

# **AWS Architecture Monthly**



January 2020

**AWS re:Invent**



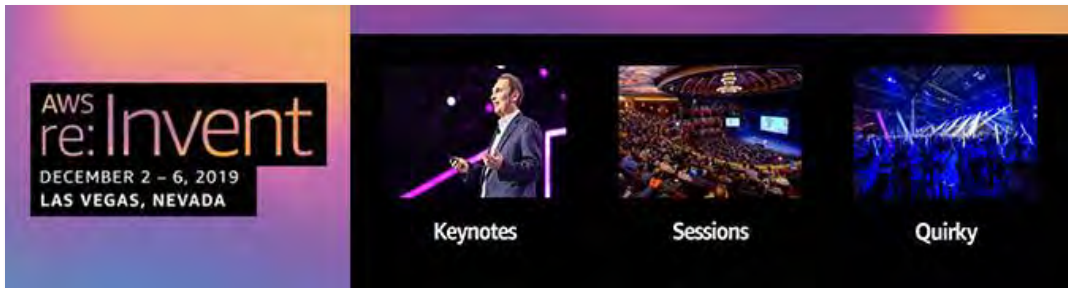
## Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents current AWS product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers or licensors. AWS products or services are provided “as is” without warranties, representations, or conditions of any kind, whether express or implied. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

© 2020 Amazon Web Services, Inc. or its affiliates. All rights reserved.

# Editor's Note

This month's magazine looks back at AWS re:Invent 2019, held in Las Vegas, Nevada, December 2-6, 2019.



Throughout the magazine, we've sprinkled quotes from real AWS customers and partners who filmed This Is My Architecture (<https://amzn.to/tma>) videos (in the studio and live, on Twitch) at re:Invent. We loved getting to meet them and learn more about how these companies are innovating with AWS technologies.

We hope you'll find this edition of Architecture Monthly useful, and we'd like your feedback. Please give us a star rating and your comments on Amazon. You can also reach out to [aws-architecture-monthly@amazon.com](mailto:aws-architecture-monthly@amazon.com) anytime. You can also view past issues at <https://aws.amazon.com/whitepapers/kindle/>.

"We come to re:Invent because you see the bigger picture, and how people are using AWS. It really opens the mind. It's good to contribute, as well, and good to share with a This is My Architecture video." — *Thomas Coombs, Architect and Team Lead, Dentsu Tracking*

## In January's issue:

- **Ask an Expert:** George Mao, AWS Specialist Solutions Architect
- **Key Notes:** Watch Amazon leaders as they take to the various stages at re:Invent
- **Announcements:** We call out a few of our favorites announcements.
- **This Is My Architecture LIVE:** Accolade
- **Blogs:** Cost Management at re:Invent: Takeaways
- **Training:** Free Digital Training Series on AWS Snowball Edge
- **Podcasts:** Our regular AWS Podcaster, Simon Elisha, recaps all of re:Invent 2019

*Annik Stahl, Managing Editor*

# Ask an Expert:

George Mao,  
Specialist Solutions Architect, Serverless

I had the opportunity to sit down with George Mao, AWS Specialist Solutions Architect, and ask him a few questions about re:Invent.

## **Why do you think customers and partners flock to re:Invent—more and more of them each year?**

One of the most unique things about re:Invent is that it's not a sale conference. re:Invent is designed to provide customers with a week of nearly unlimited learning activities. It's an amazing environment where customers have the opportunity to spend a week with AWS technical experts and partners to learn how to leverage the AWS cloud platform. Customers love coming because they hear from AWS specialists on deep-dive topics such as "Lambda performance tuning" or "Advanced DynamoDB Design." Customers who are looking to get hands-on experience have the opportunity to participate in live workshops and builder sessions to enhance their skills and learn best practices from AWS Solutions Architects.

## **Did you have any really interesting conversations with any customers about their company and how they use AWS services?**

I spoke with Accolade on how they're using serverless to both save nearly 50% of their costs and reduce infrastructure management time, not including the overhead they would have needed to manage a "serverful" architecture. They're building amazing microservices on API Gateway and Lambda at half the cost of their previous solution. Accolade is proud of their work, and they filmed a really great This is My Architecture session and Customer Testimonial video during re:Invent. They have now invested in Amazon Connect to replace their legacy call center solution.

"I want to entice other folks to realize that you know, Lockheed is not this old stodgy corporation. Yeah, we've been around for 100 years, but we're introducing new technologies and being disruptive and adding a new digital work experience. And that's really cool." — *Gregg Doppelheuer, Principal System Architect, Internet of Things, Lockheed Martin*

## **What do you believe is the one or two things customers can get out of re:Invent, and what are people who don't go missing out on?**

I think one of the best things about re:Invent is the opportunity for customers to meet AWS teams in person to discuss roadmaps and provide AWS with detailed feedback on their needs. Customer feedback is one of the most important things our service teams use to

understand customer requirements. Attendees get to see the newest launches and innovations right in the Expo Hall on the Launchpad and at our amazing keynotes from Andy and Werner.

Finally, beyond the technical learning, there's always the "fun" part! Customers look forward to spending time with their colleagues and business partners at events such as Midnight Madness and re:Play. It's an excellent team-building environment with customer teams and the extended AWS team.

### **Anything you'd like to add?**

I have been with AWS for nearly five years, and I currently lead the technical field for Serverless at AWS. It is amazing to meet customers all week and learn their use cases, architecture, and talk about how they are using AWS in ways I would have never imagined. Customers attend re:Invent to learn, but we are also learning from our customers all week long.



George Mao is a Specialist Solutions Architect at Amazon Web Services, focused on the Serverless platform. He is responsible for helping customers design and operate Serverless applications using services like AWS Lambda, Amazon API Gateway, Amazon Cognito, and Amazon DynamoDB. George is a software engineer and enjoys contributing to open source projects, delivering technical presentations at technology events, and working with customers to design their applications in the Cloud. George holds a Bachelor's of Computer Science and Master's of IT from Virginia Tech.

# re:Invent 2019: Keynotes

At our eighth annual re:Invent, more than 60,000 AWS partners and customers joined 6,000+ AWS employees in Las Vegas. Below are three top keynote addresses from senior AWS leaders.

## Andy Jassy, CEO of Amazon Web Services



Andy Jassy delivers his AWS re:Invent 2019 keynote, featuring the latest news and announcements, including the launches of M6g, R6g, C6g instances, Inf1 instances, Amazon EKS for AWS Fargate, Amazon S3 Access Points, Amazon Redshift RA3 instances, AQUA (Advanced Query Accelerator) for Amazon Redshift, UltraWarm for Amazon Elasticsearch Service, Amazon Managed Apache Cassandra Service, Amazon SageMaker Studio, Amazon SageMaker Notebooks, Amazon SageMaker Experiments, Amazon SageMaker Debugger, Amazon SageMaker Model Monitor, Amazon SageMaker Autopilot, Amazon Fraud Detector, Amazon CodeGuru, Contact Lens for Amazon Connect, Amazon Kendra, the general availability of AWS Outposts, AWS Local Zones, and AWS Wavelength.

<https://amzn.to/aws-rv2019-jassy>

Guest speakers include Dr. Matt Wood, of AWS; David Solomon, of Goldman Sachs; Brent Shafer, of Cerner, and Hans Vestberg, of Verizon.

## Werner Vogels, CTO of Amazon.com



Werner Vogels delivers his AWS re:Invent 2019 keynote and talks about the innovations behind AWS Nitro System, Firecracker, AWS Fargate, Amazon EBS, and more. This year's featured guests include Clare Liguori, Principal Software Engineer at AWS, who delves into AWS Fargate and how we've built this to give customers a truly unique customer experience. Also: Jeff Dowds, IT Executive at Vanguard, Sebastien de Halleux, COO of Saildrone, and Dr. Martin Hofmann, Group CIO of Volkswagen Automotive Group.

<https://amzn.to/aws-rv2019-vogel>

## Peter DeSantis, VP, AWS



Watch Peter DeSantis deliver the Monday Night Live keynote, featuring Cris Collinsworth from PFF, Rob Smedley from Formula 1, and Daphne Koller from Insitro.

<https://amzn.to/aws-rv2019-desantis>

# Announcements at re:Invent 2019

There were scores of announcements regarding new services, updates, and other newsy things. Here are five we find particularly compelling.

You can discover all of them here: <https://aws.amazon.com/new/reinvent/>

## AWS Lambda Provisioned Concurrency

**Keep functions initialized and hyper-ready to respond in milliseconds.**

We are announcing Provisioned Concurrency, a feature that provides customers greater control over performance of their serverless applications at any scale. Functions using Provisioned Concurrency execute with consistent start-up latency making them ideal for building interactive mobile or web backends, latency sensitive microservices, and synchronously invoked APIs.

Hundreds of thousands of customers have adopted AWS Lambda for benefits such as its simple programming model, built-in event triggers, automatic scaling, and fault tolerance. Provisioned Concurrency makes it easier than ever to develop highly scalable serverless applications with predictable latency on AWS Lambda. Developers can simply set their expected concurrency on any version or alias of a function. Thereafter, AWS Lambda ensures that the function begins executing developers' code within double digit milliseconds of being invoked.

**Read the full announcement:** <https://amzn.to/AWS-rv2019-announce-lambda>

## HTTP APIs (Preview) for Amazon API Gateway

**Make it easy to create, publish, maintain, monitor, and secure APIs at scale.**

Amazon API Gateway announced HTTP APIs, enabling customers to quickly build high performance RESTful APIs that are up to 71% cheaper than REST APIs also available from API Gateway. HTTP APIs are optimized for building APIs that proxy to AWS Lambda functions or HTTP backends, making them ideal for serverless workloads.

To build RESTful APIs, you can use either HTTP APIs or REST APIs from API Gateway. REST APIs offer a wide variety of features for building and managing RESTful APIs. HTTP APIs are up to 71% cheaper compared to REST APIs, but offer only API proxy functionality. HTTP



APIs are optimized for performance—they offer the core functionality of API Gateway at a lower price

**Read the full announcement:** <https://amzn.to/aws-rv2019-http-apis>

## Amazon Builders' Library

The Amazon Builders' Library is a collection of living articles that take readers under the hood of how Amazon architects, releases, and operates the software underpinning Amazon.com and AWS. The Builders' Library articles are written by Amazon's senior technical leaders and engineers, covering topics across architecture, software delivery, and operations. For example, readers can see how Amazon automates software delivery to achieve over 150 million deployments a year or how Amazon's engineers implement principles such as shuffle sharding to build resilient systems that are highly available and fault tolerant.

Visit the Amazon Builders' Library <https://amzn.to/aws-rv2019-rbuilders>

## Amazon Kinesis Video Streams Adds Support for Real-Time Two-Way Media Streaming with WebRTC

Amazon Kinesis Video Streams now supports real-time media streaming via WebRTC. WebRTC is an open-source project that enables real-time communication between web browsers, mobile applications, and connected devices via simple APIs. Kinesis Video Streams with WebRTC enables developers to build applications with real-time two-way media streaming and interactivity between applications and connected devices. Typical uses include interactive video and audio between a video doorbell or baby monitor and your mobile phone, video chat, and peer-to-peer media streaming.

**Read the full announcement:** <https://amzn.to/aws-rv2019-kinesis-webrtc>

"We're in the space as a cloud-native platform for video-related services, and we're innovating in this field. We're doing a lot of things that compete against traditional hardware, so I think it's important for us to essentially show what the industry can do and how the cloud can really benefit the media industry." —  
*Mun Wai Kong, CTO, Grabyo*

## Announcing General Availability of AWS Outposts

Available now: AWS Outposts, a new fully managed service that extends AWS infrastructure, AWS services, APIs, and tools to virtually any customer datacenter, co-location space, or on-premises facility for a truly consistent hybrid experience. AWS Outposts is ideal for

workloads that need low latency access to on-premises applications or systems, local data processing, or for local data storage needs.

You can now run Amazon EC2, Amazon EBS, container-based services such as Amazon ECS and Amazon EKS, database services such as Amazon RDS on AWS Outposts, and analytics services such as Amazon EMR locally on your Outposts. You can seamlessly connect to the broad range of services and tools available in the local AWS Region such as Amazon DynamoDB, AWS CloudFormation, Amazon CloudWatch, AWS CloudTrail, Elastic BeanStalk, Cloud 9, and others to run and manage workloads on Outposts the same way as you do in the cloud today.

**Read the full announcement:** <https://amzn.to/aws-rv2019-outposts>



Surely by now you know about This is My Architecture (TMA) (<https://amzn.to/tma>), the video series that highlights innovative cloud architectures from our partners and customers. At this year's re:Invent, we had the opportunity to welcome several partners and customers into our on-site studio to record upcoming videos while also filming more than 10 TMA Live segments streamed on the AWS Twitch channel.

In the below episode, Minhui Burket, Senior Software Engineer for Accolade (<https://www.accolade.com/>), a personalized health and benefits solution, spoke with AWS Solutions Architect Peter Tilsen.

"Because it's fully managed, we save costs because we don't need an IT team to maintain the infrastructure. AWS takes care of everything." — Minhui Burket, Senior Software Engineer, Accolade

<https://amzn.to/aws-rv2019-twitch-accolade>

You can watch all the content we made at re:Invent 2019 (and more) at <https://www.twitch.tv/aws/videos>

# Blog:

## Cost Management at re:Invent: Takeaways

Available online at: <https://amzn.to/aws-rv2019-cost-mgmt>

By Erin Carlson

And just like that, the Expo Hall portion of the re:Invent conference is officially closed for 2019! The AWS Cost Management booth staff would like to thank the hundreds of you who took the time to stop by and share anecdotes about your cloud financial management journey with us. We also appreciate your feedback on opportunities for improving our product capabilities in the future.

### Savings Plans are a Gamechanger

Many of you mentioned that you have already started purchasing Savings Plans. The best practice that we heard from many of you is to start out by purchasing EC2 Instance Family Savings Plans for the instance families that support your key workloads. From there, wait a week or so to see how your usage patterns are trending before diving into purchasing additional EC2 Instance Family Savings Plans or the more broadly-applicable Compute Savings Plans. Don't forget to shrink the lookback period on the Savings Plans recommendations page if you have recently made significant purchases or changes in your usage.

### Leveraging Cost Categories and Tagging Best Practices

In the cost allocation space, there were a couple of exciting announcements this week: Tag Policies and Cost Categories (Public Beta). After talking the features over with many of you, we recommend using the new tag policies to review and classify your resources. From there, enable the appropriate tags for cost allocation purposes from the dedicated page in the Billing Console. Once those tags are enabled, you can create custom business groupings (e.g., costs centers or departments) using tags or linked accounts.

### Accessing Multiple Levels of Data Grouping in AWS Cost Explorer

While you can only group by a single dimension in the Cost Explorer web application, you actually group by two dimensions when you use the Cost Explorer API. Some of you

mentioned that your ideal use case would be to view your costs organized by service and further broken down by usage type, but there are a lot of possibilities here. I would highly recommend looking into using the Cost Explorer API if this sounds interesting to you. You can learn more here: <https://amzn.to/aws-rv2019-cost-explorer>.

Please note that the Cost Explorer API has a cost of \$.01 per paginated request.

## Getting to Know the AWS Compute Optimizer

One of the hot topics of discussion was the brand new AWS Compute Optimizer, which uses machine learning techniques to optimize your compute resources for your workloads. These capabilities extend the Rightsizing Recommendations functionality currently available in Cost Explorer, which help you identify idle instances and to downsize under-utilized instances within the same instance family.

It is also worth mentioning that all of these product capabilities are free! Looking forward to hearing how more about your instance right-sizing strategies.

## Conclusion

I hope that you have all enjoyed this jam-packed week as much as I have! It has been an absolute pleasure getting to know many of you and I am looking forward to seeing what each of you build using the brand-new services and features.

And last, but not least, I would like to extend a special thank you to Levon Stepanian, Alee Whitman, and the rest of the amazing AWS Cost Management booth team for their tireless support and enthusiasm! Our Product Feature Request cup (Mason jar) runneth over. Until next time!

"In essence, we've migrated a multi-billion dollar company into AWS from our old data center." — *James Oravec*  
*Senior Software Engineer for Back Office Systems, Zappos*

# Training:

## Free Digital Training Series on AWS Snowball Edge

We're excited to announce a free new digital training series on AWS Snowball Edge, a petabyte-scale data transport device with on-board storage and compute capabilities. This three-course series is designed to teach you how to use the device and help you decide if it's appropriate for your data transport needs.

The first course, **AWS Snowball Edge Getting Started**, covers common use cases for AWS Snowball Edge, plus device types and job types for your use cases. It also covers the device's physical features, security features, and hardware resource capabilities. Enroll in the free digital course here: <https://amzn.to/aws-rv2019-edge-training-1>

The second course, **AWS Snowball Edge Logistics and Planning**, covers the AWS Regions in which AWS Snowball Edge devices are available. It also covers time frames for receiving and using the device and how AWS Snowball Edge pricing works. Enroll in the free digital course here: <https://amzn.to/aws-rv2019-edge-training-2>

The third course, **AWS Snowball Edge: Using AWS Snowball Edge**, teaches learners how to use the device to transfer data using the AWS Command Line Interface (AWS CLI). The course also teaches how to complete an AWS Snowball Edge job and ship the device back to AWS. Enroll in the free digital course here: <https://amzn.to/aws-rv2019-edge-training-3>

# Podcasts:

## re:Invent 2019, Episodes 1-4

### Episode 1

Simon kicks off the coverage of re:Invent 2019 with some really interesting announcements in Quantum computing and more, as well as a catch-up on all the previous week's updates.

<https://amzn.to/aws-rv2019-podcast1>

### Episode 2

Simon gives you a summary of all the new announcements from Andy Jassy's keynote.

<https://amzn.to/aws-rv2019-podcast2>

### Episode 3

In this third update show, Simon takes you through some quick updates in the domains of Compute, Data, ML, and Business Productivity.

<https://amzn.to/aws-rv2019-podcast3>

### Episode 4

Simon wraps up re:Invent 2019.

<https://amzn.to/aws-rv2019-podcast4>

"The reason we chose to work with AWS is because we want to help provide increased productivity for students. We put 2,300 Echo devices across campus with a skill that supports basic questions a returning student or one soon joining the university may have. And then, of course, they have all the power of Alexa in their room, as well. We saw students interacting over 150,000 times with the devices over the first year, using notifications and our skill, weather, reminders, and making calls." — *Lucas Guffey, St. Louis University*