BIGANIMAL

THE COMPLETE GUIDE TO OPTIMIZING CLOUD DATA SPEND

Balancing Price and Performance to Achieve Your Business Goals

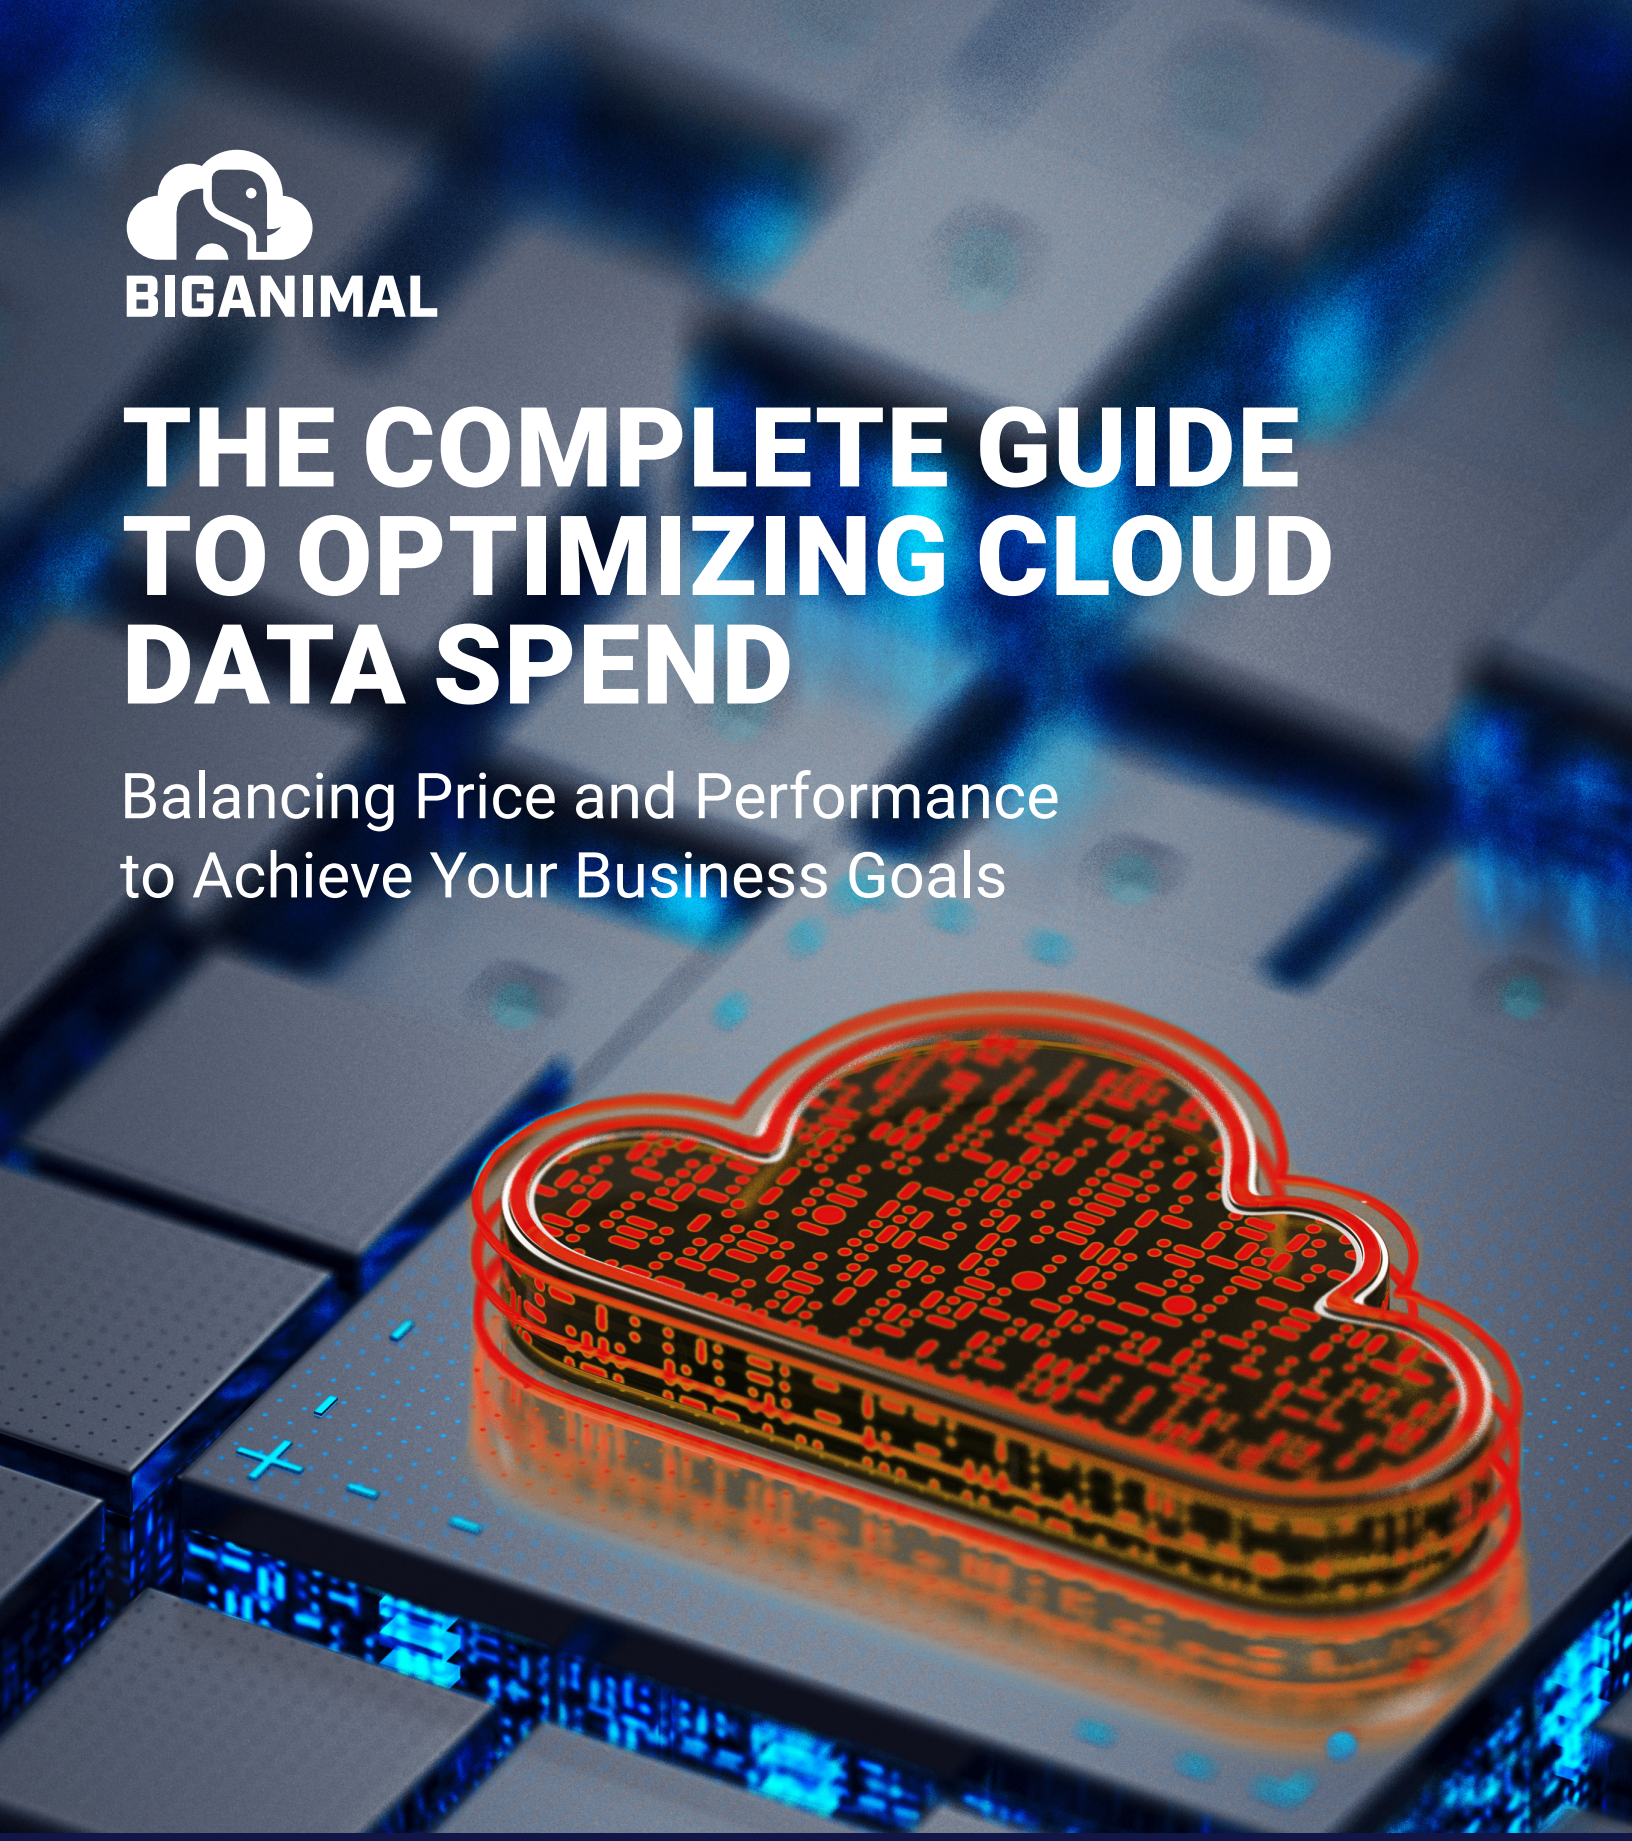

EDB™

Contents

Introduction

While the popularity of doing business in the cloud keeps growing, unfortunately, so does the cost. And not just the cost, but who is responsible for a platform costs and how it is billed. If we use history as a guide, the mainframe is much like the cloud—everything had a cost and was billed back in microscopic amounts, including CPU time (instructions), disk usage and even the number of pages that came off the printer. During the "open-system" days of real-time processing, this type of billing was simplified to a more basic charge-back model that is used today for paying for a certain number of CPU threads, memory and storage capacity. The modern database administrator had no control over the quality of CPU or even the performance characteristics of storage = that was the in the "infrastructure" group. Move to the cloud, and a few things have happened:

- How cloud vendors charge for things is by some small amount for a large number of components and how they are tuned. AWS has over 55 different 8-core computing environments to choose and the choices for storage are no less complex - it isn't just by storage quantity (generally gigabytes), but now we get to worry about inputs/outputs-per-second (IOPS) and throughput measures by megabytes-per-second. To round out costs, you have database connection pooling, monitoring tools, backup costs, egress network charges and more.

- The database team is now responsible for making more decisions within these components. During the mainframe time, the DBA made no decisions and on open-systems they could tune the database around a broad set of constraints. Now, in the cloud, most organizations put the entire responsibility in choosing all these pieces-and-parts on the database team, as they have certain scalability, security and business continuity measures to maintain.

- Cloud vendor obfuscation - AWS, Microsoft Azure and Google Cloud have made the ability to predict total cost of ownership (TCO) nearly impossible. Consider how AWS CloudWatch is billed for a database server, Azure Flex Server charges for backups and Google Cloud's database offerings make it impossible to know.

As a result, it's essential that cloud enthusiasts are able to convince decision makers of the value of moving databases and applications to their chosen platform. More so than ever, it's critical to balance performance expectations with costs—to prove that cloud spend will result in tangible ROI. In order to do so, it's imperative to determine how to get the most value for your database spend, where to invest that spend and how to convey meaningful results.

Moving forward requires collaborative planning, and this is where cloud financial operations (FinOps) comes into play. FinOps calls for aligning your finance, developer and leadership teams to reach your cloud goals and inspire confidence in cloud adoption and ongoing usage. A chasm has grown between the development and database, where optimizing code to bring greater CPU and storage efficiency has ceased to exist. Where bad code against a database used to be more tolerated in the days of simple and fixed costs for components, it should now be less tolerated.

Read on to discover the benefits of FinOps, what's needed to develop an effective FinOps strategy and how to assess the success of your investment.

01 The 2024 cloud database landscape: Why cost is more important than ever

The 2024 cloud database landscape: Why cost is more important than ever

In order to fully understand the increasing importance of cloud cost management for businesses, we must first dig into the current cloud database landscape, which offers organizations both extensive opportunities and complex pressures.

In a report published in October 2023, Gartner describes how cloud continues to outpace all other IT spending areas. Based on Gartner's research, spending on public cloud services is expected to surpass $1 trillion through 2027, representing around 18% of overall IT expenditures. In the near term, public cloud services spending is expected to grow to $678.8 billion in 2024, representing a 20.4% increase from 2023.

In short, the cloud's footprint is expanding rapidly and exponentially, becoming a bigger part of IT budgets every day. At this rate, the previously held philosophy of "cloud first" seems to be transitioning into what might be described as "cloud only."

But, this is only one half of the equation. While the cloud might be more necessary than ever, we now see a new cost-focus among organizations for two key reasons:

- The question is no longer "why the cloud?" but "how?" In the current landscape, the major differentiators between providers are less likely to be specific features—in order to be a major cloud provider, you need to have met the basic requirements with regards to performance, high availability, security, etc. Instead the difference between cloud databases lies in price, cost structure and how quickly your enterprise will begin to see the return on your cloud spend.
- While in the early days, buyers viewed the cloud as an all-around cheaper alternative to on-premise storage providers, that reputation has slowly eroded, with cloud solutions becoming more entrenched in every corner of the modern tech stack. Nowadays, the cloud isn't cheaper, just "**differently expensive**."

Now, IT teams must justify their cloud spending to decision makers and present models that communicate the value of that spending over longer periods of time. When undertaking this effort, it's important to take into account the rise in cost-consciousness and address key concerns.

02

What do cost-focused companies look out for?

What do cost-focused companies look out for?

What areas and considerations are most important for businesses looking to invest in the cloud? While this can depend on the individual organization and its needs, here are three of the most common which your teams might currently be weighing.

Inflation and economic uncertainty

There's no denying the impact of market instability on IT buying behavior in the past year. Now, many companies are employing spending gates to limit their expenses per quarter, while others are looking at tech stack ROI with greater scrutiny than ever before. When consistent revenue isn't guaranteed and the economic landscape feels volatile, it's natural for enterprises to tighten budgets and expect guaranteed returns before investing in new cloud solutions.

Growing IT budgets

Simultaneously, the growing importance of the cloud and other innovative technologies has led to financial decision makers allocating more funds for their organizations' IT budgets and more of those allocations are going towards cloud services. As a result of this influx of cash, IT leaders face greater scrutiny and higher stakes in the cloud databases they select: they need to make the most of their cloud solutions without overspending.

Cloud pricing models

As we'll discuss in-depth later, a key component in the recent understanding of the cloud as "differently expensive" is the diverse and sometimes confusing ways in which various providers price their offerings. Some might be cheaper up-front but are more likely to incur additional costs over time, while others might seem more expensive to begin with but come with services—such as remote database administration— that can help businesses save money in the long run. On top of this, some cloud providers don't have hard spending limits, and notifications for meeting data usage caps may be delayed for hours, leading to a business going far over-budget without realizing it.

03

Cloud FinOps: A modern approach to cloud spending management

Cloud FinOps: A modern approach to cloud spending management

As a result of these factors—growing cloud expectations, expanding budgets, increasing economic fluctuations—IT, developers and finance teams must align on cloud strategies, requirements and expenses to better manage cloud spend and maximize the return on a cloud database investment. Without a dynamic strategy dedicated specifically to optimizing spend, you won't be able to achieve the performance your teams demand at the price your executives desire.

That's why leading businesses are turning to cloud financial operations (FinOps).

Coined at the Cloud Economic Summit in early 2019, FinOps is currently defined as "an evolving cloud financial management discipline and cultural practice that enables organizations to get maximum business value by helping engineering, finance, technology and business teams to collaborate on data-driven spending decisions" by the **FinOps Foundation**.

Put simply, by aligning all the teams most affected by cloud purchasing decisions, FinOps aims to enable faster and more efficient usage of cloud services while maintaining better financial control.

Unsurprisingly, FinOps is only growing in popularity as businesses look to better forecast and manage cloud expenses. Given the multi-faceted nature of cloud costs, unifying necessary decision-makers before selecting the right solution is proving to be the best way to optimize cloud spend now and in the future. FinOps recognizes that these spend decisions aren't just a one-time chore, but an ongoing process that will evolve with your organization. The result: the ability to accurately benchmark and holistically compare cloud solutions with confidence.

Let's explore the specifics of how.

04

Connecting cloud FinOps to total cost of ownership

Connecting cloud FinOps to total cost of ownership

Now that we have a fundamental understanding of FinOps' role, we can dive into how to ensure your FinOps team comprehensively and effectively achieves its goal of providing a reliable and scalable picture of your cloud platform's total cost of ownership (TCO).

While a complicated process, it can be broken down into two major projects: (1) outlining the key drivers of cloud cost and (2) comparing them via a proven formula.

How to better understand your spending

In order to ascertain the true TCO of a given cloud database, there are a range of factors your FinOps team must examine.

Price vs. cost

When discussing how much your organization is going to pay for a cloud database, it's easy to get in the habit of using "price" and "cost" interchangeably. These concepts, however, are distinctly different, and their distinction can have a major impact when it comes to selecting a cloud vendor.

Price is what is listed on your billing statement. It encompasses what you pay for storage, features and more. Cost, on the other hand, is what you'll expend on your cloud database—not limited to the money that is paid to the vendor.

When choosing a database, an organization might try to minimize the price by opting out of services such as remote DBA assistance. While this might provide them with a smaller bill, they could be more likely to encounter issues like unplanned outages, service disruptions and data sprawl, which they're not as equipped to handle without an expert on-call. Such problems are expensive and damaging to revenue, leaving said organization with a low-price database that proves immensely costly over time.

Diverse pricing/subscription models

Many factors that make benchmarking difficult for FinOps teams stem from discrepancies in platform pricing. Because different providers have completely different pricing models, comparing platforms can be daunting. But the key lies in understanding your enterprise's needs and usage patterns.

A few models you might encounter include:

- **Pay-as-you-go/Consumption pricing**: If your storage usage or compute is likely to grow rapidly, this model can be problematic. As the name implies, your subscription cost changes based on your behavior and consumption. If this is the model you choose, it can help to set up alerts or configure metrics to ensure you keep track of your consumption. Ultimately, it's on your business to track that usage—lose track and you could wind up with a nasty surprise when the bill comes due.
- **Fixed pricing**: If you know exactly what you need from the get-go, fixed pricing might be for you. Here, your cloud database will have allotted caps on things like compute and storage, so when you hit those caps, you'll be cut off. It's important, however, to make sure that your provider won't let you go over, only to charge you extra later.

- **Tiered pricing**: A cousin of fixed pricing, anyone who has subscribed to a streaming service is familiar with tiered pricing. At different rates, you receive a different selection of services, capabilities, storage and other options. Businesses like this model, as it allows them to decide at the outset what works best for their needs and usage.

Finally, it's essential to consider the constraints of a provider's licensing agreement. Will you be limited in your ability to implement hybrid cloud or multi-cloud architectures going forward? Will vendor lock in prevent you from innovating and expanding on your own terms? Similar to our discussion of price vs. cost, you should look at subscription models both with regards to what you get for your money and what restrictions you have to abide by.

Compute

Once you've begun to understand the different pricing models that different vendors are using, the next cost question becomes: what am I getting for X price and does this justify the price? Will I actually see some performance value for what I'm spending?

Let's start with compute

Because the cost of compute within a given database can vary widely—AWS offers 55 different combinations of EC2 machines with 8 cores alone—FinOps teams need to critically analyze what they actually need. If they don't, they could overprovision and pay for a faster disk than they're actually using. For every business the answer to "what's my best option?" will be different, but looking closely at details such as instance types for backups will make all the difference.

CPU Chip Efficiency

CPU chip efficiency can also be a significant factor of overall cost. More efficient vCPUs can balance overall database performance and reduce the need to over-provision. As a result, you can get more work done with less vCPU cycles.

Storage

Now that we've considered compute and CPU chip efficiency, it's time to think about storage—both need and cost. Your storage considerations should include:

Defining your storage needs: Cloud-based storage has the greatest impact on cost. Each storage option has different attributes for IOPS (input/output per second) and throughput, typically measured in Mbps (megabits per second). Part of the challenge here is understanding what you get and what the "maximum up to" is. An analogy for automobiles is when the manufacturer reports estimated miles-per-gallon and notes how "your mileage may vary." How you drive, what kind of gas you use and even the weather can affect these numbers, so it can be difficult to get a true measure of performance. That's the same scenario in cloud database cost and performance estimating.

Understanding storage costs: Suppose you think you need 300 GB of storage. Depending on the cloud, you will typically have to pay for 512 GB of storage, even though you only intend to use 300 GB because the next size down only goes to 256 GB.

But you also must take into account what IOPS and Mbps you need—and, of course, these all have significant and immediate cost implications. For instance, if you look at how Azure prices the Ultradisk storage, you need to consider:

- $0.14746/mon per GB
- $0.06132/mon per IOPS
- $0.39566/mon per MBps
- Each of these is different depending on the Azure region

Then, how do you consider the number of IOs and throughput needed to optimize a workload's performance requirements and cost? Furthermore, there are 'Max' and 'Up to' limitations between storage and compute environments.

Finally, the costs between Amazon, Microsoft Azure and Google Cloud are calculated differently. If you are considering between cloud vendors, the challenge in creating a definitive price comparison is complex and time-consuming.

Database control

Cloud vendors' primary DBaaS offerings limit how much you can tune Postgres. The base Postgres configuration file includes roughly 350 parameters to control how the database uses compute, memory, network and storage. With different vendors, however, you only get to control:

- AWS RDS Postgres: 70% of those parameters
- AWS Aurora Postgres: 69 of those parameters
- Azure Flexible Server: 66% of those parameters
- Google SQL Postgres: 51% of those parameters

This greatly limits your ability to tune workloads necessary for balancing database size, workload and concurrency. Because of this, it's harder to drive optimized results.

On top of this, cloud vendors make it overly complex to change these settings and do little to help you understand the best settings for your goals. They want you to increase compute and storage sizings— removing the financial expectations of a move to the cloud in the first place.

Optimize read and write workloads

Standby database instances provide insurance against disasters; however, monetizing the standby databases is a definite "go" in Cloud economics. FinOps teams look at squeezing every vCPU cycle and every bit of memory, and why shouldn't they? This brings a significant balance in the overall TCO, starting with software, compute and other infrastructure resources. You get better read/write performance—if your read workload is significant, this would draw down your compute requirements for the primary workload, as well.

A proven formula to optimize total cloud costs

With all of the key cost factors outlined, your business can begin comparing them. At EDB, we've designed what we like to call a "scientific" method for helping FinOps teams reconcile all this information into something which can inspire meaningful decision-making.

Because you know that you want a database that provides great performance at minimal cost, you can easily eliminate any options at these extremes. This includes both hyper-performant but overpriced solutions (though some edge use cases may opt for this option, it's the exception) as well as solutions that appear cheap up front but offer limited functionality. Ultimately, you're looking for a cloud database that has the lowest cost-per-work element rate (written as $/transaction) and most linear scalability in $/transaction, based on three criteria:

- Concurrency
- Workload
- Database size

Cost-per-transaction in action

Let's look at an example of what $/transaction can look like. When you scale the compute from a 4- CPU environment to a 16-CPU environment, you should see

linear scalability in transaction throughput of 4x, which means the $/transaction would actually decrease, because compute is only part of the cost. For instance, in our research we would see roughly 167,000 transactions/min in a 4-CPU environment and roughly 680,000 in the 16-CPU environment. The overall cost per transaction decreases from $0.07 to $0.04.

Technically, the cost should be the same per transaction. EXCEPT here the database size is the same. Since storage is such a big cost, where the growth is in transaction/second against the same sized database, this means more concurrency and more work but a constant in database size.

Put plainly, look for the cloud database that not only costs the least amount of money per task or transaction, but also one whose increase in price will be the most gradual as you grow. This is a reflection of how well it takes advantage of the primary components—CPU, memory, storage and network—and allows the organization to realize the greatest "tune-ability" with the latest versions.

If we build a model to find this with as few variables as possible (like every good science experiment) and maximize our constants—component consistency, workload consistency and workload sizes—we can can get to an idea of not just who's faster or who's cheaper, but who provides the lowest cost-per-workload.

But the benefits of such a model don't stop there!

Considering dollars-per-transaction also proves a couple of other aspects of technology—do not focus on "fastest" or "cheapest." You want "'optimal'". In the cloud, this means you have to not just pick the right machine and storage type, but configure them correctly, as well. Examples especially exist on storage—the three primary clouds all have storage with tunable parameters, like read/writre performance (input/outputs per second (IOPS)) and/or throughput between the machine-type and storage device (measured in megabytes-per-second). Each one of these have a cost as previously outlined using Azure Ultradisk example—but AWS storage options and Google Cloud Hyperdisk are no less complicated. Finally, machine types have their own limits in what they can handle in IOPS and throughput. The table below shows what should be the same 8-Intel Ice Lake CPU threads and 64GB memory count. Note that the "'max" and "up-to" capacities are highly variable.

This groundwork will make it easier to predict changes in pricing and performance based on trends, automate for better accuracy, configure work and build tooling that optimizes the way your database server is run tomorrow, and the day after and the year after. That means you can find the best database for you now and ensure it remains the best database for your business at every step of your growth journey.

Cloud	Compute Name	Maximum Disk Performance	Maximum Network Performance
Azure	- E8s_v5 -	25600 IOPS / 600 MBps :: Max Burst of 40000 IOPS / 1200 MBps	Max 12.5 Gbps
AWS	r6i 2xlarge -	Up to 10 Gbps	Up to 12.5 Gbps
Google Cloud	n2-highmem-8	Not documented	Default egress 16 Gbps

05

EDB BigAnimal: The right cloud price and performance

EDB BigAnimal: The right cloud price and performance

Today, cost is more important than ever. Yet cloud power is critical, too. Armed with the strategies in this eBook, you can achieve every business' dream—a database in the cloud that's both highly-performant and cost-effective.

By employing FinOps practices: doing the research, comparing the different options and collaborating with financial, development and IT teams to choose the best database solution, you can achieve cost-efficiency now and as you grow—without sacrificing essential capabilities. It's the best of all worlds.

To support you on your cloud journey and give you more control over your database investment, we created EDB BigAnimal. This unique fully managed cloud solution combines open source PostgreSQL backed by high-quality technical assistance for an enterprise database experience unlike any other. With **EDB BigAnimal**, you can choose the subscription that works for you, without any fear of vendor lock-in that might restrict how you deploy and diversify in the cloud.

As the leading contributor to Postgres code, our advice and management services come straight from those who know Postgres inside and out. Now, you have what it takes to innovate and grow and address database issues, long before they cost you.

Within BigAnimal, EDB also allows you to:

- Offer the newest and best components in each cloud
- Bring automations to help tune Postgres to the component decisions + give you access to 98% of all Postgres-tunable parameters
- Bring a best-in-class 99.995% SLA for the highest availability for Postgres in the cloud
- Draw down against cloud commitments by running BigAnimal in your CSP account—thus maintaining consistency and budgetary controls within the same "cloud quota" systems.
- Realize future innovations that will focus on continuing to drive down that cost per workload and bring the best open-source to an organization.

We all know the future is the cloud. You deserve to make the most of it without breaking the bank.

To experience the power of EDB BigAnimal for yourself, **sign up for a free trial today!**

ABOUT EDB

EDB provides enterprise-class software and services that enable businesses and governments to harness the full power of Postgres, the world's leading open source database. With offices worldwide, EDB serves more than 1,500 customers, including leading financial services, government, media and communications and information technology organizations. As one of the leading contributors to the vibrant and fast-growing Postgres community, EDB is committed to driving technology innovation. With deep database expertise, EDB ensures extreme high availability, reliability, security, 24x7 global support and advanced professional services, both on premises and in the cloud. This empowers enterprises to control risk, manage costs and scale efficiently. For more information, visit **www.enterprisedb.com**.