



DATA CENTER

# Frontier Special Report

## Bare-Metal Computing Leads a Changing Cloud Landscape

*Written by Paul Gillin*



*Photo credit: ranjith ravindran/Shutterstock.com*

## Contents

Bare-Metal Computing Leads a Changing Cloud Landscape .....	2
What is bare metal cloud computing? .....	3
Why would I use bare-metal computing?.....	3
Why is bare-metal computing gaining momentum?.....	4
Choosing a bare-metal computing service provider .....	4
What to look for in a bare-metal computing provider .....	5
Processors matter.....	5

# Bare-Metal Computing Leads a Changing Cloud Landscape

The swift migration of enterprise workloads to cloud infrastructure shows no signs of slowing. More than 80% of respondents to [IDG's 2020 Cloud Computing Study](#) said they already have some applications in the cloud and 69% of large corporations said they have migrated most or all of their applications. [Gartner](#) [expects](#) global spending on public cloud to grow 23% in 2021 and nearly 20% in 2022 to nearly \$400 billion.

One of the key technical underpinnings of cloud environments is multi-tenancy, a software architecture that enables several users to share the same physical infrastructure and applications. Based on virtualization technology, multi-tenancy allows cloud infrastructure providers to drive high server utilization rates that yield cost savings that are passed on to their customers. The highly virtualized nature of cloud infrastructure also gives customers the flexibility to provision their virtual servers at will with very rapid deployment and usage-based pricing.

While most applications adapt well to virtualized environments, there are some performance penalties IT leaders should take into account when deciding which applications to move to the cloud. The hypervisor, which is essentially a lightweight operating system that emulates the physical machine, [consumes some resource overhead](#), particularly if the application is CPU-bound. A less predictable and potentially more disruptive performance penalty comes from the “noisy neighbor effect.” This occurs when one workload begins consuming a large number of available resources such as network capacity and memory, thus choking off other workloads from the resources they need.

## Workloads to cloud infrastructure

### Current activity and migration

- 80%** Respondents already have some applications in the cloud
- 69%** Large corporations have migrated most or all of their applications

Source: [IDG's 2020 Cloud Computing Study](#)

### Public cloud global spending

- 23%** Growth in 2021
- 20%** Near growth in 2022
- \$400 billion** Near spend by 2022

Source: [Gartner](#)

These factors are a structural element of multi-tenant environments that can't be avoided. While most workloads can tolerate some levels of variability, performance-sensitive applications like online transaction processing and artificial intelligence model training need every ounce of machine power.

Organizations that are looking to migrate performance- or location-sensitive workloads to the cloud are increasingly looking into bare metal cloud options.

## What is bare metal cloud computing?

Bare-metal cloud computing is, as the term suggests, unadorned compute, storage, and networking that customers may configure according to the needs of their workloads. It's similar to compute infrastructure you would buy to install in your data center or that of a co-location provider, but cloud computing has redefined bare-metal around a new set of constructs anchored by much higher levels of automation and customer flexibility.

---

**Because the instance is single-tenant with no virtualization overhead, there are no performance penalties or noisy neighbor problems.**

---

In an on-premises or colocation scenario, customers own the equipment and either furnish their own facilities and connectivity or license those features from the colocation provider. In a bare-metal cloud computing model, the hosting provider owns the equipment and provides access to customers on a subscription basis. All the physical server resources are dedicated to a single customer for whatever uses they require.

Because resources are not shared, no hypervisor layer is needed, allowing the full processing power to be allocated to the application. And because the instance is single-tenant with no virtualization overhead, there are no performance penalties or noisy neighbor problems.

A staged consumption model gives customers flexibility to specify how much server power they need and to upscale and downscale capacity as appropriate, paying only for what they use. Customers choose whatever elements of the operating environment they prefer with installation and configuration handled by the cloud provider. Availability is guaranteed and service-level agreements may be available.

Although relatively young, the bare metal cloud market is primed for rapid growth. [ResearchandMarkets](#) expects the global market to grow 24% annually to \$16.4 billion by 2026, up from \$4.5 billion in 2020.

## Why would I use bare-metal computing?

There are several common scenarios for bare-metal computing.

- ▶ Database-intensive processes such as OLTP, securities trading, and credit card fraud detection require sub-second response times. The unpredictability of multi-tenant cloud environments can cause unacceptable delays.
- ▶ High-performance computing tasks such as AI model training and computer-aided design are tasks best performed with hardware that incorporates such performance-boosting features as high-speed interconnects, specialized memory and storage, and multiple graphics processing units.
- ▶ Businesses in regulated industries may have to meet data locality requirements that restrict where data is kept and processed. Because many cloud providers shift and copy data between regions, a public cloud may not provide sufficient protection. A bare-metal server guarantees that data is physically kept in one location at all times.
- ▶ For businesses that require high levels of security, hardware-based assists provide an extra layer of protection that is not available in virtualized environments.

A combination of these factors influenced the [Translational Genomics Research Institute's](#) decision to make bare-metal cloud the core of its computing platform.

The non-profit TGen processes a vast amount of data as it conducts genetic research into developing early diagnostics for a variety of diseases. Much of the data is sensitive and confidential, a concern that was heightened by the organization's expansion into Europe and the accompanying need to comply with the General Data Protection Regulation.

TGen uses the Bare Metal Cloud from phoenixNAP to rapidly provision physical servers on a global scale while meeting safety and compliance requirements. Deployed in minutes, Bare Metal Cloud servers provide the hardware and security infrastructure necessary for the confidential processing of large volumes of data.

Security will soon be enhanced through the addition of the recently introduced 3rd Generation of Intel® Xeon® Scalable Processors. These powerful CPUs will give TGen access to [Intel® Software Guard Extensions \(Intel® SGX\)](#) and [Intel® Trusted Execution Environment \(TXT\)](#), providing additional data and infrastructure protection. The organization will also take advantage of phoenixNAP's Encryption Management Platform to safely store and manage its cryptographic keys in a simple, orchestrated process.



Deployed in phoenixNAP data centers in both Arizona and Amsterdam, the Bare Metal Cloud “gave us a European presence with the level of security needed to ensure GDPR compliance,” said Glen Otero, vice president of scientific computing at TGen. The combination of the next-generation Intel CPUs and the encryption management platform “further our ability to deploy and manage security tools we need, all while maintaining exceptional levels of performance.”

## Why is bare-metal computing gaining momentum?

In the 15 years since public cloud services first appeared, organizations have become far savvier about choosing their options. Many leverage a combination of on-premises, cloud, co-location, and bare-metal services that are tightly tailored to their workloads. A [Flexera survey](#) found that 92 percent of enterprises use multiple clouds with an average of 2.6 public and 2.7 private clouds each.

One of the great appeals of cloud platforms is their rapid provisioning capabilities—servers can be spun up in minutes compared to weeks or months for on-premises physical infrastructure. However, the need for dedicated servers hasn't gone away. With the emergence of bare-metal computing in the cloud, organizations now have an attractive alternative to on-premises computing that has all the benefits of cloud provisioning and management.

The rise of the encapsulated software environments called containers and the Kubernetes container orchestrator has given customers even more options about how to deploy software. Kubernetes was ported to bare-metal infrastructure in 2018 and has been shown to run up to 30% faster on bare-metal clusters than on virtualized hypervisors, giving customers the dual benefits of workload portability and high performance.

The hyperconverged infrastructure (HCI) market, which [Allied Market Research expects](#) to grow more than 30% annually through 2026, has also contributed to bare-metal's popularity. HCI's attractive benefits in workload deployment, ease of management, and scalability are possible only using the bare-metal infrastructure. Organizations can use a combination of on-premises HCI and HCI on bare-metal in the cloud and control the entire environment from a single administration point.

Finally, edge computing is one of the fastest-growing IT markets, with [Grand View Research forecasting](#) 37% annual growth through 2027. Edge servers typically require high performance for applications like stream processing, content delivery, and traffic orchestration. In such a performance-sensitive real-time environment bare-metal has a distinct advantage.

## Choosing a bare-metal computing service provider

All major cloud vendors offer bare-metal options, but most have come late to the market and their enthusiasm around bare-metal computing has been muted.

---

**A specialty bare-metal cloud provider can save customers between 45% and 80% off the price of comparable offerings from the three largest public cloud companies.**

---

There are good reasons for that. Cloud vendors specialize in operating highly automated hardware environments with high utilization rates enabled by virtual machines. They endeavor to make the underlying infrastructure as invisible to their customers as possible, which also gives them the flexibility to shift workloads between data centers and regions as they need to balance capacity.

Hosting dedicated hardware in data centers is neither an area of expertise for public cloud providers nor a compatible business model. For these and other reasons, licensing bare-metal computing from a major hyperscale cloud provider can cost significantly more than working with a vendor that specializes in bare-metal computing. One analysis [conducted by phoenixNAP](#) found a specialty bare-metal cloud provider can save customers between 45% and 80% off the price of comparable offerings from the three largest public cloud companies.



## What to look for in a bare-metal computing provider

Given the growth statistics cited earlier, it isn't surprising that many service providers now offer bare-metal options. The expertise needed to operate bare-metal and multitenant cloud environments differs in some fundamental ways, however. Look for a company that can demonstrate that bare-metal computing is a core competency, not an afterthought. Inquire about the company's history, customer base, and the variety of workloads it has managed. Check customer references and ask for examples of how the provider supported individual customer needs.

Look for the widest variety of configuration options you can find so that you won't be locked into a platform that either can't scale to meet your needs or that delivers more compute power (and cost) than you require. For example, phoenixNAP offers more than 20 pre-configured bare-metal computing instances ranging from low-cost single CPU servers with 16GB of RAM to large dual-CPU instances with 768 GB of RAM, 4 TB of NVMe storage, and 50 Gbps of network capacity. New options are being added all the time.

Some cloud providers offer an alternative to bare-metal they call "dedicated servers." While these are similar in some ways there are significant differences in areas like provisioning, billing, and customization. For example, server provisioning is fully automated in a bare-metal cloud but only partly orchestrated on a dedicated server, a difference that can translate into hours or days of wait time. A bare-metal cloud environment also offers more flexible billing and pre-configured instances for rapid deployment.

## Processors matter

The decision to use a bare-metal cloud provider is usually driven at least in part by performance considerations, making the choice of CPU a critical factor. Look for a provider that uses the latest processors, memory, solid-state storage, and interconnect technology for industry-leading performance. The 3rd Generation Intel® Xeon® Scalable processors are optimized for cloud, enterprise, high-performance computing, network, security, and edge workloads with 8 to 40 powerful cores and a wide range of frequency, feature, and power levels. Purpose-built for bare-metal computing, they are the only data center CPUs with built-in AI acceleration and power management.

The Persistent Memory 200 Series for the the 3rd Gen Intel® Xeon® processors features up to 32% more memory bandwidth and more memory per socket for cutting-edge retrieval and storage performance. The new processors also support up to six Intel UltraPath Interconnect channels for scalability and inter-CPU bandwidth to handle I/O-intensive workloads. The upgraded Intel Advanced Vector Extensions 512 are built to support even the heaviest computational workloads such as modeling and simulation, data analytics, machine learning, and digital content creation.

The 3rd Gen Xeon processors are also the most secure on the market, with in-memory application isolation, crypto acceleration, and Intel Total Memory Encryption for full physical memory encryption. They even protect against modifications to firmware.

### Bare-metal computing providers

Look for providers that employ a variety of cloud-native and cloud-like features, including the following:

- ▶ Instant provisioning,
- ▶ CPU-optimized instances for such applications a media encoding, gaming, and high-performance computing,
- ▶ GPU options,
- ▶ Flexible billing terms, including hourly and monthly rates with generous discounts for extended reservations,
- ▶ Built for DevOps with continuous integration/continuous delivery pipeline support,
- ▶ An API-first configuration that makes advanced performance features available in minutes and manages them with infrastructure-as-code tools,
- ▶ Cloud-based identity and access management,
- ▶ Hardware-based security,
- ▶ Easy scale-up and scale-down,
- ▶ Vertical CPU scaling to let organizations easily switch between preset CPU configurations and
- ▶ An ecosystem of value-added partners.

My Party Album Inc. used a bare-metal cloud service to cut its IT spending by 80% while smoothly handling scalability needs. Its platform is a mobile content-sharing app that makes it easy for people who have attended an event to exchange photos and videos privately with each other. The service requires high-performance processing, fast and scalable storage, and robust network connectivity.

The company had investigated public cloud solutions but found the costs to be unpredictable and probably more expensive than on-premises options. The company's developers also found that public cloud services didn't provide the detailed information they needed about the underlying infrastructure. Each public cloud provider had different dashboards, naming conventions, and pricing options. Support was often a self-service proposition.

My Party Album ultimately chose phoenixNAP's Bare-Metal Cloud platform. phoenixNAP's transparent and simple pricing made it easy for the team to forecast spending. There were more than 20 server instance types to choose from and options to deploy globally. Physical servers not only delivered better performance but also more control over security.

---

**My Party Album ultimately chose phoenixNAP's Bare-Metal Cloud platform. phoenixNAP's transparent and simple pricing made it easy for the team to forecast spending.**

---

As a result, My Party Album was able to get its development team up and running quickly, accelerate the delivery of new features to the market and achieve 80% lower costs compared to similar hardware offered by the big public cloud providers.

"With BMC, we found the best of both worlds: fantastic hardware performance at a lower cost," said Grant Street, My Party Album's co-founder and CEO. "Everything is upfront, transparent, and really easy to understand."

