

Automated ([00:00](#)):

This is Episode 590 of the AWS Podcast, released on May 29th, 2023. Podcast confirmed. Welcome to the official AWS Podcast.

Simon Elisha ([00:16](#)):

Hello everyone and welcome back to the AWS Podcast, Simon Elisha here with you, great to have you back. And before we get crack-a-lacking, I wanted to let you know that there's a new e-book available, the AWS Enablers e-book, which is a partnership with Accenture, talking about Accenture life trends and the power of the Cloud. This is a global world first, and Accenture's Song Life Trends are partnering with AWS to present the AWS Enablers, showcasing how Life Trends and the power of the Cloud enables business to thrive. So you can check that out. There's a link in the show notes and take a look.

([00:47](#)):

But we have some important work to be done today, which is to talk through a whole bunch of updates. And to do that, I'm joined by our Hawn Nguyen-Loughren. Hey Hawn, How you doing?

Hawn Nguyen-Loughren ([00:56](#)):

Ready for another adventure with the Three Amigos.

Simon Elisha ([00:59](#)):

Excellent. And that means of course, that Jillian Forde is here too. Good day Jillian, how you doing?

Jillian Forde ([01:04](#)):

Hello. So excited to be here.

Simon Elisha ([01:07](#)):

It's good to have you here. Well, Jillian, should we get started?

Jillian Forde ([01:10](#)):

Let's do it. All right.

([01:12](#)):

So first up we've got is the AWS marketplace, and AWS marketplace introduces new self-service listing feature for single AMI products.

([01:25](#)):

So now we're onto analytics where there's a lot of different updates. So the first one, Amazon EMR on EKS now supports vertical auto-scaling. This is a feature to automatically tune the memory and CPU resources for EMR Spark. Tuning these values was a manual process for customers that can be complex. Vertical auto-scaling automatically scales the memory and CPU allocated to an Amazon EMR Spark application based on its realtime and historic resource utilization. This simplifies the process of tuning resources and optimizing costs for an application while improving its reliability. Amazon EMR on EKS now supports self-hosted notebooks for managed endpoints. AWS Glue G.4X and G.8X instance types are now generally available. AWS Glue Crawler now supports custom JDBC drivers.

([02:25](#)):

Amazon Kinesis Data Firehose adds support for document ID that is auto-generated by Amazon Open Search Service.

[\(02:33\)](#):

We've got some more updates from Amazon QuickSight. So Amazon QuickSight adds new scatter plot options supporting additional use cases. Amazon QuickSight now supports state persistence and bookmarks for embedded dashboards. Amazon QuickSight supports VBC connections through public APIs with multi-AZ support and Amazon QuickSight launches dataset parameters to optimize slicing and dicing experiences.

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

Amazon Athena now support Apache Hudi 0.12.2.

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

Amazon MSK adds support for Apache Kafka, version 3.4.

[\(03:15\)](#):

Any Amazon MemoryDB for Redis adds support for Redis 7. Amazon MemoryDB for Redis now supports IAM authentication.

[\(03:26\)](#):

And AWS announces multi-AZ with standby for Amazon Open Search Service.

Simon Elisha [\(03:33\)](#):

Moving on the topic of application integration, Amazon SNS now supports faster automated deletion of your unconfirmed subscriptions. So now this will happen once they've been in a pending confirmation state for 48 hours. Now this reduces the time to delete your unconfirmed subscriptions from the previous 72 hour period. This applies to all new subscriptions, and does not require any onboarding.

[\(03:57\)](#):

Now Amazon SNS, if you've not come across it's a messaging service for application-to-application and application-to-person communication. And the application-to-application functionality gives you high throughput push-based many-to-many messaging between distributed systems, microservices and event-driven serverless applications. Now the A2P functionality lets you communicate with your customers via mobile text messages, so SMS, mobile push notifications, and email notifications. Now, unconfirmed SNS subscriptions are subscriptions pending a confirmation state. Because when you subscribe to an endpoint that's in another AWS account, or if the endpoint type is HTTPS or email, your subscription will be in a pending state til the subscription owner performs a configuration subscription API action to confirm the subscription.

[\(04:44\)](#):

Now, if that doesn't happen, then we want to remove that unconfirmed subscription. This helps avoid spam, et cetera.

[\(04:50\)](#):

AWS AppSync is a fully managed service that enables developers to build scalable and performant APIs. And now there is general availability of the private API support for AWS AppSync. So you can now create GraphQL APIs that can only be accessed from your VPC.

[NEW_PARAGRAPH]AWS Health now publishes service health events to Amazon Event Bridge. Previously you could only receive account specific events on this channel. Now you can also receive

events about the overall health of AWS services, which means you can get the full picture of service issues that might affect the performance of your applications.

[\(05:27\)](#):

Onto the topic of compute, you can now use X-Ray to trace and analyze your Lambda functions enabled with Lambda SnapStart. So this gives you deeper visibility into your functions performance and execution lifecycle.

[\(05:39\)](#):

We're really happy to announce the Amazon EC2 I4g storage-optimized instances. These are powered by AWS Graviton 2 processors, and they deliver the best compute price performance for a storage-optimized instance versus comparable X86-based storage-optimized instance, and the best storage performance per terabyte for a Graviton-based storage instance. Now based on AWS Nitro SSDs that are custom-built by AWS, and reduce both latency and latency variability I4g instances optimized for workloads that perform a high mix of random read/write and require very low I/O latency, think transactional databases and realtime analytics. Now I4g instances improve realtime storage performance by up to two times compared to prior generation storage-optimized instances. You get 8Gb of memory per VCPU and this lets you get a whole lot of power for a whole lot less cost.

[\(06:34\)](#):

And we're also happy to announce the launch of a customized dashboard feature on the AWS Batch Console, which gives you a single view to display resource metrics based on your specific needs.

Hawn Nguyen-Loughren [\(06:46\)](#):

Up next, customer engagement. Amazon Connect Contact Lens now provides supervisor alerts on agent performance. With this launch, supervisors can get automated notifications whenever agents on their team need additional guidance to improve their performance, such as evaluation scores. It helps meet their contact center's quality assurance requirements such as adherence to talk scripts, compliance with sensitive data collection practices, et cetera.

[\(07:11\)](#):

ASW Service Management Connector now supports provisioning with Terraform. So these Terraform products will render as ServiceNow service catalog items that can leverage approvals and business workflows within ServiceNow. Customers can now use AWS Service Catalog as a single tool to organize, govern and distribute their Terraform configuration within AWS at scale.

Jillian Forde [\(07:33\)](#):

All right, databases.

[\(07:35\)](#):

Amazon Aurora I/O-Optimized is generally available. With Aurora I/O-Optimized, there are zero charges for read and write I/O operations. You only pay for your database instances and storage uses, making it easy to predict your database spend upfront. You can now choose between two configurations, Aurora Standard or Aurora I/O-Optimized. For applications with low to moderate I/Os, Aurora Standard is a cost-effective option. And for applications with High I/Os Aurora I/O-Optimized provides improved price performance, predictable pricing, and up to 40% cost savings.

[\(08:17\)](#):

AWS Graviton3-based R7g database instances are now generally available for Amazon Aurora with Postgres compatibility and with Amazon Aurora with MySQL compatibility in Northern Virginia, Ohio, Oregon, and Ireland regions.

(08:34):

Amazon ElastiCache for Redis now supports enabling cluster mode configuration on existing clusters.

(08:42):

Amazon RDS for Postgres now supports pgvector for simplified machine learning model integration.

Simon Elisha (08:52):

Let's talk about building code, and you can't build code without developer tools. And Amazon CodeGuru Security now supports security and code quality scans for Amazon SageMaker Studio and Jupiter Notebooks. So this helps you detect your vulnerabilities like injection flaws, data leaks, weak cryptography, meaty encryption, a whole bunch of stuff. Detecting problems to the left of the development lifecycle is always an important thing.

(09:16):

And you can't have development without a front end, and AWS Device Farm is an application testing service that lets you provide web and mobile developers with desktop browsers and real mobile devices, so you can approve the quality of your apps, test your stuff. You can now have rooted Android devices in the private device fleets, so you can do all kinds of stuff including more metrics, battery benchmarking, and increased access to the file system.

Hawn Nguyen-Loughren (09:43):

And for a fun topic, Internet of Things.

[NEW_PARAGRAPH]AWS IoT Core Registry has additional AWS Cloud Formation resource types. You can now use AWS Cloud Formation templates to pre-configure and deploy AWS IoT Core Registry resources for your targeted Internet of Things, IoT solutions in a secure, efficient and repeatable way, which is pretty handy.

(10:05):

AWS IoT SiteWise now supports 15-minute intervals for automatically computing aggregated asset property values. Customers can now get aggregated values for asset properties, metrics and transforms, which were mathematical expressions that map asset property data points from one form to another for every minute, 15 minutes, hour, and day. AWS IoT SiteWise enhances optimized storage in hot path data for higher throughput. We have improved the GetAssetPropertyValueHistory and BatchGetAssetPropertyValueHistory APIs by increasing the maximum number of results for each paginated response from 250 up to 20k. With these upgrades, developers can now retrieve asset property data, including measurements, attributes, metrics, and transform at high velocity to build industrial applications.

Simon Elisha (10:57):

So it seems like everyone's talking machine learning these days. So let me join in.

(11:00):

Amazon CodeWhisperer now is available as an extension in Jupiter Lab and Amazon SageMaker Studio. So this gives you realtime code suggestions for your work in those environments. And you can write a comment in natural language, and get the code for it, and it's no additional charge.

([11:18](#)):

Amazon SageMaker Studio Notebooks and Notebook instances now support the ml.p4d and ml.p4de GPU-based instances that give you the best performance for interactive machine learning workloads, especially large language models with billions of parameters. This is really important because training your models takes lots of processing. And SageMaker has also announced ml.inf2 and ml.trn1 instances for model deployment as well. These are also super useful for things like text summarization, code generation, video, et cetera, et cetera.

([11:51](#)):

There is now also provision concurrency support for Amazon SageMaker Serverless Inference. So this lets you deploy your models on serverless endpoints with predictable performance and high scalability.

([12:02](#)):

Amazon SageMaker Canvas can now operationalize ML models into production, so just with a single click you can get up and running, and run what-if analysis, and generate single or bulk predictions.

([12:15](#)):

SageMaker Autopilot now supports training ML models with weights, and eight additional objective metrics.

([12:22](#)):

Amazon Rekognition, which lets you identify things in photos, has announced the general availability of face occlusion detection to improve face verification accuracy. The new FaceOccluded attribute in the DetectFaces and IndexFaces APIs detects if the face isn't in image is partially captured or not fully visible due to overlapping objects, clothing and body parts. And Amazon Rekognition has also improved the accuracy of content moderation for images and videos. And finally on this one, Amazon Kendra now supports content-based query suggestions. So these are seeded from index content, allowing you to provide a document-centric query auto-completion experience for your end users.

Jillian Forde ([13:05](#)):

All right, now management and governance.

([13:09](#)):

Private Access to the AWS Management Console is generally available. Private Access, this is a new security feature that allows customers to limit access to the AWS Management console from their VPC, or connected networks to a set of trusted AWS accounts and organizations. Customers can designate which networks are allowed to access their accounts, and AWS organizations from the AWS Console. It also denies attempts to access any other AWS accounts in the AWS Management Console from their network.

([13:43](#)):

AWS Console mobile application launches push notifications. Woohoo!

([13:48](#)):

AWS User Notifications are generally available. AWS User Notifications is a new service that enables you to centrally set up and view notifications from AWS Services such as AWS Health Events, Amazon CloudWatch alarms or Amazon EC2 instance state changes in a consistent human-readable format. You

can view notifications across accounts, regions, and services in a Console Notification Center, and configure delivery channels where you want to receive these notifications, maybe that's email or AWS chat bot, or AWS Console mobile application. Notifications can include URLs to direct to resources on the Console where you can then go and take additional actions.

[\(14:33\)](#):

AWS Systems Manager now allows customers to optimize compute costs of their applications.

[\(14:40\)](#):

AWS CloudTrail Lake enhances query support for Presto SQL Select functions.

[\(14:47\)](#):

Amazon CloudWatch Synthetics announces new Synthetics NodeJS runtime version 4.0.

[\(14:54\)](#):

Amazon CloudWatch Metric Streams adds support for filtering by metric name.

[\(15:00\)](#):

And AWS Resilience Hub expands AWS Trusted Advisor and Amazon DynamoDB support.

[\(15:08\)](#):

And finally, AWS Well-Architected Tool deepens its integration with AWS Service Catalog App Registry.

Hawn Nguyen-Loughren [\(15:18\)](#):

And for an entertaining topic, media services.

[\(15:21\)](#):

Monitor the health of your Amazon IVS channels with multiple hosts. Hosts are connected to each other using stage resource, which allows them to exchange audio and video with each other in realtime. This release gives developers building with Amazon Interactive Video Service, IVS the tools needed to better understand what's happening with the stage resources and for any participants, which is pretty neat.

[\(15:45\)](#):

AWS Elemental MediaConvert now supports video pass through. You can now run MediaConvert jobs that preserve the original video essence without re-encoding for intraframe-only video formats such as AVC-Intra, Apple ProRes, VC3 and JPEG 2000.

[\(16:03\)](#):

Into the topic of networking and content delivery. Amazon CloudFront announces one-click security protections. You can now secure your web applications and APIs with AWS WAF with a single click in the Amazon CloudFront Console, which is super convenient CloudFront can create and configure out-of-the-box AWS WAF protection for your application as a first line of defense against common web threats. Optionally, you can later configure additional security protection against bots and frauds or other threats specific to your applications in the AWS WAF Console. Previously, you could secure your CloudFront distribution with AWS WAF by pre-configuring AWS WAF Web Access Control List, Web ACL, containing the security rules you wanted to enable. While this approach offers flexibility, you had to decide initial security rules to enable and you needed to interact with both the CloudFront and AWS WAF management consoles. Now CloudFront handles creating and configuring AWS WAF for you out the box protection recommended by AWS for all applications.

[\(17:07\)](#):

This simple and convenient way to protect your web application and APIs is available in the CloudFront Console at the time you create or edit your distribution. Customers who prefer to use existing Web ACL may continue to select pre-configured Web ACL instead. One click security protection with a AWS WAF is available in the new web application firewall WAF section in the CloudFront Console and can be used to configure new or existing CloudFront distribution.

[\(17:35\)](#):

AWS App Mesh now supports AWS PrivateLink. You can now use AWS PrivateLink to privately access AWS App Mesh APIs from your Amazon virtual Private Cloud, VPC, without exposing your data through public Internet. This is pretty cool. AWS App Mesh standardizes how you or your services communicate, giving you end-to-end visibility and option to tune for high availability of your applications.

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

Very cool stuff.

[\(18:03\)](#):

Now let's talk about security, identity and compliance. And this is a really interesting new thing we're thrilled to introduce, Cedar, C-E-D-A-R, like the tree, which is an open-source language for access control. You can use Cedar to express fine-grained permissions as easy-to-understand policies enforced in your applications, and you can decouple access control from your application logic. Cedar supports common authorization models like role-based access controls and attribute-based access controls. It follows a new verification-guided development process to give you high assurance of Cedar's correctness and security. AWS formally models Cedar's authorization engine and other tools, proves safety and correctness properties with about them using automated reasoning and rigorously tests that the model matches the Rust implementation.

[\(18:53\)](#):

So for example, Amazon Verified Permissions uses Cedar to allow you to manage fine-grained permissions in your custom applications. With Amazon Verified Permissions, you can store Cedar policies centrally, have low latency with millisecond processing, and audit permissions across different applications.

[\(19:09\)](#):

And now with the Cedar open-source libraries, you can test and validate Cedar policies on a local computer before deploying them with Amazon Verified Permissions. You can also adapt these open source libraries for your requirements and implement use cases such as running applications that are disconnected from the network. Now this is open-sourced under the Apache License 2.0 and includes the Cedar language specification and software development kit. So you can get up and running, and we really encourage you to contribute to the repo and also the Cedar Policy Slack workspace. This is a really exciting development and puts the automated reasoning that we've spoken about many times on the podcast into your hands.

[\(19:49\)](#):

Now AWS Security Hub, which is a cloud security posture management service now features a detailed history of changes that have occurred for each finding in your environment, so you get an immutable trail of changes to see what fields were changed by whom and when.

[\(20:04\)](#):

And Amazon Inspector now allows customers to search its vulnerability intelligent database if any of the inspector scanning types is activated. Now you can retrieve the details for any vulnerability stored in the database and covered by its engine just by querying via a CVE.

(20:21):

AWS Network Firewall now allows you to override the Suricata HOME_NET variable making it easy to use AWS managed rule groups in firewalls that are deployed in a centralized deployment model.

(20:34):

And AWS Network Firewall also now supports the reject action in the stream exception policy, and this improves the performance of latency-sensitive applications.

Hawn Nguyen-Loughren (20:47):

Now for the last but not least topic, storage.

(20:50):

AWS Backup now supports AWS User Notifications. With this launch, you can view the progress of your backup, copy, and restore jobs and changes to your backup policies, vaults, recovery points, and settings from the User Notifications Notification Center, simplifying managing notifications and allowing you to respond to your backup events faster.

(21:11):

AWS VSS application-consistent backup now supports PowerShell logging. With this launch, customers will see improved compatibility between VSS and applications such as antivirus and application monitoring that require PowerShell logging to be enabled. Volume Shadow Copy Service or VSS, is a Microsoft technology that allows application data to be backed up while applications are still running by coordinating between user applications that update data on disk and those backup applications. AWS VSS application-consistent backups enables customers to take application-consistent snapshots of their Amazon EC2 Windows instances and Associate Elastic Block Store volumes, which is pretty cool.

Simon Elisha (21:54):

Yeah, backup of your stuff is always important.

Hawn Nguyen-Loughren (21:57):

Always.

Simon Elisha (21:58):

Always. Now, there were lots of updates there. In fact, there were, by my count, 62 different things that we covered in just over 20 minutes because we're experimenting with the format as we've mentioned now, we do love to get your feedback and figure out what's working for you, what's not, what should we change, what should we keep? So feel free to reach out.

(22:15):

Hawn, how do people do that? How do they get in touch with you?

Hawn Nguyen-Loughren (22:18):

You can reach out to me. I'm a Twitter handler @hawnsolo_1.

Simon Elisha ([22:23](#)):

And Jillian, how do people get in touch with you?

Jillian Forde ([22:26](#)):

missjillforde on Twitter.

Simon Elisha ([22:28](#)):

Couldn't be easier than that, There you go.

([22:30](#)):

And of course, if you like it old school like me, awspodcast@amazon.com is the email way to do it.

([22:38](#)):

And of course, until next time, keep on building.