# Accelerate Public Sector Transformation

10 ways to optimize costs
and innovate with AWS

# Table of contents

INTRODUCTION

# Optimize costs and accelerate innovation

Public sector C-suite and IT leaders are facing enormous pressure to do more with fewer resources. They are tasked with finding solutions that reduce costs, increase agility, accelerate innovation, and, ultimately, deliver transformational business value. In addition, security remains a top priority, with increasing demands around data protection and privacy. Security leaders are discovering that a secure cloud infrastructure can support and even accelerate the speed and innovation needs of their organizations.

The cloud offers nearly unlimited opportunities to meet these goals—and beyond. From streamlining operations to delivering better user experiences securely, the cloud enables organizations to yield substantial cost savings while innovating on mission delivery.

Not surprisingly, public sector leaders are looking for ways to speed up their digital transformation so they can improve operational efficiencies and invest in innovation. Wherever you are in your cloud journey, whether it's migrating enterprise applications to Amazon Web Services (AWS) or building cloud-native applications to innovate faster and create new customer experiences, enabling organizations to achieve more while lowering costs is central to how AWS builds cloud technology.

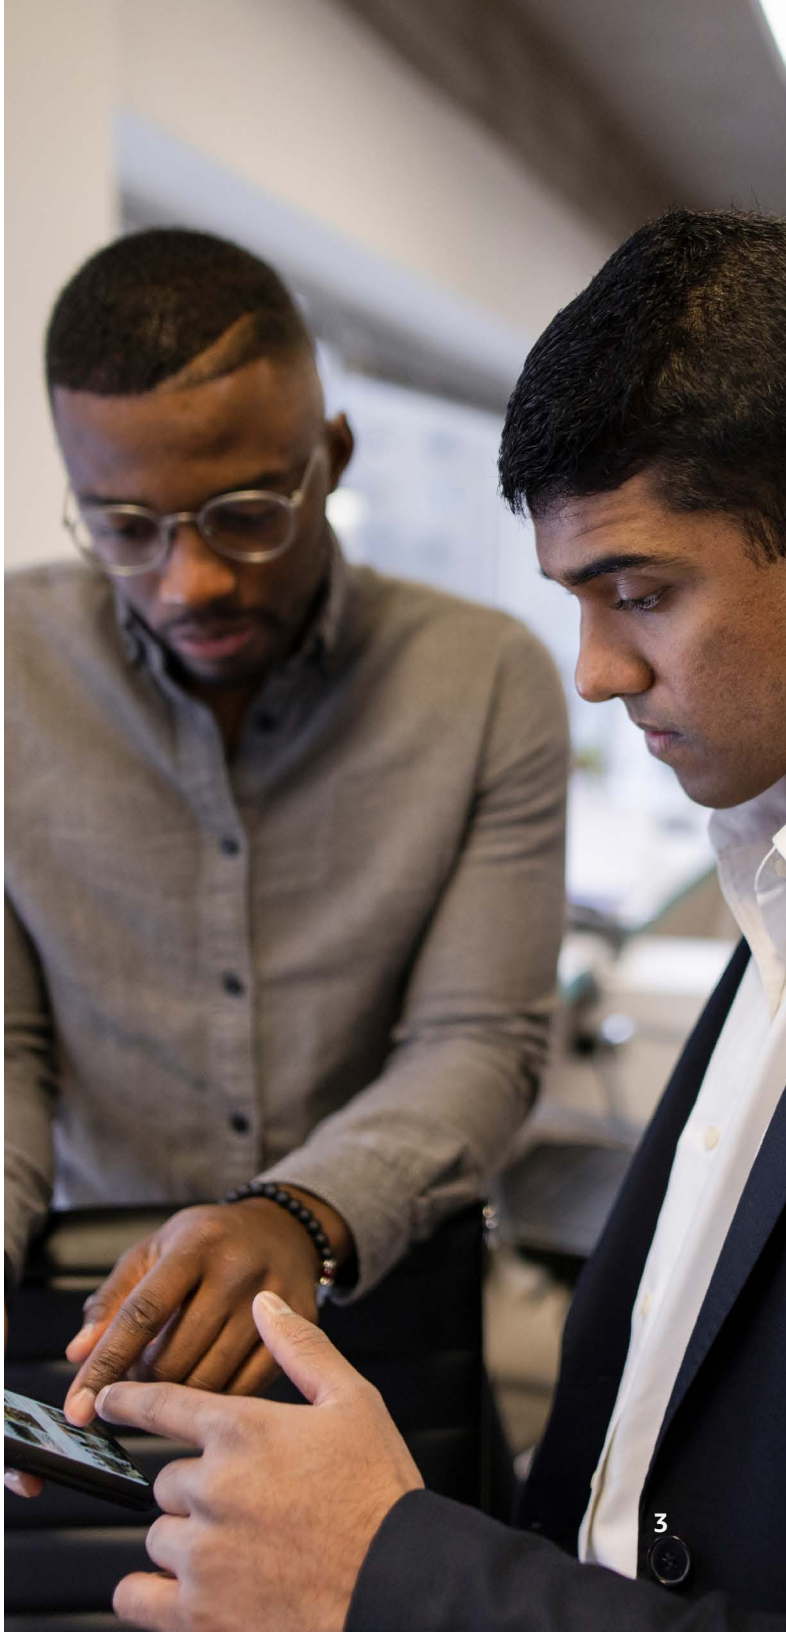

# Getting started with saving costs

Where does cost optimization start? Thinking about cloud migration as it applies to your organization's unique challenges, goals, and opportunities is a place to begin. You can start with migration to AWS for rightsizing infrastructure and reducing licensing costs with a consistent architecture across IT environments. Transformation begins with access to new services to drive agility and tools to drive deeper insights from data.

AWS enables you to take control of costs and continuously optimize your spend while building modern, scalable applications to meet your needs. Our breadth of services and pricing options offers the flexibility to effectively manage your costs and still keep the performance and capacity you require. AWS is dedicated to helping customers achieve the highest saving potential.

AWS can help you wherever you are in your cloud journey—with 10 ways to optimize costs—so you can focus on executing your mission.

# AWS: A cloud pioneer

With more than a million active customers and a global cloud presence, AWS has experience in helping organizations build, run, and scale applications in the cloud.

AWS pioneered cloud computing in 2006, years before any other cloud provider. Since the beginning, we have helped customers migrate their workloads to the cloud, offering more performance at lower costs than any other cloud provider. In fact, organizations that moved to AWS from on premises increased administrator productivity by an average of 66 percent and achieved an average of 20 percent cost savings on infrastructure.[1]

AWS relentlessly optimizes the efficiency of AWS services and has reduced prices 129 times since its inception, whether you're experimenting a new idea or scale your application as your needs grow. Helping organizations maximize savings so they can focus on innovation is at the core of what we do.

This eBook guides you through 10 ways you can reduce costs and make room for more innovation by optimizing your infrastructure investment. You will also learn how leading organizations, such as Salesforce and ENGIE, are using our differentiated solutions to achieve high performance at lower costs, driving their application innovation and maximizing the value of their cloud investments.

[1] Pastor, R., et al., "The Business Value of Migration to Amazon Web Services," The Hackett Group, February 2022

METHOD #1

# Migrate your applications to AWS

Enterprise applications (Microsoft, Oracle, SAP, VMware, and mainframes) start with migration to AWS for rightsizing infrastructure and reducing licensing costs, providing a consistent architecture across IT environments.

You can migrate any workload—applications, websites, databases, storage, physical or virtual servers, and even entire data centers—from an on-premises environment, hosting facility, or other public cloud to AWS. At every step, you can leverage our years of experience to build your organizational, operational, and technical capabilities so that you can start gaining business benefits from the beginning of the migration process.

## Saving on third-party licensing costs

AWS gives you a way to save on third-party licensing costs and run your resources more efficiently: the **AWS Optimization and Licensing Assessment** (AWS OLA).

The AWS OLA is a free program for new and existing customers. It gives you a clear overview of your current costs and potential savings by assessing your current on-premises IT environment, third-party licenses, and projections of running these workloads on AWS.

The AWS OLA will provide you with a report that will model your deployment options using existing licensing entitlements. These results can help you explore available cost savings across our flexible licensing options. The AWS OLA can also be used in combination with the **AWS Migration Acceleration Program (AWS MAP) for Windows**, providing you with tools, support, and resources during your cloud migration.

## Innovate faster, reduce costs, and enhance security

You can take advantage of better performance at lower costs just by using AWS. That's because AWS collaborates not only with major processor manufacturers to offer you the latest generation processors, but also, for nearly a decade, AWS has invested in designing and producing silicon optimized for the cloud. Our custom silicon enables us to offer you industry-leading performance, enhanced security, and faster innovation—at lower costs.

# The AWS Nitro System: Re-imagined virtualization infrastructure

The **AWS Nitro System** is the underlying infrastructure for our modern **Amazon Elastic Compute Cloud** (Amazon EC2) instances. It delivers better performance and price of any cloud provider, along with enhanced security and speed of innovation.

**Best price performance:** With the AWS Nitro system, Amazon EC2 instances can deliver more than 15 percent higher throughput performance on some workloads as compared to other major cloud providers running the same CPU. Dedicated Nitro Cards enable high-speed networking, high-speed **Amazon Elastic Block Store** (Amazon EBS), and I/O acceleration. Not having to hold back resources for management software means more savings that can be passed on to end users. Additionally, the **Elastic Network Adapter Express** (ENA Express) for Amazon EC2 instances, introduced in 2022, delivers up to 25 Gbps of single-flow bandwidth—up from 5 Gbps—at no additional cost to customers. This latest Nitro-based innovation improves price performance by providing customers with more bandwidth without incurring additional costs.

**Enhanced security:** The Nitro Security Chip provides a simple hardware-based root of trust, enabling the most secure cloud infrastructure with a minimized attack surface. Virtualization and security functions are offloaded to dedicated hardware and software. Additionally, a locked-down security model prohibits administrative access, including those of Amazon employees, eliminating the possibility of human error and tampering.

**Speed of innovation:** With the AWS Nitro System, functions are modularized. It breaks the Amazon EC2 architecture into smaller blocks by offloading the virtualization functions onto dedicated hardware. These blocks can be assembled in many ways, which delivers the flexibility to design and rapidly deliver EC2 instances with an ever-broadening number of compute, storage, memory, and networking options. The results are lower costs and the ability to bring more workloads to the cloud, enabling you to increase your pace of innovation.

## AWS firsts – innovation for your digital business

- First to offer AMD, Arm, and MacOS instances in the cloud
- First with custom-built silicon for general purpose, high performance computing (HPC), and ML workloads in the cloud
- First with a custom-built virtualization infrastructure
- AWS pioneered serverless computing with the launch of AWS Lambda

METHOD #2

# Choose the instance type that matches your application needs and budget

## AWS instance types

AWS has more than 600 instance types, exceeding any other cloud provider. Each instance type provides a choice of processor, storage, networking, and operating system, so you can choose the instance configuration that best fits your specific workload. And each instance type includes one or more instance sizes, allowing you to scale your resources to the requirements of your target workload.

AWS gives you the flexibility to change your instance type as quickly as your needs change, eliminating overhead costs for unused resources.

Amazon EC2 instances fall into six categories:

### General purpose

Our most popular instances provide a balance of CPU, memory, and network resources and are ideal for running web servers, containerized microservices, caching fleets, and development environments. One of the main distinctions within this class is between instances with fixed (**Amazon EC2 M7g**, for example) versus burstable (**Amazon EC2 T4g**, for example) performance.

### Compute optimized

Good for compute-intensive, CPU-bound, demanding applications, such as frontend fleets for high-traffic websites, on-demand batch processing, distributed analytics, video encoding, dedicated gaming servers, and high-performance science and engineering applications. These instances offer the highest ratio of virtual CPUs to memory than the other families and the lowest cost-per-virtual CPU of the EC2 instance types.

### Memory optimized

These instances are ideal for memory-intensive applications, such as real-time big data analytics, in-memory databases, enterprise-class applications that require significant memory resources, or general analytics, such as Hadoop or Spark.

### Accelerated computing

Instances in this category include additional accelerators and GPUs, FPGAs, and machine learning (ML) chips that provide massive amounts of parallel processing for tasks such as graphics processing, **ML training**, **ML inference**, and **HPC**.

### Storage optimized

Ideal for tasks that require local access to very large amounts of storage, extreme storage performance, or both. Instances are available that include both large-capacity HDD and extreme low-latency local NVMe SSDs. Choose from the industry's broadest portfolio of storage solutions optimized for your block, file, and object data.

### HPC optimized

HPC instances are purpose-built to offer the best price performance for running HPC workloads at scale on AWS. HPC instances are ideal for applications that benefit from high-performance processors, such as large, complex simulations and deep learning (DL) workloads.

METHOD #3

# Optimize costs for Windows applications

Are you looking to boost the reliability and availability of your Windows Server workloads? Or maybe you've reached the end-of-service (EOS) cutoff, where your SQL Server workloads will lose their patches and security updates.

You're not alone. Customers across multiple industries have been running Windows workloads on AWS since 2009, long before it was possible with other cloud providers. AWS supports everything you expect to build and run on Windows, including Active Directory, .NET, SQL Server, Windows desktop as a service (DaaS), and supported versions of Windows Server.

With AWS as the foundation for your Windows environment, there's no limit to the range of business benefits you can achieve, including:

**Lower total cost of ownership (TCO):** You can reduce your five-year cost of operations by 56 percent and gain 37 percent lower infrastructure costs, delivering up to 442 percent ROI over three years.

**Licensing options:** AWS has helped many customers break free from restrictive licensing scenarios. By purchasing license-included instances of Amazon EC2 or **Amazon Relational Database Service** (Amazon RDS), you get new, fully compliant SQL Server licenses from AWS.

**License optimization:** An AWS OLA can help reduce the number of Windows Server cores requiring a license by 77 percent and the number of SQL Server cores requiring a license by 45 percent when modeling workloads from on premises to AWS.

**Better performance:** Achieve two times better performance when you run your SQL Server workloads on an Amazon EC2 R5b.8xlarge instance.

**Higher reliability:** By moving your Windows workloads to AWS, you can realize a 98 percent reduction in unplanned downtime.

# Can you bring your Microsoft licenses to AWS?

The answer is a resounding YES! There are three ways to do this:

## Option 1: Bring your own license (BYOL) to AWS

You can often deploy eligible Microsoft applications on AWS by bringing your own licenses. AWS is a Microsoft Authorized License Mobility Partner, so if you have licenses with License Mobility—such as SQL Server with active Software Assurance—you can bring them to Amazon EC2 with default (shared) tenancy. Choose BYOL if you want to:

- Take advantage of cloud efficiencies and leverage your existing licensing investments
- Extend the lifecycle of your software without additional hardware costs
- Expedite your migration to the cloud by using existing virtual machine images

## Option 2: Get your Microsoft licenses from AWS

You can use AWS instances that include the cost of licensing. By using AWS-provided licenses on Amazon EC2 or Amazon RDS instances, your Windows Server, SQL Server, Office, and Visual Studio licenses will always be fully compliant. You have the option to purchase **Amazon Machine Images** (AMIs) from AWS or through the **AWS Marketplace**. Choose AWS-provided licenses if you want to:

- Pay as you go—with no upfront costs or long-term commitments
- Fully control when you are being billed for license-included instances
- Get out of the business of managing licensing compliance and leave it up to AWS—we will handle it!

## Option 3: Replace your Microsoft-licensed products and applications

You can speed up innovation with specialized built-for-the-cloud alternatives by replacing your Microsoft-licensed products and applications. Your modernization journey could include moving to Linux, porting applications from .NET Framework to .NET, decomposing monoliths into microservices, implementing DevOps techniques with container and serverless technologies, and transitioning your data tier to Amazon Aurora and other purpose-built databases. Choose this option if you want to:

- Move off legacy licensing for ultimate freedom and savings
- Break free from the punitive terms and high costs of Microsoft licenses

METHOD #4

# Move your workloads to AWS Graviton for the best price performance for a broad set of applications

**AWS Graviton processors** are designed by AWS to deliver the best performance at the lowest cost for your cloud workloads running on Amazon EC2. AWS Graviton–based instances deliver up to 40 percent better performance at lower costs versus comparable x86-based EC2 instances. They are also highly energy efficient, using up to 60 percent less energy for the same performance than comparable Amazon EC2 instances.

AWS Graviton is available in more than 25 AWS Regions, and migrating to AWS Graviton can help you increase performance, reduce costs, lower latency, and achieve better scalability.

The **AWS Graviton Fast Start** program helps you quickly move your workloads to AWS Graviton in less than four hours. Or, accelerate your adoption of AWS Graviton with the help of **AWS Graviton Partners**.

## AWS Graviton–powered managed services

AWS Graviton–based instances are also available in more than 25 popular managed AWS services. These services deliver the price performance benefits of AWS Graviton processors while providing a fully managed experience. AWS Managed Services using AWS Graviton include serverless solutions, such as **AWS Lambda** and **AWS Fargate**, and AWS Graviton–based databases, such as Amazon Aurora, Amazon RDS, and **Amazon ElastiCache**.

METHOD #5

# Lower costs with AWS machine learning accelerators and services

**1T+** **Tens of thousands of customers are using Amazon SageMaker to generate more than 1 trillion predictions per month.**

As AWS continued to drive the best price performance with AWS Graviton, we also wanted to bring price performance benefits to companies running ML workloads on AWS. These ML models are becoming increasingly large and complex, growing exponentially in size. Today, models have grown to hundreds of billions of parameters. These larger and more complex models result in increasing the cost and time for training and running large models.

**AWS Inferentia** accelerators are designed by AWS to deliver high performance at the lowest cost in Amazon EC2 for ML inference applications. The first-generation AWS Inferentia accelerator powers Amazon EC2 Inf1 instances, which deliver up to 2.3 times higher throughput and up to 70 percent lower cost per inference than comparable GPU-based instances. AWS Inferentia2 accelerator is purpose-built to deploy ML models such as large language models (LLMs) and diffusion models with more than 100 billion parameters. Inferentia2-based Amazon EC2 Inf2 instances deliver up to four times higher throughput and up to 10 times lower latency compared to Inf1 instances. They also deliver up to 40 percent better price performance than comparable GPU-based instances in Amazon EC2.

**AWS Trainium** is an ML accelerator that AWS purpose-built for high-performance, low-cost deep learning training. Amazon EC2 Trn1 instances, powered by AWS Trainium chips, offer up to 50 percent cost-to-train savings over comparable GPU-based EC2 instances. Trn1 instances provide up to 1600 Gbps (Trn1n) of Elastic Fabric Adapter (EFA) network bandwidth. They are deployed in EC2 UltraClusters that enable scaling up to 30,000 AWS Trainium accelerators, which are interconnected with a non-blocking petabit-scale network to provide up to 6.3 exaflops of compute.

**Amazon Sagemaker** is a fully managed service that offers high-performance and low-cost ML at scale. Amazon SageMaker enables more people to build, train, and deploy ML models with a choice of tools and workflows—integrated development environments (IDEs) for data scientists and engineers, no-code interface for business analysts, and support for the leading ML frameworks, toolkits, and programming languages for MLOps engineers.

To make it easier to get started, **Amazon SageMaker JumpStart** is an ML hub that provides built-in algorithms with pretrained models from model hubs, pretrained foundation models to help you perform tasks such as article summarization and image generation, and pre-built solutions to solve common use cases. In addition, you can share ML artifacts, including ML models and notebooks, within your organization to accelerate ML model building and deployment.

For customers who prefer complete control over their models, self-managed ML infrastructure offers **AWS Deep Learning AMIs** (DLAMIs) and **AWS Deep Learning Containers** (DLCs) preinstalled with optimizations for popular ML frameworks and model libraries, such as PyTorch, TensorFlow, and Hugging Face.

Visit the **AWS Machine Learning Infrastructure** webpage to learn more about how you can choose your ML infrastructure, tools, and frameworks to accelerate ML innovation.

METHOD #6

# Optimize costs and accelerate innovation with serverless computing

Serverless computing allows you to build and run applications and services without thinking about servers—so you can focus on building applications instead of managing the underlying infrastructure. It eliminates infrastructure management tasks such as server or cluster provisioning, patching, operating-system maintenance, and capacity provisioning.

Serverless computing on AWS is another way to significantly reduce your overall infrastructure costs. With our pay-for-value billing model, resource utilization is automatically optimized, and you never pay for over-provisioning. And with faster-than-ever deployments and updates, you can deliver better-than-ever user experiences.

**AWS Lambda** is an event-driven compute service that lets you run code for virtually any type of application or backend service without provisioning or managing servers. You can trigger AWS Lambda from more than 200 AWS services and software-as-a-service (SaaS) applications and only pay for what you use, resulting in 34 percent better price performance for your application compared to x86-based EC2 instances.

**AWS Fargate** is a pay-as-you-go compute engine that lets you focus on building applications without managing servers. AWS Fargate is compatible with both **Amazon Elastic Container Service** (Amazon ECS) and **Amazon Elastic Kubernetes Service** (Amazon EKS).

If serverless computing is not right for your workload, consider Amazon EC2 for secure, reliable, and cost optimized compute capacity.

## Serverless computing on AWS enables you to:

- Build and iterate quickly
- Release features and updates fast—in hours instead of days—to keep users engaged
- Scale automatically to accommodate changes in demand
- Automate security and compliance processes
- Lower TCO by requiring fewer resources
- Minimize unplanned downtime and security risks

METHOD #7

# Select the compute purchase model that best fits your budget

**$15B+** Since 2019, customers have saved $15 billion with Savings Plans

**$8B+** Since 2015, customers have saved $8 billion with Amazon EC2 Spot Instances

AWS offers you a choice of flexible, cost-effective purchase models to meet your infrastructure needs while keeping you within your budget:

## On-Demand Instances

On-Demand Instances let you pay for compute capacity by the hour or second, depending on which instances you run. No long-term commitments or upfront payments are needed. On-Demand Instances are ideal for applications with short-term, spiky, or unpredictable workloads that cannot be interrupted and applications being developed or tested on Amazon EC2 for the first time.

## Savings Plans

Savings Plans consists of flexible pricing models that can help you reduce your bill by up to 72 percent compared to On-Demand prices in exchange for a one- or three-year hourly spend commitment. AWS offers three types of plans: Compute Savings Plans, EC2 Instance Savings Plans, and Amazon SageMaker Savings Plans.

Compute Savings Plans apply to usage across Amazon EC2, AWS Lambda, and AWS Fargate. The EC2 Instance Savings Plans apply to EC2 usage, and the Amazon SageMaker Savings Plans apply to Amazon SageMaker usage. Savings Plans automatically and simultaneously apply to eligible AWS usage. They enable you to innovate faster by leveraging the newest instances, families, generations, and regions while staying on the same plan. Since the launch of Savings Plans in 2019, customers have saved more than $15 billion.

## Amazon EC2 Spot Instances

Amazon EC2 Spot Instances let you take advantage of unused EC2 capacity in the AWS Cloud. Spot Instances are available at up to a 90 percent discount compared to On-Demand prices. You can use Spot Instances for various stateless, fault-tolerant, or flexible applications, such as big data, containerized workloads, continuous integration and continuous delivery (CI/CD), web servers, HPC, and test and development workloads.

You can combine Spot Instances with other purchase models, giving you the flexibility to grow and change over time while still getting the lowest cost available on AWS. Since 2015, Spot Instances have helped our customers save more than $8 billion.

## CloudFront Security Savings Bundle

The **CloudFront Security Savings Bundle** is a flexible self-service pricing plan that helps you save up to 30 percent on your CloudFront bill in exchange for a monthly spend commitment of a one-year term. These savings are not limited to data delivered by CloudFront but apply to CloudFront usage types, including CloudFront Functions and Lambda@Edge. The CloudFront Security Savings Bundle also includes free AWS Web Application Firewall (WAF) usage up to 10 percent of your committed amount.

METHOD #8

# Optimize costs with a choice of storage services

You can minimize your TCO with AWS Storage services. These services eliminate on-premises capital equipment investment, management complexity, and infrastructure maintenance. With AWS Storage, you get the right mix of price and performance for your workloads while paying only for the storage that you use.

AWS gives you the broadest array of storage services of any cloud provider. This includes more ways to optimize storage costs through a choice of storage classes and intelligently tier data to lower cost storage and data reduction capabilities, such as compression and data deduplication.

For block-based workloads, **Amazon EBS** provides easy-to-use, high-performance block storage at any scale. You select the storage that best fits your workload, service level, and budget. EBS scales fast for your most demanding, high-performance workloads, including Microsoft, Oracle, and SAP products. **Amazon EBS Snapshots** provide a solution to maintain compliance and further reduce snapshot storage costs by up to 75 percent by using **Amazon Data Lifecycle Manager** to automatically move seldom-used snapshots to the **Amazon EBS Snapshots Archive**.

**Amazon Simple Storage Service** (Amazon S3) is the lowest-cost object storage in the cloud. No matter the size of your organization, Amazon S3 lets you store and protect any amount of data for virtually any use case, such as data lakes, cloud-native applications, and mobile apps. Amazon S3 storage classes provide the lowest-cost storage for specific access patterns. **Amazon S3 Intelligent-Tiering** has saved customers more than $1 billion in storage costs compared to Amazon S3 Standard since the launch of Amazon S3 Intelligent-Tiering in 2018. With cost-effective storage classes and easy-to-use management features, you can optimize costs, organize data, and configure fine-tuned access controls to meet specific business, organizational, and compliance requirements.

AWS offers the industry's widest portfolio of fully managed file storage. We handle the infrastructure—provisioning, patching, and backups—allowing you to choose the right file system technology to meet your workload requirements.

**Amazon Elastic File System** (Amazon EFS) is a serverless elastic file system that scales automatically as files are added, removed, and burst to higher throughput levels when necessary. You can also reduce costs by up to 92 percent by automatically tiering infrequently accessed files. The Amazon FSx file system family includes **Amazon FSx for NetApp ONTAP**, **Amazon FSx for OpenZFS**, **Amazon FSx for Windows File Server**, and **Amazon FSx for Lustre**. Amazon FSx enables you to optimize your price and performance to support a broad spectrum of use cases, from small user shares to the most demanding compute-intensive workloads. Amazon FSx file systems support a rich set of storage efficiency features, including data deduplication, compression, and usage quotas.

**Amazon S3 Glacier storage classes** are purpose-built for data archiving, providing you with the highest performance, most retrieval flexibility, and lowest-cost archive storage in the cloud.

**75%**

**Amazon EBS Snapshots Archive reduces snapshots storage costs by up to 75 percent**

**92%**

**Amazon EFS reduces storage costs by 92 percent**

**$1B**

**Since its launch in 2018, Amazon S3 Intelligent-Tiering has saved customers more than $1 billion in storage costs compared to Amazon S3 Standard**

**50%+**

**Amazon FSx family data reduction technology reduces storage costs by 50–65 percent**

METHOD #9

# Optimize your compute resources with intelligent AWS tools

**AWS Compute Optimizer** and **AWS Auto Scaling** allow you to provision with precision. These two services help lower your costs by responding in real time to changes in demand—and they're free to use:

## AWS Compute Optimizer

Over-provisioning resources can lead to unnecessary infrastructure costs, while under-provisioning can lead to poor application performance. Using ML to analyze historical utilization metrics, AWS Compute Optimizer recommends the optimal AWS resources for your workloads, further reducing your infrastructure costs. With just a few clicks, it automatically generates recommendations based on current utilization data, eliminating the need to invest time and resources to set up rule-based thresholds. Since its launch in December 2019, AWS Compute Optimizer has generated recommendations for more than 80 percent of Amazon EC2 usage and provided more than 10 billion recommendations, resulting in reduced costs and improved performance for a variety of workloads.

## AWS Auto Scaling

AWS Auto Scaling monitors your applications and automatically adjusts capacity to maintain steady, predictable performance at the lowest possible cost. You can simply set up application scaling for multiple resources across multiple services with a single intuitive interface and maintain optimal application performance and availability even when workloads are periodic, unpredictable, or continuously changing. When demand spikes, AWS Auto Scaling automatically increases the capacity of constrained resources, so you can maintain a high quality of service. When demand drops, it removes any excess resource capacity to help keep you from overspending. There's no need to add AWS Auto Scaling as a separate tool—it is already built into AWS solutions.

**10B+**

**AWS Compute Optimizer has provided more than 10 billion recommendations since launch, resulting in reduced costs and improved performance for a variety of workloads.**

METHOD #10

# Optimizing costs and productivity with hybrid cloud and edge services

With AWS hybrid cloud and edge services, you can bring AWS Cloud infrastructure and services where you need them to help your digital transformation projects get up and running faster. Securely run new and existing applications using the same proven cloud infrastructure to deliver the best possible performance and end-user experience anywhere. A consistent architecture at the edge and in the cloud enables you to extend familiar AWS services, tools, and APIs wherever you need them. And it lets you run applications that must remain on premises or at the edge due to local data processing or data residency needs.

**AWS Outposts** and VMware Cloud on AWS are cloud-native, fully managed solutions that speed application deployment to the edge, simplify application and network management, and modernize on-premises IT infrastructure and operations.

AWS Outposts is a hybrid solution that brings the same AWS infrastructure, AWS services, APIs, management tools, support, and operating models to virtually any data center, collocation space, or on-premises facility.

**VMware Cloud on AWS** enables organizations to optimize the costs of operating a consistent and seamless hybrid IT environment. Running exclusively on Amazon EC2 instances, it bypasses the challenges of moving and modernizing critical workloads to the cloud, thus enabling business agility without risk, app refactoring, or up-leveling team skills. EC2 instances for VMware Cloud on AWS deliver high networking throughput and lower latency.

You do not need to modify applications to shift to a hybrid cloud model or install custom hardware to use VMware Cloud on AWS.

**AWS Snow Family** Move petabytes of data to and from AWS, or process data at the edge

**AWS Local Zones** Run latency-sensitive applications closer to end users

**AWS Wavelength** Deliver ultra-low-latency applications for 5G devices

**AWS edge networking services** Securely deliver data processing, analysis, and storage close to your endpoints

# Customer examples

Millions of customers, including government organizations, large global enterprise, and startups, trust the capabilities, reliability, and security of AWS to run their most important applications, including mission-critical, enterprise, and cloud-native applications and the next generation of applications they need to accelerate their digital transformation. Here are a few examples:

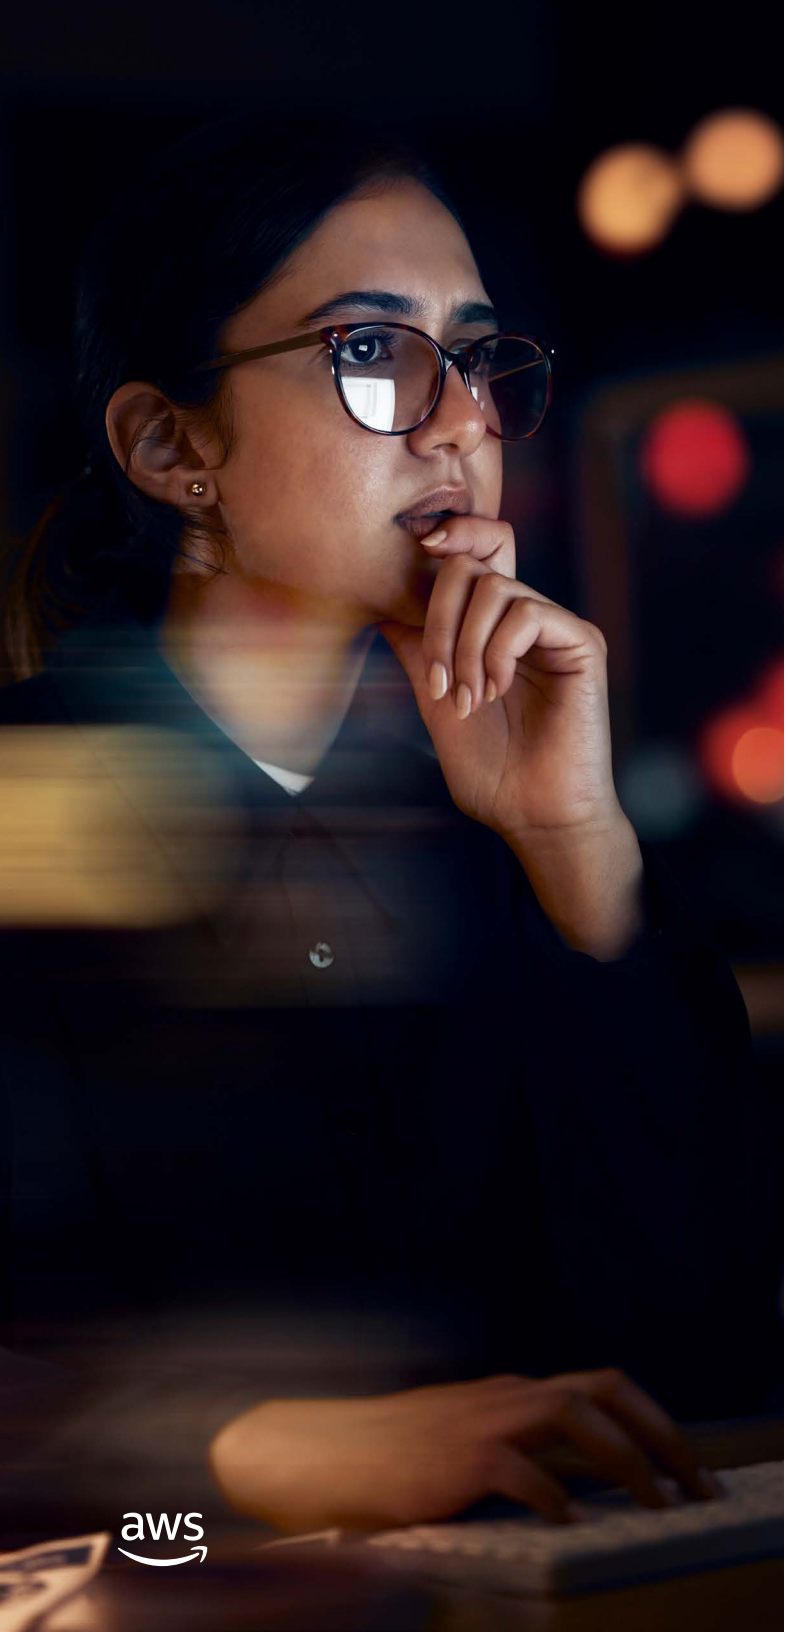

# FINRA: Scaling data processing with Amazon EMR at the speed of market volatility

Financial Industry Regulatory Authority (FINRA) Consolidated Audit Trail (CAT) is the single largest repository for US financial market events and processes billions of events from broker-dealers and exchanges using Apache Spark on Amazon EMR. Processing over 250 billion events from broker-dealers and over 1.7 trillion events from exchanges within a mandated SLA of 4 hours presents unique scaling challenges. Moreover, data volumes change each day based on market activity. In this session, the team at FINRA talks about how they improved SLAs and reduced the cost of running the system using the latest from Amazon EMR and Amazon EC2.

> **"Have a regular cadence for upgrading to newer software versions. Simply upgrading to EMR 6.5, Spark 3.1 and new instance types, we were able to get 30% performance improvement."**
>
> Sambath Velusamy, Director, Technology, FINRA

**Watch the video ›**

# Salesforce: Slashes processing times by 90% and saves $1M monthly with AWS

Using AWS for a mix of instance-provisioning models from Amazon EC2, the Salesforce team was able to build a scalable elastic compute infrastructure. With its remodeled infrastructure, it takes the company less time to process twice as much data while lowering compute costs by more than 60 percent, saving the company more than $1 million a month.

**“We use the capacity of the cloud and the wide range of Amazon EC2 instance types to do things we couldn’t do on premises. Amazon EMR Managed Scaling plays a big part in our ability to use the elastic capability of the cloud. And we significantly reduce costs just by using Spot Instances in an innovative way.”**

Eric Legault, Principal Engineer, Salesforce

**Read the story ›**

# UK Department for Work and Pensions increases workload utilization

In 2022, the Department for Work and Pensions (DWP) in the UK Government embarked on a number of initiatives that used AWS to reduce their carbon footprint. Through carrying out an AWS Well-Architected Sustainability Pillar Review and by using AWS Compute Optimizer, DWP increased some workloads' utilization from 39% to 60% in 12 months—eliminating compute infrastructure waste, increasing performance, and reducing cost.

> **"We need to move to active management to reduce overall emissions from IT activity in order to meet our Greening Government commitments. This will require a reduction of 7-8% year on year to meet the 68% reduction required by 2030. Data centre services will be a critical part of how we do that."**
>
> Tony Sudworth, Sustainability Lead at Department for Work and Pensions (DWP)

**Read the story ›**

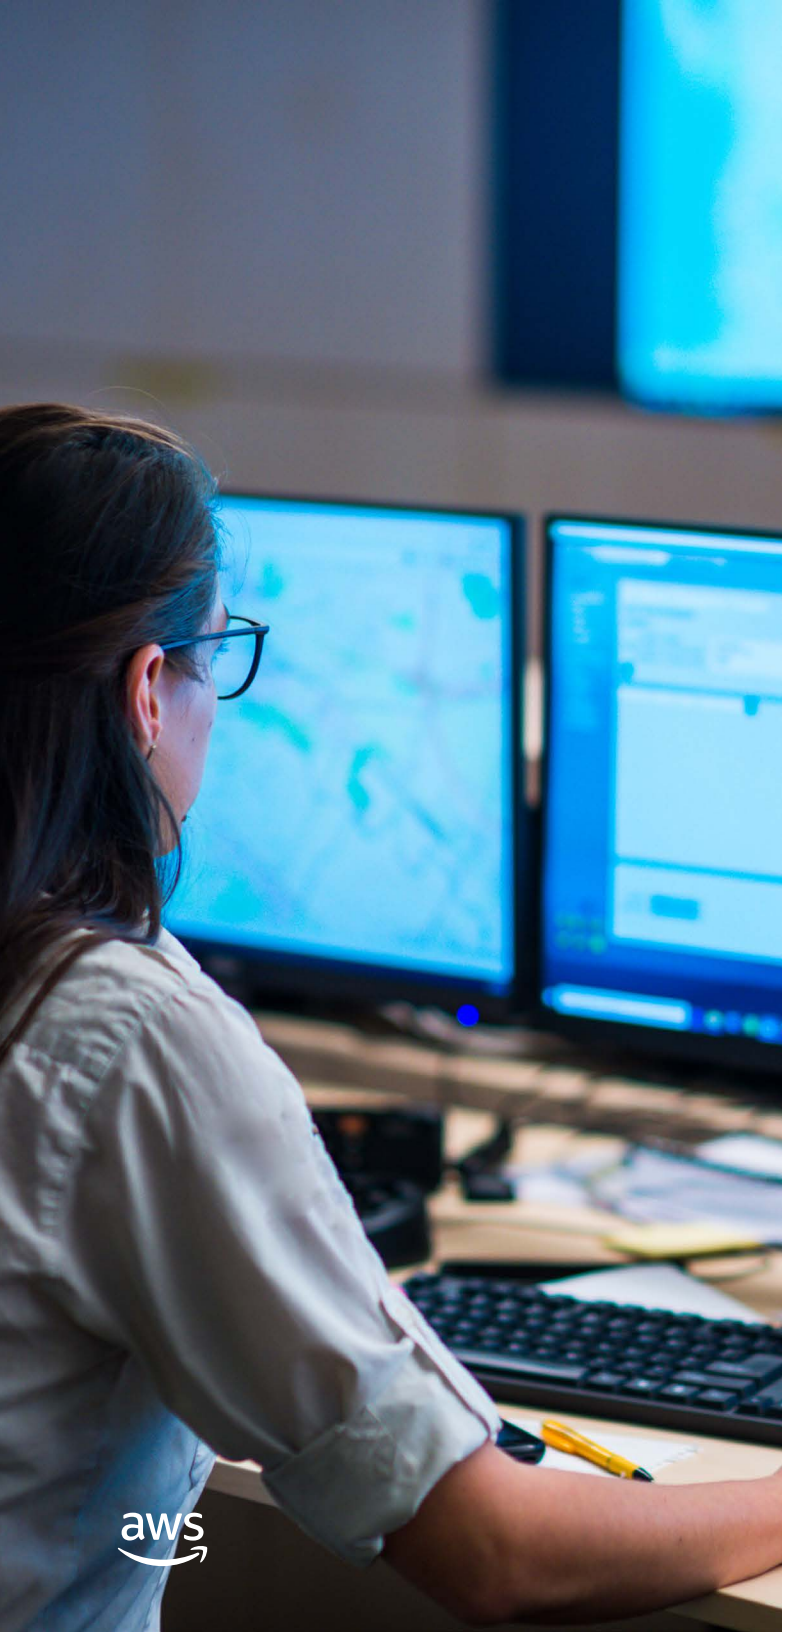

# ENGIE achieves sustainability and increases business value with AWS

ENGIE, a global energy group working to accelerate the transition to a carbon-neutral world, chose AWS to accelerate its transformation to a cloud-first, data-driven company. Using AWS technologies such as **Amazon S3**, **Amazon Redshift**, **AWS Glue**, and **Amazon SageMaker**, ENGIE saves cloud runtime costs while helping cities reduce energy consumption by up to 15 percent. The company migrated millions of dollars worth of workloads across 70 AWS accounts in eight months, avoiding a costly and risky migration project.

> **"We were convinced that AWS was a good solution for many reasons, including the cost model—and especially in terms of data storage."**
>
> Gregory Wolowiec, Technology Team Leader, ENGIE

**Read the story ›**

# Georgia Department of Human Services finds success in multi-year migration journey

The Georgia DHS's new cloud-based platform helps the Georgia DHS provide citizen support in a way it has never done before. The Gateway system serves two million citizens daily, including Medicaid and several other human services programs. The hosting environment enables the speed and agility of application delivery to services that are vital to millions of Georgians. Since migrating these services to the cloud, the Georgia DHS has seen a 50% improvement in batch processing times, and an improvement in time required to scale applications by two times; and an improvement in recovery time objective (RTO) by 14 times. Additionally, 100,000 concurrent citizens were able to log in to the application in 15 minutes. These improvements help the Georgia DHS make sure that citizens have constant and seamless access to the services they need.

> **"Our cloud strategy was to think big, start small, and scale fast with a multi-layered approach."**
>
> Sreeji Vijayan, Chief Information Officer, Georgia Department of Human Services

**Read the story ›**

# GovTech Digitizes Physical Government Payout To Gain Cost- and Time-Savings on AWS

Government Technology Agency (GovTech) is a statutory board of the Government of Singapore responsible for the delivery of digital services to the public. In 2021, GovTech launched GovWallet, a digital wallet for Singaporeans to securely receive payouts from various government agencies, including the Ministry of Home Affairs, Ministry of Defence, Health Promotion Board, and Central Provident Fund Board. Citizens can use their digital wallets to shop at over 164,000 merchants that have SGQR with PayNow and more than 43,000 merchants that have NETS QR across Singapore.

GovWallet is fully built and deployed on Amazon Web Services (AWS). Within the first nine months of its launch in November 2021, GovWallet facilitated more than 1 million disbursement transactions and about 900,000 payment transactions totaling more than S$41 million (US$30 million). “Our team’s initial target was to digitize government payout schemes for at least 50,000 recipients across three agencies by the end of 2021. The speed of innovation afforded by AWS has helped us exceed that goal, and we now have eight government payout schemes and more than 1 million users on our platform. This is a major milestone for Singapore’s Smart Nation Initiative,” said Patricia Zhao, deputy director at GovTech.

**“Our team’s initial target was to digitize government payout schemes for at least 50,000 recipients across three agencies by the end of 2021. The speed of innovation afforded by AWS has helped us exceed that goal, and we now have eight government payout schemes and more than 1 million users on our platform. This is a major milestone for Singapore’s Smart Nation Initiative.”**

Patricia Zhao, Deputy Director, Government Technology Agency

**Read the story ›**

# Innovation and Cost Optimization in the Asia Pacific Region

## Ventures Health New Zealand: Supporting innovations in virtual healthcare

Ventures is a nonprofit organization that designs and delivers smart virtual and in-person clinical care technology so healthcare organizations can quickly innovate. Ventures migrated on-premises storage, databases, and other applications to the cloud, lowering IT costs by 90% so they could invest in helping healthcare providers innovate faster and support virtualized care.

## Transport for New South Wales: Leveraging the cloud to ensure Public Transportation Safety

Transport NSW is a federal agency leading and developing safe and efficient public transportation citizens. Transport NSW implemented serverless computing with AWS Lambda, big data analytics, and machine learning to lower operational costs, better connect communities, and improve rider experience with mobile apps that track status of trains, capacity of buses, and more.

## ST Unitas: Lowering cloud costs to improve access to education

Education technology company ST Unitas in South Korea has a mission to “remove the gap between the world’s rich and poor through education.” To handle the vast amounts of video content the company produces, ST Unitas performs machine learning on AWS. ST Unitas cut expenses by 60% with Amazon EC2 reserved instances and reinvested these savings into adopting cloud database, data warehouse, and machine learning solutions that could help more students get more from their education offerings.

**Read the blog and view video testimonials ›**

CONCLUSION

# Start maximizing your savings now

With innovations in silicon and serverless technologies and a variety of flexible pricing options, AWS is committed to helping you get the most from your cloud infrastructure spend. We provide you with the capability to optimize your costs while building modern, scalable applications to meet your needs. Our breadth of services and pricing models demonstrates our commitment to giving you the compute performance and capacity you require at the lowest cost—now and as your needs evolve.

Start improving your application performance today and reach your highest savings potential with AWS.

## Explore AWS for Every Application resources ›

Discover the ways AWS can help you focus on innovation with more than 200 fully featured services, including serverless and ML capabilities.

## Take advantage of AWS Free Tier and get free, hands-on experience with AWS services ›

## Migrate your on-premises workloads to AWS ›