White Paper

Leveraging the Breadth of Storage Services and the Ecosystem at AWS — Unlock the Full Potential of Public Cloud IaaS

Sponsored by: Amazon

Deepak Mohan Laura DuBois

September 2016

IDC OPINION

Growth of public cloud infrastructure as a service (IaaS) continues to dramatically outpace that of IT infrastructure spending. However, to leverage the full potential of public cloud IaaS, organizations must evaluate the breadth of storage services available with a given cloud provider. Selecting a provider with an ecosystem of storage and adjacent infrastructure services not only drives greater economic savings but also provides agility and time-to-market, interoperability, and scalability advantages.

The evaluation criteria used to assess public cloud storage service offerings should be broad and include a comprehensive suite of offerings to satisfy a range of use cases and service levels and a strong ecosystem of value-added functionality built around these offerings. These criteria are in addition to direct assessment factors such as a reliable service and competitive pricing. Service providers with broader service and feature capabilities will allow further optimizations in operational cost and agility and drive business benefits to the organization. Organizations that successfully select and adopt public cloud services that offer such a comprehensive suite of services will be able to realize broader benefits in multiple ways, including:

- **Increased agility and reduced risk** by satisfying a range of use cases, and production and test and development (test/dev) deployment scenarios, from a common infrastructure layer
- **Optimized storage costs** by making cost performance trade-offs for each storage need
- **Faster deployment and shorter learning curve** for new services due to standardized interfaces
- **Ease and low cost of data migration** between services due to integration between services
- **Better automation and monitoring** of data operations through tooling in the ecosystem

IDC believes that Amazon, through the multitude of storage options offered and the extensive Amazon Web Services (AWS) cloud ecosystem, is well positioned to enable customers to achieve both economic and operational benefits from public cloud services. Organizations evaluating public cloud infrastructure services must evaluate the services offered by AWS and run workload-specific proof of concepts to identify the optimum AWS service mix for their needs and derive the maximum value for their business.

September 2016, IDC #US41693416

IN THIS WHITE PAPER

In this white paper, IDC examines AWS storage services and the benefits that IT organizations can leverage by using the AWS storage portfolio. Two sample scenarios that use AWS storage services are then analyzed to understand the total cost of ownership (TCO) benefits of running on AWS.

SITUATION OVERVIEW

In recent years, there has been growing recognition among IT organizations of the cost and agility benefits of using public cloud IaaS. This has been accompanied by increasing confidence in the security and reliability of public cloud services. Public cloud IaaS is now an attractive and viable infrastructure option for enterprise IT.

Storage is an essential component of the IT architecture and requires careful planning and design. Organizations have a range of data sets that they need to manage, with varying performance, availability, and resiliency needs. The storage architecture must balance performance, availability, and budget criteria to meet the varying service-level agreements (SLAs) for each workload and data set. In a 2015 IDC study of 307 storage professionals in the United States, the following were highlighted as the top storage challenges:

- Managing data growth
- Meeting the SLAs for performance, availability, and recovery
- Troubleshooting storage-related problems
- Time or budget availability to build advanced storage features
- Time to manage storage migrations/technology refreshes
- The need to quickly fulfill storage provisioning requests

Storage infrastructure owners continue to be constrained by time and budget limits even as the capacity and flexibility required of the storage architecture increase. As business needs drive pressure on the storage architecture for more functionality and reduced budget, organizations are evaluating and adopting public cloud storage services to meet this duality of needs. One example is the management of email data at fast-growing organizations — performance and existing infrastructure investments keep the production email servers on-premises, but growing backup and archival needs constantly push the bounds of available infrastructure and real estate.

In response to the aforementioned challenges, the AWS cloud offers organizations a broad set of storage services supporting multiple data organizations and storage requirements. These services are well integrated with each other and with the AWS cloud ecosystem, allowing easy operations across services. Utilizing AWS storage services helps organizations address each of the following challenges:

- Cost-effective off-premises storage options available with Amazon S3 and Amazon Glacier provide an economically effective way to deal with storage volume and growth as well as datacenter power, cooling, and floor space requirements.
- The tiers of storage classes available, and the level of integration and interoperability between services, allow changes in the storage architecture and services with minimal disruption. This allows better responses to changing storage SLA needs for specific data sets and mitigates operational risk.

- AWS reduces the storage troubleshooting effort by abstracting physical hardware and storage issues away from the customer. AWS also provides a rich set of diagnostic and monitoring tools, including AWS Trusted Advisor, AWS CloudTrail, and VPC Flow Logs.
- The constantly growing value-added feature set available on AWS provides organizations access to advanced storage features on public cloud at no marginal investment to develop and use these capabilities. Features include automated life-cycle management, resource usage monitoring, and security and compliance capabilities at a platform level.
- AWS handles back-end technology and hardware refreshes for the underlying storage infrastructure, transparent to the end user on the storage service. This abstracts the end user away from the underlying technology and storage refreshes. AWS also provides services to allow migration and integration of existing applications through the Amazon Elastic File System (Amazon EFS) and AWS Storage Gateway offerings.
- Resource provisioning on AWS is on demand and nearly instantaneous, alleviating the traditional concerns with long, manual, and resource-intensive resource provisioning lead times. Programmatic access to storage services through standard APIs also allows for automation of storage routines.

THE COST, AGILITY, AND SCALE ADVANTAGE OF USING AWS STORAGE SERVICES

AWS storage services can be broadly categorized as object, file, and block based on the data organization on the service. These services are complemented with a set of data transfer capabilities to facilitate data migration and interoperability with resources outside of AWS. Figure 1 provides a brief overview of the AWS storage and data transfer services.

FIGURE 1

AWS Storage and Data Transfer Services

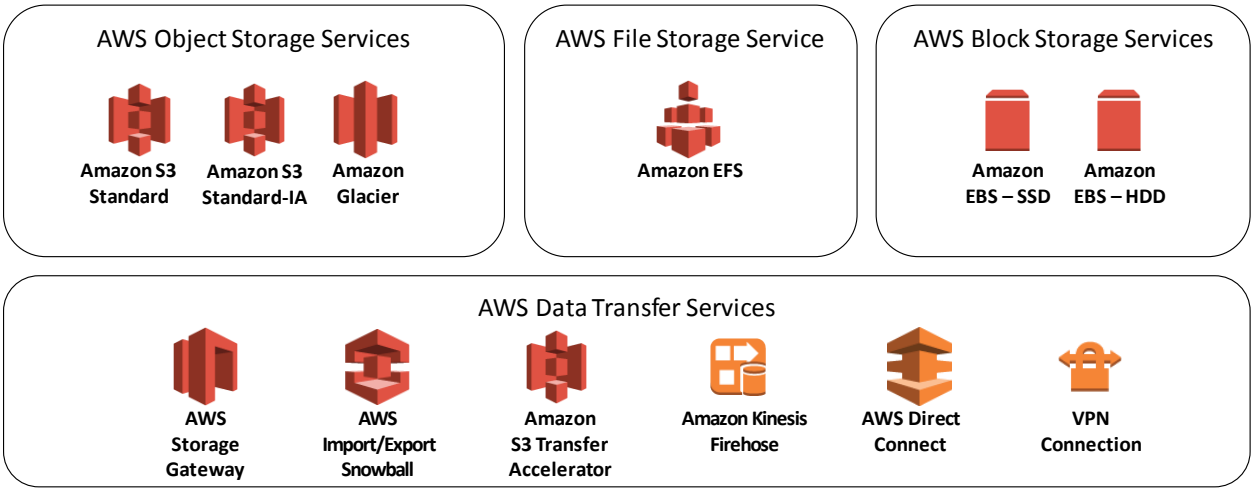

Source: IDC, 2016

Appendix A provides a brief overview of AWS storage services and the needs they address. Appendix B highlights some of the commonly used technical features available on AWS and the storage functionality provided by each feature.

Cost Benefits of Using the AWS Storage Portfolio

One of the biggest economic changes when moving from on-premises storage to the use of AWS storage services is the shift in spending from up-front investments to a variable, usage-based "pay as you go" model. The elimination of up-front capital costs reduces the barrier for testing and experimentation, making it viable to quickly run test workloads, iterate on the usage of different AWS services, and develop organizational capability around the new services through the process. The ability to change infrastructure configuration as needed, combined with AWS' track record of reliable capacity provisioning and service availability, provides the foundation for a true DevOps approach to enterprise IT infrastructure. AWS storage services allow for production workloads to be placed on hard disk drive (HDD)– and/or solid state drive (SSD)-backed block storage, as well as fixed content, archive, analytics, backup, and disaster recovery data on Amazon's S3/Glacier object storage.

For business continuity use cases such as backup and disaster recovery, AWS is able to offer organizations the high level of isolation and separation needed for a quality disaster recovery system, without the up-front investments and commitment in a remote physical datacenter. Organizations are able to meet sophisticated disaster recovery requirements by leveraging Amazon's multiregional footprint while enjoying the cost benefits possible through Amazon's scale of operation.

Example

The city of Asheville, North Carolina, used AWS alongside CloudVelox for its disaster recovery infrastructure, enabling a geographically separated backup storage location. This allowed Asheville to build a disaster recovery system that ensured data recovery and business continuity through natural disasters, despite not having the operating scale and budget for the capital investment needed for such a system. (Source: Amazon, 2016)

Interoperable and Integrated Portfolio of Storage Offerings

AWS offers a broad portfolio of storage services to meet a variety of needs. Within the highly scalable object storage service, three tiers of services are offered. While the Amazon S3 Standard - Infrequent Access (Amazon S3 Standard-IA) offering provides highly durable storage with reduced availability at approximately 40% of the standard storage cost, Amazon Glacier offers highly durable archival storage capability at approximately 23% of the standard storage cost. IDC's 2016 *Amazon Web Services (AWS) Infrastructure-as-a-Service (IaaS) User Survey* of 401 IT organizations in North America shows that backup, disaster recovery, and archiving are one of the top 3 largest data sets driving growth of AWS storage usage, and the availability of tiers of storage to address varying backup and archiving increases the storage optimization options for customers.

Within block storage, Amazon EBS also allows the end customer to choose from four distinct configurations – two SSD hardware-based options for low-latency high I/O use cases and two HDD-based options for cost-optimized low I/O high-throughput use cases. The availability of multiple storage classes and options enables organizations to select the right storage need and pay for only the functionality and storage performance required.

The newly launched Amazon EFS offers Amazon EC2 instances a standard file system interface and file system access semantics with on-demand capacity and a fully managed infrastructure. Workloads such as Big Data and analytics, media processing, content management and web serving, and home directories can be moved into the AWS cloud to eliminate the traditional pains associated with file server management and provisioning.

These services are integrated through the underlying AWS network within a region. They are also able to use a consistent framework for management such as resource access control (Amazon Identity and Access Management [IAM]), monitoring (Amazon CloudWatch), and API call audits (Amazon CloudTrail). The availability of multiple classes of storage within the AWS framework allows organizations to easily select the desired combination of storage services based on access protocol and performance-level needs, make changes to these services as needed without massive data migration overhead, and manage all the services under the AWS management framework, enabling changes and data transfer across services without the risks of managing and moving data across multiple service providers. The availability of life-cycle management policies to manage automated data transfer of objects from Amazon S3 to other tiers of storage further facilitates the integrated use of these services and allows organizations to automate and ensure optimal utilization of storage resources.

Example

SoundCloud offers a platform used by artists to collaborate on media content and manages petabytes of data on AWS. Amazon S3 is used to host and distribute the media files to end users, and all files are copied to Amazon Glacier for cost-effective cold storage. Using Amazon S3 and Amazon Glacier, SoundCloud is able to meet its production and long-term content backup requirements in a cost-effective manner. (Source: Amazon, 2016)

Agility and Time to Market

Moving from on-premises storage to public cloud infrastructure increases the speed with which new infrastructure resources can be provisioned and brought into active utilization. With traditional storage infrastructure, resource provisioning is a time-consuming, highly manual process that can take from days to weeks to complete. Large-scale increases in needs, such as those involving datacenter expansions, can even take months to complete. To accommodate this lead time, projects often plan in long buffer times for IT provisioning, adding to the time taken to deliver a project to production. With public cloud services, provisioning of resources is a near-instantaneous process. This eliminates the provisioning lead time as a bottleneck to new project delivery, and allows organizations to define, deploy, and make new projects available in production as soon as applications are ready. This is particularly useful for lean organizations in fast-paced industries where projects need to be rapidly defined and introduced to production to stay competitive.

Example

Pinterest, the web-based visual bookmarking application, has used AWS storage solutions since its launch in 2010. Using AWS has enabled Pinterest to maintain the velocity and speed of experimentation needed to constantly bring new product innovations to the market. AWS has also enabled Pinterest to meet the rapid growth in its user base as it went from a start-up to one of the most popular visual bookmarking tools on the web. (Source: Amazon, 2016)

Scalability and Efficient Resource Expansion

One of the by-products of the slow provisioning speed with traditional infrastructure is the need to plan capacity and infrastructure growth well in advance of real business needs. Business growth forecasts are used to plan infrastructure expansion projects. The long lead time, combined with early decisions on infrastructure needs, often results in underutilization of provisioned resources because of overly conservative overprovisioning. Even where planning is accurate, organizations are forced to have standby infrastructure capacity for the peak usage scenario and maintain this through the lower-usage periods. The scale-as-you-need model with public cloud allows organizations to provision

infrastructure based on real business needs, eliminate resources that are no longer needed, and rapidly scale up again when needed. This is a huge benefit for start-ups that need to be ready for rapid adoption but wish to avoid the capital investment risk and commitments associated with infrastructure.

Example

Airbnb is an online marketplace for users to list, discover, and book accommodations around the world. Airbnb has been able to take advantage of the ease and speed with which AWS allows customers to scale their infrastructure usage to grow their platform to meet their growing business needs. Airbnb uses Amazon EC2 for its applications and Amazon S3 to host backups and static files, including 10TB of user pictures. The company also uses other AWS services, including Amazon Elastic Load Balancing, Amazon Elastic MapReduce, and Amazon Relational Database Service, to easily manage its scaling and data processing needs. (Source: Amazon, 2016)

Advanced Storage Features

Besides the variety of core storage functionalities, AWS provides a rich set of advanced storage functionality. The following list highlights some of the commonly used AWS advanced storage features:

- **Versioning.** Amazon S3 provides protection of objects stored in a bucket with versioning. Versioning can be used to retrieve and restore older versions of an object in the event of corruption or deletion. Versioning provides a convenient way to protect against accidental deletions and human errors during operations.
- **Data life-cycle management.** Customers can set up rules to automatically migrate data across different tiers of storage — Amazon S3 Standard, Amazon S3 Standard-IA, or Amazon Glacier — based on the age of the data. Automated life-cycle management avoids the need for customer-initiated scrubbing and data migration operations across tiers of storage. This also enables efficient usage of storage capacity in higher-cost storage tiers, further improving the TCO.
- **Vault lock.** Amazon Glacier Vault Lock allows customers to set up specific access control rules for individual vaults in Glacier and "lock" the vault with the specified policies. The vault lock policy cannot be changed once locked. This is designed for strong enforcement of compliance and regulatory controls on archival data.
- **EBS snapshot.** Customers can back up data on EBS volumes using EBS snapshots, which are point-in-time snapshots of the volume. After an initial snapshot, subsequent snapshots are created as incremental backups. This provides a usage-optimized and cost-efficient way of maintaining a backup copy of EBS volumes.

A detailed listing of storage features is provided in Appendix C.

Datacenter Interoperability

For organizations with existing infrastructure facilities, interoperability between resources on AWS and resources in the existing facility is often important. Even for start-ups without existing infrastructure, the ability to move data in and out of the AWS cloud can be critical — to serve specific needs for their customers, for edge use cases involving collection and streaming of data from the field into AWS, or for other specialized needs involving movement of data in and out of AWS in a predictable manner. Amazon provides a number of services to enable data transfer to and from the AWS cloud, including:

- **AWS Direct Connect.** AWS was the pioneer of the Direct Connect concept, which offers customers a dedicated physical connection between their datacenter and the AWS cloud. This provides predictable throughput and costs for data transfer to and from AWS. AWS Direct Connect is suited for use cases that need regular high-throughput data transfer between the datacenter and the AWS cloud. This may be for connectivity between on-premises and cloud components during migration of applications to the public cloud, shared data resources across on-premises and cloud applications, or a reliable connection for applications using on-premises and cloud-based infrastructure such as backup or replication.
- **VPN Connection.** VPN connection provides a secure and isolated link between a customer datacenter and an Amazon VPC. The data transfer still traverses the same path as data transfer over the internet, but it is encrypted during transit through the internet. VPN connection is suited for use cases that need a secure link between on-premises and cloud resources and where the throughput requirements and data transfer volumes are not high.
- **AWS Import/Export Snowball.** Snowball is a managed data transfer service for the transfer of petabyte-scale data to and from the AWS cloud. Snowball is designed for bulk movement of large data volumes, where the available connection speed and data volume do not allow data transfer within the required time constraint. This service allows for large volumes of data to be ingested into AWS to seed the environment.
- **Amazon S3 Transfer Acceleration.** Data transfer speeds into Amazon S3 can vary depending on the customer's location. Amazon S3 Transfer Acceleration enables customers to speed up upload of data into Amazon S3 by leveraging the AWS edge infrastructure. Amazon S3 Transfer Acceleration is ideal for periodic data uploads into the AWS cloud.
- **AWS Storage Gateway.** The AWS Storage Gateway is an appliance that resides in the customer's datacenter and connects the datacenter to AWS cloud storage. It is well suited for use cases where an application running on-premises utilizes storage in the cloud. It is often used together with AWS Direct Connect for predictable performance. The AWS Storage Gateway eliminates the need for customers to invest in on-premises cloud gateways or appliances.
- **Amazon Kinesis Firehose.** The Amazon Kinesis Firehose enables loading of streaming data into AWS. It can capture and automatically load streaming data into Amazon S3 and Amazon Redshift and is well suited for ingestion of near-real-time streams of data into AWS.

Broader AWS Service Ecosystem Benefits

In addition to the cost, agility, and scale benefits enabled by the AWS storage portfolio, customers also gain access to the broad ecosystem of services, tools, and capabilities available in the AWS cloud. These include value-added services and functionality and a validated and secure infrastructure platform for customer applications:

- **The AWS ecosystem.** AWS offers a rich ecosystem of services and tools to meet the broader organizational needs around its IT resources — including needs such as resource management, access control, and development. IDC's 2016 *AWS IaaS User Survey* highlighted that "gaining the benefits of the AWS ecosystem" was the second-biggest factor driving adoption of Amazon EC2 by respondents. Commonly used services in the AWS ecosystem include a consistent IAM framework across resources; monitoring and audit features such as AWS Trusted Advisor, AWS CloudTrail, and Amazon CloudWatch; and deployment tools such as AWS CloudFormation and AWS Elastic Beanstalk.

  The availability of this ecosystem of services enables better operations and optimization of a customer's AWS resources. At the 2016 AWS Summit in Chicago, GE noted that with application migration to AWS, it observed a TCO reduction of 52%. The company also highlighted that a large enabler of the cost savings was the use of the AWS ecosystem of services on top of the core services — services such as AWS Trusted Advisor and AWS Elastic Beanstalk — enabling intelligent automation of Amazon resource usage. Utilizing the extended set of AWS services and tooling allows organizations to offload the monitoring and tracking functionality to AWS and further optimize their utilization of infrastructure resources.

- **Foundation of compliance and security.** Compliance and security are critical components for enterprise IT applications. The increase in volume and variety of cyberthreats has increased the level of attention on the need for strong controls. Certain industries require specific compliance and certifications, such as HIPAA for healthcare-related applications and PCI DSS for financial transaction processing. Even for applications that are not governed by mandatory compliance standards, business needs and internal policies require a strong focus on security.

  Designing and certifying a compliant application stack from the infrastructure up can be a tedious process. Activities such as penetration testing and arranging independent certification reviews can take time away from the core focus of the application. Amazon's cloud infrastructure is designed and certified for a number of commonly used compliance and audit standards, making it easier and faster for end customers to build, certify, and run their own compliant applications on AWS. Among the list of certifications and audits supported are PCI DSS Level 1, SOC reports, and ISO 9001. The complete list of supported compliance, certification, and regulatory guidelines is available at the Amazon Web Services website.

  Amazon takes a shared approach to ensuring security for applications running on AWS. Amazon handles "security of the cloud," and the customers are responsible for "security in the cloud." The AWS cloud is designed with security as its highest priority and is designed to meet the highest security needs. It is accredited with commonly required security certifications such as ISO 27001 and DoD SRG. The customers need to plan for and build in security at the application level. AWS offers two levels of guidance to customers that need assistance in security for their applications. The base level is through AWS support, which offers customers tools and resources to identify gaps and meet their security needs. More advanced support is available through the AWS Professional Services group, and through the AWS Partner Network, for customers with more complex and specialized security needs.

Brief Overview of Data Set Classification

AWS storage services are designed to address a variety of storage needs. To effectively utilize these services, organizations need to understand and be able to identify their storage requirements. Data managed by an organization can be broadly classified as follows:

- **Production data.** Production data refers to data sets that are accessed as part of regular operations for a business. This includes primary data sets, as well as mirror copies of current data sets maintained to increase availability. AWS storage services used for production data are Amazon S3 (for unstructured large volume data, such as the media content for Spotify), Amazon EFS (for unstructured data that needs to be shared across compute resources through a file system), and Amazon EBS (for transactional data accessed through an application running on Amazon EC2, such as the data set for a database on Amazon EC2).
- **Copy data (replication and backup).** Replication and backup copy data refers to copies of production data maintained separately to meet a target recovery time objective (RTO) or recovery point objective (RPO) in the event of a data corruption or failure. AWS storage services used for copy data are Amazon EBS, Amazon S3 Standard, Amazon S3 Standard – Infrequent Access, and Amazon Glacier. The selection of service will depend on the nature of the data and the RTO/RPO metrics acceptable for the use case.
- **Copy data (archival).** Archival data refers to copies of production data maintained for long durations in "cold storage" for infrequent data access needs — such as audit and low-speed disaster recovery. These copies are intended to be stored in a cost-efficient manner, with high durability, at the expense of slow recovery time. The AWS service used for archival data is Amazon Glacier.
- **Copy data (analytics and test/dev)**. Analytics data refers to data sets derived from an analysis of production or copy data, which is used to monitor the operational characteristics of the system. Analytics data is typically stored in Amazon S3 for offline storage and pulled into Amazon EBS for active review and evaluation. Test/dev data refers to copies of data used for testing, staging, and development purposes. The storage service used will depend on the scenario.

Appendix D provides a detailed view of how data sets in an organization can be classified based on their storage needs.

Understanding Total Cost of Ownership

Total cost of ownership is a commonly used measure of the cumulative direct and indirect costs of running an IT system over a period of time. TCO includes up-front capital costs, maintenance costs, licensing costs, additional usage charges, administrative and manpower costs, datacenter costs (power and real estate), any associated data migration charges, and other storage-related charges (such as networking). TCO is used as a quantitative financial measure of the cost of operating an IT system and can be used to compare different IT options on a linear cost scale. The components of TCO may change depending on the system and the option that is being evaluated. Comparing options with different payment schedules (such as public cloud versus on-premises infrastructure hardware) involves the use of a time period that matches the hardware refresh cycle.

Refer to Tables 1 and 2 for a comparison of the TCO of AWS services versus on-premises infrastructure for the two scenarios discussed in the section that follows. Data is derived from the AWS TCO calculator.

Sample Scenarios to Illustrate Usage of Multiple AWS Storage Products

Two sample business use cases are examined in this section. The first is a database for a consumer web application running on AWS. The second is the backup and archive for an on-premises application.

Sample Scenario #1: A Production MongoDB Database Running on AWS

In the first sample scenario, a customer runs a production database and application on AWS. The production deployment runs on us-east-1 to stay close to primary target customers. The testing, development, and staging operations are carried out in us-west-2 to take advantage of lower costs. Figure 2 illustrates the customer deployment on AWS. Specifically:

- **Production data:** To meet the availability needs of the service, the database is deployed in a replica set configuration. All data on the primary database is replicated to the secondary database, which is maintained in a different availability zone. Both databases use SSD-backed EBS volumes with provisioned IOPS to achieve a consistent and acceptable latency for database operations.
- **Copy data:** Hourly EBS snapshots are created and saved to Amazon S3 to enable a recovery point that is an hour out of date at most. Daily backups are created and stored to Amazon S3 using the built-in functions in MongoDB. Amazon S3 life-cycle policies are used to move daily backup objects older than a week into Amazon Glacier archives. Archives older than three months are deleted. Cross-region replication (CRR) is used to maintain a copy of the database in staging that is approximately a day behind the production version. This is used for daily re-creation of the staging database.

FIGURE 2

Production MongoDB Database on AWS

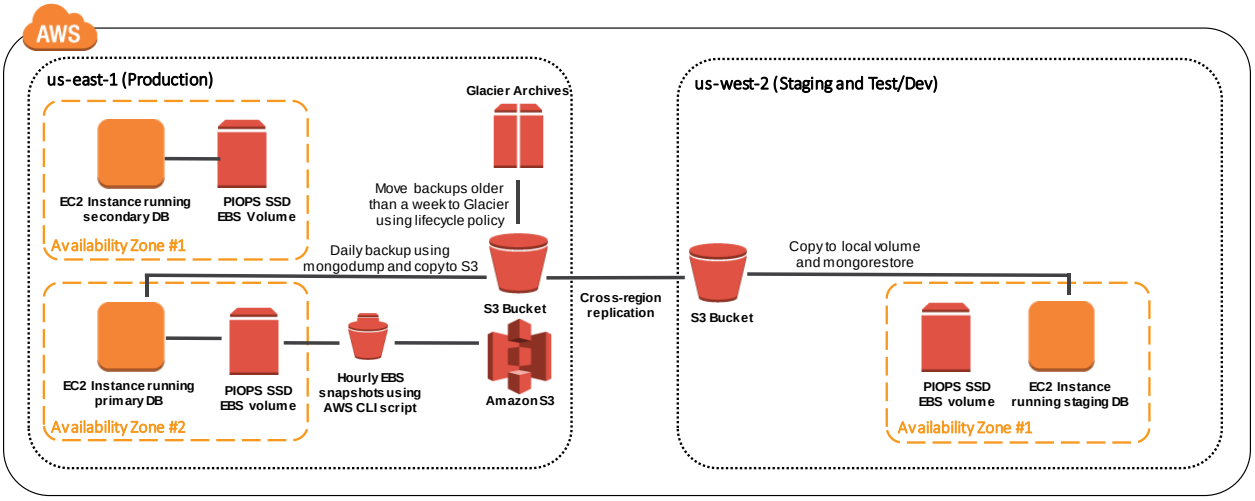

Source: IDC, 2016

For the TCO comparison, a database server with a storage size of 1TB is used (resulting in two servers, with 1TB each, due to the replicated set for high availability). This results in archival storage usage of approximately 50TB. Table 1 shows that the usage of AWS results in an annual TCO savings of 75% and a three-year savings of $270,842.

TABLE 1

Three-Year TCO Comparison Between On-Premises and AWS: Production Database on AWS

	On-Premises	AWS
Servers	$79,245	$59,787
Storage	$64,592	$3,871
Cold storage	$152,269	$13,386
Network hardware	$6,780	
IT labor	$60,000	$15,000
Three-year savings		$270,842
Annual TCO savings		75%

Note: The AWS usage costs include AWS business-level support.

Source: Amazon, 2016

Sample Scenario 2: Backup and Archive for On-Premises Microsoft Exchange

In the second sample scenario, a customer maintains an internal datacenter for its internal and external applications, including its Microsoft Exchange Server for employee email. To shift internal resources toward the growth needs of external-facing applications, the customer has moved the backup and archival storage out from its datacenter into AWS (see Figure 3).

- **Production data:** Production data resides on the Microsoft Exchange Server deployed in the customer datacenter.
- **Copy data:** The Microsoft Exchange Server is protected by a backup application that uses the AWS Storage Gateway to interact with Amazon S3 Standard and Amazon Glacier. The organization uses an AWS Direct Connect subscription to ensure a predictable throughput for the high-volume backup and restore transactions. Two different storage gateways in two different configurations are used. The first, a Storage Gateway with cached volumes, is used as the target for daily backup to local disk. This local gateway volume is copied over to Amazon S3 and can be used with a cloud instance of Microsoft Exchange. The other, a gateway in VTL mode, is a seamless tape replacement. The monthly backups are performed using existing backup software pointed at the VTL. These virtual tapes are archived to Amazon Glacier for long-term cold storage and retained for 10 years. The AWS Storage Gateway enables the customer to easily switch the target storage destination for the backup application from on-premises storage appliances and tapes to AWS storage services.

FIGURE 3

On-Premises Microsoft Exchange Server Backup to AWS

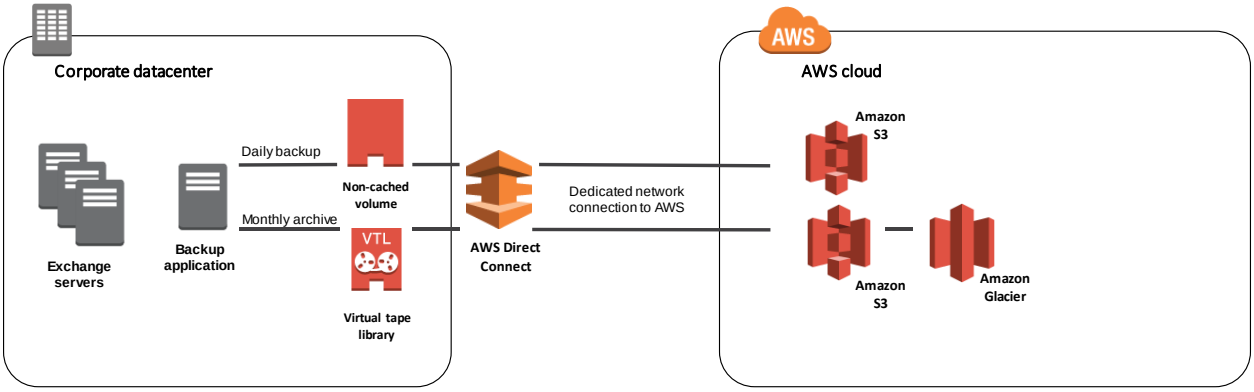

Source: IDC, 2016

For the TCO comparison, a backup storage size of 2TB and an archival storage size of 240TB are assumed for the system. For on-premises archival, tape-based long-term storage is assumed. Table 2 shows that AWS results in a three-year savings of $202,133 and an annual TCO savings of 78% on backup and archival costs on-premises.

TABLE 2

Three-Year TCO Comparison Between On-Premises and AWS: Backup and Archive Exchange Server

	On-Premises	AWS
Hot storage infrastructure	$70,627	$8,776
Cold storage infrastructure	$154,749	$33,384
Network (AWS Direct Connect)	$2,282	
IT labor	$17,365	$4,950
AWS business support	$4,220	
Three-year savings		$202,133
Annual TCO savings		78%

Source: Amazon, 2016

CHALLENGES/OPPORTUNITIES

Adoption of public cloud services provides organizations with several opportunities to improve on the cost and agility of operations. But the path to adopt public cloud resources for IT organizations needs to be well planned, especially for organizations with existing on-premises infrastructure operations. Existing investments in IT infrastructure, recent refreshes of hardware, and existing investments in training and certifications need to be kept in mind when planning the adoption path for public cloud services. Recommended practices to ensure a smooth adoption curve for public cloud include:

- Avoid a siloed approach to integration of public cloud resources across the organization. This can result in inconsistencies and missed opportunities, which are easy to address early on but hard to realign once usage patterns are established.
- Explore and utilize the datacenter interoperability features available in the ecosystem. Features like AWS Direct Connect, Amazon S3 Transfer Acceleration, and AWS Storage Gateway can offer an easy path to begin integration of cloud storage resources with existing infrastructure (for a hybrid storage model).
- Plan ahead and align migration plans with major hardware refreshes, especially tape and disk investments. This will simplify financial trade-offs and minimize the impact of sunk costs in IT infrastructure on the decision.
- Develop, communicate, and initiate the internal training and skill set upgrade plan early in the process. Do not underinvest in managing the organizational changes required.

Bringing about the organizational changes that are needed to take advantage of public cloud resources — such as skill set training, process changes, and IT governance enhancements — is crucial to effectively leverage public cloud benefits. These changes require sustained effort through the adoption process and through the early phase of public cloud inclusion to reinforce and retain the changes and practices. Mature public cloud service providers such as Amazon are in a position to help organizations manage this transition, through continued investments in partnerships, features, and training, to enable ease of adoption and integration of public cloud with traditional infrastructure.

CONCLUSION

Public cloud IaaS services offer organizations a range of economic and operational benefits. Organizations must actively evaluate and consider the breadth of features and the services available in a service provider's cloud ecosystem when evaluating and adopting a public cloud service provider. Selection of a public cloud service provider with a broad set of services, and a rich ecosystem, will allow organizations to maximize the benefits of using public cloud for their infrastructure needs.

One of the direct implications of the low up-front costs and the speed of service provisioning is that infrastructure design on public cloud can take an iterative approach to evaluation and adoption. This provides a fast and efficient way for organizations to educate themselves about the latest public cloud storage services, validate and deploy public cloud-based resources, and reduce the time to utilization of the latest capabilities released on public cloud.

The breadth and depth of AWS infrastructure services remain unparalleled in the industry and continue to drive organizations to short-list AWS public cloud services. As new releases are announced, and the infrastructure options available evolve, organizations that have moved to an iterative approach to infrastructure design will be among the first to evaluate and derive benefits from the latest services introduced in the market by AWS.

Appendix A: Overview of AWS Storage Services

TABLE 3

Overview of AWS Storage Services

Product Name	Features	Target Use Cases
Amazon Simple Storage Service (Amazon S3)	Secure, durable, and scalable cloud object storage, accessible from EC2 instances and anywhere on the web	General-purpose cloud storage, from inside or outside AWS cloud; data is stored in objects, grouped within Amazon S3 buckets
Amazon S3 Standard	99.99% availability and 99.999999999% (11-9s) of durability	Highly durable and available data storage, such as content storage and backup of production data
Amazon S3 Standard-Infrequent Access (Amazon S3 Standard-IA)	99.9% availability and 99.999999999% (11-9s) of durability	Data with less critical availability needs, such as content for internal training
Amazon Glacier	Secure, durable, and low-cost storage, optimized for infrequent access and slow retrieval (multiple hours) and accessible from anywhere on the web	Archival and long-term cold backup, such as historical database snapshots; data stored as archives and grouped within Amazon Glacier vaults
Amazon Elastic File System (Amazon EFS)	Scalable and fully managed file storage accessible from EC2 instances in the AWS cloud	Data accessed through a file system, such as the data for a legacy application that used a network-attached storage (NAS) system
Amazon Elastic Block Storage (Amazon EBS)	Persistent block storage volumes for use with EC2 instances in the AWS cloud	Latency-sensitive data that needs to be stored on an attached volume
Solid state drive (SSD)–backed Amazon EBS volumes	SSD-based volumes, with low latency and high IOPS	Highly latency-sensitive data, such as production database data
Hard disk drive (HDD)–backed Amazon EBS volumes	HDD-based volumes for high throughput and lower cost	Batch processing job data or cost-optimized data storage

Source: Amazon, 2016

Appendix B: Commonly Used Infrastructure Techniques

TABLE 4

Common Infrastructure Features and the Storage Needs They Address

Technical Features	Storage Needs Addressed
Amazon S3 — redundant copies across facilities	Durability and availability, resilience to facility failure
Amazon S3 — accessible from the web	Shared storage that can be accessed from Amazon EC2, from other AWS resources, and from outside the AWS cloud over the internet
Amazon S3 — multipart uploads	Ability to send a large file over an intermittent network connection without restarts
Amazon EBS — replication within an availability zone (AZ)[1]	Durability and availability, resilience to component failure
Amazon EBS — solid state drive (SSD) storage	Flash media–based storage for lower latency for transactional read/write operations and higher random I/O operations per second (IOPS)
Amazon EBS — provisioned IOPS with SSD storage	Consistent I/O performance for latency-sensitive and I/O-intensive applications
Amazon EBS — hard disk drive (HDD) storage	Magnetic disk media–based storage for cost-optimized local storage use cases; also enables high throughput for low IOPS sequential read operations
Amazon EFS — replication across availability zones (AZs)[1]	Durability and availability, resilience to AZ failure; also enables lower latency and higher throughput when accessing a data resource from multiple AZs
Amazon EFS — performance	Throughput and IOPS that grow with capacity while latencies stay consistent
Amazon EFS — scale	Built to handle file systems in petabytes
Amazon EFS — availability and reliability	File systems span regions, ensuring access and delivering a durable, global, and common repository for EC2 instances
Amazon EFS — elasticity	Automatic capacity provisioning so that customers pay for only what they use

[1]Please refer to the Amazon Web Services documentation for a detailed explanation of availability zones and regions in AWS.

Source: Amazon, 2016

Appendix C: List of Advanced Storage Services Available on AWS Storage

Amazon S3 features include the following:

- **Versioning.** Amazon S3 provides protection of objects stored in a bucket with versioning. Versioning can be used to retrieve and restore older versions of an object in the event of accidental corruption or deletion.
- **Multi-Factor Authentication (MFA) Delete.** To prevent accidental or malicious deletion of important data, customers can choose to add an additional layer of security to a bucket by enabling MFA Delete for a bucket. This requires all users to provide additional authentication (through an approved MFA device) before permanent deletion of an object version in the bucket.
- **Cross-region replication.** Cross-region replication allows customers to set up automatic replication of objects across buckets in different AWS regions. Once cross-region replication is enabled, changes made to objects in a source bucket can be automatically and asynchronously replicated at the destination bucket.
- **Data life-cycle management.** Customers can set up rules to automatically migrate data across different tiers of storage – Amazon S3 Standard, Amazon S3 Standard-IA, or Amazon Glacier – based on the age of the data. Automated life-cycle management avoids the need for customer-initiated scrubbing and data migration operations across tiers of storage. This also enables efficient usage of storage capacity in higher-cost storage tiers, further improving the TCO.
- **Virtual private cloud (VPC) endpoints.** VPC endpoints allow customers to create a private connection between a VPC and an Amazon S3 bucket, without requiring the VPC to have access to the public internet.
- **Encryption.** Customers can securely upload or download data from Amazon S3 using SSL-protected HTTPS data transfer. Amazon S3 can also automatically encrypt data at rest with server-side encryption. The keys used for encryption can be managed by Amazon S3, by customer encryption keys provided by the customers, or through the Amazon Key Management Service (KMS).
- **Access control.** Access to resources in Amazon S3 can be managed using the AWS IAM policies. IAM provides a consistent framework to manage access control across AWS resources for users and roles. Customers can also set access control at a resource level by using bucket policies or by using access control lists (ACLs) at a bucket or an object level. These mechanisms allow granular access control to S3 data, allowing a bucket owner to set up access to S3 data for users across multiple AWS accounts.
- **Query string authentication.** Query string authentication allows customers to share Amazon S3 objects using URLs that are valid only for a specific period of time, after which they expire.

Amazon Glacier-specific features include the following:

- **Vault inventory.** Vault inventory is an inventory of all archives in each of a user's vaults, updated approximately once a day.
- **Vault lock.** Amazon Glacier Vault Lock allows customers to set up specific access control rules for individual vaults in Glacier and "lock" the vault with the specified policies. The vault lock policy cannot be changed once locked. This is designed for strong enforcement of compliance and regulatory controls on archival data.
- **Encryption.** All data stored on Amazon Glacier is automatically encrypted on the server side using AES 256-bit block cipher. The key is managed by Glacier.

Amazon EBS-specific features include the following:

- **EBS snapshot.** Customers can back up data on EBS volumes using EBS snapshots, which are point-in-time snapshots of the volume. After an initial snapshot, subsequent snapshots are created as incremental backups.

Appendix D: Classification of Data Sets in an Organization

Table 5 classifies data sets in an IT organization into two broad categories — production data and copy data. Production data is primary or mirror data that is actively used by IT applications as part of the regular business operations of the organization. Copy data includes copies of production data that is used for various nonregular operations, such as replication, backup, archiving, and testing. This classification helps determine the storage needs for each data set. Factors such as business context will also impact storage needs. It is important to note that while the approach outlines a classification of data sets, the distinction between specific categories may be blurred in certain use cases. For example, several database deployments use physically separated mirror databases to serve as both the mirror data set and the replication data set.

Table 6 uses the classification described in Table 5 to map AWS services to sample data sets.

TABLE 5

Data Set Classification

Data Set Category	Description	Examples	Typical Needs
Production data	Data sets that are actively accessed or updated as part of regular operations for an application – such as user data, user content, and business model data		
Primary data	Data that is actively used as part of regular applications	User document listing for an online document storage service	▪ High availability ▪ Fault tolerance
Mirror data	Copies of primary data that is actively used for regular application workflows	Backup database to provide high availability for user authentication for an enterprise application	▪ High availability ▪ Fault tolerance
Copy data	Copies of production data sets that are stored and accessed for nonregular operations of an application – such as test/dev, staging, failure recovery, and audit		
Replication	Copy of current production data sets maintained separately for recovery – designed to meet a recovery time objective (RTO)	Copy of user database for an online document storage service, maintained to provide rapid recovery in the event of a failure	▪ Availability and throughput to meet the RTO
Backup	Time-stamped copy of production data, maintained separately for recovery to meet a recovery point objective (RPO)	Daily backup of user access logs for an enterprise sales support application	▪ Cost effective ▪ High durability ▪ High scalability ▪ Accessible on demand
Archival	Historical copies of production data, maintained offline for nonregular retrieval	Archived historical weekly backups of access logs for an application	▪ Cost effective ▪ High durability
Analytics	Data derived from production and copy data, used to study the operation of the system	Historical usage and retrieval pattern for backup data	▪ Cost effective ▪ High scalability
Test and development	Copy of production data or dummy data, used for development and testing	Mock database for staging and testing use case	▪ Based on the environment requirements

Source: IDC, 2016

TABLE 6

Mapping AWS Storage Services to Specific Example Data Sets

Example Data Set	Classification	Needs	Amazon Storage Product
Production database data	Production data — primary	▪ High availability ▪ Fault tolerance	Amazon EBS – SSD
Media content storage for a consumer media sharing service	Production data — primary	▪ High availability ▪ Fault tolerance ▪ Cost effective	Amazon S3
Large data set for Hadoop nodes that are part of a cluster	Production data — primary	▪ Fault tolerance ▪ High throughput ▪ Cost effective	Amazon EBS – HDD
Slave database data for a high-availability failover configuration	Copy data — replication	▪ High availability ▪ Fault tolerance	Amazon EBS – SSD
Backup copy of media content for a consumer media sharing service	Copy data — backup	▪ Fault tolerance ▪ Cost effective ▪ Durability	Amazon S3 Standard-IA
Daily backup of database, with a strict recovery time objective	Copy data — backup	▪ High availability ▪ Fault tolerance ▪ Cost effective ▪ Durability	Amazon S3

Source: IDC, 2016

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-community.com
www.idc.com