

AWS Announces Three New Analytics Capabilities

AQUA (Advanced Query Accelerator) for Amazon Redshift provides an innovative new hardware-accelerated cache that delivers up to 10x better query performance than any other cloud data warehouse

AWS Glue Elastic Views lets developers easily build materialized views that automatically combine and replicate data across multiple data stores

Amazon QuickSight Q is a machine learning-powered capability for Amazon QuickSight that lets users type questions about their business data in natural language and receive highly accurate answers in seconds

Capital One, Best Western Hotels & Resorts, NTT DOCOMO, Audible, and Panasonic Avionics among customers using new analytics capabilities

Seattle—Dec. 1, 2020—Today at AWS re:Invent, Amazon Web Services, Inc. (AWS), an Amazon.com, Inc. company (NASDAQ: AMZN), announced three new analytics capabilities that dramatically improve the performance of Amazon Redshift data warehouses, make it significantly easier for customers to move and combine data across data stores, and make it much simpler for end-users to get more value from their business data using machine learning.

- **AQUA for Amazon Redshift** accelerates querying with an innovative new hardware-accelerated cache that brings the compute to the storage and delivers up to 10x better query performance than any other cloud data warehouse, with general availability coming in January 2021.
- **AWS Glue Elastic Views** helps developers build applications that use data from multiple data stores with materialized views that automatically combine and replicate data across storage, data warehouses, and databases.
- **Amazon QuickSight Q** delivers a machine learning-powered capability for Amazon QuickSight that gives users the ability to use natural language expressions to ask business questions in the Amazon QuickSight Q search bar and receive highly accurate answers in seconds.

More data is created every hour today than in an entire year just 20 years ago. In fact, the amount of data created over the next three years will be more than the amount of data created over the past 30 years. The same old tools simply won't work in this new world of data. AWS customers use a wide variety of analytics tools for different use cases, including Amazon Athena for serverless querying, Amazon Elasticsearch Service for searching and visualizing log data, Amazon Kinesis for processing real-time data streams, Amazon Redshift for data warehousing, and Amazon EMR for running Apache Spark, Hive, Presto, and other big data frameworks. These services offer AWS customers the right tool for their needs. The new analytics capabilities announced today build on this foundation and provide faster, more cost-effective, and more accessible data analysis across all of a customer's data stores. To learn more, visit <https://aws.amazon.com/big-data/datalakes-and-analytics/>

“With the capabilities we're announcing today, we're delivering an order-of-magnitude performance improvement for Amazon Redshift, new flexible ways to more easily move data between data stores, and the ability for customers to ask natural language questions in their business dashboards and receive answers in seconds,” said Rahul Pathak, VP, Analytics, AWS. “These capabilities will meaningfully change the speed and ease of use with which customers can get value from their data at any scale.”

AQUA (Advanced Query Accelerator) for Amazon Redshift brings compute to the storage layer, delivering up to 10x faster query performance than any other cloud data warehouse

Since its launch in 2012 as the first data warehouse built for the cloud at a cost of 1/10th that of traditional data warehouses, Amazon Redshift has become the most popular cloud data warehouse. Earlier this year, AWS announced the general availability of Amazon Redshift RA3 instances, which allow customers to scale compute and storage separately and deliver up to 3x better performance than any other cloud data warehouse. However, even with the advantages offered by RA3 instances, the rapid growth of data that customers need to process in their data warehouses has led to a difficult balancing act between performance and cost-effective scaling. The prevailing approach to data warehousing has been to build out an architecture whereby large amounts of centralized storage are moved to waiting compute nodes to process the data. The challenge with this approach is that there is a lot of data movement between the shared data and compute nodes. As data volumes continue to grow at a rapid clip, this data movement saturates available networking bandwidth and slows down performance. In addition to the networking bottleneck, CPUs are not able to keep up with the faster growth in storage capabilities (SSD storage throughput has grown 6x faster than the ability of CPUs to process data from memory), which either creates a new CPU bottleneck of its own or forces more customers to over-provision compute to get their work done more quickly.

AQUA for Amazon Redshift is a distributed and hardware-accelerated cache for Amazon Redshift; an innovation that improves performance for analytics at the new scale of data. AQUA brings compute to the storage layer, so data does not have to move back and forth between the two. This enables Amazon Redshift to run up to ten times faster than any other cloud data warehouse. The AQUA cache scales out and processes data in parallel across many nodes. Each node possesses a hardware module composed of AWS-designed analytics processors that dramatically accelerate data compression, encryption, and data processing tasks like scans, aggregates, and filtering. AQUA also gives customers the added benefit of being able to do compute on their raw storage, which saves time that would otherwise be spent moving data around. With this new architecture, and the order-of-magnitude better performance it brings, Redshift customers will have more up-to-date dashboards, save development time, and their systems will be easier to maintain. AQUA's preview is now open to all customers, and AQUA will be generally available in January 2021. AQUA is available on Redshift RA3 instances at no additional cost, and customers can take advantage of the AQUA performance improvements without any code changes. To get started with AQUA, visit https://pages.awscloud.com/AQUA_Preview.html

AWS Glue Elastic Views lets developers easily build materialized views that automatically combine and replicate data across multiple data stores

Most companies are building or have already built data lakes, where they can aggregate all of the data from various silos with the right security and access controls, to make it easier to do analytics and machine learning. But for latency and operational reasons, most companies are also likely to have increasing amounts of data in purpose-built data stores outside of their data lakes. As the data in these data lakes and purpose-built data stores continue to grow, companies need an easier way to move data around.

AWS Glue Elastic Views provides developers with a new capability to easily build materialized views (also called virtual tables) that automatically combine and replicate data across multiple data stores. AWS Glue is a serverless data preparation service that makes it easy to run extract, transform, and load (ETL) jobs for analytics and machine learning. With AWS Glue Elastic Views customers can use SQL to create a materialized view of the data they want to combine from different data stores, and AWS Glue Elastic Views copies the data to create the materialized view from the different sources. For example, a

customer might create a materialized view that pulls restaurant location information from Amazon Aurora and combines it with customer reviews stored in Amazon DynamoDB to build a search engine for restaurant reviews by location on Amazon Elasticsearch Service. AWS Glue Elastic Views copies data from each source database to a target database, and automatically keeps the data in the target database up to date. Elastic Views continually monitors the source database for changes, and updates the target database within seconds. If there is a change to the data model in one of the source databases, Elastic Views proactively alerts the developers, so they can update their materialized view to adapt to the change. Customers can also use Elastic Views to copy operational data from an operational database to their data lake to run analytics in near real-time. AWS Glue Elastic Views automatically scales capacity to accommodate workloads as they ramp up or down, ensuring that the materialized views in the target databases are kept up to date. AWS Glue Elastic Views is available in preview today. To learn more, visit <http://aws.amazon.com/glue/features/elastic-views>

Amazon QuickSight Q is a machine learning-powered capability for Amazon QuickSight that lets users type natural language questions about their business data and receive highly accurate answers in seconds

Amazon QuickSight is a scalable, serverless, embeddable machine learning-powered business intelligence (BI) service built for the cloud. Amazon QuickSight provides all the benefits of a modern, interactive, self-service BI solution with capabilities that make it easy to embed dashboards in applications and cost-effectively scale to support thousands of customers. Amazon QuickSight's 'Auto-Narratives' feature provides customers with an automatically generated summary in plain language that interprets and describes what the data in a BI dashboard means, so all users have a shared understanding of the data. Customers like these human-readable narratives because it enables them to quickly interpret the data in a shared dashboard and focus on the insights that matter most. Customers have also been interested in asking business questions of their data in plain language and receiving answers in near real-time. While some BI tools and vendors have attempted to solve this challenge with Natural Language Query (NLQ), the existing approaches require that customers first spend months in advance preparing and building a model, and even then, they still have no way of asking questions that require new calculations that are not pre-defined in the data model. For example, the question, "What is our year-over-year growth rate?" requires that 'growth rate' be pre-defined as a calculation in the model. With today's BI tools, users need to work with their BI teams to update the model to account for any new calculation or data, which can take days or weeks of effort.

Amazon QuickSight Q gives users the ability to ask any question of all their data in natural language and receive an answer in seconds. To ask a question, users simply type it into the Amazon QuickSight Q search bar. As users begin typing their questions, Amazon QuickSight Q provides auto-complete suggestions with key phrases and business terms, and automatically performs spell-check and acronym and synonym matching, so users do not have to worry about typos or remembering the exact business terms for the data. Amazon QuickSight Q uses deep learning and machine learning (natural language processing, schema understanding, and semantic parsing for SQL code generation) to generate a data model that automatically understands the meaning of and relationships between business data, so users receive highly accurate answers to their business questions and do not have to wait days or weeks for a data model to be built. Because Amazon QuickSight Q eliminates the need for BI teams to build a data model, users are also not limited to asking only a specific set of questions. Furthermore, users can get more complete and accurate answers because the query is applied to all of the data, not just the datasets in a pre-determined model. Amazon QuickSight Q comes pre-trained on data from various domains and industries like sales, marketing, operations, retail, human resources, pharmaceuticals, insurance, energy, and more, so it is optimized to also understand complex business language. For

example, sales users can ask *“how are my sales tracking against quota,”* or retail users can ask *“what are the top products sold week-over-week by region?”* Amazon QuickSight Q continually improves its accuracy over time by learning from user interactions. If Amazon QuickSight Q does not understand a phrase in a question, users are prompted to select from a drop-down menu of suggested options in the search bar and Amazon QuickSight Q remembers the phrase for the next interaction. To learn more about Amazon QuickSight Q, visit <https://aws.amazon.com/quicksight/q>

Tokyo-based NTT DOCOMO is the largest mobile service provider in Japan, serving more than 80 million customers. *“Since migrating to Amazon Redshift in 2014, Amazon Redshift has been the center of our analytics environment and has allowed us to scale to over ten petabytes of uncompressed data with a 10x performance improvement over our prior on-premises system,”* said Ken Ohta, General Manager of Service Innovation Department, NTT DOCOMO. *“As customer demand for data and data volumes grow, continuous innovation from Amazon Redshift has helped us with the flexibility and ease of use needed to scale our systems. We are excited about the launch of AQUA for Amazon Redshift as we continue to increase the performance and scale of our Amazon Redshift data warehouse.”*

Intercom is a fast-growing start up with a \$1.3 billion valuation and over \$240 million in funding. *“Strong customer relationships are more important than ever, but the scale and nature of online business can make it hard to create personal connections. That’s why we created the world’s first Conversational Relationship Platform to help businesses build better customer relationships through personalized, messenger-based experiences. To make this work well, and understand our business as it explodes, we rely on an enormous amount of data—70 terabytes and growing,”* said Paul Vickers, Data Engineering Manager, Intercom. *“Our Amazon Redshift cloud data warehouse has made it easy to scale and stay on budget. We’re excited about the new AQUA capabilities in Amazon Redshift which will accelerate our queries and reduce our analysts’ time to insights. We know with AWS we can focus on our growth, without worrying about how technology will support it.”*

Accenture is a global professional services company with leading capabilities in digital, cloud and security. *“At Accenture we are committed to providing services and solutions that help customers around the world use data for real-time decision making. However, as data and the demand for insight grows at an incredible pace, it can be challenging to define, prioritize, and process the data,”* said A.K. Radhakrishnan, North American Data & AI AWS Lead, Accenture. *“AQUA for Amazon Redshift provides an innovative new way to approach data warehousing with up to 10x faster query performance. This makes it easier for us to support the goal of a data-driven enterprise.”*

ZS Associates is a professional services firm that works side-by-side with companies to help develop and deliver products that drive customer value and company results. *“AWS has always been at the forefront of innovation and is known for bringing best-in-class solutions to help its customers. Using AWS’s next-gen technologies and ZS’s deep technical as well as domain expertise, we have deployed several large scale data and analytics platforms on Amazon Redshift for customers,”* said Nishesh Aggarwal, Enterprise Architecture Lead, ZS Associates. *“With the introduction of RA3 instances for Amazon Redshift we were able to significantly improve the performance of analytical workloads while solving the data storage issue at the same time. We are really excited to explore AQUA for Amazon Redshift as it promises to further improve the performance of our most complex workloads by around 10x with no additional effort.”*

Sisense is an independent analytics platform that enables more than 2,000 customers worldwide to simplify complex data, and build and embed analytic apps. *“The strong collaboration between Sisense*

and Amazon Redshift results in an improved cloud analytics experience for our many joint customers,” said Guy Levy-Yurista, Chief Strategy Officer, Sisense. “With AQUA, we’re expecting performance boosts of up to 10x, allowing customers to optimize their Redshift data clusters. These in turn will empower our customers to quickly turn data into insights and infuse intelligence throughout their business.”

Audible is the leading producer and provider of original spoken-word entertainment and audiobooks, enriching the lives of millions of listeners every day. “At Audible, customers can search and discover original spoken-word entertainment and audiobooks across multiple categories. To power this experience, we need to quickly analyze data from a number of databases to deliver personalized results,” said Shailesh Vyas, Principal Software Development Engineer, Audible. “We look forward to trying AWS Glue Elastic Views as a serverless solution to create materialized views from data across multiple different databases in our environment. With AWS Glue Elastic Views, our developers should be able to move faster and focus more on innovating on behalf of our customers versus managing complex data integration pipelines.”

Best Western Hotels & Resorts, headquartered in Phoenix, Arizona, is a privately held hotel brand with a global network of approximately 4,700 hotels in over 100 countries and territories worldwide. Best Western offers 18 hotel brands to suit the needs of developers and guests in every market. “Amazon QuickSight’s pay-per-use pricing and serverless architecture enabled Best Western’s lean analytics team to be agile and deliver increased value to the business, faster, and at less than half the cost of our previous analytics architecture,” said Joseph Landucci, Senior Manager, Database and Enterprise Analytics, Best Western Hotels & Resorts. “With Amazon QuickSight Q, we look forward to enabling our business partners to self-serve their questions while reducing the operational overhead on our team for ad hoc requests. This will allow our partners to get answers to their critical business questions quickly by simply typing their questions in plain language.”

Founded in 1994, Capital One is a leading information-based technology company that is on a mission to help its customers succeed by bringing ingenuity, simplicity, and humanity to banking. “With Amazon QuickSight, we have been able to quickly roll out new machine learning-powered BI dashboards at scale and without any server setup or onerous capacity planning,” said Peter Tyson, Senior Data Engineer, Capital One. “Now, with the launch of Amazon QuickSight Q we look forward to making it easy for our users to quickly get answers to their ad-hoc business questions that aren’t even part of the BI dashboards.”

Panasonic Avionics Corporation is the world’s leading supplier of in-flight entertainment and communication systems. “Our cloud-based solution collects large amounts of anonymized data that help us optimize the experience for both our airline partners and their passengers,” says Anand Desikan, Director of Cloud Operations at Panasonic Avionics Corporation. “We started using Amazon QuickSight to report on in-flight Wi-Fi performance, and with its rich APIs, pay-per-session pricing, and ability to scale, we quickly rolled out Amazon QuickSight dashboards to hundreds of users. The constant evolution of the platform has been impressive: machine learning-powered anomaly detection, Amazon SageMaker integration, embedding, theming, and cross-visual filtering, and now with Amazon QuickSight Q, our users can consume insights by simply typing their business questions in the search bar and Amazon QuickSight Q interprets the business context, provides synonyms, and shows them an answer with no complex interpretation needed.”

Vyaire Medical, Inc., a global company dedicated to respiratory care, enables, improves, and extends lives with an unyielding focus on improving patient outcomes and increasing value for customers. “In less than two months we were able to pivot our old BI reporting tool into Amazon QuickSight,” said

Gopal Ramamurthi, Sr. Director Analytics & Enterprise Data Management, Vyair. “We gained so much in terms of ease of management, especially while scaling the tool to support the increase in number of BI users. Now with the launch of Amazon QuickSight Q, we are looking forward to making it easier for our executive leadership team, sales users in the field, and supervisors in the manufacturing plant to ask their data questions in plain English when the answers are unavailable in the dashboards, providing faster insights that help in making our sales and manufacturing processes even more efficient.”

About Amazon Web Services

For 14 years, Amazon Web Services has been the world’s most comprehensive and broadly adopted cloud platform. AWS offers over 175 fully featured services for compute, storage, databases, networking, analytics, robotics, machine learning and artificial intelligence (AI), Internet of Things (IoT), mobile, security, hybrid, virtual and augmented reality (VR and AR), media, and application development, deployment, and management from 77 Availability Zones (AZs) within 24 geographic regions, with announced plans for 15 more Availability Zones and five more AWS Regions in India, Indonesia, Japan, Spain, and Switzerland. Millions of customers—including the fastest-growing startups, largest enterprises, and leading government agencies—trust AWS to power their infrastructure, become more agile, and lower costs. To learn more about AWS, visit aws.amazon.com.

About Amazon

Amazon is guided by four principles: customer obsession rather than competitor focus, passion for invention, commitment to operational excellence, and long-term thinking. Customer reviews, 1-Click shopping, personalized recommendations, Prime, Fulfillment by Amazon, AWS, Kindle Direct Publishing, Kindle, Fire tablets, Fire TV, Amazon Echo, and Alexa are some of the products and services pioneered by Amazon. For more information, visit www.amazon.com/about and follow [@AmazonNews](https://twitter.com/AmazonNews).

###