

Performance Efficiency Pillar

AWS Well-Architected Framework

November 2017



© 2017, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Notices

This document is provided for informational purposes only. It represents AWS's current product offerings and practices as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS's products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

Contents

Introduction	1
Performance Efficiency	1
Design Principles	2
Definition	2
Selection	3
Compute	4
Storage	8
Database	11
Network	19
Review	23
Benchmarking	25
Load Testing	26
Monitoring	27
Active and Passive	28
Phases	29
Trade-Offs	30
Caching	31
Partitioning or Sharding	33
Compression	34
Buffering	35
Conclusion	37
Contributors	37
Further Reading	38
Document Revisions	38

Abstract

The focus of this paper is the Performance Efficiency pillar of the Amazon Web Services (AWS) [Well-Architected Framework](#). It provides guidance to help you apply best practices in the design, delivery, and maintenance of AWS environments.

This whitepaper covers the best practices for managing production environments. Designing and managing non-production environments and processes such as Continuous Integration or Delivery are out of the scope of this document.

Introduction

The [AWS Well-Architected Framework](#) helps you understand the pros and cons of decisions you make while building systems on AWS.¹ By using the Framework you will learn architectural best practices for designing and operating reliable, secure, efficient, and cost-effective systems in the cloud. It provides a way for you to consistently measure your architectures against best practices and identify areas for improvement. We believe that having well-architected systems greatly increases the likelihood of business success.

The framework is based on five pillars:

- Operational Excellence
- Security
- Reliability
- Performance Efficiency
- Cost Optimization

This paper focuses on the performance efficiency pillar and how to apply it to your solutions. In traditional, on-premises environments, achieving high and lasting performance can be challenging. By adopting the practices in this paper, you can build architectures that are efficient and deliver sustained performance over time.

This paper is intended for those in technology roles, such as chief technology officers (CTOs), architects, developers, and operations team members. After reading this paper, you'll understand AWS best practices and strategies to use when designing a performant cloud architecture. This paper doesn't provide implementation details or architectural patterns. However, it does include references to appropriate resources for this information.

Performance Efficiency

The performance efficiency pillar focuses on the efficient use of computing resources to meet requirements and how to maintain that efficiency as demand changes and technologies evolve. This paper provides in-depth, best-practice guidance for architecting for performance efficiency on AWS.

Design Principles

In the cloud, there are a number of principles that can help you achieve performance efficiency.

- **Democratize advanced technologies:** Technologies that are difficult to implement can become easier to consume by pushing that knowledge and complexity into the cloud vendor's domain. Rather than having your IT team learn how to host and run a new technology, they can simply consume it as a service. For example, NoSQL databases, media transcoding, and machine learning are all technologies that require expertise that is not evenly dispersed across the technical community. In the cloud, these technologies become services that your team can consume. This lets your team focus on product development rather than resource provisioning and management.
- **Go global in minutes:** Easily deploy your system in multiple AWS Regions around the world with just a few clicks. This allows you to provide lower latency and a better experience for your customers at minimal cost.
- **Use serverless architectures:** In the cloud, serverless architectures remove the need for you to run and maintain servers to carry out traditional compute activities. For example, storage services can act as static websites, removing the need for web servers, and event services can host your code for you. This removes the operational burden of managing these servers and can also lower transactional costs because these managed services operate at cloud scale.
- **Experiment more often:** With virtual and automatable resources, you can quickly carry out comparative testing using different types of instances, storage, or configurations.
- **Mechanical sympathy:** Use the technology approach that aligns best to what you are trying to achieve. For example, consider data access patterns when you select database or storage approaches.

Definition

Performance efficiency in the cloud is composed of four areas:

- Selection

- Review
- Monitoring
- Trade-offs

Take a data-driven approach to building a high-performance architecture. Gather data on all aspects of the architecture, from the high-level design to the selection and configuration of resource types. By reviewing your choices on a regular basis, you ensure that you are taking advantage of the continually evolving AWS Cloud. Monitoring will ensure that you are aware of any deviance from expected performance so you can take action on it. Finally, you can make tradeoffs in your architecture to improve performance, such as using compression or caching, or relaxing consistency requirements.

Selection

The optimal solution for a particular system will vary based on the kind of workload you have and often combines multiple approaches. Well-architected systems use multiple solutions and enable different features to improve performance.

With AWS, resources are virtualized and are available in a number of different types and configurations. This makes it easier to find an approach that closely aligns with your needs. You can also find options that are not easily achievable with on-premises infrastructure. For example, a managed service such as Amazon DynamoDB provides a fully managed NoSQL database with single-digit millisecond latency at any scale.

You can rely on AWS Solutions Architects, AWS Reference Architectures, and AWS Partners to help you choose an architecture based on what you have learned. However, you need to perform benchmarking or load testing to optimize your architecture. After you have identified your architectural approach, you should use a data-driven process to refine your selection of resource types and configuration options. You can obtain this data using benchmarking and load testing. For more information, see the [Review section](#) later in this whitepaper.

Your architecture can combine a number of different architectural approaches (for example, request-response, event driven; extract, transform, load (ETL); or pipeline). Your architecture's implementation will use the AWS services that can

optimize your architecture's performance. In the following sections, we look at the four main resource types that you should consider: compute, storage, database, and network.

Compute

The optimal compute solution for a particular system varies based on application design, usage patterns, and configuration settings. Architectures can use different compute solutions for various components and enable different features to improve performance. Selecting the wrong compute solution for an architecture can lead to lower performance efficiency.

When architecting your use of compute, you should take advantage of elasticity mechanisms that can ensure that you have sufficient capacity to sustain performance as demand changes. In AWS, compute is available in three forms: instances, containers, and functions. Choose which form of compute to use for each component in your architecture. Instances are generally the default option. Using containers can improve the utilization of instances. Functions are well suited to event-driven or highly parallelizable tasks.

Instances

Instances are virtualized servers and, therefore, you can change their capabilities with the click of a button or an API call. In the cloud, resource decisions are no longer fixed, so you can experiment with different server types.

Amazon Elastic Compute Cloud (Amazon EC2) virtual server instances come in different families and sizes. They offer a wide variety of capabilities, including solid-state drives (SSDs) and graphics processing units (GPUs). When you launch an EC2 instance, the instance type that you specify determines the hardware of the host computer used for your instance. Each instance type offers different compute, memory, and storage capabilities. Instance types are grouped in instance families based on these capabilities.

When selecting instance families and types also consider the configuration options that your workload needs:

- **Graphics Processing Units (GPU)**. Using general purpose computing on GPUs (GPGPU), you can build applications that benefit from the high degree of parallelism that GPUs provide by leveraging platforms such as

CUDA in the development process. Also, if your application requires 3D rendering or video compression, GPUs enable hardware-accelerated computation and encoding, making your application more efficient.

- **Field Programmable Gate Arrays (FPGA).** Using FPGAs, you can optimize your workloads by having custom hardware-accelerated execution for your most demanding workloads. You can define your algorithms by leveraging supported general programming languages such as C or Go, or hardware-oriented languages such as Verilog or VHDL.
- **Burstable instance families.** Burstable instances are designed to provide moderate baseline performance and the capability to burst to significantly higher performance when it's required by your workload. These instances are intended for workloads that don't use the full CPU often or consistently, but occasionally need to burst. They are well suited for general-purpose workloads, such as web servers, developer environments, and small databases. These instances provide CPU credits that can be consumed when the instance needs to provide performance. Credits accumulate when the instance doesn't need them.
- **Advanced computing features.** Amazon EC2 gives you access to advanced computing features such as managing the C-state and P-state registers and controlling the turbo-boost of processors. Access to co-processors allows cryptography operations offloading through AES-NI, or advanced computation through AVX extensions.

You should use data to select the optimal EC2 instance type for your workload, ensure you have the correct networking and storage options, and consider operating system settings that can improve the performance for your workload.

Containers

Containers are a method of operating system virtualization that allow you to run an application and its dependencies in resource-isolated processes.

Amazon EC2 Container Service (Amazon ECS) allows automated execution and management of containers on a cluster of EC2 instances. You can create services in Amazon ECS that will be served by containers. By using application Auto Scaling, you can define automated, metrics-based scaling for your services so that the number of containers supporting your service grows as the service needs grow.

Amazon ECS leverages Auto Scaling groups, which enable you to scale the Amazon ECS cluster by adding EC2 instances. This way you ensure that you always have the appropriate underlying server capacity to host your containers.

Amazon ECS is integrated with Elastic Load Balancing (ELB) to support load balancing your services in the container fleet dynamically. By creating an Application Load Balancer instance, your containers will automatically register themselves to the load balancer when they scale. With these features, you can host multiple containers serving the same service on a single Amazon EC2 instance, increasing the scalability of your application.

When using containers, you should use data to select the optimal type for your workload just as you use data to select your EC2 instance types. You should also consider container configuration options such as memory, CPU, and tenancy configuration.

Functions

Functions abstract the execution environment from the code you want to execute. The use of functions allows you to execute code or provide a service without running or managing an instance.

AWS Lambda lets you run code without provisioning or managing servers or containers. You just upload your code and AWS Lambda takes care of everything required to run and scale your code. You can set up your code to automatically trigger from other AWS services; you can call it directly; or you can use it with Amazon API Gateway.

Amazon API Gateway is a fully managed service that makes it easy for developers to create, publish, maintain, monitor, and secure APIs at any scale. You can create an API that acts as a “front door” to your Lambda function. API Gateway handles all the tasks involved in accepting and processing up to hundreds of thousands of concurrent API calls, including traffic management, authorization and access control, monitoring, and API version management.

To deliver optimal performance with AWS Lambda, you should choose the amount of memory you want for your function. You are then allocated proportional CPU power and other resources. For example, choosing 256 MB of memory allocates approximately twice as much CPU power to your Lambda

function as requesting 128 MB of memory. You can also control the amount of time each function is allowed to run (up to a maximum of 300 seconds).

Elasticity

Elasticity allows you to match the supply of resources you have against demand for them. Instances, containers, and functions all provide mechanisms for elasticity either in combination with Auto Scaling or as a feature of the service itself.

Optimally matching supply to demand delivers the lowest cost for a system, but you also need to plan for sufficient supply to allow for provisioning time and individual resource failures. Demand can be fixed or variable, requiring metrics and automation to ensure that management does not become a burdensome and disproportionately large cost in itself.

In AWS, you can use a number of different approaches to match supply with demand. The [Cost Optimization Pillar of the AWS Well-Architected Framework](#) whitepaper describes how to use each of these approaches:²

- Demand-based approach
- Buffer-based approach
- Time-based approach

Key AWS Services

The key AWS service for elastic compute solutions is Auto Scaling because you use it to match the supply of your resources against demand for them. Instances, containers, and functions all provide mechanisms for elasticity either in combination with Auto Scaling or as a feature of the service itself.

Resources

Refer to the following resources to learn more about AWS best practices for compute.

Documentation

- Instances: [Instance Types](#)³
- Containers: [Amazon ECS Container Instances](#)⁴

- Functions: [Compute Requirements – Lambda Function Configuration](#)⁵

Storage

The optimal storage solution for a particular system varies based on the kind of access method (block, file, or object) you use, patterns of access (random or sequential), throughput required, frequency of access (online, offline, archival), frequency of update (WORM, dynamic), and availability and durability constraints.

In AWS storage is virtualized, and there are a number of different storage types. This makes it easier to match your storage methods more closely with your needs, and it also offers storage options that are not easily achievable with on-premises infrastructure. For example, Amazon Simple Storage Service (Amazon S3) is designed for 11 nines of durability. When you use cloud storage, you can change from using magnetic hard disk drives (HDDs) to solid state drives (SSDs), and easily move virtual drives from one instance to another in seconds.

Characteristics

When you select a storage solution, you should consider the different characteristics that you require, such as ability to share, file size, cache size, latency, throughput, and persistence of data. Then match the requirements you want to the AWS service that best fits your needs: Amazon S3, Amazon Glacier, Amazon Elastic Block Store (Amazon EBS), Amazon Elastic File System (Amazon EFS), or Amazon EC2 instance store.

Performance can be measured by looking at throughput, input/output operations per second (IOPS), and latency. Understanding the relationship between those measurements will help you select the most appropriate storage solution.

Storage	Services	Latency	Throughput	Shareable
Block	EBS, EC2 instance store	Lowest, consistent	Single	Mounted on single instance, copies via snapshots
File system	EFS	Low, consistent	Multiple	Many clients
Object	S3	Low-latency	Web scale	Many clients

Storage	Services	Latency	Throughput	Shareable
Archival	Glacier	Minutes to hours	High	No

From a latency perspective, if your data is only accessed by one instance then you should use block storage, such as Amazon EBS with Provisioned IOPS. Distributed file systems such as Amazon EFS generally have a small latency overhead for each file operation, so they should be used where multiple instances need access.

Amazon S3 has features that can reduce latency and increase throughput. You can use cross-region replication (CRR) to provide lower-latency data access to different geographic regions.

From a throughput perspective, Amazon EFS supports highly parallelized workloads (for example, using concurrent operations from multiple threads and multiple EC2 instances), which enables high levels of aggregate throughput and operations per second. For Amazon EFS, use a benchmark or load test to select the appropriate performance mode.

Key AWS Services

The key AWS service for storage is Amazon S3, which provides secure, durable, highly scalable [cloud storage](#).⁶ The following services and features are also important:

- **Amazon EBS** provides persistent block storage volumes for use with EC2 instances.
- **Amazon EFS** provides simple, scalable file storage for use with EC2 instances.
- **Amazon EC2 instance store** provides temporary block-level storage for use with EC2 instances.
- **Amazon Glacier** provides an archival platform for long-term data storage and is integrated with Amazon S3 for tiered storage.

Resources

Refer to the following resources to learn more about AWS best practices for storage.

Documentation

- Amazon S3: [Request Rate and Performance Considerations](#)⁷
- Amazon Glacier: [Amazon Glacier Documentation](#)⁸
- Amazon EFS: [Amazon EFS Performance](#)⁹
- Amazon EBS: [I/O Characteristics](#)¹⁰

Videos

- [Amazon EBS Design for Performance](#)¹¹

Configuration Options

Storage solutions generally have configuration options that allow you to optimize for the type of workload.

Amazon EBS provides a range of options that allow you to optimize storage performance and cost for your workload. These options are divided into two major categories: SSD-backed storage for transactional workloads, such as databases and boot volumes (performance depends primarily on IOPS), and HDD-backed storage for throughput-intensive workloads such as MapReduce and log processing (performance depends primarily on MB/s).

SSD-backed volumes include the highest performance Provisioned IOPS SSD for latency-sensitive transactional workloads and General Purpose SSD that balance price and performance for a wide variety of transactional data.

Amazon S3 transfer acceleration enables fast transfer of files over long distances between your client and your S3 bucket. Transfer acceleration leverages Amazon CloudFront globally distributed edge locations to route data over an optimized network path. For a workload in an S3 bucket that has intensive GET requests you should use Amazon S3 with CloudFront. When uploading large files, you should use multi-part uploads with multiple parts uploading at once to help maximize network throughput.

Access Patterns

How you access data will affect how the storage solution performs. Select the storage solution that aligns best to your access patterns, or consider changing

your access patterns to align with the storage solution to maximize performance.

Creating a RAID 0 (zero) array allows you to achieve a higher level of performance for a file system than you can provision on a single volume. Consider using RAID 0 when I/O performance is more important than fault tolerance. For example, you could use it with a heavily used database where data replication is already set up separately.

Amazon EBS HDD-backed volumes include Throughput-Optimized HDD for throughput-intensive sequential workloads and the lowest cost Cold HDD for less frequently accessed data.

You can further optimize for how you use storage systems by using techniques covered in the [Partitioning or Sharding section](#) of the Trade-Offs topic in this whitepaper.

Key AWS Services

The key AWS services for storage solutions are Amazon EBS, Amazon S3, and Amazon EFS which provide block, object, and file system storage solutions. These services have configuration options that further allow you to optimize your storage solution. The access patterns your components have should influence the storage solution you select.

Resources

Refer to the following resources to learn more about AWS best practices for storage.

Documentation

- [Storage](#)¹²
- [EBS Storage Types](#)¹³

Database

The optimal database solution for a particular system can vary based on your requirements for availability, consistency, partition tolerance, latency, durability, scalability, and query capability. Many systems use different database solutions for different subsystems and enable different features to improve performance. Selecting the wrong database solution and features for a system can lead to lower performance efficiency.



Although a workload's database approach (for example, relational database management system or RDBMS, NoSQL, etc.) has significant impact on performance efficiency, it is often an area that is chosen according to organizational defaults rather than through using a data-driven approach. As with storage, it's critical to consider the access patterns of your workload, and also to consider if other non-database solutions could solve the problem more efficiently (for example, using an object store such as Amazon S3, a search engine, or a data warehouse).

In on-premises environments the tendency is to standardize on a common platform for ease of management. However, in AWS there are services for each approach. You should consider using different services for different types of data. You can leverage your own platforms that you deploy on Amazon EC2 or use fully managed services, allowing you to have the benefits of best-of-breed solutions without management effort on your side.

Efficiency Factors

There are many factors to consider when selecting your database approach. For example, your access patterns will help you choose between different types of database technologies. When selecting your database technology, keep in mind these four factors:

- Access patterns
- Characteristics
- Configuration options
- Operational effort

Access Patterns

The way that you access data will affect how the database solution performs. Select the database solution that aligns best to your access patterns, or consider changing your access patterns to align with the storage solution to maximize performance.

Optimize how you use database systems based on your access pattern (for example, indexes, key distribution, partition, or horizontal scaling).

Characteristics

Consider the different characteristics you require (for example, availability, consistency, partition tolerance, latency, durability, scalability, and query capability) so that you can select the appropriate approach to use (for example, relational, NoSQL, warehouse).

Configuration Options

Database solutions generally have configuration options that allow you to optimize for the type of workload. Consider the configuration options for your selected database approach such as storage optimization, database level settings, memory, and cache.

Operational Effort

You should consider the effort required to operate your technology choice. The effort of managing and maintaining a database can vary greatly based on your technology choice and consumption approach. You should also consider the effort required to scale these approaches as your requirements grow. For example, databases that are provided as managed services remove or reduce the effort of managing database servers and their detailed configuration. To further optimize your use of database systems, explore the techniques covered in the [Partitioning or Sharding section](#) of the Trade-Offs topic in this whitepaper.

Database Technologies

We commonly see four technology approaches used to address database requirements:

- Relational Online Transaction Processing (OLTP)
- Non-relational databases (NoSQL)
- Data warehouse and Online Analytical Processing (OLAP)
- Data indexing and searching

Relational Online Transaction Processing (OLTP)

Most traditional applications run on an existing relational database, such as Oracle, Microsoft SQL Server, MySQL, or PostgreSQL. These database engines provide high-performance, relational access to datasets. You can make complex queries that will join and aggregate data from multiple entities. The transactional model allows complex business requests.

You can run these database engines on Amazon EC2 by:

- Installing the software on your own EC2 instances running Microsoft Windows or Linux
- Using pre-configured Amazon Machine Images (AMIs), provided by AWS or by third-parties on the AWS Marketplace, with your database engine installed

There are four distinct tasks for optimizing these databases:

- Optimizing the underlying instance (compute, memory, storage, network) using the best practices presented in this white paper
- Optimizing the operating system settings such as volume management, RAID, block sizes, and settings
- Optimizing the database engine configuration and some database-specific features, such as partitioning
- Optimizing the databases themselves by managing the schemas, indexes, views and database-related options

Amazon Relational Database Service (Amazon RDS) provides a fully managed relational database. By leveraging Amazon RDS, you can keep using your existing database engine (Oracle, Microsoft SQL Server, MySQL, MariaDB, PostgreSQL) and benefit from an automated environment.

Amazon RDS also proposes a cloud-native database platform called Amazon Aurora. Aurora is compatible with MySQL or PostgreSQL, and has an execution and storage environment that is purpose-built for AWS. You will natively benefit from top-tier performance, with a convenient interface that is compatible with MySQL or PostgreSQL, and without licensing constraints.

Optimizing an Amazon RDS database comes down to:

- Selecting the appropriate configuration from the catalog of instance types, and the type of storage required
- Tuning the database engine configuration using the options and parameter groups

- Optimizing the database (schemas, indexes, views, and database-related options)

All other tasks are managed by the service. Backup, patching, and point-in-time recovery are automated.

Some database engines provide options for running read replicas. These replicas allow read-heavy workloads to use multiple servers instead of overloading a central database.

Amazon RDS provides automated read replicas (not available for Oracle or SQL Server) to scale out read performance, SSD-backed storage options, and provisioned IOPS.

Amazon Aurora has a fully redesigned, SSD-based storage layer that scales automatically as your needs grow, and supports a large number of low-latency read replicas for increased read performance.

Non-Relational Databases (NoSQL)

New applications can be built without the constraints of managing a central relational OLTP database system. These systems can leverage scalable distributed systems such as NoSQL databases. These database platforms scale out and in by adding nodes to or removing nodes from the database cluster, making them great candidates for large-scale applications. They also usually allow less structured data structures to be stored natively. This limits the burden of schema management and makes them easy to use by application developers. On the other hand, their scaling nature makes it difficult to run complex relational queries such as joining tables and entities to have an answer. This work is usually deferred to the application layer.

You can run NoSQL databases, such as Cassandra, on EC2 instances, design their storage tier, optimize the operating system, and then run the engine. By leveraging instance storage and EBS volumes, designing your clusters in multiple Availability Zones or Regions, and ensuring you have the right model, you can have high-scale, high-performance platforms for your data needs.

Amazon DynamoDB is a fully managed NoSQL database that provides single-digit millisecond latency at any scale. For internet-scale information, such as user profiles, shopping carts, or any other high volume read and write applications, Amazon DynamoDB can be configured to match your needs. You



can specify the exact amount of read and write operations per second that you expect your table to process. Storage scales up and down automatically based on your usage of the platform.

Amazon DynamoDB has advanced features that increase the performance and ease of use of the platform:

- Amazon DynamoDB Accelerator (DAX) provides a read-through/write-through distributed caching tier in front of the database, supporting the same API as Amazon DynamoDB, but providing sub-millisecond latency for entities that are in the cache.
- Amazon DynamoDB also supports throughput Auto Scaling, based on limits.

Amazon DynamoDB provides indexing features that allow you to efficiently query on any attribute (column). You can also project attributes (copied from the table into the index) to improve performance.

Amazon DynamoDB manages table partitioning for you automatically, adding new partitions if necessary and distributing provisioned throughput capacity evenly across them. To ensure that your load is evenly distributed across partitions, you should design for uniform data access across items in your tables by selecting an appropriate partition key.

Data Warehouse and Online Analytical Processing (OLAP)

OLTP and NoSQL platforms are extremely useful for application deployment, but often have limitations for large scale analytics. These use cases are better served by data warehouse platforms that are made for fast access to structured data. Some advanced technologies are particularly adapted for data warehousing, such as columnar storage of information, or Massively Parallel Processing (MPP) environments.

You can host your existing Oracle or Microsoft SQL Server environment configured for data warehousing workloads. Or you could benefit from specialized data warehousing platforms such as Pivotal Greenplum. Alternatively, you could use cloud-native environments, such as Snowflake, by finding them on the AWS Marketplace.

Amazon Redshift is a managed petabyte-scale data warehouse that allows you to change the number or type of nodes as your performance or capacity needs change. Use Amazon Redshift when you need SQL operations that will scale.

Amazon Redshift allows you to change the number or type of nodes in your data warehouse and scale up all the way to a petabyte or more of compressed user data. Dense compute (DC) nodes allow you to create very high-performance data warehouses using fast CPUs, large amounts of RAM, and SSDs.

Amazon Redshift will provide the best performance if you specify sort keys, distribution keys, and column encodings as these can significantly improve storage, I/O, and query performance.

For increased ease of use, Amazon Redshift also supports a query engine called Amazon Redshift Spectrum, providing the features of Amazon Redshift on data directly stored on Amazon S3. With Amazon Redshift Spectrum, you can request your data stored in Amazon S3 up to exabyte-scale data lake environments.

Some customers want to run analytics queries on large amounts of less structured data. Presto is a powerful platform that accepts weakly structured data. It can query information at large scale from existing storage solutions and data lakes, such as Amazon S3. Amazon Athena is a fully managed Presto service that can run queries on your data lakes. Its integration with the data cataloging features of Amazon Glue and its full serverless model give you strong query capabilities for your analytics platforms.

Data Indexing and Searching

In many cases, you may need to index, search, and consolidate information at large scale. This allows your users to access detailed information, and analyze this information following multiple angles. Technologies such as search engines are particularly efficient for these use cases. Use search engines when you want search or report with low latency and high throughput.

Apache Lucene is a search engine platform that is widely used worldwide because it has been integrated in multiple platforms, such as Elasticsearch. Elasticsearch provides an easy-to-use platform that can automatically discover and index documents at a really large scale. Elasticsearch is also at the center of a complete ecosystem that supports reporting with Kibana and document or log



aggregation with LogStash and other providers. Kibana provides a simple solution for creating dashboards and analysis on indexed data.

You can deploy Elasticsearch on EC2 instances, leveraging EBS volumes, in multiple Availability Zones for high-availability and high performance. Elasticsearch supports scale-out, which is increasing capacity by adding new nodes to the cluster.

Amazon Elasticsearch Service (Amazon ES) is a managed service in the AWS Cloud that makes it simple and cost-effective to set up, manage, and scale an Elasticsearch cluster for your cloud application.

Amazon ES offers powerful scaling features on disk size or instance configuration (CPU, memory) and a model you can scale-out by adding nodes. As your data or query volume changes, you can scale your search domain's resources up or down as needed. You can control scaling if you know that you need more capacity for bulk uploads or are expecting a surge in search traffic.

Key AWS Services

The key AWS services for database solutions are Amazon RDS, Amazon DynamoDB, and Amazon RedShift, which provide relational, NoSQL, and data warehouse solutions. These services have configuration options that further allow you to optimize your storage solution. The access patterns of your components should influence the storage solution you select.

Resources

Refer to the following resources to learn more about AWS best practices for databases.

Documentation

- [Cloud Databases with AWS](#)¹⁴
- [Amazon Aurora best practices](#)¹⁵
- [Amazon Redshift performance](#)¹⁶
- [Amazon Athena top 10 performance tips](#)¹⁷
- [Amazon Redshift Spectrum best practices](#)¹⁸

- [Amazon DynamoDB best practices](#)¹⁹
- [Amazon DynamoDB Accelerator](#)²⁰

Network

The optimal network solution for a particular system will vary based on latency, throughput requirements, and so on. Physical constraints such as user or on-premises resources will drive location options, which can be offset using edge techniques or resource placement.

In AWS, networking is virtualized and is available in a number of different types and configurations. This makes it easier to match your networking methods more closely with your needs. AWS offers product features (for example, enhanced networking instance types, Amazon EBS optimized instances, Amazon S3 transfer acceleration, and dynamic CloudFront) to optimize network traffic. AWS also offers networking features (for example, Amazon Route 53 latency routing, Amazon Virtual Private Cloud (Amazon VPC) endpoints, and AWS Direct Connect) to reduce network distance or jitter.

Location

The AWS Cloud infrastructure is built around Regions and Availability Zones. A Region is a physical location in the world having multiple Availability Zones. Availability Zones consist of one or more discrete data centers, each with redundant power, networking, and connectivity, housed in separate facilities. These Availability Zones offer you the ability to operate production applications and databases that are more highly available, fault tolerant, and scalable than would be possible from a single data center.

Choose the appropriate Region or Regions for your deployment based on some key elements:

- **Where your users are located:** Choosing a Region close to your application's users ensures lower latency when they use the application.
- **Where your data is located:** For data-heavy applications, the major bottleneck in latency is when data is transferred to the computing part of the application. Application code should execute as close to the data as possible.

- **Other constraints:** Take into account constraints such as security and compliance.

Placement Groups

Amazon EC2 provides placement groups for networking. A placement group is a logical grouping of instances within a single Availability Zone. Using placement groups with supported instance types enables applications to participate in a low-latency, 20-gigabits-per-second (Gbps) network. Placement groups are recommended for applications that benefit from low network latency, high network throughput, or both. Using placement groups has the benefit of lowering jitter in network communications.

Edge Locations

Latency-sensitive services are delivered at the edge using a global network of edge locations. These edge locations commonly provide services such as Content Delivery Network (CDN) and Domain Name System (DNS). By having these services at the edge, they can respond with low latency to requests for content or DNS resolution. These services can also provide geographic services such as Geo Targeting of content (providing different content based on end users' location), or latency-based routing to direct end users to the nearest Region (minimum latency).

CloudFront is a global CDN that can be used to accelerate both static content such as images, scripts, and videos, as well as dynamic content such as APIs or web applications. It relies on a global network of edge locations that will cache the content and provide high-performance network connectivity to our users. CloudFront also accelerates many other features such as content uploading and dynamic applications, making it a performance addition to all applications serving traffic over the internet.

Route 53 is a highly available and scalable cloud DNS web service. It's 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 numeric IP addresses, like `192.168.2.1`, that computers use to connect to each other. Route 53 is fully compliant with IPv6 as well.

You should use edge services to reduce latency and to enable caching of content. You need to ensure that you have configured cache control correctly for both DNS and HTTP/HTTPS to gain the most benefit from these approaches.

Product Features

In cloud computing networking is critical, and therefore services commonly offer features to optimize network performance. You should consider product features such as EC2 instance network capability, enhanced networking instance types, Amazon EBS-optimized instances, Amazon S3 transfer acceleration, and dynamic CloudFront to optimize network traffic.

Amazon S3 content acceleration is a feature that lets external users benefit from the networking optimizations of CloudFront to upload data to Amazon S3. This makes it easy to transfer large amounts of data from remote locations that don't have dedicated connectivity to the AWS Cloud.

EC2 instances can also leverage enhanced networking. Enhanced networking uses single root I/O virtualization (SR-IOV) to provide high-performance networking capabilities on supported instance types. SR-IOV is a method of device virtualization that provides higher I/O performance and lower CPU utilization than traditional virtualized network interfaces. Enhanced networking provides higher bandwidth, higher packet per second (PPS) performance, and consistently lower inter-instance latencies.

Amazon Elastic Network Adapters (ENA) provide further optimization by delivering 20 Gbps of network capacity for your instances within a single placement group.

Amazon EBS-optimized instances use an optimized configuration stack and provide additional, dedicated capacity for Amazon EBS I/O. This optimization provides the best performance for your EBS volumes by minimizing contention between Amazon EBS I/O and other traffic from your instance.

Networking Features

When architecting your solution in the cloud you should consider networking features to reduce network distance or jitter.

Latency-based routing (LBR) for Route 53 helps you improve your application's performance for a global audience. LBR works by routing your customers to the AWS endpoint (for EC2 instances, Elastic IP addresses, or ELB load balancers) that provides the fastest experience based on actual performance measurements of the different AWS Regions where your application is running.

AWS Direct Connect provides dedicated connectivity to the AWS environment, from 50 Mbps up to 10 Gbps. This gives you managed and controlled latency and provisioned bandwidth so your applications can connect easily and in a performant way to other environments. Using one of the AWS Direct Connect partners, you can have end-to-end connectivity from multiple environments, thus providing an extended network with consistent performance.

Amazon VPC endpoints provide reliable connectivity to AWS services (for example, Amazon S3) without requiring an internet gateway or a Network Address Translation (NAT) instance. Using VPC endpoints, the data between your VPC and another AWS service is transferred within the Amazon network, helping protect your instances from internet traffic.

When implementing a scale-out architecture where you want to use multiple instances for service content, you can leverage load balancers inside your Amazon VPC. AWS provides multiple models for your applications in the ELB service. Application Load Balancer is best suited for load balancing of HTTP and HTTPS traffic and provides advanced request routing targeted at the delivery of modern application architectures, including microservices and containers. Network Load Balancer is best suited for load balancing of TCP traffic where extreme performance is required. It is capable of handling millions of requests per second while maintaining ultra-low latencies, and it is also optimized to handle sudden and volatile traffic patterns.

Key AWS Services

The key AWS service for networking solutions is Route 53, which provides latency-based routing. In addition, the use of Amazon VPC endpoints and AWS Direct Connect can reduce network distance or jitter.

Resources

Refer to the following resources to learn more about AWS best practices for networking.

Documentation

- [Networking Products with AWS](#)²¹

Review

When you first architect your solution, you have a fixed set of options to choose from. Over time, new technologies and approaches become available that can improve the performance of your architecture. In the AWS Cloud, it's much easier to experiment with new features and services because your infrastructure is code.

Using AWS, you can take advantage of our continual innovation, which is driven by customer need. We reduce prices and release new Regions, edge locations, services, and features regularly. Any of these new releases could positively improve the performance efficiency of your architecture.

After you have identified your architectural approach, you should use benchmarking and load testing data to drive your selection of resource types and configuration options.

To adopt a data-driven approach to architecture you should implement a performance review process that considers the following:

- **Infrastructure as code:** Define your infrastructure as code using approaches such as AWS CloudFormation templates. The use of templates allows you to place your infrastructure into source control alongside your application code and configurations. This enables you to apply the same practices you use to develop software to your infrastructure so you can iterate rapidly.
- **Deployment pipeline:** Use a continuous integration/continuous deployment (CI/CD) pipeline (for example, source code repository, build systems, deployment, and testing automation) to deploy your infrastructure. This enables you to deploy in a repeatable, consistent, and low-cost fashion as you iterate.
- **Well-defined metrics:** Set up your metrics and monitoring to capture key performance indicators (KPIs). We recommend that you use both technical and business metrics. For website or mobile apps key metrics are capturing time to first byte or rendering. Other generally applicable metrics include thread count, garbage collection rate, and wait states. Business metrics, such as the aggregate cumulative cost per request, can alert you to ways to drive down costs. Carefully consider how you plan to

interpret metrics. For example, you could choose the maximum or 99th percentile instead of the average.

- **Performance test automatically:** As part of your deployment process, automatically trigger performance tests after the quicker running tests have passed successfully. The automation should create a new environment, set up initial conditions such as test data, and then execute a series of benchmarks and load tests. Results from these tests should be tied back to the build so you can track performance changes over time. For long running tests, you can make this part of the pipeline asynchronous from the rest of the build. Alternatively, you could execute performance tests overnight using Amazon EC2 Spot Instances.
- **Load generation:** You should create a series of test scripts that replicate synthetic or prerecorded user journeys. These scripts should be idempotent and not coupled, and you might need to include “pre-warming” scripts to yield valid results. As much as possible, you want your test scripts to replicate the behavior of usage in production. You can use software or software-as-a-service (SaaS) solutions to generate the load. Consider using AWS Marketplace solutions and Spot Instances. They can be cost-effective ways to generate the load.
- **Performance visibility:** Key metrics should be visible to your team, especially metrics against each build version. This allows you to see any significant positive or negative trend over time. You should also display metrics on the number of errors or exceptions to make sure you are testing a working system.
- **Visualization:** Use visualization techniques that make it clear where performance issues, hot spots, wait states, or low utilization is occurring. Overlay performance metrics over architecture diagrams-- call graphs or code can help identify issues more quickly.

This performance review process can be implemented as a simple extension of your existing deployment pipeline and then evolved over time as your testing requirements become more sophisticated. For future architectures, you should be able generalize your approach and reuse the same process and artifacts.

When architectures perform badly this is normally because a performance review process has not been put into place or is broken. If your architecture is performing badly, putting this performance review process in place will allow

you to apply Deming’s plan-do-check-act (PDCA) cycle to drive iterative improvement.

Benchmarking

Benchmarking uses synthetic tests to provide you with data on how components perform. In this section, we discuss how to use benchmarking to drive improvements in your workloads. We don’t cover the use of benchmarking to compare different vendor’s products or implementations.

Benchmarking is generally quicker to set up than load testing and is used when you want to evaluate the technology for a particular component. Benchmarking is often used at the start of a new project, when you don’t yet have a whole solution that you could load test.

For benchmarking, you should follow the performance review process. However, your deployment pipeline will just consist of the benchmark tests themselves. You can either build your own custom benchmark tests, or you can use an industry standard test, such as [TPC-DS](#) (to benchmark your data warehousing workloads).²² Industry benchmarks are helpful when you are comparing different environments. Custom benchmarks are useful for targeting specific types of operations that you expect to make in your architecture.

With benchmarking, it is generally more important to pre-warm your test environment to ensure valid results. You should run the same benchmark multiple times to be sure you’ve captured any variance over time.

Since benchmarks are generally faster to run than load tests, they can be used earlier in the deployment pipeline to provide faster feedback on performance deviations to your team. When you evaluate a significant change in a component or service, a benchmark can be a quick way to see if you can justify the effort to make the change based on the benchmarked difference. Benchmarking should be used in conjunction with load testing because load testing will tell you how your whole workload will perform in production.

Key AWS Services

The key AWS services supporting benchmarking are AWS CodeDeploy and AWS CloudFormation. Use these services to automate the testing of your infrastructure in a repeatable fashion.

Resources

Refer to the following resources to learn more about AWS best practices for benchmarking.

Videos

- [Performance Channel](#)²³
- [Performance Benchmarking on AWS](#)²⁴

Documentation

- [Amazon S3 Performance Optimization](#)²⁵
- [Amazon EBS Volume Performance](#)²⁶
- [AWS CodeDeploy](#)²⁷
- [AWS CloudFormation](#)²⁸

Load Testing

Load testing uses your *actual* workload so you can see how your whole solution performs in a production environment. Load tests should be done using synthetic or sanitized versions of production data (remove sensitive or identifying information). Use replayed or pre-programmed user journeys through your application at scale that exercise your entire architecture. As part of your delivery pipeline you can automatically carry out load tests and compare against pre-defined KPIs and thresholds to ensure that you continue to get the performance you require.

Amazon CloudWatch can collect metrics across the resources in your architecture. You can also collect and publish custom metrics to surface business or derived metrics. Use CloudWatch to set alarms that indicate when thresholds are breached to signal that a test is outside of expected performance.

Using AWS services, you can run production-scale environments to test your architecture aggressively. Since you only pay for the test environment when it is needed, you can carry out full-scale testing at a fraction of the cost of using an on-premises environment. You should take advantage of the AWS Cloud to test your workload to see where it fails to scale or scales in a non-linear way. You can

use Spot Instances to generate loads at low cost and discover bottlenecks before they are experienced in production.

When you write critical user stories for your architecture, you should include performance requirements such as specifying how quickly each critical story should execute. For these critical stories, you should implement additional scripted user journeys to ensure you have visibility into how these stories perform against your requirement.

Where load tests take considerable time to execute, you should parallelize them using multiple copies of your test environment. Your costs will be similar, but your testing time will be reduced. (It costs the same to run one EC2 instance for 100 hours as it does to run 100 instances for one hour.) You can also lower the costs of these tests by using Spot Instances and selecting Regions that have lower costs than the Regions you use for production.

The location of your load test clients should reflect the geographic spread of your end users.

Key AWS Services

The key AWS service supporting load testing is CloudWatch, which allows you to collect metrics on how your whole architecture is performing during the load test. With CloudWatch you can also create custom and business metrics.

Resources

Refer to the following resources to learn more about AWS best practices for load testing.

Documentation

- [Load Testing CloudFront](#)²⁹

Monitoring

After you have implemented your architecture you will need to monitor its performance so that you can remediate any issues before your customers are aware of them. Monitoring metrics should be used to raise alarms when thresholds are breached. The alarm can trigger automated action to work around any badly performing components.

CloudWatch provides the ability to monitor and send notification alarms. You can use automation to work around performance issues by triggering actions through Amazon Kinesis, Amazon Simple Queue Service (Amazon SQS), and AWS Lambda.

Active and Passive

Monitoring solutions fall into two types: active monitoring (AM) and passive monitoring (PM). AM and PM complement each other to give you a full view of how your workload is performing.

Active monitoring simulates user activity in scripted user journeys across critical paths in your product. AM should be continuously performed in order to test the performance and availability of a workload. AM complements PM by being continuous, lightweight, and predictable. It can be run across all environments (especially pre-production environments) to identify problems or performance issues before they affect end users.

Passive monitoring is commonly used with web-based workloads. PM collects performance metrics from the browser (non-web-based workloads can use a similar approach). You can collect metrics across all users (or a subset of users), geographies, browsers, and device types. You should use PM to understand these issues:

- **User experience performance:** PM provides you with metrics on what your users are experiencing, which gives you a continuous view into how production is working, as well as a view into the impact of changes over time.
- **Geographic performance variability:** If a workload has a global footprint and users access the application from all around the world, using PM can enable you to spot a performance problem affecting users in a specific geography.
- **The impact of API use:** Modern workloads use internal APIs and third-party APIs. PM provides the visibility into the use of APIs so you can identify performance bottlenecks that originate not only from internal APIs but also from third-party API providers.

Phases

Monitoring at AWS consists of five distinct phases, which are explained in more detail in the [Reliability Pillar of the AWS Well-Architected Framework](#) whitepaper:³⁰

1. Generation – scope of monitoring, metrics, and thresholds
2. Aggregation – creating a complete view from multiple source
3. Real-time processing and alarming – recognizing and responding
4. Storage – data management and retention policies
5. Analytics – dashboards, reporting, and insights

CloudWatch is a monitoring service for AWS Cloud resources and the workloads that run on AWS. You can use CloudWatch to collect and track metrics, collect and monitor log files, and set alarms. CloudWatch can monitor AWS resources such as EC2 instances and RDS DB instances, as well as custom metrics generated by your applications and services, and any log files your applications generate. You can use CloudWatch to gain system-wide visibility into resource utilization, application performance, and operational health. You can use these insights to react quickly and keep your application running smoothly. CloudWatch dashboards enable you to create reusable graphs of AWS resources and custom metrics so you can monitor operational status and identify issues at a glance.

Ensuring that you do not see too many false positives or are overwhelmed with data is key to an effective monitoring solution. Automated triggers avoid human error and can reduce the time it takes to fix problems. Plan for *game days*, where simulations are conducted in the production environment, to test your alarm solution and ensure that it correctly recognizes issues.

Key AWS Services

The key AWS service that supports monitoring is Amazon CloudWatch, which allows easy creation of alarms that can trigger scaling actions. The following services and features are also important:

- **Amazon S3** acts as the storage layer and allows for lifecycle policies and data management.

- **Amazon EMR** can analyze metrics to help you monitor performance.

Resources

Refer to the following resources to learn more about AWS best practices for monitoring to promote performance efficiency.

Videos

- [AWS re:Invent 2015 | \(DVO315\) Log, Monitor and Analyze your IT with Amazon CloudWatch](#)³¹
- [AWS re:Invent 2015 | \(BDT312\) Application Monitoring: Why Data Context Is Critical](#)³²

Documentation

- [CloudWatch Documentation](#)³³

Trade-Offs

When you architect solutions, think about trade-offs so you can ensure an optimal approach. Depending on your situation you could trade consistency, durability, and space versus time or latency, to deliver higher performance.

Using AWS, you can go global in minutes and deploy resources in multiple locations across the globe to be closer to your end users. You can also dynamically add read-only replicas to information stores, such as database systems to reduce the load on the primary database. AWS also offers caching solutions such as Amazon ElastiCache, which provides an in-memory data store or cache, and CloudFront, which caches copies of your static content closer to end users.

The following sections detail some of the trade-offs you can make and how you can implement them.

Technique	Applies To	Use	Gains
Caching	Read-heavy	Space (Memory)	Time
Partitioning or Sharding	Write-heavy	Size & Complexity	Time
Compression	Large data	Time	Space



Technique	Applies To	Use	Gains
Buffering	Many requests	Space & Time	Efficiency

Caching

Most workloads rely on a dependent component such as a service or database that offers a source of truth or a consolidated view of data. Generally, these architecture components are harder to scale and represent a significant proportion of the cost of the workload. You can improve performance efficiency by using caching to trade off against freshness or memory used. These techniques generally update asynchronously or periodically. The trade-off is that your data isn't always fresh and, therefore, not always consistent with the source of truth.

Application Level

You can make this trade-off at a code level by using application-level caches or memorization. When requests are cached, execution time is reduced. This provides a way to scale horizontally through the caching layer and reduces load on your most heavily used components.

In-memory and distributed caches are used in two main cases: coordinating transient data and state between distinct servers, such as user sessions in an application, or protecting databases from read-heavy workloads by serving the most requested elements directly from memory.

Platforms such as Redis, Memcached, or Varnish can be deployed on Amazon EC2 and will provide robust caching engines for your applications. Design and scale of these platforms are usually based on the memory of the instance and by designing appropriate key management in the applications when you want to activate a cluster model (for example, for Redis clustering, ElastiCache-consistent hashing).

Amazon ElastiCache is a web service that makes it easy to deploy, operate, and scale an in-memory data store or cache in the cloud. ElastiCache with Memcached supports sharding to scale in-memory cache with multiple nodes. ElastiCache for Redis includes clustering, with multiple shards forming a single in-memory key-value store that is terabytes in size, plus read replicas per shard for increased data access performance.



ElastiCache is a fully managed service that supports scaling the environment, automated failover, patching, and backup.

Database Level

Database replicas enhance performance databases by replicating all changes to the master databases to read replicas (not available for Oracle or SQL Server). This replication makes it possible to scale out beyond the capacity constraints of a single database for read-heavy database workloads.

Amazon RDS provides read replicas as a fully managed service. You can create one or more replicas of a given source database and serve high-volume application read traffic from multiple copies of your data, thereby increasing aggregate read throughput. You should also add additional indexes to the read replica, where the database engine supports it. For example, you could add more indexes to the MySQL read replica. For latency-sensitive workloads you should use the Multi-AZ feature to specify which Availability Zones the read replica should be in to reduce cross-Availability Zone traffic.

Geographic Level

Another example of caching is using a CDN, which is a good way to reduce latency for clients. CDNs should be used to store static content and to accelerate dynamic content. You should consider using a CDN for your APIs. Even dynamic content can benefit through the use of network optimization methods.

CloudFront can be used to deliver your entire website, including dynamic, static, streaming, and interactive content using a global network of edge locations. Requests for your content are automatically routed to the nearest edge location, so content is delivered with the best possible performance. CloudFront is optimized to work with other AWS services, like Amazon S3, Amazon EC2, ELB, and Route 53. CloudFront also works seamlessly with any non-AWS origin server, which stores the original, definitive versions of your files.

Key AWS Services

The key AWS services for caching solutions are ElastiCache, which provides a general-purpose application cache, and CloudFront, which allows you to cache information closer to your users.

Resources

Refer to the following resources to learn more about AWS best practices for caching.

Documentation

- [Best Practices for Implementing Amazon ElastiCache](#)³⁴
- [AWS CloudFormation Best Practices](#)³⁵

Video

- [Turbocharge your apps with Amazon ElastiCache](#)³⁶

Partitioning or Sharding

When using technologies, such as relational databases, that require a single instance due to consistency constraints, you can only scale vertically (by using higher specification instances and storage features). When you hit the limits of vertical scaling, you can use a different approach called data partitioning or *sharding*. With this model, data is split across multiple database schemas, each running in its own autonomous primary DB instance.

Amazon RDS removes the operational overhead of running multiple instances, but sharding will still introduce complexity to the application. The application's data access layer will need to be modified to have awareness of how data is split so that it can direct queries to the right instance. (You can use a proxy or routing mechanism to remove caching code from the application, or implement it in a data access layer.) In addition, any schema changes will have to be performed across multiple database schemas, so it is worth investing some effort to automate this process.

NoSQL database engines will typically perform data partitioning and replication to scale both the reads and the writes in a horizontal fashion. They do this transparently without the need of having the data partitioning logic implemented in the data access layer of your application. Amazon DynamoDB in particular manages table partitioning for you automatically, adding new partitions as your table grows in size or as read- and write-provisioned capacity changes.

Partitioning or sharding provides a way to scale write-heavy workloads, but requires that data is evenly distributed and evenly accessed across all partitions or shards. It can introduce complexity in relational database solutions, while NoSQL solutions generally trade consistency to deliver this.

Key AWS Services

The key AWS service for partitioning or sharding is Amazon DynamoDB, which manages table partitioning for you automatically.

Resources

Refer to the following resources to learn more about AWS best practices for partitioning and sharding.

Documentation

- [Best Practices for DynamoDB](#)³⁷

Video

- [AWS re:Invent 2015 | \(DAT401\) Amazon DynamoDB Deep Dive](#)³⁸

Compression

Compressing data trades computing time against space and can greatly reduce storage and networking requirements. Compression can apply to file systems, data files, and web resources such as stylesheets and images, but also to dynamic responses such as APIs.

CloudFront supports compression at the edge. The source system can serve resources in a standard fashion, and the CDN will automatically compress the resources if and only if the web clients can support it.

When transferring large quantities of information into or out of the cloud you should consider non-network based solutions. AWS Snowball is a petabyte-scale data transport solution that uses secure appliances to transfer large amounts of data into and out of the AWS Cloud. Using AWS Snowball addresses common challenges with large-scale data transfers, including high network costs, long transfer times, and security concerns.

Amazon Redshift uses compression with columnar data storage. Instead of storing data as a series of rows, Amazon Redshift organizes the data by column. Columnar data stores can be compressed much more than row-based data stores because similar data is stored sequentially on disk. Amazon Redshift employs multiple compression techniques and can often achieve significant compression relative to traditional relational data stores. When loading data into an empty table, Amazon Redshift automatically samples your data and selects the most appropriate compression scheme.

Key AWS Services

The key AWS service for compression is CloudFront, which supports compression at the edge.

Resources

Refer to the following resources to learn more about AWS best practices for compression.

Documentation

- [Amazon CloudFront: Serving Compressed Files](#)³⁹
- [Amazon RedShift: Columnar Storage](#)⁴⁰
- [AWS Snowball: What Is AWS Snowball?](#)⁴¹

Buffering

Buffering uses a queue to accept messages (units of work) from producers. For resiliency, the queue should use durable storage. A buffer is a mechanism to ensure that applications can communicate with each other when they are running at different rates over time. Messages can then be read by consumers, which allows the messages to run at the rate that meets the consumers' business requirements. By using a buffer, you can decouple the throughput rate of producers from that of consumers. You don't have to worry about producers having to deal with data durability and backpressure (where producers slow down because their consumer is running slowly).

When you have a workload that generates significant write load that doesn't need to be processed immediately, you can use a buffer to smooth out demands on consumers.

On AWS, you can choose from multiple services to implement a buffering approach. Amazon SQS provides a queue that allows a single consumer to read individual messages. Amazon Kinesis provides a stream that allows many consumers to read the same messages.

An Amazon SQS queue is a reliable, scalable, and fully managed repository for messages that are awaiting processing. Amazon SQS enables applications to quickly and reliably queue messages that one component in the application generates to be consumed by another component. For this approach, a queue is created that producers post messages to, and which resides at a well-known address.

To reduce cost, producers post messages using Amazon SQS batch API actions. Consumers read messages from the well-known queue using a fixed-sized Auto Scaling group of EC2 instances to cope with instance failures. Long polling lets you retrieve messages from your Amazon SQS queue as soon as they become available. Use long polling to reduce the cost of retrieving messages. Amazon SQS uses the message visibility feature to hide messages that have been read. Message visibility reduces the risk that you might process the same message twice, though you should be aware that by default Amazon SQS at-least-once delivery means that you can receive a message more than once.

You can use Amazon SQS first in/first out (FIFO) queues if you need exactly-once processing. Message visibility allows for messages that have not been processed to reappear (for example, in the case of an instance failure). For long-running tasks you can extend the visibility time-out. Your application will need to delete messages after they have been processed.

An alternative approach is to use Amazon Kinesis to provide buffering. It differs from Amazon SQS in that it allows multiple consumers to read the same message at any one time. However, a single message will be read using the Kinesis Client Library, or AWS Lambda, and will be delivered to one-and-only-one consumer for an *Application*, which is a name provided by the consumer when it connects to the Stream. Different *Applications* can all consume the same messages concurrently, providing a “Publish and Subscribe” model.

When architecting with a buffer keep in mind two key considerations. First, what is the acceptable delay between producing the work and consuming the work? Second, how do you plan to handle duplicate requests for work?

To optimize the speed with which work items are consumed by more consumers, you can use Spot Instances. Using Spot Instances, you can bid on spare Amazon EC2 computing capacity. To handle duplicate messages with Amazon SQS we recommend using Amazon SQS first in-first-out (FIFO) queues, which provide exactly-once processing.

Key AWS Services

The key AWS service for buffering is Amazon SQS, which allows you to decouple producers from consumers using a queue.

Resources

Refer to the following resource to learn more about AWS best practices for buffering.

Documentation

- [Best Practices for Amazon SQS](#)⁴²

Conclusion

Achieving and maintaining performance efficiency requires a data-driven approach. You should actively consider access patterns and trade-offs that will allow you to optimize for higher performance. Using a review process based on benchmarks and load tests will allow you to select the appropriate resource types and configurations. Treating your infrastructure as code will enable you to rapidly and safely evolve your architecture, while you use data to make fact-based decisions about your architecture. Putting in place a combination of active and passive monitoring will ensure that the performance of your architecture does not degrade over time.

AWS strives to help you build architectures that have performance efficiency while delivering business value. To make your architectures truly perform, you should use the tools and techniques discussed in this paper.

Contributors

The following individuals and organizations contributed to this document:

- Philip Fitzsimons, Sr Manager Well-Architected, Amazon Web Services
- Julien Lépine, Specialist SA Manager, Amazon Web Services
- Ronnen Slasky, Solutions Architect, Amazon Web Services

Further Reading

For additional help, please consult the following sources:

- [AWS Well-Architected Framework](#)⁴³

Document Revisions

Date	Description
November 2017	Refreshed the whitepaper to reflect changes in AWS
November 2016	First publication

Notes

- 1 <https://aws.amazon.com/architecture/well-architected/>
- 2 <https://d1.awsstatic.com/whitepapers/architecture/AWS-Cost-Optimization-Pillar.pdf>
- 3 <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/instance-types.html>
- 4 http://docs.aws.amazon.com/AmazonECS/latest/developerguide/ECS_instances.html
- 5 <http://docs.aws.amazon.com/lambda/latest/dg/lambda-introduction-function.html#resource-model>
- 6 <https://aws.amazon.com/what-is-cloud-storage/>
- 7 <http://docs.aws.amazon.com/AmazonS3/latest/dev/request-rate-perf-considerations.html>
- 8 <http://docs.aws.amazon.com/amazonglacier/latest/dev/introduction.html>
- 9 <http://docs.aws.amazon.com/efs/latest/ug/performance.html>
- 10 <http://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/ebs-io-characteristics.html>
- 11 https://www.youtube.com/watch?v=2wKgha8CZ_w
- 12 <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Storage.html>
- 13 <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html>
- 14 <https://aws.amazon.com/products/databases/>
- 15 <http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Aurora.BestPractices.html>
- 16 http://docs.aws.amazon.com/redshift/latest/dg/c_challenges_achieving_high_performance_queries.html
- 17 <https://aws.amazon.com/blogs/big-data/top-10-performance-tuning-tips-for-amazon-athena/>

18 <https://aws.amazon.com/blogs/big-data/10-best-practices-for-amazon-redshift-spectrum/>

19

<http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/BestPractices.html>

20 <https://aws.amazon.com/dynamodb/dax/>

21 <https://aws.amazon.com/products/networking/>

22 <http://www.tpc.org/tpcds/default.asp>

23

<https://www.youtube.com/playlist?list=PLhr1KZpdzukfOfdstfSQVAvUrjBvnGd4i>

24 <https://www.youtube.com/watch?v=sHxLpuC2CUI>

25

<http://docs.aws.amazon.com/AmazonS3/latest/dev/PerformanceOptimization.html>

26

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSPerformance.html>

27 <https://aws.amazon.com/codedeploy/>

28 <https://aws.amazon.com/cloudformation/>

29

<http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/load-testing.html>

30 <https://d1.awsstatic.com/whitepapers/architecture/AWS-Reliability-Pillar.pdf>

31 <https://youtu.be/ZaOR-ybLJFO>

32 <https://youtu.be/aoN7bRyGA0g>

33

<http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/WhatIsCloudWatch.html>

34

<http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/BestPractices.html>

35 <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/best-practices.html>

36 <http://www.slideshare.net/AmazonWebServices/video-turbocharge-your-apps-with-amazon-elasticache>

37

<http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/BestPractices.html>

38 https://www.youtube.com/watch?v=ggDIat_FZtA

39

<http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/ServingCompressedFiles.html>

40

http://docs.aws.amazon.com/redshift/latest/dg/c_columnar_storage_disk_mem_mgmt.html

41

<http://docs.aws.amazon.com/AWSImportExport/latest/ug/whatisSnowball.html>

42

<http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-best-practices.html>

43 <https://aws.amazon.com/well-architected>