Simon Elisha (00:00):

Hello everyone, and welcome back to the AWS Podcast. So much I want to share with you, great to have you back. And of course I'm joined by my cohost, Jillian Forde. G'day Jillian, how are you doing?

Jillian Forde (00:07):

G'day, I'm always excited to be here for the update show.

Simon Elisha (00:10):

It's an important show, it helps keep everyone-

Jillian Forde (00:12):

Very.

Simon Elisha (00:12):

... across everything that's going on. Now, a quick programming-

Jillian Forde (00:15):

Including me.

Simon Elisha (00:16):

... if I'm sounding different, I had an allergic reaction that meant my tongue expanded, which is not a good thing.

Jillian Forde (00:22):

I can't believe that you're doing this.

Simon Elisha (00:25):

But it's de-expanded, but there's a little bit of carryover, so I sound a little different. But you know, the show must go on, as they say.

Jillian Forde (00:34):

That's right.

Simon Elisha (00:35):

So let's start with the topic of analytics. AWS has announced the general availability of Amazon Aurora MySQL, zero-ETL integration with Amazon Redshift. Now, my favorite kind of ETL is zero-ETL, because ETL is very, very hard. Now, you don't need to build and maintain complex data pipelines to perform, extract, transform and load operations, Amazon Aurora MySQL, zero-ETL integration with Amazon Redshift is now available for Amazon Aurora Service V2, and provisioned, and Amazon Redshift Serverless, and RO3 instance types. Now, zero-ETL integration helps you derive holistic insights across many applications, and it helps break down those data silos in your organizations. You can enhance your data analytics with the rich analytic capability of Amazon Redshift, things like high-performance SQL, built-in ML and Spark integrations, materialized views, and so much more.

(01:31):

Glue Studio visual now supports interactive data previews, so you can now see what you're going to get as you're trying it out. And Amazon Redshift has launched MaxRPU, which is a new compute cost-control setting for Redshift Serverless, where you can optionally specify an upper-compute threshold. Amazon Relational Database Service, database activity streams, capture events from the database, encrypts them and uploads them to an Amazon kinesis data stream within your account, and it is now available in on-demand mode. So starting now, once you enable database activity streams we'll create a kinesis data stream in on-demand mode instead of provision mode. This change will only apply to any new database activity stream enablements, and will not affect anything existing.

(02:14):

And if you wish to move from on-demand to provisioned you can, just through the console or the API. Another quick announcement, Amazon Kinesis video streams, web RTC ingestion is now generally available, so this is really useful for any web RTC stance-compliant cameras, IOT devices and browsers. And QuickSight has launched the float data type support for Spice datasets. Before we enable that Spice-supported fixed decimal type, so now you have the choice between exact mathematical operations with less decimal places, or float, which brings its own wonderful, wonderful processing, thinking process. Amazon MSK now adds a check for too many partitions to AWS Trusted Advisor, so this helps you optimize your cost and improve your performance. And Amazon MWAA now supports Apache Airflow Version 2.7, and deferrable operators as well.

(03:05):

Reminder, update your stuff. And Amazon OpenSearch Service has introduced Neural Search, now this is in OpenSearch 2.9, which enables builders to create and operationalize semantic search applications with a reduced, undifferentiated and heavy lifting, our old friend. For years, customers have been building semantic search applications on OpenSearch KNN, but they've been burdened with building middleware to integrate text-inverting models into search and ingest pipelines. Amazon open cert service customers can now power Neural Search through integrations with Amazon Sagemaker and Amazon Bedrock, enabling semantic search pipelines run on cluster. So with Neural Search, customers can use OpenSearch APIs to run queries via human language, and deliver more relevant results using text embeddings that incorporate semantic understanding and similarity. And customers can integrate text embeddings hosted on Amazon Sagemaker, or managed by Amazon Bedrock, and you can get up and running super quick.

Jillian Forde (04:00):

Up next, bringing all the things together with application integration. So, Amazon SQS announces support for JSON, enabling lower-latency and improved performance for SQS customers. These [inaudible 00:04:15] performant tests for a five kilobit message payload, JSON protocol for Amazon SQS reduces end-to-end message processing latency by up to 23%, and reduces application client-side CPU and memory usage. The latest STK version achieves these latency gains by upgrading the default communication protocols used by clients to JSON wire protocol when they make SQS API requests. Amazon SQS announces Amazon EventBridge Pipes with a console integration. AWS Lambda now supports up to five times faster polling scale uprate, for Spiky Lambda workloads configured with Amazon SQS as an event source, using Lambda Event Source mapping, or Amazon EventBridge pipes.

(05:06):

This enables customers building event-driven applications using Lambda and SQS cues to achieve more responsive scaling during a sudden burst of messages in their queues, and reduces the need to duplicate Lambda functions or SQS queues to achieve faster message processing. Amazon SNS increases the default FIFO topic throughput by 10x to 3,000 messages per second, and Amazon EventBridge now

supports over 20 new CloudWatch metrics for event buses. Up next we've got compute, AWS Lambda launches Auto Scaling improvements for Amazon MSK and self-managed Apache Kafka as event sources. With this launch, customers will now experience faster scaling and more consistent throughput while processing Kafka messages with Lambda functions. AWS Lambda adds support for Amazon Linux 2023.

(06:02):

AWS Lambda now makes it easy to monitor and troubleshoot Lambda function failures by reporting errors and timeouts during initialization and restore phases of the Lambda function lifecycle to Amazon Cloudwatch logs. Now Lambda automatically captures and sends logs about each phase of the Lambda execution environment lifecycle to Cloudwatch logs. With this launch, if an error or timeout occurs during the init or restore phase, Lambda sends log messages containing error or error timeout details to Cloudwatch logs. This makes it easy for you to more effectively monitor and troubleshoot failures that occur during every phase of the Lambda function lifecycle. AWS CodeBuild customers can now use AWS Lambda to build and test their software packages, Amazon EKS Anywhere Enterprise subscriptions can now be purchased through the Amazon EKS console, APIs and AWS CLI.

(07:00):

You can purchase Amazon EKS Anywhere Enterprise subscriptions for 24/7 support from AWS [inaudible 00:07:07] Experts, and access to Amazon EKS Anywhere curated packages. Amazon EKS Anywhere is an AWS-supported Kubernetes offering, with automation tooling to simplify running Kubernetes in on-premises and Edge environments. When purchases are submitted, licenses are automatically issued in AWS License Manager and access is granted to Amazon EKS Anywhere-curated packages. Amazon Cloudwatch Container Insights now delivers enhanced observability for Amazon EKS with out-of-the-box detailed health and performance metrics, including container level EKS performance metrics, cube state metrics, and EKS control plane metrics for faster problem isolation and troubleshooting.

(07:51):

Customers running applications with more than one container on Amazon ACS with Fargate can now leverage seekable OCI to lazily load specific container images within the Amazon ECS task definition. This eliminates the need to generate seekable OCI indexes for smaller container images within the task definition, while still getting the benefits of SQL OCI with larger container images, therefore improving the overall application deployment and scale out time. EC2 Image Builder adds Amazon Linux 2023 and Ubuntu 22.04, LTS-managed image support. AWS adds ability for shared Windows AMIs, users to enable faster launching. AWS batch array jobs now support array size environment variable. And the last one is you can now deploy SAP applications programmatically using AWS Launch Wizard APIs. That was a lot of updates.

Simon Elisha (08:46):

There was a lot of cool things there, some great performance things too, which, I always love those performance kicks. It's like, "Yes."

Jillian Forde (08:53):

Oh, absolutely.

Simon Elisha (08:53):

"Go faster, do more." Let's talk a bit about cost management. AWS Cost Management now provides purchase recommendations for Amazon MemoryDB reserved nodes. So, MemoryDB reserved nodes allow you to save up to 55% over on-demand node prices, in exchange for usage commitment over a

one or three-year term. So now with these recommendations, you can get that done really, really quickly and know exactly what is going to make the most sense for you. Now, let's discuss customer engagement. A few updates here in the world of Amazon Connect, firstly we've reduced telephony pricing across South America, I do love myself a good price cut. Amazon Connect has also reduced the outbound telephony pricing across Europe, so yet another one. And Amazon Connect prompts configuration page provides CloudTrail coverage, so now you can see what's going on from your CloudTrail records.

(09:42):

And Amazon Connect now supports representative numbers in South Korea, this includes all five tiers of representative numbers ranging from two concurrent calls for the smallest representative number, I should say, up to 5,000 concurrent calls for the large-scale services. Amazon Connect now makes it easier to deliver persistent chat experiences for your customer, so persistent chats enable customers to resume previous conversations with the context metadata and transcript carried over, which means that you don't have to repeat yourself all the time, and it means agents can have an understanding of what the conversation history was. And finally on this topic, Amazon SES now offers 60 days of metric history in Virtual Deliverability Manager, this means you can understand what's been going on.

(10:24):

Now let's talk databases, Amazon Aurora Global Database for PostgreSQL now lets you forward write requests from a secondary to a primary region, simplifying writing from multiple regions and making it easier to develop globally-distributed applications. With this launch, write forwarding is now available on both Aurora Global Database for MySQL and PostgreSQL. We're also announcing support for Amazon Aurora PostgreSQL Optimized Reads, which is a new price performance capability available on the new R6GD and R6ID instances that give you up to eight times' improved query latency and up to 30% cost savings compared to instances without it. So, this is really useful for applications with large datasets that exceed the memory capacity of a database instance.

(11:12):

Optimized reads uses the local NVME-based SSD block level-storage, which reduces the data accesses to and from network-based storage, which helps your read latency and your throughput. Amazon RDS Performance Insights now supports exporting dashboards into CloudWatch, and Amazon RDS Proxy, which is a fully-managed, highly-available database proxy for Amazon RDS, now supports RDS Multi-AZ deployments with two readable standbys. Amazon RDS Proxy allows applications to pull and share connections established with the database, which improves database efficiency and application scalability. And Amazon RDS Proxy also now support PostgreSQL extended query protocol, and Amazon RDS for MySQL now delivers up to three times' higher write throughput at no additional charge. So here we've got improved performance and no cost.

Jillian Forde (12:00):

I love it.

Simon Elisha (12:01):

This is good. So, this starts with RDS for MySQL version 8.0.35, so you should upgrade it and get the benefit.

Jillian Forde (12:11):

All the benefits. And now for more benefits, developer tools. AWS CodeBuild customers can now use AWS Lambda to build and test their software packages. AWS Fault Injection Simulator announces scenarios and scheduled experiments. You're probably wondering, what are these scenarios? Well, scenarios define events or conditions that customers can apply to test the resilience of their applications, such as the interruption of compute resources on which the application is running in. Scenarios are created and owned by AWS, providing you with pre-defined targets and fault actions. Less work that you have to do, I love it. AWS Resilience Hub announces enhancements to its resilience score capability, which now provides more actionable recommendations customers can implement to improve their application resilience and increase their score. Combined with the exclude recommendation capability, customers can emphasize different parts of the resiliency score that are most relevant to their application.

Simon Elisha (13:15):

Quick update in the world of end user computing, NICE DCV releases version 23.1 with support for web orphan redirection. So this is really useful, and a bunch of other features as well such as a new Windows display driver which optimizes graphics pipeline, and reduces overall CPU usage. If you're not familiar with this, NICE DCV is a high-performance remote display protocol that's designed to help customers securely access remote desktop or application sessions, including 3D graphics applications hosted on servers with high-performance GPUs.

Jillian Forde (13:46):

I was wondering what this Nice is, I just like that it's called NICE.

Simon Elisha (13:54):

It's nice, it's so nice.

Jillian Forde (13:54):

You're so nice. Now for front end web and mobile, we've got one quick update. AWS Device Farm announces upgraded test environments for Android devices.

Simon Elisha (14:04):

And a quick update on the Internet of Things, there is now extended protocol support for AWS IoT SiteWise Edge through the EasyEdge, so this provides data ingestion from 10 additional industrial protocols, including Modbus EtherNet/IP, Siemens S7, KNX LoRaWAN and MQTT PROFINET, et cetera, et cetera.

Jillian Forde (14:25):

I don't even know what that was that you just said at the end.

Simon Elisha (14:25):

Oh, the IT world and the industrial is a whole other world.

Jillian Forde (14:31):

It really is. And speaking of whole other worlds, machine learning, for those who want to sprinkle a little AI in your application. AWS Audit Manager customers can now access a pre-built standard framework to

help gain visibility into how their generative AI implementation on Amazon Bedrock is working against AWS-recommended best practices, this is so cool. With this framework, Amazon Bedrock customers can now start auditing their generative AI usage and automating evidence collection, providing a consistent approach for tracking AI model usage and permissions, flagging sensitive data, and alerting on issues. Amazon Bedrock customers can access a generative AI best practices framework in the AWS Audit Manager framework library, where they can identify and select from 110 new controls focused on generative AI best practices in domains such as governance and oversight, data security and privacy, incident management and business continuity planning.

(15:40):

AWS experts in AI, compliance and security assurance developed this framework with additional partners reviewed by the global audit and assurance firm, Deloitte, as an AWS partner. That sounds super handy, definitely check that out if you're using Bedrock. Next is Amazon Comprehend launches new trust and safety features, Amazon Lex introduces selective conversation log capture for Amazon Lex version 2. And last, Amazon Sagemaker Canvas now supports advanced model build configurations, such as selecting the training method and algorithms, customizing the training or validation data split ratio, and setting limits on AutoML iterations and job runtime. This allows customers to customize the model-building configurations without having to write a single line of code.

Simon Elisha (16:30):

Nice.

Jillian Forde (16:31):

This flexibility can provide... I know, right? My favorite is the code I don't write.

Simon Elisha (16:35):

No code.

Jillian Forde (16:35):

Yes. This flexibility can provide more robust and insightful model development. Non-technical stakeholders can use the no code features with default settings, while data scientists can experiment with various machine learning algorithms and techniques, helping them understand which methods work best for their data.

Simon Elisha (16:55):

Now let's talk about management and governance, and AWS Health has introduced new features to help you manage planned lifecycle events. So things like Amazon EKS Kubernetes version end of standard support, or Amazon RDS certificate rotations, and end-of-support for other open source software. AWS Health is the authoritative source for information about service events and scheduled changes affecting your AWS cloud resources. These new features provide timely visibility into upcoming planned lifecycle events, a standardized data format that lets you repair and take actions, as well as the ability to dynamically track the completion of required actions at the resource level. AWS Health also provides organization-wide visibility into planned lifecycle events for teams that manage workloads across the company.

(17:42):

So really, you want to make sure that you're looking at AWS Health when it comes to your own environment. AWS Control Tower now supports tagging for controls enabled in AWS Control Tower, so you can tag the things as you add, remove or list your alt tags in your environment, always tag your stuff. You can now deploy your StackSets faster with concurrency mode in AWS CloudFormation, now this one got my attention. So this is a new capability which is called concurrency mode, that lets you be fast whilst balancing deployment safety. Customers control speed and safety of their stack deployments with concurrency and failure tolerance, and customers define these parameters along with concurrency mode as part of their StackSets operation preferences prior to running a StackSets operation, I should say.

(18:29):

StackSets operation preferences is available for create stack instances, delete, the tech stack drift, import stacks, update stacks, update stack instances, and you can now change the behavior of these types of things, so you can choose strict failure tolerance or soft failure tolerance with the parameter. Strict failure tolerance is the current default behavior, and in strict failure tolerance StackSets sets the actual concurrency of the operations based on the failure tolerance threshold. Customers choose strict failure tolerances when they want a hard stop on the StackSets operation. Soft failure tolerance is a new deployment behaviors for customers who want the benefits of concurrent deployments without having that strict restriction.

(19:09):

In soft failure tolerance, StackSets sets customer-defined concurrency value as the actual concurrency of the operation. So, StackSets will continue to stop the operation when StackSets crosses the failure tolerance threshold. Sounds confusing, but it actually frees you up to move in a slightly different way when you're deploying and get things going quicker. AWS Config has launched inventory and compliance dashboards, and AWS service management Connector has introduced AWS Health and AWS Systems Manager Op Center integration in Jira Cloud.

Jillian Forde (19:40):

Now, migration and transfer. AWS Application Discovery Service adds support for communication through network proxy, and AWS Application Migration Service now provides an action for replatforming workloads using AWS App2Container. You can now use Application Migration Service to replatform applications running on migrated servers into Amazon ECS, Amazon EKS, and AWS App Runner. The new feature provides support for replatforming .NET and Java-based applications, including containerizing a .NET framework-based application to Windows containers, and replatforming JBoss, Apache Tomcat, Java, and .NET core applications to Linux containers.

Simon Elisha (20:25):

Wow, JBoss, now there's a name I haven't heard for a long time.

Jillian Forde (20:29):

Me neither, yeah.

Simon Elisha (20:31):

Probably need some replatforming right there. Let's talk networking and content delivery. Amazon CloudFront has announced the unified security dashboard, with the new security dashboard you can now enable, monitor and manage common security protections for your web applications directly from

the Amazon CloudFront console. Built for customers that need unified management of their application delivery and security, the interactive security dashboard brings AWS WAF visibility and controls directly to your CloudFront distribution, including visibility into applications' top security trends, allowed and blocked traffic, and bot activity. And investigative tools like a visual log analyzer and built-in blocking controls make it easier for you to isolate traffic patterns and block traffic without querying logs or writing security rules.

Jillian Forde (21:16):

Up next, security, identity and compliance. One quick update, so Amazon GuardDuty has incorporated new machine learning techniques to more accurately detect anomalous activities indicative of threats to your Amazon EKS clusters. This new capability continuously models Kubernetes audit log events from Amazon EKS to detect highly suspicious activity, such as unusual user access to Kubernetes secrets that can be used to escalate privileges, and suspicious container deployments with images not commonly used in the cluster or account.

Simon Elisha (21:51):

That actually sounds really cool, I like that one. Moving onto the topic of storage, Amazon Elastic Block Store now supports block public access for EBS snapshots, so this is an account-wide security setting that allows customers to block public sharing of EBS snapshots in an AWS region, so you can stop it all happening if you don't want it to happen. And finally, Amazon Data Lifecycle Manager has added support for pre-script and post-script automation of EBS snapshots. So lots of updates there, but my goodness me Jillian, we have updates coming up big time in a few weeks for our Reinvent Series, don't we?

Jillian Forde (22:25):

That's right, the big ones.

Simon Elisha (22:27):

People need to get ready for that one.

Jillian Forde (22:29):

They'll need to listen to our episode about Reinvent.

Simon Elisha (22:32):

Exactly, we'll be talking about how best to tackle Reinvent whether you're there or not there. And also of course, one of the best ways to keep up to date is we will be recording nearly every day during the week with a summary of everything that's been going on, and the good stuff, and you'll be hearing from both Jillian and myself, so it should be fun. Jillian, how do folks reach out to you?

Jillian Forde (22:52):

Ms. Jill Ford on... Oh, I can't say Twitter anymore, X. It's been like weeks now.

Simon Elisha (22:58):

I know, we need to move on.

Jillian Forde (22:59):

I'm so used to saying Twitter on this show, yeah.

Simon Elisha (23:00):

The artists formerly known as. And if you're old school like me, AWS Podcast at Amazon.com is the place to do it. And of course, until next time, keep on building.