

AWS Announces Four New Storage Services and Capabilities

New Amazon S3 Glacier storage class is designed to offer the lowest cost storage for milliseconds retrieval of archived data—also available as a new access tier in Amazon S3 Intelligent-Tiering

New Amazon FSx for OpenZFS service makes it easy to move data stored in on-premises commodity file servers to AWS

New Amazon EBS Snapshots Archive storage tier reduces the cost of storing archival snapshots by up to 75%

AWS Backup brings centralized data protection and automated compliance auditing to Amazon S3 and VMware workloads

LAS VEGAS—Nov. 30, 2021—Today, at AWS re:Invent, Amazon Web Services, Inc. (AWS), an Amazon.com, Inc. company (NASDAQ: AMZN), announced four new storage services and capabilities that deliver more choice, reduce costs, and help customers better protect their data. Amazon Simple Storage Service (Amazon S3) Glacier Instant Retrieval is a storage class that provides retrieval access in milliseconds for archive data—now available as a new access tier in Amazon S3 Intelligent-Tiering. Amazon FSx for OpenZFS is a managed file storage service that makes it easy to move on-premises data residing in commodity file servers to AWS without changing application code or how data is managed. Amazon Elastic Block Store (Amazon EBS) Snapshots Archive is a new storage tier for Amazon EBS Snapshots that reduces the cost of archiving snapshots by up to 75%. AWS Backup now supports centralized data protection and automated compliance reporting for Amazon S3, as well as VMware workloads running on AWS and on premises. The new storage innovations announced today provide customers greater flexibility in how they manage storage while lowering costs and improving data management and protection capabilities.

“Every business today is a data business. One of the most important decisions that a business will make is where to store their data,” said Mai-Lan Tomsen Bukovec, Vice President, Block and Object Storage at AWS. “As these latest storage services and capabilities show, AWS is the most powerful and lowest cost way to access and protect data. Our rapid innovation makes it the best storage choice for customers now and in the future.”

Amazon S3 Glacier Instant Retrieval offers milliseconds retrieval for archived data at the lowest cost in the cloud—also available as a new tier in Amazon S3 Intelligent-Tiering

Amazon S3 offers a wide range of storage classes that deliver the lowest cost storage for different data access patterns. Customers often need to store petabytes of data that is only accessed occasionally, but that must be highly available and immediately accessible when requested (e.g. medical records, public health research data, media content, etc.). Today, customers have several options to store infrequently accessed data. Customers with data that is rarely accessed and requires retrieval times from a few minutes to a few hours can use S3 Glacier. Customers with data that is accessed once per month on average but still requires rapid retrieval can use S3 Standard-Infrequent Access (S3 Standard-IA) for only a slightly higher price. However, some customers want a combination of the lower storage costs offered by S3 Glacier and the fast retrieval of S3 Standard-IA, so they can meet their data access needs even more cost effectively.

S3 Glacier Instant Retrieval is a new storage class that is designed to offer milliseconds access for archive data, so customers can achieve the lowest cost storage in the cloud for data that is stored long-term and rarely accessed but requires immediate retrieval when requested. With S3 Glacier Instant Retrieval, customers no longer need to choose between optimizing for retrieval time or cost. Customers who move from S3 Standard-IA to S3 Glacier Instant Retrieval can save up to almost 70% for data that is accessed only a few times per year. Customers can now choose from three archive storage classes optimized for different access patterns and storage duration—S3 Glacier Flexible Retrieval (formerly S3 Glacier), S3 Glacier Deep Archive, and now S3 Glacier Instant Retrieval. S3 Glacier Instant Retrieval is the ideal storage class for customers who are sensitive to per-GB storage costs due to growing data volumes, by providing them the same low latency and high throughput of S3 Standard-IA, at the lowest cost for archive storage in the cloud.

S3 Glacier Instant Retrieval is now also available as a new access tier in S3 Intelligent-Tiering storage class. S3 Intelligent-Tiering optimizes storage costs by automatically moving data to the most cost-effective access tier based on access frequency without performance impact, retrieval fees, or operational overhead. Customers who need instant access to data and have unknown or changing access patterns can receive the same economic benefits as S3 Glacier Instant Retrieval with the new Archive-Instant Access tier in S3 Intelligent-Tiering without having to worry about where they are storing their data. To get started with S3 Glacier Instant Retrieval, visit aws.amazon.com/s3/glacier/instant-retrieval.

Amazon FSx for OpenZFS makes it easy to move data residing in on-premises commodity file servers to AWS without changing application code or how data is managed

Organizations of all sizes are migrating their on-premises data stores to the cloud to increase agility, improve security, and reduce costs. Today, many of these organizations store their data using on-premises file storage built on commodity, off-the-shelf servers and open-source software like the popular ZFS file system. These file servers provide access to data via industry-standard protocols, offer a wide variety of data management capabilities like point-in-time snapshots, cloning, and compression, and deliver hundreds of thousands of IOPS with sub-millisecond latencies. Many storage and application administrators have developed familiarity and expertise using tools that rely on the specific capabilities and performance of these file servers when running their applications. As a consequence, when migrating these applications to the cloud, storage and application administrators have to forgo the capabilities they are familiar with, and in many cases, re-architect their applications, tools, and workflows, which takes a lot of time and effort. These administrators would prefer to run their file servers on AWS to take advantage of improved agility, security, and cost but until now have not had the option of doing so.

With Amazon FSx for OpenZFS, customers can now launch, run, and scale fully managed file systems on AWS and replace their commodity, off-the-shelf servers they run on premises to achieve better agility, security, and lower costs. Amazon FSx for OpenZFS is the newest member of the Amazon FSx family of services that provides fully-featured and highly-performant file storage powered by widely-used file systems (including Amazon FSx for Windows File Server, Amazon FSx for Lustre, and Amazon FSx for NetApp ONTAP). Amazon FSx for OpenZFS is built on the open-source OpenZFS file system, which is widely used on premises to store and manage exabytes of application data for workloads that include machine learning, electronic chip design automation, application build environments, media processing, and financial analytics, where scale, performance, and cost efficiency are of utmost importance. Powered by AWS Graviton processors and the latest AWS disk and networking technologies, Amazon FSx for OpenZFS delivers up to 1 million IOPS with latencies of hundreds of microseconds. With complete

support for OpenZFS features like instant point-in-time snapshots and data cloning, Amazon FSx for OpenZFS makes it easy for customers to move their file servers to AWS, providing all of the familiar capabilities storage and application administrators rely on, and eliminating the need to perform lengthy qualifications and change or re-architect existing applications or tools. To get started with Amazon FSx for OpenZFS, visit aws.amazon.com/fsx/openzfs.

Amazon EBS Snapshots Archive storage tier reduces cost of archival snapshots by up to 75%

Today, customers use Amazon EBS Snapshots to protect data in their EBS volumes. Amazon EBS Snapshots provide incremental storage backups, retaining only the changes made to data in an EBS volume since the last snapshot. This makes Amazon EBS Snapshots cost effective for data that needs to be kept for days or weeks and requires retrieval within minutes. However, some customers also have business needs (e.g. snapshots created at the end of projects) and compliance needs (e.g. snapshots taken to audit recoverability) that require them to retain snapshots for months or years. While snapshotting data is a necessary business practice, customers also want ways to reduce the cost of these longer-term archival snapshots. To accomplish this today, some customers use third-party tools to move EBS snapshots to different tiers in Amazon S3, which increases costs and makes it complicated to track the lineage of these archival snapshots, but most customers simply absorb the increased cost of archiving snapshots for a long period of time.

To address the cost and complexity of archiving snapshots, Amazon EBS Snapshots Archive delivers a new storage tier that saves customers up to 75% of the cost for Amazon EBS Snapshots that need to be retained for months or years. Customers can now move their snapshots to EBS Snapshots Archive with a single application programming interface (API) call and reduce the cost of archival snapshots while retaining visibility alongside other Amazon EBS Snapshots. A Snapshot Archive is a full snapshot that contains all the blocks written into the volume at the moment that the snapshot is taken. To create a volume from the snapshot archive, customers can restore the snapshot archive to the Amazon EBS Snapshot standard tier, then create a volume from the snapshot in the same way they do today. To get started with Amazon EBS Snapshots Archive, visit aws.amazon.com/efs/snapshots.

AWS Backup brings centralized data protection and automated compliance auditing to Amazon S3 and VMware workloads

Today, customers use AWS Backup to meet their business continuity and regulatory compliance needs. AWS Backup enables customers to centrally protect their application data across AWS compute, database, and file and block storage services. Using a single data protection policy, customers can configure, manage, and govern backup and restore activity on Amazon Elastic Compute Cloud (Amazon EC2), Amazon EBS, Amazon Relational Database Service (Amazon RDS), Amazon Aurora, Amazon DynamoDB, Amazon DocumentDB, Amazon Neptune, Amazon FSx, Amazon Elastic File System (Amazon EFS), and AWS Storage Gateway. To meet evolving regulatory requirements, customers can opt for automated, continuous backup compliance monitoring and generate auditor-ready reports using AWS Backup Audit Manager. To protect against accidental or malicious deletions (e.g. in the case of a ransomware attack), customers can use fine-grained access controls built into AWS Backup as well as use AWS Backup Vault Lock to make their backups immutable. In addition, AWS Backup's integration with AWS Organizations enables customers to extend their data protection policy across multi-account deployments and use its cross-Region and cross-account backup capabilities to achieve global resiliency and durability for their mission-critical data.

Now with AWS Backup support for Amazon S3 and VMware workloads, AWS is extending AWS Backup's capabilities to more cloud and on-premises workloads. Previously, administrators would write custom

scripts to combine S3 data across multiple AWS Regions and accounts with AWS Backup to get a consolidated view of their backups, as well as combine S3 reports with AWS Backup built-in reports to demonstrate compliance with application-level backup policies. Administrators would also parse through S3 data to find the point-in-time they want to restore an application. Now, with AWS Backup support for S3, customers can replace the complicated custom scripts they used to centrally manage backups of their entire applications. Customers can also now replace parsing through S3 data with AWS Backup's point-in-time restore functionality, allowing them to specify the time to restore—down to the second. To learn more about AWS Backup for Amazon S3, visit aws.amazon.com/backup.

With AWS Backup support for VMware workloads, customers can now protect their VMware workloads whether they run on premises or in the VMware Cloud on AWS (VMC). Up until now, customers running VMware workloads on premises and in the AWS Cloud maintained separate solutions to protect their data alongside AWS services already supported by AWS Backup. This required them to manage separate tools and policies to protect their data in VMware environments, as well as generate distinct reports to demonstrate compliance with backup policies. With AWS Backup for VMware, customers can extend AWS Backup's centralized data protection, governance, and compliance features they already use to protect their AWS applications to their VMware workloads, whether running in AWS or in their own datacenters. To get started with AWS Backup for VMware, visit aws.amazon.com/backup.

Capital One offers a broad spectrum of financial products and services to consumers, small businesses, and commercial clients through a variety of channels. "We wanted to find a way to quickly optimize storage costs across the largest and fastest growing S3 buckets across the enterprise. Because the storage usage patterns vary widely across our top S3 buckets, there was no clear-cut rule we could safely apply without taking on some operational overhead," Jerzy Grzywinski, Director of Software Engineering at Capital One. "The S3 Intelligent-Tiering storage class delivered automatic storage savings based on the changing access patterns of our data without impact on performance. We look forward to S3 Intelligent-Tiering's new access tier, which will allow us to realize even greater savings without additional effort."

The National Association for Stock Car Auto Racing, Inc. (NASCAR) is the sanctioning body for the No. 1 form of motorsports in the United States. "The new Amazon S3 Glacier Instant Retrieval storage class provides low storage cost for the NASCAR Library, which houses our growing media archives and enables our content creators to interact with data of any age in near real time," said Chris Wolford, Senior Director of Media & Event Technology at NASCAR. "We manage one of the largest racing media archives in the world. Our customers, who range from the NASCAR Cup series teams to producers, editors, and engineers, generate video, audio, and images, many of which are stored in perpetuity. The new storage class will help us save on our storage cost while greatly improving on our restore performance. Now, we can benefit from lower storage cost, with the resiliency of multi-AZ storage, and with immediate retrievals for any media asset!"

Epic Games is the interactive entertainment company behind Fortnite, one of the world's most popular video games with over 400 million players. Founded in 1991, Epic transformed gaming with the release of Unreal Engine—the 3D creation engine powering hundreds of games now used across industries, such as automotive, film and television, and simulation, for real-time production. "Using S3 Intelligent-Tiering, we can implement storage changes without interruptions to service and activity. Our data is automatically moved to lower-cost tiers based on data access, saving us a lot of development time in addition to reducing costs," said Joshua Bergen, Cost Management Lead at Epic Games. "With that time, my team can focus on identifying other opportunities to reduce infrastructure costs in support of our

organizational goals. The new access tier in S3 Intelligent-Tiering will help us save even more on storage costs.”

STEMCELL Technologies is a global biotechnology company that supports life sciences research with more than 2,500 specialized reagents, tools, and services. “Many of our departments, including quality and finance, have a need for long-term data retention to meet regulatory requirements for our products and services,” said Hikaru Mathieson, Senior Systems Engineer at STEMCELL Technologies. “We have a large inventory of EBS Snapshots supporting our use of services like AWS Storage Gateway and Amazon EC2. Maintaining this inventory is an important piece of our regulatory compliance and an easy transition of snapshots over to archival-tier storage has been a long-desired dream. We are excited about EBS Snapshots Archive for its ease-of-use and low cost for long-term retention of our snapshots. In the future, we plan to consolidate more of our long-term backups, such as general department file backups, into EBS Snapshots Archive.”

Zilliant, Inc. is the industry leader in intelligent B2B price optimization, price management, and sales guidance SaaS software. “We backup our active volumes into EBS Snapshots and retain them for 14 days. However, we need to retain many of our snapshots for months or years, so we use scripts to manage snapshot data lifecycle into lower-cost, colder tiers,” said Shams Chauthani, CTO at Zilliant. “Maintaining and managing these scripts is getting complex at scale. We are delighted to use EBS Snapshots Archive, as it eliminates the need to maintain scripts and enables creation of secure end-to-end flows for cost effective archival of our snapshots.”

Loews Corporation is a diversified company with businesses in the insurance, energy, hospitality, and packaging industries. “We currently have server instances running in AWS and in our data center on premises,” said Emilio Renzullo, Senior Engineer, Infrastructure Services at Loews. “We are currently using AWS Backup to protect our Amazon EC2 instances, and another product to backup our on-premises VMware virtual infrastructure to Amazon S3 for archiving. AWS Backup’s new VMware capability will enable us to streamline and centralize all our backup operations, while also cutting down on costs.”

About Amazon Web Services

For over 15 years, Amazon Web Services has been the world’s most comprehensive and broadly adopted cloud offering. AWS has been continually expanding its services to support virtually any cloud workload, and it now has more than 200 fully featured services for compute, storage, databases, networking, analytics, machine learning and artificial intelligence (AI), Internet of Things (IoT), mobile, security, hybrid, virtual and augmented reality (VR and AR), media, and application development, deployment, and management from 81 Availability Zones within 25 geographic regions, with announced plans for 27 more Availability Zones and nine more AWS Regions in Australia, Canada, India, Indonesia, Israel, New Zealand, Spain, Switzerland, and the United Arab Emirates. Millions of customers—including the fastest-growing startups, largest enterprises, and leading government agencies—trust AWS to power their infrastructure, become more agile, and lower costs. To learn more about AWS, visit aws.amazon.com.

About Amazon

Amazon is guided by four principles: customer obsession rather than competitor focus, passion for invention, commitment to operational excellence, and long-term thinking. Amazon strives to be Earth’s Most Customer-Centric Company, Earth’s Best Employer, and Earth’s Safest Place to Work. Customer reviews, 1-Click shopping, personalized recommendations, Prime, Fulfillment by Amazon, AWS, Kindle

Direct Publishing, Kindle, Career Choice, Fire tablets, Fire TV, Amazon Echo, Alexa, Just Walk Out technology, Amazon Studios, and The Climate Pledge are some of the things pioneered by Amazon. For more information, visit amazon.com/about and follow @AmazonNews.

Media Hotline

Amazon-pr@amazon.com

www.amazon.com/pr