AWS Announces Amazon DevOps Guru

New machine learning powered operations service provides tailored recommendations to improve application availability

Atlassian, PagerDuty, and SmugMug among the customers and partners using Amazon DevOps Guru

SEATTLE – **Dec. 1, 2020** – Today at AWS re:Invent, Amazon Web Services, Inc. (AWS), an Amazon.com, Inc. company (NASDAQ: AMZN), announced Amazon DevOps Guru, a fully-managed operations service that uses machine learning to make it easier for developers to improve application availability by automatically detecting operational issues and recommending specific actions for remediation. Amazon DevOps Guru applies machine learning informed by years of Amazon.com and AWS operational excellence to automatically collect and analyze data like application metrics, logs, events, and traces for identifying behaviors that deviate from normal operating patterns (e.g. under provisioned compute capacity, database I/O over-utilization, memory leaks, etc.). When Amazon DevOps Guru identifies anomalous application behavior (e.g. increased latency, error rates, resource constraints, etc.) that could cause potential outages or service disruptions, it alerts developers with issue details (e.g. resources involved, issue timeline, related events, etc.) via Amazon Simple Notification Service (SNS) and partner integrations like Atlassian Opsgenie and PagerDuty to help them quickly understand the potential impact and likely causes of the issue with specific recommendations for remediation. Developers can use remediation suggestions from Amazon DevOps Guru to reduce time to resolution when issues arise and improve application availability and reliability with no manual setup or machine learning expertise required. There are no upfront costs or commitments with Amazon DevOps Guru, and customers pay only for the data Amazon DevOps Guru analyzes. To get started with Amazon DevOps Guru, visit https://aws.amazon.com/devops-guru

As more organizations move to cloud-based application deployment and microservice architectures to globally scale their businesses and operations without the limitations of on-premises deployments, applications have become increasingly distributed to meet customer needs, and developers need more automated practices to maintain application availability and reduce the time and effort spent detecting, debugging, and resolving operational issues. Application downtime events caused by faulty code or config changes, unbalanced container clusters, or resource exhaustion (e.g. CPU, memory, disk, etc.) inevitably lead to bad customer experiences and lost revenue. Companies invest considerable money and developer time to deploy multiple monitoring tools, often managed separately, and then have to develop and maintain custom alerts for common issues like spikes in load balancer errors or drops in application request rates. Setting thresholds to identify and alert when application resources are behaving abnormally is difficult to get right, involves manual setup, and requires thresholds that must be continually updated as application usage changes (e.g. an unusually large numbers of requests during holiday shopping season). If a threshold is set too high, developers don't see

alarms until operational performance is severely impacted. When a threshold is set too low, developers get too many false positives, which ultimately get ignored. Even if developers get alerted to a potential operational issue, the process of identifying the root cause can still prove difficult. Using existing tools, developers often have difficulty triangulating the root cause of an operational issue from graphs and alarms, and even when they are able to find the root cause, they are often left without a means to fix it. Each troubleshooting attempt is a cold start where teams must spend hours or days to identify problems, and this leads to time consuming, tedious work that slows down the time to resolve an operational failure and can prolong application disruptions.

Amazon DevOps Guru's machine learning models leverage over 20 years of operational expertise in building, scaling, and maintaining highly available applications for Amazon.com. This gives Amazon DevOps Guru the ability to automatically detect operational issues (e.g. missing or misconfigured alarms, early warning of resource exhaustion, config changes that could lead to outages, etc.), provide context on resources involved and related events, and recommend remediation actions – with no machine learning experience required. With just a few clicks in the Amazon DevOps Guru console, historical application and infrastructure metrics like latency, error rates, and request rates for all resources are automatically ingested and analyzed to establish normal operating bounds, and Amazon DevOps Guru then uses a pre-trained machine learning model to identify deviations from the established baseline. When Amazon DevOps Guru analyzes system and application data to automatically detect anomalies, it also groups this data into operational insights that include anomalous metrics, visualizations of application behavior over time, and recommendations on actions for remediation. Amazon DevOps Guru also correlates and groups related application and infrastructure metrics (e.g. web application latency spikes, running out of disk space, bad code deployments, memory leaks etc.) to reduce redundant alarms and help focus users on high-severity issues. Customers can see configuration change histories and deployment events, along with system and user activity, to generate a prioritized list of likely causes for an operational issue in the Amazon DevOps Guru console. To help customers resolve issues quickly, Amazon DevOps Guru provides intelligent recommendations with remediation steps and integrates with AWS Systems Manager for runbook and collaboration tooling, giving customers the ability to more effectively maintain applications and manage infrastructure for their deployments. Together with Amazon CodeGuru – a developer tool powered by machine learning that provides intelligent recommendations for improving code quality and identifying an application's most expensive lines of code – Amazon DevOps Guru provides customers the automated benefits of machine learning for their operational data so that developers can more easily improve application availability and reliability.

"Customers have asked us to continue adding services around areas where we can apply our own expertise on how to improve application availability and learn from the years of operational experience that we have acquired running Amazon.com," said Swami Sivasubramanian, Vice President, Amazon Machine Learning, Amazon Web Services, Inc. "With Amazon DevOps Guru, we have taken our experience and built specialized machine learning models that help customers detect, troubleshoot, and prevent operational issues while

providing intelligent recommendations when issues do arise. This enables teams to immediately benefit from operational best practices Amazon has learned from running Amazon.com, saving customers the time and effort that would otherwise be spent configuring and managing multiple monitoring systems."

With a few clicks in the AWS Management Console, customers can enable Amazon DevOps Guru to begin analyzing account and application activity within minutes to provide operational insights. Amazon DevOps Guru gives customers a single-console experience to visualize their operational data by summarizing relevant data across multiple sources (e.g. AWS CloudTrail, Amazon CloudWatch, AWS Config, AWS CloudFormation, AWS X-Ray) and reduces the need to switch between multiple tools. Customers can also view correlated operational events and contextual data for operational insights within the Amazon DevOps Guru console and receive alerts via Amazon SNS. Additionally, Amazon DevOps Guru supports API endpoints through the AWS SDK, making it easy for partners and customers to integrate Amazon DevOps Guru into their existing solutions for ticketing, paging, and automatic notification of engineers for high-severity issues. PagerDuty and Atlassian are among the partners that have integrated Amazon DevOps Guru into their operations monitoring and incident management platforms, and customers who use their solutions can now benefit from operational insights provided by Amazon DevOps Guru. Amazon DevOps Guru is available in preview today in US East (N. Virginia), US East (Ohio), and US West (Oregon), Asia Pacific (Singapore), and Europe (Ireland) with availability in additional regions in the coming months.

Teams at more than 170,000 companies rely on Atlassian products to make teamwork easier, and help them organize, discuss, and complete their work. "Atlassian is proud to partner with AWS on the launch of Amazon DevOps Guru and help empower teams to deploy code and operate services with confidence," said Emel Dogrusoz, Head of Product, Opsgenie. "With our new Opsgenie and Jira Service Management integration, the right teams can be immediately notified the instant Amazon DevOps Guru predicts a potential issue, or determines an incident has occurred. Amazon DevOps Guru provides a new dimension of insight, and Atlassian ensures the fastest response."

PagerDuty, Inc. (NYSE:PD) is a leader in digital operations management. "PagerDuty was built to drive the move to a DevOps culture by automating the entire incident response lifecycle with resolution," said Jonathan Rende, CVP of Product at PagerDuty. "We're excited to continue this commitment to DevOps with our latest integration with Amazon DevOps Guru. Leveraging Amazon's decades of operational excellence and Amazon DevOps Guru's machine learning capabilities, PagerDuty provides even more real-time signal-to-action capabilities to our joint customers. Through PagerDuty's ingestion of Amazon DevOps Guru's Amazon SNS, AWS customers can take real-time action on operational issues before they become customer-impacting outages."

Thomson Reuters is one of the world's most trusted providers of answers, helping professionals make confident decisions and run better businesses. "Customer experience is vital to us. Dealing with multiple sources of alerts for availability, performance, and change requests can

be a challenge when trying to prevent and mitigate incidents impacting our customers," said Steve Thoennes, Director of Infrastructure Hosting Portfolio at Thomson Reuters. "We are excited to use Amazon DevOps Guru and leverage its machine learning insights to provide clear paths for action, allowing us to mitigate issues quickly and avoid customer impacting events. The integration with PagerDuty is a bonus, as we can have recommendations delivered to the right people timely and efficiently."

SmugMug is a paid image sharing service, image hosting service, and online video platform on which users can upload photos and videos. The company facilitates the sale of digital and print media for amateur and professional photographers. "My team follows an ops-for-life motto, and we are always on the lookout for ways to automate our manual activities," said Andrew Shieh, Operations Director at SmugMug. "With Amazon DevOps Guru, we hope to realize that goal and let AIOps take over many of our day-to-day tasks and make our workday composed of a single George-Jetson-style Easy Button, so my team can focus on IT innovation. We are now not only meeting the needs of the business but able to exceed them since we have more time to focus on what matters most – delivering value for our organization and our customers."

NextRoll helps marketplaces and marketing platforms grow revenue by empowering them to build and enhance their marketing solutions. "We run thousands of Amazon Elastic Compute Cloud (Amazon EC2) instances, and we are looking for ways to reduce my team's time spent on resolving operational issues," said Valentino Volonghi, CTO at NextRoll. "We are excited to use Amazon DevOps Guru and leverage its machine learning-powered insights to help us identify, correlate, and remediate operational issues with recommendations. This will help my team save hours and reduce our mean time to recovery."

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 just some of the products and services pioneered by Amazon. For more information visit www.amazon.com/about and follow @AmazonNews.