud deployment model options.

One of these, the public cloud deployment model, was dominant in marketing and in the quality of the cloud experience. Unless you were okay with waiting 1 to 2 years for a quality private cloud, your only option was “to move to the cloud” at AWS.

Today, there are many easy and fast ways to get cloud. You’re not limited to using one provider or even one type of cloud.

However, many people still only think of the big public cloud providers when considering cloud options for their business. We want to dispel that misconception and introduce you to models that may be a better fit for your workloads and goals.

Consider the “cloud” not as a destination or a location. It is a method of managing workloads. As we discuss cloud migration, don’t think of it as “moving”, but instead choosing a cloud deployment model that works for your business.

You’ll likely find cost savings (many have seen a 50-75% reduction in their cloud bills), gain the flexibility and control over your infrastructure you’ve been hoping for, and discover new ways to operate in the cloud that better serve your organization.

"**If you're operating at scale, the cost of [public] cloud can at least double your infrastructure bill."**

The Cost of Cloud, a Trillion Dollar Paradox by Andreessen Horowitz

CLOUD DEPLOYMENT MODEL OPTIONS

Deployment Model	Cost	Startup Time	Risk	Data Center	Network	Servers	Cloud Install
On-Premise	$	6-18 Months	High	You	You	You	You
Colocation	$$	3-6 Months	Medium	Provider	You	You	You
Dedicated Servers	$$	1-2 Months	Med-Low	Provider	Provider	Provider	You
Hosted Private Cloud	$$	1 Week	Low	Provider	Provider	Provider	Provider
Smaller Public Clouds	$$$	1 Day	Low	Provider	Provider	Provider	Provider
Mega Public Clouds	$$$$$	1 Day	Low	Provider	Provider	Provider	Provider

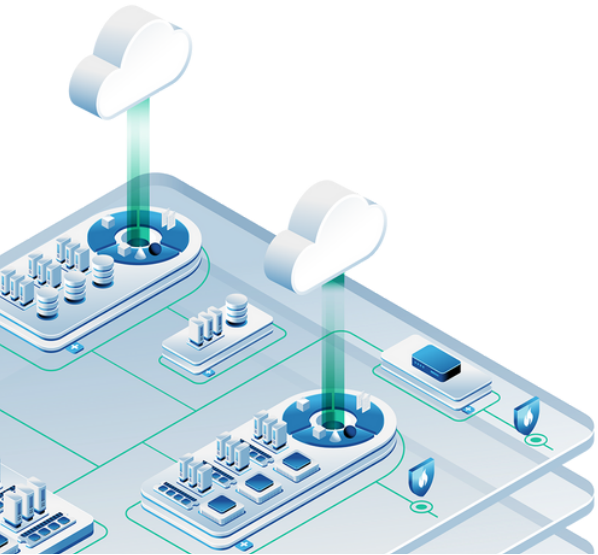

- **Cost** = Per Unit Cost of Resource (VMs, Kubernetes Clusters, Block Storage, etc.)
- **Startup Time** = From decision to when first VM is provisioned for production use
- **Risk** = Likelihood of issues that will cause significant delays or for the project to never completely
- **Data Center** = Who owns the DC
- **Network** = Who supplies the switches, routers, and internet connectivity
- **Servers** = Who supplies the physical computers in the DC
- **Cloud Install** = Who turns the hardware into a cloud native resource pool

THE DEPLOYMENT PROCESS FOR EACH MODEL

On-Premise Infrastructure (On-Prem) - $

In your data center that you build, with servers you purchase.

You automate some of your hardware management with things like IPMI, PXE booting, or even using things like standalone OpenStack Ironic.

Next you automate access to the hardware with tools like OpenStack, Ceph, Minio, CloudStack, MaaS, VMWare, Nutanix, ProxMox, OpenShift/OKD, Rancher and Longhorn, etc. These systems all generally work the same and will supply a set of APIs or GUIs for self service resource deployment.

After you supply the resulting APIs to your team, the VMs, Object Storage, Container Pods, etc. can consume these resources via automation (or GUI).

If you are going down this path we encourage you to still use OpenStack and check out our How to Build a Private Cloud on OpenStack article.

Colocation – $$

In a provider's data center, with servers you purchase.

Just like on-prem above, you automate some of your hardware management and you automate your workload deployment.

Your end "customer" consumes the VMs, Object Storage, Container infrastructure, etc. just like they do with public cloud and what was made above in on-prem.

Dedicated Servers – $$

In a provider's data center, leased hardware is delivered on demand to you, and you load on, via IPMI, OpenStack, Ceph, Minio, CloudStack, MaaS, VMWare, Nutanix, OpenShift/OKD, Rancher and Longhorn, Ceph, etc.

Then you automate your workload deployment against the APIs or worked in GUIs to deploy.

Hosted Private Cloud – $$

This is where OpenMetal sits. In a provider's data center, leased hardware is delivered on demand to you as a cloud.

You automate your workload via APIs against VMs, Object Storage, etc. GUIs are common, but infrastructure as code is more common since users of this technology are both coming from the above cloud deployment models and from the below cloud deployment models.

For System Admins, SREs, Engineers that may have been "born in the cloud", i.e., they have only used the public cloud deployment model, they often do not wish to get back into a GUI approach. So tools like OpenStack, for example, have very robust API-first approaches and have consciously deprioritized work on GUIs as that approach has less modern usage.

Check out our Hosted Private Cloud or on-demand OpenStack Private Cloud.

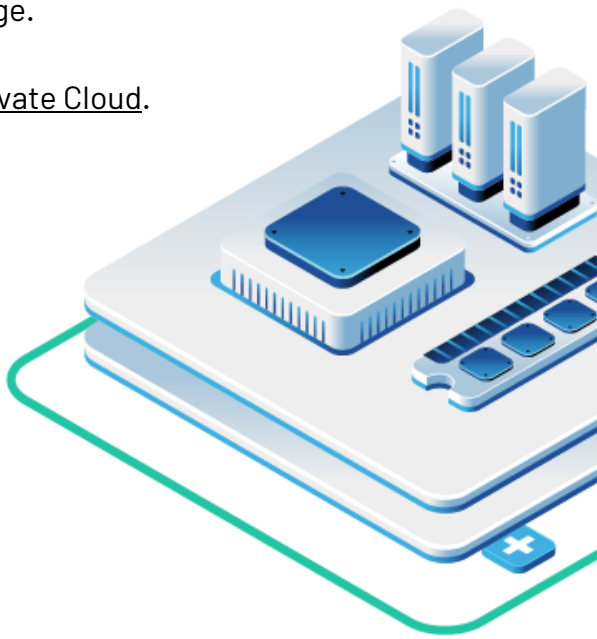

Alternative Public Clouds – $$$

Leased from providers like Digital Ocean, Vultr, and Ramnode.

You automate your workloads against VMs, Object Storage, etc.

Mega Public Clouds – $$$$$

The Big 3+, you automate your workloads against both VMs and "time saving" things like serverless.

When working in this cloud deployment model a business typically has to introduce some type of "Cloud Financial Operations" or FinOps for short. FinOps is the practice of not letting your cloud provider's bill run away.

Seems simple, but it is anything but and a whole industry and job function has been created over the last few years. In fact, the FinOps Foundation is now formally part of CNCF.

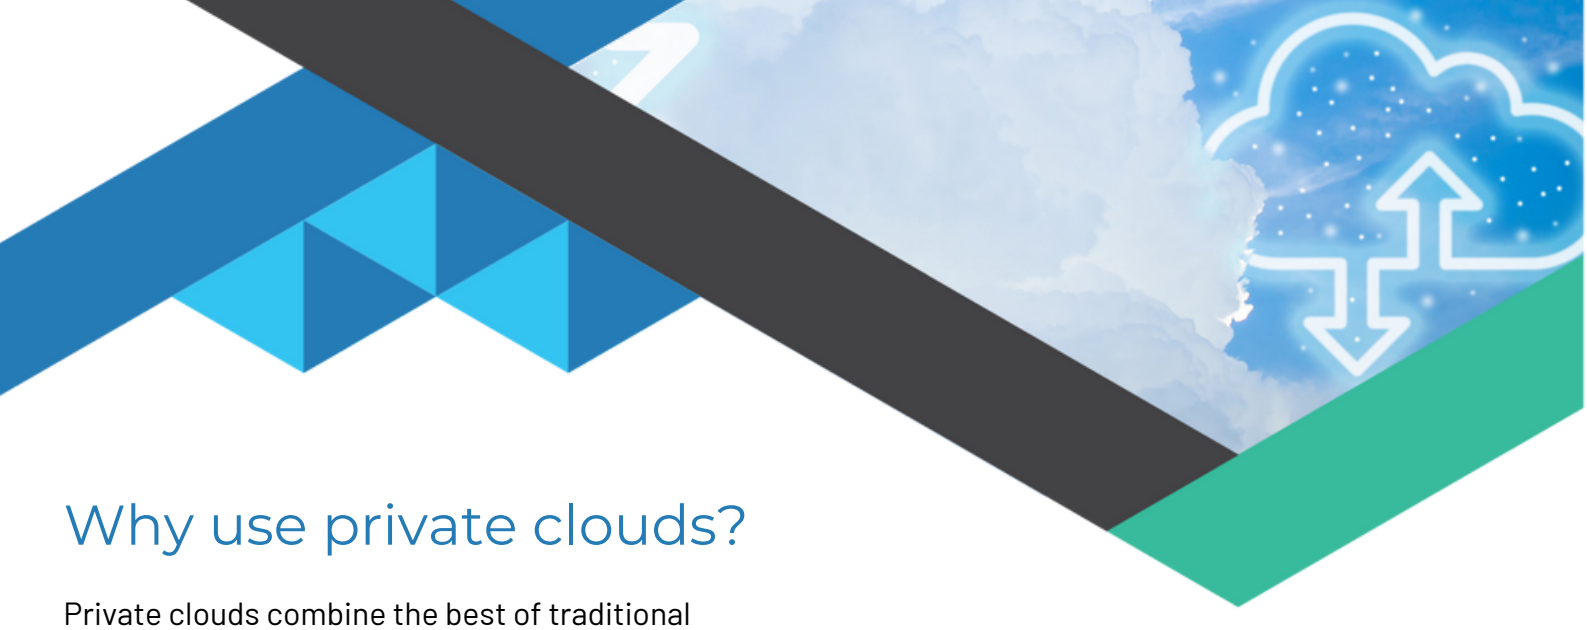

Why use private clouds?

Private clouds combine the best of traditional on-premises infrastructure and the public cloud.

Private cloud solutions, like OpenStack, give you the ability to squeeze more utilization from your hardware and infrastructure and flexibly allocate your compute and storage resources.

They are the go-to choice of enterprises because they give large teams and organizations the ability to build infrastructure customized and optimized to meet their own needs best. The one-size-fits-all architecture of public clouds means public clouds are not a perfect fit for every use case.

Private clouds are often much more cost-effective than public clouds when they reach certain economies of scale. Public clouds charge a relatively high markup on commodity hardware and resources.

Privacy and security are some of the biggest advantages private clouds have over public clouds. When you run your workloads on a public cloud, your workloads coexist alongside other customers' workloads on the same physical infrastructure. There's a minimal security risk from someone breaking out of their virtual environment and entering yours, but even that small risk is too much for sensitive work involving confidential data. Private clouds by default belong only to you; you do not share hardware with other users.

Why use public clouds?

Public clouds pioneered the consumption-based model of computing for the mass market. They gave developers, engineers, and teams of all sizes the same tools and resources as enterprises to build and serve highly available, highly redundant, easily scalable applications and infrastructure.

The public cloud provides an entry point into running and operating large infrastructure that is generally out of reach for smaller teams and organizations that cannot afford data centers, networking, hardware, and all the other pieces of physical infrastructure required to run production workloads on the Internet.

Public clouds have the advantage of requiring no setup or infrastructure to manage. You can provision resources immediately with a single click. Private clouds have historically required significant research, work, and upfront investment to build your cloud infrastructure before you even begin to deploy your workloads.

PUBLIC CLOUD ADVANTAGES?

A search on "public cloud advantages" will nearly always include some statements about public cloud being "less expensive" than alternatives. Unfortunately, this common narrative is simply not true for many situations!

Public cloud for companies with small total workloads is absolutely a great choice. If you are interested in what is "small" and ways to understand when your tipping point occurs, check our A Cost Tipping Point Guide for IT Professionals: Public vs Private Cloud. The advantage we are speaking about comes in for companies that have reached their tipping point. As you scale this advantage becomes very meaningful.

At scale the resource management advantage becomes larger and larger in raw savings. A small deployment example of just 40 VMs can give you an idea of what is going to happen. Keep in mind that many companies use thousands of mixed size VMs or hundreds of very large VMs. The costs can very quickly become significant.

In particular, we created programs to help SaaS companies and hosting and cloud providers that run on those public clouds to exit as they have found themselves with little or no profit just because of the cost of the public cloud.

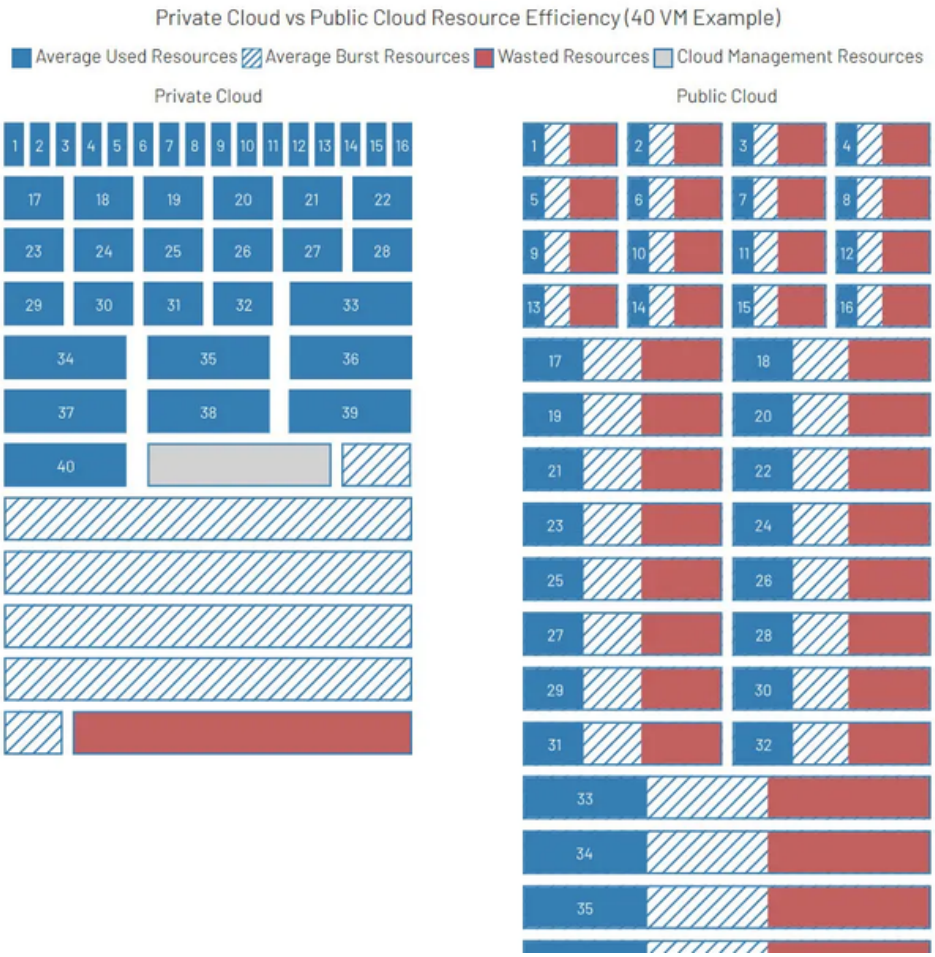

More information about alternatives to public cloud is needed to help leaders in IT explain when public cloud is right or when private cloud, bare metal, colocation, or owned data centers are the right choice.

PRIVATE CLOUD VS PUBLIC CLOUD RESOURCE EFFICIENCY

When you lease a virtual machine from a public cloud, you must pay for a VM for a workload that will use roughly 30% of the resources on average and will burst to use, when averaged out over time, no more than 30% more. When averaged, this leaves roughly 40% of the VM wasted.

Unfortunately, you must still lease this average wasted space by getting a much larger VM than your workload really needs. If you do not, your workload will run into limits periodically that will ultimately degrade your performance both at the rare, but critical, times your workload needs the full resources and at the times that the workload should already have finished but could not because it was delayed by resource constraints.

The advantage private cloud has over public in this area comes from a basic fundamental difference: you are leasing all the resources of the hardware and not just the virtual resources.

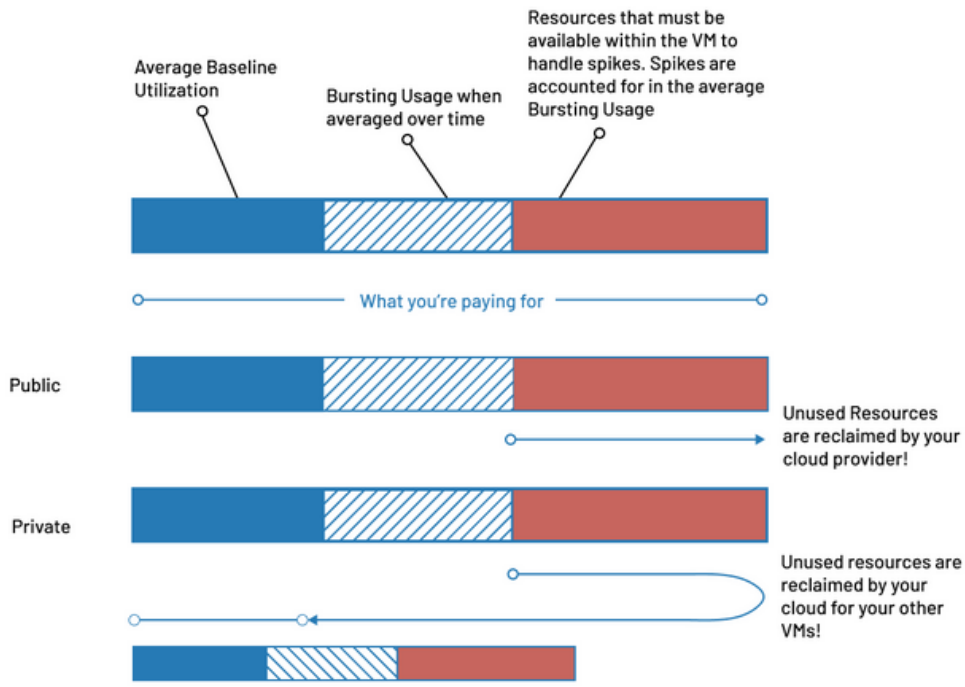

This means when you provision a VM you can provision the size you need, but during the time that VM is not spiking to consume the top level limit of the VM, those previously wasted resources are accessible via your private cloud software to be used by your other VMs!

In public cloud, your wasted resources are often designed to be reclaimed by the public cloud to be resold to other users. This is another key area for savings and a strong reason to move away from public cloud offerings.

Occasionally you may hear the idea that if there is so much waste in the VM why not add more workloads into the VM? That is certainly one of the options, but it also negates the key benefit of workload separation that is fundamental to cloud and virtualization. And your system admin team or your SREs will be very angry with you! Usage of containers in VMs is a reasonably safe way to do this and is common already but typically for related parts of a single system which has its constraints. Either way, that requires quite a shift in thinking and typically design.

FINDING THE MODERN CLOUD SOLUTION THAT WORKS FOR YOUR BUSINESS

When a company moves their workloads out of a public cloud it does not mean they move their workloads back to how they were managed before. There is no reason to move to an older infrastructure model as having an in-house cloud is relatively straightforward nowadays.

Did you use Terraform or Ansible to automate the creation of your cloud infrastructure on AWS, GCP, or Azure? Well, guess what, that same infrastructure as code mentality is supported by popular private cloud software like OpenMetal clouds.

Did your Dev team fall in love with Kubernetes? It is pretty awesome, for sure, and guess what? Kubernetes is also standard stock for any modern cloud.

Many companies have found that skills need to be reintroduced to the company to take on "cloud administration". This may be a concern for you, but the ease of running clouds has also dramatically advanced in the past 5 years. For self-managed clouds, **an experienced system administrator can learn to run clouds in as little as 40 hours**.

Relearning hardware is a valid concern, but one that has boundaries and multiple ways to de-risk including onboarding, service contracts, and management options that yield **per VM costs that are still well below mega cloud costs and can be temporary**.

A few final thoughts:

- Moving your workloads off the mega clouds does not mean you are "leaving the cloud". You are leaving the mega clouds.
- When talking about "cloud" don't define it as a location. Instead, just think of cloud as a method of managing workloads. By this we mean you don't have to "move"!
- Private cloud does not mean "on-prem" as "on-prem" doesn't even usually refer to your own data center anymore.
- Private cloud in a colocation facility is very common.
- Hosted private cloud, on-demand private cloud, or private cloud as-a-service is relatively new, but growing quickly. It is a great way to be both flexible like a traditional public cloud and cost effective like a private cloud.

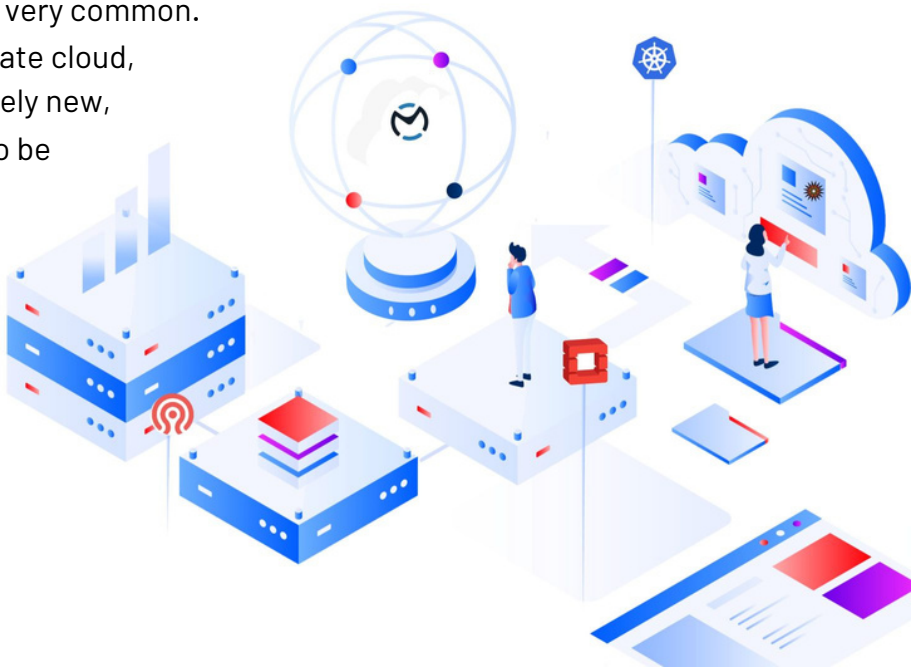

WRAPPING UP

"The solution we found with OpenMetal was perfect because we are now able to spin up new environments very cost-effectively, get to market faster with our solutions, and still create an agile environment that supports existing production customers while enabling net-new customer opportunities."
-Gary MacDougall, CTO @ Pypestream

Read Pypestream's and other customers' success stories >>

Companies are deeply entrenched, and rightly so, into the concept of "cloud". But, it is important to be crystal clear that the cloud is not a destination. It is an approach that can be implemented in many ways and not just by moving to a few mega public clouds.

We believe that reducing reliance on the hyper-scaler public cloud providers is both imperative and smart for businesses moving forward. Cloud costs have spiraled out of control and make growing your company nearly impossible once you've reached a tipping point.

On-demand hosted private cloud removes startup expense and risk and gives you the benefits of the private cloud immediately. We are excited to introduce OpenMetal and invite you to explore a new option to be "in the cloud".

REQUEST A TRIAL OF OPENMETAL CLOUD >

ABOUT OPENMETAL

OpenMetal is a leading provider of open source cloud and infrastructure-as-a-service (IaaS) solutions. By combining the strengths of traditional public cloud, private cloud, and bare metal fused into an alternative cloud platform (powered by OpenStack and Ceph), OpenMetal eases accessibility to highly complex open source systems and allows companies of all sizes to realize new opportunities in performance, productivity, and profitability. A strategic member of the Open Infrastructure Foundation (OIF), OpenMetal is committed to empowering individuals – by themselves or within teams – to meaningfully contribute to the larger open source community to foster innovation that benefits all.

Did this resonate with your business needs? Contact our team to find out how OpenMetal can help you plan the best cloud deployment method for your needs, goals, and budget.

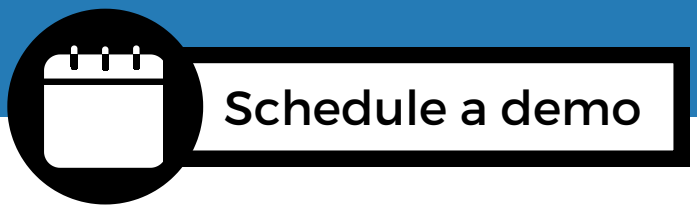

openmetal

OpenMetal.io