AWS re:Invent

ANT319

Security event and incident management using Amazon ES

Mark Carter
Product Management
Amazon Web Services

Carl Meadows
Product Management
Amazon Web Services

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Agenda

SIEM and security data lake—goals

Security data sources

Common data pipeline architectures

Analysis best practices

SIEM & security data lake—requirements

- Centralized repository of security signals
- Identify malicious activity
- Detect abnormal behavior
- Understand scope of attacks
- Multi-account, multi-region view
- 360 view of your cloud operational posture and risk level
- Low-latency, high-fidelity risk and compromise alerts
- High-speed security analyst investigation platform
- Low-cost automated long-term retention for forensics and regulatory compliance
- RBAC & SSO

Security data sources

- AWS services
  - P1—AWS CloudTrail, AWS Config, AWS Security Hub, Amazon CloudWatch Logs, Amazon Inspector, AWS WAF
  - P2—Security groups, VPC flow logs, ALB/NLB
- Compute endpoints
  - CloudWatch Logs agent, Fluentd, FireLens
  - Elastic Bits
  - AuditD, OSQuery, Wazuh, Packetbeat
  - EDP—Sophos, CrowdStrike, McAfee

Common data pipeline architectures

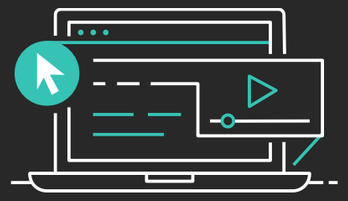
- Establish a streaming pipeline for ingest auto scale
  - Amazon Managed Streaming for Kafka
  - Amazon Kinesis Data Firehose, Amazon Kinesis Data Streams
- Parse, filter, enrich
  - AWS Lambda for streamlined use cases
  - Auto scale Logstash on Amazon Elastic Compute Cloud (Amazon EC2) to read from stream and batch to Amazon Elasticsearch Service (Amazon ES)
  - Use AWS Glue for advanced transforms to Amazon Simple Storage Service (Amazon S3)
  - Partners: Upsolver
  - Leverage AWS Glue catalog and AWS Lake Formation for security data lake
- Option: Use Amazon S3 as staging ground for Amazon ES and data lake ingest

Analysis best practices

- Pre-filter and enrich at the collection agent
- Normalize common field names across log sources to reduce joins
- Use index templates and mappings rather than auto-map
- Author and export reusable dashboards
- Use Vega graph for advanced Kibana dashboards
- Reuse community parser collection—Netflix, Empow
- Combine powers—Amazon ES for hot and warm, Amazon Athena for investigations, Amazon Redshift Spectrum
- Consider Amazon SageMaker for anomaly detection

Learn big data with AWS Training and Certification

Resources created by the experts at AWS to help you build and validate data analytics skills



New free digital course, Data Analytics Fundamentals, introduces Amazon S3, Amazon Kinesis, Amazon EMR, AWS Glue, and Amazon Redshift



Classroom offerings, including Big Data on AWS, feature AWS expert instructors and hands-on labs



Validate expertise with the **AWS Certified Big Data - Specialty** exam or the new **AWS Certified Data Analytics - Specialty** beta exam

Visit aws.amazon.com/training/paths-specialty/

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

aws training and certification

Thank you!

Mark Carter

datamark@amazon.com

Carl Meadows

carlmead@amazon.com

AWS re:Invent

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

aws



Please complete the session survey in the mobile app.

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.