



White Paper


How to operationalize FinOps to drive cost and cloud efficiency

A guide to succeeding in FinOps with AI/ML-driven automated optimization





Table of contents

- 3** Introduction
 - 5** How organizations solve for FinOps today
 - 8** A better approach to FinOps
 - 14** The benefits of operationalizing FinOps
 - 18** Conclusion
- 

Introduction

With the proliferation and acceleration of digital services, exacerbated by the pandemic, cloud environments are becoming more and more complex and growing exponentially. However, when migrating or deploying these new services to the cloud, many organizations overlook key considerations around the efficiency, cost-effectiveness, and the ROI of their investments in cloud infrastructure. They also fail to recognize how cloud costs as well as cloud procurement differ from traditional IT expenditures. Unlike on-premises environments, where procurement and finance teams closely monitor and manage infrastructure purchases, in the cloud any DevOps or ITOps team, and the code they write, can quickly spin up new resources, often without any formal processes or oversight. As a result, these teams often generate unanticipated cloud spend and waste, as well as infrastructure inefficiencies and sprawl.

96%

of cloud decision makers believe implementing FinOps practices is essential to their cloud strategy

In an uncertain and turbulent economy, organizations now face constant pressure to cut or at least control their cloud spend while supporting business growth and gaining better operational efficiencies. Additionally, many organizations are now looking to implement processes to support more efficient planning and budgeting of their cloud spend, in part to avoid nasty surprises when their cloud bill arrives at the end of the month. But the agility, flexibility, and speed of the cloud is at odds with the need for structure and processes to manage and control cloud costs and maximize ROI.

FinOps enhances the value of cloud investments, but can't work in isolation


The Cloud Financial Operations (FinOps) discipline has emerged over the past few years to bridge and solve these diverse and disparate challenges. It aims to understand and manage cloud costs while maximizing the value from cloud investments and keeping budgets under control.

But to be successful, FinOps needs the support and collaboration of multiple stakeholders, especially DevOps and ITOps teams and site-reliability engineers (SREs), on cloud spending decisions even though these teams are typically not experts in cost management nor is it their top priority. Imposing cloud cost management tasks on these already under-resourced and overstretched teams — and potentially requiring them to compromise on application innovation, performance, and reliability to lower cloud costs — can lead to conflict and frustration and hamper time to market.

FinOps and DevOps must collaborate to drive cost and resource efficiencies

This white paper details the drawbacks of the current practice of relying primarily on visibility and reporting to deliver on a FinOps strategy. It then explains how a successful FinOps strategy can be achieved through a scalable, repeatable process that continuously and consistently optimizes cloud costs while delivering infrastructure efficiencies within the dynamic nature of the cloud and how the business functions.

Thus, to truly realize the potential of FinOps, this whitepaper explains how organizations should operationalize FinOps by seamlessly embedding cost-conscious automated optimization processes within cloud operations. Such an approach will align and empower all FinOps stakeholders to achieve their common goals, instead of imposing potentially disruptive cost management tasks on overstressed DevOps and operations teams after-the-fact, separate from their existing infrastructure optimization processes.



92% of cloud decision makers struggle with FinOps, especially around reducing cloud costs



Cloud decision makers are embracing FinOps¹

- **96%** believe implementing FinOps practices is essential to their cloud strategy
- Only **9%** report having a “mature” FinOps practice
- **92%** struggle with FinOps, especially around reducing cloud costs
- Only **19%** report that their use of discounted cloud purchase options has been effective
- **59%** believe they could further reduce costs with discounted offerings for cloud resources if they had more expertise

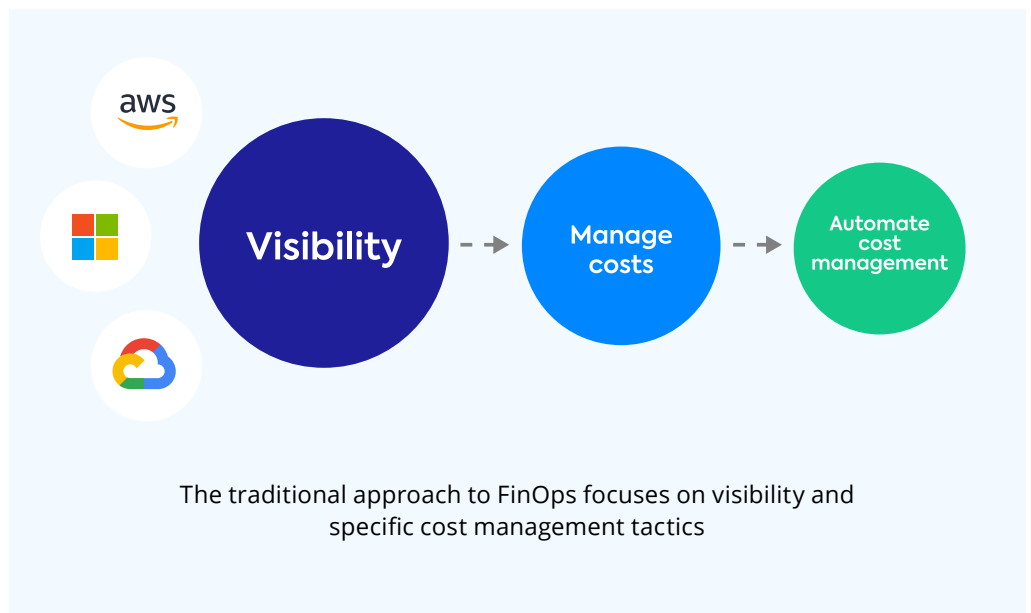
¹Source: [2023 State of CloudOps](#)

How organizations solve for FinOps today

Many organizations focus their FinOps practices and processes on gaining visibility into their cloud environment and the associated costs of cloud resources, believing that an array of reports and dashboards are enough to identify cost optimization opportunities and thus deliver on their FinOps strategy. They rely on these reports to take cost-cutting actions, such as right-sizing instances, removing unused resources, adjusting instance uptime, purchasing commitment plans, as well as allocating charges to cost centers for chargeback/showback purposes. Some organizations do try and inject task-driven automation into these processes to drive repeatability, but for the most part, they are human-driven, periodic, and manual processes.



Learn more about the complexity of managing RIs in [‘The ultimate guide to Reserved Instances’](#)



Limited visibility and cost-cutting tactics cannot solve FinOps challenges



IDC estimates that 20–30% of all cloud spending is wasted²

The reality is that the sheer scale and complexity of today's multi-cloud environment, not to mention the interdependence between costs and cloud infrastructure, make it impossible to truly gain holistic visibility and be able to accurately plan, manage, and optimize cloud spend and resources through mostly manual processes. The monthly bills from the individual cloud providers — each with hundreds of thousands if not millions of lines — require considerable time and effort to analyze and extract even the most basic insights. Experienced, highly skilled, and highly paid analysts often must comb through vast quantities of disparate cloud data, only to gain a partial snapshot of cloud spend.

Moreover, each cloud provider has its own portfolio of resource types, pricing models and tiers, and purchasing packages. Attempting to merge, manage, and optimize the different and inconsistent hodgepodge of entities for each cloud makes this analysis process even more complex, time-consuming, and prone to error. As a result, the cost management actions taken based on this visibility, such as removing idle resources or buying certain types of reserved instances (RIs) may be misguided.

With such an approach, neither visibility nor any actions taken as a result can scale easily. Nor can the organization easily adjust to changes in usage, performance, or business needs, and as such this approach may only serve short-term cost reduction needs at best.

These tasks are also so time consuming that many companies only perform cloud cost optimization reviews once or twice a year. By which time any action taken may be irrelevant or even harmful — not to mention the cost of all the wasted cloud resources the company incurs in the meantime. Nor do they prevent the previous actions that generated the waste and unnecessary spend in the first place.

²Source: IDC Future Enterprise Planning Guides, Control Cloud Costs and Expand Transparency with FinOps, IDC #US50654223, May 19, 2023

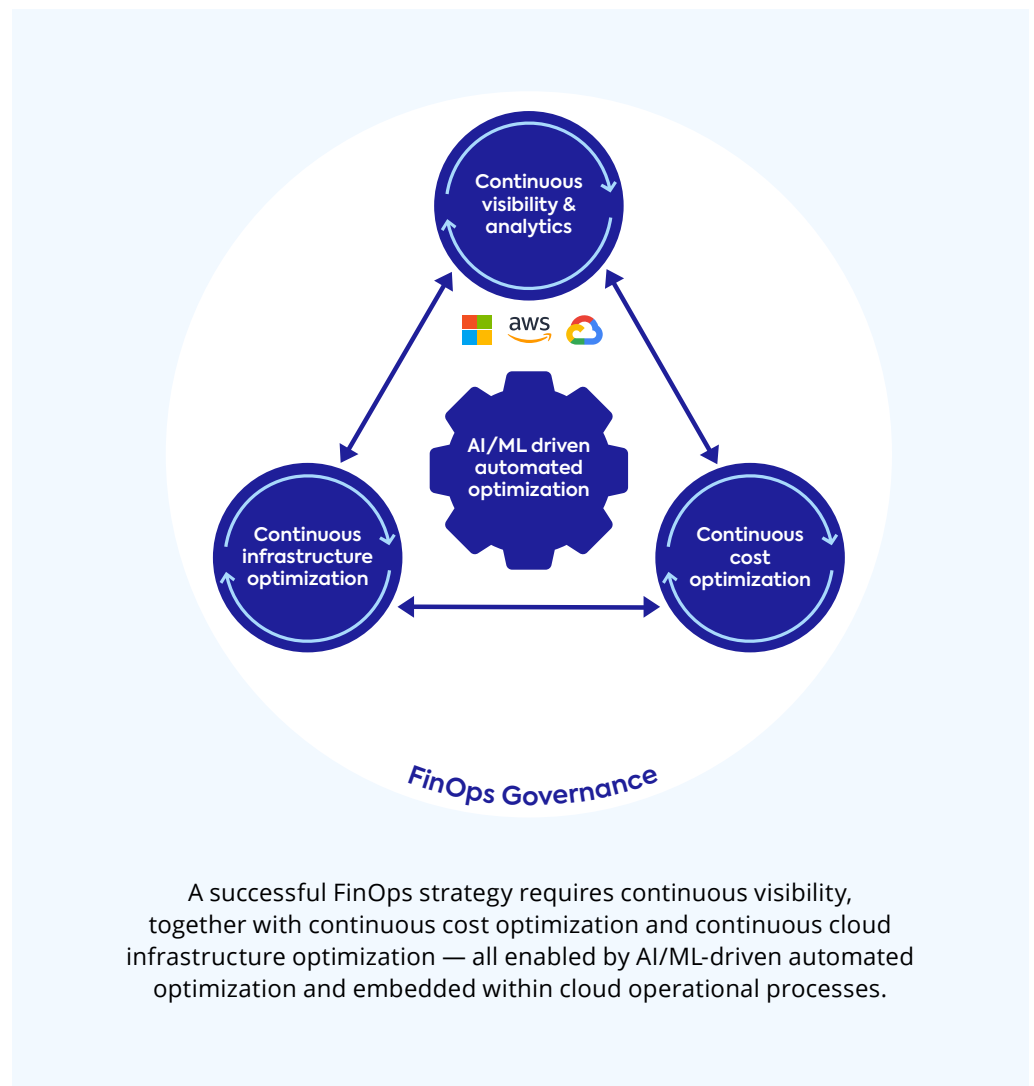
FinOps decisions can't be made without considering the infrastructure

Today, the focus of most FinOps strategies is squarely on cost management and cost reduction actions. Cost decisions, however, directly impact cloud infrastructure. Yet FinOps strategies today do not focus on infrastructure and workload requirements as an integral part of the strategy, even though they are directly tied to cloud spend.

That said, DevOps teams are not always eager to get involved in cost decisions related to infrastructure provisioning. DevOps teams primarily care about ensuring the performance and reliability of their workloads and deploying them into production quickly. Therefore, they may overprovision resources to ensure workload performance and reliability regardless of the cost. Conversely, DevOps teams don't want their deployments delayed because a late-stage cost analysis of their infrastructure forces them to retroactively change the compute selected. While incorporating cost management considerations into infrastructure provisioning *should* be part of DevOps processes, imposing them on DevOps teams often creates friction, generates waste, and may even create application performance problems and deployment delays.

A better approach to FinOps

Gaining a cost- and performance-optimized infrastructure is critical for FinOps success. Given the scale, complexity, and workload demands of multi-cloud environments, four key interlinked processes are necessary to deliver an effective FinOps strategy: continuous visibility, continuous cost optimization, and continuous cloud infrastructure optimization — all powered by automation that utilizes artificial intelligence (AI) and machine learning (ML) to predictively identify and deploy the optimal resources at the right time and at the right cost. It is important to note that these processes do not need to be implemented sequentially. Practitioners can start their FinOps strategy from any one of them and expand over time at a pace that works for their organization and where they are on their FinOps journey.



Shifting FinOps left with AI/ML-driven automated optimization


To truly deliver an effective FinOps strategy, these four processes should be embedded within cloud operations, ideally early on. That means making cloud cost visibility and optimization an integral part of the infrastructure provisioning process for a workload. These processes should be synergistic, seamless, and automated, utilizing AI/ML-driven automated optimization technology to achieve both FinOps and DevOps goals. Importantly, this should be done at the point in time where infrastructure decisions are made, instead of as an after-the-fact reactive task. Such an approach helps deliver a FinOps strategy that is consistent, scalable, and repeatable across the organization and removes the burden from DevOps, cloud operations, and SREs and friction with other stakeholders.

By its very nature, such an approach facilitates productive communication across the stakeholders involved in FinOps, from financial staff and product owners to DevOps engineers, cloud architects, SREs, and others. This, in turn, makes it much easier for everyone to collaborate to deliver a successful FinOps strategy and achieve long-term business goals, as well as drive FinOps governance across the organization. It's certainly far more effective than imposing cost optimization tactics on the operations teams when cloud spend exceeds the budget, thus creating friction, delays and disruption and may be ineffective or plain wrong.

To begin embedding FinOps into cloud operations, organizations need to understand the requirements of these process and how they connect to and impact one another:

Continuous visibility is necessary to drive informed business decisions

Visibility must provide a single source of truth for the organization's multi-cloud environment and its cost; align cloud costs to their applications, business units, or other defined cost centers; and identify meaningful opportunities for cost optimization and infrastructure efficiencies. Thus, visibility needs to be holistic and synthesize data from all cloud providers, and cloud resources must be accurately identified and tagged.




Cloud visibility must also be an active process that evolves as more data becomes available and is contextualized through continuous inputs from both cost management and infrastructure optimization processes. This will give organizations up-to-date, actionable, and prioritized intelligence that can drive effective long-term cost, infrastructure and business decisions.

Many organizations start their FinOps strategy with visibility, reasoning that if they can see it, they can fix it. Some organizations may hold back on optimizing their cloud resources until they have achieved a complete view of their cloud spend and their infrastructure. But visibility is not necessarily the only starting point for a FinOps strategy, and organizations do not need to gain holistic visibility before developing the rest of their FinOps strategy. Instead, visibility can and should evolve over time, and as organizations gain more intelligence and insights into their cloud, they can continuously improve and prioritize the opportunities for cost savings and efficiencies.

Cost optimization and infrastructure optimization must be aligned and continuous

Cost optimization is a core process of a successful FinOps strategy. Cost optimization aims to manage cloud spend efficiently and reduce costs while maximizing business value. Standard techniques include terminating or adjusting sources of wasted spend (e.g., unutilized, underutilized, or unattached resources), right-sizing instances, and achieving the best pricing and terms for cloud compute purchases by optimizing RIs and Savings Plans, for example. Other techniques can include migrating older generation resources to newer, faster, and more cost-effective ones, moving storage resources to more economical options while still meeting failover and data recovery requirements, and moving resources to more cost-effective regions if appropriate.

Many of these processes can be automated, triggered by policy-based alerts. And, given the constant changes in both cost and compute, they must be performed continuously to deliver accurate, impactful, and measurable results.



The third process is infrastructure optimization. Infrastructure optimization ensures that the optimal cloud resources are used to deliver the performance, reliability, and utilization needed to support an application workload. Historically infrastructure optimization was not a foundational component of FinOps. Yet cost optimization processes directly impact the infrastructure provisioned for application workloads and vice versa. So, while it may be cost effective to change a resource type to a cheaper one, such a change could negatively impact the performance or reliability of a critical business application. Therefore, ensuring that application workloads are optimized for both performance and cost must be integral to a successful FinOps strategy.

Aligning these optimization processes and the different teams involved in them requires a technology solution that reduces manual efforts, human error, and friction while delivering optimal outcomes for all stakeholders and the ever-changing needs of the business.

Predict and deploy the right resources at exactly the right moment

Today, many organizations use task-driven automation to scale up or down cloud resources based on prescheduled and/or repetitive events. For example, to support a daily 8:00 a.m. spike in usage, a task is set to automatically scale up a predefined set of resources to support the workload's needs and scale them down at 9:00 a.m. This is not new technology. However, it is a pre-determined, fixed process that is not suitable for the complexity of today's cloud infrastructure and the dynamic needs of application workloads.

What's required is an agile solution, embedded within cloud operational processes, that can proactively — and automatically — identify and deploy the most efficient resources for specific workloads (both containerized and non-containerized) consistently at precisely the right time and at the right cost.

AI/ML-driven automated optimization predicts the unpredictable

“

By 2027 AI-enabled automation will ensure consistent digital infrastructure configuration, performance, cost, and security by reducing the need for human operations intervention by 70% and improving service level objectives.³

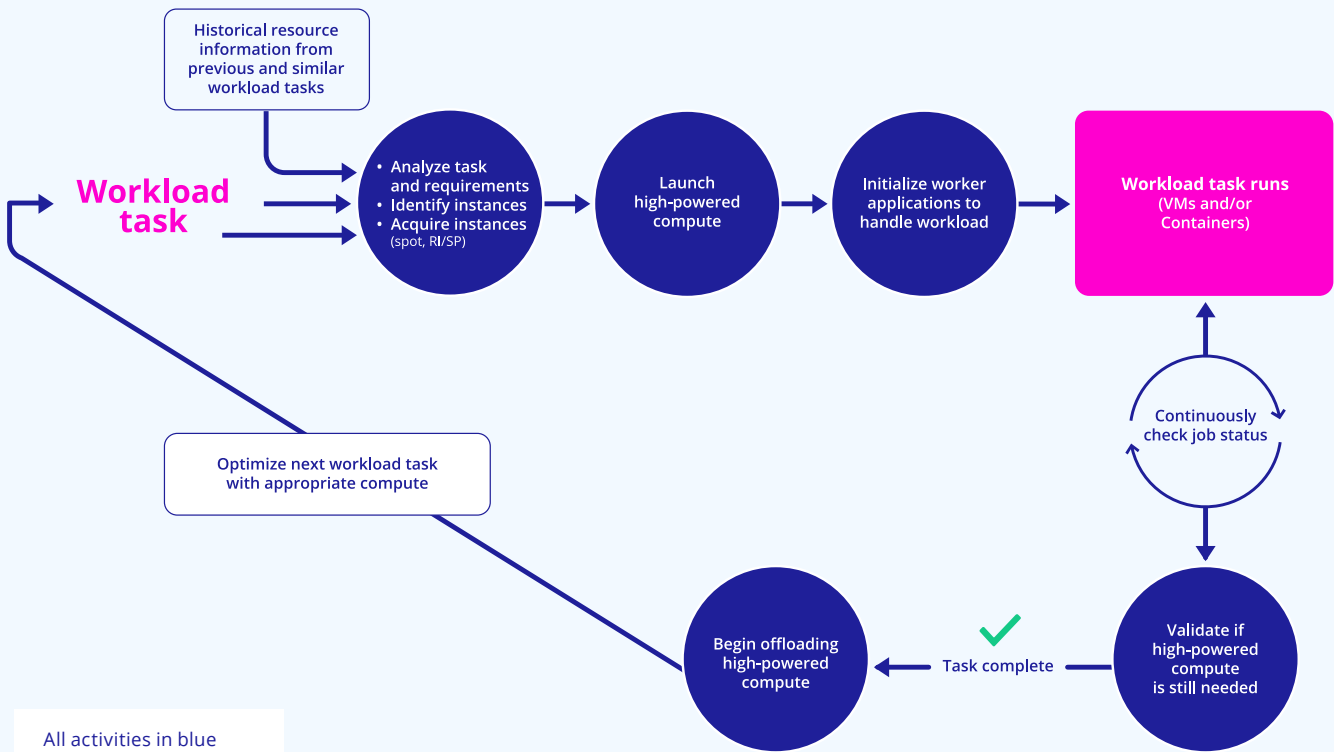
Such automated optimization technology must utilize real-time monitoring and analytics together with AI/ML technology to continuously learn usage patterns, resources, and processes for application workloads across the organization's cloud environment, and correlate insights learned from them. It should also leverage industry best practices and user-driven insights and then create data-driven models that can accurately predict the optimal cloud resources that are needed to support a specific workload's requirements at the optimal cost, at the exact moment at which they are needed.

To support this process, the automated optimization software will acquire the compute power needed for each workload task at the best possible price and offload these resources when they are no longer needed. It should also incorporate organizational policies around acceptable instance types, cost, and other parameters as well as guardrails to address anomalous usage spikes resulting from, for example, misconfigurations or workloads left on by mistake, which can lead to runaway spending and drift.

For example, for a data processing workload, the AI/ML-driven automated optimization technology works by preemptively identifying and scaling up the most efficient resources needed to handle a specific data ingestion and transformation task, such as high-performance instances designed for data-intensive workloads with high I/O requirements. Once that process has completed, the software will determine that a different set of resources, such as general compute instances, are needed to serve the data to users and will automatically scale them up at the right time. Simultaneously, the software will scale down the first set of resources as they are no longer needed. This process will continue until the workload has completed its tasks.

³Source: [IDC FutureScape: Worldwide Future of Digital Infrastructure 2023 Predictions](#).
IDC #US48376222, October 2022

Utilize AI/ML and analytics to predictively identify and deploy the optimal resources for each stage of an application workflow.



All activities in blue reflect the AI/ML-driven automated optimization processes for workload infrastructure.

“

According to Gartner

“Through 2027, 85% of workload placement decisions will need to be continually optimized, because of changing product, availability and cost requirements.”⁴



Learn how finova saved 70% on Azure cloud spend by [optimizing their infrastructure](#)

finova

⁴Source: Gartner, Predicts 2023: Collaborate, Automate and Orchestrate to Optimize Costs and Value During the Economic Crisis, Daniel Betts, Cameron Haight, Hassan Ennaciri, Chris Saunderson, Henrique Cecci, George Spafford, 1 November 2022.

The benefits of operationalizing FinOps

When continuous visibility pairs with continuous optimization of costs and infrastructure, all powered by AI/ML-driven automated optimization, FinOps becomes much more than yet another technology initiative. From finance to business leaders, product owners, DevOps, and operations, operationalizing FinOps provides a scalable, repeatable framework that's designed to consistently support the dynamic nature of the cloud and the business for the long-term.

Some of the key benefits of operationalizing FinOps include the ability to:

1. Cut cloud spend and eliminate waste without sacrificing performance

FinOps is a strategic initiative that seeks to continuously maximize the value of every dollar spent in the cloud while supporting long-term business growth. However, it's still essential to deliver cloud cost savings quickly, especially in the current economy where budgets are under significant stress.

Operationalizing FinOps helps organizations immediately gain accurate visibility into their current costs across their multi-cloud environment, and take the necessary steps to reduce cloud spend by removing waste and inefficiencies without compromising on performance — before the bill comes in at the end of the month. It also helps ensure that organizations are using the most cost-efficient purchasing models for their cloud at any given time.



Learn how Haptik saved up to 90% on [costs compared to on-demand instances](#)

2. Gain more value from cloud investments

Beyond immediate cost savings, operationalizing FinOps enables organizations to use their budget as efficiently as possible over time and doesn't continue to generate unnecessary waste and drift. It also helps ensure that the organization can maximize these efficiencies to progressively see a higher return on investment from their cloud infrastructure.

It's also important to note that cloud providers themselves do not want to generate revenue from inefficient resource utilization, misconfigurations or waste; it's bad for their business. They too seek to ensure that the cloud delivers maximum long-term value, as it increases customer stickiness.

3. Become continuously cloud efficient, right from the start

By operationalizing FinOps, organizations will become continuously cloud efficient for the long-term, consistently ensuring the availability, resiliency, and reliability of their workloads, at the optimal cost, right from the start.

Operationalizing FinOps also empowers cloud stakeholders to become more efficient themselves, giving them back time to focus on more strategic work and deliver higher quality results. This, in turn, helps to reduce the level of frustration caused by having to perform tedious manual tasks, and it improves staff productivity, retention, and job satisfaction.

4. Help achieve sustainability goals

It's worth noting that FinOps can also help reduce the environmental impact of cloud computing. While sustainability may not be the primary goal of most FinOps initiatives, operationalizing FinOps can help organizations achieve their sustainability objectives.

Creating and maintaining an efficient cloud environment and ensuring continuously efficient usage of cloud resources helps to conserve energy as well as minimizes the carbon footprint of servers in cloud data centers.



Learn how
[Choice Hotels](#)
saved over
\$1 million
a year

5. Make informed business decisions on priority initiatives

Continuous, holistic visibility gives organizations the necessary insight to identify and proactively prevent resource overprovisioning and waste while optimizing costs and infrastructure. It also empowers organizations to operate within their budget and avoid unanticipated surprises in their cloud bills.

But beyond the core benefits of visibility, companies must also be able to understand their cloud costs in a business context to assess and create competitive differentiation and drive business growth. They need to know the costs of running their applications, products, and business units — for which cloud spend is often a significant portion of the budget — and be able to assess the ROI from these investments. Continuous, holistic visibility, which aligns cloud spend to its respective cost centers, gives business leaders an accurate picture of their costs and surfaces insights into what drives them as well as their impact on application performance. This visibility can provide useful insights, for example, to help architect business applications in a way that drives greater profitability or define a product's growth strategy early in its lifecycle.

It's important to note that there are still many cloud-related business decisions that humans — not technology — must make themselves, such as whether to invest in one application over another based on market viability. But stakeholders often spend so much time attempting to understand cloud cost data and figure out how to use it that they struggle to find the time to develop an overall cloud strategy that aligns with business objectives. While visibility and automation on their own won't solve business challenges, they can help free up time to allow stakeholders to focus on more critical business decisions that only they can make.

6. Remove the burden from DevOps and drive innovation

Operationalizing FinOps removes the considerable burden of cost-efficiently optimizing infrastructure for application workloads from stakeholders, especially DevOps, freeing them to focus on innovation. So, instead of forcing cloud cost-cutting measures on these teams and micromanaging their implementation, which no doubt creates friction, DevOps teams are now empowered to deliver cost-optimized infrastructure from the get-go, which in turn helps support on-time, successful deployments. And the faster a new application or capability deploys, the quicker it can drive business productivity and growth.

7. Facilitate stakeholder collaboration and accountability, and drive governance

Many people often mistakenly assume that FinOps is a defined role that falls solely on the shoulders of finance or other departments in charge of cloud spend. However, FinOps is not a role but a cross-functional discipline that relies on input and support from across the business.

Operationalizing FinOps gives finance, business leaders, and other stakeholders an understanding of how infrastructure requirements align with cloud costs and enables them to set and enforce governance policies consistently across the organization. Meanwhile, DevOps engineers and other technical teams gain an understanding of how their work impacts budget decisions. By operationalizing FinOps, all those involved, especially DevOps, become part of the solution. It gives them agency. DevOps teams are now part of the decision-making process and thus can truly understand and accept tradeoffs.

Beyond DevOps, all stakeholders are now far more likely and willing to engage and collaborate around cloud initiatives and take ownership and responsibility not only for their cloud spend and usage but also for the efficiency of the applications that run the business. And when other organizational leaders in IT and SecOps are part of the conversation, they too can ensure that all decisions align with cloud governance, security, and compliance best practices across the organization.

8. Reduce the risk of knowledge loss

Operationalizing FinOps also helps keep the FinOps strategy on track in the event of staffing changes and organizational restructuring. When key individuals responsible for making cloud spend and/or infrastructure optimization decisions leave the organization, there's a substantial risk of losing valuable legacy, often undocumented, knowledge known only by a select group of employees. By operationalizing FinOps, this intelligence is determined by and built into existing technology and processes, thereby eliminating this risk and any potential knowledge loss.

Conclusion

Every organization today, irrespective of shape or size, should implement FinOps to manage their cloud costs. But FinOps strategies that focus on visibility and cost cutting tactics alone are not sufficient. Within an environment of ever-changing cloud architectures and workload requirements, operationalizing FinOps gives the organization control over their cloud instead of allowing the cloud and its costs to control them.

Operationalizing FinOps delivers a scalable process that enables organizations to consistently drive continuous cloud cost optimization together with infrastructure optimization, and provides actionable intelligence to help make informed decisions that support IT, financial, product, and business goals for the long term. It also helps build a culture of accountability and collaboration, drive governance, and foster creative solutions and innovation across dynamic cloud environments.

Organizations, however, should develop and apply their FinOps strategy at a pace that works for them. With the right solution for FinOps, organizations can gain early quick wins by generating immediate cost savings while building trust and laying the foundation to increase those savings, improve visibility, and gain greater and greater infrastructure efficiencies over time. Progressively incorporating AI/ML-driven automated optimization into FinOps processes, supported by checks and balances, will build confidence in the outcomes to enable the organization to continually advance their FinOps journey.

About Spot by NetApp

The Spot by NetApp portfolio enables cloud operations teams to deliver scalability, performance and security for cloud infrastructure and applications — at the lowest possible cost — through continuous automation and optimization combined with deep visibility and governance. From cloud-native startups to global enterprises and managed services providers, thousands of customers rely on Spot solutions to unlock the full value of the cloud.