

# Scalability in Cloud Computing

## Table of Contents

- **The Growing Demands of Big Data**
- **Big Data & the Cloud**
- **Cloud Migration: What Enterprises Hope for**
- **The Reality: More Complexity, Higher Costs**
- **The Pepperdata Solution**

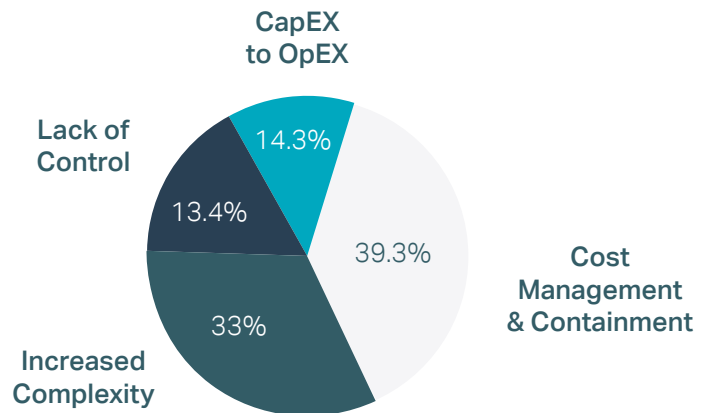


Big data stacks are being moved to the cloud, enabling enterprises to get the most value from the information they possess. But as demand for big data grows, enterprises are pressed to enhance the performance of their cloud assets.

Faced with the complexity of cloud environments, most enterprises resort to scaling up their whole cloud infrastructure, adding more compute, and running more processes.

While these measures do ensure the continuous operation of the entire cloud environment, doing so without expert guidance, optimization, and automation can lead to cloud costs far beyond initial allocation figures, thus sending any cloud initiative or strategy into a cost spiral.

This was highlighted in our [2021 Big Data Cloud Technology Report](#). Surveying multiple industries, we discovered that one in three business organizations are spending 20% to 40% more than intended in the cloud.



Effective scalability in cloud computing is crucial for businesses that rely on big data and the cloud. However, the ability to scale has to be fast, precise, cost-effective, and automated. Scaling has gone beyond human capability.

Cloud scalability, in today's world, will be powered by new technologies, including artificial intelligence and machine learning. This is the only way for businesses to meet today's requirements and standards and get the most value from their data.

## The Growing Demands of Big Data

Big data has become a ubiquitous component of modern enterprise functioning. Large sets of information from disparate business systems and sources are gathered, centralized, and thoroughly analyzed. This enables companies to derive great value and highly actionable insights to further accelerate decision-making, enhance business operations, drive more growth, and foster better relationships with their consumers, partners, and suppliers among others.

Recent [statistics](#) reveal that 52.5% of global organizations have now fully embraced big data technologies. Furthermore, 38% intend to adopt big data in the foreseeable future. The increased spending on big data tech (\$189.1 billion last year) indicates massive growth and adoption up ahead. This year, the world will create approximately [59 zettabytes](#) of data. That figure is expected to

go up exponentially in the future, to around 149 zettabytes by 2024. However, as the demand for big data intensifies, the requirements and resources required to process that data becomes more complex.

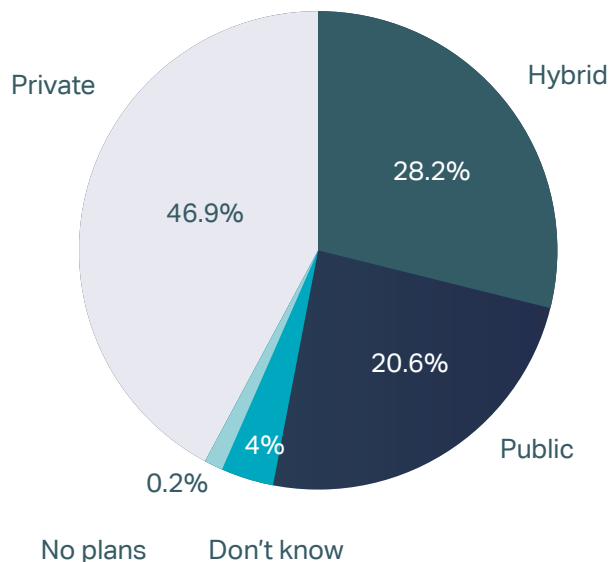
**“The tools and best practices for data infrastructure are evolving incredibly quickly. So much so that it’s difficult to get a cohesive view of how all the pieces fit together.”**

**Matt Bornstein, Martin Casado, & Jennifer Li**, *Data technology experts at Andreessen Horowitz*



## Big Data & the Cloud

Instead of taking the on-prem path, most enterprises sensibly opt to manage more and more of their big data in the cloud. A recent [Pepperdata survey](#) found that almost everyone is shifting from physical data centers to hybrid, cloud, and multi-cloud infrastructures. For example, 28.2% of business leaders are opting for a hybrid cloud strategy.



Here's the challenge: As [Matt Bornstein, Martin Casado, and Jennifer Li](#), data technology experts at Andreessen Horowitz, put it, "The tools and best practices for data infrastructure are evolving incredibly quickly. So much so that it's difficult to get a cohesive view of how all the pieces fit together."

Transitioning to the cloud requires companies to perform a comprehensive review of all business requirements, existing processes, and future needs. After a thorough review, business leaders and IT personnel have to determine which operations and systems must be transferred to the cloud and which ones have to remain where they are. They have to change their business standards and prioritize changes that bring them the greatest return on their investment in the shortest span of time.

To accommodate their cloud transformation, IT Ops need to be in total control over all the processes related to cloud computing. These include comprehensive management over cloud-based resources and network services. The former pertains to the creation, management, and control of resources. The latter focuses on improving resiliency and reducing latency.

Looking from a technology angle, companies have to determine if their existing solutions and processes are compatible with cloud technology and make the necessary changes to ensure a smooth transition and continuous operation. They also need to invest in equipping their employees with the appropriate knowledge to manage cloud systems, assets, and services.

**In short, cloud computing combines big data cloud migration and the ongoing management of cloud systems. This adds up to massive complexity. While cloud vendors trumpet the benefits of the cloud, managing cloud computing after migration is not easy. Post-cloud migration challenges constantly arise, since migration is not a one-and-done. It is a constant, ongoing journey for enterprises.**

## Cloud Migration: What Enterprises Hope for

Among the major drivers for cloud adoption among enterprises is the high cost of building, managing, and maintaining traditional data centers. While the actual cost varies, the average figure ranges from [\\$10 million to \\$25 million per year](#).

Most enterprises see the cloud as a route to agility and cost savings. This is the big pull. Cloud computing provides enterprises with a SaaS model that offers the appeal of ongoing OpEx costs in place of large upfront CapEx costs.

Alongside cost savings, enterprises typically see the cloud as a way to graduate from the complexities that come with managing and updating network hardware and software. In the cloud, you can unlock agility that isn't possible with hardware, and gain the power to scale your resources in a few clicks. With the cloud, enterprises also have access to more services than they would on their own.

Usually, a cloud-based big data infrastructure includes several products such as virtual machines, containers, database management systems, artificial intelligence, and machine learning capabilities. [By moving their big data operations to a cloud infrastructure](#), enterprises can process reams of big data in real time and harness insights almost immediately. With advanced technologies like machine learning and artificial intelligence, the quality of insights they derive from all the data they possess is greatly improved, thus more actionable.

A scalable cloud infrastructure lets you increase and decrease resources such as data storage, processing power, and networking capacity. These changes can be made immediately and easily, resulting in the maximized utilization of the cloud infrastructure, in a matter of minutes, even seconds, on-demand.

Shifting to the cloud also gives enterprises, particularly small and medium-sized organizations, access to big data technologies and modern capabilities that are otherwise unviable. With big data stacks and processes in the cloud, they can get their hands on innovations such as Kubernetes and high-speed databases.

With its highlighted benefits of cost savings, flexibility, and mobility, moving to the cloud is an attractive prospect for any enterprise. It's no surprise that migration is such a trend.

**The problem? The cost savings and new technological powers that enterprises hope will be on the other side of a migration are not nearly as easy to access as they hoped.**





## The Reality: More Complexity, Higher Costs

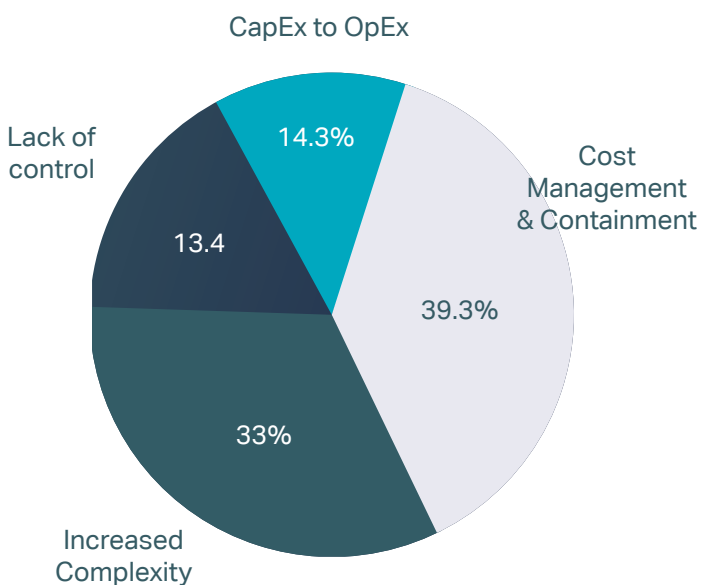
Scaling cloud infrastructure is much easier to get wrong than it is to get right.

For one, there is the new technical complexity. As companies move their typical big data processes and stacks like Kafka, Hadoop, and Spark to the cloud, they often find themselves struggling with usage and performance issues.

Then, there is dealing with the vast selection of available technologies and solutions. With so many options, enterprises are forced to re-evaluate their data architecture and pipelines and ensure the best tools are utilized for each use case, where feasible.

It's not just a matter of choosing the right cloud vendor but of having the visibility needed to see what's happening and manage the performance of these systems.

Next, there is the cost issue. Despite the promise of the cloud as a cost-cutter, according to [our survey](#), nearly 40% of business and IT leaders see cost management and containment as their most pressing concern when running big data stacks in the cloud. Why is this?



**The much-touted flexibility of the cloud gives developers and engineers unlimited leash, allowing them to scale while disregarding the costs of scaling.**

After migrating to the cloud, many organizations discover that cost is a major challenge. Each workload consumes resources such as compute and storage. In the cloud, resources mean money. The problem with most cloud computing users is that they are unaware of the effects of their workloads.

Traditionally, businesses only pay for what they need and use. With cloud computing, it's a different story. Many IT teams are found not ready for the complexities of the cloud. As companies expand their cloud footprint and increase consumption of cloud services, the associated costs go up as well.

Running more workloads may seem like a good idea if you want an application to perform better. However, the more workloads you run, the more resources you consume. Also, to address sudden spikes in performance, users tend to overprovision and place more resources in reserve. This wasn't easy in a traditional data center model. In the cloud, it's very easy. Adding to the cloud bill is the need to re-skill and upskill staff for cloud usage and management, which previous IT managers may have failed to consider.

On top of this, in the cloud, many enterprises struggle to get the insight they need. Choosing a cloud services partner (AWS? Google Cloud? Microsoft Azure?) can be complicated. Each firm brings its own unique big data stacks, packages, features, and benefits.



## Cloud services providers aren't exactly motivated to help customers keep their cloud costs down.

Expenses beyond the compute price exist, and customers may not be aware of all of them. It's very common for enterprises to get a shock when that first cloud bill comes. Given that the cloud infrastructure is wholly owned, facilitated, and monitored by the vendor, customers have minimal

control over their backend operations. Users have control over their data, applications, and services, and overall management generally falls to IT Ops.

However, IT Ops tend to limit the tools and features their users and departments across the enterprise can access. Thus, it prompts teams to purchase their own cloud tools and resources, resulting in overlapping solutions and even more increased cloud costs.

## The Solution: Observability, Autoscaling, & Financial Governance

**Observability** in big data environments is the key to taming the complexity that comes with the cloud.

**Observability refers to the capability to instantly see, track, and evaluate disruptive events and issues within the IT environment the moment they occur. Comprehensive observability tools deliver not just the "what" but also the "why."**

As more businesses adopt containerized workloads and dynamic microservice architectures, outdated practices of monitoring the infrastructure after an event are no longer feasible. It's crucial for organizations to not just monitor their applications but to also better understand their properties and performance.

Observability allows you to see all that happens across the whole big data and IT stack in real time. Your IT teams and developers turn their attention to fixing bugs, instead of spending valuable time finding them. It lets you discover issues across your platform from within a centralized location.

**True observability is knowing why something fails on top of knowing what caused it to fail.** (Check out our webinar on this [here](#).) This is especially important as complex distributed systems become permanent fixtures across the delivery pipeline.

When you opt for scaling in cloud computing, observability helps you understand your systems better so you can scale dynamically. This is why **big data observability is a must-have for modern data lakes and data platforms.** None of this is to say that the cloud can't help reduce operational spend, so long as enterprises have the right frameworks and tools in place.




## Autoscaling

Another effective way to reduce cloud spend is through autoscaling. **Autoscaling** is when your computing platform automatically scales your compute, database, and storage resources based on predefined rules.

For instance, when certain metrics like memory and network usage rates go above or below a preset threshold, the autoscaling mechanism kicks in, scaling up or down, in or out, ensuring your application performance meets SLA standards.

**With autoscaling, your application is always available and has adequate resources provisioned to prevent downtime and outages.**



To fully maximize the potential of autoscaling, you first need to ensure that each instance is fully utilized before scaling up. Default cloud configurations are proving to be more aggressive and costly than need be. Determining the optimal utilization levels and rightsizing the amount of scaling up and down is the tricky part.

There are thousands of concurrent data points to consider and act upon, making it humanly impossible. It is the job of cloud optimization solutions to analyze and act upon performance data in real time to achieve the best results.

## Financial Governance

The absence of a financial governance model means users of your cloud environment can blow through your cloud budget as there is no ceiling for cloud spend. Building and implementing a solid financial governance model for their cloud spend is a step in the right direction.

**One effective approach to reduce cloud expenses and prevent further cloud bill shock is through a chargeback model.**

**Chargeback** is used to allocate the costs of IT resources to the departments or employees that use them. It is primarily applied in situations where dedicated IT resources are involved. Chargebacks are especially relevant in situations where teams and people share IT resources.

It can be powerful for enterprises to measure and attribute the resources consumed by each cloud user or business unit. This benefits users as they become totally aware of (and therefore in control of) their cloud spend. IT administrators benefit as well, as they are able to improve utilization rates while decreasing the number of resources they need to manage.





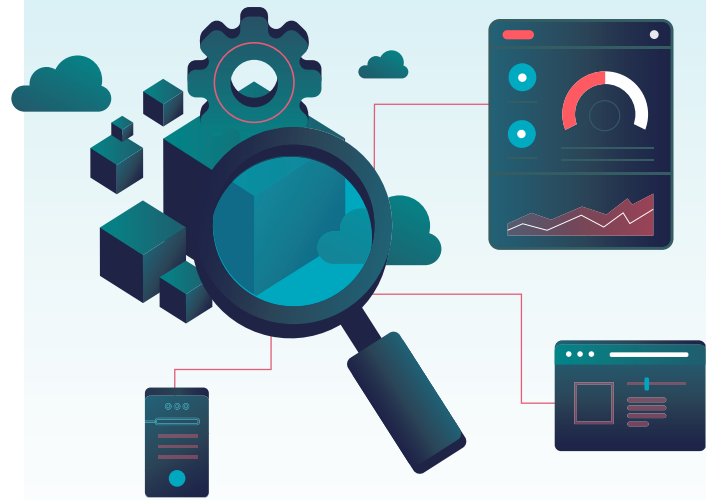


# The Pepperdata Solution

**To achieve full optimization potential, big data solutions need to collect information from all applications and associated infrastructure, analyze the data, and dynamically meet the changes in the resource requirements.**

The Pepperdata solution enables users to understand and manage the performance of complex aspects of cloud computing and big data applications. From a comprehensive and unified platform, users get a holistic and detailed picture of their cloud infrastructure and all its processes, components, applications, and more. Integrated with machine learning and artificial intelligence, Pepperdata looks into their environments; gathers relevant data, including CPU, RAM, disk I/O, and network usage metrics on every job, task, user, host, workflow, and queue; and analyzes them to identify ideal configurations to optimize performance and drive down costs within budget.

With its comprehensive performance suite, Pepperdata gives users observability into their big data stacks. Not only do they get alerted whenever an issue surfaces, but users also get a deep look into why issues and problems occur and how to address them to prevent further issues. Real-time visibility into their big data stack effectively removes complexities associated with cloud computing, allowing users to get great value from their cloud services and more so from their big data.



The Pepperdata managed autoscaling feature gives organizations superior scalability. Leveraging artificial intelligence and machine learning, Pepperdata collects information from all big data applications and analyzes the resource usage of each node in real time to determine the best configurations and apply it immediately to optimize the utilization of CPU, memory, and I/O resources on big data clusters. The net effects are optimized horizontal scaling and eliminated waste.

Superior chargeback reporting helps Pepperdata users track their cloud usage and costs. This enables them to identify and accurately charge business units for their cloud consumption. On top of that, chargeback reporting allows for the rapid and precise analysis of cloud usage trends, thus letting users effectively calculate their ROI and make strategic, data-driven IT decisions.

Visit Pepperdata to [learn more.](#)

©2021 Pepperdata Inc. All rights reserved. Pepperdata and the Pepperdata logo are trademarks or registered trademarks of Pepperdata Inc. All other trademarks are the property of their respective owners. Pepperdata reserves the right to change this document without notice. To ensure you have the latest version of this document, visit [www.pepperdata.com](http://www.pepperdata.com).



**Start a Free Trial**  
[www.pepperdata.com/trial/](http://www.pepperdata.com/trial/)

**Send an Email**  
[info@pepperdata.com](mailto:info@pepperdata.com)

3945 Freedom Circle, Suite 920  
Santa Clara, CA 95054