

Simon Elisha [\(00:02\)](#):

Hello everyone and welcome back to the AWS podcast. Simon Elisha here with you. Great to have you back. And I'm flying solo today. Totally my fault. I rescheduled a number of times, and made it impossible for Jillian to join me, so, my apologies, but she'll be back for the next one. We have a lot of updates today. In fact, we have 75, so I won't waste your time. We'll move through at a rapid clip, but we'll dive into some things, too.

[\(00:25\)](#):

Firstly, let's talk about the AWS Marketplace ...

[\(00:31\)](#):

Firstly, let's talk about the AWS Marketplace. There is now an enhanced private offer user experience for sellers. So, AWS Marketplace sellers can now directly view and manage their active agreements in the portal, and they can set custom dimensions for future reuse when creating a private offer for AMIs, containers, software as a service, and professional services products. So this means you can be even more flexible for your own customers.

[\(00:58\)](#):

Let's talk about analytics. AWS Glue now offers example jobs, drag and drop, and a preview first layout. So this gives you a much simpler UI experience, less clicks to do your extract transform and load, which means you can author, run and monitor your AWS Glue ETL jobs really, really quickly. AWS Glue for Apache Spark has also announced native connectivity for Google BigQuery and Amazon MSK, the managed streaming for Apache Kafka, has provided event bridge pipes console integration. So, now it's even easier to send events from your Apache Kafka cluster to one of over 14 AWS service targets. And Amazon MSK has also announced MSK replicator for cross region and same region streaming data replication. A couple of quick updates for Amazon QuickSight. It now supports direct connectivity to Trino and also to Starburst.

[\(01:56\)](#):

And Amazon Redshift now announces integration with AWS secrets manager to simplify the management of your Redshift administration credentials for your data warehouse. With this feature, Amazon Redshift works with AWS secrets manager to generate and manage your admin credentials when a database instance is created, modified or restored, so it supports the entire lifecycle maintenance for those admin credentials, which means you don't have to manage them yourself. So Pretty nifty. Also order complete suggestions and syntax error checking is now available in Amazon Redshift query editor v2. So, this feature empowers you to build queries more efficiently and accurately, and you can quickly access and utilize relevant database object names, without the name to remember each and every detail, by providing SQL syntax or database object names suggestions, either within the editor tab or a notebook. The suggestions are sorted by relevance, and they indicate availability as well as help with spelling difficult database object names.

[\(02:56\)](#):

Amazon Open Search Service Now supports efficient vector query filters for face. Users can now benefit from more efficient filtering...

[\(03:09\)](#):

Users can now benefit from more efficient query filtering with open searches K and N face engine, using open search 2.9 on the Amazon OpenSearch Service. Previously, exclusive to OpenSearch's loosing K and N engine, OpenSearch's efficient vector query filters capability intelligently evaluates optimal filtering

strategies, like pre-filtering with appropriate nearest neighbors, or filtering with exact K nearest neighbors, to determine the best strategy to deliver accurate and low latency vector search queries.

[\(03:40\):](#)

In early OpenSearch versions, vector queries on the face engine used post filtering techniques which allowed filter queries to scale, but potentially gave you less than the requested K number of results. And Amazon OpenSearch Service now also supports internet protocol version 6, our old friend IPV 6. And Amazon OpenSearch Service...

[\(04:05\):](#)

And Amazon OpenSearch Serverless now supports automated time-based data deletions, so you can make it easy for you to manage those data volumes. And the Amazon OpenSearch Service also has new administrative options, including the ability to restart the OpenSearch process on a node, and the ability to restart a data node. And, finally, Amazon Fin Space with managed KDB insights adds more historical database case options. This gives you more flexible choices when choosing where to store information.

[\(04:37\):](#)

Let's talk about application integration. Amazon SNS now supports in-place messaging, archiving, and replay for FIFO topics. So these are our first in first out topics, and this is a no-code in-place option to store and replay your messages, without the need to provision a separate archive or resource. This improves the durability of your event-driven applications and can help you recover from downstream failure scenarios. So, topic owners can now set an archive policy which defines a retention for the messages published to their topic. Subscribers can then set a replay policy to an individual subscription which triggers a replay of selected messages from the archive from a starting point, until an ending point. Subscribers could also set a filter policy on their subscription to further select the messages in scope for replay. So this is a very, very useful feature.

[\(05:29\):](#)

And Amazon SQS has also announced an increased throughput quota for FIFO high throughput mode. It now allows you to process up to 18,000 transactions per second, if you are doing this in the Ohio, North Virginia, Oregon, Island, Frankfurt regions. For the Mumbai, Singapore, Sydney, and Tokyo regions, the throughput quota has been increased to 9,000 transactions per API, and for London and Sao Paolo, it's been increased to four and a half thousand transactions. For all other regions, it is 2,400 transactions per second. And Amazon SNS message data protection now supports custom identifiers. So, this allows you to help security and engineering teams have real-time data protection in your applications and you can use it to protect domain specific sensitive data, things like your company's employee IDs.

[\(06:19\):](#)

A quick update in the area of blockchain, Amazon Managed Blockchain, AMB Query, now supports the Sepolia Testnet, delivering foundational balance and transactional data with full history and subsecond latency, without the need for specialized infrastructure. With this release, AMB Query now supports the Sepolia Testnet and completes its testnet and main net coverage for Bitcoin and Ethereum, helping developers from test to production.

[\(06:48\):](#)

Now it's time for some compute updates. And Amazon EC2 has enabled the Amazon machine image block public access, sorry, let me do that again...

[\(07:01\):](#)

Amazon EC2 has enabled the Amazon machine image, "block public access" setting by default for all new accounts and all existing AW accounts that have not owned a public AMI since July 15th, 2023. AMI BPA

restricts an AWS account from accidentally sharing an AMI publicly within an AWS region, and it helps improve the security and privacy posture. Previously, AMI BPA was disabled by default for all AWS accounts. With this, changes is now a default setting for all new accounts, and, as I mentioned, it is also enabled for existing accounts that have not owned a public AMI since July the 15th, 2023.

[\(07:43\):](#)

EC2 Hibernate now supports 20 additional instance families on EC2 Spot, so you can get the benefits of Hibernate and the supported instance families, whether you are using on-demand, Spot, or savings plan.

[\(07:54\):](#)

So it's a good way to have things paused and come back to life. AWS Marketplace announces a new subscription experience for the Amazon Elastic Kubernetes service console, which allows you to subscribe to Kubernetes software from leading ISVs, directly in the EKS console, without having to visit the AWS Marketplace website. Amazon EKS has added support for customer managed IAM policies, which means you can help meet regulatory and compliance requirements with very fine grain controls about those IAM permissions that your Kubernetes clusters can assume. And EKS also now allows modification of cluster subnets and security groups.

[\(08:35\):](#)

And the Amazon Elastic Container Service, ECS, has enhanced task scheduling to make customers applications even more resilient to unpredictable load spikes. Now, Amazon ECS will first start a healthier replacement for each unhealthy task that failed to pass a container or load balancer health check before terminating it. This enhancement increases the resilience of customers' applications with no additional efforts or configuration. And finally, the AWS data provider for SAP now supports Oracle Linux OS, more JDK versions, and Linux Logrotate.

[\(09:09\):](#)

A quick update on the topic of customer engagement. Amazon Connect now provides additional levels of granular access controls, so you've got things like department, retail, different teams, different types, locations, lots and lots of choices.

[\(09:23\):](#)

So, now let's talk about databases. Amazon RDS blue/green deployments now support Aurora and RDS PostgreSQL. Now blue/green deployments support a safer, simpler, and faster way to update your databases. They create a fully managed staging environment using PostgreSQL community logical replication that allows you to deploy and test your production changes, keeping your current production database safer. With just a few clicks, you can promote the staging environment to be the new production system, in as fast as a minute, with no data loss and no changes to your application, to switch database endpoints. So, use Amazon RDS blue/green deployments for deploying changes to production, like major version database engine upgrades, schema changes, maintenance updates, and scaling instances.

[\(10:10\):](#)

Blue/green deployments use built-in switchover guardrails that will time out promotion of the staging environment if it exceeds your maximum total rule downtime, the text replication errors, or identifies instance health check errors. You can also now enable Amazon RDS optimized rights using your RDS blue/green deployments as well. Amazon Relational Database Service also announces the dedicated log volume. This is for PostgreSQL, MySQL and MariaDB databases. An Amazon RDS dedicated log volume allows customers to select a configuration with the most latency sensitive components of their database, or transaction logs, are stored on a separate dedicated volume. Dedicated log volumes work

provisioned IOPS storage, and are recommended for databases with 5,000 gig or more of allocated storage.

[\(11:03\)](#):

Amazon Aurora PostgreSQL now supports federated query to MySQL and MariaDB databases. So you can use this using the foreign data wrapper library, which allows the Postgres databases to communicate with an external data source, abstracting details by connecting to the data source and obtaining data for it. Amazon Aurora PostgreSQL now captures query plans on replica instances with query plan management. So, this also includes some other improvements, like limiting the capture of query plans based on the estimated cost and other cool optimization things, and some new versions support for Amazon aurora for PostgreSQL 15.4, 14.9, 13.12, 12.16, and 11.21 are all supported. Your regular reminder from me: patch your stuff. It is the best way to protect yourself.

[\(11:52\)](#):

And Amazon Aurora PostgreSQL now supports PG vector version 0.5.0 with HNSW indexing. So this is hierarchical, navigable, small world indexing support. It also has parallelization of IV flat index builds, and improves performance of its distance functions.

[\(12:12\)](#):

Amazon Aurora MySQL 3.05, which is compatible with MySQL 8.0.32, is generally available, and we're also happy to announce the Amazon Aurora MySQL 3.04 as the long-term support release. During the lifetime, we will support it for at least three years, and we'll provide new patches to introduce fixes to critical matters such as security issues. Those patches don't include any new features. And Amazon Aurora MySQL has improved optimizations that reduce the database restart time by up to 65%, compared to what it did before. These are achieved by deferring portions of the buffer pool initialization and the validation process to occur after the database is already online and accepting connections. So this is really useful for those unplanned events like unexpected database restarts or planned operations like minor version upgrades. Amazon RDS for SQL server also now supports minor version 2019 CU 2022. Hang on.

[\(13:15\)](#):

An Amazon RDS for SQL Server now supports minor version 2019 CU22.

[\(13:21\)](#):

Let's talk a little bit about developer tools. AWS X-Ray now supports WC...

[\(13:30\)](#):

AWS X-Ray now supports W3C trace IDs, generated via open telemetry and other frameworks that can form to the W3C trace context specification, which gives customers complete end-to-end visibility for faster triaging and debugging of applications. Trace IDs are unique identifiers for each request passed from one service to another, and these are used for distributing tracing across systems Popular development tools like Nginx, Envoy Proxy, Kubernetes components and more emit traces using the open telemetry protocol and generate trace IDs in the W3C format. Customers are unable to configure these external development tools to generate x-ray format trace IDs, and view traces from those in the X-ray console, which means you can't see them. With this new extended support of those trace IDs, you can now view traces from these external development tools, along with existing X-ray format traces in the X-ray console across your distributed systems.

[\(14:27\)](#):

A few Code Pipeline, AWS code pipeline improvements, it now supports parameterized pipeline, it supports triggering on Git tags and it also supports retrying a pipeline execution from the first action in a failed stage. And we're also announcing Amazon Corretto...

[\(14:43\)](#):

And we're also announcing Amazon Corretto, October, 2023 quarterly updates, so this is the update for the Amazon Corretto long-term support versions of open JDK. Let's talk about end user computing. Amazon AppStream 2.0 has announced multi-session fleets for Windows Server OS. This new feature enables IT admins to host multiple end user sessions on a single upstream 2.0 instance, helping make better use of resources, by allowing multiple users to share compute memory and storage resources associated with a given AppStream 2.0 instance. Multi-session capability gives you improved resource utilization and reduces operational cost per user compared to that one-to-one user to instance management, it also allows you to fine tune and optimize performance and resource allocation.

[\(15:46\)](#):

Amazon workspaces now supports Windows Server 2022 bundles, so you can get the latest server operating system features, and Amazon Workspaces has introduced the graphics G4dn bundles for Ubuntu desktops. So these later run graphics intensive and compute intensive workloads on Ubuntu desktops with 4K monitors in the cloud and they give you cost-effective solutions for graphics applications that are optimized for NVIDIA GPUs, using NVIDIA libraries such as CUDA, cuDNN, OPDX, and video Kodak SDK.

[\(16:20\)](#):

Now it's time to talk about the internet of things. AWS IOT Express Link. A connectivity software that powers a range of hardware modules developed by and offered by AWS partners, has announced the release of technical specification 1.2. The updated specification provides hardware manufacturers who design and develop IOT devices, a new command set to allow host software applications to communicate wirelessly using Bluetooth low energy or BLE technology.

[\(16:49\)](#):

Using the BLE command groups offered in the new specification, makes it much easier for host processes to take control of the personal area network, the PAN, and communicate with other IOT mobile devices over Bluetooth.

[\(17:04\)](#):

And AWS IT core credential provider now supports virtual private cloud endpoints...

[\(17:13\)](#):

So, let's talk about machine learning. Introducing Amazon Code Whisperer customization capability in preview. With this new capability, you can customize Code Whisperer to generate more precise suggestions by including your organization's internal APIs, libraries, classes, methods, and best practices. With this capability, you can save developer hours that are typically spent looking at previously written code or documentation to figure out how those internal APIs work. And you can use this capability to securely connect your organization's private repositories.

[\(17:48\)](#):

Amazon Kendra has launched a feature to collapse and expand query results. What this does is, means that...

[\(18:00\)](#):

When users search for accurate answers from Amazon Kendra, they encounter situations where there are documents which are near duplicates, related to another in various ways, and need to be viewed together, or are closely linked with each other in the form of a parent-child relationship.

[\(18:14\)](#):

Users need to be able to view these documents together, as a group, for review and analysis and also next action identification. So now you can do it much more easily. Amazon Kendra has also released connectors for 11 JDBC data sources, which enables you to have structured data search, so things like Aurora MySQL compatible, RDS MySQL, RDS MySQL server, et cetera, et cetera. You can now build ML feature pipelines from custom data sources with Amazon SageMaker feature store. So, you can build richer and more varied ML features by incorporating diverse data sources, defining transformation functions to perform, and have the SageMaker feature store take care of the processing of that data into ML features. So you can now connect to streaming data sources like Amazon Kinesis and Author transforms with Spark structured scheming...

[\(19:08\)](#):

And author transforms with Spark structured streaming, which is a scalable and fault tolerance stream processing engine for real-time data processing. You can also connect to data warehouses, like Amazon Redshift, Snowflake, and Databricks for batch feature processing, and you can initiate feature processing on a schedule, or with a trigger using Amazon event bridge rules.

[\(19:33\)](#):

Amazon SageMaker has also announced a preview of MLP 5/48 XL instances for inference. So these have 80 gig of memory, per NVIDIA H100 Tencent gpu for a 640 gig total, 30 terabytes of local NVME SSD storage, 192 VCPUs, and two terabits of...

[\(19:55\)](#):

And two tebibytes of instance memory. My goodness me, these are good to get hands on. And you can personalize search results with Amazon Personalize and Amazon OpenSearch service integration. The Amazon Personalize search ranking plugin with the OpenSearch Service lets you leverage the deep learning capabilities of Amazon Personalize and add that to your search results without any ML expertise.

[\(20:21\)](#):

And AWS Neuron has added support for Llama-2 70b model and PyTorch 2.0. Now, AWS Neuron is the SDK for Amazon EC2 Inferentia and Trainium based instances, purpose-built for generative AI today, with the Neuron 2.15 release renouncing support for Llama-2 70b model training as well as PyTorch 2.0 support. Neuron integrates with popular ML frameworks like PyTorch and TensorFlow, so you can get started with minimal code changes and without vendor specific solutions. Neuron includes a compiler, runtime, profiling tools, and libraries to support high performance training of generative AR models on TRN one instances and inferences on, in two instances.

[\(21:06\)](#):

And, finally on this topic, AWS Entity Resolution has expanded its data protection capabilities. Let's talk about management and governance. AWS announces an enhanced signing...

[\(21:24\)](#):

We're happy to announce an enhanced sign-in experience in the AWS console mobile app for IOS, so it's betterer, and an enhanced sign-in experience for the app for Android is also coming soon. AWS systems manager, OpsCenter now supports deletion of operational issues, and AWS systems manager, patch manager, now supports additional Mac OS and Linux versions. AWS config now supports 19 new

resource types, so you can have extra control over your environment, and AWS service catalog has announced support for additional infrastructure as code provisioning tools. So now you can create, distribute, and launch AWS resources that are configured using third party infrastructure as code tools, like Ansible, Chef, Pulumi, Puppet, and more.

[\(22:14\)](#):

AWS Repost has introduced selections. So, this is a feature that organizes AWS knowledge for a specific topic, technical issue, or use case. So it makes it much easier for you to aggregate stuff as you're looking at things. And we're happy to also announce new options for credit sharing. The billing preferences page of the billing console now allows you to manage credit sharing preferences for your organization at a more granular level by selecting which accounts from your organization to share your credits with, and exclude others from sharing.

[\(22:47\)](#):

AWS Resource Explorer now supports 12 new resource types, and metric streams has added a quick setup experience for AWS partner destinations. So this allows partners to use Amazon CloudWatch to very quickly create a continuous near real-time stream of metrics to a destination of your choice.

[\(23:08\)](#):

A quick update in the area of media services. The Amazon Interactive Video Service, IVS, now offers audio only pricing for its real try...

[\(23:19\)](#):

The Amazon Interactive Video Service, IVS now supports audio only pricing for its real-time streaming capability at one 10th of the existing real-time streaming rate, so, got to love that. Let's move on to migration and transfer. AWS Migration Hub has added the ability to change the AWS region used to store migration data. So the AWS Migration Hub home region gives you a single repository for entire migration portfolio, regardless of the migrating application's target region, or account. This includes all the infantry discovered with the AWS application discovery service, Amazon EC2 or migration strategy recommendations and migration strategy tracking statuses set by linked services like the application management service or the database migration service.

[\(24:07\)](#):

AWS Transfer Family has announced a basic authentication for outbound IS2 messages. And the AWS application discovery service has added new actions to remove servers, data collections, and imports. Some interesting updates in the world of networking and content delivery. Firstly, we're happy to announce the general availability of multi virtual private cloud Elastic Network interface attachments for Amazon Elastic Compute Cloud instances. Say that far three times. With this capability, customers can now launch an instance with a primary ENI in one VPC and attach a secondary ENI from another VPC. With Multi-VPC ENI attachments, customers can maintain VPC level segregation between networks, whilst allowing select workloads like centralized appliances and databases to communicate between them.

[\(25:06\)](#):

Customers such as communication...

[\(25:11\)](#):

Customers such as telecommunications providers that send different types of network traffic within AWS regions, as well as between on-premises and AWS, need logical separation for their workload. Using different VPCs for different types of network traffic, such as control and data plan traffic, customers can now segment their traffic at the VPC level, while shared traffic workloads like virtual

routers, firewalls, and databases can maintain connectivity between the two. With Multi-VPC ENI attachments, customers can run connected workloads between those otherwise segmented VPCs to ensure specific functions can communicate with each VPC.

[\(25:51\)](#):

Amazon Route 53 Traffic Flow now supports geo proximity routing for AWS local zones, and you can share Amazon Route 53 application recovery controller clusters across multiple accounts. AWS Cloud WAN goes tunnelless to enable high-performance global SD-WANs. With this capability, third party SD-WAN appliances can natively integrate with Cloud WAN without the need for any specialized tunneling protocols. And finally, today we're going to talk about security, identity, and compliance. And the AWS Security Hub has launched 19 new security controls. So this increases the number up to 292. So Security Hub now supports four additional AWS services, Amazon Event Bridge, managed streaming for Apache Kafka, Amazon MQ, and Amazon Route 53.

[\(26:44\)](#):

And Security Hub has also released new controls against eight previously supported services that allow you to have an even better coverage for your environment. My goodness me, there was a lot. And I'll give you a warning right now it's November and November is reinvent month, and so lots of goodies are on the way. Please keep listening, tell others that the podcast is around, and we do love to get your feedback. AWSpodcast@amazon.com is the place to do it. And, until next time, keep on building.