

WHITE PAPER

AWS COST REDUCTION: 3 ESSENTIALS YOU MUST KNOW




In a recent [InterVision/Pulse survey](#), 92% of respondents said that their organization expects the IT team to be operating with a cloud-first approach, yet 45% were unsure or not meeting their cost objectives in the cloud.

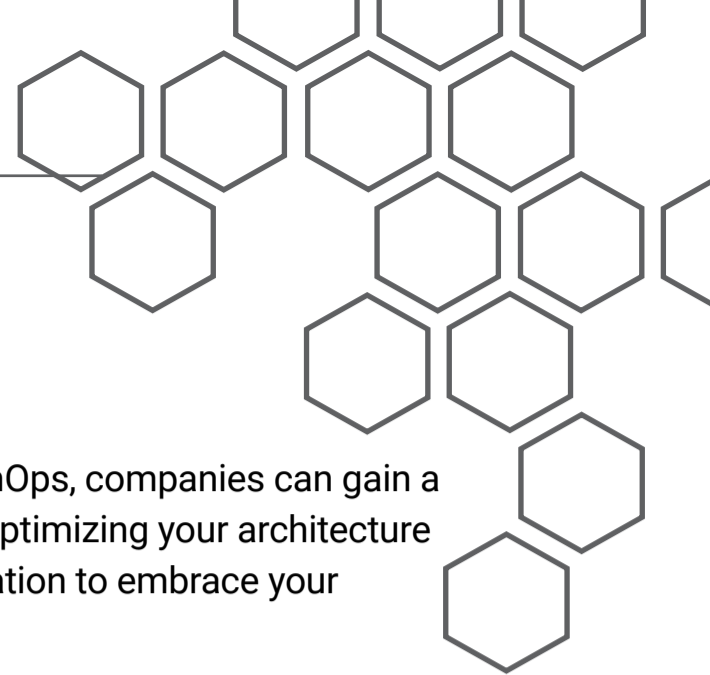
The adoption of the AWS cloud can certainly empower a business, but if migrated incorrectly and managed poorly, it can also drag a business down. The answer to success is all in how an IT team chooses to strategize and execute on their AWS cloud migration, and then optimize their AWS stance in an ongoing fashion. Indeed, migrations do have rippling impacts on total cost of ownership, both in cases of good or poor migration strategies.

There are many complexities that could create overspending in the cloud:

- Unclear cloud cost visibility
- Cloud sprawl in under-utilized or wasted resources
- Difficult-to-understand pricing and billing
- Management and process challenges with governance
- Incorrectly sized deployments on architecture
- Inaccurate forecasting

The problem is that, in addition to these commonly encountered complexities, there's a misunderstanding that going to the cloud will automatically create cost savings for an organization. The answer to cost savings in the cloud instead, lies in three key areas: 1) architecture, 2) FinOps and 3) governance. In this white paper, we'll explore these three areas of AWS optimization.





PART 1 – ARCHITECTURE AS A STEP TO LOWERING COSTS

By integrating expertise in cloud architecture, governance and FinOps, companies can gain a holistic approach that is guaranteed to save more. Focusing on optimizing your architecture in AWS is the first step to robust security while setting the foundation to embrace your cloud-first strategy, without financial complexity.

Areas in Cloud Architecture that demand attention to achieve cost optimization:

- Infrastructure:** Identify AWS services and infrastructure that are unused to eliminate unnecessary spend
- Data Transfer:** Identify and target opportunities to reduce data transfer costs based on client use cases
- Compute:** Right size compute resources and review pricing models to optimize performance and cost
- Storage:** Pinpoint opportunities to minimize AWS storage costs while ensuring availability
- Scheduling:** Identify scheduling opportunities to turn up or turn down resources, which can save up to 70% versus running them 24x7

Cloud cost management requires unique expertise in the realm of cloud architecture, and often cloud teams either don't have the skills on staff or are too busy to focus on the technical and operational processes necessary to control and reduce costs. This issue, coupled with complex billing and lack of visibility into spending trends can result in sprawled workloads, stalled migrations and eroded confidence in cloud strategies. That's why having an expert at the helm of your cloud architecture is essential.

There are several reasons why an organization considers cloud: agility, flexibility, accessibility, cost, performance, digital transformation, etc. To get at the root of your company's decision, many CIOs create a vision statement and identify the key benefits you are seeking. Creating a firm business case and realistic timeline can also allow you to accurately track the outcomes achieved following the cloud deployment.

As part of establishing your cloud strategy, you should also examine the process steps of a cloud migration to ensure alignment on roles and responsibilities as well as the desired timelines. You will need both stakeholder and executive buy-in to properly ensure a successful cloud migration, so you should embed this into your strategy. Consider how you will keep stakeholders updated on your progress. Requesting full investment from leadership shouldn't come as a surprise to them – it does demand an open dialogue though, to properly set expectations for time and cost commitments.

Begin assigning your key resources to critical planning steps. In the pre-migration planning, you should inventory all applications, datasets, and other interdependencies to be sure nothing gets let behind. As part of this inventory, your IT team should create an "app dependency map" to be sure everything interconnected is connected again in the cloud. It will also be helpful in grouping datasets by their importance or priority for the actual transition. An early-stage TCO analysis can also help with application-specific planning. This more detailed planning will help you see where potential roadblocks might occur and will keep your commitments to the business in line.

PART 2 – FINOPS: SAVINGS PLANS, RESERVED INSTANCES AND SPOT INSTANCES EXPLAINED

[TechTarget](#) describes FinOps as “a framework for managing operating expenditures (Opex) in the cloud. An important goal of FinOps is to help cross-functional teams within the same organization maintain financial accountability for cloud services.” With this definition in mind, FinOps in AWS essentially becomes your process for holding your IT team accountable to organizational spending goals.

AWS has pricing models that allow for provisioning resources in the most cost-effective way: On-Demand, Spot Instances, Reserved Instances and Savings Plans. On-Demand is what AWS offers clients without any commitments or provisioning presets—which can be costly if not leveraged with a careful approach. When driving down AWS cost, a good beginning point is leveraging native AWS tools meant for the specific optimization of cloud resources, rather than allowing for runaway storage or computing costs.


AWS Reserved Instances

Reserved Instances (RIs), for example, offer significant discounts with the commitment to a set instance type or level of on-demand compute usage, usually over the course of 1 to 3 years. Reserved Instances can be purchased for a specific Region or for a specific Availability Zone, and AWS offers no upfront, partial upfront, and all upfront pricing options, to allow clients to choose the option that works best for them.

AWS Spot Instances

In contrast, Spot Instances are best suited for temporary workload needs such as high-performance computing for big data projects. AWS defines a **Spot Instance** as a “spare EC2 compute capacity available at discounts of up to 90% off On-Demand prices with no long-term commitment required.” Integrated into multiple AWS services like EC2 ASGs, EMR, ECS or AWS Batch, Spot is an ideal cost savings tool for when you have a buffer in place or a queue where multiple requests are able to receive independent attention, such as non-critical workloads or development environments.

In the white paper, “[Cost Optimization Pillar: AWS Well-Architected Framework](#),” AWS recommends matching the maximum price for Spot Instances to the On-Demand rate, to ensure a lower-than cost; being flexible across as many instance types as possible and where your workloads will run; and designing your workloads for continuity. Together, these aspects will help you achieve the most of what Spot instances can offer.



According to AWS, “Savings Plans provide discounts for AWS Compute services such as EC2, Fargate, and Lambda. When you make the commitment, you pay that commitment amount every hour, and it is subtracted from your On-Demand usage at the discount rate.”

AWS Savings Plans


Similar to Reserved Instances (RIs) and Spot Instances, AWS Savings Plans are an easy cost-saving feature that can provide optimization in exchange for a commitment to a particular level of usage. According to AWS, “Savings Plans provide discounts for AWS Compute services such as EC2, Fargate, and Lambda. When you make the commitment, you pay that commitment amount every hour, and it is subtracted from your On-Demand usage at the discount rate.”

While AWS Savings Plans have less flexibility, they can provide clients with a higher discount rate. AWS offers two types of Savings Plans, as described below from the [AWS website](#):

Compute Savings Plans provide the most flexibility and help to reduce your costs by up to 66%. These plans automatically apply to EC2 instance usage regardless of instance family, size, AZ, region, OS or tenancy, and also apply to Fargate and Lambda usage. For example, with Compute Savings Plans, you can change from C4 to M5 instances, shift a workload from EU (Ireland) to EU (London), or move a workload from EC2 to Fargate or Lambda at any time and automatically continue to pay the Savings Plans price.

EC2 Instance Savings Plans provide the lowest prices, offering savings up to 72% in exchange for commitment to usage of individual instance families in a region (e.g. M5 usage in N. Virginia). This automatically reduces your cost on the selected instance family in that region regardless of AZ, size, OS or tenancy. EC2 Instance Savings Plans give you the flexibility to change your usage between instances within a family in that region. For example, you can move from c5.xlarge running Windows to c5.2xlarge running Linux and automatically benefit from the Savings Plans prices.

Another feature is EC2 Fleet, which according to AWS, allows you to “define a target compute capacity, and then specify the instance types and the balance of On-Demand and Spot for the fleet. EC2 Fleet will automatically launch the lowest price combination of resources to meet the defined capacity.”



PART 3 – IMPLEMENTING GOVERNANCE CONTROLS

Governance in the cloud refers to the process of IT policy making and emphasizing cybersecurity and compliance under a single umbrella. When created and applied correctly in the cloud, your AWS governance doesn't have to slow down development and operations. In fact, it can act as an accelerator, adding speed to a process that normally would require the cybersecurity team to review applications before or shortly after being released. With a strong governance approach, the security guardrails are clear, the compliance is assured, and the DevOps team can run as fast as they like.

AWS governance controls, when implemented and optimized, can do the following for a business:

- Drive greater flexibility and scalability for builders
- Provision resources and applications that are pre-approved for use
- Operate with assurance of security, allowing for the greatest speed
- Enable a strong ecosystem of third-party solutions

AWS Tools for Governance

If you've migrated and organized your AWS environment with adherence to the AWS Well-Architected Framework, you can use a lot of native AWS tools to empower your company's governance, which include the following:

- AWS Control Tower
- AWS Organizations
- AWS Budgets
- AWS License Manager
- AWS Cloud Formation
- AWS Service Catalog
- AWS OpsWorks
- AWS CloudWatch
- AWS CloudTrail
- AWS X-Ray

WRAPPING IT ALL TOGETHER FOR TRUE OPTIMIZATION

Indeed, the best practice to achieving cloud cost savings in AWS lies in finding the right mixture of what you need, approaching each AWS featured resource and tool from a holistic viewpoint, which helps maximize workload cost optimization with performance. Furthermore, by taking a larger view of your cloud objectives, you can better see the forest through the trees when it comes to cloud spending, thereby allowing you to make better decisions when it comes to your unique AWS environment and organizational goals.

However, if you find your IT team burdened with too many tasks to dedicate attention to cloud optimization, you might want to seek help from a trusted expert. In [InterVision's Pulse survey](#), 51% noted that IT personnel lacked time to focus on cost management in the cloud. If this is the case for your organization, this means that your overall strategy could be suffering, hindered with the weight of overspending.

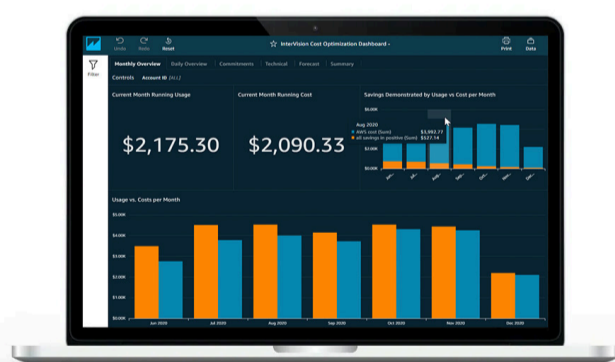
InterVision's Cloud Cost Optimization Service

InterVision is proud to offer a [Cloud Cost Optimization Service](#), which aims to help midsize and enterprise organizations create the most ideal AWS environment for their business goals. With an SLA-backed guarantee to identify 30% or more savings on monthly AWS costs (restrictions apply, [contact us](#)), this service delivers more savings for clients than any other offering in the market.

Delivered as part of our [Cloud Migration Lifecycle Assurance \(CMLA\)](#) program or as a stand-alone offering, InterVision's Cloud Cost Optimization Service delivers immediate savings for clients while helping them establish a long-term cost savings culture within their organization.

Benefits of InterVision's Cloud Cost Optimization Service:

- Access Guaranteed Savings:** Reduce AWS service costs by 30% or more
- Realize Near Immediate ROI:** Get savings within the first month with our quick start process
- Offload Cost Management Burdens:** Relieve your high-demand cloud staff from complex cost management tasks
- Create an Optimized Foundation:** Establish a strategy for ongoing savings for your cloud adoption



One of the trickiest areas of cloud computing is knowing how much to provision where and when, what your spending rate is, and whether anything needs attention—all in one place for easy visibility and fast action. This is where InterVision's proprietary, customizable dashboard that comes with our Cloud Cost Optimization Service can help to provide both business and technical level insights.

The dashboard provides insights in six key areas:

- Monthly and Daily Overview
- Commitments
- Reserved Instance Coverage
- Technical Breakdown
- Billing Breakdown
- Forecasted Spending



As an APN Premier Consulting Partner, InterVision's has more than a decade of experience helping organizations simplify their cloud operations. By integrating our expertise in cloud architecture, governance and FinOps, we deliver a holistic approach that is guaranteed to save you more.