



Agile Solutions G-Cloud 12 Service Definition Document Amazon Web Services (AWS)



Service family at a glance...

What We Offer

- Broad range of AWS services to meet all data and analytics needs
- Strong experience in delivering AWS solutions to the Public Sector
- Planning, support, development and training in AWS products

Services Included

- Amazon EC2
- Amazon EBS, Autoscaling, Route 53, ELB
- Amazon ECS, Amazon ECR
- AWS Lambda
- AWS VPC
- Amazon S3
- Amazon Glacier
- Amazon RDS
- Amazon Aurora
- Amazon Redshift
- Dynamo DB
- IAM
- CloudTrail
- KMS, SNS, SES, SQS
- CloudWatch
- CloudFormation
- Quicksight, Kinesis
- Glue
- Step Functions
- Athena
- Lake Formation

Our Differentiators

- AWS Partners
- UK Public Sector experience with suite of AWS solutions
- End-to-end AWS delivery experience
- SC and DV cleared resources available
- Experienced in delivering to GDS standards

Table of Contents

1	Related Service Families	5
2	Individual Service Details	6
2.1	Amazon EC2	6
2.2	Amazon Elastic Block Storage (EBS), Auto Scaling, Route 53, Elastic Load Balancer	8
2.3	Amazon EC2 Elastic Container Service (ECS) and Amazon Elastic MapReduce (EMR)	11
2.4	AWS Lambda	13
2.5	AWS Virtual Private Cloud (VPC)	13
2.6	Amazon Simple Storage Service (Amazon S3).....	14
2.7	Amazon Glacier.....	15
2.8	Amazon Relational Database Service (RDS).....	16
2.9	Amazon Aurora	17
2.10	Amazon Redshift	19
2.11	Amazon Dynamo DB	20
2.12	AWS Identity and Access Management (IAM).....	21
2.13	AWS CloudTrail	22
2.14	Key Management Service (KMS), Simple Notification Service (SNS), Simple Email Service (SES), Simple Queue Service (SQS).....	23
2.15	Amazon CloudWatch	26
2.16	CloudFormation	28
2.17	Amazon Quicksight, Amazon Kinesis.....	29
2.18	AWS Glue.....	31
2.19	AWS Step Functions.....	32
2.20	Amazon Athena	32
2.21	Lake Formation.....	33
3	Where We Have Delivered Before	35
4	About Agile Solutions.....	36

Cloud Strategy and Planning



Understand your data with Analytics

Agile Solutions and AWS are here to help build your on-premise infrastructure to the cloud

Category	Use Cases	AWS Services
Analytics	Interactive analytics	Amazon Athena Query data in S3 using SQL
	Big data processing	Amazon EMR Hosted Hadoop Framework
	Data warehousing	Amazon Redshift Fast, simple, cost effective data warehousing
	Real-time analytics	Amazon Kinesis Analyse real time video and data streams
	Operational analytics	Amazon Elasticsearch Service Index your logs and visualize with Kibana
	Dashboards and visualizations	Amazon QuickSight Fast business analytics service
Data movement	Real-time data movement	Amazon Kinesis Data Analytics Stream and analyse data
		Amazon Kinesis Data Firehose Real time data streams into data stores and analytics
		Amazon Kinesis Data Streams Collecting and making data available in milliseconds
		Amazon Kinesis Video Streams Video streams for Analytics and Machine Learning

Who we've delivered for



1 Related Service Families

Agile Solutions offers a range of G Cloud services which complement any AWS initiative for the public sector. Please see our G Cloud 12 listings for:

- Data Lake and Data Warehouse
- Streaming
- Messaging
- Data Integration
- Cloud DevOps

2 Individual Service Details

2.1 Amazon EC2

Amazon EC2 is a web service that provides resizable compute capacity in the cloud. It is designed to make web-scale cloud computing easier for developers. Amazon EC2's simple 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 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 the tools to build failure resilient applications and isolate them from common failure scenarios.

Top benefits include:

- 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.
- Completely Controlled – You have complete control of your instances. You have root access to each one, and you can interact with them as you would any machine.
- Flexible Cloud Hosting Services – You have the choice of multiple 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.
- Designed for Use with Other AWS Cloud Services – Amazon EC2 works in conjunction with Amazon S3, Amazon RDS, and other AWS Cloud services 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 centres.
- Secure – Amazon EC2 works in conjunction with Amazon VPC to provide security and robust networking functionality for your compute resources.
- Lockdown security model prohibits administrative access, eliminating the possibility of human error and tampering. With the AWS Nitro System virtualization resources are offloaded to dedicated hardware and software minimizing the attack surface. AWS supports 89 security standards and compliance certifications including PCI-DSS, HIPAA/HITECH, FedRAMP, GDPR, FIPS 140-2, and NIST 800-171.
- 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.

We offer four different ways to buy instances, each with their own cost benefits:

- 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** – A Reserved Instance provides you with a significant discount (up to 75%) compared to On-Demand Instance pricing. There are three Reserved Instance payment options—No Upfront, Partial Upfront, and All Upfront—that enable you to balance the amount you pay up front 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 (e.g., 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 customers to bid on unused Amazon EC2 capacity and run those instances for as long as their 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 have flexibility in when your applications can run, Spot Instances can significantly lower your Amazon EC2 costs.
- **Savings Plans** - Savings Plans is a flexible pricing model that provides savings of up to 72% on your AWS compute usage. This pricing model offers lower prices on Amazon EC2 instances usage, regardless of instance family, size, OS, tenancy or AWS Region, and also applies to AWS Fargate and AWS Lambda usage. Savings Plans offer significant savings over On Demand, just like EC2 Reserved Instances, in exchange for a commitment to use a specific amount of compute power (measured in \$/hour) for a one or three year period. You can sign up for Savings Plans for a 1- or 3-year term and easily manage your plans by taking advantage of recommendations, performance reporting and budget alerts in the AWS Cost Explorer.

2.2 Amazon Elastic Block Storage (EBS), Auto Scaling, Route 53, Elastic Load Balancer

EBS

Amazon EBS provides persistent, available, and durable block-level storage volumes for use with Amazon EC2 instances in the AWS Cloud. 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.

Top benefits include:

- **Reliable, Secure Storage** – Each Amazon EBS volume is automatically replicated within its Availability Zone to protect you from component failure.
- **Consistent and Low-Latency Performance** – Amazon EBS General Purpose volumes and Amazon EBS Provisioned IOPS volumes deliver lowlatency through SSD technology and consistent I/O performance scaled to the needs of your application.
- **Backup, Restore, and Innovate** – Back up your data by taking point-in-time snapshots of your Amazon EBS volumes. Boost the agility of your business by using Amazon EBS snapshots to create new Amazon EC2 instances.
- **Geographic Flexibility** – Amazon EBS provides the ability to copy snapshots across AWS Regions, enabling geographical expansion, data centre migration, and disaster recovery.
- **Quickly Scale Up and Easily Scale Down** – Increase or decrease block storage and performance within minutes, enjoying the freedom to adjust as your needs evolve.

Auto Scaling

AWS Auto Scaling is a web service designed to launch or terminate Amazon EC2 instances automatically based on user-defined policies, schedules, and health checks Top benefits include:

- **Setup Scaling Quickly** - AWS Auto Scaling lets you set target utilization levels for multiple resources in a single, intuitive interface. You can quickly see the average utilization of all of your scalable resources without having to navigate to other consoles. For example, if your application uses Amazon EC2 and Amazon DynamoDB, you can use AWS Auto Scaling to manage resource provisioning for all of the EC2 Auto Scaling groups and database tables in your application.
- **Make Smart Scaling Decisions** - AWS Auto Scaling lets you build scaling plans that automate how groups of different resources respond to changes in demand. You can optimize availability, costs, or a balance of both. AWS Auto Scaling automatically creates all of the scaling policies and sets targets for you based on your preference. AWS Auto Scaling monitors your application and automatically adds or removes capacity from your resource groups in real-time as demands change.
- **Automatically Maintain Performance** - Using AWS Auto Scaling, you maintain optimal application performance and availability, even when workloads are periodic, unpredictable, or continuously changing. AWS Auto Scaling continually monitors

your applications to make sure that they are operating at your desired performance levels. When demand spikes, AWS Auto Scaling automatically increases the capacity of constrained resources so you maintain a high quality of service.

- **Pay Only For What You Need** - AWS Auto Scaling can help you optimize your utilization and cost efficiencies when consuming AWS services so you only pay for the resources you actually need. When demand drops, AWS Auto Scaling will automatically remove any excess resource capacity so you avoid overspending. AWS Auto Scaling is free to use, and allows you to optimize the costs of your AWS environment.

Route 53

Amazon Route 53 is a highly available and scalable cloud 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 names like `www.example.com` into the numeric IP addresses like `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 Traffic Flow makes it easy 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. Using Amazon Route 53 Traffic Flow's simple visual editor, you can easily manage how your end users are routed to your application's endpoints—whether in a single AWS Region or distributed around the globe. 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.

Top benefits include:

- **Highly Available and Reliable** – Amazon Route 53 is built using AWS's highly available and reliable infrastructure. The distributed nature of our DNS servers helps ensure a consistent ability to route your end users to your application. Features such as Amazon Route 53 Traffic Flow help you improve reliability with easy configuration of failover to re-route your users to an alternate location if your primary application endpoint becomes unavailable. Amazon Route 53 is designed to provide the level of dependability required by important applications. Amazon Route 53 is backed by the [Amazon Route 53 Service Level Agreement](#).
- **Scalable** – Amazon Route 53 is designed to automatically scale to handle very large query volumes without any intervention from you.
- **Designed for Use with Other AWS Cloud Services** – Amazon Route 53 is designed to work well with other AWS features and offerings. You can use Amazon Route 53 to map domain names to your Amazon EC2 instances, Amazon S3 buckets, Amazon CloudFront distributions, and other AWS resources. By using the AWS Identity and Access Management (IAM) service with Amazon Route 53, you get fine-grained control over who can update your DNS data. You can use Amazon Route 53 to map your zone apex (`example.com` versus `www.example.com`) to your Elastic Load Balancing instance, Amazon CloudFront distribution, AWS Elastic

Beanstalk environment, or Amazon S3 website bucket using a feature called Alias record.

- **Simple** – With self-service sign-up, Amazon Route 53 can start to answer your DNS queries within minutes. You can configure your DNS settings with the AWS Management Console or our easy-to-use API. You can also programmatically integrate the Amazon Route 53 API into your overall web application. For instance, you can use Amazon Route 53's API to create a new DNS record whenever you create a new Amazon EC2 instance. Amazon Route 53 Traffic Flow makes it easy to set up sophisticated routing logic for your applications by using the simple visual policy editor.
- **Fast** – Using a global anycast network of DNS servers around the world, Amazon Route 53 is designed to automatically route your users to the optimal location depending on network conditions. As a result, the service offers low query latency for your end users, as well as low update latency for your DNS record management needs. Amazon Route 53 Traffic Flow lets you further improve your customers' experience by running your application in multiple locations around the world and using traffic policies to ensure your end users are routed to the closest healthy endpoint for your application.
- **Cost Effective** – Amazon Route 53 passes on the benefits of AWS's scale to you. You pay only for the resources you use, such as the number of queries that the service answers for each of your domains, hosted zones for managing domains through the service, and optional features such as traffic policies and health checks, all at a low cost and without minimum usage commitments or any up-front fees.
- **Secure** – By integrating Amazon Route 53 with AWS IAM, you can grant unique credentials and manage permissions for every user within your AWS account and specify who has access to which parts of the Amazon Route 53 service.
- **Flexible** – Amazon Route 53 Traffic Flow routes traffic based on multiple criteria, such as endpoint health, geographic location, and latency. You can configure multiple traffic policies and decide which policies are active at any given time. You can create and edit traffic policies using the simple visual editor in the Amazon Route 53 console, AWS Software Development Kits (SDKs), or the Amazon Route 53 API. Traffic Flow's versioning feature maintains a history of changes to your traffic policies, so you can easily roll back to a previous version using the console or API.

Elastic Load Balancing

Elastic Load Balancing automatically distributes incoming application traffic across multiple targets, such as Amazon EC2 instances, containers, IP addresses, and Lambda functions. It can handle the varying load of your application traffic in a single Availability Zone or across multiple Availability Zones. Elastic Load Balancing offers three types of load balancers that all feature the high availability, automatic scaling, and robust security necessary to make your applications fault tolerant.

- **Highly Available** - Elastic Load Balancing automatically distributes incoming traffic across multiple targets – Amazon EC2 instances, containers, IP addresses, and Lambda functions – in multiple Availability Zones and ensures only healthy targets receive traffic. Elastic Load Balancing can also load balance across a Region, routing traffic to healthy targets in different Availability Zones. The Amazon Elastic

Load Balancing Service Level Agreement commitment is 99.99% availability for a load balancer.

- **Secure** - Elastic Load Balancing works with Amazon Virtual Private Cloud (VPC) to provide robust security features, including integrated certificate management, user-authentication, and SSL/TLS decryption. Together, they give you the flexibility to centrally manage TLS settings and offload CPU intensive workloads from your applications.
- **Elastic** - Elastic Load Balancing is capable of handling rapid changes in network traffic patterns. Additionally, deep integration with Auto Scaling ensures sufficient application capacity to meet varying levels of application load without requiring manual intervention.
- **Flexible** - Elastic Load Balancing also allows you to use IP addresses to route requests to application targets. This offers you flexibility in how you virtualize your application targets, allowing you to host more applications on the same instance. This also enables these applications to have individual security groups and use the same network port to further simplify inter-application communication in microservice-based architecture.
- **Robust monitoring & auditing** - Elastic Load Balancing allows you to monitor your applications and their performance in real time with Amazon CloudWatch metrics, logging, and request tracing. This improves visibility into the behavior of your applications, uncovering issues and identifying performance bottlenecks in your application stack at the granularity of an individual request.
- **Hybrid Load Balancing** - Elastic Load Balancing offers ability to load balance across AWS and on-premises resources using the same load balancer. This makes it easy for you to migrate, burst, or failover on-premises applications to the cloud.

2.3 Amazon EC2 Elastic Container Service (ECS) and Amazon Elastic MapReduce (EMR)

ECS

Amazon ECS is a highly scalable, high-performance container management service that supports Docker containers and allows you to easily run applications on a managed cluster of Amazon EC2 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 Balancing,

Amazon EBS volumes, and AWS IAM 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.

Top benefits include:

- **Easily Manage Clusters for Any Scale** – Amazon ECS eliminates the need to operate cluster management software or design fault-tolerant cluster architectures. There is no software to install, manage, or scale, allowing you to focus on building Dockerised applications.

- **Flexible Container Placement** – You can use Amazon ECS to schedule long running applications and batch jobs. You can also use the APIs to get up-to-date cluster state information to integrate your own custom, as well as third-party, schedulers. Amazon ECS is a shared-state, optimistic concurrency system that supports multiple schedulers on the same cluster for each business or application-specific requirement.
- **Designed for Use with Other AWS Cloud Services** – Amazon ECS is integrated with Elastic Load Balancing, Amazon EBS, Amazon VPC, IAM, and AWS CloudTrail, providing you a complete solution for running a wide range of containerised applications or services.
- **Extensible** – Amazon ECS can easily be integrated or extended through simple APIs. Amazon ECS provides complete visibility and control into your AWS resources, so you can easily integrate and use your own schedulers or connect Amazon ECS into your existing software delivery process (e.g., continuous integration and delivery systems).
- **Performance at Scale** – Amazon ECS is built on technology developed from many years of experience running highly scalable services. You can launch tens or tens of thousands of Docker containers in seconds using Amazon ECS with no additional complexity.
- **Secure** – Amazon ECS launches your containers on your own Amazon EC2 instances. No compute resources are shared with other customers. Your clusters run in an Amazon VPC, allowing you to use your own Amazon VPC security groups and network Access Control Lists (ACLs). These features provide you a high level of isolation and help you use Amazon ECS to build highly secure and reliable applications.

EMR

Amazon EMR provides a managed Hadoop framework that makes it easy, fast, and cost-effective to process vast amounts of data across dynamically scalable Amazon EC2 instances. You can also run other popular distributed frameworks such as Apache Spark, HBase, Presto, and Flink 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 a broad set of big data use cases, including log analysis, web indexing, data transformations (Extract, Transform, Load [ETL]), ML, financial analysis, scientific simulation, and bioinformatics.

Top features include:

- **Easy to Use** – You can launch an Amazon EMR cluster in minutes. You don't need to worry about node provisioning, cluster setup, Hadoop configuration, or cluster tuning. Amazon EMR takes care of these tasks so you can focus on analysis.
- **Low Cost** – Amazon EMR pricing is simple and predictable: You pay an hourly rate for every instance hour you use. Because Amazon EMR has native support for Amazon EC2 Spot and Reserved Instances, you can also save 50-80% on the cost of the underlying instances.
- **Elastic** – With Amazon EMR, you can provision one, hundreds, or thousands of compute instances to process data at any scale. You can easily increase or decrease the number of instances and you only pay for what you use.

- **Reliable** – You can spend less time tuning and monitoring your cluster. Amazon EMR has tuned Hadoop for the cloud; it also monitors your cluster, retrying failed tasks and automatically replacing poorly performing instances.
- **Secure** – Amazon EMR automatically configures Amazon EC2 firewall settings that control network access to instances, and you can launch clusters in an Amazon VPC, a logically isolated network you define. For objects stored in Amazon S3, you can use Amazon S3 server-side encryption or Amazon S3 client-side encryption with Amazon EMR File System (EMRFS), using AWS KMS or customer-managed keys.
- **Flexible** – You have complete control over your cluster. You have root access to every instance, you can easily install additional applications, and you can customise every cluster. Amazon EMR also supports multiple Hadoop distributions and applications.

2.4 AWS Lambda

AWS Lambda lets you run code without provisioning or managing servers. You pay only for the compute time you consume.

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 call it directly from any web or mobile app.

Benefits:

- No servers to manage - AWS Lambda automatically runs your code without requiring you to provision or manage servers. Just write the code and upload it to Lambda.
- Continuous scaling - AWS Lambda automatically scales your application by running code in response to each trigger. Your code runs in parallel and processes each trigger individually, scaling precisely with the size of the workload.
- Sub-second metering - With AWS Lambda, you are charged for every 100ms your code executes and the number of times your code is triggered. You pay only for the compute time you consume.
- Consistent performance - With AWS Lambda, you can optimize your code execution time by choosing the right memory size for your function. You can also enable Provisioned Concurrency to keep your functions initialized and hyper-ready to respond within double digit milliseconds.

2.5 AWS Virtual Private Cloud (VPC)

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 customise 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 centre and your VPC and leverage the AWS Cloud as an extension of your corporate data centre.

Top benefits include:

- **Multiple Connectivity Options** – A variety of connectivity options exist for your Amazon VPC. You can connect your VPC to the Internet, to your data centre, or to other VPCs, based on the AWS resources that you want to expose publicly and those that you want to keep private.
- **Secure** – Amazon VPC provides advanced security features such as security groups and network access control lists to enable inbound and outbound filtering
- **Simple** – You can create a VPC quickly and easily using the AWS Management Console.
- **All the Scalability and Reliability of AWS** – Amazon VPC provides all the same benefits as the rest of the AWS platform. You can instantly scale your resources up or down, select Amazon EC2 instance types and sizes that are right for your applications, and pay only for the resources you use—all within Amazon's proven infrastructure.

2.6 Amazon Simple Storage Service (Amazon S3)

Amazon S3 provides developers and IT teams with secure, durable, and highly scalable object storage. Amazon S3 is easy to use, with a simple interface to store and retrieve any amount of data from anywhere on the web. With Amazon S3, you pay only for the storage you actually use. There is no minimum fee and no setup cost.

Amazon S3 can be used alone or together with other AWS Cloud services such as Amazon EC2, Amazon EBS, and Amazon Glacier, as well as third-party storage repositories and gateways. Amazon S3 provides 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.

Top benefits include:

- **Durable** – Amazon S3 provides durable infrastructure to store important data and is designed for durability of 99.999999999% of objects.
- **Low Cost** – Amazon S3 allows you to store large amounts of data at a very low cost. You pay for what you need, with no minimum commitments or upfront fees.
- **Available** – Amazon S3 is designed for 99.99% availability of objects over a given year.
- **Secure** – Amazon S3 supports data transfer over SSL and automatic encryption of your data once it is uploaded.
- **Scalable** – With Amazon S3, you can store as much data as you want and access it when you need it.
- **Send Event Notifications** – Amazon S3 can send event notifications when objects are uploaded to Amazon S3.

- High Performance – Amazon S3 supports multi-part uploads to help maximise network throughput and resiliency and lets you choose the AWS Region in which to store your data, minimising network latency.
- Integrated – Amazon S3 is integrated with other AWS Cloud services to simplify uploading and downloading data from Amazon S3 and to make it easier to build solutions that use a range of AWS Cloud services.
- Easy to Use – Amazon S3 is easy to use with a web-based management console and mobile app and full Representational State Transfer (REST) APIs and SDKs for easy integration with third-party technologies.
- Access management - To protect your data in Amazon S3, by default, users only have access to the S3 resources they create. You can grant access to other users by using one or a combination of the following access management features: AWS Identity and Access Management (IAM) to create users and manage their respective access; Access Control Lists (ACLs) to make individual objects accessible to authorized users; bucket policies to configure permissions for all objects within a single S3 bucket; S3 Access Points to simplify managing data access to shared data sets by creating access points with names and permissions specific to each application or sets of applications; and Query String Authentication to grant time-limited access to others with temporary URLs. Amazon S3 also supports Audit Logs that list the requests made against your S3 resources for complete visibility into who is accessing what data.
- Write Once, Read Many - You can also enforce write-once-read-many (WORM) policies with S3 Object Lock. This S3 management feature blocks object version deletion during a customer-defined retention period so that you can enforce retention policies as an added layer of data protection or to meet compliance obligations.
- Transferring large amounts of data - AWS has a suite of data migration services that make transferring data into the AWS Cloud simple, fast, and secure. S3 Transfer Acceleration is designed to maximize transfer speeds to S3 buckets over long distances. For very large data transfers, consider using AWS Snowball, AWS Snowball Edge, and AWS Snowmobile to move petabytes to exabytes of data to the AWS Cloud for as little as one-fifth the cost of high-speed Internet.
- Query in place - Amazon S3 has a built-in feature and complimentary services that query data without needing to copy and load it into a separate analytics platform or data warehouse. This means you can run big data analytics directly on your data stored in Amazon S3. S3 Select is an S3 feature designed to increase query performance by up to 400%, and reduce querying costs as much as 80%. It works by retrieving a subset of an object's data (using simple SQL expressions) instead of the entire object, which can be up to 5 terabytes in size.

2.7 Amazon Glacier

Amazon Glacier is a secure, durable, and extremely low-cost storage service for data archiving and online backup. Customers can reliably store large or small amounts of data. To keep costs low yet suitable for varying retrieval needs, Amazon Glacier provides three options for access to archives, from a few minutes to several hours.

Top benefits include:

- Low Cost – Amazon Glacier allows you to archive large amounts of data at a very low cost.

- Secure – Amazon Glacier supports data transfer over SSL and automatically encrypts your data at rest.
- Durable – Amazon Glacier provides a highly durable storage infrastructure designed for online backup and archival. Your data is redundantly stored across multiple facilities and multiple devices in each facility.
- Simple – Amazon Glacier allows you to offload the administrative burden of operating storage infrastructure to AWS.
- Flexible – Amazon Glacier scales to meet your storage needs. There is no limit to how much data you can store, and you can choose to store your data in the AWS Region that supports your regulatory and business criteria.
- Integrated – Through Amazon S3 life cycle policies, you can optimise your storage costs by moving infrequently accessed objects from Amazon S3 to Amazon Glacier (or vice versa).

2.8 Amazon Relational Database Service (RDS)

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.

Top benefits include:

- Easy to Administer – Amazon RDS makes it easy to go from project conception to deployment. Use the AWS Management Console, the Amazon RDS CLI, or simple API calls to access the capabilities of a production-ready relational database in minutes. There is no need for infrastructure provisioning and no need for installing and maintaining database software.
- Scalable – You can scale your database's compute and storage resources with only a few mouse clicks or an API call, often with no downtime.
- Available and Durable – Amazon RDS runs on the same highly reliable infrastructure used by other AWS Cloud services. When you provision a MultiAZ Database Instance (DB Instance), Amazon RDS synchronously replicates the data to a standby instance in a different Availability Zone. Amazon RDS has many other features that enhance reliability for critical production databases, including automated backups, Database Snapshots (DB Snapshots), and automatic host replacement.
- Fast – Amazon RDS offers database server sizing choices up to 32 vCPUs and 244 GiB, as well as storage choices for a wide range of application performance requirements.
- Secure – Amazon RDS makes it easy to control network access to your database. Amazon RDS also lets you run your database instances in Amazon VPC, which enables you to isolate your database instances and to connect to your existing IT infrastructure through an industry-standard encrypted IPsec VPN. Many Amazon RDS engine types offer encryption at rest and encryption in transit.
- Inexpensive – You pay very low rates and only for the resources you actually consume. In addition, you benefit from the option of on-demand pricing with no up-front or long-term commitments or even lower hourly rates via our reserved pricing option.

2.9 Amazon Aurora

Amazon Aurora is a MySQL and PostgreSQL-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 with the security, availability, and reliability of a commercial database at one tenth the cost.

Top features include:

- Fully managed MySQL and PostgreSQL compatible relational database.
- Distributed, fault-tolerant, self-healing storage system that auto-scales.
- Amazon Aurora Global Database is designed for globally distributed applications, allowing a single Amazon Aurora database to span multiple AWS regions.
- Amazon Aurora Serverless is an on-demand, auto-scaling configuration for Amazon Aurora (MySQL-compatible and PostgreSQL-compatible editions), where the database will automatically start up, shut down, and scale capacity up or down based on your application's needs.
- Amazon Aurora machine learning enables you to add ML-based predictions to applications via the familiar SQL programming language, so you don't need to learn separate tools or have prior machine learning experience.

Top Benefits include:

- High Performance - Get 5X the throughput of standard MySQL and 3X the throughput of standard PostgreSQL. This performance is on par with commercial databases, at 1/10th the cost. Amazon Aurora uses a variety of software and hardware techniques to ensure the database engine is able to fully leverage available compute, memory and networking. I/O operations use distributed systems techniques such as quorums to improve performance consistency.
- Highly Scalable - You can easily scale your database deployment up and down from smaller to larger instance types as your needs change, or let Aurora Serverless handle scaling automatically for you. To scale read capacity and performance, you can add up to 15 low latency read replicas across three Availability Zones. Amazon Aurora automatically grows storage as needed, up to 64TB per database instance.
- Highly Secure – Amazon Aurora provides multiple levels of security for your database. These include network isolation using Amazon VPC, encryption at rest using keys you create and control through AWS Key Management Service (KMS) and encryption of data in transit using SSL. On an encrypted Amazon Aurora instance, data in the underlying storage is encrypted, as are the automated backups, snapshots, and replicas in the same cluster.
- High Availability and Durability – Amazon Aurora is designed to offer greater than 99.99% availability, replicating 6 copies of your data across 3 Availability Zones and backing up your data continuously to Amazon S3. It transparently recovers from physical storage failures; instance failover typically takes less than 30 seconds. You can also backtrack within seconds to a previous point in time, to recover from user errors. With Global Database, a single Aurora database can span multiple AWS regions to enable fast local reads and quick disaster recovery.
- Fully Managed – Amazon Aurora is fully managed by Amazon Relational Database Service (RDS). You no longer need to worry about database management tasks such as hardware provisioning, software patching, setup, configuration, or backups.

Aurora automatically and continuously monitors and backs up your database to Amazon S3, enabling granular point-in-time recovery. You can monitor database performance using Amazon CloudWatch, Enhanced Monitoring, or Performance Insights, an easy-to-use tool that helps you quickly detect performance problems.

- **MySQL and PostgreSQL Compatible** – The Amazon Aurora database engine is fully compatible with existing MySQL and PostgreSQL open source databases, and adds compatibility for new releases regularly. This means you can easily migrate MySQL or PostgreSQL databases to Aurora using standard MySQL or PostgreSQL import/export tools or snapshots. It also means the code, applications, drivers, and tools you already use with your existing databases can be used with Amazon Aurora with little or no change.
- **Migration Support** - MySQL and PostgreSQL compatibility make Amazon Aurora a compelling target for database migrations to the cloud. If you're migrating from MySQL or PostgreSQL, see our migration documentation for a list of tools and options. To migrate from commercial database engines, you can use the AWS Database Migration Service for a secure migration with minimal downtime.

2.10 Amazon Redshift

Amazon Redshift is a fast, fully managed, petabyte-scale data warehouse service that makes it simple and cost-effective to efficiently analyse all your data using your existing business intelligence tools. It is optimised for datasets ranging from a few hundred gigabytes to a petabyte or more and costs a tenth the cost of most traditional data warehousing solutions.

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 Connect Client tab of our 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 EMR, Amazon Kinesis, or any SSH-enabled host.

Amazon Redshift's data warehouse architecture allows you to automate most of the common administrative tasks associated with provisioning, configuring and monitoring a cloud 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.

Top features and benefits include:

- **Optimised for Data Warehousing** – Amazon Redshift uses a variety of innovations to obtain very high query performance on datasets ranging in size from a hundred gigabytes to a petabyte or more. It uses columnar storage, data compression, and zone maps to reduce the amount of I/O needed to perform queries. Amazon Redshift has a Massively Parallel Processing (MPP) data warehouse architecture, parallelizing and distributing SQL operations to take advantage of all available resources. The underlying hardware is designed for high performance data processing, using local attached storage to maximise throughput between the CPUs and drives, and a 10GigE mesh network to maximise throughput between nodes.
- **Scalable** – With a few clicks of the AWS Management Console or a simple API call, you can easily change the number or type of nodes in your cloud data warehouse as your performance or capacity needs change. Dense Storage (DS) nodes allow you to create very large data warehouses using HDDs for a very low price point. Dense Compute (DC) nodes allow you to create very high performance data warehouses using fast CPUs, large amounts of RAM and SSDs.
- **No Up-Front Costs** – You pay only for the resources you provision. You can choose On-Demand pricing with no up-front costs or long-term commitments, or obtain significantly discounted rates with Reserved Instance pricing.
- **Get Started in Minutes** – With a few clicks in the AWS Management Console or simple API calls, you can create a cluster, specifying its size, underlying node type, and security profile. Amazon Redshift will provision your nodes, configure the connections between them, and secure the cluster. Your data warehouse should be up and running in minutes.
- **Fully Managed** – Amazon Redshift handles all of the work needed to manage, monitor, and scale your data warehouse, from monitoring cluster health and taking backups to applying patches and upgrades. You can easily resize your cluster as your performance and capacity needs change. By handling all these time-

consuming, labour-intensive tasks, Amazon Redshift frees you up to focus on your data and business.

- **Fault Tolerant** – Amazon Redshift has multiple features that enhance the reliability of your data warehouse cluster. All data written to a node in your cluster is automatically replicated to other nodes within the cluster and all data is continuously backed up to Amazon S3. Amazon Redshift continuously monitors the health of the cluster and automatically re-replicates data from failed drives and replaces nodes as necessary.
- **Encryption** – With just a couple of parameter settings, you can set up Amazon Redshift to use SSL to secure data in transit and hardware-accelerated AES-256 encryption for data at rest. If you choose to enable encryption of data at rest, all data written to disk will be encrypted as well as any backups. By default, Amazon Redshift takes care of key management but you can choose to manage your keys using your own HSMs, AWS CloudHSM, or AWS KMS.
- **Network Isolation** – Amazon Redshift enables you to configure firewall rules to control network access to your data warehouse cluster. You can run Amazon Redshift inside Amazon VPC to isolate your data warehouse cluster in your own virtual network and connect it to your existing IT infrastructure using industry standard encrypted IPsec VPN.
- **Audit and Compliance** – Amazon Redshift integrates with AWS CloudTrail to enable you to audit all Amazon Redshift API calls. Amazon Redshift also logs all SQL operations, including connection attempts, queries and changes to your database. You can access these logs using SQL queries against system tables or choose to have them downloaded to a secure location on Amazon S3. Amazon Redshift is compliant with Service Organisation Controls (SOC) 1, SOC 2, SOC 3, and Payment Card Industry Data Security Standard (PCI DSS) Level 1 requirements. For more details, please visit the AWS Compliance Center.
- **SQL** – Amazon Redshift is a SQL data warehouse solution and uses industry standard ODBC and JDBC connections. You can download our custom JDBC and ODBC drivers from the Connect Client tab of the service console. Many popular software vendors have certified Amazon Redshift with their offerings to enable you to continue to use the tools you do today.
- **Integrated** – Amazon Redshift is integrated with other AWS Cloud services and has built in commands to load data in parallel to each node from Amazon S3 or your Amazon EC2 and on premise servers using SSH.

2.11 Amazon Dynamo DB

Amazon DynamoDB is a key-value and document database that delivers single-digit millisecond performance at any scale. It's a fully managed, multi-Region, multi-master database with built-in security, backup and restore, and in-memory caching for internet-scale applications. DynamoDB can handle more than 10 trillion requests per day and support peaks of more than 20 million requests per second.

Top features include:

- **DynamoDB Accelerator (DAX)** is an in-memory cache that delivers fast read performance for your tables at scale by enabling you to use a fully managed in-memory cache.
- **DynamoDB encrypts all customer data at rest by default.** Encryption at rest enhances the security of your data by using encryption keys stored in AWS KMS.

- DynamoDB global tables replicate your data automatically across your choice of AWS Regions and automatically scale capacity to accommodate your workloads.
- DynamoDB provides native, server-side support for transactions, simplifying the developer experience of making coordinated, all-or-nothing changes to multiple items both within and across tables.
- Amazon DynamoDB Streams captures a time-ordered sequence of item-level modifications in any DynamoDB table and stores this information in a log for up to 24 hours.

Top benefits include:

- High performance and scalability – DynamoDB is a key-value and document database that can support tables of virtually any size with horizontal scaling. This enables DynamoDB to scale to more than 10 trillion requests per day, with peaks greater than 20 million requests per second, over petabytes of storage.
- High availability and durability – DynamoDB global tables replicate your data automatically across your choice of AWS Regions and automatically scale capacity to accommodate your workloads. With global tables, your globally distributed applications can access data locally in the selected Regions to get single-digit millisecond read and write performance.
- Fully managed – With DynamoDB, there are no servers to provision, patch, or manage and no software to install, maintain, or operate. DynamoDB automatically scales tables up and down to adjust for capacity and maintain performance. Availability and fault tolerance are built in, eliminating the need to architect your applications for these capabilities. DynamoDB provides both on-demand and provisioned capacity modes so that you can optimise costs by specifying capacity per workload or by paying for only the resources you consume.
- Secure – DynamoDB encryption at rest provides an additional layer of data protection by securing your data in the encrypted table, including its primary key, local and global secondary indexes, streams, global tables, backups, and DAX clusters whenever the data is stored in durable media. Organisational policies, industry or government regulations, and compliance requirements often require the use of encryption at rest to increase the data security of your applications.
- Integration – DynamoDB integrates with AWS Lambda to provide triggers. Using triggers, you can automatically execute a custom function when item-level changes in a DynamoDB table are detected. With triggers, you can build applications that react to data modifications in DynamoDB tables. The Lambda function can perform any actions you specify, such as sending a notification or initiating a workflow.

2.12 AWS Identity and Access Management (IAM)

AWS IAM enables you to securely control access to AWS Cloud services and resources for your users. Using AWS IAM, you can create and manage AWS users and groups and use permissions to allow and deny their access to AWS resources.

Top features include:

- Manage AWS IAM Users and their Access – You can create users in AWS IAM, assign them individual security credentials (e.g., access keys, passwords, multi-factor authentication devices) or request temporary security credentials to provide users access to AWS Cloud services and resources. You can manage permissions in order to control which operations a user can perform.

- **Manage AWS IAM Roles and their Permissions** – You can create roles in AWS IAM and manage permissions to control which operations can be performed by the entity or AWS Cloud 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 in your enterprise to access the AWS Management Console, to call AWS APIs, and to access resources without the need to create an AWS IAM user for each identity.
- **Resource-Based Policies** - Resource-based policies are attached to a resource. For example, you can attach resource-based policies to Amazon S3 buckets, Amazon SQS queues, and AWS Key Management Service encryption keys.
- **Permissions Boundaries for IAM Entities** - AWS supports permissions boundaries for IAM entities (users or roles). A permissions boundary is an advanced feature for using a managed policy to set the maximum permissions that an identity-based policy can grant to an IAM entity. An entity's permissions boundary allows it to perform only the actions that are allowed by both its identity-based policies and its permissions boundaries.
- **IAM Access Analyzer** - AWS IAM Access Analyzer helps you identify the resources in your account, such as Amazon S3 buckets or IAM roles, that are shared with an external entity. This lets you identify unintended access to your resources and data, which is a security risk. Access Analyzer identifies resources that are shared with external principals by using logic-based reasoning to analyse the resource-based policies in your AWS environment.

2.13 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 AWS 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 Cloud services (such as AWS CloudFormation). The AWS API call history produced by AWS CloudTrail enables security analysis, resource change tracking, and compliance auditing.

Top features and benefits include:

- **Increased Visibility** – AWS CloudTrail provides increased visibility into your user activity by recording AWS API calls. You can answer questions such as, what actions did a given user take over a given time period? For a given resource, which user has taken actions on it over a given time period? What is the source IP address of a given activity? Which activities failed due to inadequate permissions?
- **Durable and Inexpensive Log File Storage** – AWS CloudTrail uses Amazon S3 for log file storage and delivery, so log files are stored durably and inexpensively. You can use Amazon S3 lifecycle configuration rules to further reduce storage costs. For example, you can define rules to automatically delete old log files or archive them to Amazon Glacier for additional savings.
- **Easy Administration** – AWS CloudTrail is a fully managed service; you simply turn on AWS CloudTrail for your account using the AWS Management Console, the

AWS CLI, or the AWS CloudTrail SDK and start receiving AWS CloudTrail log files in the Amazon S3 bucket that you specify.

- **Notifications for Log File Delivery** – AWS CloudTrail can be configured to publish a notification for each log file delivered, thus enabling you to automatically take action upon log file delivery. CloudTrail uses the Amazon SNS for notifications.
- **Log File Aggregation** – AWS CloudTrail can be configured to aggregate log files across multiple accounts and regions so that log files are delivered to a single bucket. For detailed instructions, refer to the Aggregating CloudTrail Log Files to a Single Amazon S3 Bucket section of the user guide.
- **Reliable and Timely Delivery** – AWS CloudTrail continuously transports events from AWS Cloud services using a highly available and fault tolerant processing pipeline. CloudTrail typically delivers events within 15 minutes of the API call.
- **Troubleshoot operational or security issues** – You can troubleshoot operational issues or perform security analysis by looking up API activity that was captured for your AWS account. Using the AWS CloudTrail console, AWS CLI, or AWS SDKs, you can quickly and easily answer questions related to API activity for the last 7 days and take immediate action.
- **Receive SNS Notifications of API activity** – AWS CloudTrail can be configured to be deliver API activity to an Amazon CloudWatch Logs log group you specify. You can then create CloudWatch Alarms to receive Amazon SNS notifications when specific API activity occurs.

2.14 Key Management Service (KMS), Simple Notification Service (SNS), Simple Email Service (SES), Simple Queue Service (SQS)

KMS

AWS 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 HSMs to protect the security of your keys. AWS Key Management Service is integrated with other AWS Cloud services including Amazon EBS, Amazon S3, Amazon RDS, and Amazon Redshift. AWS KMS is also integrated with AWS CloudTrail to provide you with logs of all key usage to help meet your regulatory and compliance needs.

Top benefits include:

- **Centralised Key Management** – AWS Key Management Service provides you with centralised control of your encryption keys. AWS KMS presents a single view into all of the key usage in your organisation. You can easily create keys, implement key rotation, create usage policies, and enable logging from the AWS Management Console, or by using the API.
- **Integrated with AWS Cloud Services** – AWS KMS is integrated with Amazon S3, Amazon EBS, Amazon Redshift, Amazon RDS, and Amazon EMR to make it easy to encrypt the data you store with these services using keys that you manage.
- **Encryption for all your applications** – AWS KMS makes it easy to manage encryption keys used to encrypt data stored by your applications regardless of where you store it. AWS KMS provides an SDK for programmatic integration of encryption and key management into your applications.

- **Built-in Auditing** – AWS KMS works with AWS CloudTrail to provide you with logs of API calls made to or by AWS KMS. These logs help you meet compliance and regulatory requirements by providing details of when keys were accessed and who accessed them.
- **Fully Managed** – AWS KMS is a fully managed service, so you can focus on the encryption needs of your applications while AWS handles availability, physical security, and hardware maintenance of the underlying infrastructure.
- **Low-cost** – There is no charge for the storage of default keys in your account. You pay only for additional master keys that you create and your key usage.
- **Secure** – AWS KMS provides you a secure location to store and use encryption keys, using hardened systems where your unencrypted keys are only used in memory. AWS KMS keys are never transmitted outside of the AWS Regions in which they were created.
- **Compliance** - The security and quality controls in AWS KMS have been certified under multiple compliance schemes to simplify your own compliance obligations. AWS KMS provides the option to store your keys in single-tenant HSMs in AWS CloudHSM instances that you control.
- **Digitally sign data** - AWS KMS enables you to perform digital signing operations using asymmetric key pairs to ensure the integrity of your data. Recipients of digitally signed data can verify the signatures whether they have an AWS account or not.

Simple Notification Service (SNS)

Amazon SNS is a fast, flexible, fully managed push notification service that lets you send individual messages or to fan-out messages to large numbers of recipients. Amazon SNS makes it simple and cost effective to send push notifications to mobile device users, email recipients or even send messages to other distributed services.

With Amazon SNS, you can send notifications to Apple, Google, Fire OS, and Windows devices, as well as to Android devices in China with Baidu Cloud Push. You can use Amazon SNS to send SMS messages to mobile device users in the US or to email recipients worldwide.

Beyond these endpoints, Amazon SNS can also deliver messages to Amazon SQS, AWS Lambda functions, or to any HTTP endpoint.

Top features include:

- Amazon SNS lets you push messages to mobile devices or distributed services via API or an easy-to-use management console. You can seamlessly scale from a handful of messages per day to millions of messages or higher.
- With Amazon SNS you can publish a message once and deliver it one or more times. So you can choose to direct unique messages to individual Apple, Google, or Amazon devices or broadcast deliveries to many mobile devices with a single publish request.
- Amazon SNS allows you to group multiple recipients using topics. A topic is an “access point” for allowing recipients to dynamically subscribe for identical copies of the same notification. One topic can support deliveries to multiple endpoint types—for example, you can group together iOS, Android and SMS recipients. When you

publish once to a topic, Amazon SNS delivers appropriately formatted copies of your message to each subscriber.

- Amazon SNS has no upfront costs and you can pay as you go.

Simple Email Service (SES)

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 SNS. With Amazon SES, you have no required minimum commitments—you pay as you go, and you only pay for what you use.

Top benefits include:

- **Inexpensive** – There are no up-front fees or fixed expenses with Amazon SES, and you benefit from the efficiencies of Amazon’s scale. You pay low charges for the number of emails sent, number of emails received, data transfer fees, and attachments.
- **Reliable** – Amazon SES runs within Amazon’s cloud network infrastructure and data centres. Multiple servers and data centres provide high availability and data durability.
- **Scalable** – Amazon SES is based on the scalable cloud-based email technology used by Amazon websites around the world to send billions of messages a year.
- **Highly Deliverable** – Amazon SES takes proactive steps to maximise the percentage of your emails that arrive in your recipients’ inboxes.
- **Highly Configurable** – When you use Amazon SES to receive your email, you have fine-grained control over what happens to your incoming messages. You can set up rules that apply to your entire domain or to specific email addresses.
- **Designed for Use with Other AWS Cloud Services** – Amazon SES integrates with Amazon SNS, Amazon Route 53, Amazon S3, Amazon WorkMail, AWS Lambda, AWS KMS, and Amazon CloudWatch. Additionally, emails you send from Amazon EC2 and AWS Elastic Beanstalk qualify for the free usage tier.

Simple Queue Service (SQS)

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. 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.

Top benefits include:

- **Reliable** – Amazon SQS runs within Amazon’s high-availability data centres, so queues will be available whenever applications need them. To prevent messages from being lost or becoming unavailable, all messages are stored redundantly across multiple servers and data centres.

- Simple – Developers can get started with Amazon SQS by using only three APIs: SendMessage, ReceiveMessage, and DeleteMessage. Additional APIs are available to provide advanced functionality.
- Scalable – Amazon SQS was designed to enable an unlimited number of messaging services to read and write an unlimited number of messages at any time.
- Secure – Authentication mechanisms are provided to ensure that messages stored in Amazon SQS queues are secured against unauthorised access.
- Inexpensive – No up-front or fixed expenses. The only costs of sending messages through Amazon SQS are small per-request handling fees and data transfer fees.

2.15 Amazon CloudWatch

Amazon CloudWatch is a monitoring and observability service built for DevOps engineers, developers, site reliability engineers (SREs), and IT managers. CloudWatch provides you with data and actionable insights to monitor your applications, respond to system-wide performance changes, optimize resource utilization, and get a unified view of operational health. CloudWatch collects monitoring and operational data in the form of logs, metrics, and events, providing you with a unified view of AWS resources, applications, and services that run on AWS and on-premises servers. You can use CloudWatch to detect anomalous behaviour in your environments, set alarms, visualize logs and metrics side by side, take automated actions, troubleshoot issues, and discover insights to keep your applications running smoothly.

Top Features include

- Collect and store logs from your resources, applications, and services in near real-time.
- Collect default metrics from more than 70 AWS services, such as Amazon EC2, Amazon DynamoDB, Amazon S3, Amazon ECS, AWS Lambda, and Amazon API Gateway, without any action on your part. Collect custom metrics from your own applications to monitor operational performance, troubleshoot issues, and spot trends.
- Amazon CloudWatch Anomaly Detection applies machine-learning algorithms to continuously analyse data of a metric and identify anomalous behaviour.
- Dashboards enable you to create re-usable graphs and visualize your cloud resources and applications in a unified view.
- Integrated with AWS Identity and Access Management (<https://aws.amazon.com/iam/>) (IAM) so that you can control which users and resources have permission to access your data and how they can access it.

Top benefits include

- Observability on a single platform across applications and infrastructure - Modern applications such as those running on microservices architectures generate large volumes of data in the form of metrics, logs, and events. Amazon CloudWatch enables you to collect, access, and correlate this data on a single platform from across all your AWS resources, applications, and services that run on AWS and on-premises servers, helping you break down data silos so you can easily gain system-wide visibility and quickly resolve issues.

- Easiest way to collect metrics in AWS and on-premises - Monitoring your AWS resources and applications is easy with CloudWatch. It natively integrates with more than 70 AWS services such as Amazon EC2, Amazon DynamoDB, Amazon S3, Amazon ECS, Amazon EKS, and AWS Lambda, and automatically publishes detailed 1-minute metrics and custom metrics with up to 1-second granularity so you can dive deep into your logs for additional context. You can also use CloudWatch in hybrid cloud architectures by using the CloudWatch Agent or API to monitor your on-premises resources.
- Improve operational performance and resource optimization - Amazon CloudWatch enables you to set alarms and automate actions based on either predefined thresholds, or on machine learning algorithms that identify anomalous behaviour in your metrics. For example, it can start Amazon EC2 Auto Scaling automatically, or stop an instance to reduce billing overages. You can also use CloudWatch Events for serverless to trigger workflows with services like AWS Lambda, Amazon SNS, and AWS CloudFormation.
- Get operational visibility and insight - To optimize performance and resource utilization, you need a unified operational view, real-time granular data, and historical reference. CloudWatch provides automatic dashboards, data with 1-second granularity, and up to 15 months of metrics storage and retention. You can also perform metric math on your data to derive operational and utilization insights; for example, you can aggregate usage across an entire fleet of EC2 instances.
- Derive actionable insights from logs - CloudWatch enables you to explore, analyse, and visualize your logs so you can troubleshoot operational problems with ease. With CloudWatch Logs Insights, you only pay for the queries you run. It scales with your log volume and query complexity giving you answers in seconds. In addition, you can publish log-based metrics, create alarms, and correlate logs and metrics together in CloudWatch Dashboards for complete operational visibility.
- Collect and aggregate container metrics and logs - Container Insights simplifies the collection and aggregation of curated metrics and container ecosystem logs. It collects compute performance metrics such as CPU, memory, network, and disk information from each container as performance events and automatically generates custom metrics used for monitoring and alarming. The performance events are ingested as CloudWatch Logs with metadata about the running environment such as the Amazon EC2 instance ID, Service, Amazon EBS volume mount and ID, etc., to simplify monitoring and troubleshooting. CloudWatch custom metrics are automatically extracted from these ingested logs and can be further analysed using CloudWatch Logs Insights' advanced query language. Container Insights also provides an option to collect application logs (stdout/stderr), custom logs, predefined Amazon EC2 instance logs, Amazon EKS/k8s data plane logs and Amazon EKS control plane logs (<https://docs.aws.amazon.com/eks/latest/userguide/control-plane-logs.html>).
- High resolution alarms - Amazon CloudWatch alarms allow you to set a threshold on metrics and trigger an action. You can create high-resolution alarms, set a percentile as the statistic, and either specify an action or ignore as appropriate. For example, you can create alarms on Amazon EC2 metrics, set notifications, and take one or more actions to detect and shut down unused or underutilized instances. Real-time alarming on metrics and events enables you to minimize downtime and potential business impact.
- Anomaly Detection - Amazon CloudWatch Anomaly Detection applies machine-learning algorithms to continuously analyse data of a metric and identify anomalous

behaviour. It allows you to create alarms that auto-adjust thresholds based on natural metric patterns, such as time of day, day of week seasonality, or changing trends. You can also visualize metrics with anomaly detection bands on dashboards. This enables you to monitor, isolate, and troubleshoot unexpected changes in your metrics.

- Automate response to operational changes with CloudWatch Events - CloudWatch Events provides a near real-time stream of system events that describe changes to your AWS resources. It allows you to respond quickly to operational changes and take corrective action. You simply write rules to indicate which events are of interest to your application and what automated actions to take when a rule matches an event.
- Compliance and Security - Amazon CloudWatch is integrated with AWS Identity and Access Management (<https://aws.amazon.com/iam/>) (IAM) so that you can control which users and resources have permission to access your data and how they can access it. Amazon CloudWatch Logs is also PCI and FedRamp compliant. Data is encrypted at rest and during transfer. You can also use AWS KMS encryption to encrypt your log groups for added compliance and security.

2.16 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 Cloud 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 visualise your templates as diagrams and edit them using a drag-and-drop interface with the AWS CloudFormation Designer.

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

Top benefits include:

- Supports a Wide Range of AWS Resources - AWS CloudFormation supports a wide range of AWS resources, allowing you to build a highly available, reliable, and scalable AWS infrastructure for your application needs.
- Easy to Use – AWS CloudFormation makes it easy to organise and deploy a collection of AWS resources and lets you describe any dependencies or special parameters to pass in at runtime. You can use one of the many CloudFormation sample templates—either verbatim or as a starting point.
- Declarative and Flexible – To create the infrastructure you want, you enumerate what AWS resources, configuration values, and interconnections you need in a template and then let AWS CloudFormation do the rest with a few simple clicks in the AWS Management Console, one command by using the AWS CLI, or a single

request by calling the APIs. You won't have to recall the details of how to create and interconnect the respective AWS resources via their service APIs; AWS CloudFormation does this for you. You also don't need to write a template from scratch if you start with one of the many sample templates that come with AWS CloudFormation.

- **Infrastructure as Code** – A template can be used repeatedly to create identical copies of the same stack (or to use as a foundation to start a new stack). You can capture and control region-specific infrastructure variations such as Amazon EC2 AMLs, as well as Amazon EBS and Amazon RDS snapshot names.
- **Templates are simple JSON-formatted text files** that can be placed under your normal source control mechanisms, stored in private or public locations such as Amazon S3, and exchanged via email. With AWS CloudFormation, you can “open the hood,” to see exactly which AWS resources make up a stack. You retain full control and have the ability to modify any of the AWS resources created as part of a stack.
- **Customised Via Parameters** – You can use parameters to customise aspects of your template at run time when the stack is built. For example, you can pass the Amazon RDS database size, Amazon EC2 instance types, database, and web server port numbers to AWS CloudFormation when you create a stack. You can also use a parameterised template to create multiple stacks that may differ in a controlled way. For example, your Amazon EC2 instance types, Amazon CloudWatch alarm thresholds, and Amazon RDS read-replica settings may differ among AWS Regions if you receive more customer traffic in the US than in Europe. You can use template parameters to tune the settings and thresholds in each region separately and still be sure that the application is deployed consistently across the regions.
- **Visualise and Edit with Drag-and-Drop Interface** – AWS CloudFormation Designer provides a visual diagram of your template with icons representing your AWS resources and arrows showing their relationships. You can build and edit templates using the drag-and-drop interface, then edit the template details using the integrated JSON text editor. AWS CloudFormation Designer allows you to spend more time designing your AWS infrastructure and less time manually coding your templates.
- **Integration Ready** – You can integrate AWS CloudFormation with the development and management tools of your choice. AWS CloudFormation publishes progress events through the Amazon SNS. With Amazon SNS you can track stack creation and deletion progress via email and integrate with other processes programmatically.

2.17 Amazon Quicksight, Amazon Kinesis

Amazon QuickSight

Amazon QuickSight is a very fast, cloud-powered Business Intelligence (BI) service that makes it easy for all employees to build visualizations, perform analysis, and quickly get business insights from their data. Amazon QuickSight uses a new, Super-fast, Parallel, In-memory Calculation Engine (“SPICE”) to perform advanced calculations and render visualizations rapidly. Amazon QuickSight integrates automatically with AWS data services, enables organisations to scale to hundreds of thousands of users, and delivers fast and

responsive query performance to them via SPICE's query engine. At one tenth the cost of traditional solutions, Amazon QuickSight enables you to deliver rich BI functionality to everyone in your organisation.

Top benefits include:

- **Get Started Quickly** – No need to spend months building complex data models and invest in complex and costly software or hardware to generate the first report. Simply log in to Amazon QuickSight, point to a data source or upload a file, and begin visualizing your data with actionable insights in about a minute.
- **Access Data from Multiple Sources** – Easily connect Amazon QuickSight to
- AWS data services, including Amazon Redshift, Amazon RDS, Amazon Aurora,
- Amazon EMR, Amazon DynamoDB, Amazon S3, and Amazon Kinesis; upload CSV, TSV, and spreadsheet files; or connect to third-party data sources such as Salesforce.
- **Easy to Use** – Amazon QuickSight automatically infers data types and relationships and provides suggestions for the best possible visualizations, optimised for your data, to help you get quick, actionable business insights.
- **Get Answers Fast** – Amazon QuickSight uses SPICE to generate answers on large data sets.
- **Tell a Story with Your Data** – Securely share your analysis with others in your organisation by building interactive stories for collaboration using the StoryBoard and annotations. Recipients can further explore the data and respond back with their insights and knowledge, making the whole organisation efficient and effective.
- **Use AWS Partner BI Solutions with Amazon QuickSight** – Amazon QuickSight provides partners a simple SQL-like interface to query the data stored in SPICE so that you can continue using your existing BI tools from AWS BI partners while benefiting from the faster performance delivered by SPICE.

Amazon Kinesis

Amazon Kinesis is a platform for streaming data on AWS, offering powerful services to make it easy to load and analyse streaming data and also providing the ability for you to build custom streaming data applications for specialised needs. Web applications, mobile devices, wearables, industrial sensors, and many software applications and services can generate staggering amounts of streaming data—sometimes TBs 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 consists of:

- Amazon Kinesis
- Amazon Kinesis Data Analytics
- Amazon Kinesis Data Firehose
- Amazon Kinesis Data Streams
- Amazon Kinesis Video Streams

Amazon Kinesis Firehose is the easiest way to load streaming data into AWS. It can capture and automatically load streaming data into Amazon S3, Amazon Redshift, and Amazon

Elasticsearch Service, 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. With Amazon Kinesis Firehose, you only pay for the amount of data you transmit through the service. There is no minimum fee or setup cost.

Amazon Kinesis Streams enables you to build custom applications that process or analyse streaming data for specialised 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 Cloud services such as Amazon S3, Amazon Redshift, Amazon EMR, and AWS Lambda.

2.18 AWS Glue

AWS Glue is a fully managed extract, transform, and load (ETL) service that makes it easy for customers to prepare and load their data for analytics. You can create and run an ETL job with a few clicks in the AWS Management Console. You simply point AWS Glue to your data stored in AWS, and AWS Glue discovers your data and stores the associated metadata (e.g., table definition and schema) in the AWS Glue Data Catalog. Once catalogued, your data is immediately searchable, queryable, and available for ETL.

The top features include:

- Integrated data catalogue – The AWS Glue Data Catalog is your persistent metadata store for all your data assets, regardless of where they are located.
- Automatic schema discovery – AWS Glue crawlers connect to your source or target data store, progress through a prioritised list of classifiers to determine the schema for your data, and then create metadata in your AWS Glue Data Catalog. The metadata is stored in tables in your data catalogue and used in the authoring process of your ETL jobs.
- Code generation – AWS Glue automatically generates the code to ETL your data. Simply point AWS Glue to your data source and target, and AWS Glue creates ETL scripts to transform, flatten, and enrich your data. The code is generated in Scala or Python and written for Apache Spark.
- Developer endpoints – If you choose to interactively develop your ETL code, AWS Glue provides development endpoints for you to edit, debug, and test the code it generates for you. You can use your favourite IDE or notebook. You can write custom readers, writers, or transformations and import them into your AWS Glue ETL jobs as custom libraries. You can also use and share code with other developers in our GitHub repository.

- Flexible job scheduler – AWS Glue jobs can be invoked on a schedule, on demand, or based on an event. You can start multiple jobs in parallel or specify dependencies across jobs to build complex ETL pipelines.

2.19 AWS Step Functions

AWS Step Functions makes it easy to coordinate the components of distributed applications and microservices using visual workflows. Building applications from individual components that each perform a discrete function lets you scale and change applications quickly. AWS Step Functions is a reliable way to coordinate components and step through the functions of your application. AWS Step Functions provides a graphical console to arrange and visualise the components of your application as a series of steps. This makes it simple to build and run multi-step applications. AWS Step Functions automatically triggers and tracks each step, and retries when there are errors, so your application executes in order and as expected. AWS Step Functions logs the state of each step, so when things do go wrong, you can diagnose and debug problems quickly. You can change and add steps without even writing code, so you can easily evolve your application and innovate faster.

AWS Step Functions manages the operations and underlying infrastructure for you to help ensure your application is available at any scale.

Top benefits include:

- Productivity: Build Applications Quickly – AWS Step Functions includes a visual console and blueprints for commonly-used workflows that make it easy to coordinate the components of distributed applications into parallel and/or sequential steps. You can build applications in a matter of minutes, and then visualise and track the execution of each step to help ensure the application is operating as intended
- Resilience: Scale and Recover Reliably – AWS Step Functions automatically triggers each step so your application executes in order and as expected. It can handle millions of steps simultaneously to help ensure your application is available as demand increases. AWS Step Functions tracks the state of each step and handles errors with built-in retry and fall-back, whether the step takes seconds or months to complete.
- Agility: Evolve Applications Easily – AWS Step Functions makes it easy to change workflows and edit the sequence of steps without revising the entire application. You can re-use components and steps without even changing their code to experiment and innovate faster. Your workflow can support thousands of individual components and steps, so you can freely build increasingly complex applications.

2.20 Amazon Athena

Amazon Athena is an interactive query service that makes it easy to analyse data in Amazon S3 using standard SQL. Amazon Athena is serverless, so there is no infrastructure to manage, and you pay only for the queries that you run.

Amazon Athena is easy to use. Simply point to your data in Amazon S3, define the schema, and start querying using standard SQL. Most results are delivered within seconds. With Amazon Athena, there's no need for complex ETL jobs to prepare your data for analysis. This makes it easy for anyone with SQL skills to quickly analyse largescale datasets.

Top benefits include:

- **Start Querying Instantly: Serverless. No ETL** – Amazon Athena is serverless. You can quickly query your data without having to setup and manage any servers or data warehouses. Just point to your data in Amazon S3, define the schema, and start querying using the built-in query editor. Amazon Athena allows you to tap into all your data in S3 without the need to set up complex processes to extract, transform, and load the data (ETL).
- **Pay Per Query: Only pay for data scanned** – With Amazon Athena, you pay only for the queries that you run. You can save from 30% to 90% on your per query costs and get better performance by compressing, partitioning, and converting your data into columnar formats. Amazon Athena queries data directly in Amazon S3. There are no additional storage charges beyond S3.
- **Open. Powerful. Standard: Built on Presto. Runs standard SQL** – Amazon Athena uses Presto with ANSI SQL support and works with a variety of standard data formats, including CSV, JSON, ORC, Avro, and Parquet. Amazon Athena is ideal for quick, ad-hoc querying but it can also handle complex analysis, including large joins, window functions, and arrays. Amazon Athena is highly available; and executes queries using compute resources across multiple facilities and multiple devices in each facility. Amazon Athena uses Amazon S3 as its underlying data store, making your data highly available and durable.
- **Fast. Really Fast: Interactive performance even for large datasets** – With Amazon Athena, you don't have to worry about having enough compute resources to get fast, interactive query performance. Amazon Athena automatically executes queries in parallel, so most results come back within seconds.

2.21 Lake Formation

AWS Lake Formation is a service that makes it easy to set up a secure data lake in days. A data lake is a centralized, curated, and secured repository that stores all your data, both in its original form and prepared for analysis. A data lake enables you to break down data silos and combine different types of analytics to gain insights and guide better business decisions.

However, setting up and managing data lakes today involves a lot of manual, complicated, and time-consuming tasks. This work includes loading data from diverse sources, monitoring those data flows, setting up partitions, turning on encryption and managing keys, defining transformation jobs and monitoring their operation, re-organizing data into a columnar format, configuring access control settings, deduplicating redundant data, matching linked records, granting access to data sets, and auditing access over time.

Creating a data lake with Lake Formation is as simple as defining data sources and what data access and security policies you want to apply. Lake Formation then helps you collect and catalog data from databases and object storage, move the data into your new Amazon S3 data lake, clean and classify your data using machine learning algorithms, and secure access to your sensitive data. Your users can access a centralized data catalog which describes available data sets and their appropriate usage. Your users then leverage these data sets with their choice of analytics and machine learning services, like Amazon Redshift, Amazon Athena, and (in beta) Amazon EMR for Apache Spark. Lake Formation builds on the capabilities available in AWS Glue.

Benefits:

Build data lakes quickly

With Lake Formation, you can move, store, catalog, and clean your data faster. You simply point Lake Formation at your data sources, and Lake Formation crawls those sources and moves the data into your new Amazon S3 data lake. Lake Formation organizes data in S3 around frequently used query terms and into right-sized chunks to increase efficiency. Lake Formation also changes data into formats like Apache Parquet and ORC for faster analytics. In addition, Lake Formation has built-in machine learning to deduplicate and find matching records (two entries that refer to the same thing) to increase data quality.

Simplify security management

You can use Lake Formation to centrally define security, governance, and auditing policies in one place, versus doing these tasks per service, and then enforce those policies for your users across their analytics applications. Your policies are consistently implemented, eliminating the need to manually configure them across security services like AWS Identity and Access Management and AWS Key Management Service, storage services like S3, and analytics and machine learning services like Redshift, Athena, and (in beta) EMR for Apache Spark. This reduces the effort in configuring policies across services and provides consistent enforcement and compliance.

Provide self-service access to data

With Lake Formation you build a data catalog that describes the different data sets that are available along with which groups of users have access to each. This makes your users more productive by helping them find the right data set to analyze. By providing a catalog of your data with consistent security enforcement, Lake Formation makes it easier for your analysts and data scientists to use their preferred analytics service.

They can use EMR for Apache Spark (in beta), Redshift, or Athena on diverse data sets now housed in a single data lake. Users can also combine these services without having to move data between silos.

3 Where We Have Delivered Before

Agile Solutions has considerable experience in the delivery of AWS solutions in both the public and private sectors. The following are some recent successes we enabled our customers to achieve.

	Context and approach	Customer value we helped achieve
Multinational Software Company	The customer wished to migrate to a new HR system. Agile designed, developed and delivered a data migration framework based upon AWS and Informatica Cloud. Our framework utilised S3 for staging data, EC2 for Informatica compute and RDS as the endpoint.	The framework allowed the business to upload HR data onto S3, where the Informatica connection listened for files in real time and triggered a transformation process of the data and inputs into RDS – ready for migration. This data pipeline enabled the customer to migrate on-time and within budget.
UK Central Government	Working within a DevOps environment to support delivery pipelines on AWS	Provides 2 nd and 3 rd line support whilst maintaining strict security standards

4 About Agile Solutions

With a history spanning 18 years, Agile Solutions is a UK based specialist Information Management Services Company with a reputation for rapid and agile delivery of large scale analytical and data governance solutions for leading organisations.

We bring together a portfolio of data related services with a wealth of relevant domain experience and technologies related to analytics, data governance, data warehousing and business intelligence.

Agile Solutions focusses on delivering holistic data-centric solutions to both the public and private sectors. We operate throughout UK & Europe with offices in Glasgow, Manchester and Milton Keynes.

We are Information Management Specialists who are:

- Focused on business agility
- Business value orientated
- Dedicated to proven, measurable customer success
- Technology and methodology agnostic
- Data centric
- UK based with three office locations for collaborative delivery

Combining our objective independent advice, the innovative Agile Information Management Framework (AIM) and dedicated team resourcing via our Agile Competency Centres (ACCs) and augmented by our Education Services we are able to respond quickly and effectively for our clients, which ensures maximum value is realised from all their data and technology investments



Get In Touch

Agile Solutions UK Locations

454 Exchange House

CBX1 Midsummer Boulevard

Milton Keynes

MK9 2EA

Ocean Chambers

190 West George Street

Glasgow

G2 2NR

9th Floor Lowry House

17 Marble Street

Manchester

M2 3AW

Telephone

01908 010618

e-mail

GCloud@agilesolutions.co.uk



www.agilesolutions.co.uk

©Agile Solutions 2020