

Session	Title	Description	Type	Attending	Location
AIM202	Automate your mortgage document processing with AWS AI/ML	It is estimated that a new mortgage loan origination application contains an average of 400–500 pages. Other processes such as mortgage refinance also involve hundreds of documents per application. There is limited automation available today to process and extract information from all the documents, especially due to varying formats and layouts. In this session, learn how organizations in the lending and mortgage industry can take advantage of some of the latest innovations in AI/ML services from AWS to automate document processing and ultimately improve productivity and scale quickly.			
AIM203	Boost user engagement with Amazon Personalize	Many consumers today expect a tailored offering as they consider, purchase, and use products and services. Every digital touchpoint is an opportunity for businesses to deliver on that expectation. In this session, learn how Chewy, an AWS customer, used Amazon Personalize to improve user engagement by providing a dynamic, high-quality shopping experience.			
AIM204	Automate insurance document processing with AI	An estimated 80 percent of data within organizations is considered to be unstructured or “dark” data that is locked inside text, emails, PDFs, and scanned documents. This data is difficult to interpret or analyze programmatically and limits how organizations can derive insights from it and serve their customers more effectively. The rapid rate of data generation means that organizations that aren’t investing in document automation risk getting stuck with legacy processes that are slow, error-prone, and difficult to scale. In this session, learn how organizations can take advantage of the latest innovations in AI and machine learning from AWS to improve the efficiency of their document-intensive claims processing use case.			
AIM207	Make better decisions with no-code ML using SageMaker Canvas, feat. Samsung	Organizations everywhere use ML to accurately predict outcomes and make faster business decisions. However, this often requires preparing, building, training, and deploying ML models. Amazon SageMaker Canvas expands access to ML by giving business analysts a visual point-and-click interface to generate accurate ML predictions on their own without requiring any ML experience or writing a single line of code. In this session, learn how you can use SageMaker Canvas to access and combine data from a variety of sources, clean data, build ML models to generate predictions with a single click, and share models across your organization to improve productivity.			
AIM208	Idea to production on Amazon SageMaker, with Thomson Reuters	Amazon SageMaker brings together fully managed infrastructure, tools, and workflows so you have everything you need to prepare data and build, train, and deploy machine learning (ML) models. SageMaker offers many innovative features such as an IDE for ML and interfaces for business analysts to generate predictions without code. In this session, discover how to use the SageMaker portfolio and learn how SageMaker capabilities work together so that you can get started quickly no matter your role—from business analyst to developer or data scientist. Also hear from Thomson Reuters about how they use SageMaker across the ML lifecycle.	Session	Watch	Wynn

Session	Title	Description	Type	Attending	Location
AIM210	Solve common business problems with AWS AI/ML services	In this session, learn how companies across industries are using artificial intelligence (AI) to address use cases that create measurable results. This session defines the results based on four categories: enhancing customer experience, helping employees and organizations make better and faster decisions, improving business operations while reducing costs, and creating completely new products and services powered by AI and machine learning (ML).	Session	Watch	Caesars Forum
AIM214	Learn how Black Knight is using AI to accelerate mortgage workflows	Explore how Black Knight, a premier mortgage analytics company, is adopting AWS artificial intelligence (AI) services for its data and domains to process documents at scale. Learn how they use Amazon Textract to reduce manual processes, mitigate regulatory risks, and deliver significant cost savings to their clients, including many of the largest US lenders. The mortgage processing industry is complex due to the loan lifecycle, regulatory requirements, and the sophisticated data and analytics required to support each process. Hear how Black Knight uses AWS AI services to help clients improve and scale their business processes with automation and a complete solutions ecosystem that supports the entire real estate and mortgage lifecycle.			
AIM307	JPMorganChase real-time agent assist for contact center productivity	Resolving complex customer issues is often time-consuming and requires agents to quickly gather relevant information from knowledge bases to resolve queries accurately. Join this session to learn how JPMorganChase built an AWS Contact Center Intelligence (CCI) real-time agent assist solution to help 75 million customers and help 8,500 servicing agents generate next best actions in the shortest time—reducing agent frustration and churn. Hear how JPMorganChase’s real-time agent assist solution uses Amazon Transcribe to provide real-time transcriptions, Amazon Kendra to provide best answers from knowledge bases during a live call, and Amazon SageMaker for machine learning and training their intent models.	Session	Watch	Encore
AIM313-R1	Scalable data preparation & ML using Apache Spark on Amazon SageMaker Studio	In this workshop, learn how to build end-to-end data preparation and machine learning (ML) workflows on Amazon SageMaker Studio. Learn how to connect to Apache Spark data processing environments running on Amazon EMR and AWS Glue interactive sessions from SageMaker Studio. Once connected, discover how to access data governed by AWS Lake Formation; interactively query, explore, and visualize data; and run and debug Spark jobs as you prepare large-scale data for use in ML. Finally, find out how you can use the prepared data to train ML models and deploy them at scale. You must bring your laptop to participate.	Workshop	Waitlisted	
AIM314	Accelerate your ML journey with Amazon SageMaker low-code tools	The machine learning (ML) journey requires continuous experimentation and rapid prototyping to be successful. In order to create highly accurate models, data scientists have to first experiment with feature engineering, model selection, and optimization techniques. These processes are traditionally time-consuming and expensive. In this session, learn how low-code tools, including Amazon SageMaker Data Wrangler, Amazon SageMaker Autopilot, and Amazon SageMaker JumpStart, make it easier to experiment faster and bring highly accurate models to production more quickly and efficiently.	Session	Watch	Venetian

Session	Title	Description	Type	Attending	Location
AIM319	Build, manage, and scale ML development with a web-based visual interface	Amazon SageMaker Studio is an integrated development environment (IDE) for data science and machine learning (ML). In this session, explore how to use SageMaker Studio to prepare data and build, train, deploy, and manage ML models from a single, web-based visual interface. Learn how to access a comprehensive set of tools for your ML development and how to quickly move between steps without leaving SageMaker Studio.			
AIM320	Boost ML development productivity with managed Jupyter notebooks	Amazon SageMaker offers two options that use fully managed Jupyter notebooks for data exploration and building ML models. In this session, discover how to use quick-start, collaborative Amazon SageMaker Studio notebooks to increase productivity across all steps in your ML development. Learn how to get started with standalone SageMaker notebook instances that offer a broad choice of cloud compute resources, including GPUs for accelerated computing and the latest versions of open-source ML packages.			
AIM321	Productionize ML workloads using Amazon SageMaker MLOps, feat. NatWest	Machine learning operations (MLOps) tools help you automate and standardize processes across the ML lifecycle to productionize ML models more quickly and maintain model quality in production. Amazon SageMaker provides a breadth of MLOps tools to train, test, troubleshoot, deploy, and govern ML models at scale. In this session, explore Amazon SageMaker MLOps features, including SageMaker Pipelines, SageMaker Projects, SageMaker Experiments, SageMaker Model Registry, and SageMaker Model Monitor, and learn how to increase automation and improve the quality of your ML workflows. Hear from NatWest Group about how they built a secure MLOps platform using Amazon SageMaker to standardize their ML model development and deployment process across the organization.	Session	Watch	Mandalay
AIM332	WaFd bank delivers enhanced self-service with AWS conversational AI	Join this session to learn how Washington Federal (WaFd) bank used AWS conversational AI solutions and Talkdesk to reduce their customer wait time in the interactive voice response system from 4.5 minutes to 28 seconds. See a live demo to understand how WaFd's chatbot goes beyond canned responses and allows customers to perform multiple online banking tasks, such as checking account balances, scheduling transfers, and more. Learn how, powered by Amazon Lex, WaFd's chatbot can interact and respond to customer queries in both voice and text. If a customer's query is unresolved, the AI-powered system connects customers to the agent best equipped to resolve the issue quickly.	Session	Reserved	Wynn
AIM333	How Taco Bell is improving digital availability with ML forecasting	Accurate demand forecasting can help a company maximize revenue and minimize inventory costs. Traditional approaches to deploying ML-based forecasting required customers to click through user interfaces or learn how to use APIs and Jupyter notebooks. Both of these increase the barrier to ongoing, sustainable production operation at scale. In this session, learn how customers like Taco Bell are deploying no-code solutions that automate the Amazon Forecast pipeline using a number of AWS services, such as AWS Step Functions and AWS CloudFormation. This solution helps deploy and manage many decoupled and concurrent enterprise-grade workloads.			

Session	Title	Description	Type	Attending	Location
AIM337-R1	Build better ML models for business decisions using Amazon SageMaker Canvas	Efficient machine learning (ML) models enable desired business outcomes across organizations. In this workshop, learn a no-code approach to ML with Amazon SageMaker Canvas. The workshop poses a specific business problem and you address it by building an ML model using SageMaker Canvas. Import your data, prepare, explore, and transform the data, and finally create an ML model without writing a line of code. Then analyze the model using different metrics in SageMaker Canvas to understand the impact of the model artifacts. Finally, share the model with ML practitioners for additional insights. Learn how you can use SageMaker Canvas to make more effective business decisions. You must bring your laptop to participate.			
ALX201	10 things every voice application should do	In this session, Jeff Blankenburg, Chief Technical Evangelist for Amazon Alexa, discusses his experience building dozens of skills for Alexa and working with hundreds of developers on their own voice applications, and he shares ten specific patterns that are common to the most successful of them all. Join this session to learn insight and acceleration tips that can help you take your voice application to the next level.			
ANT201	What's new with Amazon Redshift	Join this session to hear about important new features of Amazon Redshift—a fast, cost-optimized, and widely used cloud data warehouse. Amazon Redshift delivers powerful and secure insights from all your data at scale, with zero administration required. Learn how Amazon Redshift reinvented data warehousing to help you analyze all your data across data lakes, data warehouses, and databases with the best price performance and how it empowers you to apply machine learning quickly and easily.	Session	Watch	Encore
ANT204	Enabling agility with data governance on AWS	Data governance is the process of managing data throughout an end-to-end process, ensuring its accuracy and completeness and making sure it is accessible to those who need it. Join this session to learn how AWS is delivering comprehensive data governance— from data preparation and integration to data access, data quality, and metadata management—across analytics services.			
ANT205	Achieving your modern data architecture	A modern data strategy is a comprehensive plan or blueprint to help you manage, access, analyze, and act on your data. Many organizations are developing roadmaps to help them achieve digital transformation. This session helps you understand how deploying a modern data architecture can help you navigate your data challenges, optimize analytics processes, and deliver faster insights to the people and applications that need it.	Session	Watch	Mandalay
ANT207	What's new in AWS Streaming	Organizations across industries need to make quick, real-time decisions to reduce costs, improve inefficiencies, increase customer engagement, and drive more revenue. This session reviews recent AWS streaming service launches and how you can use them to solve a variety of streaming use cases.	Session	Watch	MGM Grand
ANT208	What's new in Amazon Athena	Amazon Athena is a highly scalable analytics service that makes it easy to analyze all your data across Amazon S3, in on-premises stores, and on other cloud platforms. Amazon Athena is serverless, so there is no infrastructure to manage, and you pay only for the queries that you run. This session offers a deep dive into the service, customer use cases, best practices, newly launched features, and what's next for Amazon Athena.	Session	Watch	Caesars Forum

Session	Title	Description	Type	Attending	Location
ANT209	Build interactive analytics applications	Spend more time on insights and less time waiting for results with Amazon Athena, where applications start instantly. In this session, learn how to run complex applications faster, more simply, and most effectively with support for new frameworks.			
ANT212	How JPMC and LexisNexis modernize analytics with Amazon Redshift	Tens of thousands of customers use Amazon Redshift every day to run data warehousing and analytics in the cloud and process exabytes of data for business insights. Join this session to hear from JPMorgan Chase and LexisNexis as they share their journeys of moving to a modern cloud data warehouse with Amazon Redshift. Learn best practices in modernizing analytics to operate at scale, with speed while keeping costs low.			
ANT220	Using Amazon AppFlow to break down data silos for analytics and ML	Data silos can create barriers to efficiency and data quality and, as a result, often hinder an organization's ability to begin their data and ML journeys in the cloud. Manual data ingestion, data preparation, and application integration are time-consuming and expensive. In this session, learn how organizations can efficiently break down data silos with low-code and no-code integration services, such as Amazon AppFlow, and speed up innovation with the help of AWS analytics, AI/ML, and integration services. Also get prescriptive guidance on how to select the best tool for your integration strategy.			
ANT223	Simplify and accelerate data integration & ETL modernization with AWS Glue	The first step in an analytics or machine learning project is to discover and prepare your data to obtain quality results. AWS Glue is a serverless, scalable data integration service that helps you discover, prepare, move, and integrate data from multiple sources. In this session, learn about the latest innovations in AWS Glue and hear how an AWS customer uses AWS Glue to enable self-service data preparation across their organization. Itau, Brazil's largest private-sector bank, also shares their AWS Glue success story.	Session	Watch	MGM Grand
ANT301	Democratizing your organization's data analytics experience	AWS analytics services empower data users, such as data scientists, analysts, and business users with diverse technical expertise, across an organization to quickly access, analyze, and gain insights from their data. In this session, learn how AWS can democratize analytics for your organization with ease of use and better price performance. Dive into how serverless AWS simplifies data preparation and makes it easier for analysts and data scientists of all skill levels to use machine learning.			
ANT302	What's new with Amazon EMR	Amazon EMR makes it easy to run and scale big data applications built using open-source frameworks such as Apache Spark, Hive, Presto, Hudi, and more. It helps users run ETL and ML jobs and accelerate data science adoption, real-time data processing, and low-latency SQL at petabyte scale. In this session, learn about the latest and greatest Amazon EMR developments, including Amazon EMR Serverless, Amazon EMR Studio, and more. Also, find out how to add support for transactions and real-time streams in data lakes with Hudi and Apache Iceberg and how to enforce fine-grained access control	Session	Watch	Wynn

Session	Title	Description	Type	Attending	Location
ANT310	Build a data mesh with AWS Lake Formation and AWS Glue	In this workshop, learn how to build a data mesh architecture on AWS. Organizations are interested in implementing this architecture to move away from centralized data lakes and toward decentralized ownership and delivery of analytics solutions across business units. Learn about data mesh fundamentals and principles, how data mesh can impact your organization, and how data mesh architecture can be implemented with AWS services. Build your own data mesh using AWS Lake Formation, AWS Glue, and AWS CDK, and share data across business units. You must bring your laptop to participate.			
ANT317	Self-service analytics with Amazon Redshift Serverless	Amazon Redshift Serverless lets you get started in seconds and run data warehousing and analytics workloads at scale, without worrying about data warehouse management. In this session, learn how Amazon Redshift automatically provisions data warehouse capacity and intelligently scales the underlying resources to deliver consistently high performance and simplified operations for even the most demanding and volatile workloads.			
ANT324	Modernize your data warehouse	As data volumes grow and the demand to get faster insights increases, organizations face challenges with performance, scalability, and rigidity of traditional data warehouse systems. In this session, explore how you can use the new capabilities of Amazon Redshift to modernize your data warehouse by gaining access to all your data—real-time streaming data, transactional data, semi-structured data, data in open formats, or even third-party data—and share consistently across the organization through a petabyte-scale serverless engine. Learn how Amazon Redshift allows you to build powerful analytics and ML applications for a wide range of use cases and the most demanding workloads.			
ANT325	How Disney processes clickstream data on Amazon EMR Serverless	Amazon EMR Serverless makes it easy to run large-scale distributed data processing jobs using open-source frameworks like Apache Spark and Hive without worrying about configuring, managing, and securing clusters. Learn about running Amazon EMR Serverless jobs cost-effectively, discover new features, and hear Disney discuss their migration to Amazon EMR Serverless. Disney+'s data platform, Wakanda, is built on domain-driven federated architecture using Amazon EMR Serverless to process Adobe clickstream data. Wakanda provides clean, mapped, and standardized views across media networks that power Disney+'s analytics, data science, BI, and reporting. Learn how Disney+ processes large volumes of clickstream data for various workloads performantly with fine-grained scaling and job-specific authorization roles using Amazon EMR Serverless.			
ANT326	How BMW, Intuit & Morningstar are transforming with AWS and Amazon Athena	Data science and machine learning are transforming products and services in every industry. Whether it's improving customer experience at BMW, personalizing software in real time at Intuit, or accelerating insights and innovation at Morningstar, data science and ML techniques are being used to improve business processes, speed product development, and enhance customer engagement. In this session, hear firsthand from AWS customers using Amazon Athena alongside other AWS open-source services to modernize their businesses and bring transformative change to their industries.			

Session	Title	Description	Type	Attending	Location
ANT335	How Disney used AWS Glue as a data integration and ETL framework	Disney Parks, Experiences and Products is one of the world's leading providers of family travel and leisure experiences. Disney Parks, Experiences and Products uses AWS Glue—a serverless data integration service—as a key component to replace thousands of Apache Hadoop, Spark, and Sqoop jobs. In this session, Disney and AWS Glue experts discuss some ways to scale AWS Glue beyond the traditional setup and how they configure AWS Glue for job monitoring and performance.			
ANT336	Building data mesh architectures on AWS	Data-first organizations are increasingly curious about data mesh architectures. In this session, learn how to design, build, and operationalize a data mesh architecture on AWS so you can navigate data challenges, optimize analytics processes, and deliver insights to the business faster. Hear from GoDaddy's Chief Data Officer Travis Muhlestein about how they delivered their data mesh using AWS.	Session	Watch	Wynn
ANT337	Migrating to Amazon EMR to reduce costs and simplify operations	Do you run Apache Spark, Presto, Hive, or Flink applications on premises? Running open-source data analytics frameworks requires hard-to-find skills, knowledge, and experience. In this session, Thomson Reuters and LendingClub share how they reduced costs and simplified operations by migrating to Amazon EMR.			
ANT339	How Samsung modernized architecture for real-time analytics	In this session, Samsung SmartThings shares how they modernized their streaming data analytics architecture for real-time analytics. Originally Samsung developers spent most of their time managing infrastructure. After migration to Amazon Kinesis Data Analytics, developers were able to focus on delivering business value without needing to bother about infrastructure management. Learn how they moved from self-managed Apache Spark to a fully managed stream-processing platform with Amazon Kinesis Data Analytics for Apache Flink and how they securely streamed data at scale with fully managed Kinesis Data Streams and Amazon MSK.			
ANT340	How Sony Orchard accelerated innovation with Amazon MSK	In this session, The Orchard, a subsidiary of Sony Music Entertainment, shares how they built a high-performing data synchronization solution using Amazon MSK. By using this highly available managed service, Sony Orchard reduced the time to synchronize data from 10 hours to 30 minutes, enabled an event-driven architecture paradigm, and improved search capabilities. Discover how Amazon MSK helped Sony Orchard innovate faster and provide access to near real-time data for both employees and customers.			
ANT341	How Riot Games processes 20 TB of analytics data daily on AWS	Riot Games ingests about 20 TB of data every day on AWS. This data powers a wide range of services including game matchmaking, in-game personalization, analytics, security, and player behavior management. Until recently, this data was only queryable for up to 6 hours after it was produced. With Amazon MSK, the team reduced time down to 5 minutes. Amazon MSK also allowed Riot Games to lower TCO and deprecate an aging map-reduce-based pipeline. In this session, Riot Games describes the before state, migration path, current state and future state of their data ingestion pipeline			

Session	Title	Description	Type	Attending	Location
ANT342	How Poshmark accelerates growth via real-time analytics & personalization	In this session, learn how Poshmark is focusing on achieving top-line growth through personalization and enhanced user experience. The initial approach of using batch processing for personalization and security did not meet expectations for customer experience improvement. To enhance customer experience, Poshmark designed real-time personalization using real-time event capture with Amazon MSK and real-time data enrichment with Amazon Kinesis Data Analytics. Amazon SageMaker and Kinesis Data Firehose supported real-time personalization by saving copies of the personalizations to Amazon S3. This solution improved customer experience, reduced security risks, and allowed end users to more confidently interact with the Poshmark app.			
API303-R1	Application integration patterns for microservices	Virtually all integration approaches come with trade-offs. However, loosely coupled integration can help you design independent systems that can be developed and operated individually and can also increase the availability and reliability of the overall system landscape—particularly by using asynchronous communication. While there are many approaches for integration and conversation scenarios, it’s not always clear which approach is best for a given situation. Join this session to learn about foundational patterns for integration and conversation scenarios with an emphasis on loose coupling and asynchronous communication. Explore real-world use cases architected with cloud-native and serverless services, and get guidance on choosing integration technology.			
API307-R1	Designing event-driven integrations using Amazon EventBridge	Event-driven architecture can bring alignment to an organization’s business capabilities. Amazon EventBridge enables application integration across business applications and improves modern application development practices. In this session, learn about the roles and responsibilities of stakeholder groups and discover foundational concepts for designing and building a sustainable, scalable, serverless event bus for your organization.			
API308	Are you integrating or building distributed applications?	Integrating systems provides many benefits, from seamless user experiences to consolidated data for better insights to interactions with partner ecosystems through APIs. Despite having connected systems for decades, fundamental concepts like partial failure, eventual consistency, or idempotency still challenge many developers. In this session, reflect on two decades of building connected systems, ponder the nuances of messages versus events, learn that loose coupling isn’t always better, and discover how serverless cloud automation can change the way you think about integration.			
API309	Advanced serverless workflow patterns and best practices	Are you an experienced serverless developer building applications with AWS Step Functions workflows? Do you want a guide for unleashing the full power of Step Functions for your production workloads? Are you wondering whether to choose an Express or a Standard workflow? Or whether to use a Map state, a parallel state or a nested workflow? In this session, learn about architectural best practices and repeatable patterns for building workflows and cost optimizations, and discover handy cheat codes that you can use to build secure, high-scale, high-performance serverless applications. Real production scenarios illustrate the benefits.			

Session	Title	Description	Type	Attending	Location
API311-R	Building next-gen applications with event-driven architectures	Event-driven architectures can solve many difficult problems in modern enterprise workloads. For example, it can be challenging to work with large amounts of data in different data stores and locations. Teams building microservices architecture often find that integration with other applications and external services can make their workloads more monolithic and tightly coupled. In this chalk talk, learn how to use event-based architecture to decouple and decentralize application components. Discover how you can use AWS messaging services to connect microservices and coordinate data flow using minimal custom code.			
API314-R1	What's possible with AWS Step Functions	Building distributed solutions with functions and microservices requires orchestrating workflows. You might need a workflow for engineers to request cloud resources and managers to approve requests. Or a data processing workflow that involves coordinating multiple services like AWS Lambda, Amazon S3, and Amazon Athena. Or maybe a security workflow. AWS Step Functions lets you to build and execute workflows using simple yet powerful visual tools. In this chalk talk, review the options in Step Functions Workflow Studio and learn about the multitude of supported AWS SDK integrations. See a workflow built from scratch and discuss common use cases that span development, operations, and security.	Chalk Talk	Reserved	MGM Grand
ARC205	Establishing your cloud foundation on AWS—beyond your landing zone	After working with thousands of customers, AWS has determined a minimum baseline that customers need to have in their environment as they mature through their cloud journey. Having all the components in your environment to deploy, operate, and govern your workloads allows you to efficiently meet your customers' needs and focus on key differentiators for your business. In this session, learn how to use a set of foundational capabilities in your cloud environment, including the people, technology, and processes you need to succeed. Walk through an example of one of the capabilities and how to establish it in your environment.			
ARC207	Modern cloud applications: Do they lock you in?	Building on the cloud is amazing—you have instant access to fully managed services, elastic on-demand scaling, and fine-grained access control. If you build your applications on a serverless ecosystem, the cloud even takes care of deploying and managing all of the underlying infrastructure. But can letting the cloud do more for you lock you in? What if one day you have to take your application somewhere else? This session discusses trade-offs for modern cloud applications, explains why “lock-in” isn't just the result of the services you use, and shows why the best abstraction layers are the ones in your head.			
ARC301-R1	Building a monitoring strategy	In this workshop, follow Amazon best practices to help organizations build a consistent approach to workload monitoring and observability. The goal is to create a strategy that aligns business and operational metrics. The workshop identifies the key metrics that matter most to delivering successful business outcomes. Customer organizations are most successful when they develop a monitoring strategy that starts with the client experience and utilizes the appropriate technologies. You must bring your laptop to participate.	Workshop	Waitlisted	

Session	Title	Description	Type	Attending	Location
ARC306-R	Multi-Region design patterns and best practices	In this session, dive deep on the topic of multi-Region design and explore different ways to implement such architectures. Consider how to think about multi-Region architectures and what is included from a cost, operations, and engineering perspective. Also learn about AWS service features that can help you build multi-Region solutions, including Amazon Route 53 Application Recovery Controller, Amazon CloudFront Functions, Amazon Aurora Global Database, Amazon DynamoDB global tables, and more.			
ARC311-R	Build production-ready prototypes rapidly using serverless patterns	Applying serverless patterns to develop a new prototype reduces development time and offers many advantages in the operational phase as well. In this session, learn about applying serverless patterns to satisfy non-functional requirements step by step. The final implemented result is a prototype ready for production that delivers security, automation, high availability, and operational excellence. For reference, the functional requirement of this prototype is to develop an API service that analyzes sentences. However, as a non-functional requirement, the analysis results must be visualized, while being designed to be secure and available. In addition, all implementations should be provided in a standardized, automated form using IaC.	Session	Watch	MGM Grand
ARC313-R1	Building modern data architectures on AWS	The modern data architecture is an evolution from data warehouse and data lake-based solutions, and it allows you to query data across your data warehouse, data lake, and purpose-built analytics services to gain faster and deeper insights that would not be possible otherwise. In this session, learn how to design, create, and operate a modern data architecture on AWS by using AWS purpose-built data services. Also learn about modern data architecture concepts, why customers are adapting to it, modern data architecture pillars, reference architectures, and best practices while building modern data architectures on AWS for optimal performance and cost effectiveness.	Session	Watch	Venetian
ARC314	Publishing real-time financial data feeds using Kafka	This session describes how to offer a real-time financial data feed as a service on AWS. By using Amazon Managed Streaming for Apache Kafka (Amazon MSK), you can use Kafka to allow your customers to subscribe to message topics containing the financial data of interest. Consider connectivity best practices to distribute data at a global scale, security options to ensure a secure architecture, and lessons learned from customers that are using AWS to distribute financial data on AWS.			
ARC315	How to understand and improve resiliency for serverless systems	If AWS Lambda is highly available automatically, you might wonder, what is left for architects to do? The answer is: plenty. However, in a serverless world, designing resilient serverless applications can be a little different. In this session, learn to understand and improve resiliency for serverless systems throughout your application. This includes anticipating, monitoring, responding, and learning. Also discover how to validate your resiliency with AWS Resilience Hub and complete experiments with AWS Fault Injection Simulator.			

Session	Title	Description	Type	Attending	Location
BIZ201	Improving multichannel communications	Your communication strategy directly impacts customer satisfaction, revenue, growth, and retention. This session covers how the customer communication landscape is changing and how you can shift your product strategies to impact business outcomes. Learn how AWS Communication Developer Services (CDS) improve customer communication through CPaaS channels like email, SMS, push notifications, telephony, audio, and video. Explore how building a multichannel architecture can shift user habits and increase satisfaction. This session is tailored for chief marketing officers, chief product officers, chief innovation officers, marketing technologists, and communication builders.			
BOA203-R1	Evolution of the machine learning development environment	Software engineering has changed dramatically over the years with better collaborative development practices, migration from monolithic applications to microservices-based deployments, and many others. But what about machine learning? How has machine learning-based development evolved over the years? In this chalk talk, through examples of each approach, discuss early machine learning setup and practices and how it evolved into today's collaborative data science and streamlined machine learning deployment and MLOps best practices.			
BOA309	Build a managed analytics platform for your ecommerce business	With the increase in popularity of online shopping, building an analytics platform for ecommerce is important for any organization, as it provides insights about the business, trends, and customer behavior. But, more importantly, it can uncover hidden insights that can trigger revenue-generating business decisions and actions. In this session, learn how to build a complete analytics platform in batch and real-time mode. The real-time analytics pipeline also shows how to detect distributed denial of service (DDoS) and bot attacks, which is a common requirement for such use cases.			
BOA310	Building observable applications with OpenTelemetry	In this demo-focused session, review how OpenTelemetry works and how you can use it to instrument your applications. The session starts with an existing application running on containers and serverless functions and demos how to make it observable using AWS Distro for OpenTelemetry to collect telemetry data such as logs, metrics, and traces. Then, observe a simulation of a few issues to learn how you can use telemetry data for troubleshooting and root cause analysis. Finally, to collect and visualize telemetry data, see how to compare different options using tools like Amazon CloudWatch, AWS X-Ray, Amazon Managed Service for Prometheus, and Amazon Managed Grafana			
BSI205-R1	Build stunning customized dashboards with Amazon QuickSight	Want to grow your dashboard building skills? In this workshop, the Amazon QuickSight team demonstrates the latest authoring functionality designed to empower you to build beautiful layouts and robust interactive experiences with other applications, right from within your dashboard. You must bring your laptop to participate.	Workshop	Reserved	Wynn
BSI207	Get clarity on your data in seconds with Amazon QuickSight Q	Amazon QuickSight Q is a machine learning-powered natural language capability that empowers business users to ask questions about all of their data using everyday business language and get answers in seconds. Q interprets questions to understand their intent and generates an answer instantly in the form of a visual without requiring authors to create graphics, dashboards, or analyses. In this session, the Amazon QuickSight Q team provides an overview and demonstration of Q in action.			

Session	Title	Description	Type	Attending	Location
BWP102	The story of headless ecommerce and the future of online storefronts	Headless services are rapidly changing the ecommerce landscape. Learn how you can leverage headless services and AWS to both optimize your costs and architecture flexibility and build world-class shopping experiences. Also learn about Buy with Prime, a new invite-only service from Amazon that supports ecommerce sellers in offering the Prime shopping experience to their customers beyond Amazon.com.			
BWP201	Build with prototypes: Be a customer-obsessed developer	To succeed in the software as a service market, it is crucial to quickly go to the market and learn what customers love about your product. Prototypes can help you achieve rapid illustration of the final product and validation of your design ideas. Join this session to learn about the Amazon Buy with Prime (BWP) team's journey of building prototypes and shaping product development with a focus on customers. Discover how they were able to develop and deploy prototypes at scale and speed using AWS managed services and AWS CDK infrastructure as code.			
BWP302	How Alloy synced Amazon data in real time	Buy with Prime, one of Amazon's newest services, allows online retailers to take advantage of the benefits of Amazon's massive fulfillment service. In this session, dive deep into how Amazon and Alloy worked together to create a data backfilling and synchronization system using Amazon EventBridge, AWS Lambda, Amazon API Gateway, and the Alloy Automation platform.			
CMP225	What's new in Amazon EC2	Amazon EC2 provides secure, resizable compute capacity in the cloud and makes web-scale computing easier. It is a foundational service for AWS and offers a wide variety of compute instances that are well suited for virtually every imaginable use case, from static websites to on-demand supercomputing, and it's available with flexible pricing options. This session provides an overview of what's new in the Amazon EC2 portfolio, including updates on capabilities, instance families, storage and networking functionality, and edge and hybrid offerings.	Session	Watch	MGM Grand
CMP310	Best practices for growing your business on Amazon Lightsail	Whether you need a simple website, blog, ecommerce store, or other applications for your business, Amazon Lightsail gives you the tools you need to quickly and easily meet your requirements. In this workshop, launch and configure multiple products similar to those you may use in your own business. Learn about architecture and best practices when running your applications on Lightsail and how to avoid some common pitfalls. Gain confidence in your ability to successfully run your business applications online. You must bring your laptop to participate.	Workshop	Reserved	Wynn
COM207	Competition of the modern workloads: Serverless vs. Kubernetes on AWS	How do you choose the right infrastructure for you? In this session, witness a competition between the serverless and Kubernetes approaches, examining use cases from each speaker's experience. After an overview of each architecture and the AWS services that are a part of it, compare maintenance and compliance, scaling, developer experience, cost, monitoring and logging, and ecosystem. For each category, consider the advantages and disadvantages of each architecture side by side, taking into consideration various aspects and restrictions different companies have when choosing an infrastructure. Vote on which approach wins each round during the session.	Session	Watch	Mandalay

Session	Title	Description	Type	Attending	Location
COM301	Unleash developer productivity with infrastructure from code	A new category of developer tools is emerging that challenges the legacy-centric approach to building modern cloud applications. The complexity of configuring services, implementing best practices, and securing workloads has led to challenges with developer productivity, cost, performance, and time to market. In the current cloud-native landscape, many organizations are left with skill gaps and unrealized cloud initiatives. Infrastructure from code is a new, configuration-free approach that interprets application code then automatically deploys, configures, and optimizes infrastructure to support your workloads. In this session, explore these tools and see how they can be used to build better and faster on AWS.	Session	Watch	Mandalay
CON210	What's new and what's next with Amazon ECS	Amazon ECS is a fully managed container orchestration service that makes it easy to run highly secure, reliable, and scalable containers. The Amazon ECS team continues to innovate for their users, delivering powerful features that deeply integrate with the rest of AWS. Join this session to hear about the latest advancements with Amazon ECS. Discover what's new since last year's launch of Amazon ECS Anywhere, new features of AWS Fargate, and a look ahead at the exciting enhancements to Amazon ECS.	Session	Watch	MGM Grand
CON309	Build your application easily & efficiently with serverless containers	In this session, explore concurrency and scaling across AWS App Runner, AWS Lambda containers, and AWS Fargate. Consider value and performance differences with a focus on helping you make the best platform decisions for your use case.			
CON325-R	Building containers for AWS	You have many choices when deploying containers on AWS. Whether you're building with AWS Lambda, Amazon ECS, or Kubernetes with Amazon EKS, the process starts with building your images. This session breaks down guidance on building with containers into two parts. First, take a look at some popular choices for development tools, and discover the easiest way to get started with a fully open-source development environment. Then, learn about development patterns and hear practical advice for container development on AWS.			
CON326	Are you ready? Essential strategies for Kubernetes adoption	Like any other software infrastructure solution, Kubernetes deployments need planning. In this session, learn about the best strategies and practices for building resilient applications even when the control plane encounters issues.			
CON327	Serving 4.8 billion railway customers a year: The story of Deutsche Bahn	The national railway company of Germany, Deutsche Bahn, is the second-largest transport company in the world. With the Vendo project, Deutsche Bahn is modernizing their core B2B and B2C booking portal that serves 4.8 billion passengers per year. The ambitious project includes 500 project participants across 22 domains and runs hundreds of microservices on Amazon EKS. In this session Ralf Gernhold, CTO at Deutsche Bahn Vertrieb and program manager for Vendo, shares Vendo's journey since 2018. Learn in detail how Deutsche Bahn has used Amazon EKS at scale to implement a platform in a highly available, secure, cost efficient, and resilient way.			
COP215-R1	Visualizing AWS Config and Amazon CloudWatch usage and costs	In this session, explore dashboards that you can deploy into your own account to get a real-time view of some of the typical main contributors to AWS Config and Amazon CloudWatch costs. The dashboards are designed to help you identify high-cost areas and see the impact of any changes made over time. You can deploy the dashboards into your own account and explore how to create and modify them for your own needs.	Chalk Talk	Waitlisted	

Session	Title	Description	Type	Attending	Location
DAT207-R1	Choosing the right databases for your application	In this chalk talk, learn how to evaluate the right database for your application. Often when organizations are moving from monolithic databases they have to evaluate a daunting number of factors to determine the best option for their workloads. Discover the most common and influential qualification criteria that will help you determine if a relational database, non-relational database, or mix of both would be most appropriate for your use case.	Chalk Talk	Reserved	Caesars Forum
DAT211	How Capital One accelerates innovation with AWS databases	Capital One has been a disrupter in the financial services industry since 1994, using technology to transform banking and payments. As part of their digital transformation journey, Capital One adopted several purpose-built AWS databases to help manage data at a large scale and accelerate innovation. Join this session to learn how Capital One uses databases such as Amazon DocumentDB (with MongoDB compatibility), Amazon DynamoDB, and Amazon RDS to scale the digital experiences their customers want. Also learn best practices and migration guidance along the way.			
DAT212	How ADP & Twilio realize business vision with purpose-built databases	In this session, learn how ADP, the largest payroll company in the world, delivered their next-generation products and achieved goals to increase performance and reduce cost with superior scalability using AWS purpose-built databases. ADP will share their best practices for achieving digital transformation to match business ambitions. Learn how Twilio transformed their data infrastructure by using AWS purpose-built databases to successfully address the scalability requirements associated with hypergrowth, helping them to achieve more than 50% YoY growth. Twilio will highlight best practices for designing a scalable storage layer for your application.			
DAT213	How Fidelity Investments and Reltio modernized with Amazon DynamoDB	When two industry leaders started their application modernization journey, both chose Amazon DynamoDB. In this session, learn how their journeys were unique and how they embraced Amazon DynamoDB as their non-relational, serverless database solution. Learn how Fidelity modernized their live trade execution and order routing workload from Oracle to DynamoDB, and increased resiliency through a multi-Region architecture with global tables. Discover how Reltio, a cloud-native real-time master data management SaaS company, adopted Amazon DynamoDB to meet growth demands for a subset of their microservices, followed by the migration of all relevant workloads from Apache Cassandra.			
DAT217	Modernize apps with purpose-built databases, featuring Liberty Mutual	Modernizing your applications is more than moving them to the cloud—that is simply the first step. By moving your applications to the cloud, you gain access to the latest database technologies, allowing you to modernize your applications to better serve your customers, stakeholders, or constituents. In this session, learn how to use open-source relational databases and new technologies such as key-value, document, and graphs to enhance performance of your applications. Learn why Liberty Mutual adopted AWS services and how they will reduce future total cost of operations while positively impacting the environment by shifting to AWS's more energy-efficient data centers.			

Session	Title	Description	Type	Attending	Location
DAT218	What's new with database migrations	The benefits of migrating applications and databases to the cloud are well known, but knowing how to get started can be difficult. In this session, learn about AWS migration services and tools, including the latest innovations, to get you started on your modernization journey. Learn about best practices, see examples, and watch demos using AWS DMS and AWS SCT, covering everything from simple lift-and-shift migrations to ETL conversions and data lake feeds. Join this session to discover tools to start your next migration with AWS.	Session	Watch	Wynn
DAT219	What's new with Amazon DocumentDB	Amazon DocumentDB (with MongoDB compatibility) is a fully managed, native JSON database designed for scaling enterprise workloads. Amazon DocumentDB provides you the flexibility of a document database without the time-consuming database administration tasks, freeing you to focus on innovation. In this session, learn about the latest Amazon DocumentDB innovations, features, and updates that will push past the scalability limits seen in traditional databases.	Session	Watch	Caesars Forum
DAT221	Dive deep on AWS managed PostgreSQL databases	In recent years, PostgreSQL has become a top database choice for new application development and migration from other commercial databases. This session covers recent innovations for Amazon RDS for PostgreSQL and Amazon Aurora PostgreSQL-Compatible Edition. In this session, review general guidelines for common user operations and activities, such as performance and security, for PostgreSQL instances. Join this session to dive deep into the capabilities of the service and review the latest available features, which give you the background to solve different technical challenges.			
DAT319	S&P Dow Jones Indices' AWS journey to Amazon Aurora	In this session, learn how migrating to Amazon Aurora PostgreSQL-Compatible Edition and Amazon Aurora MySQL-Compatible Edition helped S&P Dow Jones Indices (SPDJI) scale, accelerate, and modernize their workloads. Discover the challenges SPDJI encountered during the migration journey, how they were able to solve them, and the best practices they developed along the way. Learn how using AWS SCT and AWS DMS can streamline and accelerate migrations and improve migration outcomes			
DAT320	Deploy modern and effective data models with Amazon DynamoDB	Amazon DynamoDB is a fully managed NoSQL database that provides consistent performance at any scale. Customers like Disney+, Zoom, and Snap use DynamoDB to handle some of the largest applications on the planet. In this session, AWS Data Hero Alex DeBrie and DynamoDB Senior Principal Engineer Amrith Kumar walk you through key data modeling concepts for DynamoDB and share why DynamoDB architecture and implementation can help your applications scale seamlessly from 30 to 300 million customers while maintaining consistent, single-digit millisecond performance.			

Session	Title	Description	Type	Attending	Location
DAT325	Managed Oracle Databases with Amazon RDS: New features & best practices	Many organizations are looking to take advantage of the power of managed databases in the cloud. Amazon RDS for Oracle is a fully managed relational database service that automates time-consuming database administration tasks while cost-effectively running Oracle Database workloads. Amazon RDS Custom for Oracle offers a shared responsibility model ideal for custom, packaged, and third-party applications that require customization of the underlying operating system and database environment. In this session, learn about new features and best practices to make Oracle Database deployments highly available, scalable, secure, and manageable. Discover how companies, like Jefferies, are transforming their businesses with Amazon RDS.			
DAT326	Deep dive into Amazon Aurora and its innovations	With an innovative architecture that decouples compute from storage as well as offering advanced features like Global Database and low-latency read replicas, Amazon Aurora reimagines what it means to be a relational database. The result is a modern database service that offers performance and high availability at scale, fully open-source MySQL- and PostgreSQL-compatible editions, and a range of developer tools for building serverless and machine learning-driven applications. In this session, dive deep into some of the most exciting features Aurora offers, including Aurora Serverless v2 and Global Database. Also learn about the latest performance and availability innovations, which give you the background to solve technical and operational challenges.			
DOP201	AWS infrastructure as code: A year in review	Join this session to learn about the new features and improvements for AWS infrastructure as code with AWS CloudFormation and AWS CDK.	Session	Watch	Caesars Forum
DOP202	Building on AWS with AWS tools, services, and SDKs	AWS provides a wide range of services, patterns, and tools to empower developers and accelerate software development. In this session, learn the ways developers can set up secure development environments and use their favorite IDEs to interact with, and deploy to, the AWS Cloud. Then, dive into the AWS CLI tools and the SDKs we offer for popular programming languages. Also see how organizations can provide dev accounts to their teams securely and at scale.	Session	Watch	Mandalay
DOP205	Build it & run it: Streamline your DevOps capabilities with machine learning	While organizations have vastly improved how they deliver and operate software, development teams still run into issues when performing manual code reviews, looking for hard-to-find defects, and uncovering security-related problems. Developers have to keep up with multiple programming languages and frameworks and often have their productivity impaired when they have to search online for code snippets. Additionally, developers now require expertise in observability to successfully operate the applications they build. Join this session to learn how to use machine learning-powered tools like Amazon CodeWhisperer, Amazon CodeGuru, and Amazon DevOps Guru to boost your applications' availability and write software faster and more reliably.	Session	Watch	Mandalay
DOP206	Continuous integration and delivery (CI/CD) on AWS	With its continuous integration/continuous delivery (CI/CD) tools, AWS provides one place where you can plan work; collaborate on code; and build, test, and deploy applications. In this session, learn how setting up a CI/CD pipeline on AWS can help automate steps in your software delivery process	Session	Watch	Venetian

Session	Title	Description	Type	Attending	Location
DOP207	AWS Well-Architected best practices for DevOps on AWS	In this session, learn about all of the components required to align your DevOps practices to the pillars of the AWS Well-Architected Framework. Review organization adoption, development lifecycle, quality assurance, automated governance, and observability.			
DOP301	The Amazon Builders' Library: 25 years of operational excellence at Amazon	Amazon has practiced DevOps, where builders build but also test and operate almost every aspect of their services, since before it had a name. In this session, hear how Amazon practices have changed and improved over time and what we've learned as builders and as operators. What has it been like to be on call at Amazon? And what's it like now to be on a call at Amazon? We'll share real-world experience and answers.			
DOP302	How to reuse patterns when developing infrastructure as code	Explore AWS Cloud Development Kit (AWS CDK) constructs and AWS CloudFormation modules and how they can help make it easier to build applications on AWS. Learn best practices from United Airlines, who shares how they have used AWS CDK to create an internal library of infrastructure components that their development team can reuse to accelerate application delivery.			
DOP314	Governance and security with infrastructure as code	In this session, learn how to use AWS CloudFormation and the AWS CDK to deploy cloud applications in regulated environments while enforcing security controls. Find out how to catch issues early with cdk-nag, validate your pipelines with cfn-guard, and protect your accounts from unintended changes with CloudFormation hooks.			
DOP317	Accessing AWS from your development environment	Many developers find it more productive to work in an integrated development environment (IDE). AWS provides all of the tools you need to code, test, and debug for virtually any project. As a result, you can quickly set up a cloud development environment, switch between projects, and replicate the development workspace configuration across team members. In this session, learn how to use your development environment to work on code stored in your project's source repositories with a supported IDE.			
ENT213	How to migrate, modernize, and grow using the AWS MAP	AWS has helped hundreds of thousands of businesses migrate and modernize their digital operations through the AWS Migration Acceleration Program (MAP), a proven framework for cost-efficient, resilient, and agile transformation. In this session, Faraz Shafiq, AWS Global Head for MAP, shares what makes the MAP so successful—specifically, how it's designed to support businesses not only as they migrate to AWS but also beyond, as they embrace new technologies, modernize digital operations, and look toward the future. Krishna Srinivasan, Director of Cloud Architecture and DevOps at United Airlines, joins Faraz to share how the MAP continues to help one of the world's largest airlines better serve their customers.			
ENT221	Building a smarter organization powered by machine learning	Every decade comes with the introduction of technology that is "going to change the world," but this leads to disappointment when it becomes too challenging to adopt the new technology. For several decades, this has been the story behind artificial intelligence (AI) and machine learning (ML), and yet we've reached a tipping point where some organizations are transforming their businesses and others are struggling to adapt. In this session tailored for business leaders, AWS enterprise strategists discuss how organizations can embrace AI and ML. They also share stories on how some organizations are driving profound societal and business results using AI and ML.			

Session	Title	Description	Type	Attending	Location
ENT311	Migrate and modernize Oracle E-Business Suite on AWS in 5 steps	In this chalk talk, experts walk you through 5 steps to migrate your Oracle E-Business Suite to AWS. Learn how to modernize the underlying database to Amazon RDS Custom, and discover strategies to optimize E-Business Suite performance, scale, and RTO and RPO objectives. Also learn about the architectural patterns for enhanced database performance, increased availability, and disaster recovery across single-Availability Zone (AZ) and multi-AZ environments.			
FSI201	Goldman Sachs: Accelerating time to value in data analytics	Financial institutions rely on data to create new products, manage risk, and make informed business decisions. However, it can be challenging and time-consuming to access data that is dispersed across on-premises and public cloud data stores and distributed by third-party vendors. In this session, learn how Goldman Sachs engineered its Goldman Sachs Financial Cloud for Data—a cloud-native, modular, and integrated collection of services—to run distributed server-side analytics and enrich data in real time and stream and analyze time-series data (also in real time) by ingesting relational data using AWS Data Exchange, Amazon Redshift, AWS Glue, and the FINOS Legend open-source platform.			
FSI304-R1	Your blueprint to secure customer-facing applications on AWS	When deploying public endpoints, financial services organizations require advanced firewall, inspection, and detection capabilities. This chalk talk walks you through some of the most common patterns that our customers with the largest and most sensitive workloads use to protect their customer-facing applications. For this talk, you should have a foundational understanding of Amazon VPC concepts and traditional network security patterns.			
FWM308	What's new for frontend web and mobile developers with AWS Amplify	AWS Amplify helps frontend web and mobile developers build full-stack applications in hours with minimal cloud expertise. In this session, explore the capabilities of Amplify, including how you can easily configure a backend with authentication, data, and storage; create a frontend UI with Amplify Studio; and host static and server-side rendered web apps with Amplify Hosting. Learn about recent Amplify updates, including Next.js support; Amplify Flutter support for web, desktop, and iOS/Android; and extensibility features that let you use the breadth of 175+ AWS services with AWS CDK and modify deployment operations as your needs evolve.	Session	Reserved	Encore
FWM309	Design to code: Build an app frontend UI in minutes with AWS Amplify Studio	Frontend developers use AWS Amplify to quickly build and deploy web and mobile apps on AWS with a common backend. Amplify's zero-configuration workflows help developers get started on AWS faster because they can easily configure robust backend capabilities for their apps. Backends can be extended to add capabilities from 175+ AWS services as use cases expand, and Amplify Studio extends these benefits to UI development. With Amplify Studio, developers can now visually build their apps front to back, from exporting prototypes in Figma to frontend code, to configuring the app backend and connecting them together. Developing full-stack apps on AWS has never been easier, and this session shows you how.	Session	Watch	MGM Grand

Session	Title	Description	Type	Attending	Location
FWM310	Building a multi-Region serverless application with AWS AppSync	As organizations grow, they often need to serve geographically dispersed users with low latency, prompting them to have a distributed global infrastructure in the cloud. AWS AppSync is a serverless GraphQL and pub/sub API service that simplifies building modern web and mobile applications, allowing you to optimize data access with fewer network calls and built-in real-time features. In this session, discover how you can connect to a GraphQL API endpoint closest to you, automatically access the same data replicated between AWS Regions, and eliminate infrastructure management tasks so you can focus on writing application code that serves your customers.			
FWM315	Create real-time, event-driven apps with Amazon EventBridge & AWS AppSync	Today, real-time digital experiences (such as real-time sports scores and financial portfolio updates, interactive learning, location tracking, and chat capabilities) are ubiquitous, and consumers have come to expect them. Join this session to learn how event-driven and publish/subscribe (pub/sub) architectures, implemented with Amazon EventBridge and serverless WebSockets from AWS AppSync, help application developers use real-time data to create differentiated digital experiences for their customers.	Session	Watch	Wynn
FWM316	Accelerate GraphQL API app development & collaboration with AWS AppSync	Customers use AWS AppSync to build serverless GraphQL APIs because it helps them develop applications faster by giving developers the ability to access data from multiple backend services in a single endpoint with fewer network calls. Join this session to learn about exciting new AWS AppSync features that make it even easier for API developers to realize the benefits of GraphQL for their application development.			
HYB311	How to simplify your hybrid cloud migration strategy with AWS	Complex legacy applications can be difficult to lift and shift, or to refactor and replatform, without disrupting business-critical operations. Traditionally, organizations have kept these applications on premises, and they miss out on the benefits of the cloud. AWS hybrid cloud solutions help solve this problem by extending AWS infrastructure and services to wherever applications need to run: from the cloud to on premises and at the edge in locations such as large metropolitan areas. In this session, dive deep into how AWS hybrid cloud and edge computing services, AWS Local Zones, AWS Outposts, and the AWS Snow Family are helping organizations simplify their hybrid cloud migrations.			
KEY001	Monday Night Live with Peter DeSantis		Keynote	Watch	Venetian
KEY002	Adam Selipsky Keynote		Keynote	Watch	Venetian
KEY003	Swami Sivasubramanian Keynote		Keynote	Watch	Venetian
KEY004	AWS Partner Keynote with Ruba Borno		Keynote	Watch	Venetian
KEY005	Dr. Werner Vogels Keynote		Keynote	Watch	Venetian
LDR001	Executive Cloud Insights		Leadershi	Watch	Venetian
LDR002	AWS Storage		Leadershi	Watch	Venetian
LDR003	AI/ML		Leadershi	Watch	Venetian
LDR004	Analytics		Leadershi	Watch	Venetian
LDR005	Cloud Operations		Leadershi	Watch	Venetian

Session	Title	Description	Type	Attending	Location
LDR006	Compute		Leadershi	Watch	Venetian
LDR007	Developer		Leadershi	Watch	Venetian
LDR008	Hybrid Cloud		Leadershi	Watch	Venetian
LDR009	Kubernetes		Leadershi	Watch	Venetian
LDR010	Serverless		Leadershi	Watch	Venetian
OPN306	AWS Lambda Powertools: Lessons from the road to 10 million downloads	AWS Lambda Powertools is an open-source library to help organizations discover and incorporate serverless best practices early and quickly. In two years, Powertools went from a tiny pilot program to a fast-growing project. This rapid growth led to challenges ranging from balancing new features with operational excellence, triaging bug reports and RFCs, and scaling and redesigning documentation, to lowering the bar for contribution and providing a public road map. In this session, learn about the current state of Lambda Powertools, how this growth was supported, key lessons learned in the past two years, and what's next on the horizon.			
PRT222	Why operationalizing data mesh is critical for operating in the cloud (sponsored by Capital One)	As companies look to scale in the cloud, they face new and unique challenges related to data management. Data mesh offers a framework and a set of principles that companies can adopt to help them scale a well-managed cloud data ecosystem. In this session, learn how Capital One approached scaling its data ecosystem by federating data governance responsibility to data product owners within their lines of business. Also hear how companies can operate more efficiently by combining centralized tooling and policy with federated data management responsibility. This presentation is brought to you by Capital One, an AWS Partner.			
PRT303	The digital transformation journey: Santander UK	Join this session to discover how Santander is powering its digital transformation for both customers and colleagues by using AWS as the foundation for building an omnichannel API platform, contact center replacement, and Customer 360, a data lake architecture. This presentation is brought to you by Salesforce, an AWS Partner.			
SAS305-R	SaaS architecture patterns: From concept to implementation	SaaS architecture varies across domains, stacks, and customer requirements. However, there are well-defined patterns that must be addressed by every SaaS solution. Identity, onboarding, tenant isolation, data partitioning, tiering—these are among the patterns that, while implemented differently on AWS stacks and services, are core to most SaaS solutions. In this session, take a detailed look at the landscape of these SaaS patterns. For each one, move from concept to implementation, looking at the permutations of architecture and code that are used to bring these patterns to life with different compute, storage, identity, networking, and management constructs.			
STG303-R1	AWS storage for containers and serverless applications	Modern applications have revolutionized application architecture. Builders can now use containers for greater portability and/or building applications on serverless infrastructure to optimize compute resources. In this session, learn about different storage services, including Amazon EFS, Amazon S3, and Amazon EBS, to help you determine which storage is the best fit for your modern applications.			

Session	Title	Description	Type	Attending	Location
SVS204	What's new with serverless	Join this session to learn what new features in AWS Lambda, Amazon EventBridge, AWS Step Functions, and AWS SAM you might have missed this year. Explore the latest patterns and what has been made possible. Learn practical tips to make use of the latest features designed to help you build better serverless applications.	Session	Watch	Mandalay
SVS209	Get started building your first serverless, event-driven application	Emily Shea, Head of Application Integration Go-To-Market at AWS, shares firsthand experience and practical steps for getting started building serverless web applications. Emily wanted to build a solution to an everyday problem: regularly practicing a second language. This session shares the process of evolving that solution from an AWS Lambda function into a fully featured application that uses Amazon API Gateway, AWS Lambda, Amazon EventBridge, AWS Step Functions, Amazon DynamoDB, Amazon QuickSight, the AWS Amplify Console, and more. Learn why serverless is a great tool set for experimenting with new ideas, and hear how the extensibility and modularity of serverless applications allow you to start small and quickly add features.			
SVS302-R	Architecting secure serverless applications	This session explores how to think about security from the front to the back of a typical serverless application. How do you configure AWS serverless services to provide least-privileged access while helping to ensure functionality? How should you think about managing IAM policies for your AWS Lambda functions? This session covers all of this and more, leaving you with concrete examples applicable to almost any workload.			
SVS303	Building real-world serverless applications with AWS SAM	Treating your serverless infrastructure and application as code can improve efficiency, build in flexibility, and save time. In this session, learn about AWS SAM, an open-source framework for building serverless applications, and discover how to use AWS SAM templates to manage serverless infrastructure as code. From serverless-specific resources to pseudo parameters, learn how to construct highly flexible and reusable AWS SAM templates to manage your serverless applications. Find out how to use AWS SAM CLI to create, develop, and deploy serverless applications. Additionally, learn about the latest AWS SAM CLI features that can help you build deployment pipelines and test your code in the cloud.	Session	Watch	Wynn
SVS308-R1	AWS serverless developer experience: A day in the life of a developer	AWS serverless developer experience: A day in the life of a developer	Workshop	Reserved	Wynn
SVS311	Practical experience with a serverless-first strategy at Capital One	In this session, Capital One shares concrete examples of how they have implemented serverless patterns, challenges they have faced, and the best practices and tips they have learned from operating AWS Lambda at scale in production. Dive deep on scaling Lambda with various event sources, considerations of event source integrations, observability at scale, and failure-handling mechanisms.			

Session	Title	Description	Type	Attending	Location
SVS312	Building Serverlesspresso: Creating event-driven architectures	Serverlesspresso is an event-driven, serverless workload that uses Amazon EventBridge and AWS Step Functions to coordinate events across microservices and support thousands of orders per day. This session explores the design decisions made in building this application, how new features influenced the development process, and lessons learned in creating a production-ready application using this approach. Explore useful patterns and options for extensibility that helped the team design a robust, scalable solution that costs about 1 dollar a day to operate. This session includes examples you can apply to your serverless applications and teaches you how to combine EventBridge and Step Functions to solve complex architectural challenges in larger applications.			
SVS401-R1	Best practices for advanced serverless developers	Are you an experienced serverless developer? Do you want a helpful guide for unleashing the full power of serverless architectures for your production workloads? Are you wondering whether to choose a stream or an API as your event source, or whether to have one function or many? This session provides architectural best practices, optimizations, and useful cheat codes that you can use to build secure, high-scale, and high-performance serverless applications and uses real customer scenarios to illustrate the benefits.			
SVS404-R1	A closer look at AWS Lambda	Serverless computing allows you to build and run applications and services without needing to provision, scale, or manage any servers. When you take a closer look at a serverless application, there is a sophisticated architecture that takes care of all the undifferentiated heavy lifting for the developer. Join this session to learn how AWS architected AWS Lambda. Explore how Lambda takes care of virtually everything required to run and scale your code in a reliable, highly available manner.			
TNC108	Introduction to serverless computing and AWS Lambda	Serverless computing allows you to build and run applications and services without relying on servers. In this session, explore the benefits of serverless computing and learn the basics of AWS Lambda.			
TNC306	Build a notification system with Amazon Translate & Amazon Pinpoint	In this lab scenario, you need to communicate with customers who use your application globally to inform them about announcements. The communication should be customized by a customer's preferred language and should be delivered as either an email, text message, or voicemail. Your task is to build an application using serverless technologies, Amazon Translate, and Amazon Pinpoint that allows you to enter data, translate it to a specific language, and deliver it to your customers in their preferred formats and languages.	Lab	Waitlisted	
XNT303	Building AI-enriched applications with .NET on AWS	With AWS AI services, you can add capabilities like image and video analysis, natural language processing, personalized recommendations and translation, virtual assistants, and speech recognition. You can use each service standalone, or you can use them together to build a sophisticated AI-capable application. In this demo-heavy session, discover how to iteratively enrich your existing .NET applications with AI capabilities using AWS SDK for .NET.			

Session	Title	Description	Type	Attending	Location
XNT304	Purpose-built databases for modern .NET applications on AWS	Previously, the only database choice that .NET developers had was either Microsoft SQL Server or Oracle. Today, developers are moving away from one-size-fits-all monolithic databases and are increasingly building modern cloud-native applications using multiple purpose-built databases. AWS offers a choice of relational, key-value, wide-column, document, in-memory, graph, time-series, and ledger databases, with each purpose-built database solving a specific problem or usage scenario. In this chalk talk, learn about AWS purpose-built databases for .NET developers that meet the scale, performance, and manageability requirements of modern applications built for the cloud.			
XNT402	Cloud-native .NET applications on AWS	Does your future include building cloud-native applications? In this session, learn how to design well-architected, modern .NET applications. This includes building secure, scalable, low-cost serverless applications with tracing, debugging, testing, monitoring, CI/CD pipelines, and purpose-built databases. Review a sample application, including full source code in GitHub, that showcases a reference architecture and how to design a .NET application on AWS. This sample code includes Blazor, .NET Core, Amazon Cognito, Amazon ECS, Amazon API Gateway WebSocket APIs, AWS Step Functions with AWS Lambda, and Amazon DynamoDB. Leave this session prepared for your future building modern cloud-native .NET applications on AWS.			