International Journal of Engineering Trends and Technology
ISSN: 2231–5381 / https://doi.org/10.14445/22315381/IJETT-V73I6P129

Volume 73 Issue 6, 350-370, June 2025

Review Article

A Weighted Scoring Model Analysis of Cloud Storage Services: Comparing AWS, Azure, and Google Cloud Platform

Hasnae NOUHAS[1*], Abdessamad BELANGOUR[1], Mahmoud NASSAR[2]

[1]*Laboratory of Artificial Intelligence and Systems LAIS, Hassan II University, Faculty of Sciences Ben M'Sik, Casablanca, Morocco.*
[2]*IMS Team, ADMIR Laboratory, ENSIAS, Rabat IT Center, Mohammed V University in Rabat, Rabat, Morocco.*

[*]*Corresponding Author : hasnae_nouhas@um5.ac.ma*

Received: 17 March 2025 Revised: 12 May 2025 Accepted: 07 June 2025 Published: 28 June 2025

Abstract - *Cloud storage solutions have become integral to modern data management for scalability, Security, and efficiency. The current study has assessed offerings by Azure, AWS, and the Google Cloud Platform for storage services using the WSM, extending a first systematic review to object, block, file, archive, hybrid, backup, and edge/offline storage. This review considers critical criteria: performance, Availability, scalability, Security, cost efficiency, integration, usability, and data management. In this study, sharp differences emerge among providers. Google Cloud Platform is exceptionally strong in performance and scalability-particularly block and archive storage-which makes it a perfect fit for latency-sensitive and integration-intensive applications. AWS outshines most on performance grounds, usability, and scalability; it leads in file, object, and hybrid storage classes and can be adapted into use in any given scenario. Azure is driven by cost-effectiveness and seamlessly integrates with ecosystem engagement, offering tangible value to budget-conscious organisations and tactically leveraging the Azure cloud. This research provides methods to evaluate cloud storage, spanning the depth that generally seems to lie between a rich set of technological capabilities and specific business needs. The results reinforce that storage decisions must be aligned with specific operational needs to realise the most efficiency and value. Implications range from guiding organisations in adopting fitted storage strategies to contributing to the greater understanding of relative cloud storage performance and usability within the cloud computing domain.*

Keywords - *AWS, cloud storage evaluation, Google Cloud platform, Azure, WSM.*

1. Introduction

Cloud computing allows users to utilise computing resources such as networks, servers, storage, applications, and services over the Internet at any time. These resources are pooled together in a multi-shared virtual environment that allows users to access and release services quickly and with minimal interaction. This model has transformed how organisations handle and store information by providing flexible, scalable, and cost-efficient solutions [1]. Storage plays a foundational role in cloud computing services by allowing organisations to control, store, and retrieve vast volumes of data across distributed environments. Cloud storage solutions offer high availability, durability, and elasticity, making them essential for data-driven operations.

As the demand for cloud-native architectures rises, choosing the most suitable storage solution becomes a strategic priority. Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP) are the dominant players in this field, each offering various storage services tailored to different use cases ranging from object and block storage to edge and hybrid solutions. While these providers advertise high performance, security, and integration features, comparative studies often lack a systematic and quantitative approach that assists businesses in making informed decisions [2]. While several previous studies have attempted comparisons, many are limited in scope, relying on qualitative methods or focusing on a narrow subset of storage types. Such approaches lack the comprehensive, data-supported evaluation necessary for informed enterprise decision-making.

A clear research gap exists in quantitatively assessing cloud storage solutions using a standardised scoring model that reflects technical and usability criteria. This paper fills that gap by introducing a Weighted Scoring Model (WSM) that evaluates cloud storage services across eight dimensions: performance, Availability and durability, scalability, Security, cost efficiency, integration, usability, and data management. The study provides a reproducible and practical benchmark for assessing AWS, Azure, and GCP storage services through this structured approach. The paper is organised as follows: The section on Related Work reviews the prior studies on cloud storage comparisons. The section named Background provides an overview of each cloud platform and the different types of cloud storage services provided. Comparison contains a detailed analysis of storing anything on AWS, Azure, and Google Cloud. The discussion presents implications from the findings, followed by the Conclusion, where a summary is made of the findings of the comparison study.

2. Related work

In [3], This book chapter presents an acategorised comparison of cloud storage solutions based on diverse application demands, such as archival needs, real-time data access, and high-availability scenarios. It uses a descriptive methodology to classify storage types-object, block, and file-according to workload sensitivity and evaluates their scalability, performance, cost-effectiveness, and suitability for different use cases. The discussion acknowledges the strengths of each major provider: AWS for its scalability and flexibility, Azure for its seamless integration with Microsoft services and cost-efficient tiered storage, and GCP for its user-friendly and simplified approach. While the study provides practical insights into aligning storage types with application needs, the providers' specific benchmarking foundational scoring or ranking approach is absent. This limits its utility for rigorous platform selection across multiple cloud vendors, especially when objective evaluation criteria are required. This limits its utility for rigorous platform selection across multiple cloud vendors, especially when objective evaluation criteria are required.

In [4], the authors concentrated on a technical comparison between Amazon S3 and Azure Blob Storage, examining aspects such as latency, Availability or data replication. The paper is worth reading to understand the differences between the top two object storage services. It is, however, very limited as it does not include other types of storage or other storage providers, and the quantitative decision-making model has not been developed. Further, due to fast-moving cloud services, the age of the study (2018) limits its relevance.

In[5], the paper presents a robust architectural and performance comparison of AWS, GCP, and Azure in the context of IoT solutions. It discusses cloud infrastructure, latency metrics, and throughput efficiency, offering valuable quantitative performance data. However, the storage component is only briefly discussed and treated as a supporting service. Thus, while the paper is strong on networked application architecture, it does not provide depth for evaluating cloud storage specifically. Paper [2] offers an in-depth overview of AWS's tools, best practices, and technical ecosystem, covering various computing, storage, and Security services. It provides extensive insight into AWS's internal mechanisms, service orchestration, and cost management strategies. In particular, it highlights the flexibility and power of AWS storage offerings-such as Amazon S3 for object storage, Amazon EBS for block storage, and Amazon EFS for file storage-demonstrating how these services effectively address varied application requirements. The paper emphasises best practices in cost control, Security, and performance optimisation, illustrating AWS's capability to adapt to different operational needs in cloud environments. However, the study focuses exclusively on AWS and lacks comparative analysis with other cloud providers. It also does not offer empirical benchmarks or decision-making frameworks for evaluating storage services specifically, which limits its applicability for cross-platform storage selection.

Paper [6] offers a useful overview of cloud storage services from the three leading providers and comparisons, as well as latency, redundancy, pricing, integration, and compliance features. While informative, the analysis remains qualitative and lacks a structured evaluation model. It highlights key service characteristics but does not provide a weighted framework to guide decision-making based on operational priorities. This reveals a gap in the existing literature: the need for a comprehensive, data-driven comparison method that can adapt to different use cases.

In [7], the paper provides an accessible overview of major cloud storage providers, focusing on their architectural design, core components, and service offerings. It compares AWS, Azure, and GCP based on storage structure, integration capabilities, scalability, and cost-effectiveness. While the paper is valuable for its broad and practical insights, it primarily adopts a descriptive approach and does not employ standardised evaluation metrics or quantitative analysis. It also lacks a structured method for comparing performance or prioritising features based on enterprise requirements.

Building on these prior studies, it becomes evident that while many offer valuable descriptive insights, there remains a notable gap in the literature: a comprehensive, data-driven comparison method that can adapt to diverse application needs. To address this, our research introduces a Weighted Scoring Model (WSM) that enables a more rigorous andcustomisable evaluation of cloud storage services. Our approach bridges the gap between narrative comparisons and systematic evaluation by assessing multiple storage types-object, block, file, archive, hybrid, edge/offline, and backup-across eight well-defined criteria. The result is a reproducible and practical benchmark designed to guide decision-makers in selecting optimal storage solutions across AWS, Azure, and GCP.

3. Background

Until recently, cloud storage solutions were categorised into distinct types, each designed to address specific data management needs. Object storage is tailored for scalability and durability, particularly for unstructured data like backups and multimedia. Block storage delivers high-performance, low-latency storage at the block device level, which is ideal for applications such as databases and virtual machines. File storage supports shared access for multiple applications using traditional file systems, making it suitable for applications that require collaboration and legacy systems [8]. Archive storage saves data for a long time without paying too much. Hybrid and edge storage solutions connect on-premises and cloud environments, and backup storage protects your data and helps you recover if something goes wrong [9]. These storage types can be used for many things, like business operations, IoT, and following rules. Table 1 shows the services offered by cloud providers for each storage type.

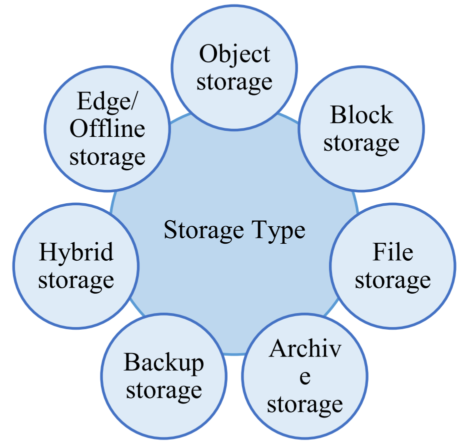

Fig. 1 Storage types provided by cloud providers

Table 1. Storage services per storage type

Storage Type	Cloud Provider	Service Name	Definition
Object Storage	Azure	Azure Blob Storage	A cloud object storage service specifically built to manage and store huge amounts of unstructured data, like text files or binary content [10].
	AWS	Amazon S3	A scalable object storage solutionrecognised for its exceptional durability, high Availability, strong security measures, and consistent performance [11].
	GCP	Google Cloud Storage	A secure and scalable object storage solution tailored for managing large volumes of unstructured data [12].
Block Storage	Azure	Azure Managed Disks	Are Azure-managed block storage volumes specifically designed to support Azure Virtual Machines (VMs) [13].
	AWS	AWS Elastic Block Store (EBS)	A high-performance block storage is built to work seamlessly with Amazon EC2 instances [14].
	GCP	Google Persistent Disks	Google Persistent Disks are durable, high-performance block storage options for virtual machine (VM) instances on Google Cloud [15].
File Storage	Azure	Azure Files	A fully managed cloud file share service that enables seamless file sharing across cloud and on-premises environments [16].
	AWS	Amazon EFS	A fully managed and scalable NFS file system that automatically adjusts capacity as needed, supporting both AWS cloud services and on-premises environments [17].
	GCP	Google Filestore	Google Filestore is a fully managed file storage service on Google Cloud, providing scalable and high-performance Network File System (NFS) file shares for applications [18].
Archive Storage	Azure	Azure Archive Storage	It offers a low-cost, secure, and highly durable cloud storage option, ideal for storing rarely accessed data with flexible latency needs [19].
	AWS	AWS Glacier	A reliable, secure, and cost-effective storage solution built explicitly for archiving data and long-term backups [20].
	GCP	Google Cloud Archive Storage	It provides a budget-friendly and highly durable option for long-term data storage and archiving needs [21].
Hybrid Storage	Azure	Azure Data Box Gateway	A cloud storage service that facilitates smooth and secure data transfer between on-premises systems and the Azure cloud [22].
	AWS	AWS Storage Gateway	A hybrid cloud storage solution that enables on-premises applications to seamlessly connect with virtually unlimited cloud storage [23].

	GCP	Google Anthos Storage	Google Cloud's Anthos is a modern application platform that enables consistent development and operations across hybrid and multi-cloud environments. While Anthos is not a storage service, it integrates with various solutions to support application data needs [24].
Edge/Offline Storage	Azure	Azure Stack Edge	Azure Stack Edge is a managed hardware-as-a-service solution from Microsoft that brings Azure's compute, storage, and intelligence capabilities to edge locations. It enablesorganisations to process data locally, run machine learning models, and transfer data efficiently between edge sites and Azure [25].
	AWS	AWS Snow Family	Physical devices like Snowball for edge computing and offline data transfer [26].
	GCP	Google Transfer Appliance	A high-capacity storage device designed to facilitate the secure and efficient transfer of large datasets to Google Cloud [27].
Backup Storage	Azure	Azure Backup	A cloud-based solution that offers an easy, secure, and budget-friendly way to back up data and restore it directly from the Microsoft Azure cloud [28].
	AWS	AWS Backup	A fully managed service that centralises and automates data backup across both AWS services and on-premises systems [29].
	GCP	Google Cloud Backup	Google Cloud Backup and DR (Disaster Recovery) Service is a managed solution that provides centralised, application-consistent data protection for workloads running in Google Cloud and on-premises environments [30].

3.1. Object Storage

Object storage is built to handle unstructured data like media files, backups, and logs. Unlike conventional file systems, it stores data as individual objects, each assigned a unique identifier and accompanied by metadata for easy access and management [31]. The capabilities and use cases of object storage solutions are exemplified by Azure Blob Storage, Amazon S3 and Google Cloud Storage, as detailed in Table 2, which reviews their strengths in scalability and durability for managing unstructured data.

Table 2. Review of object storage services: Azure blob storage, Amazon S3, and Google cloud storage

Criteria	Azure Blob Storage	Amazon S3	Google Cloud Storage
Performance	designed to handle high-throughput and low-latency data. In most regions, standard storage accounts can handle up to 20,000 requests per second. Certain regions support up to 40,000 requests per second [32].	Designed for high throughput and low latency. It supports concurrent data access and can easily handle large volumes of data[11]. Amazon S3 automatically adjusts to handle high volumes of requests. For instance, an application can perform at least 3,500 PUT, COPY, POST, or DELETE operations or 5,500 GET and HEAD requests per second for each partitioned S3 prefix. For instance, employing 10 prefixes can scale read performance to approximately 55,000 GET requests per second. Small Object Retrieval: Consistent latencies of approximately 100–200 milliseconds [33].	Offers low latency and high throughput. It provides consistent performance across various storage classes [21]. Cloud Storage buckets start with a default input/output capacity of 1,000 write and 5,000 read operations per second. Cloud Storage automatically scales as demand increases to support higher request rates [34].
Scalability	capable of storing petabytes of data [32].	It provides virtually unlimited storage capacity [11].	Scale seamlessly, accommodating any amount of data [12].
Availability & Durability	Designed for 99.9%-99.99% availability for geo-redundant storage (GRS) [35].	Designed for 99.999999999% (11 nines) durability and 99.99% availability annually[11].	>99.99% availability in multi-regions and dual-regions [21].
Security	Microsoft Entra ID for authentication supports role-based access control	AWS Identity and Access Management (IAM). Data encryption, and S3 Block Public	Google Cloud's Identity and IAM for fine-grained access control. Encryption

	(RBAC), including encryption and advanced threat protection [36].	Access [11].	[37].
Cost Efficiency	offers multiple storage tiers to optimise costs based on data access patterns. The pay-as-you-go price for the first 50 terabytes (TB) / month for Hot Data storage is $0.018 per GB [36].	Amazon S3 offers a range of storage classes tailored to different access patterns and cost requirements. The pay-as-you-use price for the first 50 terabytes (TB) / month for S3 Intelligent – Tiering, Frequent Access Tier is $0.023 per GB [38].	Google Cloud Storage provides multiple storage, allowing you tooptimise costs based on data access patterns. Standard storage (per GB per Month) is $0.023 [21].
Integration	Azure Data Factory, Azure Databricks, and Azure Synapse Analytics [39].	Amazon EC2, AWS Lambda, and Amazon EMR [11].	BigQuery and Vertex AI [12].
Usability	Azure portal, Azure Storage Explorer, REST APIs, PowerShell, Azure CLI and Azure Storage client libraries (.NET, Java, Node.js, Python and Go) [39].	The management console, SDKs (Java, Python, Ruby, .NET, iOS, Android) and a REST API [11].	Google Cloud Console, command-line tools, RESTful APIs, and client libraries (C++, C#, Go, Java, Node.js, PHP, Python, and Ruby) [37].
Data Management	data lifecycle management, versioning, soft deleting, and immutable storage [36].	Versioning, lifecycle policies, and cross-region replication [11].	Object versioning, lifecycle management, and Object Change [37].

3.2. Block Storage

Block storage divides data into fixed-size "blocks" that are handled separately. Blocks are updatable independently with high performance and minimal latency.

Storage of this type is typically used for direct data access applications such as virtual machines and databases. Block storage is suitable for workloads with high throughput and low latency and is applicable for transactional databases, enterprise applications, and HPC systems [31].

The unique features of block storage systems such as Azure Managed Disks, AWS EBS, and Google Persistent Disks are elaborated in Table 3, as well as their high performance and scalability for computationally demanding workloads.

Table 3. Review of block storage services: Azure managed disks, AWS EBS, and Google persistent disks

Criteria	Azure Managed Disks	AWS Elastic Block Store (EBS)	Google Persistent Disks
Performance	offer various types to cater to different performance needs. For Ultra Disk Storage, it provides: IOPS: Up to 400,000 IOPS per disk. Throughput: Up to 10,000 MB/s per disk. Latency: Sub-millisecond latencies [40].	Amazon EBS offers multiple volume types to cater to diverse workload requirements [41] For EBS Provisioned IOPS SSD (io2 Block Express): IOPS: Up to 256,000 IOPS per volume. Throughput: Up to 4,000 MB/s per volume. Latency: sub-millisecond [42].	Disks offer consistent performance, with data distributed across multiple physical disks [15] For Extreme Persistent Disk (A3 VMs): IOPS: Up to 400,000 IOPS per disk. Throughput: Up to 8,000 MB/s per disk. Latency: Sub-millisecond latencies [43].
Scalability	support up to 50,000 VM disks per region per subscription [13].	Supports volumes up to 64 TB. The Elastic Volumes feature enables you to increase capacity, adjust performance, and change volume types without downtime [41].	It can create persistent disks up to 64 TB in size and attach multiple disks to a single VM [44].
Availability &	Designed for 99.999%	Provide 99.8% to 99.999%	Provide 99.99% to

Durability	availability and least 99.9999999999% (12 9's) of durability [13].	durability and 99.9% to 99.99% availability [14].	99.9999% durability [44].
Security	Azure Active Directory for access control encryption at rest [13].	IAM for access control encryption of data at rest and in transit [14].	Data is encrypted at rest and in transit-customer-managed encryption keys for additional control [15].
Cost Efficiency	allows selection based on performance and cost requirements. Standard SSDs are priced at approximately $0.0959 per GB per month [45].	It allows you to balance performance and cost according to your workload requirements. You pay only for the storage you provision. Pricing starts at $0.10 per GB per month for general-purpose SSD volumes [41].	It allows you to balance performance and cost. Standard Persistent Disks are priced at $0.04 per GB per month [46].
Integration	Virtual Machine Scale Sets and Availability Sets [13].	Amazon EC2, AWS Backup and AWS Data Lifecycle Manager [14].	Compute Engine and Kubernetes Engine [15].
Usability	Azure Storage REST API, Azure CLI, Azure PowerShell, and Azure storage client library [10].	AWS CLI, AWS Management Console and AWS SDKs [41].	Google Cloud Console, gcloud CLI, or REST API [47].
Data Management	Supports snapshots and images and Shared disks [13].	Supports point-in-time snapshots for data backup. Integration with AWS Data Lifecycle Manager [14].	Supports snapshots for data backup and recovery [48].

3.3. File Storage

File storage has a hierarchical folder and file structure with shared access to data across users or systems. It accommodates legacy protocols such as SMB and NFS and is suitable for enterprise workloads and collaborative environments. File storage supports efficient file sharing with numerous users or applications with concurrent access through locking [31]. It is commonly used for home directories, media rendering workloads, and hosting applications. A good example of the flexibility of file storage services is Azure Files, Amazon EFS, and Google Filestore. Table 4 compares their ability to provide shared access to structured data across platforms.

Table 4. Review of file storage services: Azure files, Amazon EFS, and Google filestore

Criteria	Azure Files	Amazon EFS	Google Filestore
Performance	Azure Files offers multiple performance tiers-Premium, TransactionOptimised, Hot, and Cool-to cater to various workload requirements. These tiers provide flexibility in balancing performance and cost [16].	Amazon EFS offers high throughput and low latency, supporting hundreds of thousands of I/O operations per second (IOPS) and tens of gigabytes per second of throughput. It provides two performance modes: General Purpose, suitable for latency-sensitive use cases, and Max I/O, designed for highly parallelised workloads [49].	Filestore offers multiple service tiers-Basic, Zonal, Regional, and Enterprise-to cater to various performance needs. The Zonal tier, for instance, supports capacities up to 100 TiB with throughput up to 25 GB/s and 920K IOPS, suitable for high-performance computing and data-intensive applications [50].
Scalability	can handle large-scale deployments, supporting up to 100 TiB per share [16].	Scales as you add or remove files, accommodating petabytes of data without manual intervention [17].	Support for capacities ranging from 1 TiB to 100 TiB per instance [50].
Availability & Durability	ensures a high Availability of 99.9% with geo-redundancy [16].	Regional file systems are designed for 99.99% availability and 99.999999999% durability [17].	Offers a 99.99% regional availability [18].
Security	The service integrates Azure Integration with Active Directory	IAM, encryption, and Active Directory support [17].	IAM for fine-grained access control. Data is encrypted at

	RBACnd, which supports rest and transit encryption[16].		rest and in transit [51].
Cost Efficiency	With multiple pricing tiers, Azure Files allows you to optimise costs based on your performance and access requirements [52]. Standard SSD premiums are priced at $0.16 per GB per month.	Cost varies based on EFS storage classes [17]. Standard storage is priced at $0.3 per GB per month	Filestore allows you to choose the appropriate balance between performance and cost [50]. Basic SSD (Premium) storage costs $0.3 per GB per month.
Integration	integrates seamlessly with Azure services such as Azure Kubernetes Service (AKS) and Azure Virtual Desktop. It also supports hybrid scenarios through Azure File Sync, enabling on-premises access and synchronisation of file shares [16].	Integrates seamlessly with Amazon EC2, AWS Lambda, Amazon ECS, and Amazon EKS. It also supports on-premises access via AWS Direct Connect or VPN [17].	Integrates seamlessly with Google Kubernetes Engine (GKE). It also supports applications running on Compute Engine VMs and other Google Cloud services [51].
Usability	Azure Files supports SMB and NFS protocols, allowing simultaneous access from cloud and on-premises Windows, Linux, and macOS systems. This ensures seamless integration and easy cross-platform file sharing [16].	As a fully managed service, EFS eliminates the need for provisioning, deploying, patching, or maintaining file system infrastructure. It supports standard file system semantics, making it easy to use with existing applications [17].	Managed through the Google Cloud Console, gcloud CLI, or REST API. Its compatibility with standard NFS protocols facilitates easy integration into existing workflows [53].
Data management	Features such as share snapshots and Azure File Sync provide robust data management capabilities, including backup, recovery, and on-premises caching [16].	EFS supports features like lifecycle management. It also integrates with AWS Backup for centralised backup management and supports replication [17].	It supports features like backups and snapshots. You can schedule regular backups and restore data as needed [54].

3.4. Archive Storage

The storage of archives is optimised for long-term preservation at minimal cost. This storage is suitable for data accessed infrequently but needed for preservation for regulatory compliance purposes. Solutions for archives are for "write-once, read-rarely" operations with retrieval times of minutes to hours depending on the service and tier. Compliance files, historical data, and media archives are typical uses [31]. These platforms leverage advanced data compression and deduplication to provide maximum storage efficiency. They also have backup and disaster recovery processes for secure, resilient, and cost-effective long-term storage. Solutions such as Azure Archive Storage, AWS Glacier, and Google Cloud Archive Storage are some of the secure, resilient, and scalable architectures required for infrequent data access. Table 5 explains their architecture for infrequent data access and compliance with archiving standards.

Table 5. Review of archive storage services: Azure archive storage, AWS glacier, and Google cloud archive storage

Criteria	Azure Archive Storage	AWS Glacier	Google Cloud Archive Storage
Performance	Azure Archive Storage offers flexible latency, typically requiring several hours for data retrieval [55].	Retrieval times range from minutes to hours [20].	Archive Storage provides millisecond access latency [21].
Scalability	provides virtually unlimited scalability [56].	It provides virtually unlimited scalability [57].	Archive Storage offers virtually unlimited scalability [21].
Availability & Durability	99.9% SLA for Archive Storage; 99.999999999% durability [35].	Designed for 99.99% availability and 99.999999999% (11 nines) durability [57].	Designed for 99.999999999% (11 nines) durability. It offers availability SLAs of up to 99.95% depending on the chosen storage location type [21].

Security	Data is encrypted both at rest and in transit. Azure Archive Storage integrates with Azure Active Directory for access control [56].	S3 Glacier integrates with AWS CloudTrail to log, monitor, and retain storage API call activities for auditing purposes. It supports multiple forms of encryption to protect data at rest and in transit [57].	Data is encrypted both at rest and in transit. Integration with Google Cloud's IAM allows fine-grained access control [21].
Cost Efficiency	provides storage at $0.00099 per GB per month [58].	Provides storage at $0.0036 per GB / Month [59].	Storage costs start at $0.0012 per GB per month [60].
Integration	Seamlessly integrates with Azure services such as Azure Blob Storage [55].	S3 Glacier integrates seamlessly with most AWS services [20].	Integrates with other Google Cloud services, including BigQuery and Vertex AI [12].
Usability	Managed through the Azure portal, Azure CLI, PowerShell, REST APIs, and SDKs [19].	Managed through the AWS Management Console, AWS CLI, and SDKs [20].	Managed through the Google Cloud Console, CLI, and APIs [61].
Data Management	Supports features like bulk data archiving and rehydration, enabling efficient management of large datasets [19].	Supports features like object tagging and S3 Lifecycle configurations [20].	It supports features like object versioning and lifecycle policies [61].

3.5. Hybrid Storage

Hybrid storage solutions bridge the gap between cloud and on-premises environments to enable easy management and unification of data across environments-hybrid systems enable organisations to satisfy data residency regulations and enjoy cloud scalability [62].

Hybrid systems are most appropriate for workloads with local data processing requirements such as financial services, healthcare, and IoT.

Hybrid solutions enable single management interfaces with consistent policies for cloud and on-premises environments. Hybrid storage optimises performance and cost by enabling low-latency access to highly accessed data and offloading less critical workloads to the cloud. Azure Data Box Gateway, AWS Storage Gateway, and Google Anthos Storage are some of the most widely used hybrid storage systems that integrate on-premises environments with the cloud. Table 6 describes their hybrid integration features and performance.

Table 6. Review of hybrid storage systems: Azure data box gateway, AWS storage gateway, and Google anthos storage

Criteria	Azure Data Box Gateway	AWS Storage Gateway	Google Cloud Anthos Storage
Performance	Data Box Gateway facilitates efficient data transfer to Azure Storage by supporting standard protocols such as SMB and NFS. It includes a local cache to accommodate high data ingestion rates during peak business hours, ensuring that data is uploaded to Azure without overwhelming network resources [22].	AWS Storage Gateway offers low-latency access to data by caching frequently accessed data locally, ensuring quick response times for on-premises applications. It supports standard storage protocols such as NFS, SMB, and iSCSI [23]. Performance varies based on the protocol used and the number of concurrent threads: SMBv3 Protocol: Single thread: Up to 265 MiB/s (approximately 2.2 Gbps) read throughput. Eight threads: Up to 780 MiB/s (approximately 6.5 Gbps) read throughput. NFSv3 Protocol: Single thread: Up to 220 MiB/s (approximately 1.8 Gbps) read	Anthos supports high-performance storage solutions through integration with Kubernetes-native storage APIs, allowing for dynamic provisioning of volumes. This setup enables applications to achieve low-latency and high-throughput access to data, depending on the underlying storage system used. For instance, integrating high-performance storage backends can support workloads requiring substantial IOPS and throughput [24].

		throughput. Eight threads: Up to 570 MiB/s (approximately 4.8 Gbps) read throughput [63].	
Scalability	The virtual appliance supports continuous data ingestion, making it ideal for large-scale transfers. It handles one-time bulk transfers and ongoing incremental uploads, offering flexible data transfer options.[22].	The service provides virtually unlimited cloud storage capacity [23].	Anthos facilitates scalable storage by supporting dynamic provisioning and management of storage resources across clusters [24].
Availability & Durability	Data Box Gateway integrates with Azure Storage, which offers high availability and durability for stored data [22].	AWS Storage Gateway stores data in Amazon S3, which is designed for 99.999999999% (11 nines) durability [23].	The Availability and durability of storage in Anthos depend on the underlying storage solutions integrated with the platform [24].
Security	The solution supports authentication to control access to the device and data, with data-in-flight encrypted using AES-256-bit encryption [22].	The service integrates with AWS IAM for access control and supports data encryption at rest using AWS KMS and in transit using SSL. It also supports compliance with various industry standards and regulations [23].	Anthos integrates, including (IAM) and encryption mechanisms, with support for encrypted data at rest and in transit, as well as fine-grained access controls [64].
Cost Efficiency	Data Box Gateway follows a subscription model with a monthly fee for service usage. Storage and transaction costs are billed separately, allowing organisations to manage expenses based on data transfer and storage needs [65].	AWS Storage Gateway helps reduce costs by minimising the need for on-premises storage infrastructure. You pay only for the storage [23].	Pay-as-you-go model, reducing costs by integrating hybrid cloud and on-prem solutions. The cost efficiency of storage within Anthos is influenced by the choice of integrated storage solutions and their respective pricing models [24].
Integration	The solution integrates seamlessly with Azure services, enabling data transfer to Azure Blob Storage and Azure Files. It supports standard SMB and NFS protocols, facilitating easy integration with existing on-premises systems [22].	The service integrates seamlessly with various AWS services, including Amazon S3, Amazon EBS, Amazon S3 Glacier, AWS Backup, Amazon CloudWatch, and AWS CloudTrail [23].	Anthos is designed to integrate seamlessly with various storage solutions, both on-premises and in the cloud [24].
Usability	Data Box Gateway is deployed as a virtual appliance in your virtualised environment or hypervisor. Management is facilitated through a local web UI for initial setup and diagnostics and the Azure portal for day-to-day management, providing a user-friendly experience [22].	AWS Storage Gateway is easy to deploy and manage, offering a consistent management experience through the AWS Management Console. It supports deployment as a virtual machine, hardware appliance, or in AWS as an Amazon EC2 instance, providing flexibility to suit different environments [23].	Managed using Google Cloud Console, CLI, and APIs for hybrid environments [64].
Data Management	supports continuous data ingestion and can be used with Azure Data Box for initial bulk transfers followed by	It supports automated backups, data archiving to Amazon S3 Glacier and Amazon S3 Glacier Deep Archive, and integration	Supports data management features such as automated policy enforcement and configuration management across clusters. By

	incremental transfers. It also allows for data refresh, enabling local files to be updated with the latest from the cloud [22].	with AWS Backup for centralised backup management. It also provides monitoring and logging capabilities through Amazon CloudWatch and AWS CloudTrail [23].	integrating with storage solutions that offer capabilities like snapshots and backups, Anthos enables comprehensive data protection and management strategies [24].

3.6. Edge/Offline Storage

Edge storage brings compute and storage resources closer to the data source, reducing latency and enabling offline operations. It is particularly useful for remote setups or IoT environments. Edge storage facilitates real-time data processing in environments with minimal connectivity, such as oil rigs, ships, and remote construction sites. It is closely integrated with edge computing devices so that data is pre-processed before synchronising with core cloud repositories [62]. Edge and offline storage solutions such as Azure Stack Edge, AWS Snow Family, and Google Transfer Appliance address special challenges in remote and disconnected environments. Their functionalities are compared in Table 7 to highlight their applicability in edge computing.

Table 7. Review of edge and offline storage solutions: Azure stack edge, AWS snow family, and Google transfer appliance

Criteria	Azure Stack Edge	AWS Snowball Edge StorageOptimised	Google Transfer Appliance
Performance	Azure Stack Edge devices are equipped with hardware accelerators, such as NVIDIA GPUs or Intel VPUs, to facilitate high-performance computing and machine learning inference at the edge [25].	High-performance storage for disconnected operations and edge computing. The device supports high-speed data transfer with 10 Gbit RJ45, 25 Gbit SFP28, and 100 Gbit QSFP28 interfaces. It includes 40 vCPUs and 80 GiB of memory for efficient edge computing and data processing [26].	The appliance uses all-SSD storage for reliability and offers 10 Gbps RJ45 and 40 Gbps QSFP+ interfaces for fast data transfer. While 300 TB might take 9 months over a 100 Mbps network, the appliance completes transfer and cloud upload in under 50 days without using outbound bandwidth [27].
Scalability	The solution offers various device models, including Azure Stack Edge Pro Series and Azure Stack Edge Mini Series, to cater to different workload requirements. These devices can be deployed individually or in clusters, providing scalability to meet organisational needs [25].	Snowball Edge Storage Optimised devices can be clustered, enabling scalable storage and compute capacity to meet varying workload demands. This clustering capability allows for flexible scaling in response to data growth and processing requirements [66].	Transfer Appliance supports scalability by allowing the use of multiple appliances to increase transfer speed, accommodating large-scale data migration needs [27].
Availability & Durability	Azure Stack Edge devices are designed for high availability, with options for ruggedised models suitable for harsh environments. Data can be cached locally and synchronised with Azure Storage, ensuring durability and resilience [25].	The device is designed for rugged environments, ensuring data integrity during transit. Data is stored redundantly across multiple devices when clustered, enhancing durability and Availability [66].	The appliance is built for secure data transfer, with Google Cloud ensuring durability. EU customers receive and return appliances in Belgium, where data is uploaded to a chosen Cloud Storage region [27].
Security	The solution integrates with Azure's security features, including Azure Active Directory for identity management and encryption for data at rest and in transit. This integration ensures that data processed and stored on Azure Stack Edge devices adhere to stringent security standards [25].	Snowball Edge Storage Optimized integrates with AWS IAM for access control and supports encryption of data at rest and in transit. It also features physical security measures such as tamper-evident seals and Trusted Platform Modules (TPM) to ensure data security during transit [66].	Tamper Resistance: The device has a secure casing and tamper-evident tags for detecting unauthorised access. TPM Chip: Ensures the core system remains unchanged and secure. Hardware Attestation: Verifies the device's integrity before connection. Encryption: Uses AES-256 with

			customer-managed keys for complete security control. Data Erasure: Wipes data per NIST 800-88 standards, with a certificate available upon request [27].
Cost Efficiency	Azure Stack Edge uses a subscription model with monthly fees based on device configuration, helping organisations control costs by selecting devices that meet their needs [67].	By physically transporting data, Snowball Edge Storage Optimised reduces the time and cost associated with large-scale data transfers over the Internet. Pricing varies based on the device and duration of use [66].	Transfer Appliance provides a budget-friendly way to move large datasets, ideal for limited bandwidth or tight deadlines. Shipping the device speeds up transfers and reduces network costs [27].
Integration	The solution works smoothly with Azure services, allowing the deployment of containerised apps and virtual machines at the edge. It also supports Azure IoT Edge for IoT solutions [25].	The device integrates seamlessly with AWS services, enabling data to be imported into Amazon S3 or used with Amazon EC2 instances for edge computing workloads [66].	The appliance connects easily with Google Cloud, which allows data to be uploaded to your chosen Cloud Storage bucket. It supports NFS for Linux/macOS and SCP or SSH for Windows transfers [27].
Usability	Azure Stack Edge devices are designed for ease of deployment and management. They can be ordered through the Azure portal and are managed using standard Azure tools, allowing for a consistent experience across cloud and edge environments [25].	WS provides AWS OpsHub, a graphical user interface, to simplify the management and monitoring of Snowball Edge devices, making it easier to deploy edge computing workloads and migrate data [66].	Selection: Google helps you choose the right appliance. Data Upload: Linux/macOS uses NFS, while Windows uses SCP or SSH. Shipping: After transfer, you return the sealed appliance to Google. Data Availability: Google uploads your data and securely wipes the device[27].
Data Management	It supports local caching, bandwidth throttling, and automatic data synchronisation [25].	Snowball Edge Storage Optimized supports the clustering of devices for increased storage capacity and durability and integrates with AWS services for data processing and analysis [66].	Enabling online mode allows for streaming data directly to your Cloud Storage bucket after copying it to the appliance, facilitating quicker data transfers with low latency [27].

3.7. Backup Storage

Backup storage offers protection and recovery of data in the event of failures and disasters. It maximises recovery point objectives (RPOs) and recovery time objectives (RTOs) with automated snapshots and replications. It offers flexible retention policies, tiering for cost optimisation, and encryption for data security.

Use cases are disaster recovery, ransomware protection, and archiving of confidential business information [68]. The critical role of backup storage is demonstrated through Azure Backup, AWS Backup, and Google Cloud Backup, which have centralised data protection and recovery solutions. Their functionalities are elaborated in Table 8, focusing on performance, Security, and integration.

Table 8. Review of backup storage services: Azure backup, AWS backup, and Google cloud backup

Criteria	Azure Backup	AWS Backup	Google Cloud Backup
Performance	Azure Backup provides efficient data transfer and storage with compression, encryption, long retention, and low maintenance features. It supports backups for Azure VMs, SQL Server databases,	AWS Backup offers efficient, policy-driven backups with features like incremental backups, which capture only changes since the last backup, optimising both performance	The service uses "incremental-forever" backups with changed-block tracking to speed up backups and reduce system impact. This

	and on-premises resources [28].	and storage utilisation [29].	ensures efficient storage use and supports low Recovery Point Objectives (RPOs) [30].
Scalability	The service is designed to handle increasing data volumes, providing scalable short-term and long-term data retention solutions without deploying complex on-premises backup solutions [28].	Designed to handle growing data needs, AWS Backup supports a wide range of AWS services and on-premises resources, allowing for seamless scaling as your backup requirements evolve [29].	Designed to handle growing data needs, Google Cloud Backup and DR support a wide range of workloads, allowing for seamless scaling as your backup requirements evolve [30].
Availability & Durability	99.9% SLA for Azure Backup [69].	99.99% availability with cross-region replication for backups [29].	Geo-redundant backup storage with 99.9% SLA for Google Cloud Backup [30].
Security	The service provides Security for your backup environment when your data is in transit and at rest. It includes features like multi-user authorisation for critical operations, alerting and monitoring, and recovery point encryption using customer-managed keys [70].	AWS Backup encrypts data at rest and in transit, integrating with AWS Key Management Service (KMS) for encryption key management. It also offers features like AWS Backup Vault Lock for write-once-read-many (WORM) protection [29].	Data is encrypted both at rest and in transit, ensuring robust Security. Integration with Google Cloud's Identity and Access Management (IAM) allows fine-grained access control [30].
Cost Efficiency	Azure Backup offers a pay-as-you-go pricing model, allowing you to pay only for the storage you consume. It provides options for storage redundancies and supports features like compression and incremental backups to optimise storage costs [69].	With a pay-as-you-go pricing model, AWS Backup allows you to manage costs effectively. You pay only for the backup storage you use, data transferred between AWS Regions, data restored, and backup evaluations. There are no minimum fees or setup charges [69].	With a pay-as-you-go pricing model, Google Cloud Backup and DR allows you to manage costs effectively. Features like incremental backups help reduce storage expenses, and options for different storage classes can provide additional savings [30].
Integration	The service smoothly connects with multiple Azure offerings, such as Azure Site Recovery for disaster recovery and Azure Monitor for tracking performance and sending alerts [28].	It effortlessly connects with multiple AWS services, such as Amazon EBS, Amazon RDS, Amazon DynamoDB, Amazon EFS, and AWS Storage Gateway [29].	Seamlessly integrates with various Google Cloud services, including Google Cloud Compute Kubernetes [30].
Usability	managed through the Azure portal. It also offers PowerShell and REST APIs for automation and scripting [71].	Managed through the AWS Management Console, AWS CLI, and SDKs [29].	Managed through the Google Cloud Console, CLI, and APIs [30].
Data Management	The service supports long-term retention, application-consistent backups, and centralised monitoring and management [28].	AWS Backup supports features like lifecycle management. It also provides centralised monitoring and reporting capabilities [29].	Supports features like lifecycle management. It also provides centralised monitoring and reporting capabilities [30].

4. Weighted Scoring Model

Weighted Scoring Model (WSM) is one of the methods of comparing and selecting from options based on predefined criteria. It is commonly applied to compare different algorithms. Application of the method typically follows a series of steps, as depicted in Figure 2 [72]:

- Step 1: Identify the criteria to be applied for each option.
- Step 2: Assign weights for the criteria as percentages of their priority.
- Step 3: Create a table listing the nominal values for all criteria across different options.
- Step 4: Create a table reflecting weighted values for each criterion with weights as percentages, ensuring the total weight sums to 100%.
- Step 5: Calculate weighted scores as the product of the nominal values and the weights for each option.

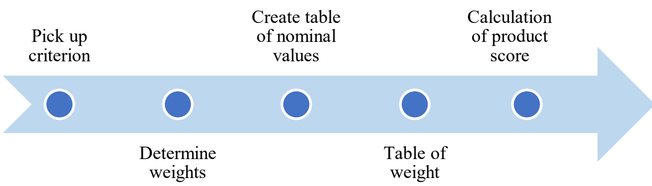

Fig. 2 WSM process [72]

5. Comparison Of Multiple Storage Types

5.1. Comparison Criteria

The criteria for comparing various cloud storage types were thoughtfully established through detailed analyses of their functionalities and performance metrics. They come from common challenges in cloud storage, enabling a comprehensive and practical evaluation to identify the most appropriate storage solution for specific requirements.

- Performance: Measures the speed at which data can be read and written, latency, and Input/Output Operations Per Second (IOPS). This criterion determines how fast and responsive the storage service is for data operations [62].
- Scalability: Indicates the service's ability to efficiently handle large volumes of data and its capacity to scale elastically based on demand without compromising performance [62].
- Durability and Availability: Refers to the likelihood that data will not be lost (durability) and the percentage of time the service is available (Availability), often expressed through Service Level Agreements (SLAs) and replication techniques [62].
- Security: Involves the measures taken to protect data, including encryption, role-based access control (RBAC), and compliance with industry regulations or certifications (e.g., GDPR, HIPAA) [62].
- Cost Efficiency: Refers to the pricing structure for storage, data transfer, and operations such as requests (GET, PUT). It evaluates the service's affordability based on the data volume and the operations frequency [62]. The cost section includes only the cloud provider's storage price, excluding operation fees.
- Integration: Evaluate the service's ability to integrate smoothly with other cloud services, analytics platforms, machine learning tools, or enterprise workloads to enable comprehensive solutions [62].
- Usability: Focuses on the accessibility and simplicity of managing the service, including the Availability of APIs, Software Development Kits (SDKs), and user-friendly management tools.
- Data Management: organising, storing, securing, and optimising data accessibility and usage in scalable cloud environments, ensuring compliance, reliability, and efficient lifecycle management.

5.2. Application of Weighted Scoring Model

The Weighted Scoring Model compares cloud storage solutions against scores for various factors, with each factor being assigned weights according to their priority, as shown in Table 9. Performance and Security are each weighted at 20% due to their priority in delivering fast access to data and protecting confidential information, respectively. Scalability, Availability, and durability are weighted at 15% due to their priority in allowing the solution to scale and be reliable. Cost Efficiency, Integration, Usability, and Data Management are weighted less due to their lower priority. This makes it easy to closely analyse cloud storage solutions so that businesses can select the best solution for their needs. With the weighted scoring model in Table 10, each value shall be scored according to their priority using a scale of 1 to 5 (n/a = 0). After scoring, each is multiplied by their corresponding weights, as outlined in Table 9. The final values are shown in Table 11. Scoring in Table 10 is based on the data in Tables 2, 3, 4, 5, 6, 7, and 8.

Table 9. Associated weight to each criterion

Criterion	Abbreviation	Proposed weight
Performance	PR	20%
Scalability	SC	15%
Availability & Durability	AD	15%
Security	SE	20%
Cost Efficiency	CE	10%
Integration	IN	5%
Usability	US	5%
Data Management	DM	10%

Table 10. Storage type comparison

Cloud solution	Storage type	PR	SC	AD	SE	CE	IN	US	DM
Azure Blob Storage	Object storage	4	5	4	5	5	5	5	5
Amazon S3	Object storage	5	5	5	5	4	5	4	5

Google Cloud Storage	Object storage	4	5	5	5	4	5	5	5
Azure Disk Storage	Block storage	5	4.5	5	4.5	4	5	5	4.5
Amazon EBS	Block storage	4.5	5	4	5	4	5	5	5
Google Persistent Disk	Block storage	5	5	4.5	4.5	5	5	4.5	4.5
Azure Files	File storage	4	4	4	4.5	4	4.5	4.5	4.5
Amazon EFS	File storage	5	5	5	4.5	3.5	5	5	4
Google Filestore	File storage	5	4	4.5	4.5	3.5	4.5	4.5	4
Azure Blob Storage (Archive)	Archive storage	3	5	4.5	4	5	3	5	4
Amazon S3 Glacier	Archive storage	3	5	5	5	4	4	4	4
Google Cloud Archive	Archive storage	4	5	4.5	4	4.5	4	5	5
Azure Data Box Gateway	Hybrid storage	3	4	4	5	4	5	3	4
AWS Storage Gateway	Hybrid storage	4	5	4.5	5	4	5	3.5	5
Google Cloud Anthos Storage	Hybrid storage	3	4	4.5	5	4	4	4	5
Azure Backup	Backup storage	3	5	4	5	4	4	5	5
Amazon Backup	Backup storage	3	5	5	5	4	5	5	5
Google Cloud Backup and DR	Backup storage	3	5	5	5	4	4	5	5
Azure Stack Edge	Edge/ offline Storage	4	3	4	5	4	5	5	4
AWS Snowball Edge StorageOptimised	Edge/ offline Storage	5	4	4	5	4	5	5	4
Google Transfer Appliance	Edge/ offline Storage	4	4	4	5	4	5	4	3.5

5.3. Comparison

All solutions perform well, as shown in Figure 3, especially in the criteria of scalability, Security, and data management, where they all received top scores of 5/5. With a 5/5 rating for cost-effectiveness and usability, Azure Blob Storage is the most economical option. Additionally, Google Cloud Storage has a usability rating of 5/5, whereas Amazon S3 has a slightly lower rating of 4/5. Amazon S3 is the fastest option because of its superior performance. Although both Google Cloud Storage and Amazon S3 have availability ratings 5/5, Google Cloud Storage is comparatively more costly. Each of the three solutions integrates well with other services.

With a 5/5 rating for both performance and cost-effectiveness, Figure 4 shows Google Persistent Disk as the top option. It is a well-rounded choice because it produces excellent scalability, integration, and general usability results.

With high ratings for usability, Availability, and durability, as well as strong performance in other areas, particularly performance and integration, Azure Disk Storage performs admirably. With consistently high ratings in every category, Amazon EBS exhibits a well-rounded profile, thriving in scalability, Security, and data management. Because of its adaptability can be used for various tasks that call for robust Security and dependable integration.

Figure 5 demonstrates that Amazon EFS is the best file storage option overall, with flawless performance, scalability, Availability and durability, and usability ratings. Although it has a cost efficiency score of 3.5, indicating a higher price point, its robustness is influenced by its strong integration and high Security. With impressive Security, usability, data management, and cost-effectiveness outcomes, Azure Files is a well-rounded, affordable, and practical substitute. Unlike Amazon EFS, Google Filestore has a lower cost efficiency score, making it slightly more expensive than Azure Files despite its strong performance, Security, usability, and integration.

Figure 6 shows that Amazon S3 Glacier is the best durable and long-term archival storage solution. It received perfect scores in scalability, Availability, durability, and Security. However, its performance and usability scores are slightly lower at 3 and 4, respectively. Azure Archive Storage stands out for its cost efficiency and ease of use, which scored 5/5, making it an ideal option for budget-conscious users. Google Cloud Archive offers a balanced solution, with high scores of 5/5 in data management and usability. It also outperforms the others in performance, earning a score of 4 and making it the fastest archival option. All three solutions demonstrate excellent scalability but differ in usability and pricing, allowing users to choose based on their priorities.

Figure 7 illustrates the ranking of hybrid storage solutions, with AWS Storage Gateway scoring a 5 for scalability, data management, and Security and a 4.5 for Availability and durability. Its usability score of 3.5 is slightly lower, but its performance and integration capabilities balance it out. Azure Data Box Gateway is the most cost-efficient, scoring 4/5 and 5/5 for Security and integration, respectively.

However, its usability is the lowest among the three, with a score of 3/5, which could pose challenges during implementation. Google Cloud Anthos Storage scores a perfect 5 in Security and data management, demonstrating strength in these areas. It also has solid usability and Availability but scores lower in integration at 4, which reduces the flexibility to provide seamless interconnection of services.

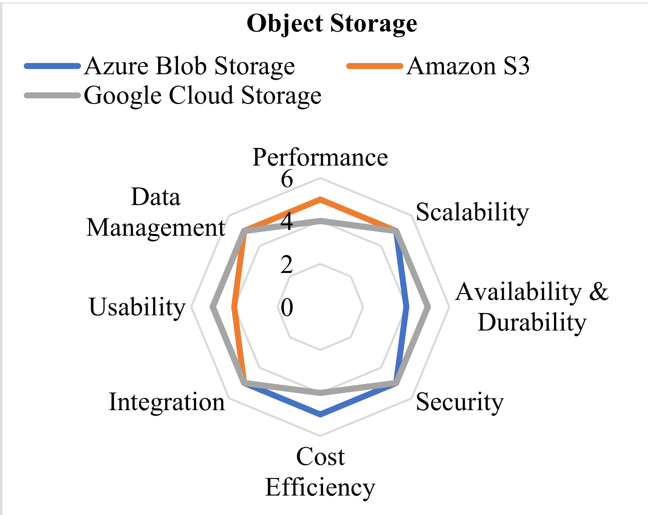

Fig. 3 Object storage comparison across key evaluation criteria

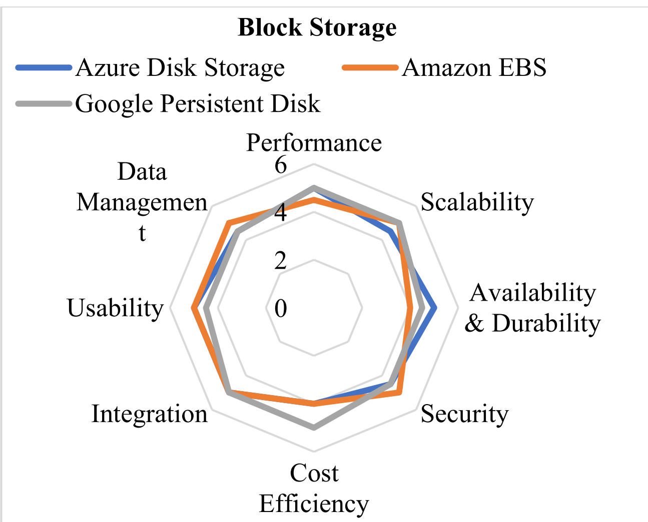

Fig. 4 Block storage comparison across key evaluation criteria

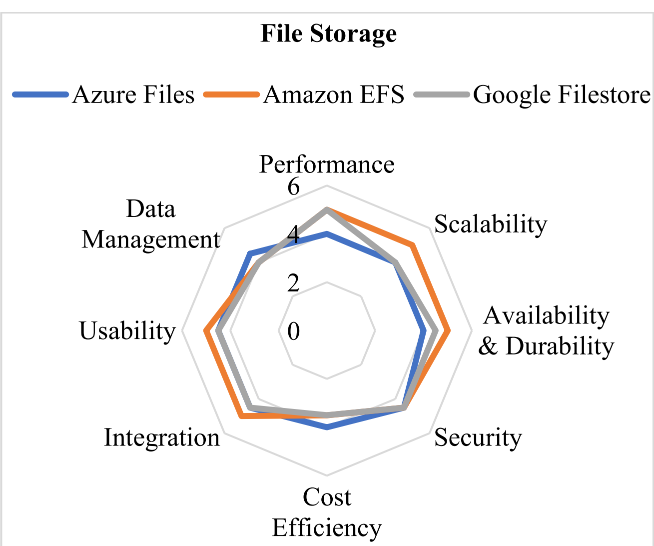

Fig. 5 File storage comparison across key evaluation criteria

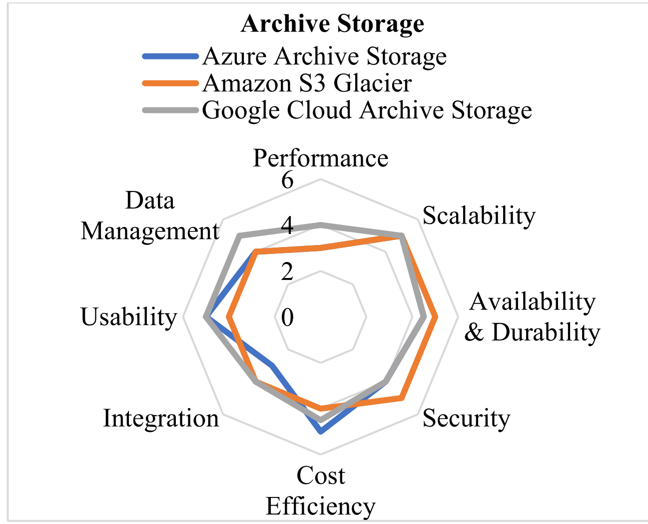

Fig. 6 Archive storage comparison across key evaluation criteria

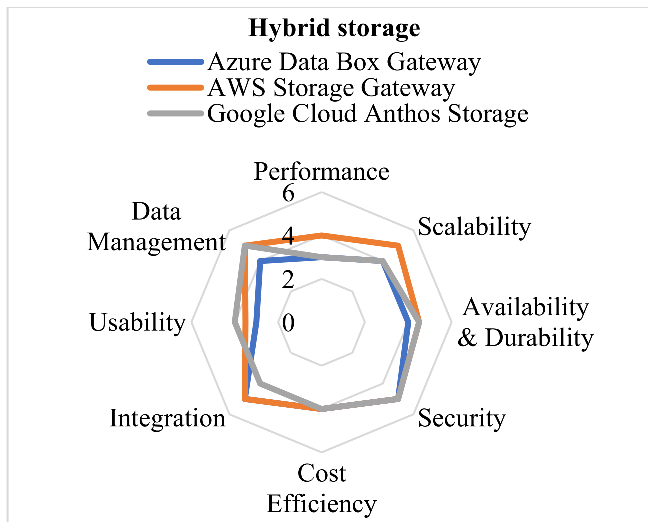

Fig. 7 Hybrid storage comparison across key evaluation criteria

Figure 8 shows that Azure Stack Edge is the top performer in Security and usability. It scores a perfect 5/5 in both categories, making it a particularly appealing option for secure, user-friendly edge deployments. However, its scalability score 3/5 may limit its effectiveness in highly dynamic environments. AWS Snowball Edge Storage Optimized offers a well-balanced profile. It leads in performance with a score of 5 and achieves solid results in scalability, Security, usability, and integration, all of which score 4 or above. Google Transfer Appliance stands out in Security and integration, scoring a 5. It maintains consistent scores across most other criteria. However, it falls slightly behind in data management with a score of 3.5. All three solutions are cost-efficient and well-suited for edge computing scenarios. Each solution has specific strengths depending on user priorities, such as performance, Security, or ease of use. Figure 9 ranks Amazon Backup as the top backup storage solution, giving it perfect scores in the categories of Availability and durability, integration, usability, and data management. Its strong scalability and security performance make it a highly reliable, well-integrated option. Azure Backup, Google Cloud Backup, and DR demonstrate strong performance across key criteria. Both services achieve top scores in scalability, Security, usability, and data management. However, they score slightly lower in integration, each receiving a 4/5 compared to Amazon Backup's 5/5. All three solutions offer consistent cost efficiency and meet high standards for backup reliability. However, Amazon Backup stands out for its seamless integration capabilities.

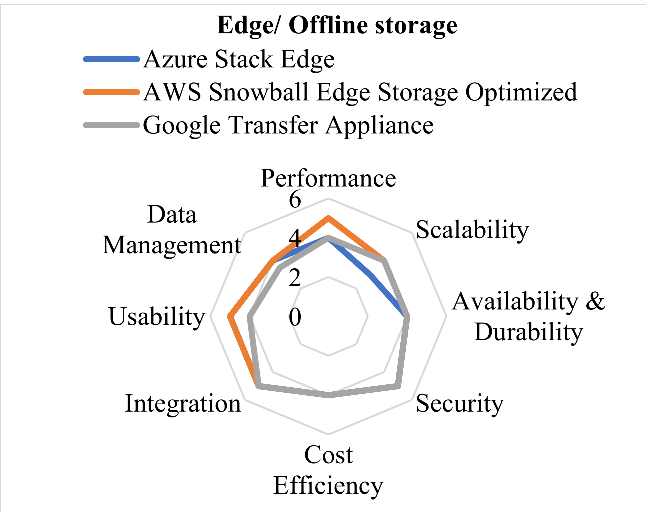

Fig. 8 Edge/ offline storage comparison across key evaluation criteria

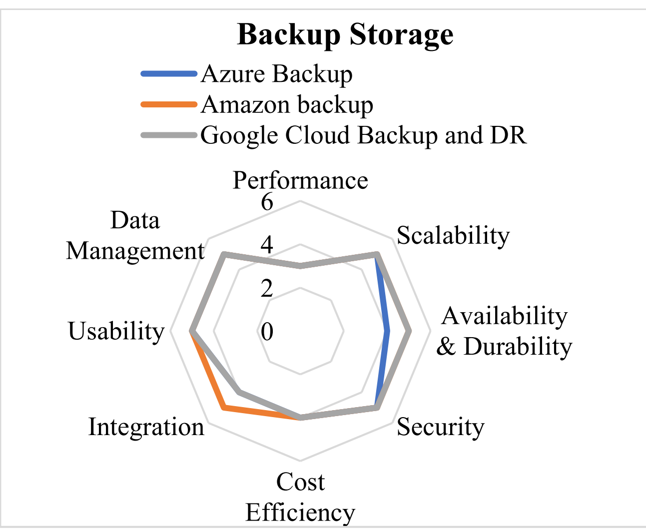

Fig. 9 Backup storage comparison across key evaluation criteria

Table 11. WSM results

Storage Type	Cloud service	WSM Score	WSM Score (%)
Object Storage	Azure Blob Storage	4.65	93
	Amazon S3	4.85	97
	Google Cloud Storage	4.7	94
Block Storage	Azure Disk Storage	4.675	93.5
	Amazon EBS	4.65	93
	Google Persistent Disk	4.75	95
File Storage	Azure Files	4.2	84
	Amazon EFS	4.65	93
	Google Filestore	4.375	87.5
Archive Storage	Azure Archive Storage	4.125	82.5
	Amazon S3 Glacier	4.3	86
	Google Cloud Archive Storage	4.425	88.5
Hybrid Storage	Azure Data Box Gateway	4	80
	AWS Storage Gateway	4.55	91
	Google Cloud Anthos Storage	4.175	83.5
Edge/ Offline Storage	Azure Stack Edge	4.15	83
	AWS Snowball Family	4.5	90
	Google Transfer Appliance	4.2	84
Backup Storage	Azure Backup	4.3	86
	Amazon back up	4.5	90
	Google Cloud Backup and DR	4.45	89

6. Results and Discussion

Comparing different storage services from Azure, AWS, and Google Cloud helps understand clearly their strong and weak points. This analysis looks at several storage types: object storage, archive storage, hybrid and edge storage, backup solutions, and file and block storage. It covers performance, scalability, Availability, Security, cost, integration, ease of use, and data management.

6.1. Object Storage Services

As shown in Figure 3, all three object storage services perform well. Amazon S3 ranks highest with a WSM score of 4.85 (97%), thanks to its excellent performance, scalability, and availability. However, its cost efficiency is slightly lower at 4/5, indicating a trade-off between speed and affordability. Google Cloud Storage follows with a 4.7 (94%) score, offering strong availability, usability, and integration. It performs well overall but scores slightly lower in cost efficiency and scalability. Azure Blob Storage stands out for its cost efficiency and usability, both of which receive a score of 5/5. Although affordable and user-friendly, its performance score is slightly lower (4/5), which could affect speed-critical applications. These results confirm that object storage is well-suited for unstructured data and that each provider offers specific strengths depending on performance, cost, or ease of use.

6.2. Block Storage Services

As shown in Figure 4, all three block storage services perform exceptionally well. Google Persistent Disk ranks highest with a WSM score 4.75 (95%). Its top-tier performance and cost efficiency make it ideal for high-throughput, latency-sensitive workloads. Azure Disk Storage follows closely behind with a score of 4.675 (93.5%), standing out for its strong Availability, usability, and integration. However, its scalability score of 4.5/5 suggests that it may not be ideal for extremely large-scale environments. Amazon EBS offers a well-balanced solution with a WSM score 4.65 (93%). It is particularly strong in scalability and Security but has a slightly lower cost efficiency rating of 4/5, indicating higher potential costs for extended usage. These results demonstrate that, although all three options are robust, the optimal choice depends on specific requirements, such as prioritising performance, cost, or seamless integration.

6.3. File Storage Services

As shown in Figure 5, Amazon EFS is the leader in the file storage category, earning a WSM score of 4.65 (93%). It offers excellent performance and scalability for enterprise-grade workloads. However, its cost-efficiency score of 3.5/5 is a potential drawback for budget-conscious users. Google Filestore follows closely behind with a score of 4.375 (87.5%). It performs well in performance and integration, particularly for use cases like media processing and collaborative environments. However, it is also limited by moderate cost efficiency. Azure Files ranks third with a WSM score 4.2 (84%). It stands out for its strong integration with Microsoft services and better cost efficiency. However, it scores lower in scalability and performance, 4 and 4.5, respectively. These results suggest that Amazon EFS best suits high-performance shared file systems. At the same time, Azure Files is a better option for applications with moderate performance needs and cost-sensitive setups.

6.4. Archive Storage Services

Figure 6 shows archive storage solutions designed for low-cost, long-term data retention reveal distinct trade-offs. Google Cloud Archive Storage takes the lead with a WSM score of 4.425 (88.5%). It benefits from robust data management and integration features, making it ideal for archival scenarios requiring occasional analysis or retrieval. Amazon S3 Glacier follows with a score of 4.3 (86%). It excels in durability and Security but has a lower integration score (4/5), which could affect hybrid or multi-cloud deployments. Azure Archive Storage ranks third with a score of 4.125 (82.5%). It is the most cost-effective option but scores lower in performance and usability (3/5). This limits its practicality for workflows that require frequent access or easy manageability. These results demonstrate that, although all three platforms support cost-effective long-term storage, integration can pose challenges for complex workflows, particularly with Azure Archive Storage.

6.5. Backup Storage Services

Figure 9 shows that backup storage solutions perform well overall. Amazon Backup achieved the highest WSM score of 4.5 (90%) due to its strong Availability, integration, and usability features. Google Cloud Backup and Disaster Recovery (DR) closely follow with a score of 4.45 (89%). They offer robust policy-based management and reliable encryption. However, their cost efficiency (4/5) may be a concern for large-scale or frequent backups. Azure Backup is the most cost-effective, with a 4.3 (86%) score, but the lowest performance (3/5). This may limit its suitability for environments requiring high-speed or frequent data restoration. These results indicate that, although all platforms provide secure backup capabilities, Amazon Backup is ideal for enterprises requiring high flexibility and automation. On the other hand, Azure offers a budget-friendly solution with some performance limitations.

6.6. Hybrid Storage Services

Figure 7 shows hybrid storage solutions, which enable seamless data movement between on-premises and cloud environments, have different strengths. AWS Storage Gateway is the leader with a WSM score of 4.55 (91%), thanks to its strong scalability, Availability, and data management capabilities. However, its usability rating 3.5/5 suggests that setup and management may be challenging. Google Cloud Anthos Storage follows with a score of 4.175 (83.5%). It offers solid security and integration features, but its performance score is lower (3/5), which could affect latency-sensitive operations. Azure Data Box Gateway is the most cost-effective option, with a WSM score 4.0 (80%). However, it scores the lowest in performance and scalability (3/5), which limits its effectiveness in dynamic or high-demand hybrid workloads. These results suggest that AWS is best suited for fully integrated hybrid infrastructures, while Azure and Google Cloud are more appropriate for basic hybrid use cases or phased migrations.

6.7. Edge/Offline Storage Services

As shown in Figure 8, edge and offline storage solutions are designed for environments with limited or intermittent connectivity. These solutions offer localised processing and data retention. The AWS Snowball Family leads with a WSM score of 4.5 (90%), excelling in offline computing capabilities and durability. However, its cost efficiency score of 4/5 suggests it may be costly for long-term or large-scale use. Google Transfer Appliance follows with a score of 4.2 (84%). It is well-suited for bulk offline data migration but scores lower in data management (3.5/5). This indicates limited flexibility in handling complex workflows. Azure Stack Edge performs well in Security and usability with a WSM score of 4.15 (83%), but its lower scalability score of 3/5 could restrict its use in rapidly expanding edge deployments. While all three platforms address key edge and offline needs, AWS is optimal for compute-intensive edge use cases. Google excels in high-volume transfers, and Azure offers a practical solution for smaller-scale, intelligent edge applications.

6.8. Comparative Evaluation with Existing Techniques

In this subsection, we discuss the key improvements to our approach that helped us achieve better results than previous methods in the literature.

Our approach delivers a comprehensive and objective assessment of cloud storage services by applying the Weighted Scoring Model (WSM), extending the evaluation across multiple critical criteria, and rigorously quantifying each aspect. Unlike prior studies, which often relied on fewer criteria or generalised qualitative comparisons, our method incorporates a broad set of essential factors, including performance, scalability, Availability, Security, cost efficiency, integration capabilities, usability, and data management. We tailored the weights to reflect realistic priorities commonly observed in industry use cases, ensuring the evaluation remains practically relevant. Additionally, using recent official documentation and verified benchmarks significantly enhanced the accuracy and reliability of our evaluation. We employed graphical tools, such as radar charts, to present the results interpretably, enabling intuitive comparisons of each platform's strengths and weaknesses. These methodological improvements allow for a more accurate, practical, and actionable analysis than previously reported in the literature.

7. Conclusion

Using the Weighted Scoring Model (WSM) to assess cloud storage services from Azure, AWS, and Google Cloud Platform helps identify the strengths and weaknesses of each provider. This analysis gives insight into how their storage solutions fit different business needs. Google Cloud Platform performed very well in block storage, with Google Persistent Disk standing out in speed, scalability, and cost-effectiveness. Google Cloud Archive Storage also led in integration and data management, making it a good option for complex workflows and long-term storage. However, some Google services, like Google Filestore, were less cost-efficient, which might be a drawback for businesses with tight budgets. Overall, Google is an excellent choice for companies that value high performance and smooth integration.

Amazon Web Services (AWS) performed strongly in key areas like speed, ease of use, scalability, and Security. It excelled in file storage (Amazon EFS), object storage (Amazon S3), and hybrid storage (AWS Storage Gateway), making it a dependable option for a variety of workloads, including those needing fast access. AWS also led edge and offline storage (AWS Snowball Edge Storage Optimized). However, AWS services can sometimes be expensive, especially for large or long-term use. Still, it remains a strong option for businesses focusing on performance and scalability.

Azure showed solid results across multiple storage types, including block storage (Azure Disk Storage), file storage (Azure Files), and object storage (Azure Blob Storage). Its main advantages are affordability and easy integration, making it a good fit for businesses already using other Azure services. However, Azure scored lower in hybrid and edge storage performance, which may not be ideal for demanding applications. Despite this, Azure is a budget-friendly and well-integrated choice for businesses within its ecosystem.

To sum up, Google Cloud Platform is best for performance and integration; AWS is the top choice for performance, scalability, and ease of use, while Azure is the most cost-effective and integrates well with its services. Businesses should choose a provider based on their priorities-speed, scalability, or cost savings-to find the best fit for their needs and budget.

References

[1] Amanpreet Kaur Sandhu, "Big Data with Cloud Computing: Discussions and Challenges," *Big Data Mining and Analytics*, vol. 5, no. 1, pp. 32-40, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[2] Praveen Borra, "Comprehensive Survey of Amazon Web Services (AWS): Techniques, Tools, and Best Practices for Cloud Solutions," *International Research Journal of Advanced Engineering and Science*, vol. 9, no. 3, pp. 24-29, 2024. [Google Scholar] [Publisher Link]

[3] Antara Debnath Antu et al., "Comparative Analysis of Cloud Storage Options for Diverse Application Requirements," *Cloud Computing-CLOUD 2021*, vol. 12989, pp. 75-96, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[4] Zouheir Daher, and Hassan Hajjdiab, "Cloud Storage Comparative Analysis Amazon Simple Storage vs. Microsoft Azure Blob Storage," *International Journal of Machine Learning and Computing*, vol. 8, no. 1, pp. 85-89, 2018. [CrossRef] [Google Scholar] [Publisher Link]

[5] Paola Pierleoni et al., "Amazon, Google and Microsoft Solutions for IoT: Architectures and a Performance Comparison," *IEEE Access*, vol. 8, pp. 5455-5470, 2020. [CrossRef] [Google Scholar] [Publisher Link]

[6] Omar Ahmad Alqahtani, and Mahmmed Mahmoud Alsandouny, "Comprehensive Comparison of Cloud Storage: GCP, AWS, AND AZURE," *International Journal of Engineering Research and Applications*, vol. 15, no. 3, pp. 27-35, 2025. [Publisher Link]

[7] Talluri Durvasulu Mohan Babu, "Navigating the World of Cloud Storage: AWS, Azure, and More," *International Journal of Multidisciplinary Research in Science, Engineering and Technology*, vol. 2, no. 8, pp. 1667-1673, 2019. [CrossRef] [Google Scholar] [Publisher Link]

[8] Goutham Sabbani, "Choosing the Right Cloud Storage: Object, Block, or File," *Journal of Scientific and Engineering Research*, vol. 9, no. 11, pp. 287-291, 2022. [CrossRef] [Publisher Link]
[9] Compare storage on Azure and AWS, Microsoft Learn, 2025. [Online]. Available: https://learn.microsoft.com/en-us/azure/architecture/aws-professional/storage
[10] About Blob (object) Storage - Azure Storage, Microsoft Learn, 2022. [Online]. Available: https://learn.microsoft.com/en-us/azure/storage/blobs/storage-blobs-overview
[11] What is Amazon S3?, Amazon Simple Storage Service, 2025. [Online]. Available: https://docs.aws.amazon.com/AmazonS3/latest/userguide/Welcome.html
[12] Cloud Storage, Google Cloud, 2025. [Online]. Available: https://cloud.google.com/storage/?hl=en
[13] Azure Virtual Machines, Overview of Azure Disk Storage, Microsoft Learn, 2025. [Online]. Available: https://learn.microsoft.com/en-us/azure/virtual-machines/managed-disks-overview
[14] What is Amazon Elastic Block Store?, Amazon EBS, 2025. [Online]. Available: https://docs.aws.amazon.com/ebs/latest/userguide/what-is-ebs.html
[15] Compute Engine Documentation, About Persistent Disk, Google Cloud, 2016. [Online]. Available: https://cloud.google.com/compute/docs/disks/persistent-disks
[16] Introduction to Azure Files, Microsoft Learn, 2025. [Online]. Available: https://learn.microsoft.com/en-us/azure/storage/files/storage-files-introduction
[17] Cloud File Storage, Amazon Elastic File System (EFS) Features, AWS Amazon, 2025. [Online]. Available: https://aws.amazon.com/efs/features/
[18] Filestore: High-performance, Fully Managed File Storage, Google Cloud, 2024. [Online]. Available: https://cloud.google.com/filestore?hl=en
[19] Azure Storage, Archive a Blob, Microsoft Learn, 2023. [Online]. Available: https://learn.microsoft.com/en-us/azure/storage/blobs/archive-blob?tabs=azure-portal
[20] Understanding S3 Glacier Storage Classes for Long-Term Data Storage, Amazon Simple Storage Service, 2025. [Online]. Available: https://docs.aws.amazon.com/AmazonS3/latest/userguide/glacier-storage-classes.html
[21] Storage Classes, Google Cloud, 2025. [Online]. Available: https://cloud.google.com/storage/docs/storage-classes
[22] Microsoft Azure Data Box Gateway Overview, Microsoft Learn, 2021. [Online]. Available: https://learn.microsoft.com/en-us/azure/databox-gateway/data-box-gateway-overview
[23] AWS Storage Gateway Features, Amazon Web Services, 2025. [Online]. Available: https://aws.amazon.com/storagegateway/features/
[24] Rayn Veerubhotla, and Manu Batra, Announcing the Anthos Ready Storage Qualification, Google Cloud Blog, 2020. [Online]. Available: https://cloud.google.com/blog/topics/hybrid-cloud/announcing-the-anthos-ready-storage-qualification
[25] Microsoft Azure, Azure Stack Edge, Microsoft Learn, 2025. [Online]. Available: https://azure.microsoft.com/en-us/products/azure-stack/edge/
[26] AWS Snowball Features, Amazon Web Services, 2024. [Online]. Available: https://aws.amazon.com/snowball/features/?nc1=h_ls
[27] Transfer Appliance, Overview, Google Cloud, 2025. [Online]. Available: https://cloud.google.com/transfer-appliance/docs/4.0/overview
[28] Azure Backup, What is Azure Backup?, Microsoft Learn, 2025. [Online]. Available: https://learn.microsoft.com/en-us/azure/backup/backup-overview
[29] Centralized Cloud Backup, AWS Backup Features, 2025. [Online]. Available: https://aws.amazon.com/backup/features/
[30] Backup and DR Service, Google Cloud, 2025. [Online]. Available: https://cloud.google.com/backup-disaster-recovery?utm_source=chatgpt.com&hl=en
[31] Andreas Wittig, and Michael Wittig, *Amazon Web Services in Action, An In-depth Guide to AWS*, 3rd ed., Manning, 2023. [Google Scholar] [Publisher Link]
[32] Azure Storage, Scalability and Performance Targets for Standard Storage Accounts, Microsoft Learn, 2024. [Online]. Available: https://learn.microsoft.com/en-us/azure/storage/common/scalability-targets-standard-account/
[33] Best Practices Design Patterns: Optimizing Amazon S3 Performance, Amazon Simple Storage Service, 2025. [Online]. Available: https://docs.aws.amazon.com/AmazonS3/latest/userguide/optimizing-performance.html
[34] Cloud Storage, Request Rate and Access Distribution Guidelines, Google Cloud, 2016. [Online]. Available: https://cloud.google.com/storage/docs/request-rate
[35] Licensing Documents, Microsoft Learn, 2025. [Online]. Available: https://www.microsoft.com/licensing/docs/view/Service-Level-Agreements-SLA-for-Online-Services
[36] Microsoft Azure, Azure Blob Storage, Microsoft Learn, 2025. [Online]. Available: https://azure.microsoft.com/en-us/products/storage/blobs/
[37] Product Overview of Cloud Storage, Google Cloud, 2025. [Online]. Available: https://cloud.google.com/storage/docs/introduction
[38] Amazon S3 Storage Classes, Amazon Web Service, 2025. [Online]. Available: https://aws.amazon.com/s3/storage-classes/

[39] Azure Storage, Introduction to Blob (Object) Storage, Microsoft Learn, 2023. [Online]. Available: https://learn.microsoft.com/en-us/azure/storage/blobs/storage-blobs-introduction
[40] Managed Disks, Azure Virtual Machines, Select a Disk Type for Azure IaaS VMs, Microsoft Learn, 2024. [Online]. Available: https://learn.microsoft.com/en-us/azure/virtual-machines/disks-types
[41] Amazon EBS Features, Amazon Web Services, 2025. [Online]. Available: https://aws.amazon.com/ebs/features/
[42] Amazon EBS Volume Types, Amazon Web Services, 2025. [Online]. Available: https://aws.amazon.com/ebs/volume-types/
[43] Compute Engine Documentation, Performance Limits for Block Storage, Google Cloud, 2024. [Online]. Available: https://cloud.google.com/compute/docs/disks
[44] Compute Engine Documentation, Choose a Disk Type, Google Cloud, 2016. [Online]. Available: https://cloud.google.com/compute/docs/disks/
[45] Pricing, Managed Disks, Microsoft Azure, 2025. [Online]. Available: https://azure.microsoft.com/en-us/pricing/details/managed-disks/
[46] Compute Engine Documentation, About Persistent Disk Performance, Google Cloud, 2016. [Online]. Available: https://cloud.google.com/compute/docs/disks/performance
[47] Compute Engine Documentation, Create a New Persistent Disk Volume, Google Cloud, 2016. [Online]. Available: https://cloud.google.com/compute/docs/disks/add-persistent-disk
[48] Compute Engine Documentation, Modify a Persistent Disk, Google Cloud, 2016. [Online]. Available: https://cloud.google.com/compute/docs/disks/modify-persistent-disk
[49] Amazon EFS Performance Specifications, Amazon Elastic File System, 2025. [Online]. Available: https://docs.aws.amazon.com/efs/latest/ug/performance.html
[50] Filestore, Service Tiers, Google Cloud, 2024. [Online]. Available: https://cloud.google.com/filestore/docs/service-tiers
[51] Barak Shein, and Harry Sadoyan, What's new with Filestore: Enhancing your stateful workloads on GKE, Google Cloud, 2023. [Online]. Available: https://cloud.google.com/blog/products/containers-kubernetes/filestore-features-for-stateful-workloads-on-gke
[52] Azure Files Pricing, Microsoft Azure, 2025. [Online]. Available: https://azure.microsoft.com/en-us/pricing/details/storage/files/
[53] Filestore Documentation, Google Cloud, 2025. [Online]. Available: https://cloud.google.com/filestore/docs/
[54] Filestore, Backups Overview, Google Cloud, 2025. [Online]. Available: https://cloud.google.com/filestore/docs/backups
[55] Azure Storage, Access Tiers for Blob Data, Microsoft Learn, 2025. [Online]. Available: https://learn.microsoft.com/en-us/azure/storage/blobs/access-tiers-overview
[56] Data Management, Azure Archive Storage, Microsoft Azure, 2024. [Online]. Available: https://azure.microsoft.com/en-us/products/storage/
[57] Secure Storage, Amazon S3 Glacier Storage Classes, Amazon, 2025. [Online]. Available: https://aws.amazon.com/s3/storage-classes/glacier/
[58] Azure Blob Storage Pricing, Microsoft Azure, 2025. [Online]. Available: https://azure.microsoft.com/en-us/pricing/details/storage/blobs/
[59] Natalia. S, Amazon Glacier Pricing Guide, MSP360 Blog, 2017. [Online]. Available: https://www.msp360.com/resources/blog/amazon-glacier-pricing-explained/
[60] Geoffrey Noer, Archive Storage Class for Coldest Data Now Available, Google Cloud Blog, 2020. [Online]. Available: https://cloud.google.com/blog/products/storage-data-transfer/archive-storage-class-for-coldest-data-now-available
[61] Cloud Storage, Object Lifecycle Management, Google Cloud, 2019. [Online]. Available: https://cloud.google.com/storage/docs/lifecycle
[62] Pravin Mishra, *Cloud Computing with AWS, Everything You Need to Know to be an AWS Cloud Practitioner*, 1st ed., Apress Berkeley, CA, 2023. [CrossRef] [Google Scholar] [Publisher Link]
[63] Performance and Optimization, AWS Storage Gateway, 2025. [Online]. Available: https://docs.aws.amazon.com/filegateway/latest/files3/Performance.html
[64] Anthos Powers Enterprise Container Platforms, Google Cloud, 2025. [Online]. Available: https://cloud.google.com/anthos?hl=en
[65] Azure Data Box Gateway Pricing, Microsoft Azure, 2025. [Online]. Available: https://azure.microsoft.com/en-us/pricing/details/databox/gateway/
[66] What is AWS Snowball Edge?, AWS Snowball Edge Developer Guide, 2025. [Online]. Available: https://docs.aws.amazon.com/snowball/latest/developer-guide/whatisedge.html
[67] Azure Stack Edge Pricing, Microsoft Azure, 2025. [Online]. Available: https://azure.microsoft.com/en-us/pricing/details/azure-stack/edge/
[68] Backup and Disaster Recovery, Google Cloud, 2025. [Online]. Available: https://cloud.google.com/solutions/backup-dr?hl=en
[69] Azure Backup Pricing, Microsoft Azure, 2019. [Online]. Available: https://azure.microsoft.com/en-us/pricing/details/backup/
[70] Azure Backup, Overview of Security Features, Microsoft Learn, 2025. [Online]. Available: https://learn.microsoft.com/en-us/azure/backup/security-overview

[71] Azure Backup, Azure Backup Service Documentation, Microsoft Learn, 2025. [Online]. Available: https://learn.microsoft.com/en-us/azure/backup/

[72] Sara Bouraya, and Abdessamad Belangour, “A WSM-based Comparative Study of Vision Tracking Methodologies,” *International Journal of Advanced Computer Science and Applications*, vol. 12, no. 8, pp. 87-98, 2021. [CrossRef] [Google Scholar] [Publisher Link]