# Application Observability

[Real life Schrodinger's cat]







### We are One-stop-shop for [DevOps, Cloud & Kubernetes]



25 Top cloud experts **100+** Kubernetes clusters AWS PARTNER Advanced Tier Services

Solution Provider
Amazon EKS Delivery
Amazon RDS Delivery
Well-Architected Partner Program

- Amazon OpenSearch Service Delivery
- Migration and Modernization Services Competency



### **Our mission**

[Transforming businesses to become cloud native and future ready]



### About me



#### Adam Hamšík

#### CEO & Co-Founder

#### Personal Profile

One of 3 founders of Labyrinth Labs.

#### Profession

I really enjoy working on Open Source projects & being able to help.

#### Work Experience

I've been a DevOps engineer for several years, enjoying some kernel hacking with NetBSD in my free time.

#### Hobbies

Family, hiking, lacrosse + basketball (sports)



### WTF is Schrodinger's cat?



LABYRINTH



### T1;DR



LABYRINTH LABS

- Current state Infrastructure Architecture
  - Distributed systems vs Dynamic systems
  - Horizontal vs Vertical Scaling
  - Availability vs Reliability
  - Microservices vs Monoliths vs Event Driven Architecture
  - Serverless
  - Chaos Engineering
- What does observability mean?
  - Why is it important?
- Observability Stack
  - Description and short intro
- How do we do it then ?



#### LABYRINTH LABS

### **Quick Wins**

• Use libraries for **logging**, **metrics and traces** 

- No deploys to production on Fridays and holidays
- Avoid shiny new object/framework syndrome
- printf != Logging / Debugger

# Part I [Modern Infrastructure Architecture ]

### Examples how we run apps these days.



### **Distributed systems**



LABYRINTH LABS

#### **Distributed System**

A system where components are spread across multiple physical or virtual machines and communicate over a network.

- Multiple machines work as a single logical entity
- Each **machine** (node) may handle a specific task.
- Helps with **scalability and fault tolerance.**

### **Dynamic systems**



LABYRINTH LABS

#### Dynamic System

A system that changes its structure or behavior automatically in response to conditions (like load or failure).

- Automatically scales up/down based on demand.
- Can **self-heal** if a part fails
- **Adapts** in **real time** = cost-effective + resilient.

### **Horizontal vs Vertical Scaling**



- Using bigger machines and bigger machines (Hulk)
- Less reliable
  - One machine fails and you are done.

#### Horizontal

- Adding more machines (Naturo)
- More complicated architecture
- Can survive failure of multiple machines





LABYRINTH

I ABS

ക

### Availability vs Reliability



- The likelihood of a service failing
- To increase this you have to remove single points of failure.

#### Availability

- Availability = Reliability x Maintenance
- The percentage of time a service is operational

LABYRINTH

LABS

ക

# Monolith as software architecture pattern



LABYRINTH LABS

A **monolith** is a **single**, **tightly** coupled application where all features are **packaged** and **deployed** as one unit.

#### Pros

- Easier to start
- No network latency between components
- Easy for developers to contribute to.

#### Cons

- Can be **harder** to scale later
- Longer deployment cycles
- Tight coupling = harder to test or isolate bugs

### Microservices



LABYRINTH LABS

**Microservices** are an architectural pattern where an application is **split** into **small**, **independent services**, each responsible for a specific business capability.

#### Pros

- Independent deployments
- Better Scalability per service
- Fault isolation.
- Required if you have multiple product teams.

#### Cons

- Increased complexity
- Requires DevOps & automation maturity
- Harder debugging
- API Versioning, Compatibility

### **Microservices vs Monolith**



LABYRINTH LABS



https://medium.com/startlovingyourself/microservices-vs-monolithic-architecture-c8df91f16bb4



https://devcamp.com/site\_blogs/monolith-vs-microservice-rails-applications





LABYRINTH LABS

### SO THAT EVERY OUTAGE COULD BE MORE LIKE A MURDER MYSTERY

### **Event Driven Architecture**



LABYRINTH LABS

**Event-Driven Architecture** (EDA) is a software pattern where components communicate by producing and reacting to **events** (messages), rather than **calling each other directly**.

#### Pros

- Highly decoupled
- Scalable
- Async by default
- Required if you have multiple product teams.

#### Cons

- Hard debugging
- API Versioning, Compatibility
- Event ordering can be tricky
- Visibility is tough without good observability tools

### **Microservices vs Monolith**



LABYRINTH LABS



https://hazelcast.com/foundations/event-driven-architecture/event-driven-architecture/

### **Serverless Computing**



LABYRINTH LABS

**Serverless** is a cloud-native architecture where developers write and deploy code without managing servers. Cloud providers handle all the infrastructure, scaling, and runtime management.

#### Pros

- No server management
- Automatic scaling based on usage
- Great for event-driven use cases and microservices

#### Cons

- Limited execution time
- Not ideal for long-running or stateful processes
- Harder to debug locally



"These servers run our serverless computing service. We call it Schrödinger's Server."

https://deploy.equinix.com/blog/schrdingers-server/

### **Chaos Engineering**



LABYRINTH

Practice of **intentionally** breaking your system in **controlled ways** to understand how it behaves under failure.

"If we know things will break, let's test them before they break for real."

E.g. Inject failures (e.g., kill services, delay responses, drop network)



# Part II [Observability Theory]

**Some Theory** 





### **Observability Definition**

**Observability** is a measure of how well internal **states of a system** can be inferred from knowledge of its **external outputs**.

In **control theory**, the **observability** and **controllability** of a linear system are mathematical duals.

The concept of observability was introduced by the Hungarian-American engineer Rudolf E. Kálmán for linear dynamic systems. [1][2]

### **Observability Software Definition**



**Observability** is the ability to **collect** data about programs' **execution**, modules' internal states, and the communication among components.

To improve observability, software engineers use a wide range of **logging** and **tracing techniques** to **gather telemetry** information, and tools to analyze and use it.

**Observability** is **foundational** to site reliability engineering, as it is the first step in triaging a service outage.



### **Observability Definition**

Concept	Control Theory	Software Observability
System State	Internal physical variables (e.g., speed, position)	Application behavior and health (e.g., request rate, error rate)
Outputs	Measurable outputs (e.g., RPM, temperature)	Logs, metrics, traces from app and infra
Feedback	Adjust inputs to control system	Scale up, restart, alert, throttle



### **Observability Definition**

Concept	Control Theory	Software Observability
System State	Internal physical variables (e.g., speed, position)	Application behavior and health (e.g., request rate, error rate)
Outputs	Measurable outputs (e.g., RPM, temperature)	Logs, metrics, traces from app and infra
Feedback	Adjust inputs to control system	Scale up, restart, alert, throttle



### **Metrics**

 $\bullet$ 







# Part III [Observability Stack]

### Metrics, Logs, Traces, Profiling ...





### **Metrics**

#### "Numerical values gathered from your system aggregated over time"

### **Metrics {example}**



**node\_load5**{app\_kubernetes\_io\_name="node-exporter",cluster="prod", instance="192.168.1.1:9100", kubernetes\_cluster="prod", kubernetes\_namespace="monitoring", monitoring\_scope="system", namespace="monitoring""} 2.81

- node\_load5 -> metric name
- kubernetes\_cluster="prod" -> key value pair aka label
- 2.81 -> **Actual value**

### **Metrics {definitions}**



LABYRINTH LABS

#### • Metric types

- **Counter** only goes up (e.g., requests total)
- **Gauge** goes up/down (e.g., memory usage)
- **Histogram** measures latency buckets
- **Summary** like histogram, but sampled
- **Cardinality** in observability refers to the number of **unique** combinations of **labels** (or dimensions) attached to a metric.

http\_requests\_total{ method="GET", status="200", user\_id="123456", uri="/auth/login"}

## **RED, USE & Golden Signals**



LABYRINTH LABS

- **RED** (Rate, Errors, Duration)
  - for services especially User experience
- **USE** (Utilization, Saturation, Errors)
  - Best for Hardware & infra health
- Golden Signals (Google SRE)
  - latency, traffic, errors, saturation
  - Fast insight into service health

# Metrics

#### Prometheus

- Time series collection happens via a pull model over HTTP
- Kubernetes service discovery
- Instrumentation
  - client libraries
  - exporters

#### Metric examples

- How much traffic is entering and leaving your network
- How many deployments are you making each day
- How much CPU is your service consuming



# Metrics

#### Prometheus

- Time series collection happens via a pull model over HTTP
- Kubernetes service discovery
- Instrumentation
  - client libraries
  - exporters

#### Metric examples

- How much traffic is entering and leaving your network
- How many deployments are you making each day
- How much CPU is your service consuming


# Metrics

#### Prometheus

- Time series collection happens via a pull model over HTTP
- Kubernetes service discovery
- Instrumentation
  - client libraries
  - exporters

#### Metric examples

- How much traffic is entering and leaving your network
- How many deployments are you making each day
- How much CPU is your service consuming



#### Prometheus

- Prometheus collects and **stores** its metrics as time series data
- Metadata are attached to  $\bullet$ metrics in optional key-value pairs called labels
- Has a **multi-dimensional** data model.
- Flexible query language PromQL



LABYRINTH

# What is Prometheus [intentionally] lacking



- No methods of data deduplication
- No methods for global view

#### Long term retention

- No concept of distributed data storage/sharding
- Keeping data long term is not feasible

A

I ARVRINTH

I ABS



### Challenges

- Prometheus needs to run in a Highly available setup
- We need to keep metrics indefinitely and store them efficiently
- We need a single observer point for all clusters

## Challenges



LABYRINTH LABS

Monitoring 5000+ pods in multiple clusters

# 19 TBs of metric data3 TBs ingested each month



#### Thanos

- Open source, highly available Prometheus setup with long term storage capabilities.
- Requires just object store AWS S3
- Global view, unlimited retention, high availability





#### Logs

#### "Records of events that have occurred within a software application or infrastructure."





47.29.201.179 - - [28/Feb/2019:13:17:10 +0000] "GET /?p=1 HTTP/2.0" 200 5316 "https://domain1.com/?p=1" "Mozilla/5.0 (Windows NT 6.1) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/72.0.3626.119 Safari/537.36" "2.75"

log\_format custom '\$remote\_addr - \$remote\_user [\$time\_local] "\$request" \$status \$body\_bytes\_sent "\$http\_referer"
"\$http\_user\_agent" "\$gzip\_ratio"';

We can see what **type** of request was sent to which URL, from which ip and what was our status code we used to respond.



#### LABYRINTH LABS

## **Log Collection**

#### Multiple ways to collect logs

- Application runtime
- Log collector sidecars
- Log collector DaemonSets

#### Many log collectors...

- fluentd
- fluentbit
- logstash
- Promtail
- vector.dev



#### LABYRINTH LABS

# **Log Collection**

#### **Customer Example**

- Storage 2TB/logov daily
- Throughput 51 MB/s
- During peak 500k logs/sec

Please do not enable debug logs in production. That can lead to serious performance issues.

### **Amazon OpenSearch Service**



LABYRINTH LABS

#### fully managed

native integration AWS services

search and analytics engine for use cases such as log analytics, real-time application monitoring, and clickstream analysis

AWS CloudTrail integration

	Log files					
	Messages			[	$\rightarrow$	Application & infrastructure monitoring
	Metrics	]				Security info & event management (SIEM)
	Config info		Amazon OpenSearch Service			Search
	Documents and lists		Dashboards and Kibana (1.5 to 7.10 versions)		$\rightarrow$	<b>Output</b> Search, analyze, and visualize logs to get real-time insights
1	Input Capture, process, and load data nto Amazon OpenSearch Servic	e				



### **Grafana Loki**

- Only indexes labels and igodolmetadata
- Simple query language LogQL  $\bullet$
- Multi-tenancy ullet
- Similar alerting concept to ulletPrometheus

А	(Loki)					
Query pa	tterns 🗸	Explain 🔵	Raw query 🦲	曰 Give feedback		
Labels service	e ~ = ·	~ web_app_	1 × × +			
Json		→ Label filter	expression	+ Operations	hint: add level_label_format()	
+ Exp	pression	Label Operator Value	status_code = ~ 200			
Raw quer {servic	<b>y</b> e="web_ap	p_1"}   json	status_code = `200	3,		
> Optic	ons Type: I	Range				
- Add qu	iery (	inspector				
						Л



#### LABYRINTH LABS

## **Grafana Loki**

- Only indexes labels and metadata
- Simple query language LogQL
- Multi-tenancy
- Similar alerting concept to Prometheus

	ideis			: нии reacned, received logs cover нь.4и% (утіп 53sec) от your selected time range (In) <b>нотанрутез processed</b> : 4.87 Ми
022-	89-2		:12:22 {	lb494d6abee562", , (Windows NT 6.1) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/47.0.2526.106 Safari/537.36"
1.00	lab	.1.0		
I		Θ	agent	Mozilla/5.0 (Windows NT 6.1) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/47.0.2526.106 Safari/537.36
al	Ð	Q	error_level	INFO
	Ð	Q	error_level_extracted	INFO
	€		http_method	PATCH
	€		http_method_extracted	PATCH
	€		path	php?foobar
	€		service	web_app_1
			status_code	200
			status_code_extracted	299
			tempo_trace_id	f84b494d6abee562
Dete	cte	d fi	elds 💿	
			Time	1664277142660
			agent	"Mozilla/5.0 (Windows NT 6.1) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/47.0.2526.106 Safari/537.36"
			error_level	*INFO*
			http_method	"Patch"
			path	"php?foobar"
			status_code	200
			tempo_trace_id	"f84b494d6abee562"
			tsNs	166427714266000000





#### Record of user requests across services, networks, and protocols to create a complete picture of how your distributed system works.



## **Distributed Tracing**

Tracks user requests across services, networks, and protocols and creates a complete picture of how your distributed system performed while processing a user request.

- Grafana Tempo
- OpenTelemetry, Jaeger, Zipkin



## **Distributed Tracing**

- Each activity (segment or span) is recorded as it through and across services
- Allows us to find and isolate bottlenecks quickly
- Identify the root cause of unseen problems



## AWS X-Ray

- Fully managed
- Creates a map of services used by your application with trace data that you can use to drill into specific services or issues
- X-Ray SDK captures metadata for requests made to MySQL and PostgreSQL databases (self-hosted, Amazon RDS, Amazon Aurora), and Amazon DynamoDB



Name	Res.	Duration	Status	0.0ms	100ms I	200ms I	300ms I	400ms	500ms I	600ms I	700ms I	800ms I	900ms I	
dev-portal-api-AppSyncApi AWS::AppSync::Graphe	QLAPI													
dev-portal-api-AppSyncApi	200	801 ms										POST	Z3rtgfxze5bk5fy3yu6e4hkiea.apps	sync-api
/categories	-	0.0 ms		5										
requestMappingTemplateEvaluation	-	0.0 ms												
/companies	-	0.0 ms		<b>-</b>										
requestMappingTemplateEvaluation	-	0.0 ms												
Query.categories	-	719 ms		51										
Lambda	200	693 ms		 								Invoke: d	lev-portal-api-portal-AppSyncApiLa	ambdaQ
Query.companies	-	720 ms		Ly	-									
Lambda	200	699 ms		 								Invok	e: dev-portal-api-portal-AppSyncA	piLambo
responseMappingTemplateEvaluation	-	0.0 ms		L								- 1		
responseMappingTemplateEvaluation	-	0.0 ms												
dev-portal-api-portal-AppSyncApiLambdaQue	ycatego	r-vi4tFhQkto	0 AWS::La	ımbda										
dev-portal-api-portal-AppSyncApiLambdaQueryca	200	661 ms												
dev-portal-api-portal-AppSyncApiLambdaQue	ycompa	ni-R3PM2AV	cwmJI AV	VS::Lambda										
dev-portal-api-portal-AppSyncApiLambdaQueryco	200	668 ms												



#### LABYRINTH LABS

### Grafana Tempo

- Distributed tracing backend
- Cost-efficient, needs just object storage
- Multiple tracing protocols
- Allows us to find and isolate bottlenecks quickly
- Identify the root cause of unseen problems

🔫 Tempo 🗸			Close 🔠 Add to dashboard		ର <mark>ଓ</mark> ୪	<u></u>	
Query type Search TraceID J	SON file Service Graph						
Trace ID 5be939944da6e15a							
+ Add query 💿 Inspector							
Trace View				Find			
for a to a for the state of the							
frontend: /product 5be939944da6e15a							
Trace Start: 2022-09-27 11:02:46.763 Durat	ion: 1.32s Services: 14 Depth: 4 Te	otal Spans: <b>50</b>					
0µs	329.5ms	659ms	988.5ms		1.32s		
	1 1944 - C						
		000 5	650	000 5	1.00-		
Service & Operation $\checkmark$ > $\Rightarrow$ »	Uµs	329.5ms	659ms	988.5ms	1.325		
rrontend /product (1.32s)							
spanFiller9 /9 (0µs)	l Oµs						
span-liler8 /8 (0µs)	I 0µs						
spanEiller() /0 (0uo)	126ms						
spanFiller6 /6 (0us)	l Ouro						
> recommendationservice /GetRecom	r ops	415ms					
	(a a		attender bei der	0	00.46.070)		
	/GetRecommendat	IONS Service: recommend	ationservice   Duration: 415ms	Start Time: 110ms (11: Chil	02:46.873) d Count: 11		
	> Attributes: http.method = GE	T http.status_code = 200 http	.url = http://recommendationservic	e/GetRecommendations			
	> Resource: host.name = syntl	netic-load-generator-7ffc9cdc65-k4	lc4x   ip = 10.147.107.169   open	census.exporterversion =	Jaeger-Ja		
				ලා SpaniD: a87	ifaf1a53bc5c3		
spanFiller4 /4 (0µs)	l 0μs   spanFiller4::/4						
spanFiller1 /1 (0µs)	l Oµs						
spanFiller5 /5 (0µs)	l Oµs						
spanFiller2 /2 (0µs)	l Oµs						
spanFiller3 /3 (0µs)	l Oµs						
spanFiller7 /7 (0µs)							
> adservice /AdRequest (961ms)	961ms						



### **Visualization & Analysis**

#### How it all comes together...



### **Visualization with Grafana**

- Single point of view for all metrics, logs and traces
  - Mixed datasources
  - Alerting
  - Open-source (self-hosted)

















## Alerting

#### An alert is a last-resort signal that tells you:

#### The system has gone outside the boundaries you defined as acceptable — someone needs to act."



#### LABYRINTH

# Alerting

#### Boundaries comes from

- SLIs/SLOs
  - latency must stay below 300ms
- Infrastructure rules
  - CPU > 90% for 5 minutes
- Business impact thresholds
  - $\circ$  0 checkouts in 3 mins
- Hardware based limits



I ARVRINTH

# Alerting

**Really**, **really** hard to do

- Normal is **relative** 
  - Today's traffic baseline ≠ tomorrow's
- Dynamic infrastructure
  - Pods restart, services autoscale, server die
- **Infrastructure** noise can mask or trigger false alerts

https://sre.google/books/chapters/alerting-on-slos/

https://landing.google.com/sre/sre-book/chapters/alerting-introduction/

### **Continuous Profiling**



- Find performance bottlenecks
- Optimize hot code paths (functions burning CPU or memory)
- Reduce cloud costs
- Catch performance regressions

LABYRINTH

I ABS

A

### **Continuous Profiling**



LABYRINTH

LABS

ഷ



#### LABYRINTH LABS

## Summary

- Pillars of observability
  - Metrics
    - Prometheus
    - Thanos
  - $\circ$  Logs
    - OpenSearch
    - Grafana Loki
  - Traces
    - AWS X-Ray
    - Grafana Tempo

# Part VI [Observability Prax]

#### **Implementation & lessons learned**




### Labyrinth Labs [Approach]

 $\bullet$ 

73

# Part IV [Q&A]





LABYRINTH LABS



<u>slido.com</u> -> 3822198





LABYRINTH LABS

Without good **Observability** your application is both **working** and **crashed** at the same time until ....

Your **customer calls** your boss.



#### **Contact us**

Adam Hamšík CEO

adam@lablabs.io www.linkedin.com/in/adam-hamsik/



77



#### 









## #NOBULLSH\*T







