

Speaker 1 ([00:00](#)):

This is episode 633 of the AWS Podcast, released on October 30th, 2023.

Simon Elisha ([00:08](#)):

Hello everyone, and welcome back to the AWS Podcast. Simon Elisha here with you. Great to have you back. Of course, I'm joined by my cohost, Jillian Forde. G'day, Jillian. How are you doing?

Jillian Forde ([00:16](#)):

So awesome to be here, Simon.

Simon Elisha ([00:19](#)):

It is fun, and we get to talk about cool stuff, and we're going to start with some cool stuff in the topic of analytics. Amazon DataZone is now generally available. Now, this data management service is designed to catalog, discover, analyze, share, and govern data at scale across organizational boundaries with governance and access control. It provides data visibility and helps data producers and consumers across business units more securely share data.

([00:46](#)):

Now, data producers can publish, manage, and govern their data with enriched business context in the Amazon DataZone business data catalog. This means that data consumers can more easily find and understand the data. Data producers get business-friendly asset naming suggestions from an artificial intelligence powered assistant which accelerates the process to enrich that metadata. Data consumers then search and subscribe to those data assets using Amazon DataZone data projects. This is used to create business use case based groupings of teams, analytics tools, and data assets.

([01:20](#)):

Now, consumers can instantly collaborate, share, and work with their peers using data projects and environments to seamlessly invoke tools like Amazon Redshift query editor, Amazon Athena query editor, or the tool of your choice programmatically. The integrated publishing and subscription workflow gives you systematic governance for all the metadata processes, interior AI usage with governed data sets using the same consistent project framework.

([01:45](#)):

Amazon EMR on EKS interactive endpoints is now generally available. Now, what this means is that Amazon EMR or EKS customers can run an interactive workload using an integrated development environment like EML Studio. For customers that require control on their execution environment, you can use self-hosted Jupyter Notebooks as another mechanism to run your interactive workloads via those interactive endpoints. Now, these enhance the flexibility when running interactive with EMR on EKS, and customers will be able to build a custom image with the application dependencies and use that image in running their interactive jobs using those endpoints.

([02:22](#)):

If you have updates for Amazon QuickSight, it now supports running queries as an IAM role for RedShift data sources. It also announces predictive analytics using Amazon SageMaker Canvas. AWS Glue now supports GitLab and Bitbucket in its kit integration feature, so you can integrate all the things.

([02:40](#)):

Something new for OpenSearch. OpenSearch 2.9 has a new feature called Search Pipelines. Now, this makes it easier to build query end result processing pipelines, which means you can build them as a composition of modular processing steps without complicating your own application software, which is pretty nice. Amazon OpenSearch service has also announced the Amazon OpenSearch service integrations, so you can take advantage of new schemer standards like OpenTelemetry and build dashboards based on an agreed on schemer between your ingestion pipeline and the OpenSearch service.

[\(03:13\)](#):

The OpenSearch service now supports OpenSearch version 2.9, which has several improvements to search observability, security analytics, and ML capabilities. That version itself, 2.9, has expanded geospatial aggregation support as well, so you can do more statistical analysis on your data, make it easier to draw conclusions, and interpret that information.

[\(03:36\)](#):

Finally, the Amazon OpenSearch team has been very busy. That service now supports alerting and anomalies on dashboards, so you can now create new or associate existing learning monitors and anomaly detectors from the dashboard launcher.

Jillian Forde [\(03:50\)](#):

Now we've got application integration. AWS Step Functions announces an optimized integration for Amazon EMR Serverless, adding support for the Run a Job integration pattern with six EMR Serverless API actions. EMR Serverless is a serverless option in Amazon EMR that makes it easy for data analysts and engineers to run OpenSource big data analytics like Apache Spark and Apache Hive without configuring, managing, and scaling clusters or servers. AWS Step Functions is a visual workflow service that makes it easy to compose AWS services into scalable, reliable, and resilient application components.

[\(04:35\)](#):

All right. The next one is definitely handy. Amazon EventBridge rules now support wildcard filters which enable you to match any character or sequence of characters within a string in your even payload. For example, you can use wildcards to match against values to end a specific file type in a directory or contain a specific word. Support for wildcards allow you to more precisely specific the types of events you want to consume from an EventBridge event bus, opening new use cases, and helping to optimize your event consumers.

[\(05:12\)](#):

Amazon EventBridge Event Bus is a serverless event router that allows you to create highly scalable event-driven applications by rounding events between your own applications, third party SaaS applications, and other AWS services. You can set up routing rules to determine where to send your events, allowing for application architectures that react to changes in your systems as they occur. Event Buses make it easier to build event-driven applications by facilitating event ingestion, delivery, security, authorization, and error handling.

[\(05:50\)](#):

AWS Health now aggregates health events in your organization in Amazon EventBridge. Amazon SQS binding extension support in the Core Windows Communications Foundation, also known as Core WCF's open-source project, is now generally available. Core WCF binding decide how the service endpoint should communicate with clients. As part of this open-source contribution, we are publishing two new

gate packages. One for server side, and one for client side. With Amazon SQS transport bindings, customers can send SOAP messages to Amazon SQS via the WCF client, and run core WCF services to receive and process those messages without changing any contract or service implementations.

[\(06:45\)](#):

This transport provides the extensibility to attach callback methods that can trigger additional actions, such as Amazon SMS notifications and AWS Lambda invocations once a message is processed. Last, Amazon SMS now connects AWS CloudTrail data event logging for the publish and publish batch API actions.

Simon Elisha [\(07:07\)](#):

Nice. Lots of updates there, and more updates in the topic of compute. It's a big list this week. We're happy to announce AWS Lambda's support for IPv6 for outbound connections in VPC. We now are providing the capability for Lambda functions to access resources in Dual-stack of VPC, outbound connections over IPv6 at no additional cost.

[\(07:30\)](#):

With this launch, along with Lambda support for public IPv6 endpoint inbound connections in 2021, Lambda enables you to scale your application without being constrained by the limited number of IPv4 addresses in your VPC, and you don't need to have translation mechanisms. Now, previously Lambda functions configured within IPv4 only, or Dual-stack VPC, could access VPC resources only over IPv4, but that made it tricky. We don't want it to be tricky, so now you can do it over IPv6.

[\(07:59\)](#):

Lambda test events are now available in AWS SAM CLI. If you use SAM CLI to author your serverless application, and it does make it pretty easy, test events are adjacent objects that mock the structure of request emitted by AWS services to invoke a Lambda function and return an execution event, which means you can validate a successful operation or identify errors. Previously, Lambda test events were only available in the console. Now it's in SAM CLI, so test all the things.

[\(08:26\)](#):

Amazon EC2 now supports setting AMIs to a disabled state, so now you can disable your unused or obsolete Amazon Machine Image, or A-M-I, or AMI. That's always the big discussion. What do you call it, Jillian? AMI or A-M-I?

Jillian Forde [\(08:40\)](#):

A-M-I.

Simon Elisha [\(08:40\)](#):

Yeah. Me too.

Jillian Forde [\(08:42\)](#):

I think we need the listeners to come and just send us one of those of how they pronounce it.

Simon Elisha [\(08:46\)](#):

It's right up there with Tabs and Spaces, let me tell you.

Jillian Forde [\(08:48\)](#):

Yep.

Simon Elisha ([08:49](#)):

Now, disabling an AMI changes a state to disabled which makes the AMI private if it was previously shared, and it prevents any new EC2 instances launching from that disabled AMI. This really helps your workflow because you can still keep it if you need to, just for regulatory reasons, but you wouldn't use it.

([09:07](#)):

So two new instance types are now available. We're happy to announce the general availability of Amazon EC2 R7a bare metal instances. These are powered by fourth generation AMD EPYC processors. These were the Genoa codenamed processors with a maximum frequency of 3.7 gigahertz, and they give you up to 50% higher performance compared to our R6a instances. Now, these are SAP certified, and they offer new processor capabilities like AVX-512, VNNI, and BFlight16. They have DDI5 memory, so you get much faster performance. In fact, 2.25 times more memory bandwidth compared to R6a instances, so these are good for latency sensitive workloads.

([09:50](#)):

Another new kind is the EC2 C7a instances. These are also powered by those Genoa CPU family with 3.7 gigahertz. Again, 50% higher performance than the previous generation. Similar improvements in terms of the processor capabilities and the use of DDR as well. Lots of different sizes. These also allow you to attach up to 128 EBS volumes to an AC2 instance. The C6as, rather than this C7a, could only do 28. Now you can do 128, so if you want to attach lots of volumes that's a good option.

([10:24](#)):

Amazon Linux has now announced support for Ansible and Corretto 21 in Amazon Linux 2023.2. Amazon EKS extended support for Kubernetes version is now available in preview. Now, I know this is really interesting for a lot of folks. Starting with Kubernetes version 1.23, you can now use a Kubernetes version for up to 26 months from the time the version is generally available from Amazon EKS. Now, Amazon EKS follows the Kubernetes project release cycle. For the first 14 months after a minor version is released as generally available, it's covered under standard support.

([11:00](#)):

Now, Amazon EKS extended support gives an additional 12 months of support for each Kubernetes minor version, so you get 26 months of support. During this time, Amazon EKS will continue to provide security patches for the Kubernetes control plan, and also release critical patches for the Amazon VPC CNI, Kube-Proxy, and core DNS add ons out of its published EKS-optimized APIs for Amazon Linux, and Bottlerocket, and EKS Fargate notes.

([11:28](#)):

Amazon ECR Public has introduced new navigation and search public in the ECR Public gallery, so new filters to help you find things from well known publishers like Docker and Amazon, and some new landing page highlights to help you find what's frequently used. Amazon EC2 Hibernate now supports more operating systems, so you can now hibernate the Microsoft Windows Server 2022, Red Hat Enterprise Linux 9, and Amazon Linux 2023 operating systems.

[NEW\_PARAGRAPH]Now, as a reminder, hibernation is an Amazon EC2 feature. It helps lower costs and achieve faster startup times by enabling customers to pause and resume their running instances at scale. Customers are not billed for compute time whilst instances are suspended, and your applications will resume right from where they left off. There are all the cool updates for compute.

Jillian Forde ([12:17](#)):

Now we've got four quick updates on customer engagement, all from the Amazon Connect team. Amazon Connect now supports multiple out of the box chat widgets per instance. There's a new contact and agent metrics in the API, and the API's been enhanced for accessing historical agent and contact metrics. The last is Amazon Connect Contact Lens now provides manager alerts on realtime metrics.

([12:45](#)):

Up next is database. We've got a lot of these Well-Architected related updates, and here's another one. PG Active active-active replication extension for Postgres is available.

Simon Elisha ([12:58](#)):

That's a lot of active.

Jillian Forde ([12:59](#)):

It really is. Yeah.

Simon Elisha ([12:59](#)):

That's a whole lot of active.

Jillian Forde ([13:04](#)):

Is available for MM RDS4 Postgres. PG Active just lets you use a synchronous active-active replication for streaming data between database instances to provide additional resiliency and flexibility in moving data between database instances, including writers located in different AWS regions for the purposes of maintaining availability for operations like switching right traffic to a different instance. Customers use a synchronous active-active replication to allow applications designed for eventual consistency to write to two or more database instances to help maximize available and reduce right latency.

([13:47](#)):

To support active-active replication, PG Active includes features such as node synchronization, automatic conflict detection and resolution, replication lag metrics, and can continuously copy data in up to 16 different writer instances. Wow. That's a lot. PG Active provides several automatic conflict resolution strategies in case two writers simultaneously modify the same database record, such as last update wins, first update wins, or customizable rules. Amazon RDS4 Postgres now supports PG Vector version 0.5.0 with HNSW indexing. Last is Amazon RS4 Postgres, MySQL, and MariaDB supports the M6IN, M6IDN, R6IN, and R6IDN instances.

Simon Elisha ([14:45](#)):

So let's talk about developer tools. AWS CodeDeploy now supports multiple load balances for Amazon EC2 applications. Before, you could only deploy one application per classic load balance at ALB or NLB. CodeDeploy now supports up to 10 load balancer target groups per application on EC2. AWS has also launched more encryption options for AWS app config feature flags, so now you can use the AWS key management services customer manager keys, and you can use keys that you yourself manage to encrypt and decrypt configuration data hosted by AWS app config.

[NEW\_PARAGRAPH]AWS App Runner has added support for monorepos and has also launched improvements for using custom domains. There are some new features for CDK8s, which includes synthesize and import helm charts and resolve cloud tokens. Now, this makes life a bit easier because

helm stands out as a widely embraced solution for the deployment and management of Kubernetes applications by converging CDK8s and helm, users can enjoy a unified workflow for creating and deploying Kubernetes manifests. With the recent addition of the CDK8s synth command, you can transform a CDK8s app directly into a helm chart ready to be integrated with helm deployments.

Jillian Forde ([16:05](#)):

Next up is end-user computing, and two quick updates. The first is AWS WorkSpaces now provides bundle selection guidance in the AWS console. The second one is Amazon WorkSpaces services are pleased to announce the availability of new applications, including Microsoft Office LTSC Professional Plus 2021, Microsoft Visio LTSC Professional 2021, Microsoft Project Professional 2021, Microsoft Office LTSC Standard 2021, Microsoft Office Visio LTSC Standard 2021, and Microsoft Project Standard 2021 on Workspace instances.

Simon Elisha ([16:48](#)):

It's all the Offices.

Jillian Forde ([16:51](#)):

Yes, that's right.

Simon Elisha ([16:53](#)):

Let's talk quickly about front-end web and mobile. AWS amplifies GraphQL API capabilities are now available as an AWS CDK construct, so you can build things beautifully, easily, and automatedness. Amazon Location Services added bounding box search for device position, so this allows developers to retrieve the latest position of their devices within a specific polygon. With bounding box search, developers can efficiently find devices within an area of interest to support use cases like searching for tracked assets near a location, showing device positions on a map at different zoom levels, and personalizing communication to customers within specific geographic areas.

Jillian Forde ([17:35](#)):

That's super cool. I'd have no idea how to actually do that without an API like Amazon Location Services.

Simon Elisha ([17:40](#)):

It would be hard.

Jillian Forde ([17:41](#)):

Yeah, it would be very hard. Up next, we've got internet of things. AWS IOT Device Management, this is a fully managed cloud service that helps you register, organize, monitor, and remotely manage IOT devices at scale. Now, we've announced the general availability of the software package catalog feature. With this new feature, you can now track and monitor software package version across Fleet, gain valuable insights from Essential Dashboard, and perform targeted updates on devices running on a specific software version.

([18:15](#)):

Using the software package catalog, you can register and maintain an inventory of up to 10,000 software package versions and related metadata in a centralized place. The feature works closely with AWS IOT device management Fleet indexing, making it easier for you to index, search, and group devices

based on software package versions. You can gain an overview of package version distribution and aggregated metrics across device Fleet. Moreover, you can use AWS IOT named Shadow to record and maintain package version information that's installed on each device. This information can derive necessary metrics to support granular analysis. Similarly, the integration with AWS IOT device management jobs gives you the flexibility to perform targeted actions based on the software versions.

Simon Elisha ([19:10](#)):

Let's talk about machine learning. Amazon Titan Embeddings is now generally available. Now, Amazon Titan Embeddings is a text embeddings model that converts natural language text, including single words, phrases, or even large documents into numerical representations that can be used to power use cases like search, personalization, and clustering based on semantic similarity.

([19:33](#)):

Now, this is optimized for text retrieval to enable retrieval augmented generation, or RAG, use cases. Amazon Titan Embeddings enables you to first convert your text data into numerical representations or vectors, and then use those vectors to accurately search for relevant passages from a vectored database, which allows you to make the most of your proprietary data in combination with other foundation models.

([19:56](#)):

Now, Titan Embeddings supports more than 25 languages, including English, Chinese, and Spanish. You can input up to 8,192 tokens, which means it's well suited to work with single word phrases or entire documents. The model returns output vectors of 1,536 dimensions, giving it a high degree of accuracy, whilst also optimizing for low latency and cost-effective results. Because Titan Embeddings is available via Amazon Bedrocks Serverless experience, you can easily access it using a single API without managing any infrastructure.

([20:28](#)):

You can now deploy ML models built in SageMaker Canvas to SageMaker realtime endpoints. Until now, SageMaker Canvas provided the ability to evaluate an ML model, generate fault predictions, and run what-if analyses within its interactive workspace, but now it can also deploy the models to SageMaker endpoints for realtime inferencing. It's just quicker to get things done.

([20:50](#)):

Some other updates for Amazon SageMaker. Amazon SageMaker Canvas has expanded content summarization and information extraction capabilities, so you can do this without having to write any code. Amazon SageMaker Canvas has also expanded its support for ready to use models to include foundation models, so now you can access and evaluate foundation models to generate and summarize content using those models, in addition to the ready to use models that already exist.

([21:16](#)):

Amazon SageMaker Feature Store now supports a fully managed in memory online store, which enables you to retrieve features for models serving in realtime throughout your ML applications. The new online store is powered by ElastiCache for Redis, which is a blazing fast in memory data store built on open source Redis.

([21:36](#)):

Amazon SageMaker Modeler Registry also announces support for private modeler repositories, so you can give them to yourself. You can now enhance the accuracy of pre-trained recognition moderation models with custom moderation. This makes it easier to find and remove content that is inappropriate,

unwanted, or offensive at scale. Customers across industries, such as social media, gaming, and advertising, often use recognition's content moderation capabilities to protect their brand reputation and enable safe user communities. With custom moderation, customers can now enhance the accuracy of the moderation deep learning model on their business specific model by training an adapter with as few as 20 annotated images in less than an hour.

[\(22:20\)](#):

Amazon Textract has launched custom queries to improve information extraction for business specific documents. Now, Amazon Textract is a machine learning service that automatically extracts printed text, handwriting, and data from any document or image. This new feature lets customers adapt the queries feature and improve extraction accuracy for their particular documents, so it allows you to extra specific pieces of information using natural language questions, which makes life a bit easier. Finally, Amazon Kendra has released the Drupal Connector to enabled Drupal document search from that service.

Jillian Forde ([22:56](#)):

Now onto management and governance. AWS customers now have the option to enabled systems manager and configure permissions for EC2 instances in an organization that has configured to use AWS organizations with a single action using default host management configuration. This feature provides a method to help customers ensure core systems manager capabilities, such as patch manager, session manager, and inventory, are available for all new and existing instances.

[\(23:31\)](#):

Default host management configuration is recommended for all EC2 customers. It offers a simple scalable process to standardize the availability of system manager tools. AWS systems manager, Fleet manager, RDP, now supports up to 10 ADP resolution. AWS systems application manager now supports SAP HANA. We're also introducing a new CloudWatch metric called attach EBS status check to monitor if one or more Amazon EBS volumes attached to your EC2 instances are reachable and able to complete IO operations. With this new metric, you can now quickly detect and respond to any EBS impairments that may potentially be impacting the performance of your applications running on Amazon EC2 instances.

[\(24:26\)](#):

You can now launch Amazon CloudWatch Internet Monitor directly from the VPC console. Internet Monitor provides visibility into how internet issues impact the performance and availability of your application hosted on AWS. To use internet monitor, you create a monitor and associate with one or more resources, which are VPCs, network load balancers, Amazon CloudFront distributions, or Amazon WorkSpaces directories.

[\(24:55\)](#):

AWS Control Tower releases 22 proactive controls and 12 AWS security hub detective controls. AWS is pleased to announce an update to the AWS Well-Architected framework. If you haven't heard of it, you should definitely check it out. Which will provide customers and partners with more prescriptive guidance on building and operating in the cloud, and enable them to stay up to date on the latest architectural best practices, and a constantly evolving technological landscape.

[\(25:26\)](#):

The enhanced implementation guidance for the new and updated best practices in this release are more prescriptive, including enhanced recommendations and steps on reusable architecture patterns, targeting specific business outcomes in the AWS cloud. Last is, now in the AWS Well-Architected Tool,

you can effortlessly craft review templates to answer questions, update notes, and even incorporate custom lenses across your workloads.

Simon Elisha ([25:58](#)):

Onto the topic of media services. Amazon IVS introduces in-console broadcasting for low latency streaming. Once you create at least one channel, you can initiate your first stream from the console itself. Previously, you could use broadcast SDKs or third party streaming software. Now you can do it yourself inside the console. I think I'm going to give that a go. AWS Elemental MediaTailor now supports ad break discovery in channel assembly, and AWS Elemental MediaConvert now supports video source replacement.

Jillian Forde ([26:31](#)):

Let's talk about networking and content delivery. AWS Network Load Balancer now supports availability zonal DNS affinity, disable connection termination for unhealthy targets, and UDP connection terminations by default. I know that's a lot, so definitely want to break that down.

Simon Elisha ([26:49](#)):

Dive into that one, yeah.

Jillian Forde ([26:50](#)):

So let's start with the availability zone.

Simon Elisha ([26:53](#)):

There's a lot there.

Jillian Forde ([26:53](#)):

Yep, so the availability zonal DNS affinity. This capability resolves DNS such that clients resolving the network load balancer DNS receive the load balancer IP addresses in the same availability zone that they're in. This enables customers to build zonally independent application stacks and improve application performance by reducing latency. You'll still have the ability to build your application across multiple availability zones for redundancy.

([27:28](#)):

The next is the disabled connection termination for unhealthy targets. With this feature, you can choose to maintain or terminate active connections to targets that fail health checks. By disabling the default network load balancer behavior, you can prevent client reconnect storms in the event of an outage.

([27:50](#)):

The last is the UDP connection termination by default. Network load balancers now terminate UDP connections at the end of the deregistration timeout by default for newly created UDP target groups. With this new feature, you can gracefully terminate long-lived UDP connections, improving the performance of your application.

([28:15](#)):

The last update is the allocation load balancer and network load balancer now support registering instances as targets when addressed by IPV6 to simplify load balancer configuration. This launch enables you to directly register instances using IPV6 addresses as targets, removing the need to track individual

IP addresses. The feature also enables the use of EC2 auto-scaling groups with your target groups, which provides automatic registration of targets based on the scaling needs of your application.

Simon Elisha ([28:51](#)):

Let's talk about security, identity, and compliance. AWS Verified Access, which is a service that helps you provide secure VPN-less access to your corporate applications, now supports custom managed keys, or CMKs for KMS, for the encryption of data at rest, make it easier to meet the compliance and regulatory requirements of your organizations. By default, AWS Verified Access has also provided encryption for all data, including trust provider information, group policy, and endpoint policy using AWS owned KMS keys when stored at rest.

([29:25](#)):

Now, you can also have the option to use custom managed keys to encrypt data, including your trust provider information, group policy, and endpoint policy. You can begin using custom managed keys with just a single click in the console, or through the verified access APIs.

([29:40](#)):

Amazon Macie has introduced new managed data identifiers to expand its capabilities for discovering and identifying Stripe API keys, Google Cloud API keys, driver's license numbers for India, and national identification numbers in India as well, in this simple storage service. Understanding the presence and location of such data in your S3 storage helps you better plan data security, governance, and privacy of your organization. With over 100 managed data identifiers, Macie helps protect your sensitive data at scale.

Jillian Forde ([30:13](#)):

And the last one in today's episode is supply chain. AWS supply chain Demand Planning now supports product lineage. Creating an Account Demand plan requires Demand Planners to include sales history of prior models, or similar products as input to forecast. With this release, planners can establish links between products and their predecessor versions, or alternate products incorporating rules around the extent of history to be used for forecasting purposes, referred to as product lineage.

([30:47](#)):

AWS supply chain Demand Planning will leverage this information to generate a surrogate history for the product forming the basis for more accurate demand predictions. This ensures that the forecast generated through product lineage is inherently more precise, thereby minimizing the need for manual adjustments to forecasts. I like the sound of that.

Simon Elisha ([31:10](#)):

Yep, sounds handy. Sounds handy. It's funny. Whenever we do these update episodes, there's always one or two things I go, "Yep, going to try that. Yep, going to have a go at that." There's always something useful.

Jillian Forde ([31:21](#)):

I agree, yeah. I'm loving all of the Well-Architected related types of updates today.

Simon Elisha ([31:27](#)):

If you're not using it as an organization, you should. It really is. It's a catalyst for talking to each other about how you build and run your solution software services, et cetera. It's amazing the things that come up during those conversations.

Jillian Forde ([31:44](#)):

It really is, and then you don't know what you don't know. That happens every single time.

Simon Elisha ([31:46](#)):

Yep. Yeah. There's often the, "I thought you were looking after that," and the, "Well, I thought you were looking after that."

Jillian Forde ([31:51](#)):

Yeah, so true.

Simon Elisha ([31:51](#)):

Awkward.

Jillian Forde ([31:51](#)):

It's so true.

Simon Elisha ([31:52](#)):

It's interesting. As we start to get to sort of the end of October into November, our thoughts of course turn to Reinvent.

Jillian Forde ([31:58](#)):

That's right.

Simon Elisha ([32:00](#)):

We're at a crossroad, aren't we, Jillian? We're covering that thing from top to bottom this year.

Jillian Forde ([32:05](#)):

That's going to be exciting. All the latest of what happens at Reinvent, we're going to cover all the exciting announcements.

Simon Elisha ([32:11](#)):

Yep, every day there'll be an episode to keep you up to that note. I know many folks listen to the podcast just for that to stay up to date for that big week of stuff, so we're going to be working hard together to get you all the content each and every day on the day it happens. That's our commitment to you.

Jillian Forde ([32:28](#)):

That's right. Simon, you're even going to be up at like 2:00 in the morning or something? I mean, that is serious commitment.

Simon Elisha ([32:34](#)):

Well, these are the sacrifices we make to deliver a quality product, but yes. I'll be delivering things as quickly as possible, and you too, even though you're in a friendlier time zone. We'll make it happen, because we care about our listeners. We want to make sure you have the-

Jillian Forde ([32:34](#)):

That's right.

Simon Elisha ([32:47](#)):

Latest and greatest in your hands.

Jillian Forde ([32:48](#)):

That is so true.

Simon Elisha ([32:49](#)):

It's very cool. Jillian, how do folks reach out to you?

Jillian Forde ([32:52](#)):

On Twitter @MissJillForde.

Simon Elisha ([32:54](#)):

Awesome. If you're old school like me, AWS Podcast at Amazon.com is the place to do it. Until next time, keep on building.