

Keep it Flexible: How Cloud Makes it Easier and Harder to Detect Bad Stuff

Lily Lee | Staff Security Specialist

April 29, 2019



Forward-Looking Statements

During the course of this presentation, we may make forward-looking statements regarding future events or the expected performance of the company. We caution you that such statements reflect our current expectations and estimates based on factors currently known to us and that actual events or results could differ materially. For important factors that may cause actual results to differ from those contained in our forward-looking statements, please review our filings with the SEC.

The forward-looking statements made in this presentation are being made as of the time and date of its live presentation. If reviewed after its live presentation, this presentation may not contain current or accurate information. We do not assume any obligation to update any forward-looking statements we may make. In addition, any information about our roadmap outlines our general product direction and is subject to change at any time without notice. It is for informational purposes only and shall not be incorporated into any contract or other commitment. Splunk undertakes no obligation either to develop the features or functionality described or to include any such feature or functionality in a future release.

Splunk, Splunk>, Listen to Your Data, The Engine for Machine Data, Splunk Cloud, Splunk Light and SPL are trademarks and registered trademarks of Splunk Inc. in the United States and other countries. All other brand names, product names, or trademarks belong to their respective owners. © 2017 Splunk Inc. All rights reserved.

#whoami > Lily Lee

GIAC GCIH




Lily | lily@splunk.com
Staff Security Specialist

- ▶ Based out of Splunk HQ in San Francisco
- ▶ 15+ years in IT and security
- ▶ Work with Fortune 500 companies, government agencies, education
- ▶ Focus on security and Splunk for security
 - Including Splunk and AWS
- ▶ Founding member of WiCyS Silicon Valley

Goals

1. Understand what AWS services your on-premises technologies correspond to
2. Understand what data is security-relevant; and where and how to get that data
3. Understand how that data can be used to detect malicious activity

130.60.4 - [07/Jan 18:10:57:153] "GET /category.screen?category_id=GLFT&SESSIONID=SD15L4FF10ADFF10 HTTP 1.1" 404 720 "http://buttercup-shopping.com/cart.do?action=view&itemId=EST-6&product_id=FE-5W63-4069/2-20
128.241.230.82 - [07/Jan 18:10:57:123] "GET /product.screen?product_id=FL-DSH-01&SESSIONID=SD55L7FFGADFF9 HTTP 1.1" 404 3322 "http://buttercup-shopping.com/category.screen?category_id=GLFT&SESSIONID=SD15L4FF10ADFF10 HTTP 1.1" 404 720 "http://buttercup-shopping.com/cart.do?action=purchase&itemId=EST-6&product_id=FE-5W63-4069/2-20
317.27.160.0.0 - [07/Jan 18:10:57:123] "GET /product.screen?product_id=FL-DSH-01&SESSIONID=SD55L7FFGADFF9 HTTP 1.1" 200 1318 "http://buttercup-shopping.com/cart.do?action=purchase&itemId=EST-1&product_id=FE-5W63-4069/2-20
www.nts.1: svt: [07/Jan 18:10:57:153] "GET /category.screen?category_id=GLFT&SESSIONID=SD15L4FF10ADFF10 HTTP 1.1" 404 720 "http://buttercup-shopping.com/category.screen?category_id=GLFT&SESSIONID=SD15L4FF10ADFF10 HTTP 1.1" 404 720 "http://buttercup-shopping.com/category.screen?category_id=GLFT&SESSIONID=SD15L4FF10ADFF10 HTTP 1.1" 404 720 "http://buttercup-shopping.com/cart.do?action=purchase&itemId=EST-6&product_id=FE-5W63-4069/2-20
/buttercup-EST-1&product_id=RP-LI-02" 468 125.17 "GET /oldlink?item_id=EST-26&SESSIONID=SD55L9FF1ADFF3 HTTP 1.1" 200 1318 "http://buttercup-shopping.com/cart.do?action=purchase&itemId=EST-1&product_id=FE-5W63-4069/2-20
/buttercup-shopping.com/n-02" 468 125.17 "GET /category.screen?category_id=GLFT&SESSIONID=SD15L4FF10ADFF10 HTTP 1.1" 404 720 "http://buttercup-shopping.com/category.screen?category_id=GLFT&SESSIONID=SD15L4FF10ADFF10 HTTP 1.1" 404 720 "http://buttercup-shopping.com/category.screen?category_id=GLFT&SESSIONID=SD15L4FF10ADFF10 HTTP 1.1" 404 720 "http://buttercup-shopping.com/cart.do?action=purchase&itemId=EST-6&product_id=FE-5W63-4069/2-20

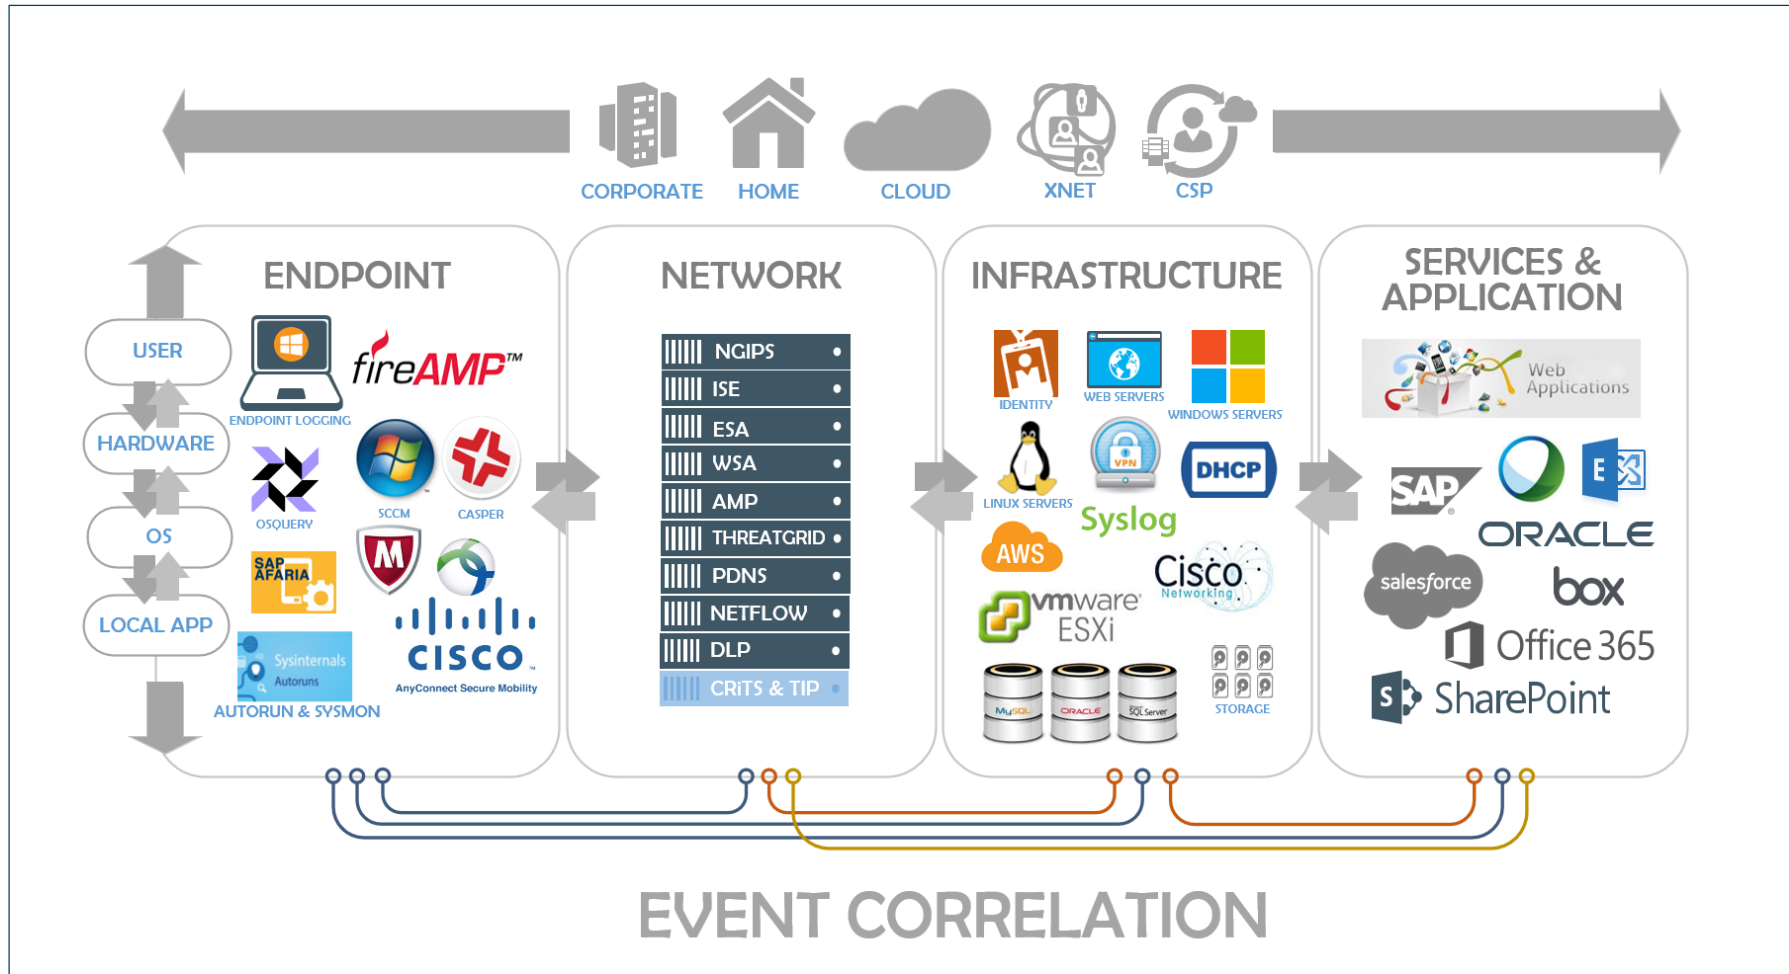


On-prem vs. AWS

Infrastructure comparison



“We’re moving to the cloud, but we don’t know where to get security data.”



Environment Visibility

Data parity

AWS SECURITY DATA SOURCES

ENDPOINT

NETWORK

INFRASTRUCTURE

SERVICES & APPS

ON PREM DATA CENTER

PHYSICAL RESOURCES



ANALYTICS ENGINE

INFRASTRUCTURE AS A SERVICE



VIRTUALIZED RESOURCES



INFRASTRUCTURE



APIs



Environment Visibility

Data parity

AWS SECURITY DATA SOURCES

ENDPOINT

NETWORK

INFRASTRUCTURE

SERVICES & APPS



ON PREM DATA CENTER


PHYSICAL RESOURCES



ANALYTICS ENGINE

INFRASTRUCTURE AS A SERVICE

VIRTUALIZED RESOURCES



INFRASTRUCTURE



Environment Visibility

Data parity

AWS SECURITY DATA SOURCES

ENDPOINT

NETWORK

INFRASTRUCTURE

SERVICES & APPS



ON PREM DATA CENTER

PHYSICAL RESOURCES



ANALYTICS ENGINE

INFRASTRUCTURE AS A SERVICE

VIRTUALIZED RESOURCES



APIs



INFRASTRUCTURE



Environment Visibility

Data parity

AWS SECURITY DATA SOURCES

ENDPOINT

NETWORK

INFRASTRUCTURE

SERVICES & APPS



ON PREM DATA CENTER

PHYSICAL RESOURCES



ANALYTICS ENGINE

INFRASTRUCTURE AS A SERVICE

VIRTUALIZED RESOURCES



APIs



INFRASTRUCTURE



Environment Visibility

Data parity

AWS SECURITY DATA SOURCES

ENDPOINT

NETWORK

INFRASTRUCTURE

SERVICES & APPS



ON PREM DATA CENTER

PHYSICAL RESOURCES




splunk>

ANALYTICS ENGINE

INFRASTRUCTURE AS A SERVICE

VIRTUALIZED RESOURCES



INFRASTRUCTURE



Environment Visibility

Data parity

AWS SECURITY DATA SOURCES


ENDPOINT



NETWORK



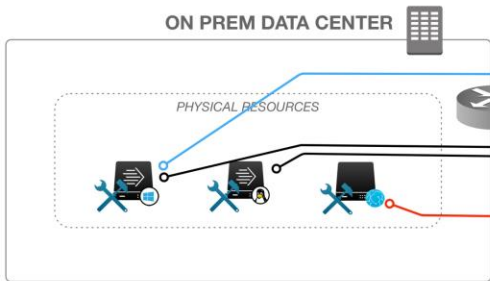
INFRASTRUCTURE



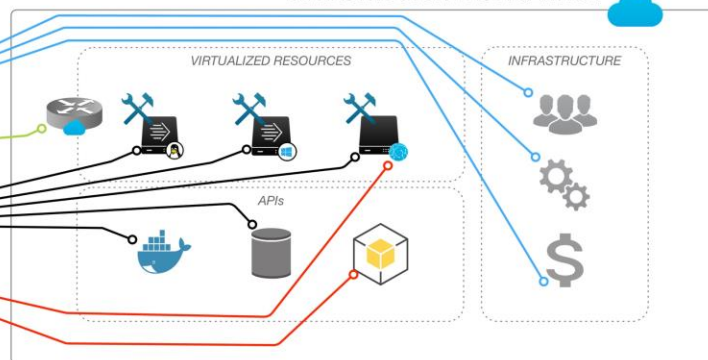
SERVICES & APPS



ON PREM DATA CENTER



INFRASTRUCTURE AS A SERVICE



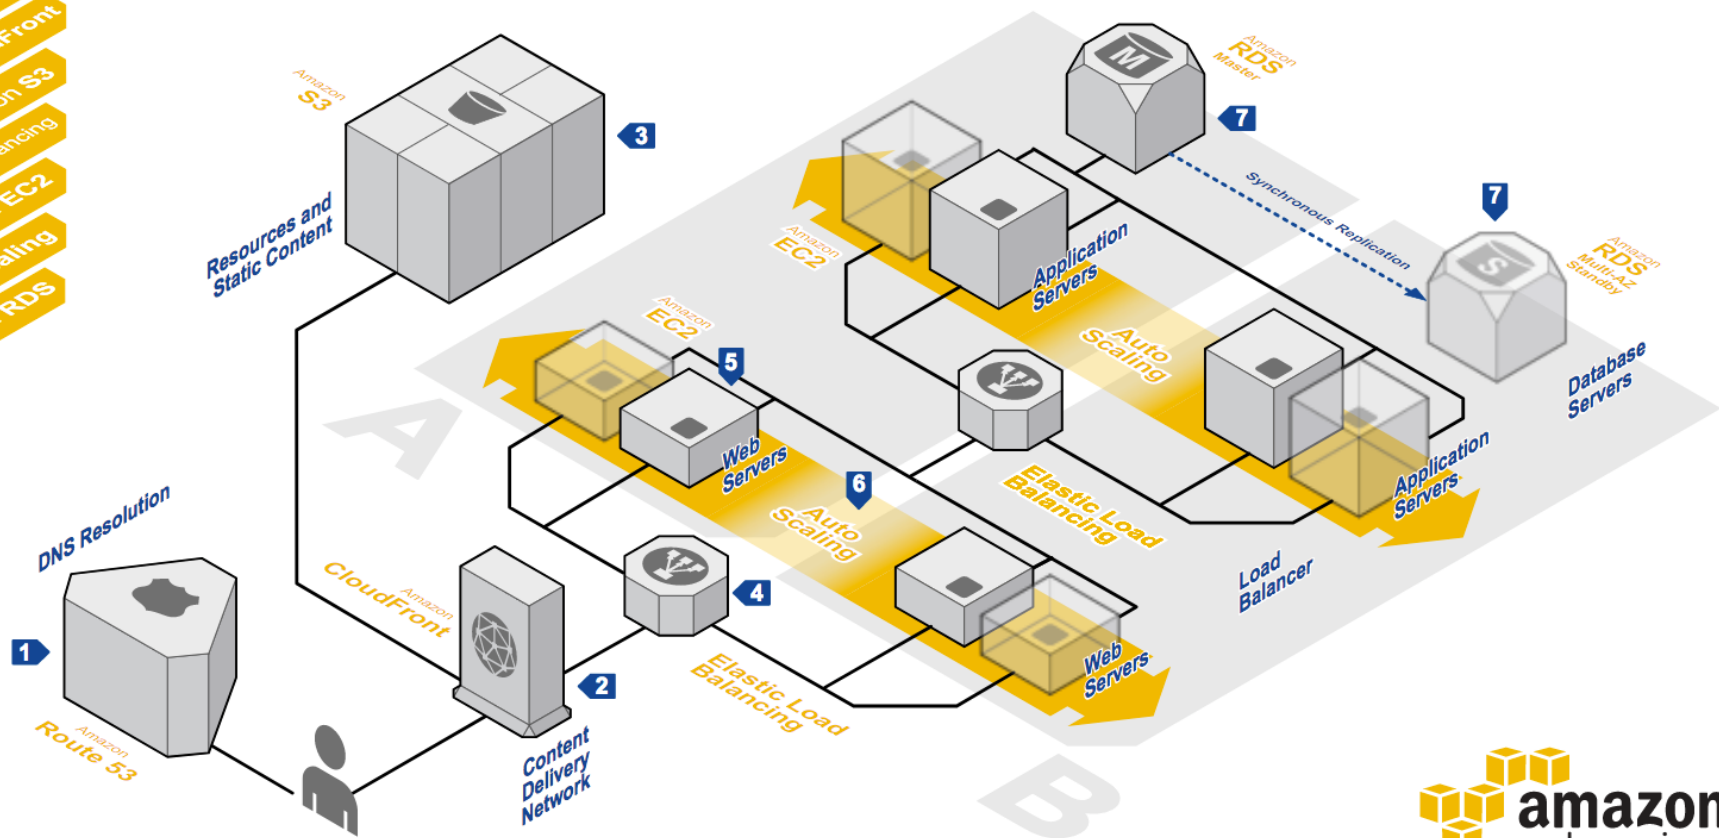
splunk>

ANALYTICS ENGINE

WEB APPLICATION HOSTING

Highly available and scalable web hosting can be complex and expensive. Dense peak periods and wild swings in traffic patterns result in low utilization of expensive hardware. Amazon Web Services provides the reliable, scalable, secure, and high-performance infrastructure required for web applications while enabling an elastic, scale-out and scale-down infrastructure to match IT costs in real time as customer traffic fluctuates.

- AWS Reference Architectures
- Amazon Route 53
- Amazon CloudFront
- Amazon S3
- Elastic Load Balancing
- Amazon EC2
- Auto Scaling
- Amazon RDS



Source: https://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_web_01.pdf



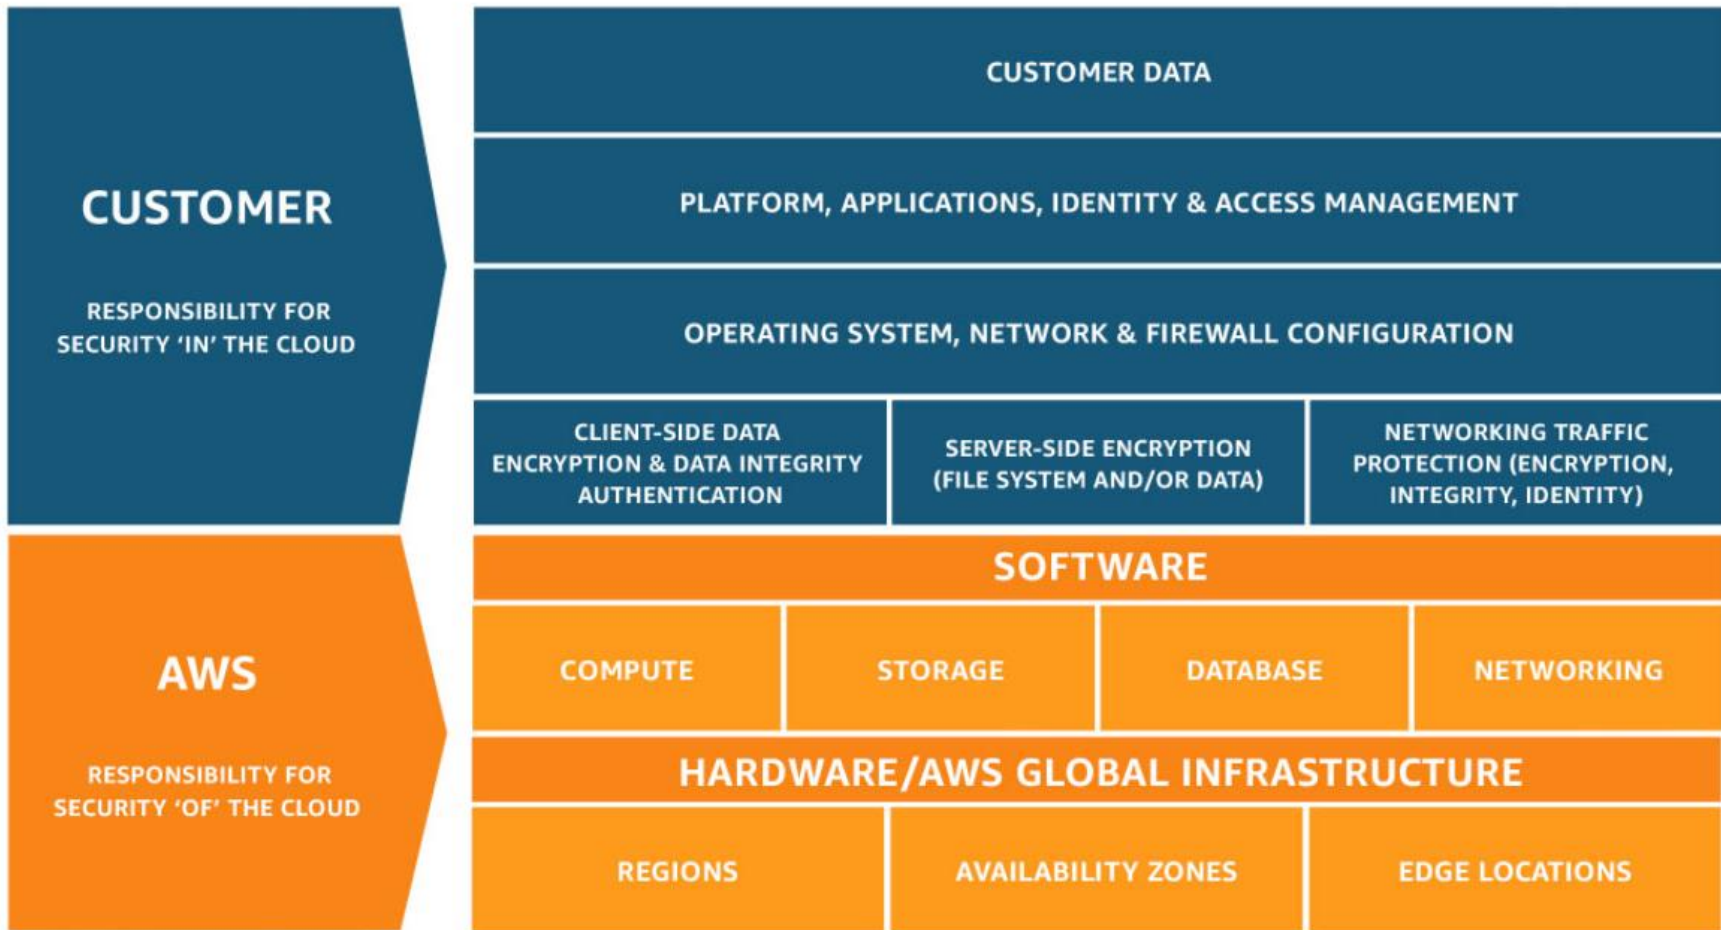
ON-PREMISES INFRASTRUCTURE MAPPED TO AWS

TECHNOLOGY	ON-PREMISES SOLUTION	AWS
Archiving	Tape library, off site tape storage	Glacier
Caching	Memcached, Redis	ElastiCache
Computer	Hardware, virtualization	Elastic Compute Cloud (EC2)
Containers	Docker, Kubernetes	Elastic Container Service (ECS), Elastic Container Service for Kubernetes (EKS)
Content delivery	CDN solutions	CloudFront
Data centers	Data centers	Availability Zones
Data warehousing	Specialized hardware and software solutions	RedShift
Databases	MS SQL, MySQL, Oracle, PostgreSQL	DynamoDB, Relational Database Service (RDS)
Deployment	Ansible, Chef, Fabric, Puppet, SaltStack	Amazon Machine Images (AMIs), CloudFormation, Beanstalk, OpsWorks
Disaster recovery	Multi-site data centers	Multi-region
Domain name services	DNS providers	Route 53
Email	Email software	Simple Email Service (SES)
Identity management	LDAP	Directory Service, Identity and Access Management (IAM)
Load balancing	Hardware and software load balancers, HA Proxy	Elastic Load Balancing
Management and monitoring	Performance and user monitoring solutions	CloudTrail, CloudWatch, Kinesis, Simple Notification Service (SNS)
Messaging and workflow	Messaging and workflow software	Simple Notification Service (SNS), Simple Queue Service (SQS), Simple Workflow Service (SWF)
Network	MPLS, VPN	Direct Connect, Virtual Private Cloud (VPC)
Scaling	Hardware and software clustering, Apache ZooKeeper	Auto Scaling
Security	Firewalls, NACLs, routing tables, disk encryption, SSL, IDS, IPS	CloudHSM, Key Management Service (KMS), security groups
Storage	DAS, NAS, SAN, SSD	EC2 Instance storage (SSD), Elastic Block Store (EBS), Simple Cloud Storage Service (S3)



AWS Shared Responsibility Model





Source: <https://aws.amazon.com/compliance/shared-responsibility-model/>

WORKING WITH YOUR AWS DATA

DATA TYPE	WHAT IT CAN TELL YOU	SPLUNK SOURCETYPE	SAMPLE SECURITY DETECTIONS	KEY DATA ATTRIBUTES
Billing	Your configured billing reports, including historical bills and capacity planning information	aws:billing aws:billing:cur	<ul style="list-style-type: none"> - Determine how AWS resources are being consumed - Scenarios where an IAM key is compromised, and used to spin up many resources 	AvailabilityZone ItemDescription Operation PayerAccountId ProductName ReservedInstance ResourceId
CloudTrail	The AWS CloudTrail service provides a record of management and change events via the API calls	aws:cloudtrail	<ul style="list-style-type: none"> - Escalations of privileges - Uses of exposed credentials - Publicly accessible S3 buckets - Activity trail before and after an incident - Suspicious login activities - Suspicious provisioning activities - Spike in API activity 	eventTime awsRegion eventName eventSource sourceIPAddress userIdentity requestParameters
CloudWatch	Performance and billing metrics are available from the AWS CloudWatch service	aws:cloudwatch	<ul style="list-style-type: none"> - Determine how AWS resources are being consumed - Identify if charges exceed the normal usage 	account_id metric_name metric_dimensions period SampleCount Average Minimum Maximum
CloudWatch Logs	VPC logs available from the AWS CloudWatch Logs service capture IP traffic flow data for the network interfaces in your account	aws:cloudwatchlogs aws:cloudwatchlogs:vpctflow	<ul style="list-style-type: none"> - From where is the blocked traffic originating - Protocol used to send the data - Identify potential botnet activity - Web scanner looking for vulnerable software 	account-id interface-id srcaddr dstaddr srcport dstport packets protocol bytes start end action
Config	Configuration snapshots, historical configuration information and change notifications can show when changes were made—which can be valuable when troubleshooting	aws:config aws:config:notification	<ul style="list-style-type: none"> - Changes to resources configuration 	action resource
Config Rules	Config rules data give you information on status and compliance	aws:config:rule	<ul style="list-style-type: none"> - Is the change that just occurred to a resource compliant? 	ConfigRuleArn ConfigRuleId ConfigRuleName Description

WORKING WITH YOUR AWS DATA


DATA TYPE	WHAT IT CAN TELL YOU	SPLUNK SOURCETYPE	SAMPLE SECURITY DETECTIONS	KEY DATA ATTRIBUTES
Config Rules	Config rules data give you information on status and compliance	aws:config:rule	- Is the change that just occurred to a resource compliant?	ConfigRuleArn ConfigRuleId ConfigRuleName Description
GuardDuty	GuardDuty produces security findings to help you protect your AWS accounts and workloads	aws:cloudwatch:guardduty	- Potentially compromised accounts and instances - Reconnaissance by attackers - Unauthorized deployments	account time region detail
Inspector	Data from the Amazon Inspector service can help you improve the security and compliance of your AWS-hosted application	aws:inspector	- Vulnerabilities or deviations from best practices	arn assetAttributes assessmentTemplate attributes description recommendation severity
S3	Log data that is sent to S3 from AWS services, such as access logs for S3	aws:s3 aws:s3:accesslog	- Know which files are most accessed - Detect publicly accessible S3 buckets	Bucket Owner Bucket Time Remote IP Requester Operation Request URI HTTP Status User-Agent
S3 – CloudFront	CloudFront access logs provide insights into traffic and error metrics about your content delivery network (CDN) service	aws:cloudfront:accesslogs	- Traffic patterns - Charges	date time c-ip cs-method cs-uri-stem cs(User-Agent) cs-uri-query cs(Cookie) x-host-header x-forwarded-for
S3 – ELB (and ALB)	Access logs capture detailed information about requests sent to your load balancer(s)	aws:elb:accesslogs	- Unusually long or large requests - Track suspicious behavior from IPs and user agents - Excessive errors - Web scanner looking for vulnerable software	timestamp client:port elb_status_code request user_agent
SQS	Message queues	aws:sqs		attributes body

“How do we detect attacks in the cloud?”



Playbook Methodology

- ▶ WHAT ARE YOU TRYING TO **PROTECT**?
- ▶ WHAT ARE THE **THREATS**?
- ▶ HOW DO YOU **DETECT** THOSE THREATS?
- ▶ HOW DO YOU **RESPOND**?




Threat-Based Monitoring Plan

AWS Web Application Hosting reference architecture

REFERENCE ARCHITECTURE

WEB APPLICATION HOSTING

Highly available and scalable web hosting can be complex and expensive. Serverless services and micro services in traffic patterns result in the utilization of serverless hardware. Amazon Web Services provides the reliable, scalable, secure, and high-performance infrastructure required for web applications while enabling an elastic, horizontal and serverless architecture in multi-10 zones in the time as customer SaaS Solutions.



amazon
web services

INFRA COMPONENT

THREAT

AWS SERVICE

DETECTION LOGIC

INFRA COMPONENT	THREAT	AWS SERVICE	DETECTION LOGIC
RDS (7)	Access to sensitive data	RDS	<ul style="list-style-type: none"> Access to MySQL "SYSTEM"; /etc/mysql/* Access by admins to user tables (requires knowledge of schema?) Access to encrypted fields?
	New accounts created with superuser privileges	RDS	<ul style="list-style-type: none"> "GRANT ALL PRIVILEGES" "GRANT INSERT, UPDATE, DELETE, CREATE, DROP"
	Abnormally large response sets	RDS	<ul style="list-style-type: none"> "SELECT * FROM" Response size
	Connection pooling account deobfuscation	RDS	<ul style="list-style-type: none"> Correlate application activity with RDS audit information
	SQLi (baseline activity & anomaly detection)	RDS	<ul style="list-style-type: none"> xp_cmdshell execution Outlier query formats Queries that don't use an index
	Unusually high volumes of slow SQL queries	RDS	<ul style="list-style-type: none"> Std deviation from norm over x time
	Failed DB connection attempts	RDS	<ul style="list-style-type: none"> All failed login for users w/ admin privileges Threshold-based for non-admin
AUTHORITATIVE DNS (1)	Query Spikes (SYN Floods & queries cost \$\$)	Route53	<ul style="list-style-type: none"> Route53 query log usage spike Route53 health check failure
	Zone enumeration	Route53	<ul style="list-style-type: none"> Look for xfer queries
CDN (2)	Dynamic Content Floods	CloudFront	<ul style="list-style-type: none"> CDNs often protect static content, while dynamic content is forwarded to backend infra. Look for randomized/dynamic strings in the URI to identify CDN pass-through requests
	SSL Attacks	CloudFront	<ul style="list-style-type: none"> Pass-through floods SSL resource exhaustion (spikes in SSL requests)
	Cache attacks	VPC Flows Splunk Stream	<ul style="list-style-type: none"> Amplification: packet count/size disparity
	CDN subversion attack	access_combined	<ul style="list-style-type: none"> HTTP requests directly to IPs
DATA STORAGE (3)	Public Buckets	S3 Access Logs AWS Config / CloudTrail	<ul style="list-style-type: none"> Detect ACL PUT to everyone or similar group Detect new open S3 buckets Config Rule: S3 global read & write
	Deleted buckets	S3 Access Logs	<ul style="list-style-type: none"> DELETE_OBJECT
	Unencrypted data transmission to/from buckets	AWS Config	<ul style="list-style-type: none"> Config Rule: Encrypted buckets
WEB INFRA (6,5)	Instance created by an unusual user	CloudTrail	<ul style="list-style-type: none"> First time a user provisioned an instance
	Unusual amount of modifications to ACLs	CloudTrail	<ul style="list-style-type: none"> Sudden change in number of ACL modifications
	First time instance started in a new region	CloudTrail	<ul style="list-style-type: none"> Provisioning activity from unusual country
	Common Web Attacks	ELB, Apache	<ul style="list-style-type: none"> OWASP Top 10 web attacks
			<ul style="list-style-type: none"> SQLi Directory traversal XSS
<ul style="list-style-type: none"> Brute force 			

Web Infrastructure Example

INFRA COMPONENT	THREAT	AWS SERVICE	DETECTION LOGIC
WEB INFRA (6,5)	<i>Instance created by an unusual user</i>	CloudTrail	First time a user provisioned an instance
	<i>Unusual amount of modifications to ACLs</i>	CloudTrail	• Sudden change in number of ACL modifications
	<i>First time instance started in a new region</i>	CloudTrail	• Provisioning activity from unusual country
	<i>Common Web Attacks</i>	ELB, Apache	<ul style="list-style-type: none"> • OWASP Top 10 web attacks • SQLi • Directory traversal • XSS • Brute force




An In-depth Look



Two examples

CloudTrail



“A service that enables governance, compliance, operational auditing, and risk auditing of your AWS account.”

Source: <https://aws.amazon.com/cloudtrail/>

Working with CloudTrail Log Files

▶ Log file name format:

- <https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-log-file-examples.html>
- AccountID_CloudTrail_RegionName_YYYYMMDDTHHmZ_UniqueString.FileNameFormat

Example

111122223333_CloudTrail_us-east-2_20150801T0210Z_Mu0KsOhtH1ar15ZZ.json.gz

▶ Record contents:

- <https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-event-reference-record-contents.html>
- Determine the requested action as well as when and where the request was made



i	Time	Event
	2018-08-20T09:16:12.000AM	<pre>{ [-] awsRegion: us-east-1 eventID: e55a41a8-a425-448e-8a95-e0bb22696d1c eventName: GetCallerIdentity eventSource: sts.amazonaws.com eventTime: 2018-08-20T09:16:12Z eventType: AwsApiCall eventVersion: 1.05 recipientAccountId: 622676721278 requestID: 135e0154-9093-11e8-9dd7-db171eecd903 requestParameters: null responseElements: { [-] account: 622676721278 arn: arn:aws:iam::622676721278:user/web_admin userId: AIDAJNUCQVD57VVGYEFTQ } sourceIPAddress: 35.153.154.221 userAgent: Boto3/1.7.44 Python/2.7.12 Linux/4.4.0-1063-aws Botocore/1.10.44 userIdentity: { [-] accessKeyId: AKIAJOGCDXJ5NW5PXUPA accountId: 622676721278 arn: arn:aws:iam::622676721278:user/web_admin principalId: AIDAJNUCQVD57VVGYEFTQ type: IAMUser userName: web_admin } }</pre>

WHERE?

WHEN?

WHAT?

WHO?

Sample of Requested Actions

eventName	count
CreateAccessKey	1
CreateDefaultVpc	15
CreateUser	1
DeleteAccessKey	1
DescribeAccountAttributes	1
DescribeInstances	15
DescribeKeyPairs	15
GetCallerIdentity	16
GetSessionToken	1
GetUser	1
ListAccessKeys	2
ListBuckets	1
RunInstances	576

Does Anything Stand Out?

eventName	count
CreateAccessKey	1
CreateDefaultVpc	15
CreateUser	1
DeleteAccessKey	1
DescribeAccountAttributes	1
DescribeInstances	15
DescribeKeyPairs	15
GetCallerIdentity	16
GetSessionToken	1
GetUser	1
ListAccessKeys	2
ListBuckets	1
RunInstances	576

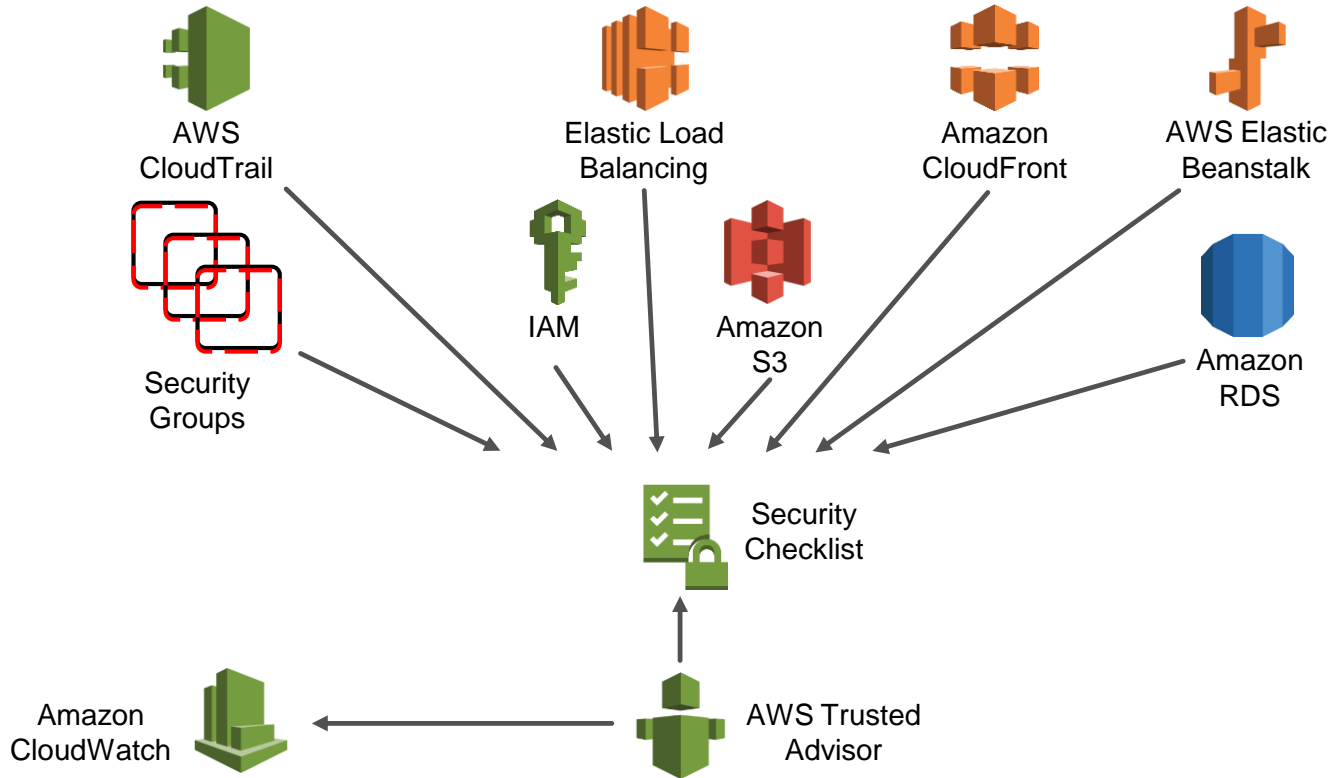


AWS Trusted Advisor

Best practice checks

How It Works

<https://aws.amazon.com/premiumsupport/technology/trusted-advisor/>



AWS Trusted Advisor Best Practice Checks

<https://aws.amazon.com/premiumsupport/technology/trusted-advisor/best-practice-checklist/>

AWS Trusted Advisor Best Practice Checks

Select a Support Plan Try Trusted Advisor Now

Trusted Advisor Best Practice Checks

AWS Trusted Advisor offers a rich set of best practice checks and recommendations across five categories: cost optimization; security; fault tolerance; performance; and service limits.

Cost Optimization

See how you can save money on AWS by eliminating unused and idle resources or making commitments to reserved capacity.

- Amazon EC2 Reserved Instances Optimization
- Low Utilization Amazon EC2 Instances
- Idle Load Balancers
- Underutilized Amazon EBS Volumes
- Unassociated Elastic IP Addresses
- Amazon RDS Idle DB Instances
- Amazon Route 53 Latency Resource Record Sets

Seven (7) core checks – Security

- ▶ S3 Bucket Permissions
- ▶ Security Groups - Specific Ports Unrestricted
- ▶ IAM Use
- ▶ MFA on Root Account
- ▶ EBS Public Snapshots
- ▶ RDS Public Snapshots

Service Limits

A large iceberg is shown floating in a blue ocean under a blue sky with scattered white clouds. The top of the iceberg is above the water line, while the vast majority of its mass is submerged below the surface. The text 'THE TIP OF THE ICEBERG' is overlaid on the image. 'THE' is in white, 'TIP' is in white with a blue shadow, 'OF THE' is in light blue, and 'ICEBERG' is in light blue.

THE
TIP

OF THE

ICEBERG

Key Takeaways

1. Become familiar with AWS data
2. Understand where to find security-relevant data in AWS
3. Understand how to use that data to detect malicious activity



SECURITY

CloudTrail - Digital Breadcrumbs for AWS



SECURITY

Go With the Flow - Network Telemetry (VPC Data) in AWS



This blog post talked quite a bit about Bezos in the experience of

Supplemental Reading



This blog post is part twenty-three of the "[Hunting with Splunk: The Basics](#)" series. That's right, it's a Brave New World of Cloud and we in Splunk Security are doubling down on it. We were thinking of releasing some ASMR podcasts on the subject, but apparently sleep teaching is actually prohibited in England (we assume something to do with GDPR.) Once again, [Matt Valites](#) brings us some nuggets from the world of AWS.

Thank You

