Auditing, Logging, and Observability Cloud Security Part II

CS 40 | February 7, 2024

More (harder) Cloud Security Best Practices

Open Policy Agent

- Enforce security policies on cloud deployments when using IaC
  - e.g., “ensure that S3 buckets are not exposed to the public”

- Can be embedded in deployment pipelines or IaC state management systems (e.g., Terraform Cloud)

Example: Open Policy Agent

```
fail contains msg if {
   buckets := [bucket | bucket := input.Resources[_]; bucket.Type == "AWS::S3::Bucket"]
   configs := buckets[_].Properties.PublicAccessBlockConfiguration
   not checkBucketRestricted(configs)

   msg := sprintf("S3 buckets should block public access”, [])
}
```

From the Assignment 2 autograder

Cloud Security Products

- **Cloud Security Posture Management** (CSPM): software that scans cloud resources and IaC to spot misconfigurations and insecurities

- **Cloud Native Application Protection Platform** (CNAPP): CSPM plus more active application monitoring (e.g. with agents on compute resources)

Observability

Observability *is the ability to ask arbitrary questions about a system without having to know ahead of time what to ask.*

Why is observability important?

Print statements (plus more) for deployed applications

- **Debugging**: If “something” in your deployed application doesn’t work
  - Where in the chain did something go wrong?
  - Isolate the behavior of the failed component + potential logical assumptions surrounding it
- **Performance**: If “something” in your deployed application feels slow
  - Profile the slowest components of the application, to know where optimizations are needed
- **Security**: If an attacker was able to exploit “something” in your application
  - Where did the exploit originate?
  - How much damage was the attacker able to do?

Logging

Logging from an Application Perspective

Application event logs

- For each web request, trace:
  - Handoffs between services
  - Control flow abnormality
  - Errors, exceptions, warnings
- Log levels:
  - **DEBUG**
  - **INFO**
  - **WARNING**
  - **ERROR**
- Generated by application itself

Access and security logs

- For each web request, trace:
  - IP address of client
  - Path requested
  - Response status
- Usually generated by underlying web server or load balancer

Example: Application Event Logs

Timestamp (UTC-08:00)	Message	Container
February 07, 2024 at 10:57 (UTC-8:00)	INFO: 10.0.5.56:26024 - "GET /api/v1/users/profile/37078c93-e20f-46bf-b046-4ab1bd6a288b HTTP/1.1" 200 OK	yoctogram-app-container
February 07, 2024 at 10:57 (UTC-8:00)	INFO: 10.0.5.56:26008 - "GET /api/v1/users/profile/a595dfb2-3402-4369-9009-5aab259e83af HTTP/1.1" 200 OK	yoctogram-app-container
February 07, 2024 at 10:57 (UTC-8:00)	INFO: 10.0.4.146:31826 - "GET /api/v1/users/profile/ba88f5de-ee05-4028-a224-169fb03c804b HTTP/1.1" 200 OK	yoctogram-app-container
February 07, 2024 at 10:57 (UTC-8:00)	INFO: 10.0.5.56:26008 - "GET /api/v1/users/profile/d4f3e0b7-edf0-4591-a0dd-7a766c6729dc HTTP/1.1" 200 OK	yoctogram-app-container
February 07, 2024 at 10:57 (UTC-8:00)	INFO: 10.0.4.146:11860 - "GET /api/v1/health HTTP/1.1" 200 OK	yoctogram-app-container
February 07, 2024 at 10:57 (UTC-8:00)	INFO: 10.0.5.56:25992 - "GET /api/v1/health HTTP/1.1" 200 OK	yoctogram-app-container
February 07, 2024 at 10:57 (UTC-8:00)	INFO: 10.0.4.146:31826 - "GET /api/v1/feed/latest?before=2024-02-08T18:56:16.764Z&after=1970-01-01T00:00:00.000Z HTTP/1.1" 200 OK	yoctogram-app-container
February 07, 2024 at 10:57 (UTC-8:00)	INFO: 127.0.0.1:43920 - "GET /api/v1/health/ HTTP/1.1" 307 Temporary Redirect	yoctogram-app-container
February 07, 2024 at 10:57 (UTC-8:00)	INFO: 10.0.4.146:59258 - "GET /api/v1/health HTTP/1.1" 200 OK	yoctogram-app-container
February 07, 2024 at 10:57 (UTC-8:00)	INFO: 10.0.5.56:26942 - "GET /api/v1/health HTTP/1.1" 200 OK	yoctogram-app-container
February 07, 2024 at 10:57 (UTC-8:00)	INFO: 127.0.0.1:51926 - "GET /api/v1/health/ HTTP/1.1" 307 Temporary Redirect	yoctogram-app-container
February 07, 2024 at 10:57 (UTC-8:00)	INFO: 10.0.5.56:45390 - "GET /api/v1/feed/latest?before=2024-02-08T18:56:14.725Z&after=1970-01-01T00:00:00.000Z HTTP/1.1" 200 OK	yoctogram-app-container
February 07, 2024 at 10:56 (UTC-8:00)	INFO: 10.0.5.56:45378 - "GET /api/v1/health HTTP/1.1" 200 OK	yoctogram-app-container
February 07, 2024 at 10:56 (UTC-8:00)	INFO: 10.0.4.146:31028 - "GET /api/v1/health HTTP/1.1" 200 OK	yoctogram-app-container
February 07, 2024 at 10:55 (UTC-8:00)	INFO: 127.0.0.1:33890 - "GET /api/v1/health/ HTTP/1.1" 307 Temporary Redirect	yoctogram-app-container

Example: Access and Security Logs

```
221.178.143.70 - - [07/Feb/2024:18:01:09 +0000] "GET / HTTP/1.1" 301 169 "-" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10_11) AppleWebKit/601.1.27 (KHTML, like Gecko) Chrome/47.0.2526.106 Safari/601.1.27" "-"
221.178.143.70 - - [07/Feb/2024:18:01:10 +0000] "GET / HTTP/1.1" 200 2370 "-" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10_11) AppleWebKit/601.1.27 (KHTML, like Gecko) Chrome/47.0.2526.106 Safari/601.1.27" "-"
45.33.87.154 - - [07/Feb/2024:18:01:13 +0000] "GET / HTTP/1.1" 301 169 "-" "Mozilla/5.0 zgrab/0.x" "-"
172.104.11.4 - - [07/Feb/2024:18:01:28 +0000] "GET / HTTP/1.1" 301 169 "-" "Mozilla/5.0 (Macintosh; Intel Mac OS X 13_1) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/108.0.0.0 Safari/537.36" "-"
221.178.143.70 - - [07/Feb/2024:18:01:30 +0000] "GET /favicon.ico HTTP/1.1" 301 169 "-" "Dalvik/2.1.0 (Linux; U; Android 9.0; ZTE BA520 Build/MRA58K)" "-"
221.178.143.70 - - [07/Feb/2024:18:01:31 +0000] "GET /favicon.ico HTTP/1.1" 404 150 "-" "Dalvik/2.1.0 (Linux; U; Android 9.0; ZTE BA520 Build/MRA58K)" "-"
36.150.60.24 - - [07/Feb/2024:18:01:44 +0000] "GET / HTTP/1.1" 301 169 "-" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10_11) AppleWebKit/601.1.27 (KHTML, like Gecko) Chrome/47.0.2526.106 Safari/601.1.27" "-"
36.150.60.24 - - [07/Feb/2024:18:01:45 +0000] "GET / HTTP/1.1" 200 2370 "-" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10_11) AppleWebKit/601.1.27 (KHTML, like Gecko) Chrome/47.0.2526.106 Safari/601.1.27" "-"
36.150.60.24 - - [07/Feb/2024:18:01:45 +0000] "GET / HTTP/1.1" 200 2370 "-" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10_11) AppleWebKit/601.1.27 (KHTML, like Gecko) Chrome/47.0.2526.106 Safari/601.1.27" "-"
36.150.60.24 - - [07/Feb/2024:18:01:48 +0000] "GET /favicon.ico HTTP/1.1" 301 169 "-" "Dalvik/2.1.0 (Linux; U; Android 9.0; ZTE BA520 Build/MRA58K)" "-"
36.150.60.24 - - [07/Feb/2024:18:01:49 +0000] "GET /favicon.ico HTTP/1.1" 404 150 "-" "Dalvik/2.1.0 (Linux; U; Android 9.0; ZTE BA520 Build/MRA58K)" "-"
221.178.143.70 - - [07/Feb/2024:18:02:09 +0000] "GET / HTTP/1.1" 301 169 "-" "Dalvik/2.1.0 (Linux; U; Android 9.0; ZTE BA520 Build/MRA58K)" "-"
221.178.143.70 - - [07/Feb/2024:18:02:10 +0000] "GET / HTTP/1.1" 200 2370 "-" "Dalvik/2.1.0 (Linux; U; Android 9.0; ZTE BA520 Build/MRA58K)" "-"
36.150.60.24 - - [07/Feb/2024:18:02:13 +0000] "GET / HTTP/1.1" 200 2370 "-" "Dalvik/2.1.0 (Linux; U; Android 9.0; ZTE BA520 Build/MRA58K)" "-"
221.178.143.70 - - [07/Feb/2024:18:02:22 +0000] "GET / HTTP/1.1" 200 2370 "-" "Mozilla/5.0 (Linux; U; Android 6.0.1; zh-CN; Redmi Note 3 Build/MMB29M) AppleWebKit/537.36 (KHTML, like Gecko) Version/4.0 Chrome/57.0.2987.
108 UCBrowser/11.8.9.969 Mobile Safari/537.36" "-"
36.150.60.24 - - [07/Feb/2024:18:02:28 +0000] "GET / HTTP/1.1" 301 169 "-" "Dalvik/2.1.0 (Linux; U; Android 9.0; ZTE BA520 Build/MRA58K)" "-"
36.150.60.24 - - [07/Feb/2024:18:02:29 +0000] "GET / HTTP/1.1" 200 2370 "-" "Dalvik/2.1.0 (Linux; U; Android 9.0; ZTE BA520 Build/MRA58K)" "-"
167.94.138.50 - - [07/Feb/2024:18:02:30 +0000] "GET / HTTP/1.1" 200 2370 "-" "-" "-"
167.94.138.50 - - [07/Feb/2024:18:02:33 +0000] "GET / HTTP/1.1" 200 2370 "-" "Mozilla/5.0 (compatible; CensysInspect/1.1; +https://about.censys.io/)" "-"
167.94.138.50 - - [07/Feb/2024:18:02:33 +0000] "GET /favicon.ico HTTP/1.1" 404 150 "-" "Mozilla/5.0 (compatible; CensysInspect/1.1; +https://about.censys.io/)" "-"
221.178.143.70 - - [07/Feb/2024:18:03:01 +0000] "GET / HTTP/1.1" 200 2370 "-" "Dalvik/2.1.0 (Linux; U; Android 9.0; ZTE BA520 Build/MRA58K)" "-"
36.150.60.24 - - [07/Feb/2024:18:03:37 +0000] "GET / HTTP/1.1" 200 2370 "-" "Mozilla/5.0 (Linux; U; Android 8.1.0; zh-CN; CLT-AL00 Build/HUAWEICLT-AL00) AppleWebKit/537.36 (KHTML, like Gecko) Version/4.0 Chrome/57.0.298
7.108 UCBrowser/12.1.3.993 Mobile Safari/537.36" "-"
36.150.60.24 - - [07/Feb/2024:18:03:42 +0000] "GET /favicon.ico HTTP/1.1" 404 150 "-" "Dalvik/2.1.0 (Linux; U; Android 9.0; ZTE BA520 Build/MRA58K)" "-"
36.150.60.24 - - [07/Feb/2024:18:03:58 +0000] "GET / HTTP/1.1" 301 169 "-" "Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.1) Opera 7.10 [fr]" "-"
221.178.143.70 - - [07/Feb/2024:18:05:17 +0000] "GET / HTTP/1.1" 301 169 "-" "Mozilla/5.0 (iPhone; CPU iPhone OS 11_2_1 like Mac OS X) AppleWebKit/604.4.7 (KHTML, like Gecko) Mobile/15C153 MicroMessenger/6.7.1 NetType/4
G Language/zh_CN" "-"
221.178.143.70 - - [07/Feb/2024:18:05:37 +0000] "GET /favicon.ico HTTP/1.1" 404 150 "-" "Dalvik/2.1.0 (Linux; U; Android 9.0; ZTE BA520 Build/MRA58K)" "-"
221.178.143.70 - - [07/Feb/2024:18:07:10 +0000] "GET / HTTP/1.1" 200 2370 "-" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10_11) AppleWebKit/601.1.27 (KHTML, like Gecko) Chrome/47.0.2526.106 Safari/601.1.27" "-"
143.110.222.166 - - [07/Feb/2024:18:19:14 +0000] "GET / HTTP/1.1" 301 169 "-" "Mozilla/5.0 (iPhone; CPU iPhone OS 16_1 like Mac OS X) AppleWebKit/605.1.15 (KHTML, like Gecko) Version/16.1 Mobile/15E148 Safari/604.1" "-"
95.214.235.169 - - [07/Feb/2024:18:33:51 +0000] "GET /.env HTTP/1.1" 301 169 "-" "Mozilla/5.0 (X11; Linux x86_64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/81.0.4044.129 Safari/537.36" "-"
95.214.235.169 - - [07/Feb/2024:18:33:52 +0000] "POST / HTTP/1.1" 301 169 "-" "Mozilla/5.0 (X11; Linux x86_64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/81.0.4044.129 Safari/537.36" "-"
```

Application Log Management

ELK (Elastic) stack, open-source: Elasticsearch, Logstash, Kibana*

1. Logstash ingests incoming application logs
2. Elasticsearch allows easy searching and analytics of logs
3. Kibana helps create visualizations from logs

Elastic Licensing Drama

- In April 2021, Elastic (ELK stack parent company/developer) relicensed Elasticsearch and Kibana from *Apache License 2.0* (open-source) to *Server-Side Public License* (source-available)

- Why this is bad: SSPL forces anyone (e.g. cloud providers) offering ELK as a service to open-source *all* supporting code – which is infeasible

- In response, AWS forked Elasticsearch and Kibana to create OpenSearch, which is still Apache License 2.0

Service Logging

- Application logging isn’t always enough

- Sometimes, need visibility into underlying infrastructure to debug
  - “Did my web request make it through the load balancer to my container?”

- AWS CloudWatch unifies application and service logs into a single place

AWS CloudWatch

- Log aggregation service for AWS resources
- Each resource forwards logs to a *log group*
  - Both application and service logs
- Logs are sharded into *log streams*
  - Representing *log events* from same logical source – e.g. individual containers

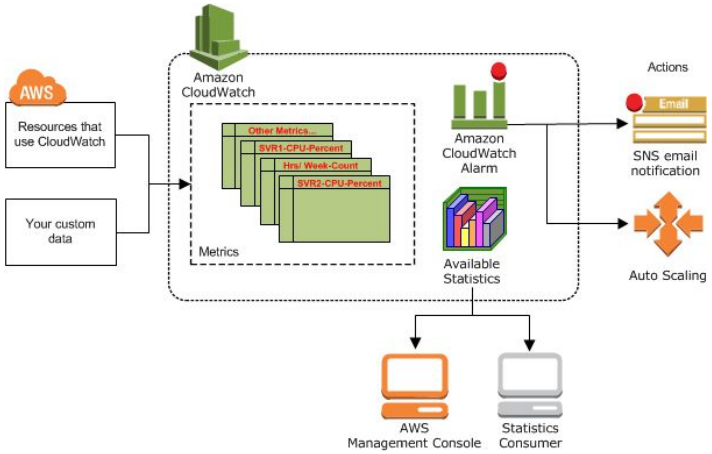

Pros and Cons of CloudWatch

Pros:

- Unify AWS application and service logging in the same place
- Integrate with other AWS services for alarms and visualizations

Cons:

- UI makes it difficult to trace individual events and find issues
- **Pricing**

CloudWatch Pricing

- Ingestion: $0.50/GB
  - Footgun: this gets charged at raw data size, even when the data is transmitted compressed!
- Retention: $0.03/GB
- Querying: $0.005/GB scanned

This gets expensive when dealing with many resources all logging to CloudWatch.

Security Information and Event Management

- Log management plus network information collection with a security focus
- *Anomaly detection* to find and alert to potential security events like intrusions

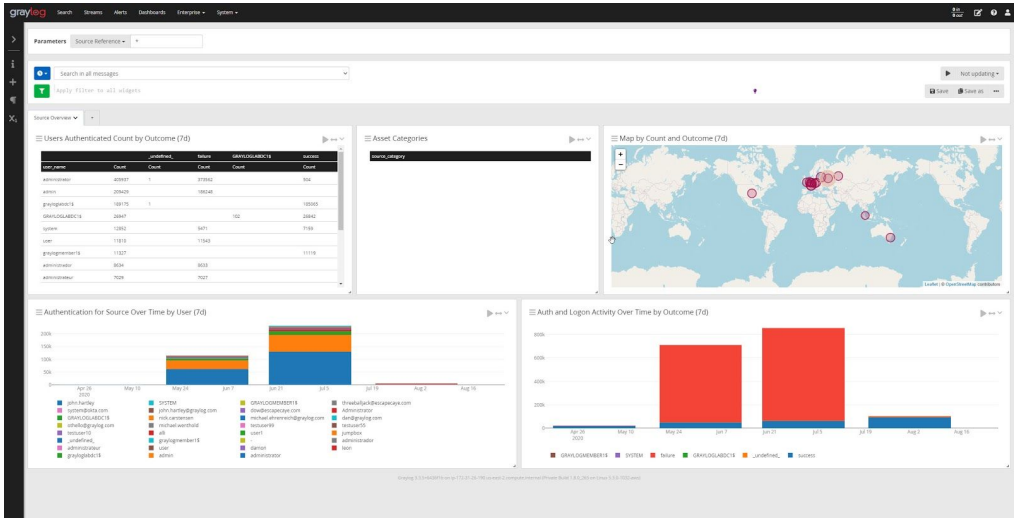

Using Logs in Practice

- Goal: Isolate the source of the problem by understanding where it is *not*

- Possible methods:
  - Filter logs to only those of the affected users
  - Identify the component causing the issue; use logs to discover which parts of the pipeline are working properly
  - Use your intuition to identify why the problem is occurring

Logs usually don't tell you what's going wrong directly – but they yield important context.

Using Logs in Practice

***Error resolution scenario**: some users are unable to access the website*

Using Logs in Practice

***Security incident scenario**: you find a big bill and unknown ECS clusters created*

Metrics & Monitoring

Motivation

- Proactively **and** reactively observe the state of a deployed system
  - To know what changes may need to be made for continued reliability
  - To anticipate future demand and scaling

- Goal: decrease **mean time to recovery** – the time it takes to return to normal operation following an incident
  - Alert to start the incident response process as soon as an issue is detected

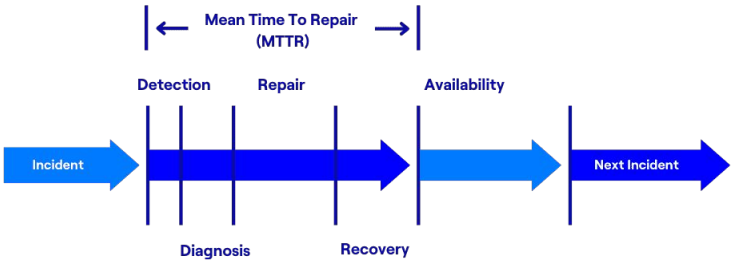

What should be monitored?

- **Latency**: how long does it take to service a request?
- **Traffic**: how much demand is being placed on the system?
- **Errors**: what requests fail and why, and context surrounding failures
- **Saturation**: how much demand are compute and storage resources under?
    - e.g. CPU & memory usage, I/O saturation

Google SRE – The Four Golden Signals, 2017

Metric Granularity

- Metric collection windows are contextual
  - CPU load should be observed at ~seconds frequencies: utilization spikes don’t last long
  - But probing for storage saturation or web server errors can be less frequent

- Overcollecting metrics can be costly!

Tracing

- *Motivation*: like a stack trace for distributed processes, with performance profiling

- This gives you more details and context around both errors and latency events

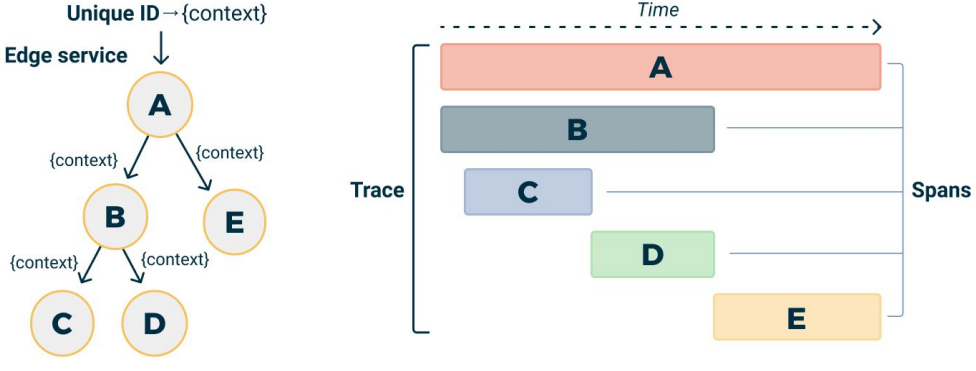

Prometheus & Grafana

Open-source metrics collection and management.

Comprehensive Observability Platforms

Integrate logging and metrics into one platform – mostly a commercial space

Why is observability so expensive?

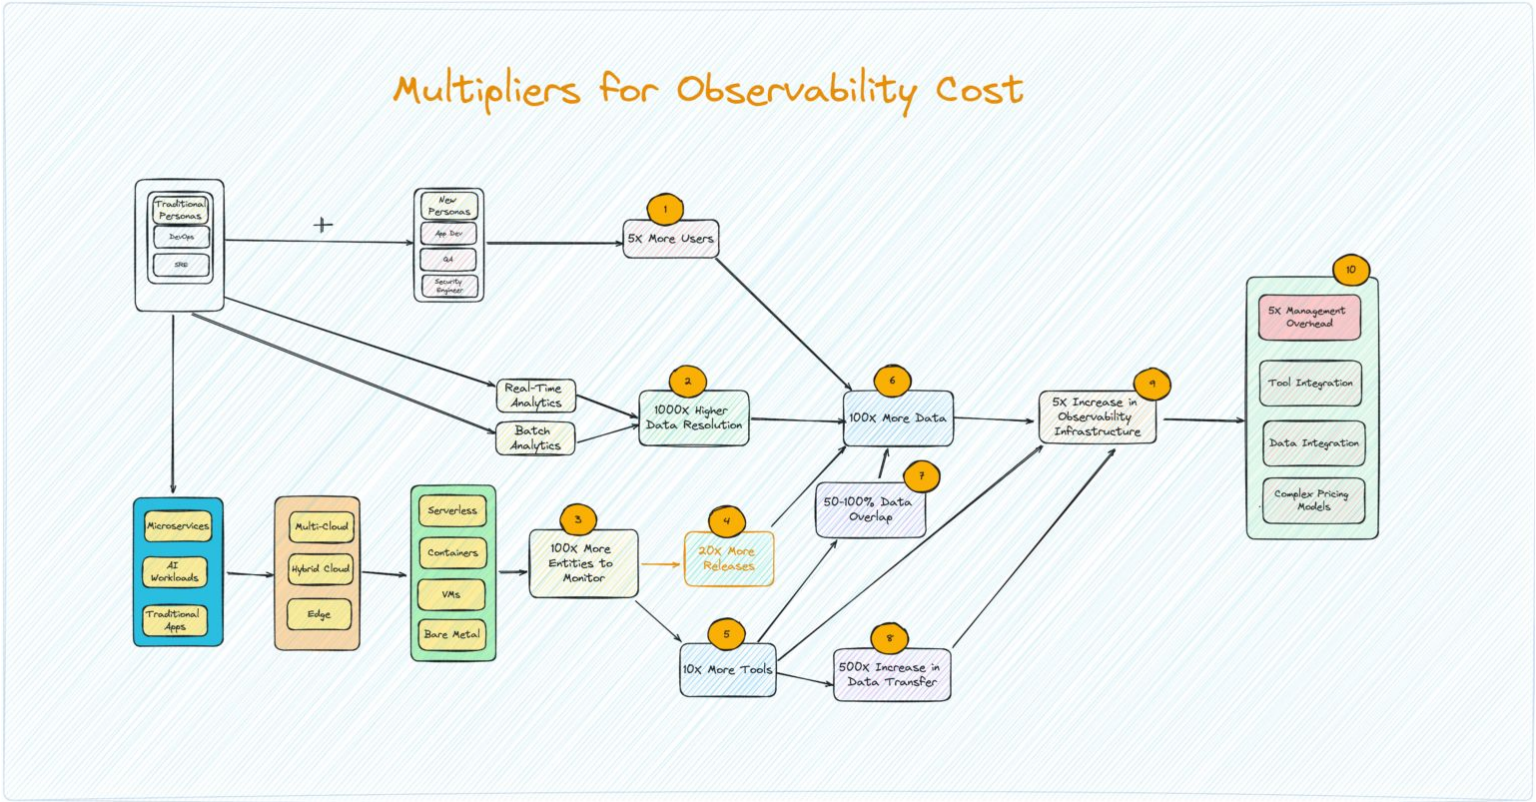

Next Lecture: Serverless Compute (2/12)