Detecting Malicious Cloud Account Behavior:

A Look at the New Native Platform Capabilities

@bradgeesaman

Bio

Previously

- Network Security Engineer
- Penetration Tester/Security Consultant

Twitter: @bradgeesaman

Past 8+ Years

- Cloud Infrastructure Administrator
- “DevOps” practitioner *
- Ethical Hacking Educator
  - CTF Scenario design
  - Running CTF competition workloads inside public clouds using containers **

Past Two Years

- Researching Cloud Security Issues with Containers and Container Orchestrators
  - *Hacking and Hardening Kubernetes Clusters by Example*: https://youtu.be/vTgQLzeBfRU
- Independent Consulting - Securing Containers and Kubernetes

* Sorry
** Not recommended. It seemed like a good idea then.

Detecting Malicious Cloud Account Behavior

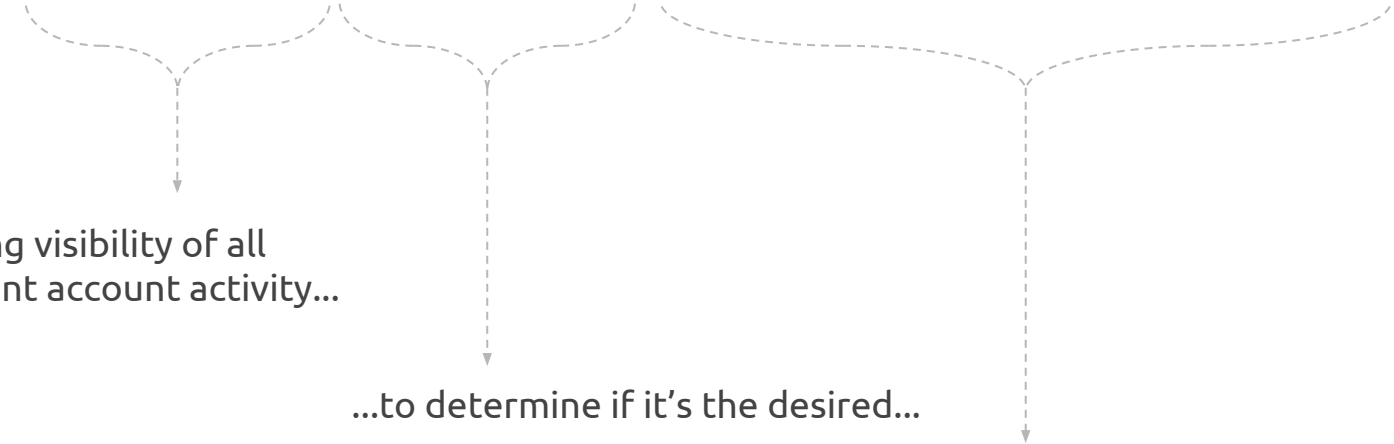

Getting visibility of all relevant account activity...

...to determine if it's the desired...

...usage of data, resources, workloads, and APIs inside a public cloud environment.

This Talk is Aimed at

Attackers

How malicious activity is being detected with the latest services enabled.

Defenders

Know more about cloud-specific attack indicators and how to gain better visibility of that activity.

Business Leaders

Understand the cloud-specific threat landscape, the cloud shared responsibility model, and where to focus detection efforts.

Security Architects/Ops/Builders

How and when to best leverage their cloud provider’s security service offerings.

Roadmap

Cloud Detection Challenges and Example Scenario

- Differences from Traditional Environments
- The Cloud Shared Responsibility Model
- The Cloud-Specific Attack Lifecycle
- Public Cloud Detection Data Sources
- Example Attack Scenario

The Latest “Native” Cloud Security Services

- Microsoft Azure Security Center
- Amazon GuardDuty (and CloudTrail) - DEMO!
- Google Cloud Security Command Center

Key Takeaways

- Benefits of the New Capabilities
- Areas for Improvement
- Adoption Recommendations
- Parting Perspectives

https://flic.kr/p/afRuwn

What makes detecting malicious behavior in the cloud different from traditional environments?

Cloud Environments Change Fundamental Assumptions

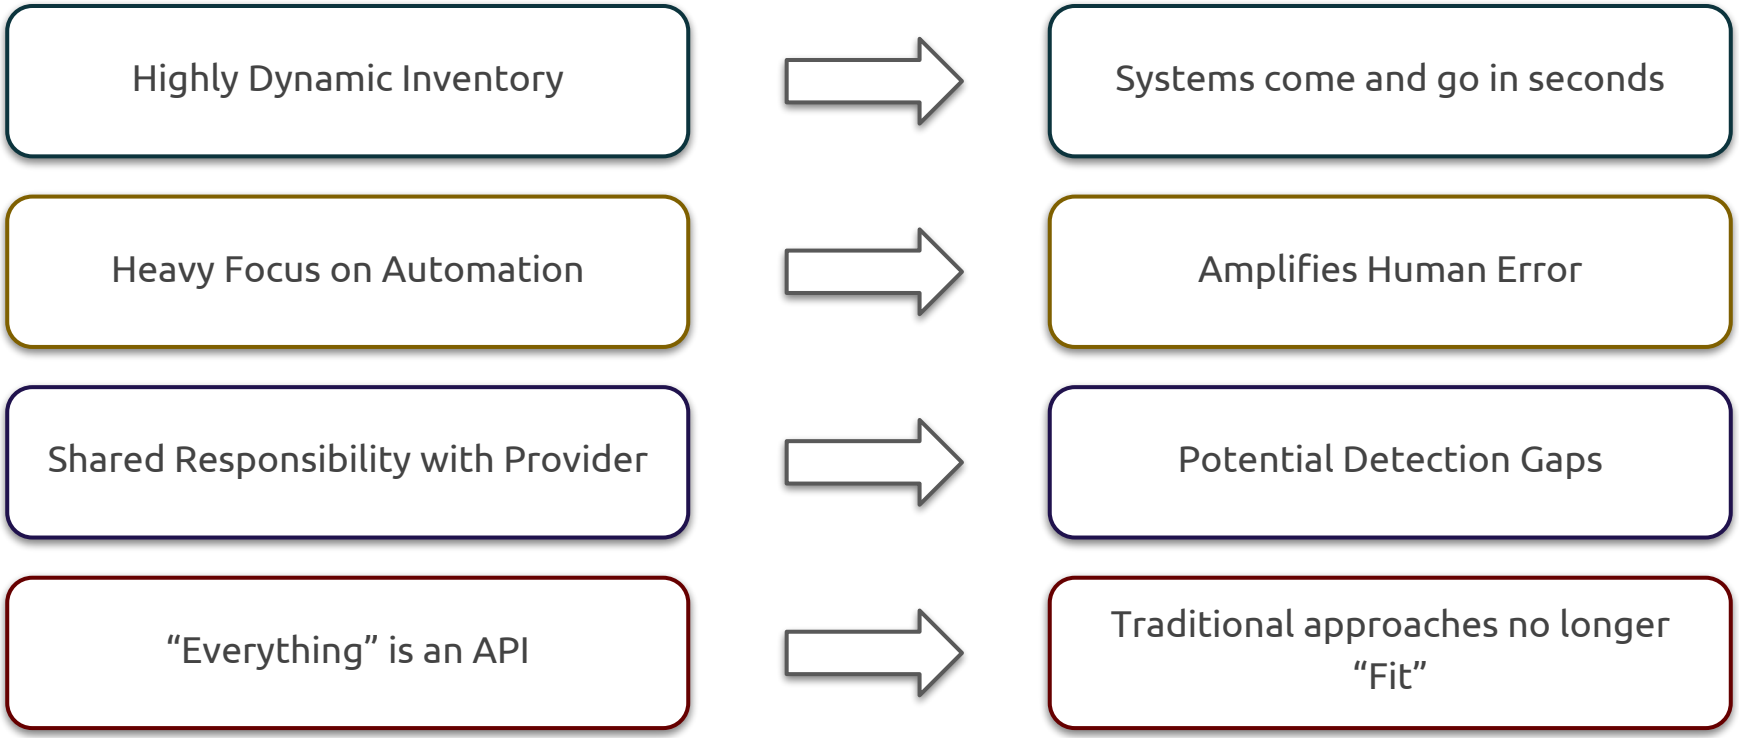

Pace of Innovation Leaves A Wake

Increasing business competition	→	Focus on shipping features first, outsourcing non-core capabilities.
An explosion of cloud services	→	What “Perimeter”?
A renaissance of infrastructure and deployment tooling	→	New environments with new security models and attack surfaces.
Always a security expertise shortage	→	Amplified with all the newly released features and services.

Understanding the responsibility boundary

AWS Shared Responsibility Model

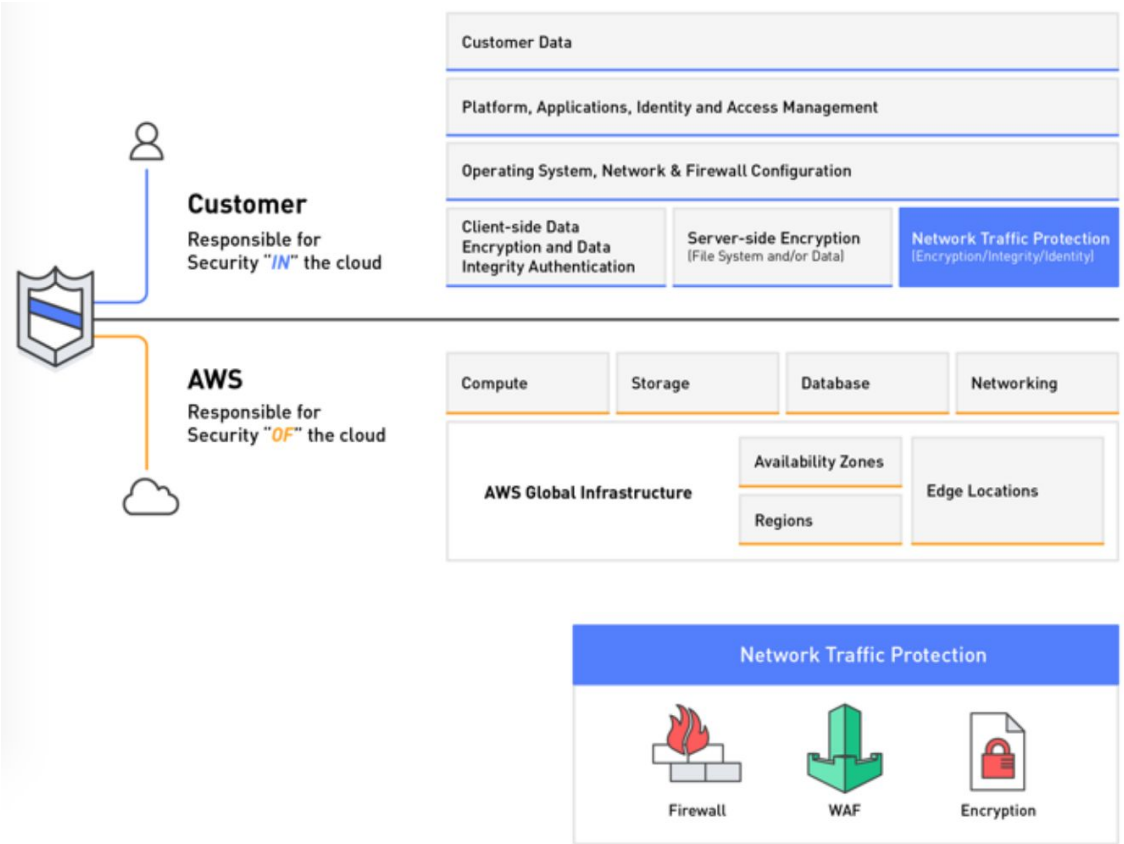

https://aws.amazon.com/mp/scenarios/security/malware/

AWS Shared Responsibility Model (Adapted)

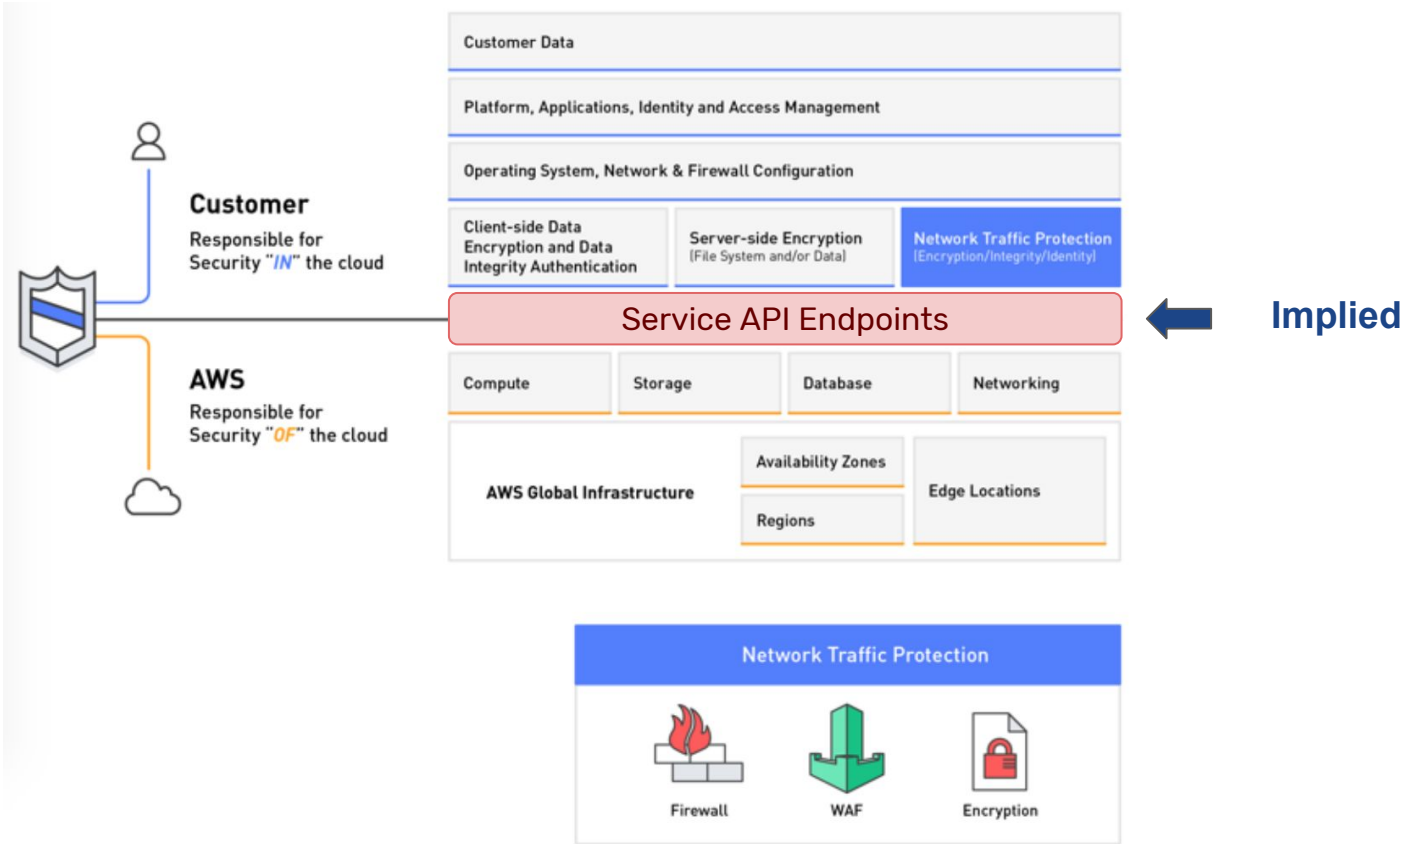

Adapted from https://aws.amazon.com/mp/scenarios/security/malware/

Shared Responsibility Model - Customer's View

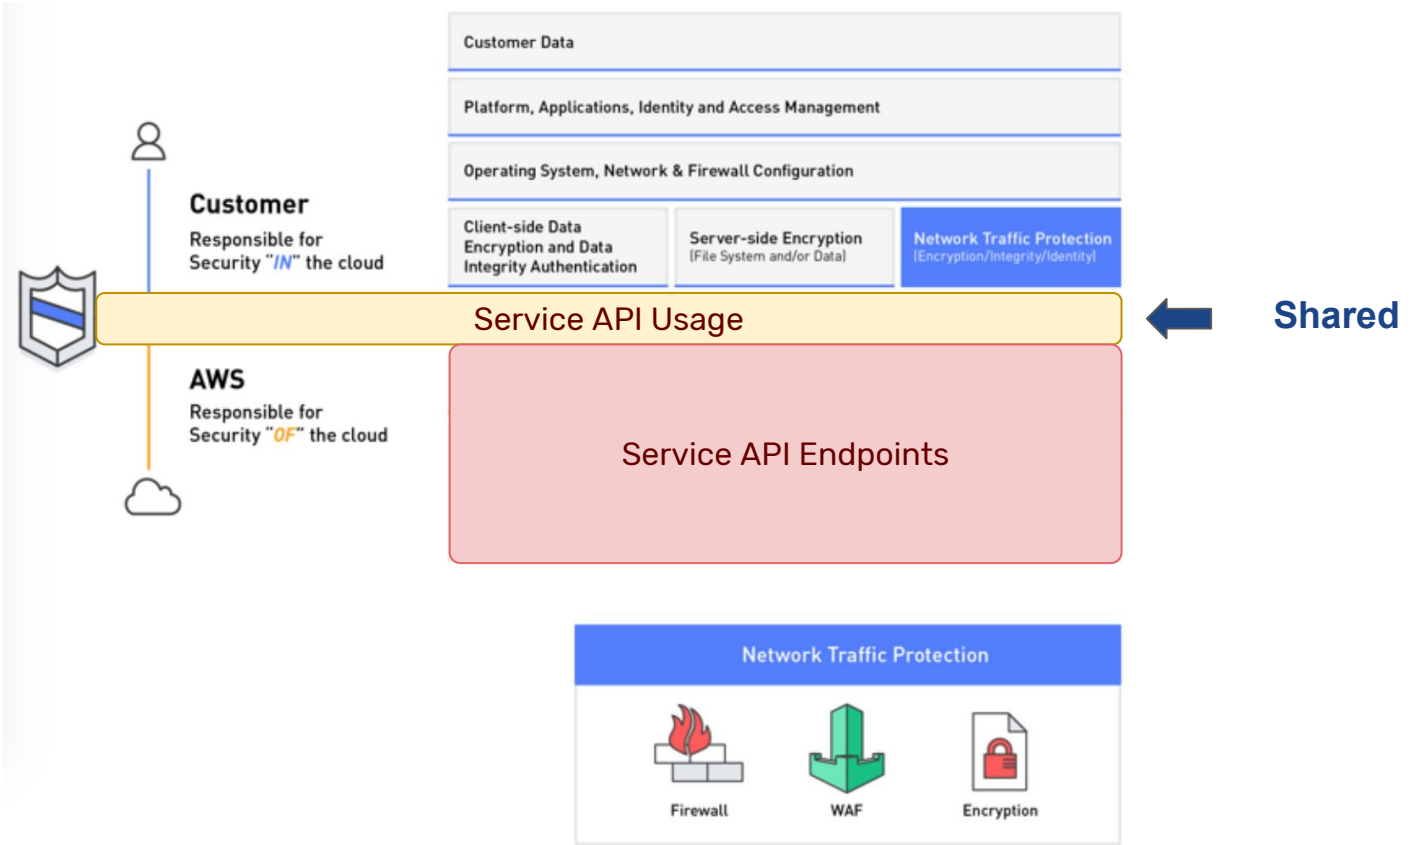

Adapted from https://aws.amazon.com/mp/scenarios/security/malware/

Shared Responsibility Model - Cloud Provider's View

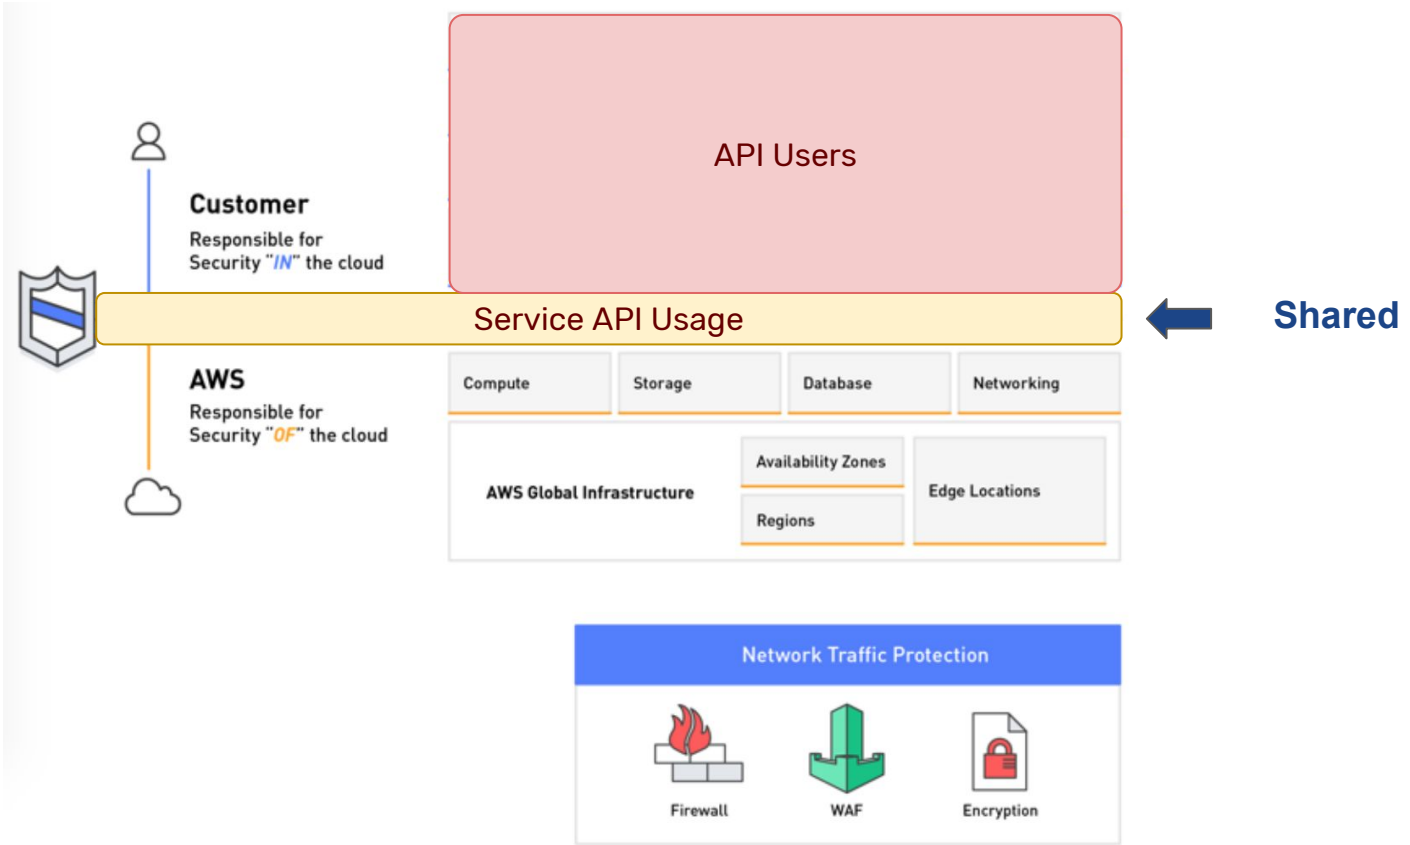

Adapted from https://aws.amazon.com/mp/scenarios/security/malware/

Protecting those shared APIs are challenging and nuanced,
but very necessary.

What does an attack lifecycle look like in a cloud environment?

Traditional Attack Path/Lifecycle (Simplified)

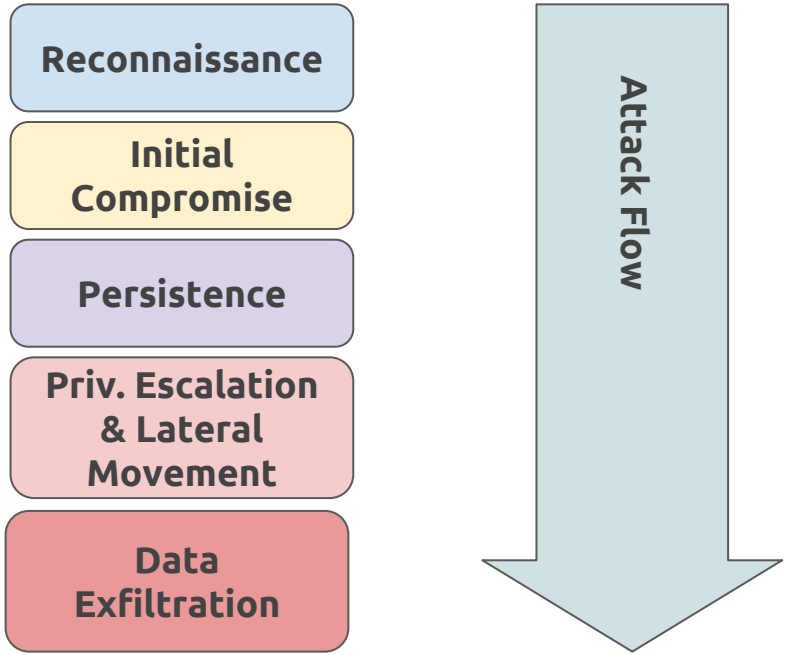

Cloud Attack Path/Lifecycle (Adapted)

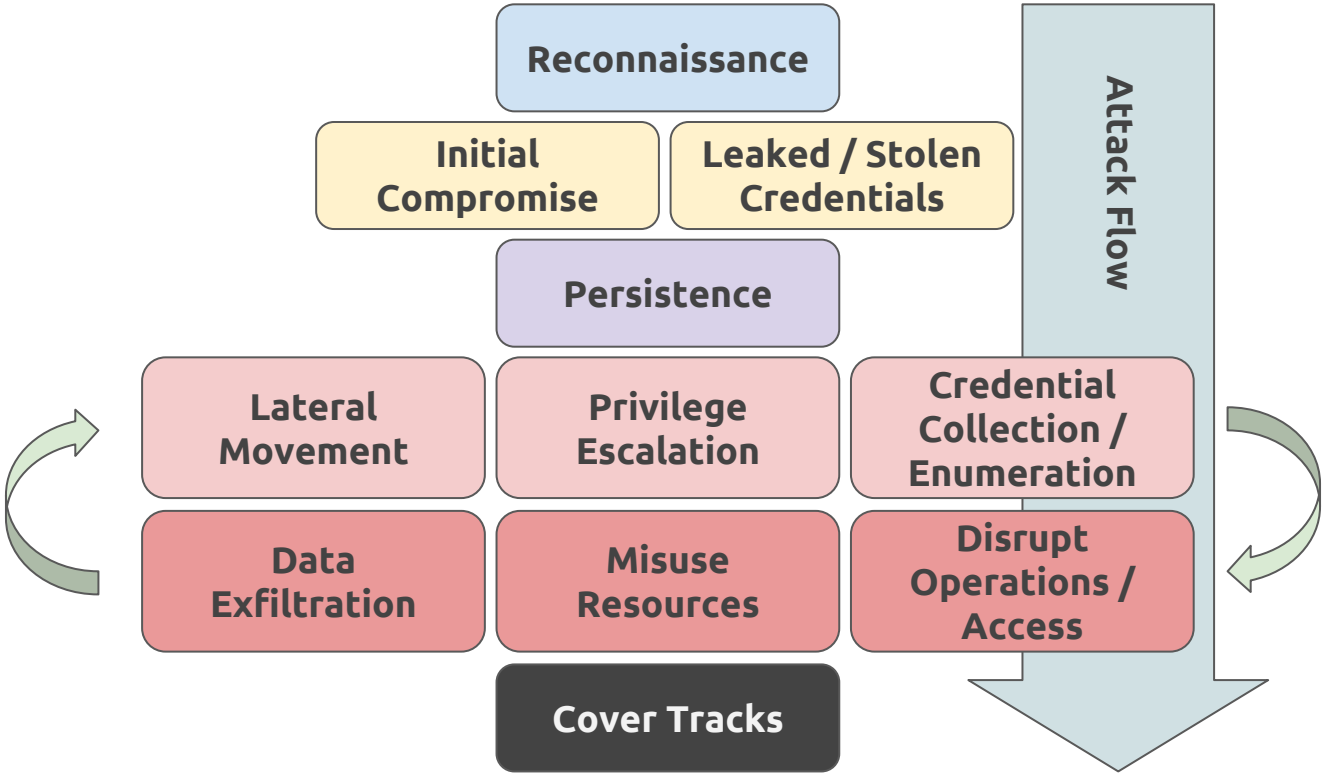

Escalation, Enumeration, Persistence, Covering Tracks

Escalation

- https://rhinosecuritylabs.com/aws/aws-privilege-escalation-methods-mitigation/

Enumeration/Exploration

- https://danielgrzelak.com/exploring-an-aws-account-after-pwning-it-ff629c2aae39

Persistence

- https://danielgrzelak.com/backdooring-an-aws-account-da007d36f8f9

Covering Tracks

- https://danielgrzelak.com/disrupting-aws-logging-a42e437d6594

* Concepts apply to all cloud providers

What detection methods are available?

Cloud Account Behavior Data Detection Sources*

Network

- Activity to/from **known-bad IPs**
- Unusual changes to traffic patterns
- Unusual outbound port usage

DNS

- Queries to **known-bad domains** (CnC, bots, malware, crypto-mining, etc) or embed data in the lookup

Host-based

- **OS, Application, Security/Audit logs**
- **Endpoint security events**

Network-Device based

- FW/IDS/IPS “drop-in” solution logs/alerts

Cloud Provider API Activity

- Multiple failed logins
- Simultaneous API access from different countries
- Attempted activity from terminated accounts/credentials/keys
- Uncommon service/API usage
- Credential/permission enumeration
- Changes to user accounts/logging/detection configurations
- Sensitive changes to user permissions
- Internal IAM credentials used from external sources

Service Access Logs

- Web/User Access logs

* Not an exhaustive list.

Cloud Account Behavior Data Detection Sources*

Network

- Activity to/from **known-bad IPs**
- Unusual changes to traffic patterns
- Unusual outbound port usage

DNS

- Queries to **known-bad domains** (CnC, bots, malware, crypto-mining, etc) or embed data in the lookup

Host-based

- **OS, Application, Security/Audit logs**
- **Endpoint security events**

Network-Device based

- FW/IDS/IPS “drop-in” solution logs/alerts

Cloud Provider API Activity

- Multiple failed logins
- Simultaneous API access from different countries
- Attempted activity from terminated accounts/credentials/keys
- Uncommon service/API usage
- Credential/permission enumeration
- Changes to user accounts/logging/detection configurations
- Sensitive changes to user permissions
- Internal IAM credentials used from external sources

Service Access Logs

- Web/User Access logs

* Not an exhaustive list.

Example Attack Walkthrough

An Electric Car Manufacturer

Exposed Kubernetes Dashboard

- Kubernetes Cluster on AWS
- Installed CPU-throttled **crypto-mining** workers
- Tight integration with AWS Access Keys led to **S3 data exfiltration**
- Masked their sources behind a CDN

Not Alone

- Multiple other Companies had the same issue

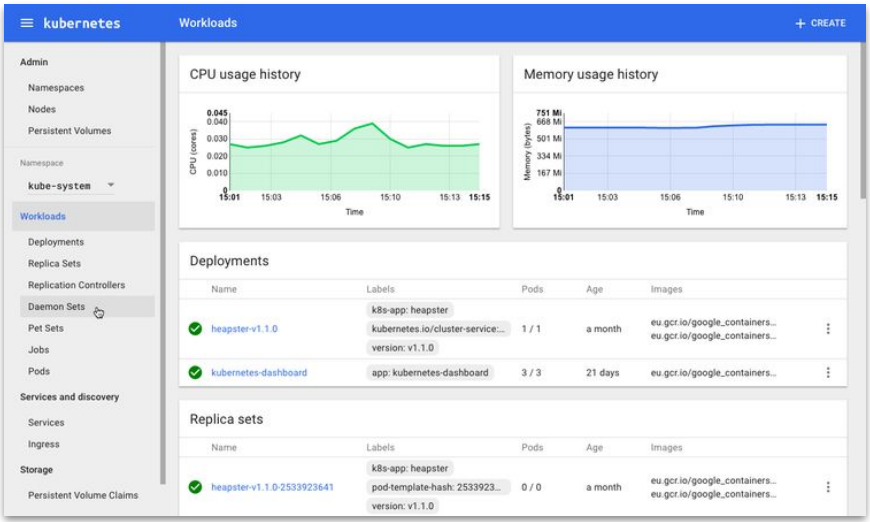

http://fortune.com/2018/02/20/tesla-hack-amazon-cloud-cryptocurrency-mining/

An Electric Car Manufacturer

Possible Detection Methods

- **Instance IAM credentials usage** from a non-cloud instance
- **DNS logs** of malware/crypto-mining software
- Dashboard **Application Logs**
- **Netflow Logs** of Docker image download
- **Netflow Logs** of reports into mining pool

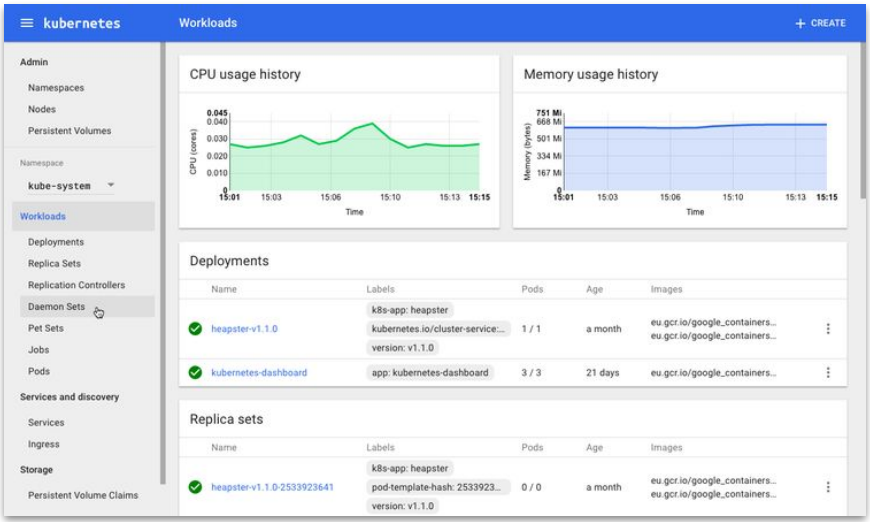

http://fortune.com/2018/02/20/tesla-hack-amazon-cloud-cryptocurrency-mining/

Note: A Direct Compromise May Not Be Needed

Credential theft

- **Phishing**
- Malware
- **Backdoored libraries/tools**
- Password guessing/weak passwords

Malicious Outsiders

- **Compromise of 3rd Party Services with integrated access**
  - Source Control
  - CI/CD
  - Mail Delivery
- Failure to disable, delete, rotate credentials post termination

Credential Leaks

- **Checked into source code**
- Technical support tickets
- Public Q&A Tech Help chat/forums
- Applications transmit keys in headers, messages, or logs of API calls

The Latest “Native” Cloud Security Services

Services In Scope

Microsoft Azure Security Center, Advanced Threat Protection

AWS GuardDuty (and CloudTrail, CloudWatch)

Google Cloud Security Command Center

Service Launch Dates

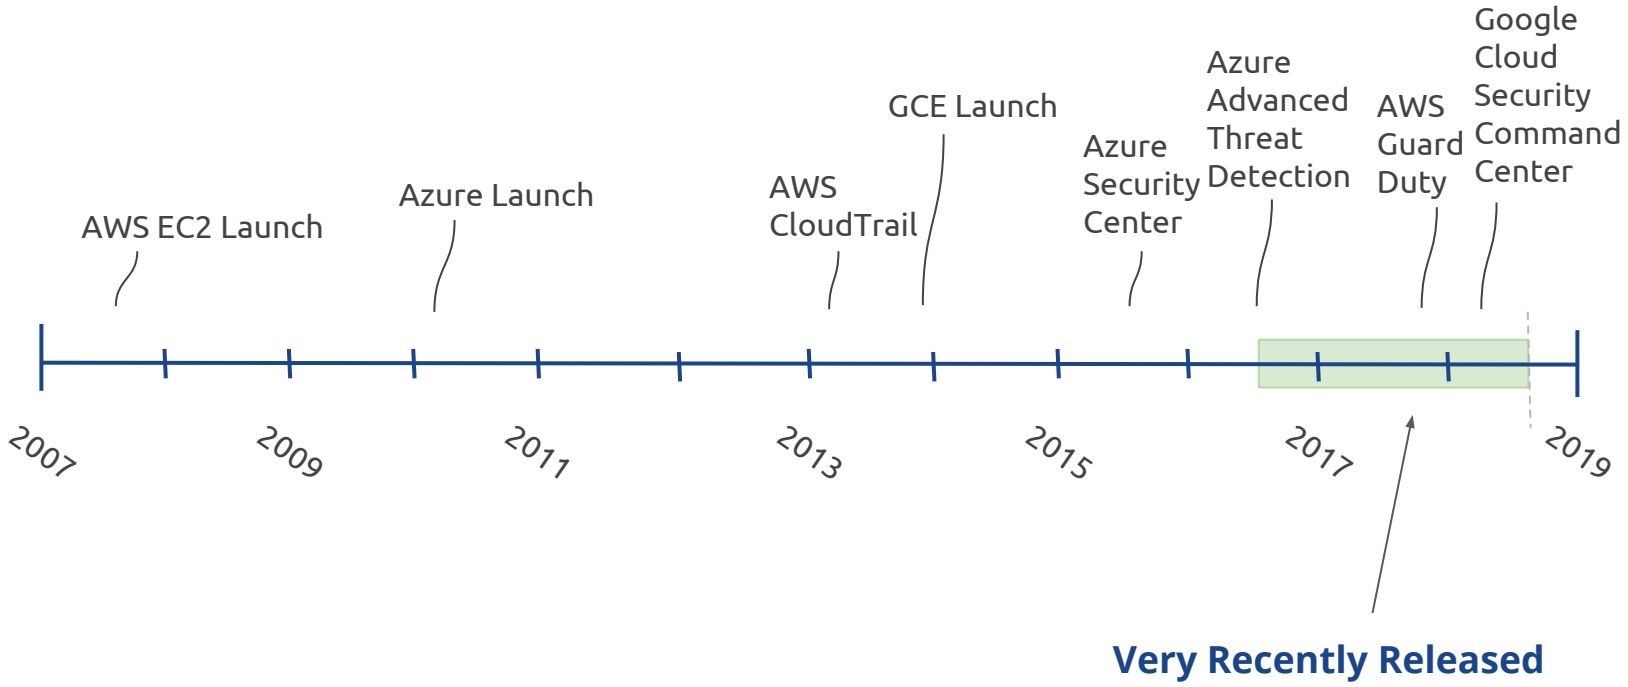

Questions Asked During This Review

- What **data sources** do they use?
- How do they operate on that data?
- What **visibility** does that data provide?
- What is **not covered** in the service?
- What is needed for onboarding?
- What's the **cost** structure?
- How does it **integrate** with other internal services and partners?
- How accessible are these services to **customization**?
- How do you validate the detection capabilities?

Different Questions for Different Roles

Attackers

- What methods and tactics need to change to remain undetected?

Defenders

- What can be covered?
- What still isn't covered?

Business Leaders

- What's my exposure?
- What's the ROI?

Security Architects/Ops/Builders

- How does this change my infrastructure design?
- What do I no longer have to build?

Azure Security Center

Azure Security Center

Released

- Initial - Fall 2015
- Generally Available - Spring/Summer 2016
- Advanced Threat Detection - Summer 2017

Description

- Azure Security Center provides **unified security management** and advanced threat protection **across hybrid cloud workloads**. With Security Center, you can apply security policies across your workloads, limit your exposure to threats, and detect and respond to attacks.
- Cost: $15/system/month

Links and Documentation

- https://docs.microsoft.com/en-us/azure/security-center/

Key Features

Unified / Hybrid Security Dashboard

- Common **Windows-style management experience** in the cloud and on-premise in a single place.

Security Recommendation Engine

- Suggests **security hygiene** items to address proactively. Offers customizable policy (XML) for user-supplied checks.

Microsoft Provided Agent

- **OS, Application, Security/Audit logs, missing patches, weak configurations** and more supplement network-based detections. Can be automatically enabled for all VMs.

Key Features (Cont'd)

Third-Party Security Tool Integration Marketplace

- **Centrally integrate** your choice of multiple security endpoint solutions, host-based vulnerability management agents, and network-security devices with a few clicks and some license keys.

Custom Alert Rules

- Custom queries on all log event types to trigger **notification alerts**.

File Integrity Monitoring (Preview)

- Validates the integrity of Windows files, Windows registry, and Linux files

REST API

- **Integration with your existing security systems** and workflows for inserting and pulling events.

Detection Data Sources

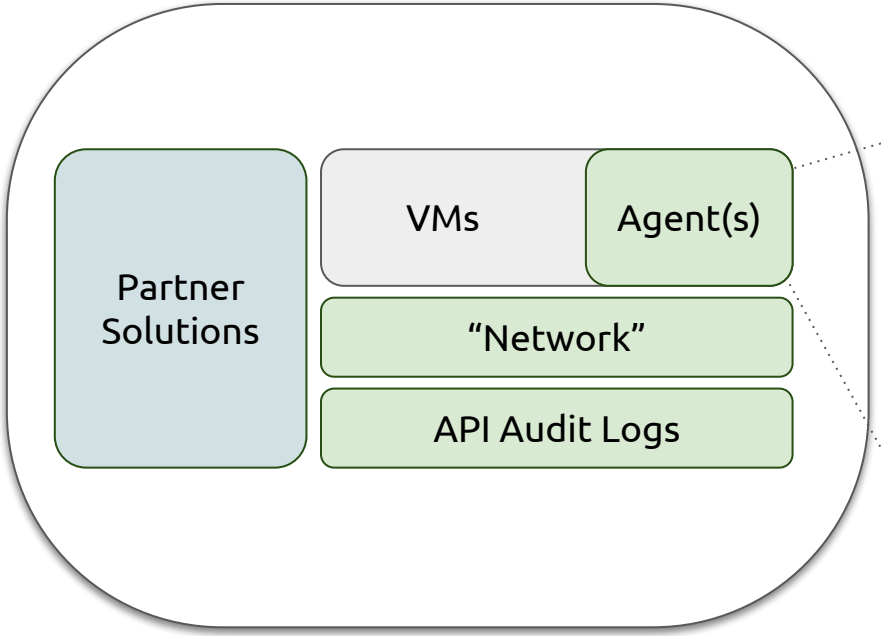

Microsoft Agent Operating Systems

- Windows Server (of course)
- Amazon Linux 2012.09 --> 2017
- CentOS Linux 5,6, and 7
- Oracle Linux 5,6, and 7
- Red Hat Enterprise Linux Server 5,6 and 7
- Debian GNU/Linux 6, 7, 8, and 9
- Ubuntu 12.04, 14.04, 16.04 LTS
- SUSE Linux Enterprise Server 11/ 12

Simplified Architecture

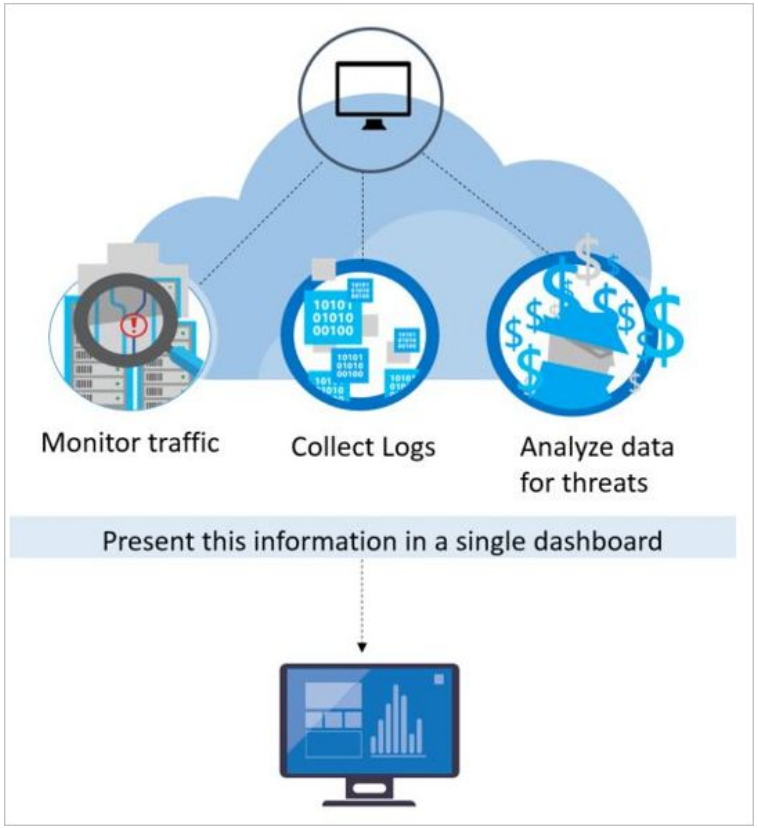

https://docs.microsoft.com/en-us/azure/security-center/security-center-detection-capabilities

Detections

Threat Intelligence

- Outbound communication to a **malicious IP address/Domains**
- Threat intelligence monitoring and signal sharing across all their services

Behavioral Analytics

- Suspicious process execution: models processes behaviors and **monitors process executions** to detect outliers
- Hidden malware and exploitation attempts: **memory analysis, crash dump analysis**
- Lateral movement and internal reconnaissance: monitors process and **login activities** such as remote command execution network probing, and **account enumeration**
- Malicious PowerShell Scripts: inspects **PowerShell activity** for evidence of suspicious activity
- Outgoing attacks: take part in brute force, scanning, DDoS, and Spam sending campaigns

Anomaly Detection

- Inbound RDP/SSH brute force attacks

https://docs.microsoft.com/en-us/azure/security-center/security-center-detection-capabilities

Dashboard

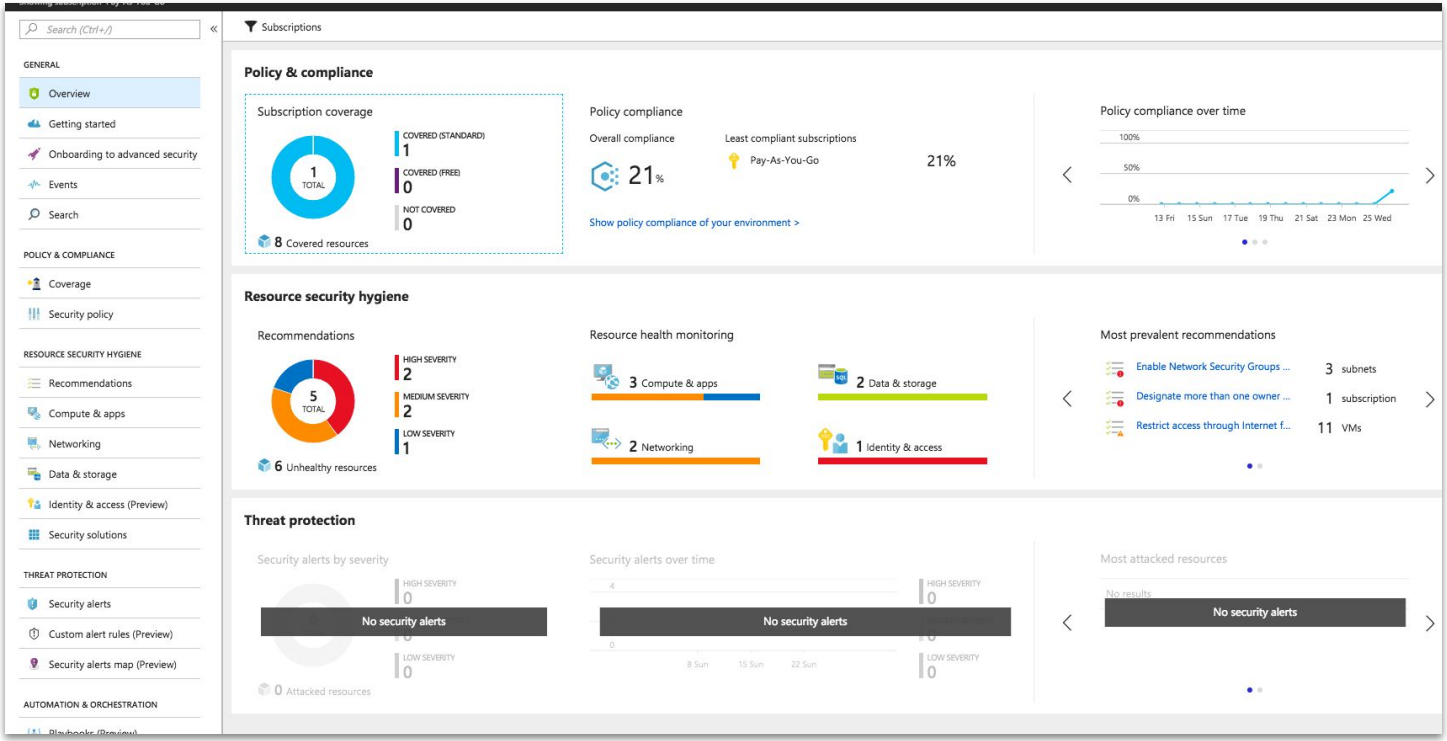
Search (Ctrl+/)
Subscriptions
GENERAL
Overview
Getting started
Onboarding to advanced security
Events
Search
POLICY & COMPLIANCE
Coverage
Security policy
RESOURCE SECURITY HYGIENE
Recommendations
Compute & apps
Networking
Data & storage
Identity & access (Preview)
Security solutions
THREAT PROTECTION
Security alerts
Custom alert rules (Preview)
Security alerts map (Preview)
AUTOMATION & ORCHESTRATION
Policy & compliance
Subscription coverage
1 TOTAL
COVERED (STANDARD) 1
COVERED (FREE) 0
NOT COVERED 0
8 Covered resources
Policy compliance
Overall compliance
21%
Least compliant subscriptions
Pay-As-You-Go 21%
Show policy compliance of your environment >
Policy compliance over time
100%
50%
0%
13 Fri 15 Sun 17 Tue 19 Thu 21 Sat 23 Mon 25 Wed
Resource security hygiene
Recommendations
5 TOTAL
HIGH SEVERITY 2
MEDIUM SEVERITY 2
LOW SEVERITY 1
6 Unhealthy resources
Resource health monitoring
3 Compute & apps
2 Data & storage
2 Networking
1 Identity & access
Most prevalent recommendations
Enable Network Security Groups ... 3 subnets
Designate more than one owner ... 1 subscription
Restrict access through Internet f... 11 VMs
Threat protection
Security alerts by severity
HIGH SEVERITY 0
No security alerts
LOW SEVERITY 0
0 Attacked resources
Security alerts over time
HIGH SEVERITY 0
No security alerts
LOW SEVERITY 0
8 Sun 15 Sun 22 Sun
Most attacked resources
No results
No security alerts

Dashboard

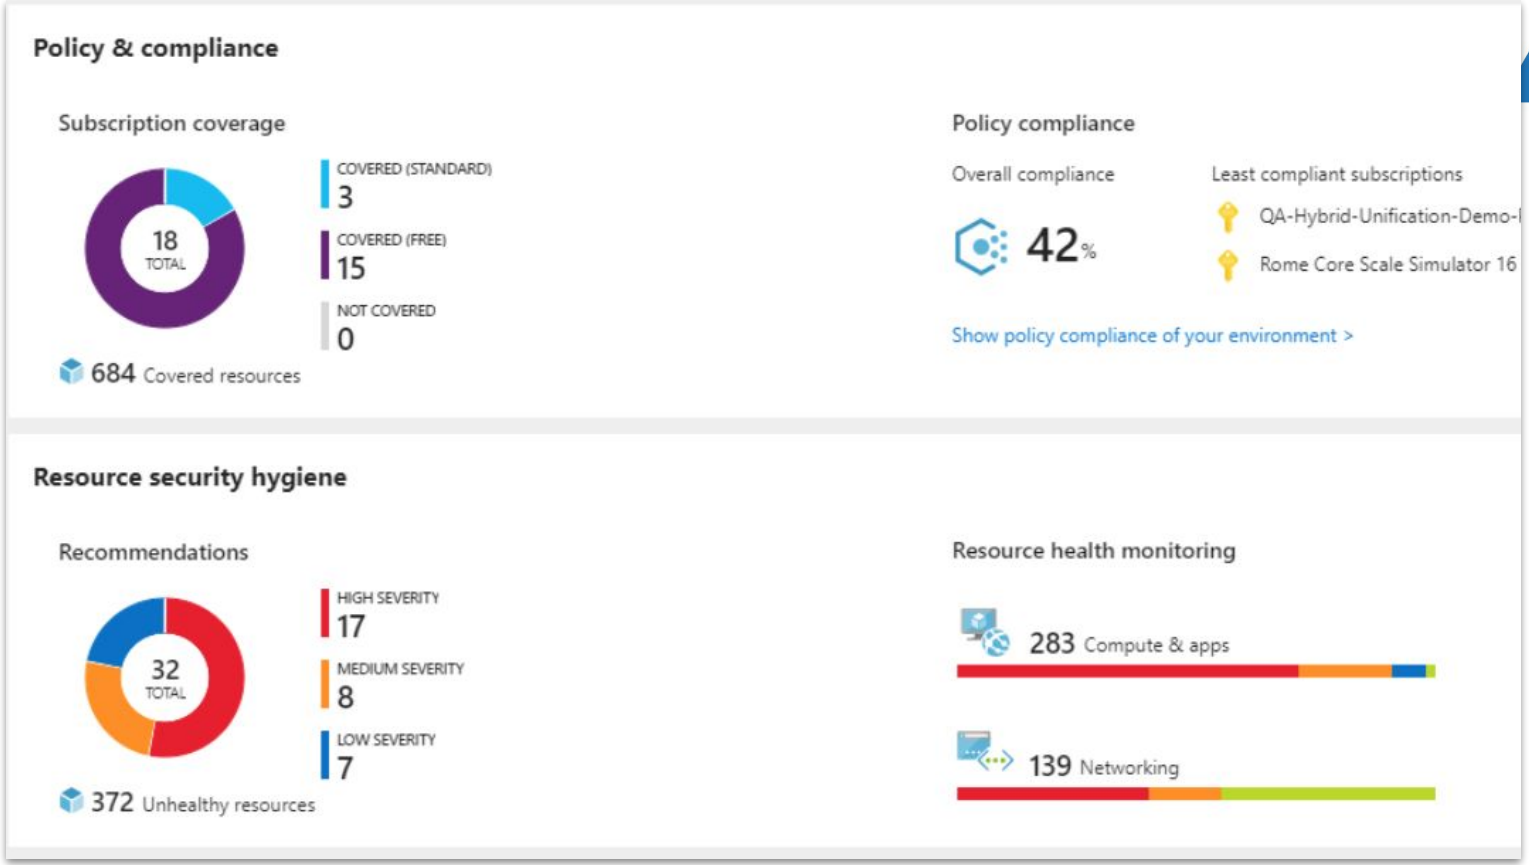
Policy & compliance
Subscription coverage
18
TOTAL
COVERED (STANDARD)
3
COVERED (FREE)
15
NOT COVERED
0
684 Covered resources
Policy compliance
Overall compliance
42%
Least compliant subscriptions
QA-Hybrid-Unification-Demo-
Rome Core Scale Simulator 16
Show policy compliance of your environment >
Resource security hygiene
Recommendations
32
TOTAL
HIGH SEVERITY
17
MEDIUM SEVERITY
8
LOW SEVERITY
7
372 Unhealthy resources
Resource health monitoring
283 Compute & apps
139 Networking

Dashboard

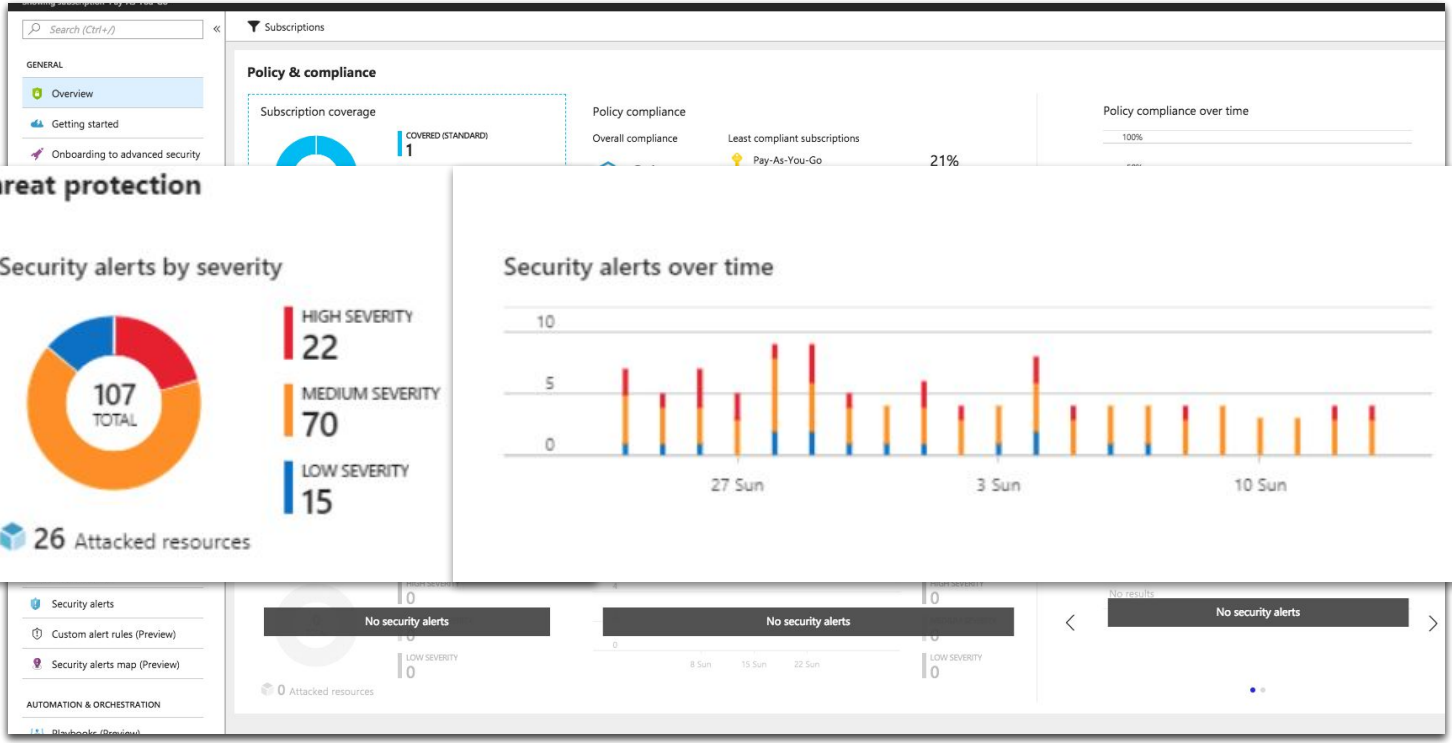
Search (Ctrl+/)
Subscriptions
GENERAL
Overview
Getting started
Onboarding to advanced security
Policy & compliance
Subscription coverage
COVERED (STANDARD)
1
Policy compliance
Overall compliance
Least compliant subscriptions
Pay-As-You-Go
21%
Policy compliance over time
100%
Threat protection
Security alerts by severity
107
TOTAL
HIGH SEVERITY
22
MEDIUM SEVERITY
70
LOW SEVERITY
15
26 Attacked resources
Security alerts over time
10
5
0
27 Sun
3 Sun
10 Sun
HIGH SEVERITY
22
MEDIUM SEVERITY
70
LOW SEVERITY
15
Security alerts
Custom alert rules (Preview)
Security alerts map (Preview)
AUTOMATION & ORCHESTRATION
No security alerts
No security alerts
No security alerts
0 Attacked resources

Agent Reports Missing Patches

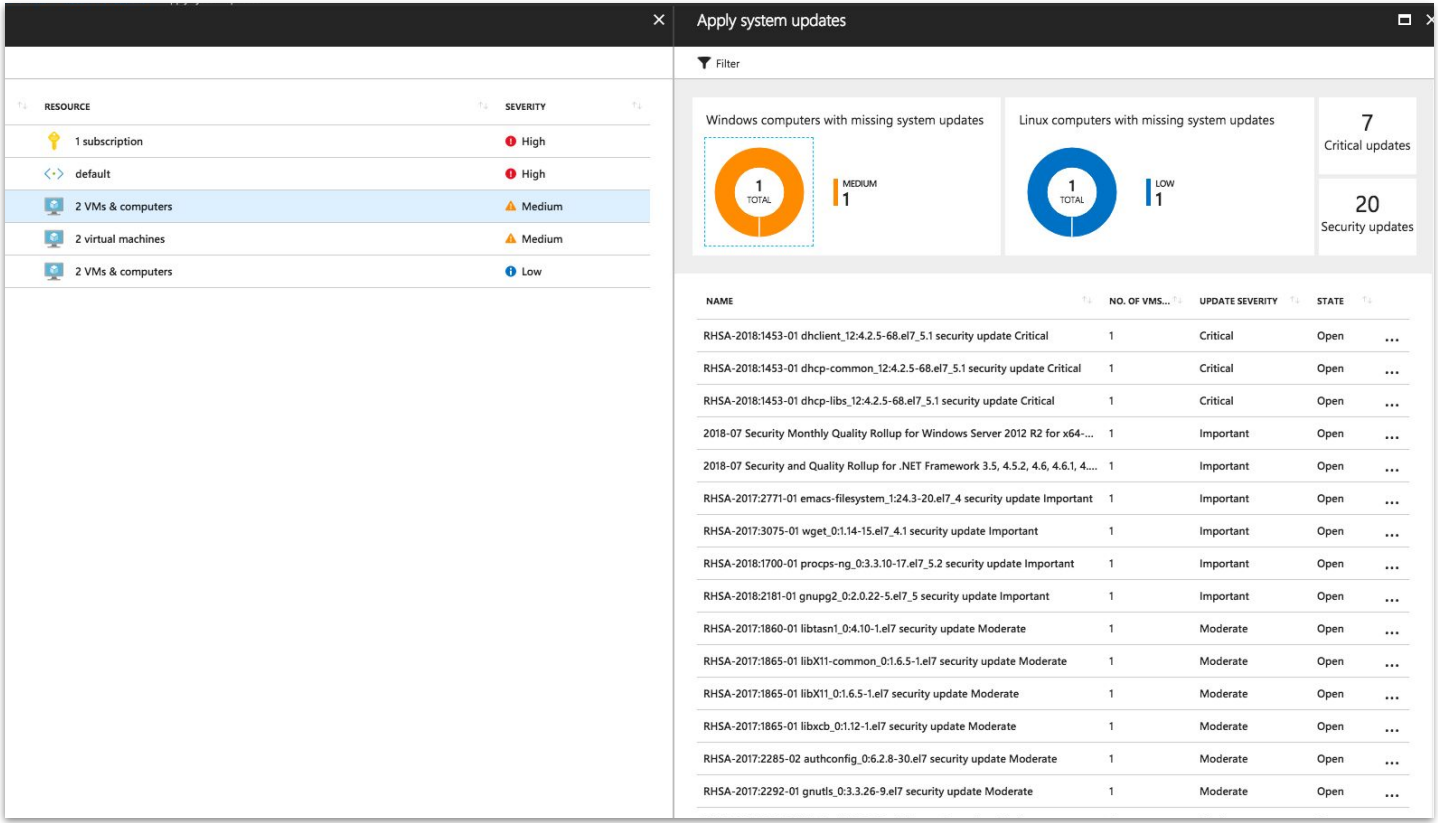

RESOURCE	SEVERITY
1 subscription	High
default	High
2 VMs & computers	Medium
2 virtual machines	Medium
2 VMs & computers	Low

Apply system updates

Filter

Windows computers with missing system updates

1 TOTAL

MEDIUM 1

Linux computers with missing system updates

1 TOTAL

LOW 1

7 Critical updates

20 Security updates

NAME	NO. OF VMS...	UPDATE SEVERITY	STATE
RHSA-2018:1453-01 dhclient_12:4.2.5-68.el7_5.1 security update Critical	1	Critical	Open
RHSA-2018:1453-01 dhcp-common_12:4.2.5-68.el7_5.1 security update Critical	1	Critical	Open
RHSA-2018:1453-01 dhcp-libs_12:4.2.5-68.el7_5.1 security update Critical	1	Critical	Open
2018-07 Security Monthly Quality Rollup for Windows Server 2012 R2 for x64-...	1	Important	Open
2018-07 Security and Quality Rollup for .NET Framework 3.5, 4.5.2, 4.6, 4.6.1, 4....	1	Important	Open
RHSA-2017:2771-01 emacs-filesystem_1:24.3-20.el7_4 security update Important	1	Important	Open
RHSA-2017:3075-01 wget_0:1.14-15.el7_4.1 security update Important	1	Important	Open
RHSA-2018:1700-01 procps-ng_0:3.3.10-17.el7_5.2 security update Important	1	Important	Open
RHSA-2018:2181-01 gnupg2_0:2.0.22-5.el7_5 security update Important	1	Important	Open
RHSA-2017:1860-01 libtasn1_0:4.10-1.el7 security update Moderate	1	Moderate	Open
RHSA-2017:1865-01 libX11-common_0:1.6.5-1.el7 security update Moderate	1	Moderate	Open
RHSA-2017:1865-01 libX11_0:1.6.5-1.el7 security update Moderate	1	Moderate	Open
RHSA-2017:1865-01 libxcb_0:1.12-1.el7 security update Moderate	1	Moderate	Open
RHSA-2017:2285-02 authconfig_0:6.2.8-30.el7 security update Moderate	1	Moderate	Open
RHSA-2017:2292-01 gnutls_0:3.3.26-9.el7 security update Moderate	1	Moderate	Open

Agent Reports Missing Patches

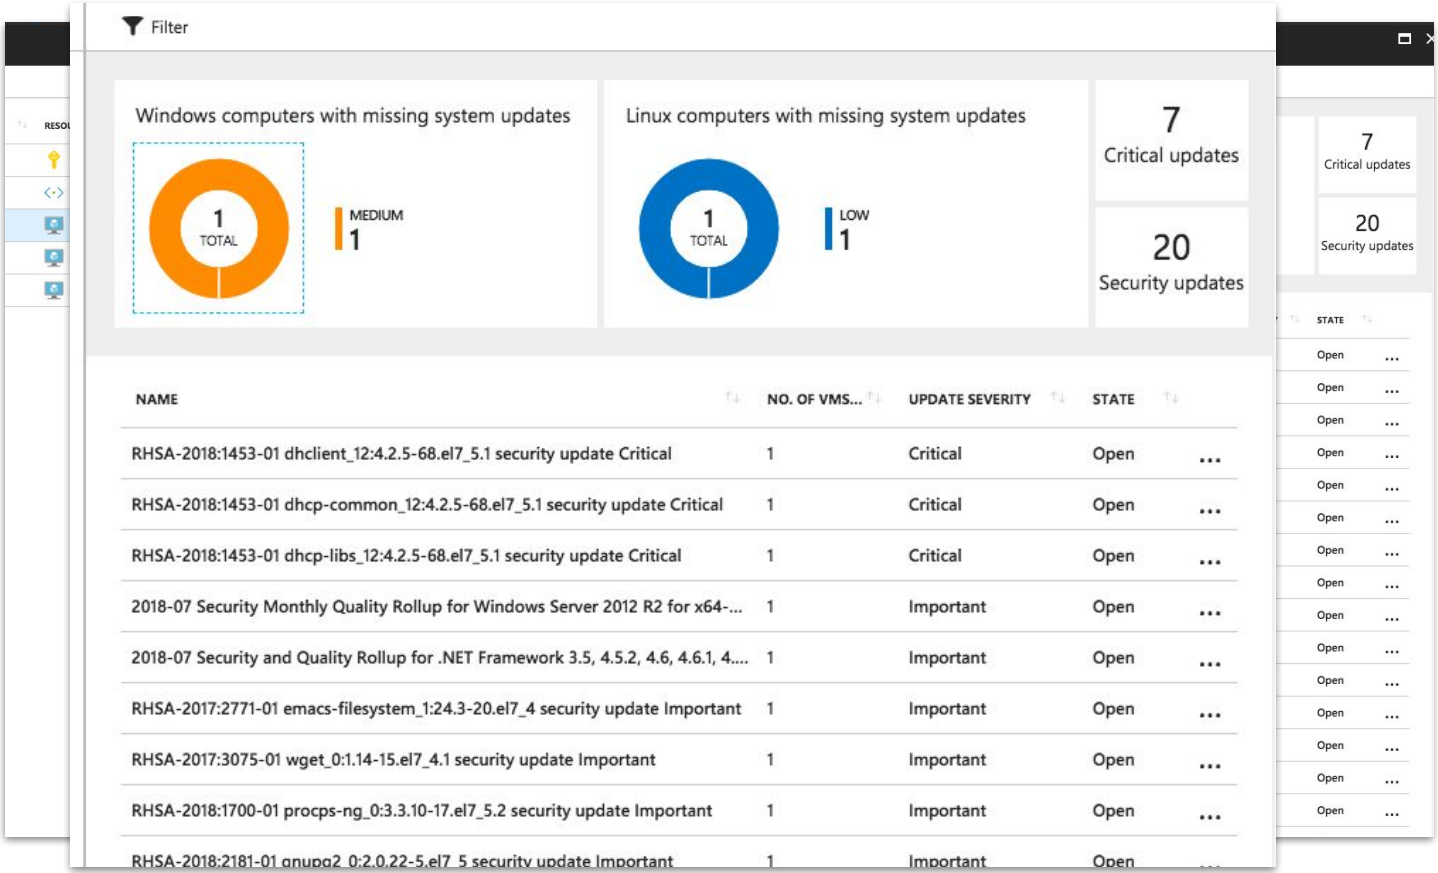

All Agent Logs are Searchable

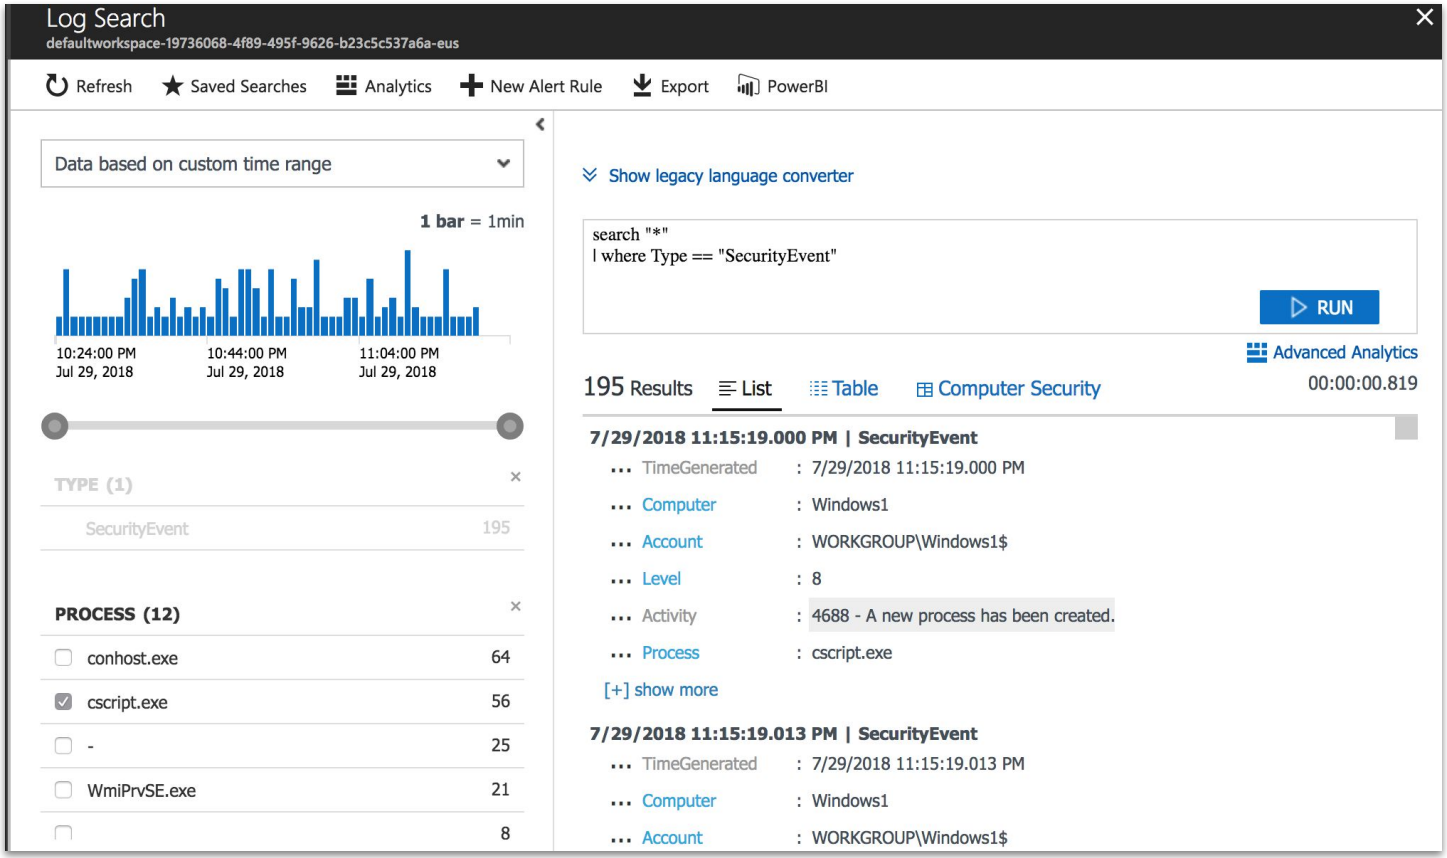

Value Added

Hybrid-first approach

- Leverages the vast amount of **enterprise management features** and capabilities applied to Azure resources.

Provides a Microsoft-supported Windows/Linux Agent

- Supported OSes get enhanced detection capabilities (logs, process monitoring, crash dump analysis)

Integrated, Self-Service Partner Marketplace

- Adding a solution is a few clicks and a license away in many cases.

Leverages the Azure Log Analytics Service

- Mature integrations, advanced querying, and **full-featured REST API**

Areas for Improvement

Areas for Improvement

- A detailed list of anomalous detection capabilities is not yet available.
- Ability to **tune detection** parameters.
- Potential delay from agent deployment to it reporting in the Dashboard.
- The ability to supply custom threat/IP feeds to aid in improving detection accuracy.

Amazon GuardDuty et al

Amazon GuardDuty et al

Released

- AWS CloudTrail: Spring 2013
- AWS VPC Flow Logs: Summer 2015
- Amazon GuardDuty: Winter 2017

Description

- Amazon GuardDuty offers threat detection that enables you to continuously monitor and protect your AWS accounts and workloads.
- 30-day free trial.
- North America: $0.25-$1 per GB of VPC/DNS, $4 per 1M Cloudtrail Events

Links and Documentation

- https://aws.amazon.com/guardduty/

Key Features

Watches Data Streams

- AWS CloudTrail Events, Amazon VPC Flow Logs, and DNS Logs.

Integrates Threat Intelligence Feeds and Machine Learning

- Feeds with **known malicious IP addresses and domains**.
- **Environment specific** baselining.
- You can supply your **own IP lists** for “good” and “bad” hosts.

Generates Findings

- Creation action creates CloudWatch events useful for triggering Lambda functions for further processing and sending notifications.

Cross-Account Visibility

- **Events can be centralized** across multiple “member” accounts to a centralized “master” account.

How CloudTrail Works

Amazon CloudTrail

Account activity occurs

CloudTrail records a CloudTrail Event

You can view/download your activity in the CloudTrail Event History

You can set up CloudTrail and define an Amazon S3 bucket for storage

A log of CloudTrail Events is delivered to an S3 bucket and optionally to CloudWatch Logs and CloudWatch Events

How GuardDuty Works

Amazon GuardDuty

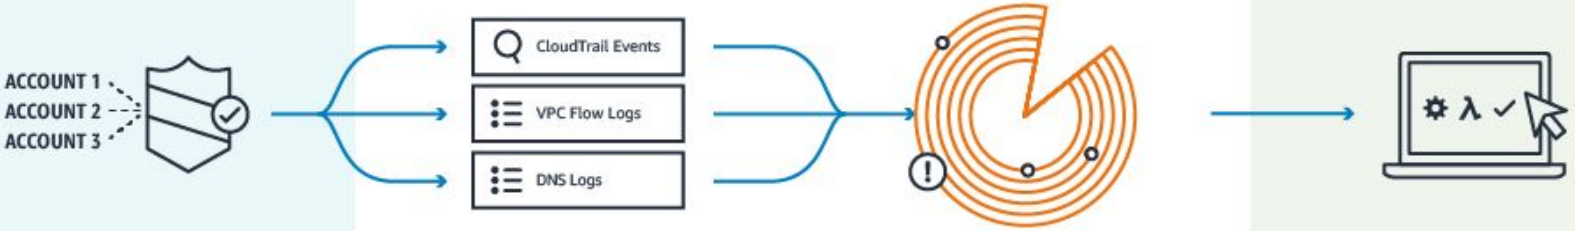

Enable GuardDuty

With a few clicks in the console, monitor all your AWS accounts without additional security software or infrastructure to deploy or manage

Continuously analyze

Automatically analyze network and account activity at scale, providing broad, continuous monitoring of your AWS accounts

Intelligently detect threats

GuardDuty combines managed rule-sets, threat intelligence from AWS Security and 3rd party intelligence partners, anomaly detection, and ML to intelligently detect malicious or unauthorized behavior

Take action

Review detailed findings in the console, integrate into event management or workflow systems, or trigger AWS Lambda for automated remediation or prevention

Detections

“Threat Purposes” (Types of Findings)

- **Backdoor** - Compromised AWS resource contacting its **C&C** server.
- **Behavior** - Activity patterns that are different from the established baseline.
- **Cryptocurrency** - Detecting software that is associated with **cryptocurrencies**.
- **Pentest** - Potential attack activity generated by known **pen testing tools**.
- **Persistence** - An IAM user is **behaving differently** from the established baseline.
- **Recon** - Reconnaissance attack underway probing ports, listing users, database tables, etc.
- **Resource Consumption** - An IAM user is **behaving differently from the established baseline** to create new resources, such as EC2 instances.
- **Stealth** - Detects attacks leveraging an **anonymizing** proxy server, disguising the true nature of the activity.
- **Trojan** - Malicious activity associated with certain **Trojan applications**.
- **Unauthorized Access** - A **suspicious activity pattern** by an unauthorized individual.

https://docs.aws.amazon.com/guardduty/latest/ug/guardduty_finding-types.html

Dashboard

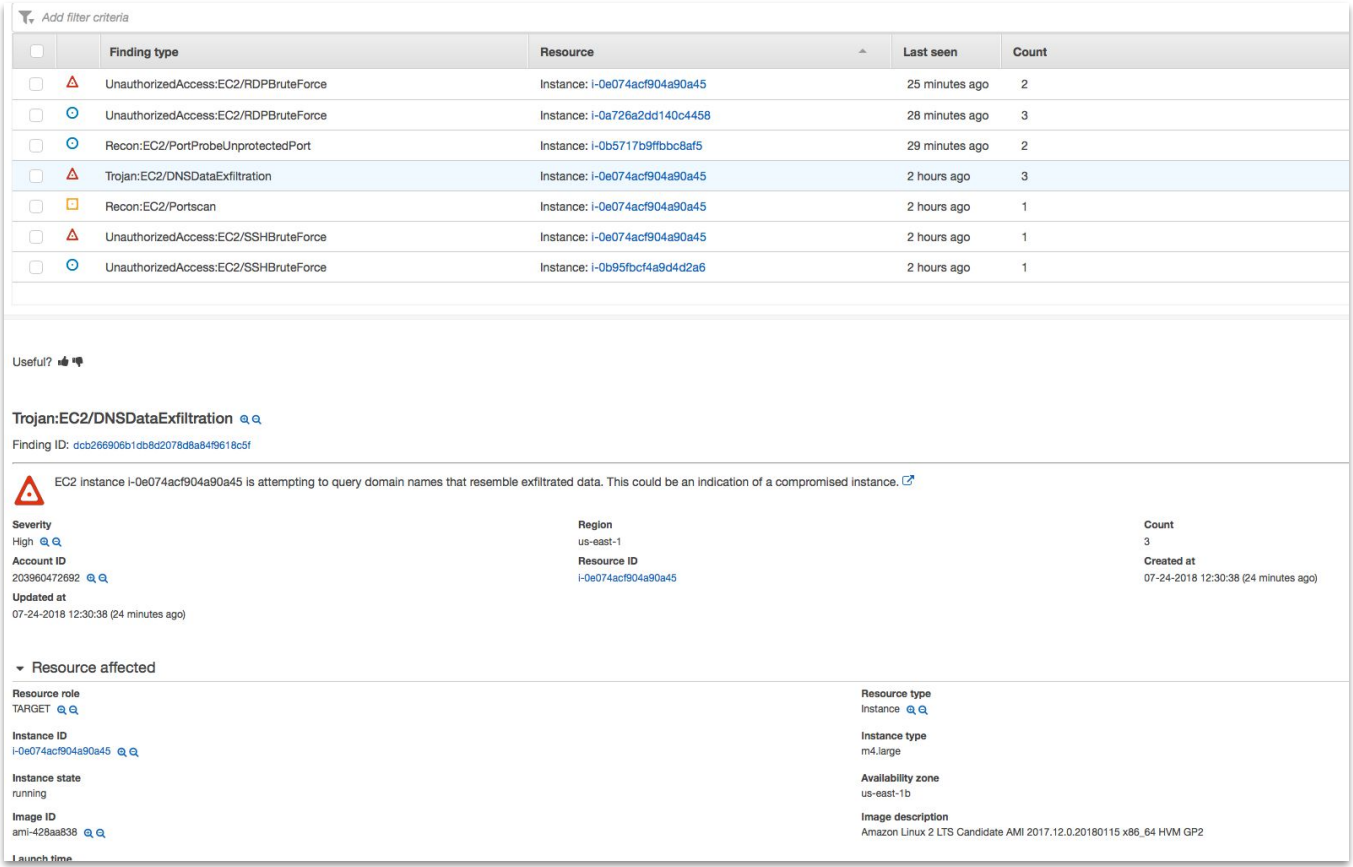
Add filter criteria
Finding type
Resource
Last seen
Count
UnauthorizedAccess:EC2/RDPBruteForce
Instance: i-0e074acf904a90a45
25 minutes ago
2
UnauthorizedAccess:EC2/RDPBruteForce
Instance: i-0a726a2dd140c4458
28 minutes ago
3
Recon:EC2/PortProbeUnprotectedPort
Instance: i-0b5717b9ffbbc8af5
29 minutes ago
2
Trojan:EC2/DNSDataExfiltration
Instance: i-0e074acf904a90a45
2 hours ago
3
Recon:EC2/Portscan
Instance: i-0e074acf904a90a45
2 hours ago
1
UnauthorizedAccess:EC2/SSHBruteForce
Instance: i-0e074acf904a90a45
2 hours ago
1
UnauthorizedAccess:EC2/SSHBruteForce
Instance: i-0b95fbcf4a9d4d2a6
2 hours ago
1
Useful?
Trojan:EC2/DNSDataExfiltration
Finding ID: dcb266906b1db8d2078d8a84f9618c5f
EC2 instance i-0e074acf904a90a45 is attempting to query domain names that resemble exfiltrated data. This could be an indication of a compromised instance.
Severity
High
Region
us-east-1
Count
3
Account ID
203960472692
Resource ID
i-0e074acf904a90a45
Created at
07-24-2018 12:30:38 (24 minutes ago)
Updated at
07-24-2018 12:30:38 (24 minutes ago)
Resource affected
Resource role
TARGET
Resource type
Instance
Instance ID
i-0e074acf904a90a45
Instance type
m4.large
Instance state
running
Availability zone
us-east-1b
Image ID
ami-428aa838
Image description
Amazon Linux 2 LTS Candidate AMI 2017.12.0.20180115 x86_64 HVM GP2
Launch time

Dashboard

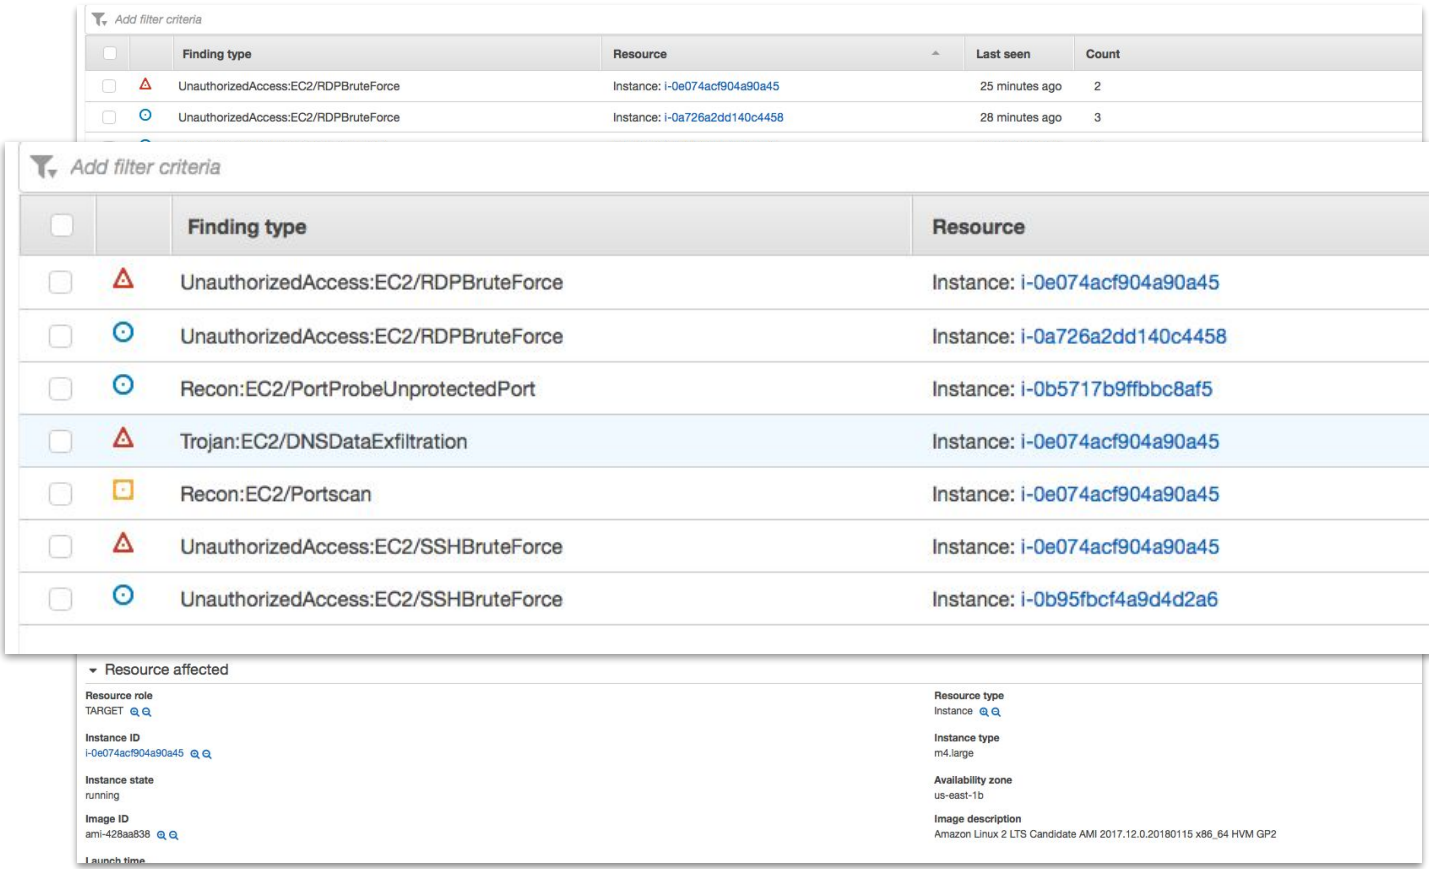

Add filter criteria

Finding type	Resource	Last seen	Count
UnauthorizedAccess:EC2/RDPBruteForce	Instance: i-0e074acf904a90a45	25 minutes ago	2
UnauthorizedAccess:EC2/RDPBruteForce	Instance: i-0a726a2dd140c4458	28 minutes ago	3
Recon:EC2/PortProbeUnprotectedPort	Instance: i-0b5717b9ffbbc8af5	29 minutes ago	2
Trojan:EC2/DNSDataExfiltration	Instance: i-0e074acf904a90a45	2 hours ago	3
Recon:EC2/Portscan	Instance: i-0e074acf904a90a45	2 hours ago	1
UnauthorizedAccess:EC2/SSHBruteForce	Instance: i-0e074acf904a90a45	2 hours ago	1
UnauthorizedAccess:EC2/SSHBruteForce	Instance: i-0b95fbcf4a9d4d2a6	2 hours ago	1

Resource affected

Resource role TARGET	**Resource type** Instance
Instance ID i-0e074acf904a90a45	**Instance type** m4.large
Instance state running	**Availability zone** us-east-1b
Image ID ami-428aa838	**Image description** Amazon Linux 2 LTS Candidate AMI 2017.12.0.20180115 x86_64 HVM GP2
Launch time	

Dashboard - Finding

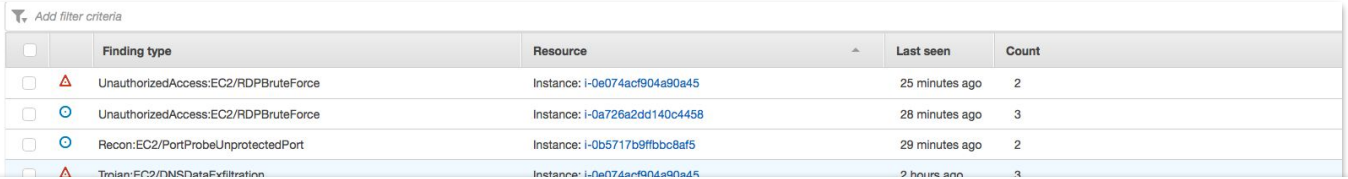

Trojan:EC2/DNSDataExfiltration

Finding ID: dcb266906b1db8d2078d8a84f9618c5f

EC2 instance i-0e074acf904a90a45 is attempting to query domain names that resemble exfiltrated data. This could be an indication of a compromised instance.

Severity
High

Region
us-east-1

Count
49

Account ID
203960472692

Resource ID
i-0e074acf904a90a45

Created at
07-24-2018 09:30:38 (15 days ago)

Updated at
07-24-2018 09:30:38 (15 days ago)

Dashboard - Finding

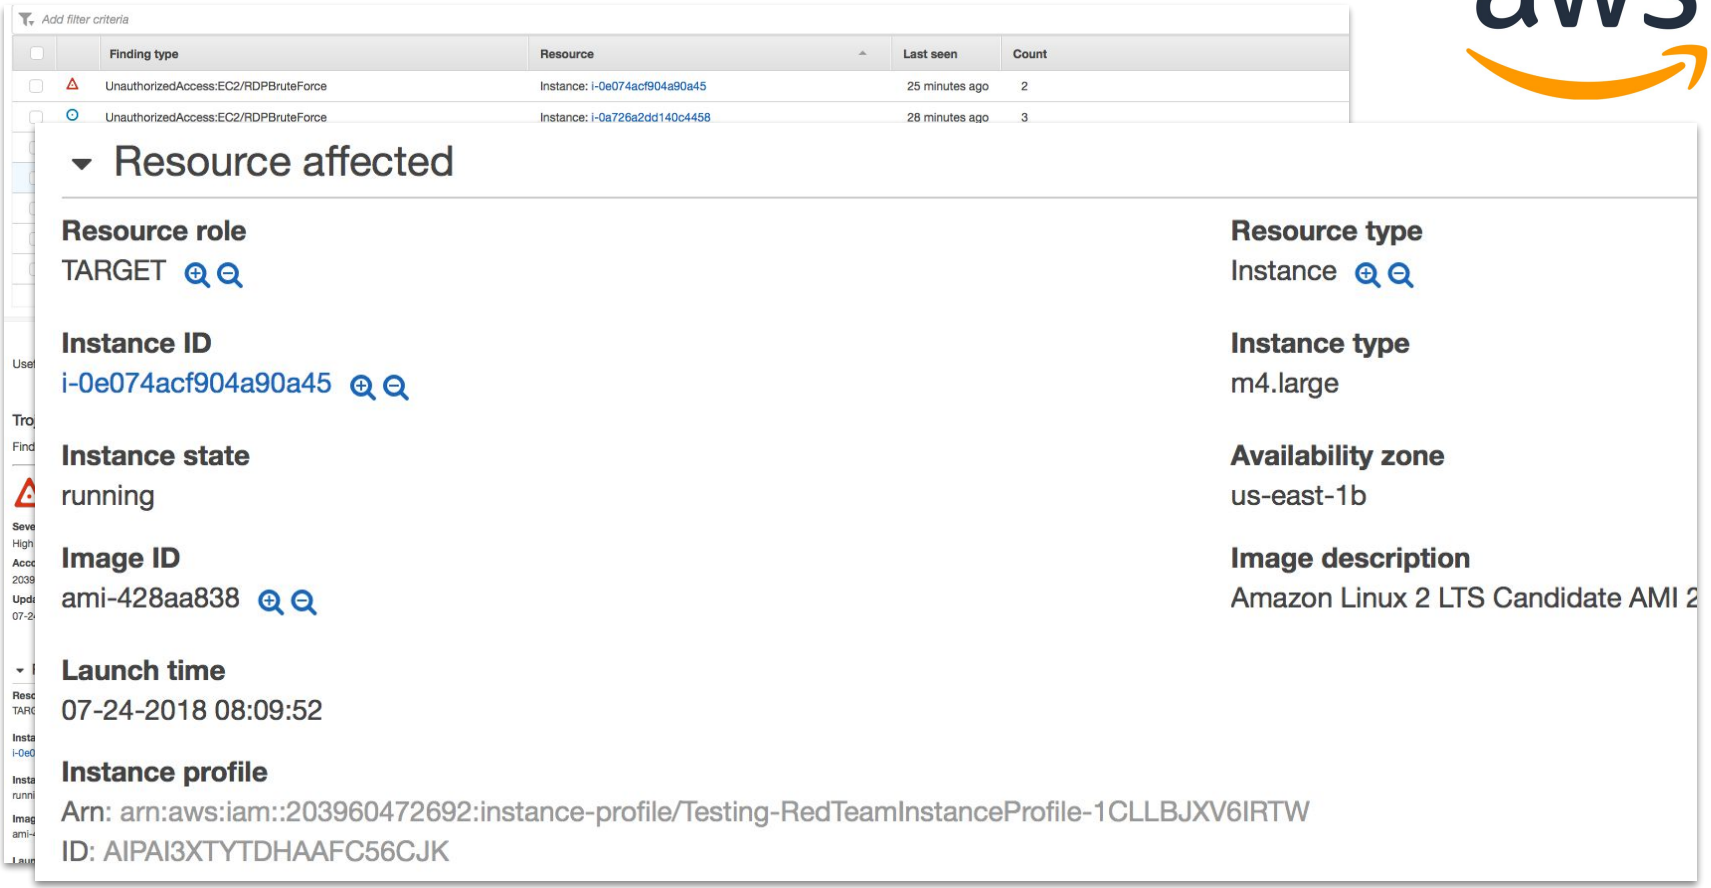

Dashboard - Finding

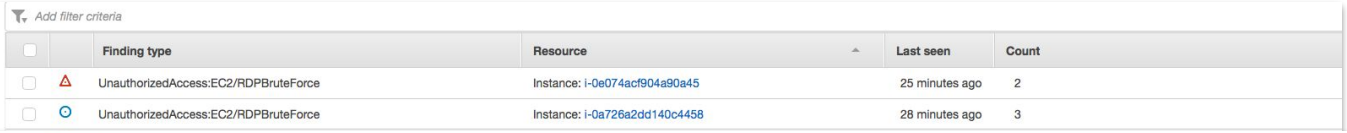

Action

Action type
DNS_REQUEST

Protocol
0

Blocked
false

First seen
07-24-2018 08:18:28 (15 days ago)

Last seen
07-24-2018 08:18:28 (15 days ago)

Actor

Domain
0qhbrenpjyobgcbvgighmcmbwyfo9dwaffxmqmx5cwcb1besgbs252n3gccydo...

Demo!

Rhino Security Labs - Cloud Goat (Slightly Modified)

Safe, practice environment for learning how to collect keys, move laterally, escalate privileges, and more.

- Blog https://rhinosecuritylabs.com/aws/cloudgoat-vulnerable-design-aws-environment/
- Code https://github.com/RhinoSecurityLabs/cloudgoat

Demo!

Attack Path Steps

1. Server-Side Request Forgery to steal EC2 IAM instance credentials
2. Enumerate API access unsuccessfully
3. Escalate to Administrator using `attach-role-policy`
4. Enumerate API access successfully
5. Find and exfiltrate PII from an S3 bucket
6. Add a permanent Administrative user
7. Cover our tracks
8. Review the IAM, Cloudtrail, and GuardDuty logs/alerts

Value Added

Zero-Impact Setup

- Nearly a “one-click” installation process.

Clear Listing of GuardDuty Detections

- You know what AWS is monitoring for you.

Broad Partner Ecosystem

- Many options to choose from in many different areas of security, not just detection.

Detects Multiple Forms of API Misuse

- Several key detections for behaviors associated with compromised credentials.

Areas for Improvement

Areas for Improvement

- Ability to **tune** parameters for all settings and detections
- Ability to add **custom detections** into the native analytics engine/flow
- API ability to create **custom findings**, not just view them.
- **Unified security dashboard** and workflow for all AWS Security services
  - AWS Config
  - AWS Inspector
  - AWS CloudTrail
  - AWS GuardDuty

Google Cloud Security Command Center

Google Cloud Security Command Center

Released

- Google StackDriver: Spring 2016
- Google Cloud VPC Flow logs: Spring 2018
- Google Cloud Security Command Center (Alpha): Spring 2018

Description

- The Cloud Security Command Center (Cloud SCC) is the canonical security and data risk database for Google Cloud Platform (GCP). Cloud SCC enables you to **understand your security and data attack surface** by providing asset inventory, discovery, search, and management.

Links and Documentation

- https://cloud.google.com/security-command-center/

Key Features

Asset Discovery/Inventory

- Across App Engine, Compute Engine, Cloud Storage, and Cloud Datastore

Anomaly Detection

- Identifies threats like **botnets**, **cryptocurrency** mining, anomalous reboots, and **suspicious network traffic**.
- Cost: Unknown. Free during Alpha period.

Centralized Finding Dashboard

- Web application vulnerability scans - Cloud Security Scanner
- Sensitive data on storage bucket scans - DLP API
- Access control and policy scans - Forseti
- All third party security solution findings/results

Key Features (Cont’d)

Real-Time Notifications

- Receive Cloud SCC alerts via Gmail, SMS, and Jira with Cloud Pub/Sub notification integration.

REST API

- Integration with your existing security systems and workflows.

Detection Data Flow

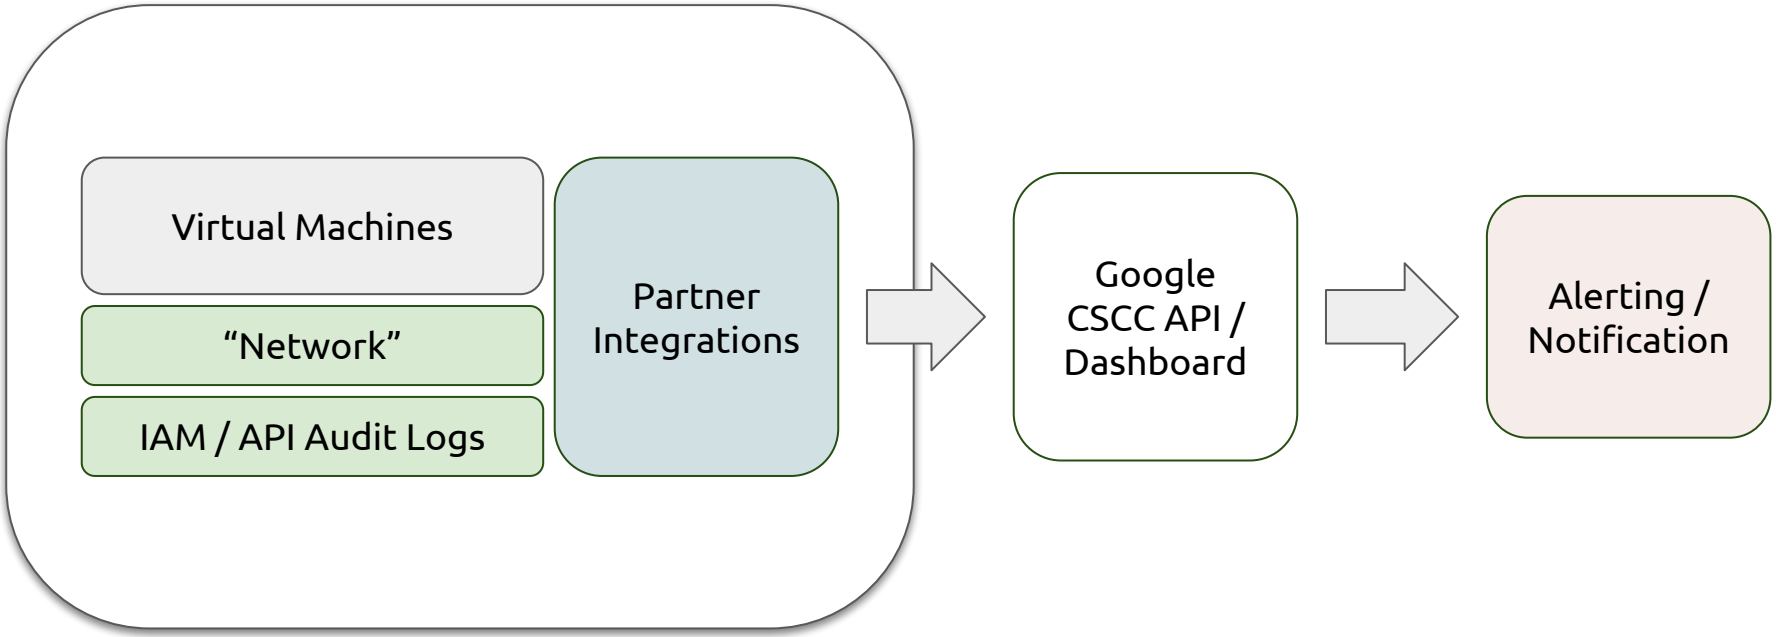

Detections

As Listed but not Detailed

- Botnets
- Cryptocurrency mining
- Anomalous reboots
- Suspicious/anomalous network traffic

Dashboard

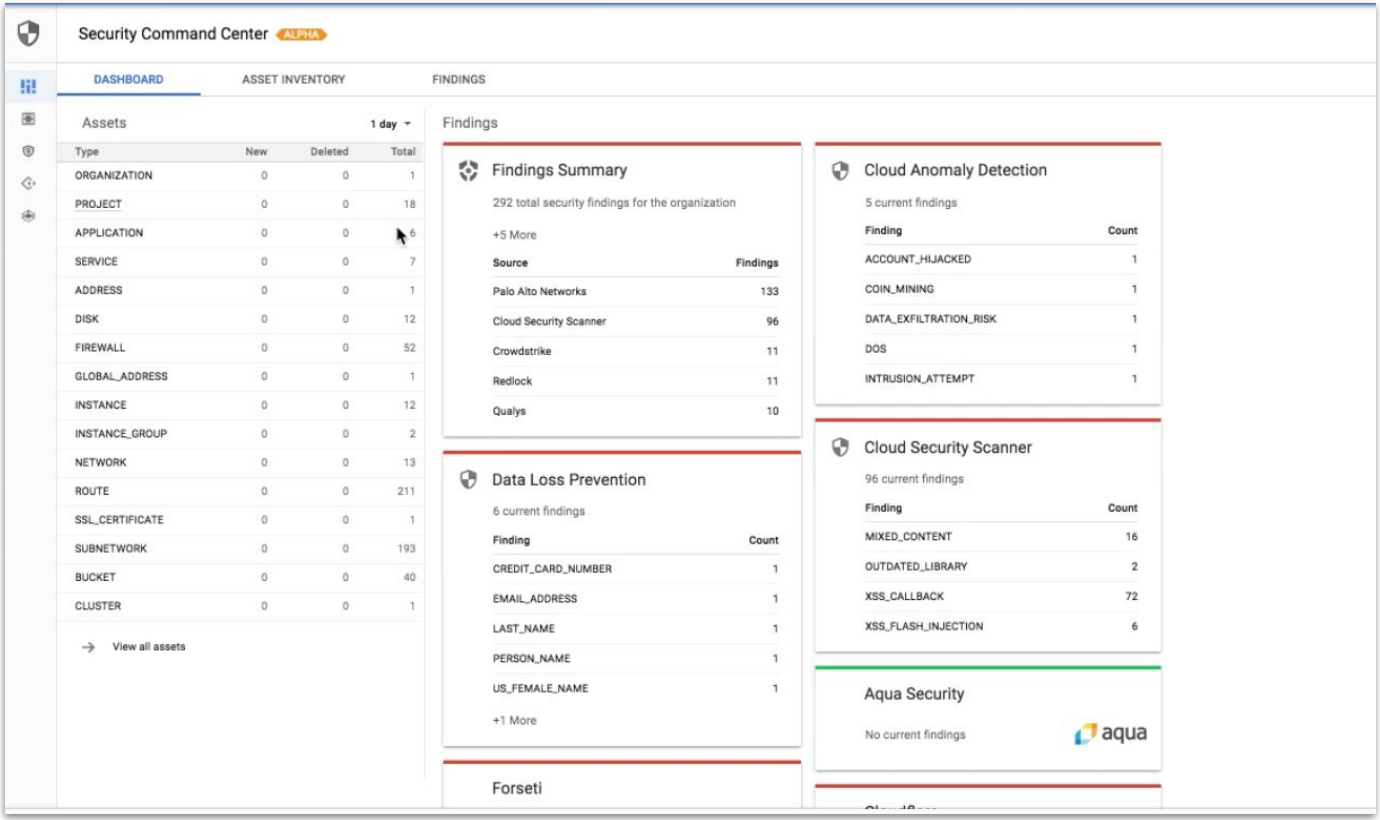
Security Command Center
ALPHA
DASHBOARD
ASSET INVENTORY
FINDINGS
Assets
1 day
Type New Deleted Total
ORGANIZATION 0 0 1
PROJECT 0 0 18
APPLICATION 0 0 6
SERVICE 0 0 7
ADDRESS 0 0 1
DISK 0 0 12
FIREWALL 0 0 52
GLOBAL_ADDRESS 0 0 1
INSTANCE 0 0 12
INSTANCE_GROUP 0 0 2
NETWORK 0 0 13
ROUTE 0 0 211
SSL_CERTIFICATE 0 0 1
SUBNETWORK 0 0 193
BUCKET 0 0 40
CLUSTER 0 0 1
View all assets
Findings
Findings Summary
292 total security findings for the organization
+5 More
Source Findings
Palo Alto Networks 133
Cloud Security Scanner 96
Crowdstrike 11
Redlock 11
Qualys 10
Data Loss Prevention
6 current findings
Finding Count
CREDIT_CARD_NUMBER 1
EMAIL_ADDRESS 1
LAST_NAME 1
PERSON_NAME 1
US_FEMALE_NAME 1
+1 More
Forseti
Cloud Anomaly Detection
5 current findings
Finding Count
ACCOUNT_HIJACKED 1
COIN_MINING 1
DATA_EXFILTRATION_RISK 1
DOS 1
INTRUSION_ATTEMPT 1
Cloud Security Scanner
96 current findings
Finding Count
MIXED_CONTENT 16
OUTDATED_LIBRARY 2
XSS_CALLBACK 72
XSS_FLASH_INJECTION 6
Aqua Security
No current findings
aqua

Dashboard

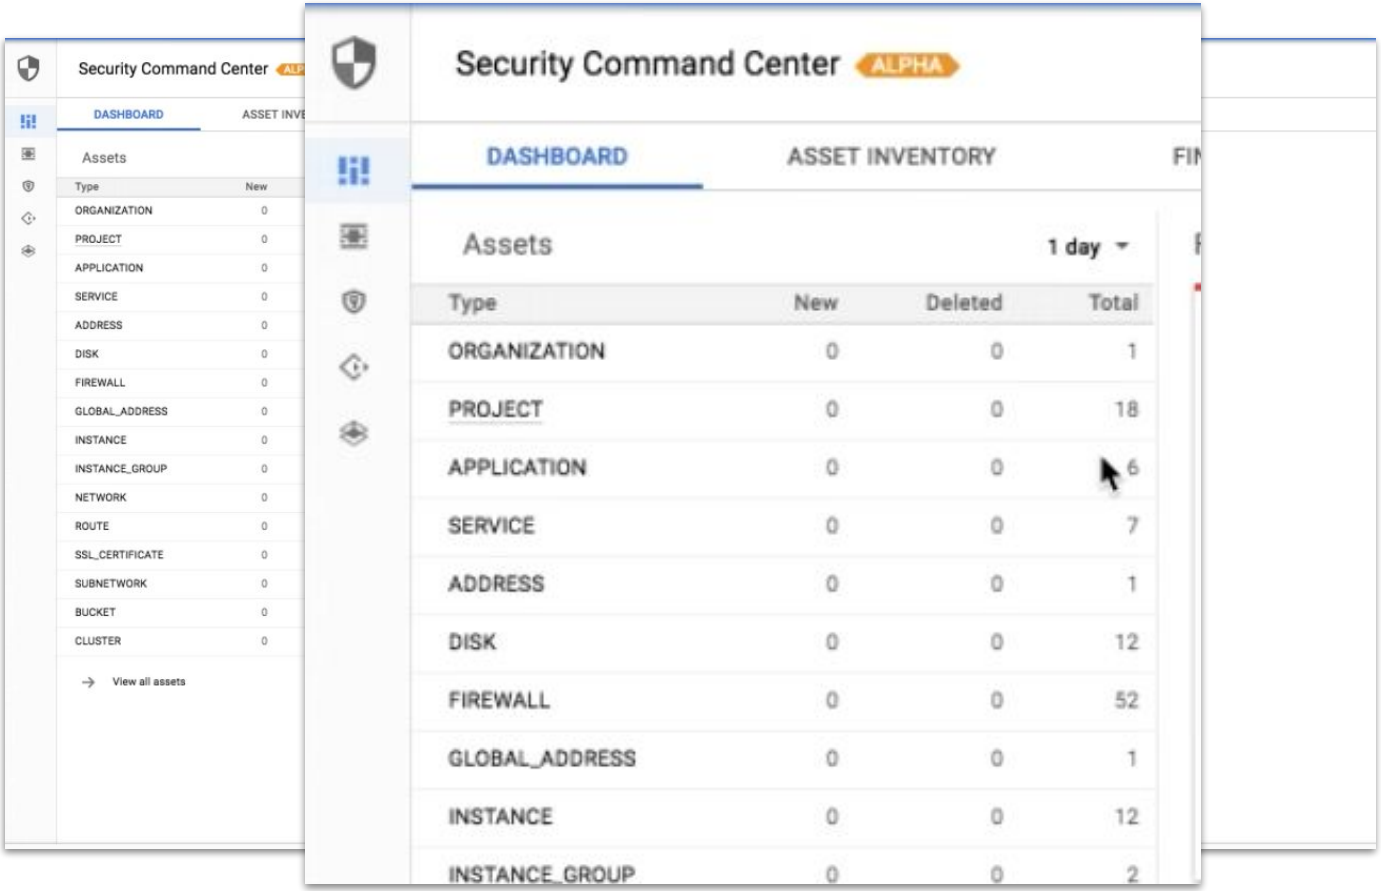
Security Command Center
DASHBOARD
ASSET INVENTORY
Assets
1 day
Type
New
Deleted
Total
ORGANIZATION 0 0 1
PROJECT 0 0 18
APPLICATION 0 0 6
SERVICE 0 0 7
ADDRESS 0 0 1
DISK 0 0 12
FIREWALL 0 0 52
GLOBAL_ADDRESS 0 0 1
INSTANCE 0 0 12
INSTANCE_GROUP 0 0 2
Security Command Center
DASHBOARD
Assets
Type
New
ORGANIZATION 0
PROJECT 0
APPLICATION 0
SERVICE 0
ADDRESS 0
DISK 0
FIREWALL 0
GLOBAL_ADDRESS 0
INSTANCE 0
INSTANCE_GROUP 0
NETWORK 0
ROUTE 0
SSL_CERTIFICATE 0
SUBNETWORK 0
BUCKET 0
CLUSTER 0
View all assets

Dashboard

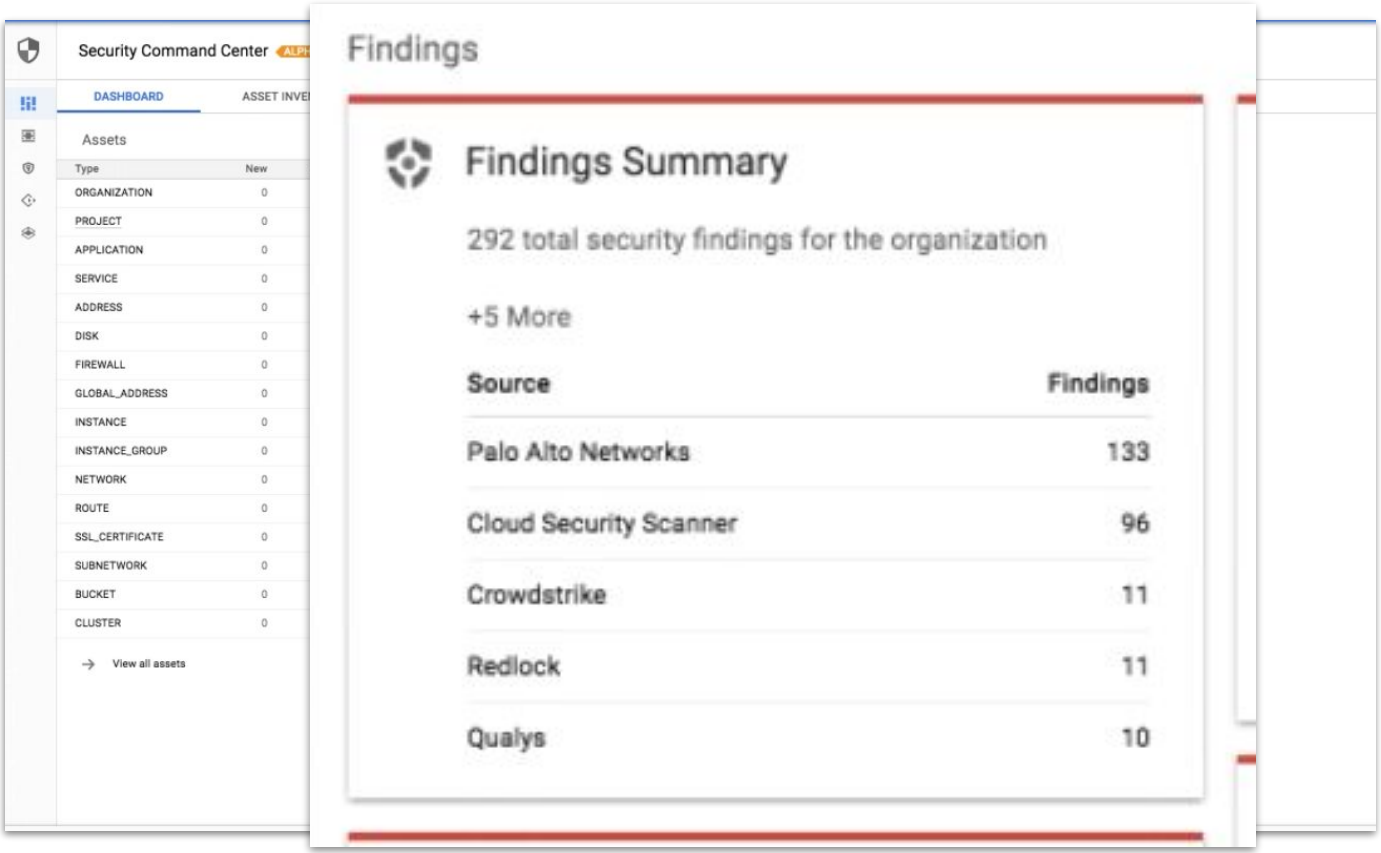
Security Command Center
DASHBOARD
ASSET INVE
Assets
Type
New
ORGANIZATION 0
PROJECT 0
APPLICATION 0
SERVICE 0
ADDRESS 0
DISK 0
FIREWALL 0
GLOBAL_ADDRESS 0
INSTANCE 0
INSTANCE_GROUP 0
NETWORK 0
ROUTE 0
SSL_CERTIFICATE 0
SUBNETWORK 0
BUCKET 0
CLUSTER 0
View all assets
Findings
Findings Summary
292 total security findings for the organization
+5 More
Source
Findings
Palo Alto Networks 133
Cloud Security Scanner 96
Crowdstrike 11
Redlock 11
Qualys 10

Dashboard

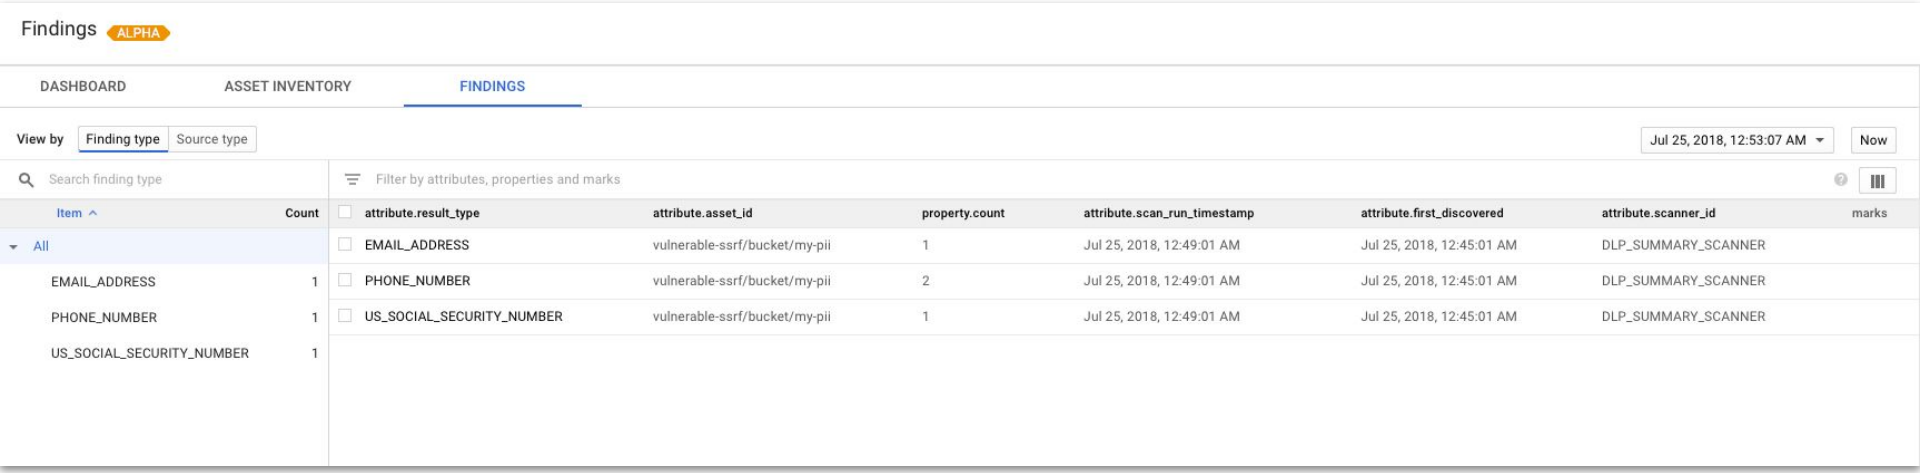
Findings
ALPHA
DASHBOARD
ASSET INVENTORY
FINDINGS
View by
Finding type
Source type
Jul 25, 2018, 12:53:07 AM
Now
Search finding type
Filter by attributes, properties and marks
Item
Count
All
EMAIL_ADDRESS 1
PHONE_NUMBER 1
US_SOCIAL_SECURITY_NUMBER 1
attribute.result_type
attribute.asset_id
property.count
attribute.scan_run_timestamp
attribute.first_discovered
attribute.scanner_id
marks
EMAIL_ADDRESS vulnerable-ssrf/bucket/my-pii 1 Jul 25, 2018, 12:49:01 AM Jul 25, 2018, 12:45:01 AM DLP_SUMMARY_SCANNER
PHONE_NUMBER vulnerable-ssrf/bucket/my-pii 2 Jul 25, 2018, 12:49:01 AM Jul 25, 2018, 12:45:01 AM DLP_SUMMARY_SCANNER
US_SOCIAL_SECURITY_NUMBER vulnerable-ssrf/bucket/my-pii 1 Jul 25, 2018, 12:49:01 AM Jul 25, 2018, 12:45:01 AM DLP_SUMMARY_SCANNER

Dashboard

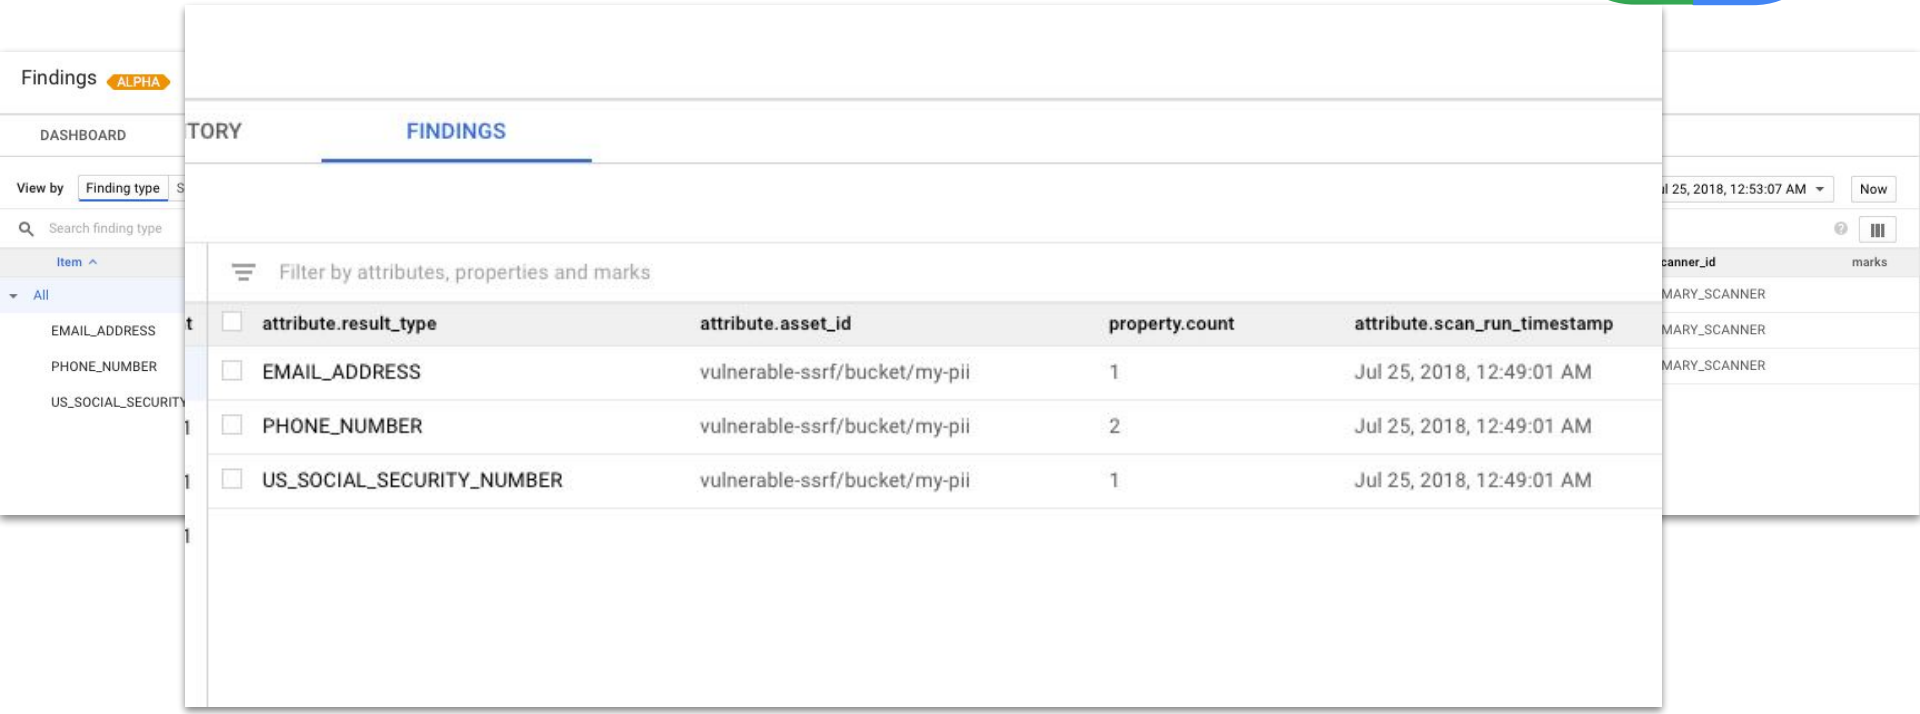
Findings ALPHA
DASHBOARD
TORY
FINDINGS
View by
Finding type
Search finding type
Item
All
EMAIL_ADDRESS
PHONE_NUMBER
US_SOCIAL_SECURITY
Filter by attributes, properties and marks
attribute.result_type
attribute.asset_id
property.count
attribute.scan_run_timestamp
EMAIL_ADDRESS
vulnerable-ssrf/bucket/my-pii
1
Jul 25, 2018, 12:49:01 AM
PHONE_NUMBER
vulnerable-ssrf/bucket/my-pii
2
Jul 25, 2018, 12:49:01 AM
US_SOCIAL_SECURITY_NUMBER
vulnerable-ssrf/bucket/my-pii
1
Jul 25, 2018, 12:49:01 AM
Jul 25, 2018, 12:53:07 AM
Now
marks

Dashboard

Finding Details

Summary

Finding type	First discovered	Most recently seen
US_SOCIAL_SECURITY_NUMBER	Jul 25, 2018 12:45 AM (8 minutes ago)	Jul 25, 2018 12:49 AM (4 minutes ago)

Security marks

No marks

Attributes

Asset Id	vulnerable-ssrf/bucket/my-pii
Configuration Id	bf32ed7a46ef88f385768ee87b816e6b5b9b5e893edc93ee8e05693be4d15a54
First Discovered	July 25, 2018 at 12:45:01 AM UTC-4
Id	80fbe57de6ce511e0095a90d0a7f502f992dd8b9bb570f329a0f1358806a6805
Result Type	US_SOCIAL_SECURITY_NUMBER
Scan Run Id	projects/vulnerable-ssrf/dlpJobs/i-5405955858454632083
Scan Run Timestamp	July 25, 2018 at 12:49:01 AM UTC-4
Scanner Id	DLP_SUMMARY_SCANNER
Update Time	July 25, 2018 at 12:48:56 AM UTC-4

Properties

Count	1

Dashboard

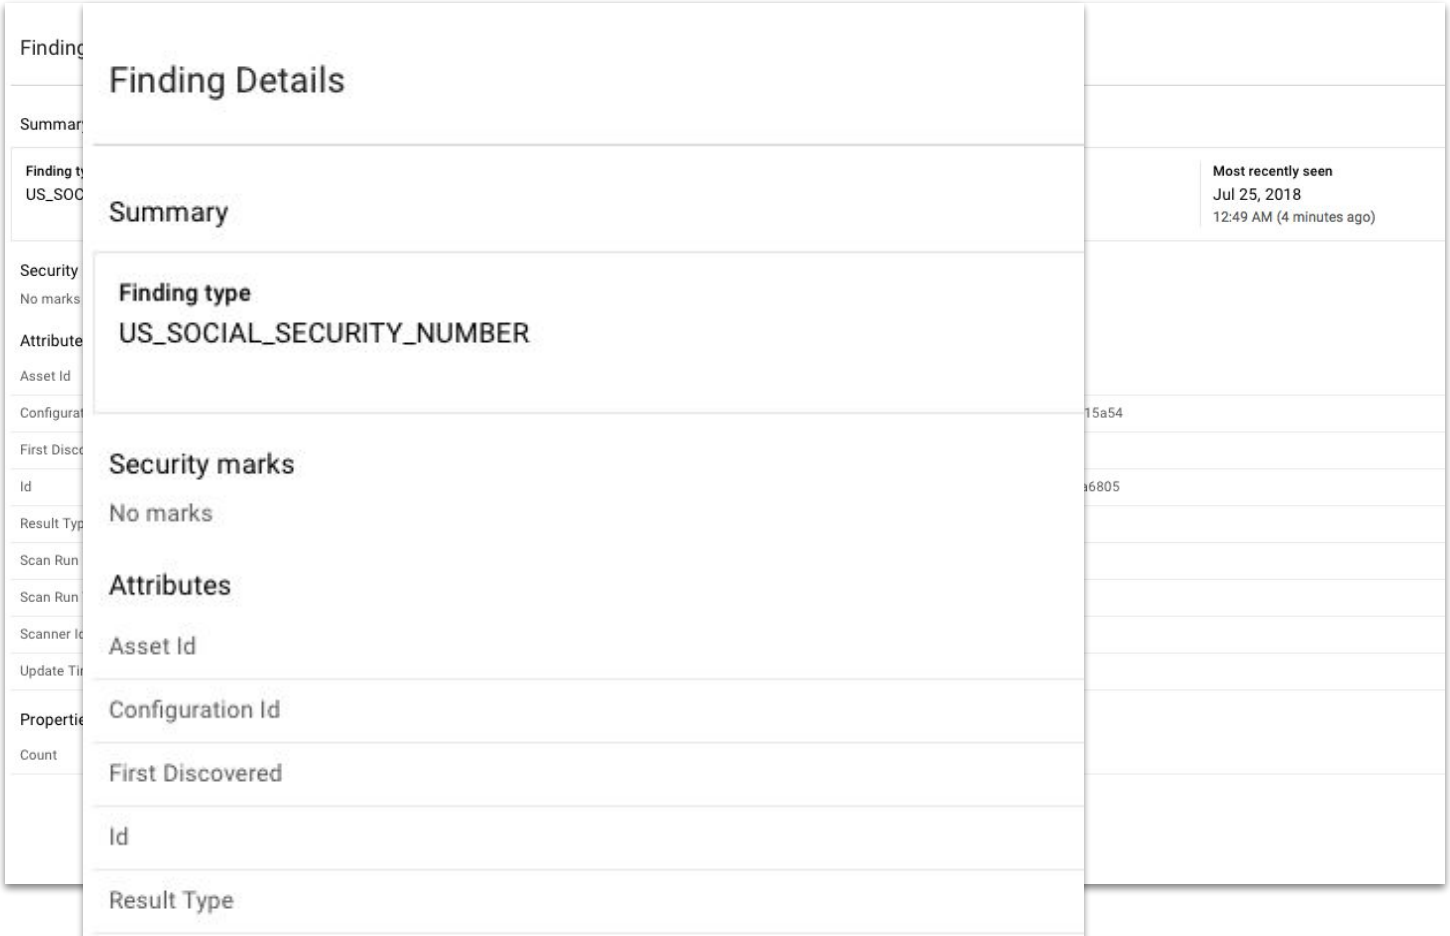
Finding Details
Summary
Finding type
US_SOCIAL_SECURITY_NUMBER
Security marks
No marks
Attributes
Asset Id
Configuration Id
First Discovered
Id
Result Type
Most recently seen
Jul 25, 2018
12:49 AM (4 minutes ago)

Dashboard

Finding Details

Summary

Finding type	First discovered	Most recently seen

Security marks

No marks

Attributes

Asset Id	vulnerable-ssrf/bucket/my-pii
Configuration Id	bf32ed7a46ef88f385768ee87b816e6b5b9b5e893edc93ee8e05693be4d15a54
First Discovered	July 25, 2018 at 12:45:01 AM UTC-4
Id	80fbe57de6ce511e0095a90d0a7f502f992dd8b9bb570f329a0f1358806a6805
Result Type	US_SOCIAL_SECURITY_NUMBER
Scan Run Id	projects/vulnerable-ssrf/dlpJobs/i-5405955858454632083
Scan Run Timestamp	July 25, 2018 at 12:49:01 AM UTC-4
Scanner Id	DLP_SUMMARY_SCANNER
Update Time	July 25, 2018 at 12:48:56 AM UTC-4

Properties

Value Added

Zero-Impact Setup

- Setup does not affect any running workflows.

Partner Focus

- The API and Interface feature partner solutions and integrate their output streams into a single management interface.

Framework-Oriented

- Similar to the Stackdriver logging service in that it’s a framework for handling all security events across all applicable services.

Limitations and Suggestions

Limitations

- Still in Alpha, so anomalous detection capabilities are still in the early stages.
- Not yet a comprehensive or detailed list of detection capabilities.

Suggestions

- Ability to **tune** all settings and detections
- Ability to add **custom detections** into the native flow
- Integrated security detections for all managed GCP services
- Integrate **native notification and alerting** functionality

Key Takeaways and Looking Ahead

Common Areas for Improvement

Detections

- Visibility dependent on implementation
- Detection capability listings
- Customization / Tuning
- ML/AI in use, but how exactly?

Integrations

- Wide range of ease of integration
- A small selection of vendors that integrate natively into the new services.

Education

- Clearer guidelines needed.

Are the provider-native threat detection services all I need?

Should I Adopt These Services Now?

The *Framework* is Important

Wherever possible, avoid
undifferentiated heavy lifting

Watch this space closely

Security solution vendors -- Take note

Additional Learning and Exploration

Cloud Goat - Rhino Security Labs

- https://github.com/RhinoSecurityLabs/cloudgoat
- "Vulnerable by Design" AWS infrastructure setup and testing environment

FlAWS.Cloud - Scott Piper

- http://flaws.cloud

Detecting Credential Compromise in AWS - Will Bengston

- https://www.blackhat.com/us-18/briefings.html#detecting-credential-compromise-in-aws

Cloud Security Trends Reports

- https://info.redlock.io/cloud-security-trends-may2018
- https://start.paloaltonetworks.com/cloud-security-report-2018


black hat®
USA 2018
AUGUST 4-9, 2018
MANDALAY BAY / LAS VEGAS
Thank you!
Questions?
@bradgeesaman
#BHUSA / @BLACK HAT EVENTS