
Overview of Amazon Web Services

Sajee Mathew

November 2014



Contents

Contents	2
Introduction	4
What Is Cloud Computing?	4
Six Advantages of Cloud Computing	4
Trade capital expense for variable expense	4
Benefit from massive economies of scale	5
Stop guessing about capacity	5
Increase speed and agility	5
Stop spending money running and maintaining data centers	5
Go global in minutes	5
Global Infrastructure	6
Security and Compliance	7
Security	7
Compliance	8
Amazon Web Services Cloud Platform	9
Compute	9
Networking	13
Storage and Content Delivery Network	14
Database	18
Analytics	21
Application Services	22
Deployment and Management	24
Administration and Security	26
Mobile Services	29
Enterprise Applications	30
Conclusion	31
Get Started	31

Introduction

In 2006, Amazon Web Services (AWS) began offering IT infrastructure services to businesses in the form of web services—now commonly known as cloud computing. One of the key benefits of cloud computing is the opportunity to replace up-front capital infrastructure expenses with low variable costs that scale with your business. With the cloud, businesses no longer need to plan for and procure servers and other IT infrastructure weeks or months in advance. Instead, they can instantly spin up hundreds or thousands of servers in minutes and deliver results faster.

Today, AWS provides a highly reliable, scalable, low-cost infrastructure platform in the cloud that powers hundreds of thousands of businesses in 190 countries around the world.

This white paper is an introduction to the AWS cloud computing platform. It discusses the advantages of cloud computing and the fundamentals of AWS. It provides an overview of the AWS services that comprise the platform.

What Is Cloud Computing?

Cloud computing is the on-demand delivery of IT resources and applications via the Internet with pay-as-you-go pricing. Whether you run applications that share photos to millions of mobile users or you support the critical operations of your business, the cloud provides rapid access to flexible and low-cost IT resources. With cloud computing, you don't need to make large upfront investments in hardware and spend a lot of time managing that hardware. Instead, you can provision exactly the right type and size of computing resources you need to power your newest bright idea or operate your IT department. With cloud computing, you can access as many resources as you need, almost instantly, and only pay for what you use.

Cloud computing provides a simple way to access servers, storage, databases, and a broad set of application services over the Internet. Cloud computing providers such as AWS own and maintain the network-connected hardware required for these application services, while you provision and use what you need using a web application.

Six Advantages of Cloud Computing

Trade capital expense for variable expense

Building on-premises infrastructure can be slow and expensive. You need to order, pay for, install, and configure expensive hardware, long before you actually need it. With cloud computing, you simply pay for the resources you consume on a variable basis. Instead of investing heavily in data centers and servers, you can take advantage of

cloud computing and only pay when you consume computing resources, and only pay for the amount you consume.

Benefit from massive economies of scale

Because the usage of hundreds of thousands of customers is aggregated in the cloud, cloud computing providers such as Amazon Web Services can achieve high economies of scale. These economies translate into lower pay-as-you-go prices, plus they can help you reduce IT labor costs and give you access to a highly distributed, full-featured platform at a fraction of the cost of traditional infrastructure. AWS's efficiency improvements allow us to continually lower prices, and you can use our multiple pricing models to optimize costs for your variable and stable workloads.

Stop guessing about capacity

Predicting how customers are going to adopt your new application is difficult, and it's hard to get right. Eliminate guessing on your infrastructure capacity needs. When you make a capacity decision prior to deploying an application, you often end up sitting on expensive idle resources or dealing with limited capacity. With cloud computing, these problems go away. You can access as much or as little as you need, and scale up and down as required with only a few minutes notice.

Increase speed and agility

With traditional infrastructure, it can take weeks to get a server procured, delivered, and running. In a cloud computing environment, new IT resources are only a click away, which means that you can reduce the time it takes to make those resources available to your developers from weeks to just minutes. You can deploy hundreds or even thousands of servers in minutes. This self-service environment changes the speed at which you can develop and deploy applications and allows your team to experiment more quickly and more frequently. This results in a dramatic increase in agility for the organization, since the cost and time it takes to experiment and develop is significantly lower.

Stop spending money running and maintaining data centers

Cloud computing lets you shift resources away from data center investments and operations and move them to innovative new projects. Specifically, you can have your scarce IT and engineering resources focus on projects that grow your business instead of having them focus on IT infrastructure—which is important but hardly ever differentiates your business. Cloud computing lets you focus on your own customers, rather than on the heavy lifting of racking, stacking, and powering servers.

Go global in minutes

Whether you are a large global company or small start-up, you might have potential customers around the world. With traditional infrastructure, it's hard to deliver great

performance to a broadly distributed user base. Most companies focus on one geographic region at a time to save costs and time. With cloud computing, you can easily deploy your application in multiple regions around the world with just a few clicks. This means you can provide lower latency and better experience for more customers at minimal cost.

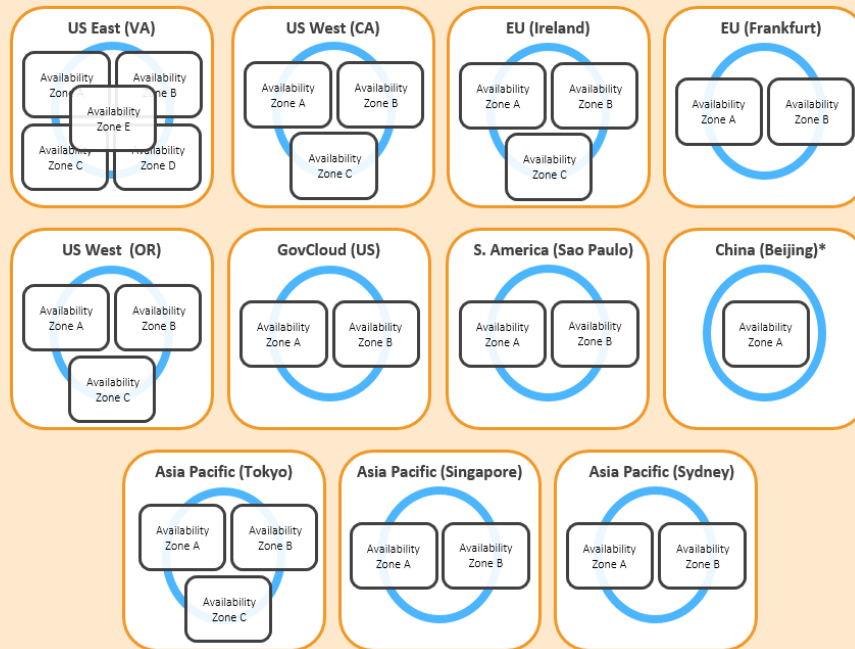
Global Infrastructure

Amazon Web Services serves hundreds of thousands of customers in more than 190 countries. We are steadily expanding our [global infrastructure](#) to help our customers achieve lower latency and higher throughput, and to ensure that your data resides only in the region you specify. As our customers grow their businesses, AWS will continue to provide infrastructure that meets your global requirements.

AWS is available in multiple locations worldwide. These locations are composed of regions and Availability Zones. A region is a named set of AWS resources in the same separate geographic area. Each region has multiple, isolated locations known as Availability Zones. AWS enables the placement of resources, such as instances, and data in multiple locations. Resources aren't replicated across regions unless you chose to do so.

Each region is completely independent and is designed to be completely isolated from the other regions. This achieves the greatest possible fault tolerance and stability. Each Availability Zone is isolated, but the Availability Zones in a region are connected through low-latency links. Availability Zones are physically separated within a typical metropolitan region and are located in lower risk flood plains (specific flood zone categorization varies by region). In addition to utilizing discrete uninterruptable power supply (UPS) and onsite backup generators, they are each fed via different grids from independent utilities to further reduce single points of failure. Availability Zones are all redundantly connected to multiple tier-1 transit providers.

AWS Regions



*A limited preview of the China (Beijing) Region is available to a select group of China-based and multinational companies with customers in China. These customers are required to create a AWS Account, with a set of credentials that are distinct and separate from other global AWS Accounts.

As of this writing, there are eleven AWS regions: US East (Northern Virginia), US West (Oregon), US West (Northern California), AWS GovCloud (US), EU (Ireland), EU (Frankfurt), Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Sydney), South America (Sao Paulo), and China (Beijing).

Security and Compliance

Security

The AWS cloud infrastructure has been architected to be one of the most flexible and secure cloud computing environments available today. It provides an extremely scalable, highly reliable platform that enables customers to deploy applications and data quickly and securely.

With the AWS cloud, not only are infrastructure headaches removed, but so are many of the security issues that come with them. AWS's world-class, highly secure data centers utilize state-of-the-art electronic surveillance and multi-factor access control systems. Data centers are staffed 24x7 by trained security guards, and access is authorized strictly on a least privileged basis. Environmental systems are designed to minimize the impact of disruptions to operations. And multiple geographic regions and Availability

Zones allow you to remain resilient in the face of most failure modes, including natural disasters or system failures.

The AWS virtual infrastructure has been designed to provide optimum availability while ensuring complete customer privacy and segregation. For a complete list of all the security measures built into the core AWS cloud infrastructure, platforms, and services, please read our [Overview of Security Processes](#) whitepaper.

Compliance

AWS Compliance enables our customers to understand the robust controls in place at AWS to maintain security and data protection. As you build systems on top of the AWS cloud infrastructure, the compliance responsibilities will be shared: AWS Compliance provides assurance related to the underlying infrastructure, and your organization owns the compliance initiatives related to anything placed on the AWS infrastructure. The information provided by AWS Compliance helps you to understand our compliance posture and to assess your organization's compliance with your industry and/or government requirements.

The IT infrastructure that AWS provides to its customers is designed and managed in alignment with best security practices and a variety of IT security standards, including:

- SOC 1/SSAE 16/ISAE 3402 (formerly SAS 70 Type II)
- SOC 2
- SOC 3
- FISMA, DIACAP, and FedRAMP
- PCI DSS Level 1
- ISO 27001
- ISO 9001
- ITAR
- FIPS 140-2

In addition, the flexibility and control that the AWS platform provides allows customers to deploy solutions that meet several industry-specific standards, including:

- HIPAA
- Cloud Security Alliance (CSA)
- Motion Picture Association of America (MPAA)

AWS provides a wide range of information regarding its IT control environment to customers through white papers, reports, certifications, accreditations, and other third-party attestations. More information is available in the [Risk and Compliance whitepaper](#) and the [AWS Security Center](#).

Amazon Web Services Cloud Platform

AWS consists of many cloud services that you can use in combinations tailored to your business or organizational needs. This section introduces the AWS services in the following categories: compute, networking, storage and content delivery, databases, analytics, application services, deployment and management, mobile and applications.

Compute

Amazon EC2

[Amazon Elastic Compute Cloud](#) (Amazon EC2) is a web service that provides resizable compute capacity in the cloud. It is designed to make web-scale computing easier for developers and system administrators.

Amazon EC2's simple web service interface allows you to obtain and configure capacity with minimal friction. It provides you with complete control of your computing resources and lets you run on Amazon's proven computing environment. Amazon EC2 reduces the time required to obtain and boot new server instances (called Amazon EC2 instances) to minutes, allowing you to quickly scale capacity, both up and down, as your computing requirements change. Amazon EC2 changes the economics of computing by allowing you to pay only for capacity that you actually use. Amazon EC2 provides developers and system administrators the tools to build failure resilient applications and isolate themselves from common failure scenarios.

Benefits

Elastic Web-Scale Computing

Amazon EC2 enables you to increase or decrease capacity within minutes, not hours or days. You can commission one, hundreds or even thousands of server instances simultaneously. Of course, because this is all controlled with web service APIs, your application can automatically scale itself up and down depending on its needs.

Completely Controlled

You have complete control of your Amazon EC2 instances. You have root access to each one, and you can interact with them as you would any machine. You can stop your Amazon EC2 instance while retaining the data on your boot partition, and then subsequently restart the same instance using web service APIs. Instances can be rebooted remotely using web service APIs. In addition, you can use the [AWS Management Console](#), a simple, web-based user interface, to access and manage your Amazon EC2 instances.

Flexible Cloud Hosting Services

You can choose among multiple Amazon EC2 instance types, operating systems, and software packages. Amazon EC2 allows you to select a configuration of memory, CPU, instance storage, and the boot partition size that is optimal for your choice of operating system and application. For example, your choice of operating systems includes numerous Linux distributions and Microsoft Windows Server.

Designed for use with other Amazon Web Services

Amazon EC2 works in conjunction with Amazon Simple Storage Service (Amazon S3), Amazon Relational Database Service (Amazon RDS), Amazon DynamoDB, and Amazon Simple Queue Service (Amazon SQS) to provide a complete solution for computing, query processing, and storage across a wide range of applications.

Reliable

Amazon EC2 offers a highly reliable environment where replacement instances can be rapidly and predictably commissioned. The service runs within Amazon's proven network infrastructure and data centers. The Amazon EC2 Service Level Agreement commitment is 99.95% availability for each Amazon EC2 region.

Secure

Amazon EC2 works in conjunction with Amazon Virtual Private Cloud ([Amazon VPC](#)) to provide security and robust networking functionality for your compute resources.

- Your compute instances are located in an Amazon Virtual Private Cloud with an IP address range that you specify. You decide which instances are exposed to the Internet and which remain private.
- Security groups and networks ACLs allow you to control inbound and outbound network access to and from your instances.
- You can connect your existing IT infrastructure to resources in your Amazon VPC using industry-standard encrypted IPsec VPN connections.
- You can provision your Amazon EC2 resources as [Dedicated Instances](#). Dedicated Instances are Amazon EC2 instances that run on hardware dedicated to a single customer for additional isolation.
- For more information on Amazon EC2 security refer to our [Amazon Web Services: Overview of Security Process white paper](#).

Inexpensive

Amazon EC2 passes on to you the financial benefits of Amazon's scale. You pay a very low rate for the compute capacity you actually consume. See Amazon [EC2 Instance Purchasing Options](#) for a more detailed description.

- On-Demand Instances—On-Demand instances let you pay for compute capacity by the hour with no long-term commitments. This frees you from the costs and complexities of planning, purchasing, and maintaining hardware and transforms what are commonly large fixed costs into much smaller variable costs. On-Demand Instances also remove the need to buy “safety net” capacity to handle periodic traffic spikes.
- Reserved Instances—[Reserved Instances](#) give you the option to make a low, one-time payment for each instance you want to reserve and in turn receive a significant discount on the hourly charge for that instance. There are three Reserved Instance types (Light, Medium, and Heavy Utilization Reserved Instances) that enable you to balance the amount you pay upfront with your effective hourly price. The [Reserved Instance Marketplace](#) is also available, which provides you with the opportunity to sell Reserved Instances if your needs change. For example, you might want to move instances to a new AWS region, change to a new instance type, or sell capacity for projects that end before your Reserved Instance term expires.
- Spot Instances—[Spot Instances](#) allow you to bid on unused Amazon EC2 capacity and run those instances for as long as your bid exceeds the current Spot Price. The Spot Price changes periodically based on supply and demand, and customers whose bids meet or exceed it gain access to the available Spot Instances. If you can be flexible about when your applications need to run, Spot Instances can significantly lower your Amazon EC2 costs.

Auto Scaling

[Auto Scaling](#) allows you to scale your Amazon EC2 capacity up or down automatically according to conditions you define. With the Auto Scaling service, you can ensure that the number of Amazon EC2 instances you're using increases seamlessly during demand spikes to maintain performance, and decreases automatically during demand lulls to minimize costs. Auto Scaling is particularly well suited for applications that experience hourly, daily, or weekly variability in usage.

Elastic Load Balancing

[Elastic Load Balancing](#) (ELB) automatically distributes incoming application traffic across multiple Amazon EC2 instances. It enables you to achieve even greater fault tolerance in your applications, seamlessly providing the amount of load balancing capacity needed in response to incoming application traffic. Elastic Load Balancing detects unhealthy instances and automatically reroutes traffic to healthy instances until the unhealthy instances have been restored. Customers can enable Elastic Load Balancing within a

single [Availability Zone](#) or across multiple zones for even more consistent application performance.

AWS Lambda

[AWS Lambda](#) is a compute service that runs your code in response to events and automatically manages the compute resources for you, making it easy to build applications that respond quickly to new information. AWS Lambda starts running your code within milliseconds of an event such as an image upload, in-app activity, website click, or output from a connected device. You can also use AWS Lambda to create new back-end services where compute resources are automatically triggered based on custom requests. With AWS Lambda you pay only for the requests served and the compute time required to run your code.

Amazon EC2 Container Service

[Amazon EC2 Container Service](#) is a highly scalable, high performance container management service that supports Docker containers and allows you to easily run distributed applications on a managed cluster of Amazon EC2 instances. Amazon EC2 Container Service lets you launch and stop container-enabled applications with simple API calls, allows you to query the state of your cluster from a centralized service, and gives you access to many familiar Amazon EC2 features like security groups, Amazon EBS volumes and IAM roles. You can use Amazon EC2 Container Service to schedule the placement of containers across your cluster based on your resource needs, isolation policies, and availability requirements. Amazon EC2 Container Service eliminates the need for you to operate your own cluster management and configuration management systems or worry about scaling your management infrastructure.

VM Import/Export

[VM Import/Export](#) enables you to easily import virtual machine images from your existing environment to Amazon EC2 instances and export them back to your on-premises environment. This offering allows you to leverage your existing investments in the virtual machines that you have built to meet your IT security, configuration management, and compliance requirements by bringing those virtual machines into Amazon EC2 as ready-to-use instances. You can also export imported instances back to your on-premises virtualization infrastructure, allowing you to deploy workloads across your IT infrastructure.

To import your images, use the Amazon EC2 API tools, or if you use the VMware vSphere virtualization platform, the Amazon EC2 VM Import Connector to target a virtual machine (VM) image in your existing environment. You then specify which Availability Zone and instance type you want to run in Amazon EC2, and VM Import/Export will automatically transfer the image file and create your instance. Once you have imported your VMs, you can take advantage of Amazon's elasticity, scalability and monitoring via offerings like Auto Scaling, Elastic Load Balancing and CloudWatch to support your imported images. Your instance will be up and running in Amazon EC2 in as little time as it takes to upload your image.

You can export previously imported EC2 instances using the Amazon EC2 API tools. You simply specify the target instance, virtual machine file format and a destination Amazon S3 bucket, and VM Import/Export will automatically export the instance to the Amazon S3 bucket. You can then download and launch the exported VM within your on-premises virtualization infrastructure.

Networking

Amazon VPC

[Amazon Virtual Private Cloud](#) (Amazon VPC) lets you provision a logically isolated section of the AWS cloud where you can launch AWS resources in a virtual network that you define. You have complete control over your virtual networking environment, including selection of your own IP address range, creation of subnets, and configuration of route tables and network gateways.

You can easily customize the network configuration for your Amazon VPC. For example, you can create a public-facing subnet for your web servers that has access to the Internet, and place your backend systems such as databases or application servers in a private-facing subnet with no Internet access. You can leverage multiple layers of security (including security groups and network access control lists) to help control access to Amazon EC2 instances in each subnet.

Additionally, you can create a hardware virtual private network (VPN) connection between your corporate data center and your Amazon VPC and leverage the AWS cloud as an extension of your corporate data center.

AWS Direct Connect

[AWS Direct Connect](#) makes it easy to establish a dedicated network connection from your premises to AWS. Using AWS Direct Connect, you can establish private connectivity between AWS and your data center, office, or co-location environment, which in many cases can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than Internet-based connections.

AWS Direct Connect lets you establish a dedicated network connection between your network and one of the AWS Direct Connect locations. Using industry standard 802.1Q virtual LANS (VLANs), this dedicated connection can be partitioned into multiple virtual interfaces. This allows you to use the same connection to access public resources, such as objects stored in Amazon S3, using public IP address space, and private resources, such as Amazon EC2 instances running within an Amazon VPC, using private IP address space, while maintaining network separation between the public and private environments. Logical connections can be reconfigured at any time to meet your changing needs.

Amazon Route 53

[Amazon Route 53](#) is a highly available and scalable Domain Name System (DNS) web service. It is designed to give developers and businesses an extremely reliable and cost-

effective way to route end users to Internet applications by translating human readable names, such as `www.example.com`, into the numeric IP addresses, such as `192.0.2.1`, that computers use to connect to each other.

Amazon Route 53 effectively connects user requests to infrastructure running in AWS—such as Amazon EC2 instances, Elastic Load Balancing load balancers, or Amazon S3 buckets—and can also be used to route users to infrastructure outside of AWS. You can use Amazon Route 53 to configure DNS health checks to route traffic to healthy endpoints or to independently monitor the health of your application and its endpoints. Amazon Route 53 makes it possible for you to manage traffic globally through a variety of routing types, including Latency Based Routing, Geo DNS, and Weighted Round Robin—all of which can be combined with DNS Failover in order to enable a variety of low-latency, fault-tolerant architectures. Amazon Route 53 also offers Domain Name Registration—you can purchase and manage domain names such as `example.com` and Amazon Route 53 will automatically configure DNS settings for your domains.

Storage and Content Delivery Network

Amazon S3

Amazon Simple Storage Service ([Amazon S3](#)) provides developers and IT teams with safe, secure, highly-scalable object storage.

Amazon S3 provides a simple web-services interface that can be used to store and retrieve any amount of data, at any time, from anywhere on the web. Amazon S3 can be used alone or together with Amazon EC2/EBS, Amazon Glacier, and third-party storage repositories and gateways to provide cost-effective object storage for a wide variety of use cases including cloud applications, content distribution, backup and archiving, disaster recovery, and big data analytics.

Amazon S3 stores data as objects within resources called *buckets*. You can store as many objects as you want within a bucket, and write, read, and delete objects in your bucket. Objects can be up to 5 terabytes in size.

You can control access to the bucket (for example, who can create, delete, and retrieve objects in the bucket), view access logs for the bucket and its objects, and choose the AWS region where Amazon S3 will store the bucket and its contents.

Use Cases

Backup

Amazon S3 offers a highly durable, scalable, and secure solution for backing up and archiving your critical data. You can use Amazon S3's versioning capability to provide even further protection for your stored data. You can also define lifecycle rules to archive sets of Amazon S3 objects to Amazon Glacier, an extremely low-cost storage service.

Content Storage and Distribution

Amazon S3 provides highly durable and available storage for a variety of content. It allows you to offload your entire storage infrastructure into the cloud, where you can take advantage of Amazon S3's scalability and pay-as-you-go pricing to handle your growing storage needs. You can distribute your content directly from Amazon S3 or use Amazon S3 as an origin store for delivery of content to your Amazon CloudFront edge locations.

Big Data Analytics

Whether you're storing pharmaceutical or financial data, or multimedia files such as photos and videos, Amazon S3 is the ideal big data object store. AWS offers a comprehensive portfolio of services to help you manage big data by reducing costs, scaling to meet demand, and increasing the speed of innovation.

Static Website Hosting

You can host your entire static website on Amazon S3 for a low-cost, highly available hosting solution that scales automatically to meet traffic demands. With Amazon S3, you can reliably serve your traffic and handle unexpected peaks without worrying about scaling your infrastructure.

Cloud-native Application Data

Amazon S3 provides high performance, highly available storage that makes it easy to scale and maintain cost-effective mobile and Internet-based apps that run fast. With Amazon S3, you can add any amount of content and access it from anywhere, so you can deploy applications faster and reach more customers.

Disaster Recovery

Amazon S3's highly durable, secure, global infrastructure offers a robust disaster recovery solution designed to provide superior data protection. Whether you're looking for disaster recovery in the cloud or from your corporate data center to Amazon S3, AWS has the right solution for you.

Key Features

Security and Access Management

Amazon S3 provides several mechanisms to control and monitor who can access your data as well as how, when, and where they can access it.

Lifecycle Management

Amazon S3 provides a number of capabilities to manage the lifecycle of your data, including automated archival using the lower-cost Amazon Glacier.

Versioning

Amazon S3 allows you to enable versioning so you can preserve, retrieve, and restore every version of every object stored in an Amazon S3 bucket.

Encryption

You can securely upload/download your data to Amazon S3 via SSL-encrypted endpoints. Amazon S3 also provides multiple options for encryption of data at rest, and allows you to manage your own keys or have Amazon S3 manage them for you.

Cost Monitoring and Controls

Amazon S3 has several features for managing and controlling your costs, including bucket tagging to manage cost allocation and integration with Amazon CloudWatch to receive billing alerts.

Choice of AWS Region

Amazon S3 is available globally in multiple AWS regions. You can choose the region where a bucket is stored to optimize for latency, minimize costs, or address regulatory requirements.

Programmatic Access Using the AWS SDKs

Amazon S3 is supported by the AWS SDKs for Java, PHP, .NET, Python, Node.js, Ruby, and the AWS Mobile SDK. The SDK libraries wrap the underlying REST API, simplifying your programming tasks.

Transfer Data to and from Amazon S3 with Ease

AWS supports several methods for uploading and retrieving data in Amazon S3 including the public Internet, [AWS Direct Connect](#), and the [AWS Import/Export](#) service. And [AWS Storage Gateway](#) automatically backs up on-premises data to Amazon S3.

Flexible Storage Options

Amazon S3 is designed for 99.999999999% durability and 99.99% availability of objects over a given year. There is also a low-cost Reduced Redundancy Storage option for less critical data and Amazon Glacier for archiving cold data at the lowest possible cost.

Amazon Glacier

[Amazon Glacier](#) is an extremely low-cost cloud archive storage service that provides secure and durable storage for data archiving and online backup. In order to keep costs low, Amazon Glacier is optimized for data that is infrequently accessed and for which retrieval times of several hours are suitable. With Amazon Glacier, customers can reliably store large or small amounts of data for as little as \$0.01 per gigabyte per month, a significant savings compared to on-premises solutions.

Companies typically over-pay for data archiving. First, they're forced to make an expensive upfront payment for their archiving solution (which does not include the ongoing cost for operational expenses such as power, facilities, staffing, and maintenance). Second, since companies have to guess what their capacity requirements will be, they understandably over-provision to make sure they have enough capacity for data redundancy and unexpected growth. This set of circumstances results in under-utilized capacity and wasted money. With Amazon Glacier, you pay only for what you use. Amazon Glacier changes the game for data archiving and cloud backup as you pay nothing upfront, pay a very low price for storage, and can scale your usage up or down as needed, while AWS handles all of the operational heavy lifting required to do data retention well. It only takes a few clicks in the AWS Management Console to set up Amazon Glacier and then you can upload any amount of data you choose.

Amazon Elastic Block Store

[Amazon Elastic Block Store](#) (Amazon EBS) provides persistent block level storage volumes for use with Amazon EC2 instances in the AWS cloud. Each Amazon EBS volume is automatically replicated within its Availability Zone to protect you from component failure, offering high availability and durability. Amazon EBS volumes offer the consistent and low-latency performance needed to run your workloads. With Amazon EBS, you can scale your usage up or down within minutes—all while paying a low price for only what you provision.

AWS Storage Gateway

[AWS Storage Gateway](#) is a service connecting an on-premises software appliance with cloud-based storage to provide seamless and secure integration between an organization's on-premises IT environment and AWS's storage infrastructure. The service allows you to securely store data in the AWS cloud for scalable and cost-effective storage. The AWS Storage Gateway supports industry-standard storage protocols that work with your existing applications. It provides low-latency performance by maintaining frequently accessed data on-premises while securely storing all of your data encrypted in Amazon S3 or Amazon Glacier.

Amazon CloudFront

[Amazon CloudFront](#) is a content delivery web service. It integrates with other services in AWS to give developers and businesses an easy way to distribute content to end users with low latency, high data transfer speeds, and no minimum usage commitments.

Amazon CloudFront can be used to deliver your entire website, including dynamic, static, streaming, and interactive content using a global network of edge locations. Requests for your content are automatically routed to the nearest edge location, so content is delivered with the best possible performance. Amazon CloudFront is optimized to work with other services in AWS, such as Amazon S3, Amazon EC2, Elastic Load Balancing, and Amazon Route 53. Amazon CloudFront also works seamlessly with any non-AWS origin server that stores the original, definitive versions of your files. Like other services in AWS, there are no long-term contracts or minimum

monthly usage commitments for using Amazon CloudFront—you pay only for as much or as little content as you actually deliver.

AWS Import/Export

[AWS Import/Export](#) accelerates moving large amounts of data into and out of the AWS cloud using portable storage devices for transport. AWS Import/Export transfers your data directly onto and off of storage devices using Amazon's high-speed internal network and bypassing the Internet. For significant data sets, AWS Import/Export is often faster than Internet transfer and more cost effective than upgrading your connectivity.

AWS Import/Export supports data transfer into and out of Amazon S3 buckets in the US East (N. Virginia), US West (Oregon), US West (Northern California), EU (Ireland), and Asia Pacific (Singapore) regions. The service also supports importing data into Amazon EBS snapshots in the US East (N. Virginia), US West (Oregon), and US West (Northern California) regions. In addition, AWS Import/Export supports transferring data into Amazon Glacier in the US East (N. Virginia), US West (Oregon), US West (Northern California), and EU (Ireland).

Database

Amazon RDS

[Amazon Relational Database Service](#) (Amazon RDS) is a web service that makes it easy to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity while managing time-consuming database management tasks, freeing you up to focus on your applications and business.

Amazon RDS gives you access to the capabilities of a familiar [MySQL](#), [Oracle](#), [SQL Server](#), [PostgreSQL](#) or [Amazon Aurora](#) relational database management system. This means that the code, applications, and tools you already use today with your existing databases can be used with Amazon RDS. Amazon RDS automatically patches the database software and backs up your database, storing the backups for a user-defined retention period and enabling point-in-time recovery. You benefit from the flexibility of being able to scale the compute resources or storage capacity associated with your Database Instance (DB Instance) via a single API call.

Amazon RDS DB Instances can be provisioned with either standard storage or Provisioned IOPS storage. Amazon RDS Provisioned IOPS is a storage option designed to deliver fast, predictable, and consistent I/O performance, and is optimized for I/O-intensive, transactional (OLTP) database workloads.

In addition, Amazon RDS makes it easy to use replication to enhance availability and reliability for production workloads. Using the [Multi-AZ](#) deployment option you can run mission critical workloads with high availability and built-in automated fail-over from your primary database to a synchronously replicated secondary database in case of a failure. Amazon RDS for MySQL also enables you to scale out beyond the capacity of a single database deployment for read-heavy database workloads. As with all Amazon Web

Services, there are no up-front investments required, and you pay only for the resources you use.

Amazon Aurora

[Amazon Aurora](#) is a MySQL-compatible, relational database engine that combines the speed and availability of high-end commercial databases with the simplicity and cost-effectiveness of open source databases. Amazon Aurora provides up to five times better performance than MySQL at a price point one tenth that of a commercial database while delivering similar performance and availability. Amazon Aurora joins MySQL, Oracle, Microsoft SQL Server, and PostgreSQL as the fifth database engine available to customers through Amazon RDS. Amazon RDS handles routine database tasks such as provisioning, patching, backup, recovery, failure detection, and repair.

Amazon DynamoDB

[Amazon DynamoDB](#) is a fast and flexible NoSQL database service for all applications that need consistent, single-digit millisecond latency at any scale. It is a fully managed database and supports both document and key-value data models. Its flexible data model and reliable performance make it a great fit for mobile, web, gaming, ad-tech, the Internet of things (IoT), and many other applications.

Benefits

Fast, Consistent Performance

Amazon DynamoDB is designed to deliver consistent, fast performance at any scale for all applications. Average service-side latencies are typically single-digit milliseconds. As your data volumes grow and application performance demands increase, Amazon DynamoDB uses automatic partitioning and solid state drive (SSD) technologies to meet your throughput requirements and deliver low latencies at any scale.

Highly Scalable

When creating a table, simply specify how much request capacity you require. If your throughput requirements change, simply update your table's request capacity using the AWS Management Console or the Amazon DynamoDB API. Amazon DynamoDB manages all the scaling behind the scenes, and you are still able to achieve your prior throughput levels while scaling is underway. Go to the Amazon DynamoDB FAQ to [learn more](#).

Flexible

Amazon DynamoDB supports both document and key-value data structures, giving you the flexibility to design the best architecture that is optimal for your application.

Fine-grained Access Control

Amazon DynamoDB integrates with AWS Identity and Access Management (IAM) for fine-grained access control for users within your organization. You can assign unique security credentials to each user and control each user's access to services and resources.

Fully Managed

Amazon DynamoDB is a fully managed cloud NoSQL database service—you simply create a database table, set your throughput, and let the service handle the rest. You no longer need to worry about database management tasks such as hardware or software provisioning, setup and configuration, software patching, operating a reliable, distributed database cluster, or partitioning data over multiple instances as you scale.

Amazon Redshift

[Amazon Redshift](#) is a fast, fully managed, petabyte-scale data warehouse solution that makes it simple and cost-effective to efficiently analyze all your data using your existing business intelligence tools. You can start small with no commitments or upfront costs and scale to a petabyte or more.

Amazon Redshift delivers fast query performance by using columnar storage technology to improve I/O efficiency and parallelizing queries across multiple nodes. Amazon Redshift uses standard PostgreSQL JDBC and ODBC drivers, allowing you to use a wide range of familiar SQL clients. Data load speed scales linearly with cluster size, with integrations to Amazon S3, Amazon DynamoDB, Amazon Elastic MapReduce, Amazon Kinesis, or any SSH-enabled host.

We've automated most of the common administrative tasks associated with provisioning, configuring and monitoring a data warehouse. Backups to Amazon S3 are continuous, incremental, and automatic. Restores are fast; you can start querying in minutes while your data is spooled down in the background. Enabling disaster recovery across regions takes just a few clicks.

Security is built-in. You can encrypt data at rest and in transit using hardware-accelerated AES-256 and SSL, isolate your clusters using Amazon VPC, and even manage your keys using hardware security modules (HSMs). All API calls, connection attempts, queries, and changes to the cluster are logged and auditable.

Amazon ElastiCache

[Amazon ElastiCache](#) is a web service that makes it easy to deploy, operate, and scale an in-memory cache in the cloud. The service improves the performance of web applications by allowing you to retrieve information from fast, managed, in-memory caches, instead of relying entirely on slower disk-based databases. ElastiCache supports two open-source in-memory caching engines:

- [Memcached](#) : A widely adopted memory object caching system. ElastiCache is protocol compliant with Memcached, so popular tools that you use today with existing Memcached environments will work seamlessly with the service.

-
- [Redis](#) : A popular open-source in-memory key-value store that supports data structures such as sorted sets and lists. ElastiCache supports Redis master / slave replication which can be used to achieve redundancy across Availability Zones.

Amazon ElastiCache automatically detects and replaces failed nodes, reducing the overhead associated with self-managed infrastructures. It provides a resilient system that mitigates the risk of overloaded databases, which slow website and application load times. Through integration with [Amazon CloudWatch](#), Amazon ElastiCache provides enhanced visibility into key performance metrics associated with your Memcached or Redis nodes.

Analytics

Amazon EMR

[Amazon Elastic MapReduce](#) (Amazon EMR) is a web service that makes it easy to quickly and cost-effectively process vast amounts of data.

Amazon EMR uses Hadoop, an open source framework, to distribute your data and processing across a resizable cluster of Amazon EC2 instances. Amazon EMR is used in a variety of applications, including log analysis, web indexing, data warehousing, machine learning, financial analysis, scientific simulation, and bioinformatics. Customers launch millions of Amazon EMR clusters every year.

Amazon Kinesis

[Amazon Kinesis](#) is a fully managed service for real-time processing of streaming data at massive scale. Amazon Kinesis can collect and process hundreds of terabytes of data per hour from hundreds of thousands of sources, allowing you to easily write applications that process information in real-time, from sources such as website click-streams, marketing and financial information, manufacturing instrumentation and social media, and operational logs and metering data.

With Amazon Kinesis applications, you can build real-time dashboards, capture exceptions and generate alerts, drive recommendations, and make other real-time business or operational decisions. You can also easily send data to a variety of other services such as Amazon S3, Amazon DynamoDB, or Amazon Redshift. In a few clicks and a couple of lines of code, you can start building applications which respond to changes in your data stream in seconds, at any scale, while only paying for the resources you use.

AWS Data Pipeline

[AWS Data Pipeline](#) is a web service that helps you reliably process and move data between different AWS compute and storage services as well as on-premises data sources at specified intervals. With AWS Data Pipeline, you can regularly access your data where it's stored, transform and process it at scale, and efficiently transfer the

results to AWS services such as Amazon S3, Amazon RDS, Amazon DynamoDB, and Amazon EMR.

AWS Data Pipeline helps you easily create complex data processing workloads that are fault tolerant, repeatable, and highly available. You don't have to worry about ensuring resource availability, managing inter-task dependencies, retrying transient failures or timeouts in individual tasks, or creating a failure notification system. AWS Data Pipeline also allows you to move and process data that was previously locked up in on-premises data silos.

Application Services

Amazon SQS

[Amazon Simple Queue Service](#) (Amazon SQS) is a fast, reliable, scalable, fully managed message queuing service. Amazon SQS makes it simple and cost-effective to decouple the components of a cloud application. You can use SQS to transmit any volume of data, at any level of throughput, without losing messages or requiring other services to be always available.

With Amazon SQS, you can offload the administrative burden of operating and scaling a highly available messaging cluster, while paying a low price for only what you use.

Amazon SWF

[Amazon Simple Workflow Service](#) (Amazon SWF) Amazon SWF helps developers build, run, and scale background jobs that have parallel or sequential steps. You can think of SWF as a fully-managed state tracker and task coordinator in the cloud. If your application's steps take more than 500 milliseconds to complete, you need to track the state of processing. If you need to recover, or retry if a task fails, Amazon SWF can help you.

Amazon SWF promotes a separation between the control flow of your background job's stepwise logic and the actual units of work that contain your unique business logic. This allows you to separately manage, maintain, and scale the *state machinery* of your application from the core business logic that differentiates it. As your business requirements change, you can easily change application logic without having to worry about the underlying state machinery, task dispatch, and flow control.

Amazon SWF runs within Amazon's high-availability data centers, so the state tracking and task processing engine is available whenever applications need them. Amazon SWF redundantly stores the tasks, reliably dispatches them to application components, tracks their progress, and keeps their latest state.

Amazon SWF replaces the complexity of custom-coded workflow solutions and process automation software with a fully managed cloud workflow web service. This eliminates the need for developers to manage the infrastructure plumbing of process automation and allows them to focus their energy on the unique functionality of their application.

Amazon SWF seamlessly scales with your application's usage. No manual administration of the workflow service is required when you add more cloud workflows to your application or increase the complexity of your workflows.

Amazon SWF lets you write your application components and coordination logic in any programming language and run them in the cloud or on-premises.

Amazon AppStream

[Amazon AppStream](#) is a flexible, low-latency service that lets you stream resource-intensive applications from the cloud. Amazon AppStream deploys and renders your application on AWS infrastructure and streams the output to mass-market devices, such as personal computers, tablets, and mobile phones. Because your application is running in the cloud, it can scale to handle vast computational and storage needs, regardless of the devices your customers are using. You can choose to stream either all or parts of your application from the cloud. Amazon AppStream enables use cases for applications that wouldn't be possible running natively on mass-market devices. Using Amazon AppStream, your applications are no longer constrained by the hardware in your customers' hands.

Amazon AppStream includes a SDK that currently supports streaming applications from Microsoft Windows Server 2008 R2 to devices running FireOS, Android, Chrome, iOS, Mac OS X, and Microsoft Windows.

Amazon SES

[Amazon Simple Email Service](#) (Amazon SES) is a cost-effective outbound-only email-sending service built on the reliable and scalable infrastructure that Amazon.com has developed to serve its own customer base. With Amazon SES, you can send transactional email, marketing messages, or any other type of high-quality content, and you only pay for what you use.

Along with high deliverability, Amazon SES provides easy, real-time access to your sending statistics and built-in notifications for bounces, complaints, and deliveries to help you fine-tune your cloud-based email-sending strategy.

Amazon Elastic Transcoder

[Amazon Elastic Transcoder](#) is media transcoding in the cloud. It is designed to be a highly scalable, easy-to-use, and cost-effective way for developers and businesses to convert (or *transcode*) media files from their source format into versions that will play back on devices like smartphones, tablets, and PCs.

Amazon Elastic Transcoder manages all aspects of the transcoding process for you transparently and automatically. There's no need to administer software, scale hardware, tune performance, or otherwise manage transcoding infrastructure. You simply create a transcoding job by specifying the location of your source video and how you want it transcoded. Amazon Elastic Transcoder also provides transcoding presets

for popular output formats, which means that you don't need to guess about which settings work best on particular devices. All these features are available via service APIs and the AWS Management Console.

Amazon CloudSearch

[Amazon CloudSearch](#) is a managed service in the AWS cloud that makes it easy to set up, manage, and scale a search solution for your website or application. Amazon CloudSearch supports 34 languages and popular search features such as highlighting, autocomplete, and geospatial search.

With Amazon CloudSearch, you can quickly add custom search capabilities to your website or application without having to become a search expert or worry about hardware provisioning, setup, and maintenance. With a few clicks in the AWS Management Console, you can create a search domain and upload the data you want to make searchable, and Amazon CloudSearch automatically provisions the required resources and deploys a highly tuned search index.

As your volume of data and traffic fluctuates, Amazon CloudSearch seamlessly scales to meet your needs. You can easily change your search parameters, fine-tune search relevance, and apply new settings at any time without having to re-upload your data.

Deployment and Management

AWS Elastic Beanstalk

[AWS Elastic Beanstalk](#) is an easy-to-use service for deploying and scaling web applications and services developed with Java, .NET, PHP, Node.js, Python, Ruby, and Docker on familiar servers such as Apache, Nginx, Phusion Passenger, and Internet Information Services (IIS).

You simply upload your code, and then Amazon Elastic Beanstalk automatically handles the deployment, from capacity provisioning, load balancing, and auto scaling to application health monitoring. At the same time, you retain full control over the AWS resources powering your application and can access the underlying resources at any time.

AWS OpsWorks

[AWS OpsWorks](#) is a cloud application management service that makes it easy to deploy and operate applications of all shapes and sizes. You can define the application's architecture and the specification of each component including package installation, software configuration, and resources such as storage. Start from templates for common technologies like application servers and databases, or build your own to perform any task that can be scripted. AWS OpsWorks includes automation to scale your application based on time or load and dynamic configuration to orchestrate changes as your environment scales.

AWS CloudFormation

[AWS CloudFormation](#) gives developers and systems administrators an easy way to create and manage a collection of related AWS resources, provisioning and updating them in an orderly and predictable fashion.

You can use AWS CloudFormation's sample templates or create your own templates to describe the AWS resources, and any associated dependencies or runtime parameters, required to run your application. You don't need to figure out the order for provisioning AWS services or the subtleties of making those dependencies work. AWS CloudFormation takes care of this for you. After the AWS resources are deployed, you can modify and update them in a controlled and predictable way, in effect applying version control to your AWS infrastructure the same way you do with your software.

You can deploy and update a template and its associated collection of resources (called a *stack*) by using the AWS Management Console, AWS Command Line Interface, or APIs. AWS CloudFormation is available at no additional charge, and you pay only for the AWS resources needed to run your applications.

AWS CodeDeploy

[AWS CodeDeploy](#) is a service that automates code deployments to Amazon EC2 instances. AWS CodeDeploy makes it easier for you to rapidly release new features, helps you avoid downtime during deployment, and handles the complexity of updating your applications. You can use AWS CodeDeploy to automate deployments, eliminating the need for error-prone manual operations, and the service scales with your infrastructure so you can easily deploy to one Amazon EC2 instance or thousands.

AWS CodeCommit

[AWS CodeCommit](#) is a secure, highly scalable, managed source control service that hosts private Git repositories. CodeCommit eliminates the need for you to operate your own source control system or worry about scaling its infrastructure. You can use CodeCommit to store anything from code to binaries, and it supports the standard functionality of Git allowing it to work seamlessly with your existing Git-based tools. Your team can also use CodeCommit's online code tools to browse, edit, and collaborate on projects.

AWS CodePipeline

[AWS CodePipeline](#) is a continuous delivery and release automation service that aids smooth deployments. You can design your development workflow for checking in code, building the code, deploying your application into staging, testing it, and releasing it to production. You can integrate 3rd party tools into any step of your release process or you can use CodePipeline as an end-to-end solution. CodePipeline enables you to rapidly deliver features and updates with high quality through the automation of your build, test, and release process.

Administration and Security

AWS Identity and Access Management

[AWS Identity and Access Management](#) (IAM) is a service that you can use to securely control user access to AWS services and resources. Using IAM, you can create and manage AWS users and groups and use permissions to allow and deny their access to AWS resources. IAM allows you to do the following:

- [Manage IAM users](#) and [their access](#)—You can create users in IAM, assign them individual security credentials (i.e., access keys, passwords, and [multi-factor authentication](#) devices) or request temporary security credentials to provide users access to AWS services and resources. You can manage permissions in order to control which operations a user can perform.
- [Manage IAM roles](#) and their [permissions](#)—You can create roles in IAM, and manage permissions to control which operations can be performed by the entity, or AWS service, that assumes the role. You can also define which entity is allowed to assume the role.
- [Manage federated users](#) and their [permissions](#)—You can enable identity federation to allow existing identities (e.g., users) in your enterprise to access the AWS Management Console, to call AWS APIs, and to access resources, without the need to create an IAM user for each identity.

AWS Key Management Service

[AWS Key Management Service](#) (KMS) is a managed service that makes it easy for you to create and control the encryption keys used to encrypt your data, and uses Hardware Security Modules (HSMs) to protect the security of your keys. AWS Key Management Service is integrated with other AWS services including Amazon EBS, Amazon S3, and Amazon Redshift. AWS Key Management Service is also integrated with AWS CloudTrail to provide you with logs of all key usage to help meet your regulatory and compliance needs.

AWS CloudTrail

[AWS CloudTrail](#) is a web service that records AWS API calls for your account and delivers log files to you. The recorded information includes the identity of the API caller, the time of the API call, the source IP address of the API caller, the request parameters, and the response elements returned by the AWS service.

With CloudTrail, you can get a history of AWS API calls for your account, including API calls made via the AWS Management Console, AWS SDKs, command line tools, and higher-level AWS services (such as AWS CloudFormation). The AWS API call history produced by CloudTrail enables security analysis, resource change tracking, and compliance auditing.

Amazon CloudWatch

[Amazon CloudWatch](#) is a monitoring service for AWS cloud resources and the applications you run on AWS. You can use Amazon CloudWatch to collect and track metrics, collect and monitor log files, and set alarms. Amazon CloudWatch can monitor AWS resources such as Amazon EC2 instances, Amazon DynamoDB tables, and Amazon RDS DB Instances, as well as custom metrics generated by your applications and services, and any log files your applications generate. You can use Amazon CloudWatch to gain system-wide visibility into resource utilization, application performance, and operational health. You can use these insights to react and keep your application running smoothly.

AWS Directory Service

[AWS Directory Service](#) is a managed service that allows you to connect your AWS resources with an existing on-premises Microsoft Active Directory or to set up a new, stand-alone directory in the AWS Cloud. Connecting to an on-premises directory is easy and once this connection is established, all users can access AWS resources and applications with their existing corporate credentials. You can also launch managed, Samba-based directories in a matter of minutes, simplifying the deployment and management of Windows workloads in the AWS Cloud.

AWS Config

[AWS Config](#) is a fully managed service that provides you with an AWS resource inventory, configuration history, and configuration change notifications to enable security and governance. With AWS Config you can discover existing AWS resources, export a complete inventory of your AWS resources with all configuration details, and determine how a resource was configured at any point in time. These capabilities enable compliance auditing, security analysis, resource change tracking, and troubleshooting.

Trusted Advisor

[AWS Trusted Advisor](#) acts like your customized cloud expert, and it helps you provision your resources by following best practices. Trusted Advisor inspects your AWS environment and finds opportunities to save money, improve system performance and reliability, or help close security gaps.

All of the AWS Trusted Advisor checks are available to customers with Business-level and Enterprise-level support, and the best-practice recommendations can be integrated into your application through the [AWS Support API](#).

AWS Service Catalog

[AWS Service Catalog](#) is a service that allows administrators to create and manage approved catalogs of resources that end users can then access via a personalized portal. You can control which users have access to which applications or AWS resources to enable compliance with your business policies, while users can easily browse and launch products from the catalogs you create. AWS Service Catalog can

enable your organization to benefit from increased agility and reduced costs as end users can find and launch only the products they need from a catalog that you control.

AWS CloudHSM

The [AWS CloudHSM](#) service helps you meet corporate, contractual, and regulatory compliance requirements for data security by using dedicated Hardware Security Module (HSM) appliances within the AWS cloud.

AWS and AWS Marketplace partners offer a variety of solutions for protecting sensitive data within the AWS platform, but for applications and data subject to rigorous contractual or regulatory requirements for managing cryptographic keys, additional protection is sometimes necessary. Until now, your only option was to store the sensitive data (or the encryption keys protecting the sensitive data) in your on-premises data centers. Unfortunately, this either prevented you from migrating these applications to the cloud or significantly slowed their performance. The AWS CloudHSM service allows you to protect your encryption keys within HSMs designed and validated to government standards for secure key management. You can securely generate, store, and manage the cryptographic keys used for data encryption such that they are accessible only by you. AWS CloudHSM helps you comply with strict key management requirements without sacrificing application performance.

The AWS CloudHSM service works with Amazon VPC. AWS CloudHSMs are provisioned inside your Amazon VPC with an IP address that you specify, providing simple and private network connectivity to your Amazon EC2 instances. Placing AWS CloudHSMs near your Amazon EC2 instances decreases network latency, which can improve application performance. AWS provides dedicated and exclusive access to AWS CloudHSMs, isolated from other AWS customers. Available in multiple regions and Availability Zones, AWS CloudHSM allows you to add secure and durable key storage to your Amazon EC2 applications.

AWS Management Console

You can use the [AWS Management Console](#), a simple and intuitive web-based user interface, to access and manage your AWS services. When you're on the go, you can also use the [AWS Console mobile app](#) to quickly view resources.

The AWS Management Console is a single destination for managing all your AWS resources, from Amazon EC2 instances to DynamoDB tables. Use the console to perform any number of tasks, from deploying new applications to monitoring the health of your application.

The console enables cloud management for all aspects of your AWS account, including accessing your monthly spending by service, managing security credentials, or setting up new IAM users.

The AWS Console mobile app for iOS and Android lets you quickly and easily view your existing resources, including CloudWatch alarms, and perform operational tasks from your mobile device. Our mobile app can now be downloaded from Amazon Appstore, Google Play, or iTunes.

The console supports all [AWS regions](#) and lets you provision resources across multiple regions. Choose a region that best suits your business needs to optimize your customer's experience.

AWS Command Line Interface

The [AWS Command Line Interface](#) (CLI) is a unified tool to manage your AWS services. With just one tool to download and configure, you can control multiple AWS services from the command line and automate them through scripts.

Mobile Services

Amazon SNS

[Amazon Simple Notification Service](#) (Amazon SNS) is a fast, flexible, fully managed push messaging service. Amazon SNS makes it simple and cost-effective to push notifications to Apple, Google, Fire OS, and Windows devices, as well as Android devices in China with Baidu Cloud Push. You can also use Amazon SNS to push notifications to Internet-connected smart devices, as well as other distributed services.

Besides pushing cloud notifications directly to mobile devices, Amazon SNS can also deliver notifications by SMS text message or email to Amazon SQS queues, or to any HTTP endpoint.

To prevent messages from being lost, all messages published to Amazon SNS are stored redundantly across multiple Availability Zones.

Amazon Cognito

[Amazon Cognito](#) is a service that makes it easy to save user data, such as app preferences or game state, in the AWS cloud without writing any backend code or managing any infrastructure. You can save data locally on users' devices allowing your applications to work even when the devices are offline. You can also synchronize data across a user's devices so that their app experience will be consistent regardless of the device they use. With Amazon Cognito, you can focus on creating great app experiences instead of having to worry about building and managing a backend solution to handle identity management, network state, storage, and sync.

Amazon Mobile Analytics

[Amazon Mobile Analytics](#) is a service that lets you easily collect, visualize, and understand app usage data at scale. Many mobile app analytics solutions deliver usage data several hours after the events occur. Amazon Mobile Analytics is designed to deliver usage reports within 60 minutes of receiving data from an app so that you can act on the data more quickly.

Amazon Mobile Analytics is built to scale with your app, allowing you to collect and process billions of events per day from millions of users. It's easy to get started with Amazon Mobile Analytics. You simply add the [AWS Mobile SDK](#) to your app and publish the app using your existing distribution mechanism (such as the iTunes Store, Google Play, or Amazon Appstore), and you can start accessing reports in the [AWS Management Console](#). Amazon Mobile Analytics automatically starts to collect metrics on active users, sessions, and retention, and you can add reporting on in-app revenue or any custom event you choose.

Amazon Mobile Analytics helps you spend more time on creating great apps rather than the undifferentiated heavy lifting of setting up and managing an analytics system. Amazon Mobile Analytics is included in the [AWS Mobile SDK](#), which supports iOS, Android, and Fire OS, or you can use the [Amazon Mobile Analytics REST API](#) directly.

AWS Mobile SDK

The [AWS Mobile SDK](#) helps you build high quality mobile apps quickly and easily. It provides access to AWS Mobile services, mobile-optimized connectors to popular AWS data and storage services, and easy access to a broad array of other AWS services.

The AWS Mobile SDK includes libraries, code samples, and documentation for iOS, Android, and Fire OS so you can build apps that deliver great experiences across devices and platforms.

Enterprise Applications

Amazon WorkSpaces

[Amazon WorkSpaces](#) is a managed desktop computing service in the cloud. Amazon WorkSpaces allows customers to easily provision cloud-based desktops that allow end-users to access the documents, applications and resources they need with the device of their choice, including laptops, iPad, Kindle Fire, Android tablets, and zero clients. With a few clicks in the AWS Management Console, customers can provision a high-quality cloud desktop experience for any number of users at a cost that is highly competitive with traditional desktops and half the cost of most virtual desktop infrastructure (VDI) solutions.

Amazon Zocalo

[Amazon Zocalo](#) is a fully managed, secure enterprise storage and sharing service with strong administrative controls and feedback capabilities that improve user productivity.

Users can comment on files, send them to others for feedback, and upload new versions without having to resort to emailing multiple versions of their files as attachments. Users can take advantage of these capabilities wherever they are, using the device of their choice, including PCs, Macs, and tablets. Amazon Zocalo offers IT administrators the option of integrating with existing corporate directories, flexible sharing policies, audit logs, and control of the location where data is stored.

Conclusion

AWS provides building blocks that you can assemble quickly to support virtually any workload. With AWS you'll find a complete set of highly available services that are designed to work together to build sophisticated scalable applications. You have access to highly durable storage, low-cost compute, high-performance databases, and management tools. All this is available without up-front cost, and you pay for only what you use.

Get Started

To learn how you can reinvent how you work with IT, the [AWS Free Tier](#) allows you to gain hands-on experience with a broad selection of AWS products and services. Within the AWS Free Tier you can test workloads and run applications to learn more and build the right solution for your organization.

By [signing up for AWS](#), you have access to Amazon's cloud computing services. Note: The sign-up process requires a credit card, which will not be charged until you start using services. There are no long-term commitments and you can stop using AWS at any time.

To help familiarize you with AWS, view these [short introductory videos](#) that cover topics like creating an account, launching a virtual server, storing media and more.

Notices

© 2014, Amazon Web Services, Inc. or its affiliates. All rights reserved. This document is provided for informational purposes only. It represents AWS's current product offerings as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS's products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. The responsibilities and liabilities of AWS to its

customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.