

Simon Elisha (00:00):

Hello everyone. Welcome back to the AWS podcast, Simon here with you. Great to have you back and I'm joined by my co-host Jillian Ford. Good day, Jillian, how are you doing?

Jillian Forde (00:07):

Fantastic. But Simon, there is a huge weight I feel like with Hawn not being here. Such a responsibility.

Simon Elisha (00:14):

There's something missing. And Hawn is mid-flight at the moment on her way to visit a customer and of course she'd planned it to land in plenty of time and be at her hotel and be able to participate. But of course airline schedules, what am I going to say?

Jillian Forde (00:31):

I know. But I feel she brings such a lightheartedness and fun and playfulness and I'm so serious.

Simon Elisha (00:37):

That's true.

Jillian Forde (00:38):

I feel like I got to step my game up.

Simon Elisha (00:40):

Yeah, totally, totally. You have to overcompensate now.

Jillian Forde (00:43):

That's right.

Simon Elisha (00:44):

Well let's get into it because we have quite a lot to cover today actually. It's probably a bigger episode than we've had the last couple of times. So without any further ado, let's of course start with the topic of analytics and AWS has open sourced the Amazon Redshift integration for Apache Spark to help Apache Spark developers seamlessly build and run Apache Spark applications on Amazon Redshift data. So with this release, Amazon Redshift open source is the Amazon Redshift contributions and integration for Apache Spark, which means you can now review the source code, extend it, contribute features, and or make modifications to meet your own needs, which is pretty cool.

[NEW_PARAGRAPH]Amazon EMR Serverless has added API support for the application UIs and Amazon EMR Serverless also now supports retrieving secrets from AWS Secrets Manager. And AWS Lake Formation has launched the READONLY administrator role. We're also happy to introduce two new capabilities for AWS Clean Rooms that gives customers flexibility to generate breach insights, custom analysis rule and analysis templates. These capabilities enables customers to bring their own custom SQL queries into an AWS Clean Room collaboration based on their specific use case. And with the custom analysis rule, customers can create their own queries using advanced SQL constructs as well as review queries prior to their collaboration partners running them. So this gives you more built-in control

about how data is used in collaborations upfront in addition to reviewing query logs after analysis are complete.

[\(02:19\)](#):

Using analysis templates, customers can create queries with parameters that provide reusability and flexibility to those running queries in a collaboration, so this helps you expand and automate the types of analysis you run frequently with multiple partners and minimizes the need to write a new SQL code when analyzing collective data sets. And it is yet another reminder of my constant refrain, learn SQL, it will last for you for your entire career.

Jillian Forde [\(02:45\)](#):

This is a really good reminder. And next we've got application integration. Amazon Event Bridge Scheduler now supports the ability to set schedules to automatically delete upon completion of the last invocation. This can be used for one time, cron and rate schedules with an end date. Event Bridge Scheduler allows you to schedule millions of tasks that can invoke more than 270 AWS services and over 6,000 API operations. Amazon SQS announces increased throughput quota for FIFO high throughput mode and Amazon SQS announces a preview of JSON protocol support, enabling lower latency and improved performance for SQS customers. Based on AWS performance tests for a five kilobyte message payload, JSON protocol for SQS reduces end-to-end message processing latency by up to 23% and reduces application client side CPU and memory usage. Amazon MQ now supports Rabbit MQ version 3.11 and new operator policies and Amazon MQ now supports customer managed configuration for Rabbit MQ brokers.

[\(04:06\)](#):

Next up is blockchain. So Amazon Managed Blockchain Access provides a new serverless offering for non-mining full blockchain nodes. Developers can instantly call standard remote procedure calls to interact with digital assets and distributed applications built on multiple blockchains starting with Bitcoin without specialized infrastructure. Amazon Managed Blockchain Access for Bitcoin compliments the existing dedicated offering for Ethereum. And Amazon Managed Blockchain Query delivers standardized and formatted blockchain data from multiple blockchains starting with Bitcoin and Ethereum. And these applications can consume this data quickly without customers having to perform complex data transformations or paying expensive storage costs. Now the next one is compute. We've got a lot of big ones here.

Simon Elisha [\(05:03\)](#):

Some [inaudible 00:05:04] big ones here.

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

It really is, yeah. Take your notes. AWS announces the general availability of Amazon EC2 M7i-flex and EC2 M7i instances powered by custom fourth gen Intel Zion scalable processors. Woo-hoo. These custom processors available only on AWS offer up to 15% better performance over comparable X86 based intel processors utilized by other cloud providers. M7i-flex instances are the easiest way for you to get price performance benefits for a majority of general purpose workloads. They deliver up to 19% better price performance compared to M6i instances. M7i-flex instances offer the most common sizes from large to eight x large and are a great first choice for applications that don't fully utilize all compute resources. With M7i-flex instances, you can seamlessly run web and application servers, virtual desktops, batch processing, microservices, enterprise applications, databases, gaming services, CPU

based machine learning video streaming and more. That's pretty much almost everything, right, that you can put on the internet.

Simon Elisha ([06:32](#)):

Yeah. The things you want to do.

Jillian Forde ([06:36](#)):

Yeah, M7i instances offer larger instance sizes up to 48 x large and the latest DDR five memory. M7i instances will soon offer multiple bare metal sizes for high transaction and latency sensitive workloads. M7i bare metal sizes also support built-in intel accelerators for data streaming accelerator, in memory analytics accelerator and quick asset technology. These accelerators are used to facilitate efficient offload and acceleration of data operations and optimize performance for workloads. Another big one we've got coming up is that the availability of Amazon EC2 C7gd instances and M7gd instances and R7gd instances with up to 3.8 terabytes of local NVMe based SSD block level storage.

Simon Elisha ([07:39](#)):

That's crazy.

Jillian Forde ([07:41](#)):

I know, there's so much.

Simon Elisha ([07:44](#)):

Numbers have been big.

Jillian Forde ([07:45](#)):

They have up to 45% improved real-time NVMe storage performance then comparable Graviton2 based instances. These instances are powered by AWS Graviton3 processors delivering up to 25% better performance over Graviton2 based instances. These Graviton based instances feature the latest DDR five memory, which provides 50% more memory bandwidth than DDR four. These instances are also built on the AWS nitro system and are a great fit for applications that need access to high speed, low latency local storage, including those that need temporary storage for data scratch base, temporary files and caches. C7gd instances are ideal for high performance computing, CPU based machine learning inference and ad serving. Now the M7gd instances, these are ideal for general purpose workloads such as application servers, microservices, and game services. And your R7gd instances are ideal for memory intensive applications such as open source databases and real-time big data analytics. But wait, there's even more.

([09:09](#)):

So with the launch of M7i instances, customers can attach up to 128 EBS volumes to an EC2 instance. By comparison, the previous generation M6i instances allow up to 28 EBS volume attachments to an EC2 instance. The higher volume attachment limit in enables customers to increase storage density per EC2 instance so that they can host larger applications per instance or reduce total compute cost by improving their resource utilization. With this launch, customers can also take crash consistent snapshots for up to 128 EBS volumes attached to a single EC2 incense.

Simon Elisha ([09:58](#)):

The funny thing with that one is that I remember when I was working with a customer years ago and 28 EBS volumes wasn't enough for the use case. And so, it's like people need this. So 128, that's pretty cool.

Jillian Forde ([10:13](#)):

It is really cool and we've studied even more on compute. So VT1 instances. These are the first EC2 instances that deliver dedicated hardware acceleration for video transcoding and are optimized for workloads such as live-streaming, video conferencing, video library optimization, and just-in-time asset transcoding. VT1 is powered by the AMD Xilinx Alveo U30 media accelerator and can support streams up to 4K ultra-high def resolution at 60 frames per second and can transcode up to 64 simultaneous 1080p 60 streams in real time. Two from AWS batch. So AWS Batch now supports price capacity optimized allocation strategy for spot instances and AWS Batch on Fargate now supports Linux arm 64 and Windows X86 containers in the console. AWS Fargate for Amazon EKS now lets customers configure the size of ephemeral storage for their workloads up to a maximum of 175 gigabytes. And Amazon EKS supports the Amazon EFS container storage interface CSI driver as an EKS add-on. Amazon EC2 adds templates for Amazon ECS to the Spot Blueprints portfolio.

([11:44](#)):

With Spot Blueprints, you can generate a reusable quick start template as CloudFormation or Terraform based on your workload and use the template to configure a compute environment using EC2 spot instances. This saves you time to use spot instances with services like ECS and start accessing the savings and scale that spot instances offer for interruptible workloads. AWS Lambda add support for Python 3.11 and the Amazon ECS console now has enhanced support for task definition workflows. And there are three new features for AWS SimSpace Weaver. So first is support for custom container images. Developers can now use their own AL2 based images stored in Amazon ECR and the second is an additional setting for their simulation tick rate. So developers can now enable SimSpace Weaver to advance the simulation as fast as possible by setting their simulation tick rate to unlimited.

([12:53](#)):

Using this setting can help reduce simulation time for scenarios that do not need real-time visualization or interaction. And the third is developers can now launch simulations with multiple spatial domains to separate different types of simulation logic and better allocate compute resources. Developers can use multiple domains for more advanced simulations that combine multiple sets of behaviors in a single environment. So let's say you have an application simulating pedestrians walking on a sidewalk and cars on a road. A simulation can now allocate a fixed number of compute resources for pedestrians and another fixed number of compute resources for cars.

Simon Elisha ([13:38](#)):

So now that Jillian is in a exhausted heap on the ground after all that upside.

Jillian Forde ([13:38](#)):

That's right, that's a lot.

Simon Elisha ([13:46](#)):

A lot of stuff. Let's move on to customer engagement and all of our updates for customer engagement relate to Amazon Connect. So Amazon Connect scheduling now offers automated flexible days scheduling. It also now supports flow only attributes, so it supports restricting the use and access of attributes to a single flows. This lets you have granular control when an attribute is associated with a

contact and is only accessible via a particular flow. And Amazon Connect now also supports custom flow block titles that make it easy to identify and distinguish blocks within your flow. So you could rename things like a play Prompt to a welcome message or get customer input to hotel booking Lex bot. It gives you lots of options and makes it easier to diagnose where errors may be occurring.

[\(14:32\)](#):

Connect now offers automatic activity scheduling based on shift duration and scheduling agent time off balance and group allowance support. It's also launched Flows UI toolbar and also the ability to add notes and you can also archive and delete flows from the UI. And Connect now supports customer profiles rule-based identity resolution. This allows you to match and merge similar profiles which help companies maintain accurate customer information and deliver more personalized customer interactions via agents and automated experiences. A lot of quality of life improvements there that if you're a user of Amazon Connect, you'll probably enjoy.

[\(15:12\)](#):

Now let's move on to databases and Amazon RDS has simplified the connectivity between an RDS or Aurora database and AWS Lambda. So you now have the option to automatically set up connectivity between RDS or Amazon Aurora databases and AWS Lambda function using Amazon RDS proxy. From the Amazon RDS console, you can choose an action to set up a connection between your Amazon RDS or Amazon Aurora database and a Lambda function using an existing RDS proxy or using an auto created new Amazon RDS proxy in your account, Amazon RDS automatically sets up your related network settings to enable a secure and scalable connection between the Lambda function and the Amazon RDS or Amazon Aurora database.

[\(15:57\)](#):

Now this automation minimizes the need for additional manual networking tasks like setting up VPC, security groups, subnet, ingress, egress rules to get a connection between your serverless application and a database. So this is really good for improving productivity for new users and application developers who can get up and running within minutes. I'm a big fan of this function because it's less work to get going. Amazon Aurora supports local right forwarding for Amazon Aurora MySQL. And speaking of Amazon Aurora MySQL, there is now version MySQL 3.04, which is compatible with MySQL 8.0.28 ga, and also MySQL 2.12, which is compatible with MySQL 5.7.40, which is also generally available. And speaking of Amazon Aurora and now supports automatic backup of deleted clusters. So both MySQL compatible and PostgreSQL compatible editions support that automated backup. So you can now choose to retain automatic backups when deleting an Aurora cluster and later use it to restore a new cluster for it.

[\(16:57\)](#):

So Aurora automatic backup takes a continuous backup of your database as you make changes allowing you to restore to any point in time in your configured backup retention period. Now, previously when you deleted your cluster, the automated backup was also deleted. Now you can decide to keep it for the same duration as the retention period for the actual cluster and you can restore the data from the automated backup to a new cluster at any point in time during the backup retention period. Backups retained for a day are free of charge. Beyond that, they billed in a per gig per month rate based on your region. Amazon RDS for SQL Server now supports secondary host metrics in the enhanced monitoring and RDS custom for SQL Server now supports local time zones. Amazon RDS for Oracle now supports read and mounted replicas for instances on the multi-tenant container database or CDB architecture running in a single-tenant configuration.

[\(17:54\)](#):

Amazon RDS for Oracle replicas fully manage the configuration of Oracle data guard to create and maintain replicas in the same or different AWS region as the primary DB instance. And Amazon ElastiCache now supports the M7g and R7g Graviton3 based nodes. These are really great because you can get up to 21% improved P99 latency compared to an R6g x four large. And Amazon RDS custom for Oracle now supports AWS CloudFormation templates and they're not done because Amazon RDS for SQL Server now supports a minor version 2019 CU21. This gives performances and security fixes my regular reminder to patch your stuff. And Amazon RDS now supports a progress indicator for improved visibility into the storage optimization process. With the progress indicator, you get better visibility into the progress of storage optimization processes such as scaling up of storage sizes and changing the storage volumes.

[\(18:52\)](#):

And finally, let's talk serverless databases and we're happy to announce the general availability of AWS database encryption SDK, which is an upgrade to the existing Amazon Dynamo DB encryption client, which enables you to include client side encryption in your Dynamo DB workloads. Now the AWS database encryption SDK allows you to more easily perform attribute level encryption, enabling you to encrypt specific attribute values before storing them in your Dynamo DB table. This lets you protect sensitive data in transit and at rest as data cannot be exposed unless decrypted by your application. It also lets you easily search on encrypted attributes without decrypting the entire database beforehand. This lets you find the right information quickly to download to your application whilst your data remains securely encrypted in the database.

[\(19:38\)](#):

Now the AWS database encryption SDK makes it easy to let your customers bring their own encryption key to your application, giving them direct ownership over their data by controlling the encryption key. Designed with multi-tenancy in mind, you can use different encryption key providers across a single database table to safely isolate data. In conjunction with the AWS Key Management Service, KMS, you can use KMS key policies to enforce clear separation between the authorized users who can access specific encrypted attributes and those who cannot. Now this SDK is compatible with the Amazon Dynamo DB database and is available in Java at the GitHub repository. Links in the show notes.

Jillian Forde ([20:18](#)):

Up next, developer tools. AWS application composer introduces three new updates to help developers build and share modern applications faster. With these capabilities, developers can more easily iterate on their application, export an image of the application and integrate application composer with their local development environment whenever they're ready. And we've announced support for deploying to Kubernetes clusters running on EKS for an Amazon Code Catalyst workflow. A workflow is an automated procedure that describes how to build, test and deploy your code as part of continuous integration and continuous delivery system. A workflow defines a series of steps or actions to take during a workflow run. This launch allows you to add a code deploy to Kubernetes cluster action to a workflow, providing a simple and repeatable process for deploying updates to Kubernetes clusters running on EKS.

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

Now onto end user computing. So AWS announces the availability of Microsoft 365 apps for enterprise on Amazon workspaces, cloud-based virtual desktop services. Your end users can now benefit from Microsoft 365 apps on workspaces services along with their other devices. To run Microsoft 365 apps for enterprise on Amazon workspaces services, you can bring your own Microsoft 365 licenses that meets Microsoft's licensing requirements and activate the applications at no additional costs. Eligible licensed plans include Microsoft 365 E3/E5, Microsoft 365, A3/A5, or Microsoft 365 business Premium.

Supported workspaces instances include those with Windows server and bring your own license Windows desktop operating systems running on dedicated instances. Eligible applications include Microsoft Word, Microsoft Excel, Microsoft PowerPoint, Microsoft Outlook, Microsoft OneDrive and more. Customers can use Microsoft 365 apps on workspaces services can also bring licenses to install and use additional Microsoft applications, including Microsoft Project, Microsoft Visio, and Microsoft Power Automate.

[\(22:52\)](#):

Microsoft 365 apps for enterprise can be installed and used on primary workspaces instances and failover workspaces instances using multi-region resiliency. AWS Wickr now offers data retention that extends to both internal and external communications in a Wickr network. A wicker network administrator can configure and apply data retention to all in network conversations. This includes shared media and files as well as conversations with guest users, external teams, and other partner networks to help meet internal legal and regulatory requirements.

Simon Elisha [\(23:35\)](#):

So let's talk a little bit about frontend web and mobile. And we're really excited to announce the launch of a new CloudWatch logger feature with AWS Amplify, which is available now for SWIFT and Android developers. This feature empowers developers to log errors from the Amplify Libraries to CloudWatch, which enhances your ability to detect production issues. It also enables developers to write custom logs to detect failures in different parts of their applications. So using this gives you far more visibility, understanding issues that affect your production apps. And let's face it, sometimes strange things happen in production and logs are the only way to figure it out.

Jillian Forde [\(24:12\)](#):

All right, game developers we're talking about game talk. So we are excited to announce that Amazon GameLift now supports both the dot net standard and dot net framework profiles in unity with the latest update to the Amazon GameLift server SDK. Amazon GameLift is a fully managed service that allows developers to quickly manage and scale dedicated game servers for multiplayer games. With this release, Amazon GameLift server SDK now supports unity 2022.3, unreal 4.26, unreal 5.1, Golang and custom C++ and C-Sharp engines.

Simon Elisha [\(24:58\)](#):

Game on ISO. Let's talk about some internet of things things. AWS IoT Core has announced the support of four new algorithms for certificate signing and key generation expanding the list of already supported asymmetric X509 client certificates and certificate signature schemes. AWS IoT Core is a managed service that allows customers to connect billions of IoT devices to AWS and uses X509 certificates as one of the means to authenticate client and device connections to the AWS cloud. The support for Revest-Shamir-Edelman's signature scheme with appendix based on the probabilistic signature scheme, or RSASSA-PSS signing and the P521 elliptic curve key algorithms provide developers a more flexible option to strengthen the security posture of their IT solutions and comply with organization specific cryptographic standards compliance requirements.

[\(25:54\)](#):

Now let's discuss some machine learning things and funnily enough, a fair bit's been happening in this space too. Amazon SageMaker training jobs now support the MLP five instances, which are powered by NVIDIA H100 chips, which are purpose built for high performance ML training applications in the cloud. You can use MLP five instances on SageMaker to train some of the most demanding models, which

includes large language models and diffusion models powering some of the most demanding generative AI applications. The MLP five instances currently feature up to eight of the NVIDIA H100 Tensor Core GPUs and P5 instances compliment the NVIDIA H100 Tensor Core GPUs with two times higher CPU performance, two times higher system memory, and four times higher local storage as compared to the previous generation.

Jillian Forde ([26:43](#)):

I know so many people are like, woo-hoo, I've been getting so many emails about, what are the P5s? I want to use them.

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

We want them, we want them. Give me the good stuff.

Jillian Forde ([26:52](#)):

Yep.

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

Amazon SageMaker has announced a new direct integration with Salesforce Data Cloud and SageMaker Studio announces prebuilt docker SageMaker distribution for machine learning. Also, the SDXL foundation model from stability AI is now available in Amazon SageMaker Jumpstart. Now you can deploy and use the SDXL 1.0 with just a few clicks in this studio or programmatically through the SageMaker Python SDK. Now SDXL 1.0 is the evolution of stable diffusion and the next frontier for generative AI for images. The model is capable of generating images with complex concepts like various art styles, including photorealism at quality levels that exceed the best image models available today.

([27:35](#)):

Like the original stable diffusion series SDXL 1.0 is highly customizable with simpler prompting requirements true to the intention of the prompts. It also exhibits the ability to generate difficult concepts such as hands and text or spatially arranged objects and personas. It's funny, in the world of drawing, hands are the hardest thing to make look natural. This is the reason why the Simpsons only have four fingers on their hands. There you go. The things you learn about as podcasts.

Jillian Forde ([28:01](#)):

No way.

Simon Elisha ([28:01](#)):

Yeah, it was just too hard to draw and make it look good.

Jillian Forde ([28:05](#)):

Well, my drawing skills stopped at kindergarten, so I can't really relate. But now with generative AI.

Simon Elisha ([28:10](#)):

Exactly, I welcome a new generative AI drawing overlords to help me draw when I can't.

Jillian Forde ([28:16](#)):

I need that.

Simon Elisha ([28:18](#)):

Amazon Translate enables language detection for real-time document translation and Amazon Polly has launched a new Belgian French female voice called Isabelle and the Claude two foundation model from philanthropic is also now available in Amazon Bedrock. So you can get the latest version of this. It can take up to 100,000 tokens in each prompt, it's a lot, which means it can work over hundreds of pages of text or even an entire book. It could also write longer documents in the order of a few thousand tokens compared to its prior version.

([28:50](#)):

Also, happy to announce AWS Health Scribe in preview. This is a new HIPAA eligible service designed to help healthcare software vendors build clinical applications that automatically generate preliminary clinical notes by analyzing patient clinician conversations. With AWS Health Scribe, you can integrate conversational and generative AI into your application to accelerate clinical documentation workflow and enhance the consultation experience with no machine learning experience required.

Jillian Forde ([29:19](#)):

That one's a really cool one for all you that are listening. Definitely check out a demo of it. It is really cool. I can't believe it. I mean this is like a heavy hitter episode and we are not even on.

Simon Elisha ([29:30](#)):

I know there's more to come.

Jillian Forde ([29:31](#)):

There's more to come. Now for all things serious management and governance. AWS CloudFormation announces looping capability with a four each intrinsic function. With the `forEach` function, you can replicate parts of your templates with minimal lines of code. You can use the `forEach` function to simplify your template layout and make it easier and faster for you and your peers to review your code. This helps reduce human errors such as updating wrong properties or missing out on updating multiple target properties in your template. With this launch, you can use the `forEach` function in your resources, resource properties, conditions, and output sections of your templates. For example, you can declare five instances of a Dynamo DB table with unique identifier names using the `forEach` function instead of replicating the same lines of code. Additionally, you can use nested `forEach` function loops and combine them with other intrinsic functions such as the `join` function to generate advanced templates.

([30:39](#)):

For example, you can combine resources such as VPCs, subnets and transit gateways as nested loops to create a networking stack. And AWS CloudFormation launches a new deletion policy called `RetainExceptOnCreate` that provides customers additional control over CloudFormation behavior when resources are deleted from stacks. Customers use deletion policies to preserve or backup resources during accidental stack deletions. Today customers can mark resources with a deletion policy of `retain`, `snapshot` or `delete`. Now customers can use `RetainExceptOnCreate` to protect in-use resources from accidental deletion without preserving resources that have not served traffic.

([31:29](#)):

This allows customers to reduce manual intervention during reach rise of stack Operations. AWS config now supports 19 new resource types. Starting today you can use Amazon CloudWatch alarms to trigger

an instant refresh rollback. Amazon CloudWatch Internet Monitor now allows you to customize thresholds for when to trigger a health event for the internet facing traffic between your end users and your applications hosted on AWS. AWS Resilience Hub add support for application resilience drift detection, allowing you to opt into notifications that your application is no longer meeting the recovery objectives set by your business. Resilience Hub provides a single place to define, validate, and track the resilience of your applications so that you can avoid unnecessary downtime caused by software infrastructure or operational disruptions. And now you can view your balance due information and the invoice amount in the AWS billing console.

[\(32:40\)](#):

Now we've got one update on media services, AWS Thinkbox deadline 10.3.0.9 is generally available with a deadline command application plugin, updated support for third party integrations and an upgrade to Python 3.10. And migration and transfer. AWS database migration service makes homogenous migrations simpler with built-in native database tooling. Today in addition to supporting MySQL and Postgres, this feature now supports MariaDB built in native database tooling with homogenous data migrations provides simple and performant like to like migrations with minimal downtime.

Simon Elisha [\(33:28\)](#):

Let's now talk about networking and content delivery. And there is now a new database public IPv4 address charge and some public IP insights which we'll get into. So we're introducing a new charge for public IPv4 addresses. So effective February the first 2024, there'll be a charge of half of 1 cent per IP per hour for all public IPv4 addresses, whether attached to a service or not. Now there's already a charge for public IPv4 addresses you allocate in your account that you don't attach to an EC2 instance. Now, as you probably know, IPv4 addresses are an increasingly scarce resource and the cost to acquire a single public IPv4 address has risen by almost 300% over the past five years. This change reflects our own costs and is also intended to encourage you to be a bit more frugal with your use of public IPv4 addresses and to think about accelerating your adoption of IPv6 as a modernization and conservation measure.

[\(34:29\)](#):

Now this change applies to all AWS services including EC2, RDS and EKS nodes and other services that have a public IPv4 address allocated and attached in all AWS regions. There is lots of detail in a great blog post by Jeff Barr that you can have a look at. And again, this doesn't come into effect till February the first of 2024, so you have plenty of time to organize yourself to optimize for that. And related to that, AWS has announced public IP Insights, which is a new feature of VPC IP address Manager, which makes it easier for you to monitor, analyze and audit your public IPv4 addresses so you can see everything you're using.

[\(35:09\)](#):

And Amazon VPC now supports primary IPv6 addresses on an elastic network interface, which enables you to make the first IPv6 address associated with your ENI immutable. And Amazon VPC IP address manager or IPAM now supports three new CloudWatch metrics, VPC IP usage, subnet IP usage, and public IPv4 pool usage, which allows you to identify underutilized or near full capacity IP address ranges, which means you can optimize your usage on AWS effectively. And Amazon Route 53 has added support for 14 additional top level domains. So previously there were 315 top level domains and now there's 14 new ones. You can get dot actor, dot AirForce, dot BET, dot Degree, dot Games, dot Gives, dot GMBH, dot Group, dot Press, dot Site, dot Space, dot Tech, dot Store and dot Wine. You can really get a lot more, a lot more. The old days when there were five domain names. There you go.

[\(36:15\)](#):

A quick update in the topic of security identity and compliance. Amazon Inspector now provides enhanced vulnerability intelligence as part of its findings, including lots more information like the name of malware kits used, the Mitre attack framework when the vulnerability was added to Caesar's KVEC list, the exploit prediction scoring system score, and evidence of public events associated with a vulnerability. So this gets you way more information when you're looking into an incident.

[\(36:45\)](#):

And finally, let's discuss storing stuff. Let's talk about storage. Amazon EFS now supports up to 55,000 IOPS per system for reads and up to 25,000 write IOPS, making it even easier to power your machine learning research, financial simulations, big data and other heavy duty stuff. So as a reminder, Amazon EFS is serverless, fully elastic file system that gets you up and running quickly. This increase is a 1.6 times increase on the read capacity and a 3.6 times increase on the write capacity. So that's a pretty significant increase. And finally, AWS backup for Amazon S3 now improves the speed of backups by up to 10 times for buckets with more than 300 million objects.

[\(37:31\)](#):

So this speed improvement enables you to do your initial S3 backup a lot quicker, particularly if you have lots and lots of objects. Again, one of those improvements that happens behind the scenes that the service teams are always working hard for our customers to make things better and you just get the benefit, pay more for it. You don't have to select it, you just get it, which is kind of nice. Jillian, there had to be something for everyone in this episode today.

Jillian Forde [\(37:55\)](#):

There really was the CloudFormation, those P5 instances that I keep hearing that everyone wants these days.

Simon Elisha [\(38:02\)](#):

There's just a bunch of different things. I know for a lot of customers, the Microsoft 365 apps on workspaces is kind of a big deal. So with many of these updates, it's really interesting as we talk to listeners about, some things are relevant, some things are not, and everyone has slightly different focuses in their world, but it's kind of nice to know all these things are happening at the same time. I love the JSON Protocol support for SQS. It's kind of interesting at a very plumbing level. I mean, there's lots, so it's been fun.

Jillian Forde [\(38:31\)](#):

There is. There's lots for everyone to just geek out over. I love it.

Simon Elisha [\(38:34\)](#):

Exactly. Jillian, how do people get in touch with you?

Jillian Forde [\(38:37\)](#):

Miss Jill Ford on Twitter.

Simon Elisha [\(38:38\)](#):

That's the way. And we of course love getting all kinds of feedback. awspodcast@amazon.com is the place to do it if you want to go old school. And until next time, keep on building.