MRI for the Cloud Workloads

How Network Data Can Power Visibility, Detection, and Response Programs

Edward Wu
ExtraHop Networks

Disclaimer

The opinions expressed in this presentation are the presenter's own and do not reflect the view of my employer.

About Me

- Senior Principal Data Scientist, leading AI/ML and detection at Extrahop Networks
  - Also spearheading product's expansion to Cloud Workload Security
- Previously worked on automated binary analysis and software defenses at UC Berkeley and UW Seattle
- Fun fact: built the first working exploit of Zeus Bot a decade ago

Cloud Security Challenges

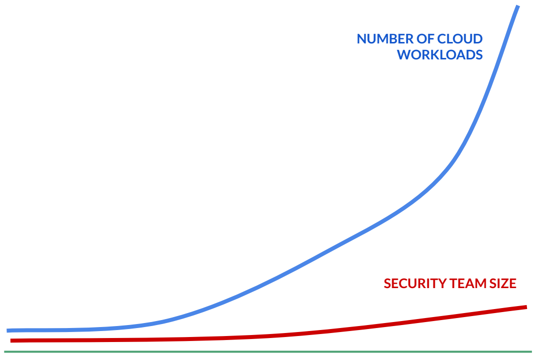

Security teams unable to keep pace with exponential growth in cloud workloads

+

Workload sprawl

Anatomy of Cloud Workload

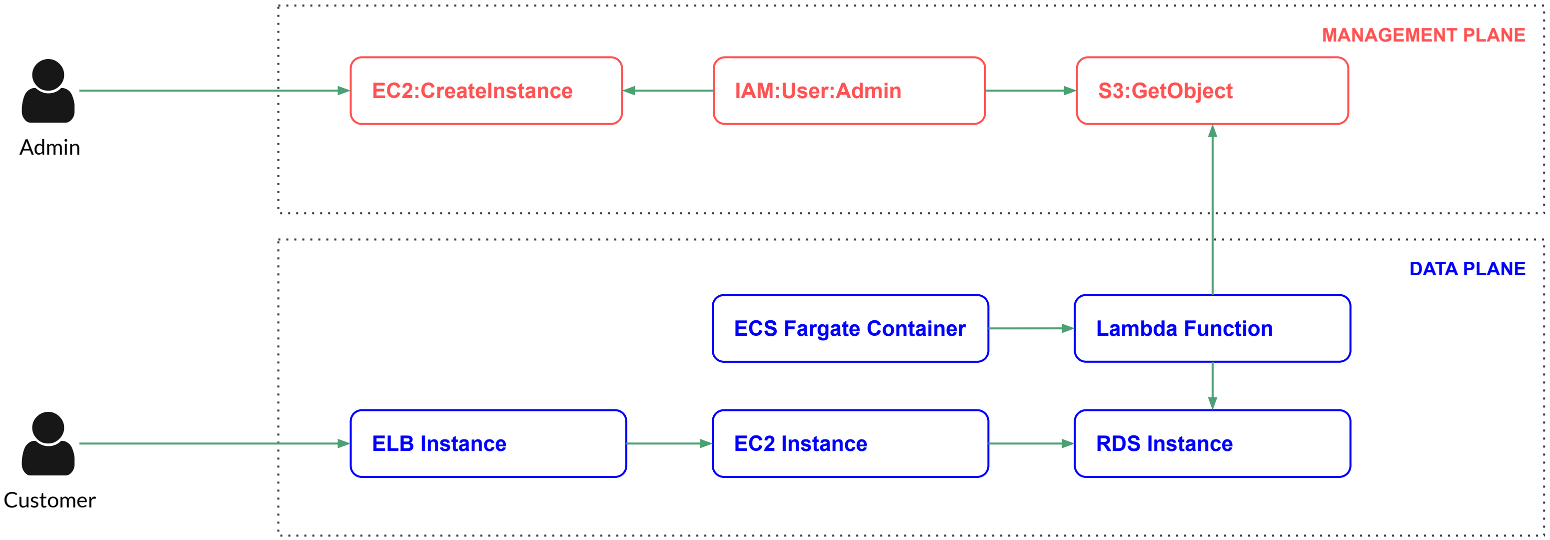

Two Behaviors Planes

Cloud workloads operate in 2 parallel behavior planes:

- Management plane: consists of cloud service provider (CSP) management APIs that enable organizations to create, modify, and manage compute, storage capacity, and infrastructure
- Data plane: where different workloads communicate on the network, similar to traditional on-prem data center workloads

Two Planes of Attack

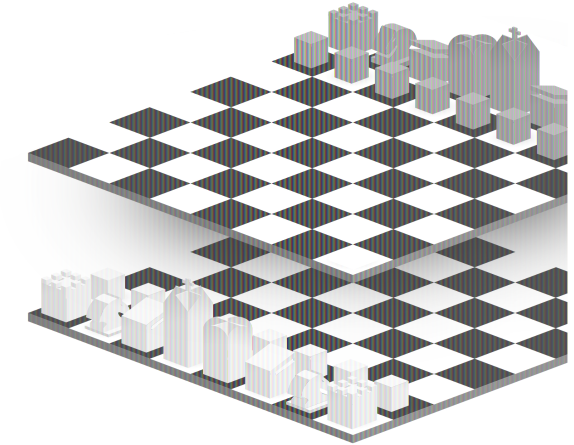

Given cloud workloads span 2 behavior planes, the attackers could also operate in these 2 planes:

- Management plane: leaked/compromised credentials, abuse of over privileged policies, CSP management software vulnerabilities
    - Good coverage from existing CSP and third-party tools
- Data plane: the same battleground for the traditional on-premises data center and corporate networks
    - Infrequently covered by existing security tools

MITRE ATT&CK Cloud Workload Matrix

Initial Access	Execution	Persistence	Privilege Escalation	Defense Evasion	Credential Access	Discovery	Lateral Movement	Collection	Exfiltration	Impact
Exploit Public-Facing Application	User Execution	Account Manipulation	Valid Accounts	Impair Defenses	Brute Force	Account Discovery	Use Alternate Authentication Material	Automated Collection	Transfer Data to Cloud Account	Data Destruction
Trusted Relationship		Create Account		Modify Cloud Compute Infrastructure	Forge Web Credentials	Cloud Infrastructure Discovery		Data from Cloud Storage Object		Data Encrypted for Impact
Valid Accounts		Implant Internal Image		Unused/Unsupported Cloud Regions	Multi-Factor Authentication Request Generation	Cloud Service Dashboard		Data from Information Repositories		Defacement
		Valid Accounts		Use Alternate Authentication Material	Network Sniffing	Cloud Service Discovery		Data Staged		Endpoint Denial of Service
				Valid Accounts	Unsecured Credentials	Cloud Storage Object Discovery				Network Denial of Service
Management Plane						Network Service Discovery				Resource Hijacking
Data Plane						Network Sniffing				
Both						Password Policy Discovery				
						Permission Groups Discovery				
						Software Discovery				
						System Information Discovery				
						System Location Discovery				
						System Network Connections Discovery				

Cross plane attacks

- In addition to moving on each plane, attackers could also weave between 2 planes similar to 3d chess
  - Frodo jumping between spiritual world and physical world via the Ring
- Example of Data plane -> management plane pivot: credential harvesting to gain access to additional credentials that provide expanded management plane privilege
- Examples of Management plane -> data plane pivot:
  - “airdrop" workloads of their control directly behind defenses in the data plane
  - Inject malicious code into existing cloud workloads from the management plane via existing tooling like AWS Systems Manager Agent or User data

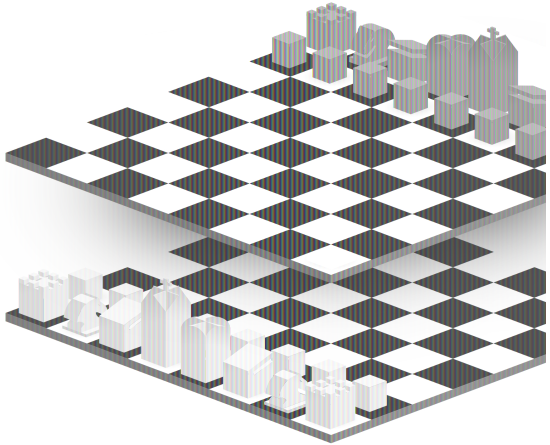

Network Data

Data extracted and derived from the actual network communications between entities on the network

- Empirical
  - Observed instead of self-reported
  - Can not be turned off or bypassed
- Breath of coverage
  - Can be enabled without any change or consent to the entities being monitored
  - Can be enabled on any entity that communicates on the network
- High Signal-To-Noise Ratio
  - Normalized and consistent across different applications, workloads, and OS

Great fit for cloud workload security due to its transparent deployment model and broad coverage, compared to agents and logs

- Passive/non-intrusive to devs
- Covers wide range of Cloud workloads from IaaS, PaaS, containerized, to serverless workloads

Use cases of network data

- Visibility
  - Behavior context
    - inspecting the behavior of a workload is often the best way to understand its role and purpose
  - Asset inventory and dependency mapping
    - “One can not defend something he/she can’t see”
- Posture management
  - Unexpected public facing assets
  - Network micro-segmentation

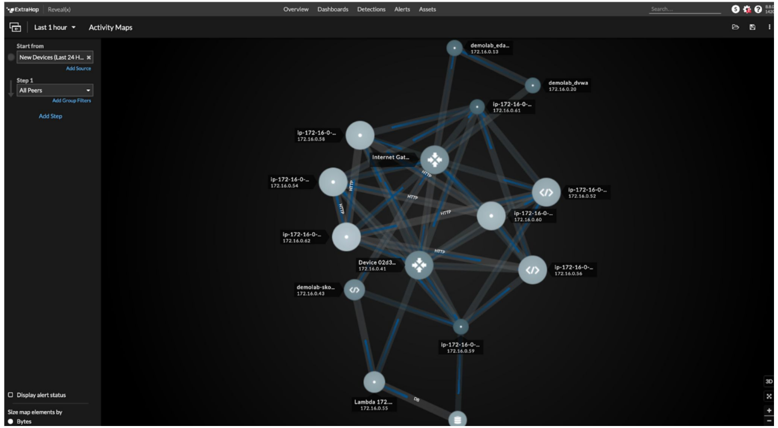

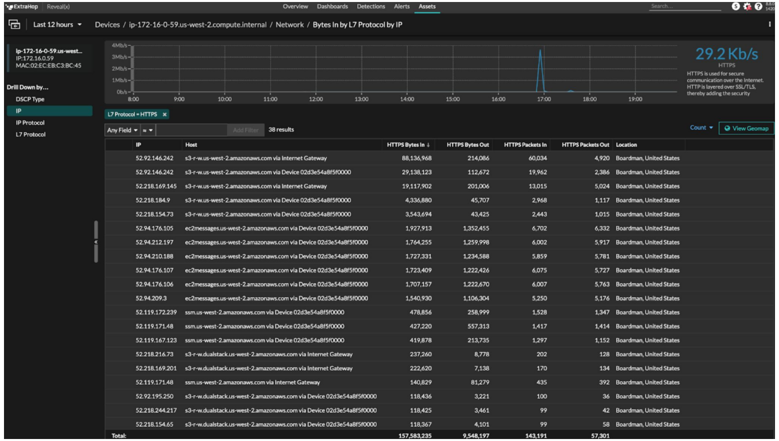

Use cases of network data

- Detections and Investigation
  - Known attack techniques
    - Brute force
    - C2
    - Data exfiltration
  - Unknown unknown attacks based on unusual network connections
- Forensics
  - Identify Root Cause and scope of impact
  - Demonstrate Proof

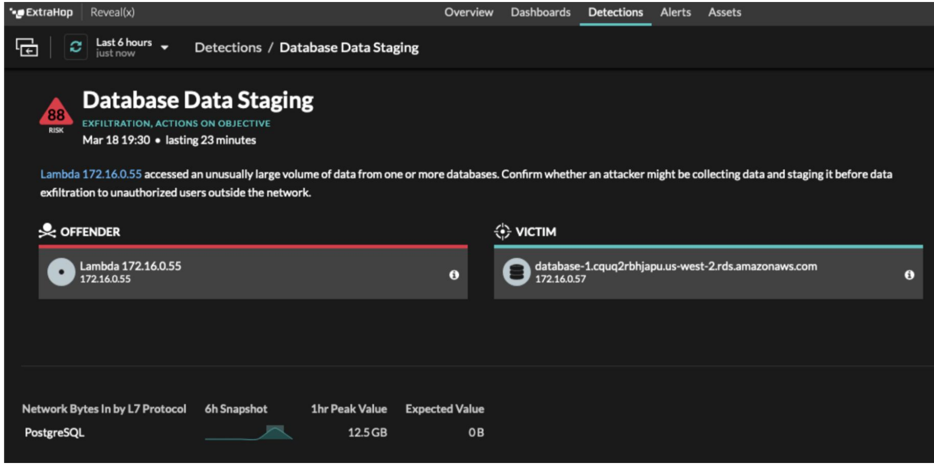

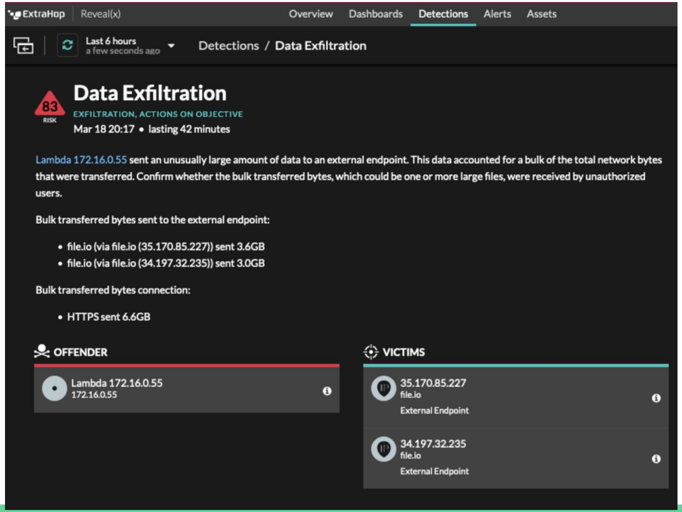

2 types of network data

- Flow logs
  - Aggregated metadata about network connections at L3
    - Source IP address/port number
    - Destination IP address/port number
    - IANA protocol (e.g., TCP, UDP)
    - # of bytes and packets
  - Similar to mobile phone call logs
- Full packets
  - Full payloads (L2-L7), could be processed to extract a variety of metadata, including:
    - SNI of HTTPS connections
    - URI of HTTP requests
    - SQL statements being issued
  - Similar to full recording of phone calls
  - A superset of information compared to flow logs
  - Require additional network sensor to transform into useable metadata

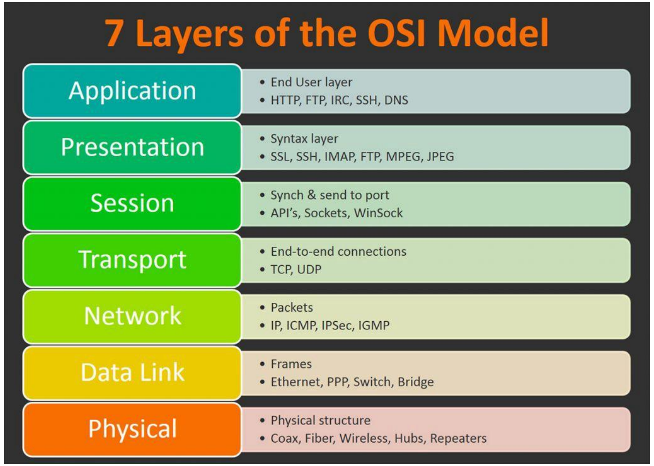

Example of Flog logs

SSH traffic (destination port 22, TCP protocol) to network interface eni-1235b8ca123456789 in account 123456789010 was allowed:

2 123456789010(account-id) eni-1235b8ca123456789(interface-id) 172.31.16.139(srcaddr) 172.31.16.21(dstaddr) 20641(srcport) 22(dstport) 6(protocol) 20(packets) 4249(bytes) 1418530010(start) 1418530070(end) ACCEPT OK

Example of Metadata available in Full packets


```
"http": {
    "hostname": "test.co.uk",
    "url":"\/test\/file.json",
    "http_user_agent": "<User-Agent>",
    "http_content_type": "application\/json",
    "http_refer": "http:\/\/www.test.com\/",
    "http_method": "GET",
    "protocol": "HTTP\/1.1",
    "status":"200",
    "length":310,
    "request_headers": [
        {
            "name": "User-Agent",
            "value": "Wget/1.13.4 (linux-gnu)"
        },
   .......
```

Example of Metadata available in Full packets


```
"tls": {
    "subject": "C=US, ST=California, L=Mountain View, O=Google Inc, CN=*.google.com",
    "issuerdn": "C=US, O=Google Inc, CN=Google Internet Authority G2",
    "serial": "0C:00:99:B7:D7:54:C9:F6:77:26:31:7E:BA:EA:7C:1C",
    "fingerprint": "8f:51:12:06:a0:cc:4e:cd:e8:a3:8b:38:f8:87:59:e5:af:95:ca:cd",
    "sni": "calendar.google.com",
    "version": "TLS 1.2",
    "notbefore": "2017-01-04T10:48:43",
    "notafter": "2017-03-29T10:18:00"
}
```

Flow logs vs full packets

Flow logs aggregate L3 network metadata over time but lose all the information in the content of the transactions

- Flow logs could see an outbound connection to server port 80, but full packets can tell exactly whether the connection was using HTTP or SSH
- L7 application layer metadata are must-have for many analysis:
  - Status codes
  - Errors
  - Usernames
  - URIs
  - Certificates

Flow logs vs full packets

Flow logs have many practical advantages:

- Easier to acquire than full packets, enabled on the network level instead of individual workload
- Cheaper and significantly lower volume than full packets
- Cover more types of workloads than full packets due to how networking is implemented in CSPs

	IaaS	PaaS	Containerized	FaaS	Network Infrastructure
Flow logs	Yes	Yes	Yes	Yes	Yes
Full packets	Yes	Sometimes	Sometimes	No	Sometimes

How to get started - Data Acquisition

- Flow logs:
  - Can be turned on at network level, immediately granting visibility to large chunks of workloads
  - AWS VPC flow log, Azure NSG flow log, GCP VPC flow log
- Full packets:
  - Generally needs to be individually enabled on each workload
    - Could be automated with additional tooling
  - AWS Traffic mirroring, Azure Virtual network tap (beta), GCP traffic mirroring

How to get started - Analytics

- Flow logs can be directly used for analytics
  - Could be enriched with auxiliary DNS logs to annotate IP addresses in flow logs with hostnames
- Full packets requires deployment of separate software sensor to extract relevant metadata and generate structured logs first
  - Security Onion, Arkime, Suricata, Zeek
- Structured metadata from full packets and flow logs are a good fit for a wide range of analytics platforms ranging from generic columnar data stores to SIEMs
  - For example, in AWS, one can directly query VPC flow logs using SQL: https://docs.aws.amazon.com/athena/latest/ug/vpc-flow-logs.html
- Conversion to time series data is another way to explore behaviors over time
  - Number of inbound connections on a specific server
  - Number of HTTP 404s for a specific API endpoint

How to get started - Multi-account Deployment

1. One sensor per account
   a. Requires some footprint in every monitored account
   b. Sensor overhead/cost could be nontrivial if there are a lot of small accounts
2. Centrally-deployed small pool of sensors to process network data from all accounts
   a. Might need to pay for cross account traffic depending on the CSP
   b. Overlapping network segments could confuse sensors

Example Multi-account Deployment Architecture

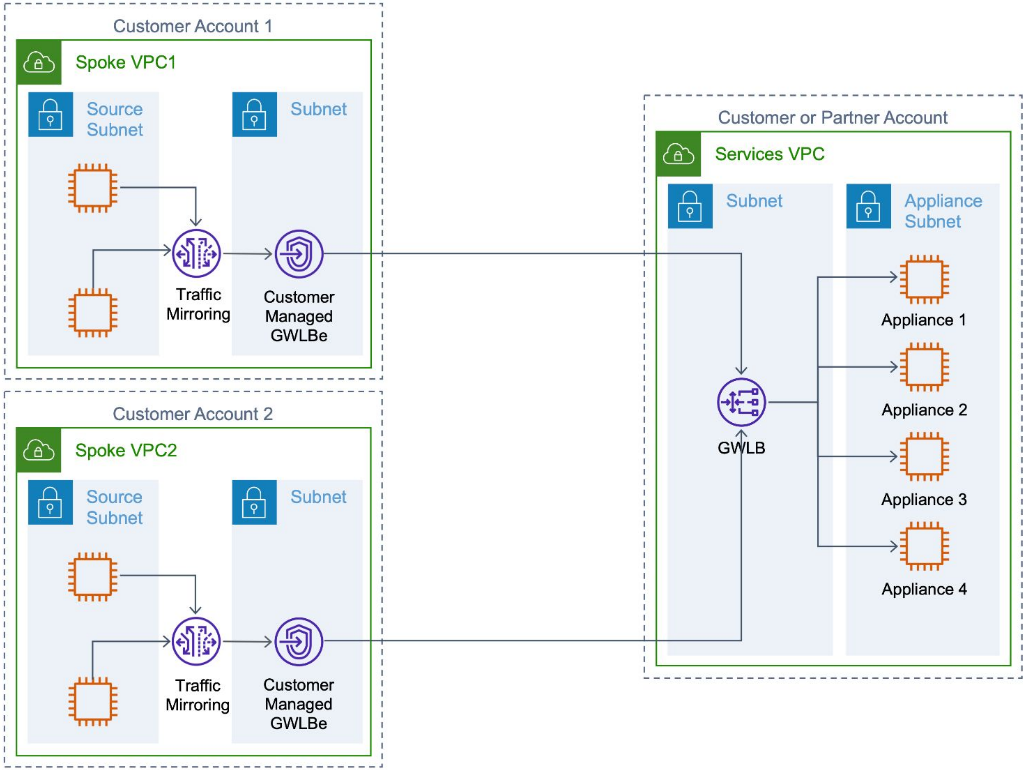

Conclusion

- Data plane visibility is often overlooked
  - Provides behavior context of different workloads
  - As CSP management plane security levels up and stops being the weakest link, attackers are expected incorporate more data plane attack techniques that are invisible to CSP management plane logs
- Network data is the single biggest ONE STEP jump to situational awareness from near-total unawareness in the data plane
  - Passive deployment model and broad coverage fit really well with fast moving cloud application development teams
- Flow logs have broad coverage, are easier to get started, but offer lower fidelity data
- Full packets are more expensive to acquire and utilize, but offer the ultimate data fidelity, which can power more sophisticated detection and analytics