

Overview of Amazon Web Services

December 2015



© 2015, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Notices

This document is provided for informational purposes only. It represents AWS's current product offerings and practices 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.

Contents

Contents	3
Introduction	4
What Is Cloud Computing?	4
Six Advantages of Cloud Computing	5
Global Infrastructure	5
Security and Compliance	6
Security	6
Compliance	7
Amazon Web Services Cloud Platform	7
Compute	8
Storage and Content Delivery	12
Database	14
Networking	15
Developer Tools	17
Management Tools	17
Security and Identity	19
Analytics	21
Internet of Things	24
Mobile Services	25
Application Services	26
Enterprise Applications	27
Next Steps	28
Conclusion	29
Contributors	29
Document Revisions	29
Notes	29

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, Amazon Web Services 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

Instead of having to invest heavily in data centers and servers before you know how you're going to use them, you can pay only when you consume computing resources, and pay only for how much you consume.

Benefit from massive economies of scale

By using cloud computing, you can achieve a lower variable cost than you can get on your own. Because usage from hundreds of thousands of customers is aggregated in the cloud, providers such as AWS can achieve higher economies of scale, which translates into lower pay as-you-go prices.

Stop guessing about capacity

Eliminate guessing on your infrastructure capacity needs. When you make a capacity decision prior to deploying an application, you often end up either 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 capacity as you need, and scale up and down as required with only a few minutes' notice.

Increase speed and agility

In a cloud computing environment, new IT resources are only a click away, which means that you reduce the time to make those resources available to your developers from weeks to just minutes. 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

Focus on projects that differentiate your business, not the infrastructure. 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

Easily deploy your application in multiple regions around the world with just a few clicks. This means you can provide lower latency and a better experience for your customers at minimal cost.

Global Infrastructure

AWS serves over a million active customers in more than 190 countries. We are steadily expanding [global infrastructure](#)¹ to help our customers achieve lower

latency and higher throughput, and to ensure that their data resides only in the region they specify. As our customers grow their businesses, AWS will continue to provide infrastructure that meets their 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.

Security and Compliance

Security

[Cloud security](#)² at AWS is job zero. All AWS customers benefit from data center architecture and network architecture built to satisfy the requirements of the most security-sensitive organizations. AWS and its partners offer hundreds of tools and features to help you meet your security objectives for visibility, auditability, controllability, and agility. This means that you can have the security you need, but without the capital outlay, and with much lower operational overhead than in an on-premises environment.

As an AWS customer you inherit all the best practices of AWS policies, architecture, and operational processes built to satisfy the requirements of our most security sensitive customers. Get the flexibility and agility you need in security controls.

AWS provides you with guidance and expertise through online resources, personnel, and partners. AWS provides you with advisories for current issues,

plus you have the opportunity to work with AWS when you encounter security issues.

You get access to hundreds of tools and features to help you to meet your security objectives. AWS provides security-specific tools and features across network security, configuration management, access control, and data encryption.

And finally, AWS environments are continuously audited, with certifications from accreditation bodies across geographies and verticals. In the AWS environment, you can take advantage of automated tools for asset inventory and privileged access reporting.

Compliance

[AWS Cloud Compliance](#)³ enables customers to understand the robust controls in place at AWS to maintain security and data protection in the cloud. As systems are built on top of AWS cloud infrastructure, compliance responsibilities will be shared. By tying together governance-focused, audit-friendly service features with applicable compliance or audit standards, AWS Compliance enablers build on traditional programs. This helps customers to establish and operate in an AWS security control environment.

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. The following is a partial list of the many certifications and standards with which AWS complies:

- SOC 1/ISAE 3402, SOC 2, SOC 3
- FISMA, DIACAP, and FedRAMP
- PCI DSS Level 1
- ISO 9001, ISO 27001, ISO 27018

AWS provides a wide range of information regarding its IT control environment to customers in whitepapers, reports, certifications, accreditations, and other third-party attestations. More information is available in the [Risk and Compliance whitepaper](#)⁴ and the [AWS Security Center](#)⁵.

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 major AWS services by category. To access the services, you can use the AWS Management Console or the Command Line Interface.

AWS Management Console

Access and manage Amazon Web Services through the [AWS Management Console](#)⁶, a simple and intuitive user interface. You can also use the [AWS console mobile app](#)⁷ to quickly view resources on the go.

AWS Command Line Interface

The [AWS Command Line Interface](#)⁸ (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.

Compute

Amazon EC2

[Amazon Elastic Compute Cloud](#)⁹ (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.

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. Instances can be

rebooted remotely using web service APIs. You also have access to console output of your instances.

Flexible Cloud Hosting Services

You can choose among multiple instance types, operating systems, and software packages. Amazon EC2 allows you to select the memory configuration, CPU, instance storage, and boot partition size that are optimal for the operating system and applications that you choose. For example, your choice of operating systems includes numerous Linux distributions and [Microsoft Windows Server¹⁰](#).

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 (SLA) commitment is 99.95% availability for each Amazon EC2 region.

Secure

Amazon EC2 works in conjunction with Amazon Virtual Private Cloud ([Amazon VPC¹¹](#)) to provide security and robust networking functionality for your compute resources.

- Your compute instances are located in a VPC with an IP address range that you specify. You decide which instances are exposed to the Internet and which remain private.
- Security groups and network access control lists (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 VPC using industry-standard encrypted IPsec virtual private network (VPN) connections.

- You can provision your Amazon EC2 resources as [Dedicated Instances](#)¹². Dedicated Instances are Amazon EC2 instances that run on hardware dedicated to a single customer for additional isolation.

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](#)¹³ 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](#)¹⁴ allow you to reserve Amazon EC2 computing capacity for 1 or 3 years, in exchange for a significantly discounted hourly rate (up to 75%), compared to On-Demand Instance pricing. You can use the [Reserved Instance Marketplace](#)¹⁵ 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](#)¹⁶ 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](#)¹⁷ helps you maintain application availability and allows you to scale your Amazon EC2 capacity up or down automatically according to conditions you define. You can use Auto Scaling to help ensure that you are running your desired number of Amazon EC2 instances. Auto Scaling can also automatically increase the number of Amazon EC2 instances during demand spikes to maintain performance and decrease capacity during lulls to reduce costs. Auto Scaling is

well suited both to applications that have stable demand patterns and applications that experience hourly, daily, or weekly variability in usage.

Elastic Load Balancing

[Elastic Load Balancing¹⁸](#) (ELB) automatically distributes incoming application traffic across multiple Amazon EC2 instances in the cloud. It enables you to achieve greater levels of fault tolerance in your applications, seamlessly providing the required amount of load balancing capacity needed to distribute application traffic.

AWS Lambda

[AWS Lambda¹⁹](#) lets you run code without provisioning or managing servers. You pay only for the compute time you consume—there is no charge when your code is not running. With Lambda, you can run code for virtually any type of application or backend service—all with zero administration. Just upload your code, and Lambda takes care of everything required to run and scale your code with high availability. You can set up your code to automatically trigger from other AWS services, or you can call it directly from any web or mobile app

Amazon EC2 Container Service

[Amazon EC2 Container Service²⁰](#) (Amazon ECS) is a highly scalable, high-performance container management service that supports Docker containers. It allows you to easily run applications on a managed cluster of [Amazon EC2²¹](#) instances. Amazon ECS eliminates the need for you to install, operate, and scale your own cluster management infrastructure. With simple API calls, you can launch and stop Docker-enabled applications, query the complete state of your cluster, and access many familiar features like security groups, [Elastic Load Balancing²²](#), Amazon Elastic Block Store (Amazon [EBS²³](#)) volumes, and AWS Identity and Access Management ([IAM²⁴](#)) roles. You can use Amazon ECS to schedule the placement of containers across your cluster based on your resource needs and availability requirements. You can also integrate your own scheduler or third-party schedulers to meet business- or application-specific requirements.

AWS Elastic Beanstalk

[AWS Elastic Beanstalk²⁵](#) is an easy-to-use service for deploying and scaling web applications and services developed with Java, .NET, PHP, Node.js, Python, Ruby, Go, and Docker on familiar servers such as Apache, Nginx, 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.

VM Import/Export

[VM Import/Export](#)²⁶ 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.

Storage and Content Delivery

Amazon S3

Amazon Simple Storage Service ([Amazon S3](#)²⁷) provides developers and IT teams with safe, secure, and highly-scalable object storage. Amazon S3 is easy to use. It has a simple web service interface for storage and retrieval of any amount of data from anywhere on the web.

Amazon S3 offers a range of storage classes designed for different use cases, including Amazon S3 Standard for general-purpose storage of frequently accessed data, Amazon S3 Standard - Infrequent Access (Standard - IA) for long-lived, but less frequently accessed data, and Amazon Glacier for long-term archiving. Amazon S3 also offers configurable lifecycle policies for managing your data throughout its lifecycle. Once a policy is set, your data will automatically migrate to the most appropriate storage class without any changes to your applications.

Amazon Glacier

[Amazon Glacier](#)²⁸ is a secure, durable, and extremely low-cost storage service for data archiving and long-term backup. Customers can reliably store large or small amounts of data for as little as \$0.007 per gigabyte per month, a significant savings compared to on-premises solutions. To keep costs low, Amazon Glacier is optimized for infrequently accessed data where a retrieval time of several hours is suitable.

Amazon Elastic Block Store

[Amazon Elastic Block Store](#)²⁹ (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.

Amazon Elastic File System

[Amazon Elastic File System](#)³⁰ (Amazon EFS) is a shared file storage service for Amazon EC2 instances. Amazon EFS is easy to use and provides a simple interface that allows you to create and configure file systems quickly and easily. With Amazon EFS, storage capacity is elastic, growing and shrinking automatically as you add and remove files, so your applications have the storage they need, when they need it.

AWS Storage Gateway

[AWS Storage Gateway](#)³¹ 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 the AWS 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](#)³² 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 Snowball

[AWS Import/Export Snowball](#)³³ is a petabyte-scale data transport solution that uses secure appliances to transfer large amounts of data into and out of AWS. Using Snowball addresses common challenges with large-scale data transfers

including high network costs, long transfer times, and security concerns. Transferring data with Snowball is simple, fast, secure, and can be as little as one-fifth the cost of high-speed Internet.

Database

Amazon RDS

[Amazon Relational Database Service](#)³⁴ (Amazon RDS) makes it easy to set up, operate, and scale relational databases in the cloud. Amazon RDS frees you up to focus on your application by managing time-consuming database administration tasks including backups, software patching, monitoring, scaling and replication. It's a fully managed cost-efficient service that allows you to launch secure, highly available, fault tolerant production ready databases in minutes. You can scale your database's compute and storage resources with only a few clicks or an API call, often with no downtime. Amazon RDS provides you six familiar database engines to choose from, including [Amazon Aurora](#)³⁵, [Oracle](#)³⁶, [Microsoft SQL Server](#)³⁷, [PostgreSQL](#)³⁸, [MySQL](#)³⁹, and [MariaDB](#)⁴⁰. For commercial database engines like Oracle and SQL Server, you can bring your existing licenses or pay for the license as part of the service with the "License Included" option.

Amazon Aurora

[Amazon Aurora](#)⁴¹ is a MySQL-compatible, relational database engine that combines the reliability, 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 that is one-tenth of a commercial database while delivering similar performance and availability. Amazon Aurora is available as a managed database service through Amazon RDS.

AWS Database Migration Service

[AWS Database Migration Service](#)⁴² helps you migrate databases to AWS easily and securely while the source database remains fully operational, minimizing downtime to applications that rely on the database. It can migrate your data to and from all widely used commercial and open-source databases, and it supports homogenous migrations such as Oracle to Oracle, as well as heterogeneous migrations between different database platforms, such as Oracle to Amazon Aurora.

Amazon DynamoDB

[Amazon DynamoDB](#)⁴³ 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, Internet of Things, and many other applications.

The [downloadable version of DynamoDB⁴⁴](#) makes it easy to develop and test applications on your laptop or in an EC2 instance. Once your solution is ready, easily scale your application on the cloud with DynamoDB.

Amazon Redshift

[Amazon Redshift⁴⁵](#) is a fast, fully managed, petabyte-scale data warehouse service 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 petabytes for as little as \$1000/terabyte/year.

Amazon Redshift delivers fast query performance by using columnar storage technology to improve I/O efficiency and parallelizing queries across multiple nodes. Amazon Redshift has custom JDBC and ODBC drivers that you can download from the AWS Management Console, allowing you to use a wide range of familiar SQL clients. You can also use standard PostgreSQL JDBC, and ODBC drivers. 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.

Amazon ElastiCache

[Amazon ElastiCache⁴⁶](#) 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:

- [Memcached⁴⁷](#) 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.

[Redis⁴⁸](#) A popular open-source in-memory key-value store that supports data structures such as sorted sets and lists. ElastiCache supports master/slave replication and Multi-AZ, which can be used to achieve cross-AZ redundancy.

Networking

Amazon VPC

[Amazon Virtual Private Cloud⁴⁹](#) (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 Connect⁵⁰](#) 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. Virtual interfaces can be reconfigured at any time to meet your changing needs.

Amazon Route 53

[Amazon Route 53⁵¹](#) 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.

Developer Tools

AWS CodeCommit

[AWS CodeCommit⁵²](#) is a fully managed source control service that makes it easy for companies to host secure and highly scalable private Git repositories. AWS CodeCommit eliminates the need to operate your own source control system or worry about scaling its infrastructure. You can use AWS CodeCommit to securely store anything from source code to binaries, and it works seamlessly with your existing Git tools.

AWS CodeDeploy

[AWS CodeDeploy⁵³](#) is a service that automates code deployments to any instance, including Amazon EC2 instances and instances running on-premises. AWS CodeDeploy makes it easier for you to rapidly release new features, helps you avoid downtime during application deployment, and handles the complexity of updating your applications. You can use AWS CodeDeploy to automate software deployments, eliminating the need for error-prone manual operations, and the service scales with your infrastructure so you can easily deploy to one instance or thousands.

AWS CodePipeline

[AWS CodePipeline⁵⁴](#) is a continuous delivery service for fast and reliable application updates. AWS CodePipeline builds, tests, and deploys your code every time there is a code change, based on the release process models you define. This enables you to rapidly and reliably deliver features and updates. You can easily build out an end-to-end solution by using our pre-built plugins for popular third-party services like GitHub or integrating your own custom plugins into any stage of your release process.

Management Tools

Amazon CloudWatch

[Amazon CloudWatch⁵⁵](#) 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 CloudFormation

[AWS CloudFormation⁵⁶](#) 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 the sample templates in AWS CloudFormation 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 also visualize your templates as diagrams and edit them using a drag-and-drop interface with the [AWS CloudFormation Designer⁵⁷](#).

AWS CloudTrail

[AWS CloudTrail⁵⁸](#) 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.

AWS Config

[AWS Config⁵⁹](#) 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.

AWS Config Rules is a new set of cloud governance capabilities that allow IT administrators to define guidelines for provisioning and configuring AWS resources and then continuously monitor compliance with those guidelines. AWS Config Rules lets you choose from a set of pre-built rules based on common AWS best practices or custom rules that you define. For example, you can ensure that Amazon EBS volumes are encrypted, EC2 instances are properly tagged, and Elastic IP addresses (EIPs) are attached to instances. AWS Config Rules can continuously monitor configuration changes to your AWS resources and provides a new dashboard to track compliance status. Using Config Rules, an IT administrator can quickly determine when and how a resource went out of compliance.

AWS OpsWorks

[AWS OpsWorks](#)⁶⁰ is a configuration management service that helps you configure and operate applications of all shapes and sizes using Chef. 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 Service Catalog

[AWS Service Catalog](#)⁶¹ allows organizations to create and manage catalogs of IT services that are approved for use on AWS. These IT services can include everything from virtual machine images, servers, software, and databases to complete multi-tier application architectures. AWS Service Catalog allows you to centrally manage commonly deployed IT services, and helps you achieve consistent governance and meet your compliance requirements, while enabling users to quickly deploy only the approved IT services they need.

AWS Trusted Advisor

[AWS Trusted Advisor](#)⁶² acts like your customized cloud expert, and it helps you provision your resources by following best practices. AWS Trusted Advisor inspects your AWS environment and finds opportunities to save money, improve system performance and reliability, or help close security gaps.

Security and Identity

AWS Identity and Access Management

[AWS Identity and Access Management](#)⁶³ (IAM) enables you to securely control access to AWS services and resources for your users. 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 users⁶⁴](#) and [their access⁶⁵](#)—You can create users in IAM, assign them individual security credentials (access keys, passwords, and [multi-factor authentication⁶⁶](#) 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 roles⁶⁷](#) and their [permissions⁶⁸](#)—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 users⁶⁹](#) and their [permissions⁷⁰](#)—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 Service⁷¹](#) (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 Directory Service

[AWS Directory Service⁷²](#) 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, standalone directory in the AWS cloud. Connecting to an on-premises directory is easy. 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 Linux and Microsoft Windows workloads in the AWS cloud.

Amazon Inspector

[Amazon Inspector⁷³](#) is an automated security assessment service that helps improve the security and compliance of applications deployed on AWS. Amazon Inspector automatically assesses applications for vulnerabilities or deviations from best practices. After performing an assessment, Amazon Inspector produces a detailed report with prioritized steps for remediation. To help you get started

quickly, Amazon Inspector includes a knowledge base of hundreds of rules mapped to common security compliance standards (e.g., PCI DSS) and vulnerability definitions. Examples of built-in rules include checking for remote root login being enabled, or vulnerable software versions installed. These rules are regularly updated by AWS security researchers.

AWS WAF

[AWS WAF⁷⁴](#) is a web application firewall that helps protect your web applications from common web exploits that could affect application availability, compromise security, or consume excessive resources. AWS WAF gives you control over which traffic to allow or block to your web application by defining customizable web security rules. You can use AWS WAF to create custom rules that block common attack patterns, such as SQL injection or cross-site scripting, and rules that are designed for your specific application. New rules can be deployed within minutes, letting you respond quickly to changing traffic patterns. Also, AWS WAF includes a full-featured API that you can use to automate the creation, deployment, and maintenance of web security rules.

AWS CloudHSM

The [AWS CloudHSM⁷⁵](#) 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. 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. 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. 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.

Analytics

Amazon EMR

[Amazon Elastic MapReduce⁷⁶](#) (Amazon EMR) is a web service that makes it easy to quickly and cost-effectively process vast amounts of data.

Amazon EMR simplifies big data processing, providing a managed Apache Hadoop framework that makes it easy, fast, and cost-effective for you to distribute and process vast amounts of your data across dynamically scalable Amazon EC2 instances. You can also run other popular distributed frameworks

such as Apache Spark and Presto in Amazon EMR, and interact with data in other AWS data stores such as Amazon S3 and Amazon DynamoDB.

Amazon EMR securely and reliably handles your big data use cases, including log analysis, web indexing, data warehousing, machine learning, financial analysis, scientific simulation, and bioinformatics.

Amazon QuickSight

[Amazon QuickSight⁷⁷](#) is a fast, cloud-powered business intelligence (BI) service that makes it easy for you to build visualizations, perform ad-hoc analysis, and quickly get business insights from your data. Amazon QuickSight uses SPICE—a Super-fast, Parallel, In-memory Calculation Engine—to perform advanced calculations and render visualizations rapidly. Amazon QuickSight integrates with AWS data sources, like Amazon Kinesis, Amazon Redshift, or Amazon DynamoDB; flat files such as CSV or text; and third-party sources such as Salesforce. Amazon QuickSight enables organizations to scale to hundreds of thousands of users, delivering fast and responsive query performance via SPICE. At one-tenth the cost of traditional solutions, Amazon QuickSight enables you to deliver rich BI functionality to everyone in your organization.

AWS Data Pipeline

[AWS Data Pipeline⁷⁸](#) 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.

Amazon Elasticsearch Service

[Amazon Elasticsearch Service⁷⁹](#) is a managed service that makes it easy to deploy, operate, and scale Elasticsearch in the AWS cloud. Elasticsearch is a popular open-source search and analytics engine for use cases such as log analytics, real-time application monitoring, and click stream analytics. You can set up and configure your Amazon Elasticsearch cluster in minutes from the AWS Management Console. Amazon Elasticsearch Service provisions all the resources for your cluster and launches it. The service automatically detects and replaces

failed Elasticsearch nodes, reducing the overhead associated with self-managed infrastructure and Elasticsearch software. Amazon Elasticsearch Service allows you to easily scale your cluster using a single API call or a few clicks in the AWS Management Console. With Amazon Elasticsearch Service, you get direct access to the Elasticsearch open-source API so that code and applications you're already using with your existing Elasticsearch environments will work seamlessly.

Amazon Kinesis

[Amazon Kinesis⁸⁰](#) is a platform for streaming data on AWS, offering powerful services to make it easy to load and analyze streaming data, and also providing the ability for you to build custom streaming data applications for specialized needs. Web applications, mobile devices, wearables, industrial sensors, and many software applications and services can generate staggering amounts of streaming data—sometimes terabytes per hour—that need to be collected, stored, and processed continuously. Amazon Kinesis services enable you to do that simply and at a low cost. Amazon Kinesis currently offers two services: Amazon Kinesis Firehose and Amazon Kinesis Streams.

Amazon Kinesis Firehose

[Amazon Kinesis Firehose⁸¹](#) is the easiest way to load streaming data into AWS. It can capture and automatically load streaming data into Amazon S3 and Amazon Redshift, enabling near real-time analytics with existing business intelligence tools and dashboards you're already using today. It is a fully managed service that automatically scales to match the throughput of your data and requires no ongoing administration. It can also batch, compress, and encrypt the data before loading it, minimizing the amount of storage used at the destination and increasing security. You can easily create a Firehose delivery stream from the AWS Management Console, configure it with a few clicks, and start sending data to the stream from hundreds of thousands of data sources to be loaded continuously to AWS—all in just a few minutes.

Amazon Kinesis Analytics

[Amazon Kinesis Analytics⁸²](#) is the easiest way to run standard SQL queries against streaming data.

Amazon Kinesis Streams

[Amazon Kinesis Streams⁸³](#) enables you to build custom applications that process or analyze streaming data for specialized needs. Amazon Kinesis Streams can continuously capture and store terabytes of data per hour from hundreds of thousands of sources such as website clickstreams, financial transactions, social media feeds, IT logs, and location-tracking events. With Amazon Kinesis Client Library (KCL), you can build Amazon Kinesis Applications and use streaming data to power real-time dashboards, generate alerts, implement dynamic pricing

and advertising, and more. You can also emit data from Amazon Kinesis Streams to other AWS services such as Amazon Simple Storage Service (Amazon S3), Amazon Redshift, Amazon Elastic Map Reduce (Amazon EMR), and AWS Lambda.

Amazon Machine Learning

[Amazon Machine Learning](#)⁸⁴ (Amazon ML) is a service that makes it easy for developers of all skill levels to use machine learning technology. Amazon Machine Learning provides visualization tools and wizards that guide you through the process of creating machine learning models without having to learn complex ML algorithms and technology. Once your models are ready, Amazon Machine Learning makes it easy to obtain predictions for your application using simple APIs, without having to implement custom prediction generation code or manage any infrastructure.

Amazon Machine Learning is based on the same proven, highly scalable, ML technology used for years by Amazon's internal data scientist community. The service uses powerful algorithms to create ML models by finding patterns in your existing data. Then, Amazon Machine Learning uses these models to process new data and generate predictions for your application.

Amazon Machine Learning is highly scalable and can generate billions of predictions daily, and serve those predictions in real-time and at high throughput. With Amazon Machine Learning, there is no upfront hardware or software investment, and you pay as you go, so you can start small and scale as your application grows.

Internet of Things

AWS IoT

[AWS IoT](#)⁸⁵ is a managed cloud platform that lets connected devices easily and securely interact with cloud applications and other devices. AWS IoT can support billions of devices and trillions of messages, and can process and route those messages to AWS endpoints and to other devices reliably and securely. With AWS IoT, your applications can keep track of and communicate with all your devices, all the time, even when they aren't connected.

AWS IoT makes it easy to use AWS services like AWS Lambda, Amazon Kinesis, Amazon S3, Amazon Machine Learning, and Amazon DynamoDB to build Internet of Things (IoT) applications that gather, process, analyze and act on data generated by connected devices, without having to manage any infrastructure.

Mobile Services

AWS Mobile Hub

[AWS Mobile Hub](#)⁸⁶ is the fastest way to build mobile apps powered by AWS. It lets you easily add and configure features for your apps, including user authentication, data storage, back end logic, push notifications, content delivery, and analytics. After you build your app, AWS Mobile Hub gives you easy access to testing on real devices, as well as analytics dashboards to track usage of your app—all from a single, integrated console.

Amazon Cognito

[Amazon Cognito](#)⁸⁷ is a service that makes it easy to save mobile user data, such as app preferences or game state, in the AWS cloud without writing any back end code or managing any infrastructure. Amazon Cognito offers mobile identity management and data synchronization across devices. 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 back end solution to handle user authentication, network state, storage, and sync.

AWS Device Farm

[AWS Device Farm](#)⁸⁸ is a mobile app-testing service that enables developers to improve the quality of their Android, iOS, and Fire OS apps by quickly and securely testing them on smartphones, tablets, and other devices in the cloud. Developers can upload their own test scripts or use the AWS Device Farm device compatibility test suite. A test report is updated as soon as tests are complete. In addition to a high-level summary, the reports contain detailed logs, screenshots, and performance data. Testing on multiple devices is parallelized, which enables developers to get test results faster.

Amazon Mobile Analytics

With [Amazon Mobile Analytics](#)⁸⁹, you can measure app usage and app revenue. By tracking key trends such as new versus returning users, app revenue, user retention, and custom in-app behavior events, you can make data-driven decisions to increase engagement and monetization for your app. You can view key charts in the Mobile Analytics console and automatically export your app event data to Amazon S3 and Amazon Redshift to run custom analysis.

Amazon SNS

[Amazon Simple Notification Service](#)⁹⁰ (Amazon SNS) is a fast, flexible, fully managed pub-sub messaging service. Use it as a cloud-based mobile app

notification service to send push notifications, email, and SMS messages; or as an enterprise-messaging infrastructure.

AWS Mobile SDK

The [AWS Mobile SDK⁹¹](#) 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.

Application Services

Amazon API Gateway

[Amazon API Gateway⁹²](#) is a fully managed service that makes it easy for developers to create, publish, maintain, monitor, and secure APIs at any scale. With a few clicks in the AWS Management Console, you can create an API that acts as a “front door” for applications to access data, business logic, or functionality from your back-end services, such as workloads running on Amazon EC2, code running on AWS Lambda, or any web application. Amazon API Gateway handles all the tasks involved in accepting and processing up to hundreds of thousands of concurrent API calls, including traffic management, authorization and access control, monitoring, and API version management. Amazon API Gateway has no minimum fees or startup costs. You pay only for the API calls you receive and the amount of data transferred out.

Amazon AppStream

[Amazon AppStream⁹³](#) lets you deliver your Windows applications to any device, reaching more users on more devices, without code modifications. With Amazon AppStream, your application will be deployed and rendered on AWS infrastructure and the output is streamed 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. Amazon AppStream provides an SDK for streaming your application from the cloud. You can integrate your own custom clients, subscriptions, identity, and storage solution with AppStream to build a custom streaming solution that meets the needs of your business.

Amazon CloudSearch

[Amazon CloudSearch⁹⁴](#) 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.

Amazon Elastic Transcoder

[Amazon Elastic Transcoder⁹⁵](#) 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 SES

[Amazon Simple Email Service⁹⁶](#) (Amazon SES) is a cost-effective email service built on the reliable and scalable infrastructure that Amazon.com 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 to your customers. You can also use Amazon SES to receive messages and deliver them to an Amazon S3 bucket, call your custom code via an AWS Lambda function, or publish notifications to Amazon Simple Notification Service (Amazon SNS). With Amazon SES, you have no required minimum commitments—you pay as you go, and you only pay for what you use.

Amazon SQS

[Amazon Simple Queue Service⁹⁷](#) (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 Amazon SQS to transmit any volume of data, at any level of throughput, without losing messages or requiring other services to be always available.

Amazon SWF

[Amazon Simple Workflow⁹⁸](#) (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.

Enterprise Applications

Amazon WorkSpaces

[Amazon WorkSpaces⁹⁹](#) 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 WorkDocs

[Amazon WorkDocs¹⁰⁰](#) 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, tablets, and phones. Amazon WorkDocs 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.

Amazon WorkMail

[Amazon WorkMail¹⁰¹](#) is a secure, managed business email and calendaring service with support for existing desktop and mobile email clients. Amazon WorkMail gives users the ability to seamlessly access their email, contacts, and calendars using Microsoft Outlook, their web browser, or their native iOS and Android email applications. You can integrate Amazon WorkMail with your existing corporate directory and control both the keys that encrypt your data and the location in which your data is stored.

Next Steps

To learn how you can reinvent how you work with IT, the [AWS Free Tier¹⁰²](#) 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 AWS¹⁰³](#), 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 videos¹⁰⁴](#) that cover topics like creating an account, launching a virtual server, storing media and more.

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.

Contributors

The following individuals and organizations contributed to this document:

- Sajee Mathew, AWS Principal Solutions Architect

Document Revisions

December 2015: Refreshed with numerous changes to reflect new services and features released in 2015.

Notes

¹ <https://aws.amazon.com/about-aws/global-infrastructure/>

² <https://aws.amazon.com/security/>

³ <https://aws.amazon.com/compliance/>

⁴ <https://aws.amazon.com/compliance/>

⁵ <http://aws.amazon.com/security/>

⁶ <https://aws.amazon.com/console/>

⁷ <http://aws.amazon.com/console/mobile/>

⁸ <https://aws.amazon.com/cli/>

⁹ <http://aws.amazon.com/ec2/>

¹⁰ <https://aws.amazon.com/windows/>

¹¹ <http://aws.amazon.com/vpc/>

¹² <http://aws.amazon.com/ec2/purchasing-options/dedicated-instances/>

¹³ <http://aws.amazon.com/ec2/purchasing-options/>

¹⁴ <http://aws.amazon.com/ec2/purchasing-options/reserved-instances/>

¹⁵ <http://aws.amazon.com/ec2/purchasing-options/reserved-instances/marketplace/>

¹⁶ <http://aws.amazon.com/ec2/spot/>

¹⁷ <http://aws.amazon.com/autoscaling/>

¹⁸ <http://aws.amazon.com/elasticloadbalancing/>

19 <http://aws.amazon.com/lambda/>
20 <http://aws.amazon.com/ecs/>
21 <http://aws.amazon.com/ec2/>
22 <http://aws.amazon.com/elasticloadbalancing/>
23 <http://aws.amazon.com/ebs/>
24 <http://aws.amazon.com/iam/>
25 <http://aws.amazon.com/elasticbeanstalk/>
26 <http://aws.amazon.com/ec2/vm-import/>
27 <http://aws.amazon.com/s3/>
28 <http://aws.amazon.com/glacier/>
29 <http://aws.amazon.com/ebs/>
30 <https://aws.amazon.com/efs/>
31 <http://aws.amazon.com/storagegateway/>
32 <http://aws.amazon.com/cloudfront/>
33 <http://aws.amazon.com/importexport/>
34 <http://aws.amazon.com/rds/>
35 <http://aws.amazon.com/rds/aurora/>
36 <http://aws.amazon.com/rds/oracle/>
37 <http://aws.amazon.com/rds/sqlserver/>
38 <http://aws.amazon.com/rds/postgresql/>
39 <http://aws.amazon.com/rds/mysql/>
40 <http://aws.amazon.com/rds/mariadb/>
41 <https://aws.amazon.com/rds/aurora/>
42 <https://aws.amazon.com/dms/>
43 <http://aws.amazon.com/dynamodb/>
44 <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Tools.DynamoDBLocal.html>
45 <http://aws.amazon.com/redshift/>
46 <http://aws.amazon.com/elasticache/>
47 <http://www.memcached.org/>
48 <http://redis.io/>
49 <http://aws.amazon.com/vpc/>
50 <http://aws.amazon.com/directconnect/>
51 <http://aws.amazon.com/route53/>
52 <https://aws.amazon.com/codecommit/>
53 <https://aws.amazon.com/codedeploy/>
54 <https://aws.amazon.com/codepipeline/>
55 <http://aws.amazon.com/cloudwatch/>
56 <http://aws.amazon.com/cloudformation/>
57 <http://aws.amazon.com/cloudformation/details/#designer>
58 <http://aws.amazon.com/cloudtrail/>
59 <https://aws.amazon.com/config/>
60 <https://aws.amazon.com/opsworks/>
61 <http://aws.amazon.com/servicecatalog/>
62 <https://aws.amazon.com/premiumsupport/trustedadvisor/>

63 <http://aws.amazon.com/iam/>
64 <http://aws.amazon.com/iam/details/manage-users/>
65 <http://aws.amazon.com/iam/details/managing-user-credentials/>
66 <http://aws.amazon.com/iam/details/mfa/>
67 <http://aws.amazon.com/iam/details/manage-roles/>
68 <http://aws.amazon.com/iam/details/manage-permissions/>
69 <http://aws.amazon.com/iam/details/manage-federation/>
70 <http://aws.amazon.com/iam/details/manage-permissions/>
71 <http://aws.amazon.com/kms/>
72 <http://aws.amazon.com/directoryservice/>
73 <https://aws.amazon.com/inspector/>
74 <https://aws.amazon.com/waf/>
75 <http://aws.amazon.com/cloudhsm/>
76 <http://aws.amazon.com/elasticmapreduce/>
77 <https://aws.amazon.com/quicksight/>
78 <http://aws.amazon.com/datapipeline/>
79 <https://aws.amazon.com/elasticsearch-service/>
80 <http://aws.amazon.com/kinesis/>
81 <http://aws.amazon.com/kinesis/firehose/>
82 <http://aws.amazon.com/kinesis/analytics/>
83 <http://aws.amazon.com/kinesis/streams/>
84 <https://aws.amazon.com/machine-learning/>
85 <https://aws.amazon.com/iot/>
86 <https://aws.amazon.com/mobile/>
87 <http://aws.amazon.com/cognito/>
88 <https://aws.amazon.com/device-farm/>
89 <http://aws.amazon.com/mobileanalytics/>
90 <http://aws.amazon.com/sns/>
91 <http://aws.amazon.com/mobile/sdk/>
92 <https://aws.amazon.com/api-gateway/>
93 <https://aws.amazon.com/appstream/>
94 <http://aws.amazon.com/cloudsearch/>
95 <http://aws.amazon.com/elastictranscoder/>
96 <http://aws.amazon.com/ses/>
97 <http://aws.amazon.com/sqs/>
98 <http://aws.amazon.com/swf/>
99 <http://aws.amazon.com/workspaces/>
100 <http://aws.amazon.com/workdocs/>
101 <https://aws.amazon.com/workmail/>
102 <http://aws.amazon.com/free/>
103 <https://aws.amazon.com/premiumsupport/signup/>
104 https://aws.amazon.com/training/intro_series/