



Comprehensive Observability via Distributed Tracing

Chinmay Gaikwad, Technical Evangelist

Twitter: @epsagon



Hello!

- Software Engineer, Applications Engineer, Technical Marketing Engineer: Intel, IBM, early stage startups
- Traveling, Soccer, Restaurants, Video Games



What we'll discuss today

- Microservices: The New Normal and New Challenges
- Troubleshooting Distributed Environments
- Benefits of Distributed Tracing

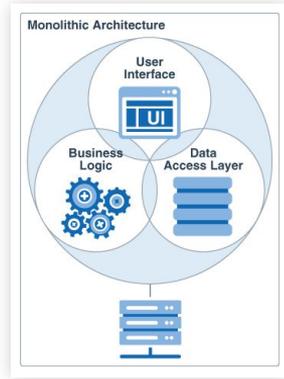
} WHY?

} HOW?

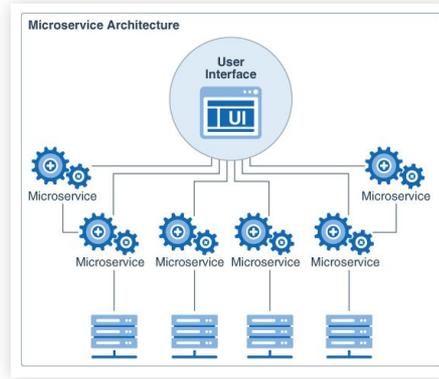


Microservices: The New Normal and New Challenges

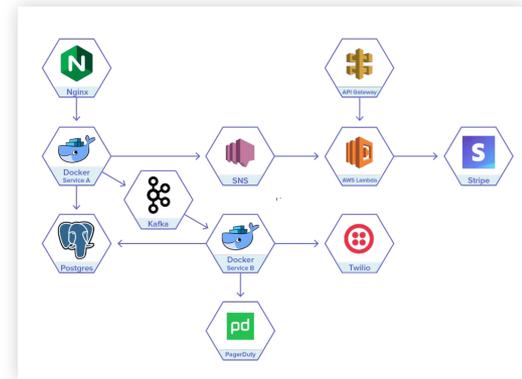
The Rise of Microservices



Host-based
Monolithic



Host-based
Distributed

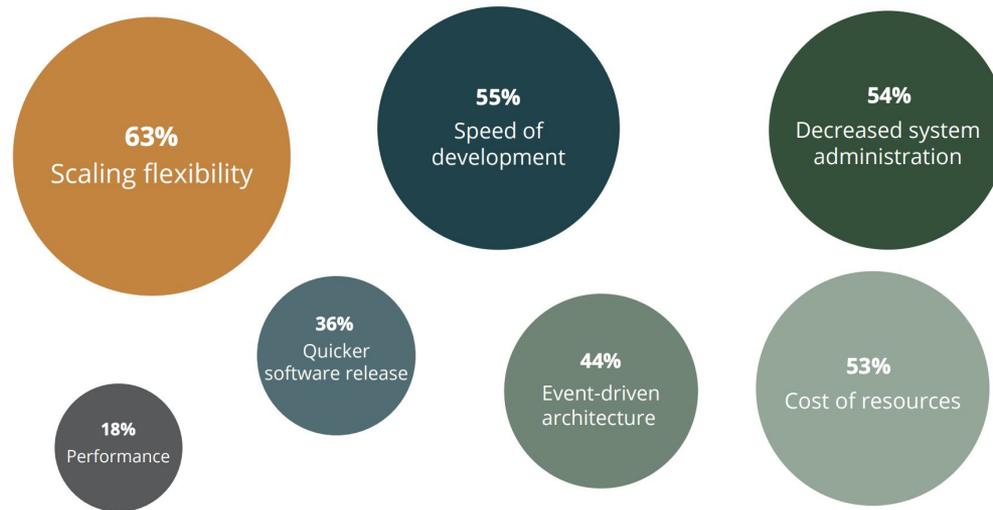


Abstracted-host
Highly Distributed

Extremely hard to monitor and troubleshoot

Why Microservices?

What are the biggest benefits of using Serverless for your organization?



Source: State of Serverless Report 2020 - CodingSans

New Paradigm, New Challenges



Difficulty Identifying & Troubleshooting Issues

Customer-facing impacts (downtime, latency)
Decreased velocity of new feature releases



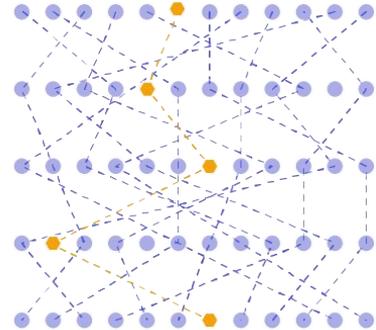
Traditional Monitoring from Multiple Sources

Lack of application insights & visibility into errors
Difficulty correlating data



Incomplete Data Insights

Sampling, resulting in gaps
No visibility into payloads





Troubleshooting Distributed Applications

Observability: Overview

- Observability: Actively debug a system
- Monitoring: Watch and understand the state of a system
- Monitoring and observability is one of a set of capabilities that drive higher software delivery and organizational performance
- Who is monitoring and observability for? Everyone!

Source: DORA research

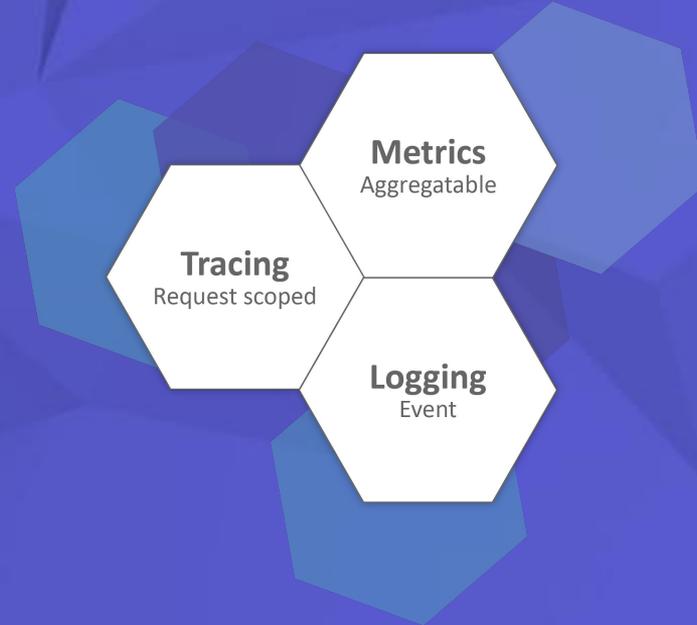
Achieving Observability in Microservices

Combining metrics, logs, and traces for observability is the only way to understand complex environments

Metrics tell us the “**what**”

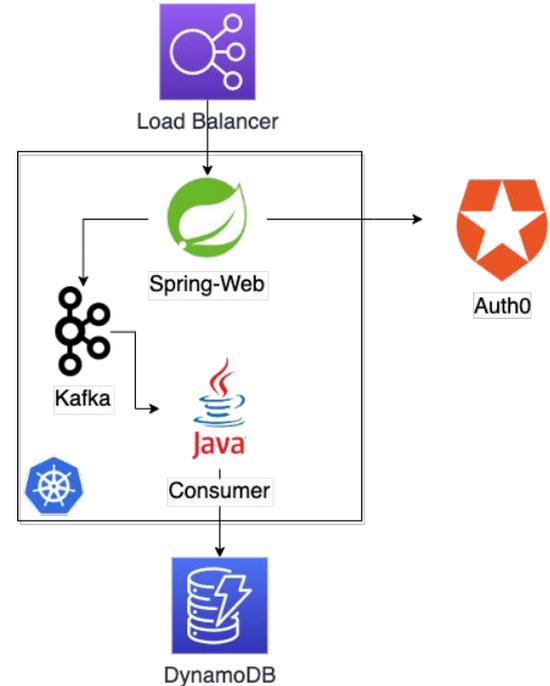
Logs tell us the “**why**”

Traces tell us the “**where**”



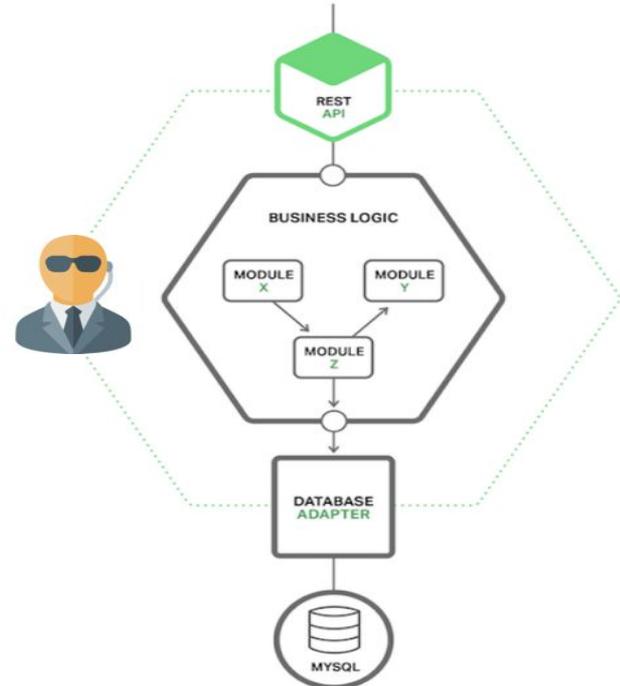
Distributed Troubleshooting Use Case

- The service implements a simple **virtual shop**, where users can send orders for items
- The **HTTP server** authenticates requests using **Auth0 API** (3rd party) and pushes them to a **Kafka** stream
- Another **Java container** polls the stream and updates the orders on a **DynamoDB** table
- Both containers run + Kafka stream runs on **Kubernetes**
- Users complain about orders that were sent but not handled



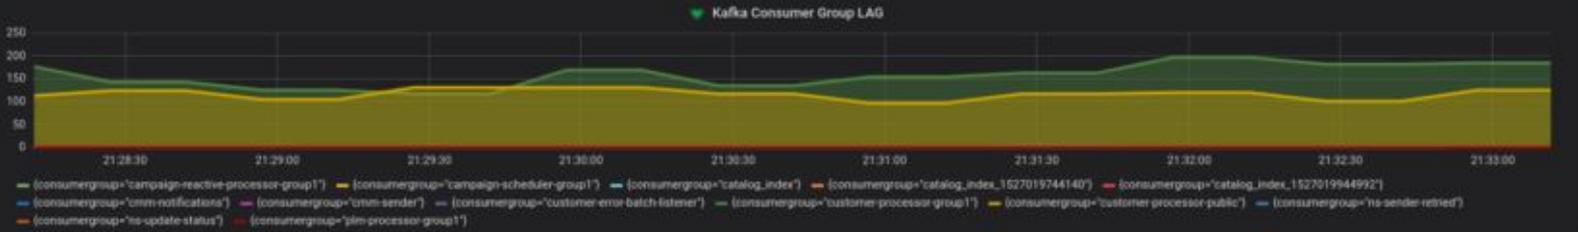
Old School Monitoring

- Heavy Instrumentation
- Collects only host data
- Collects only metrics



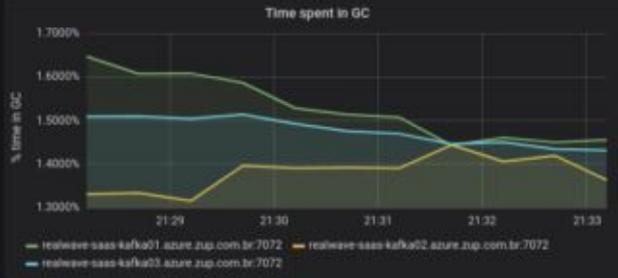
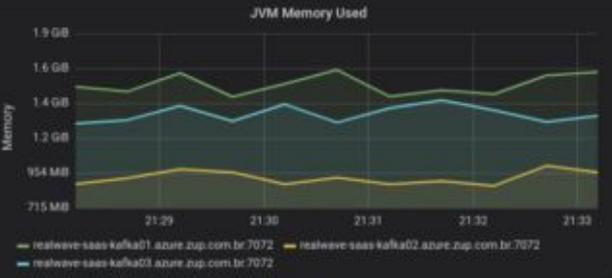
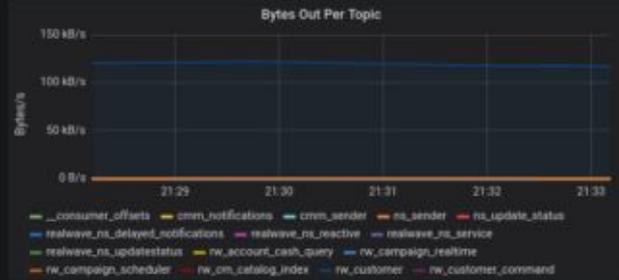
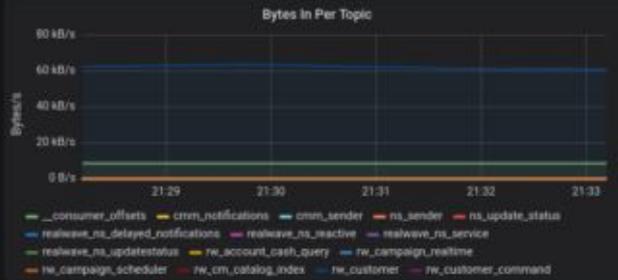
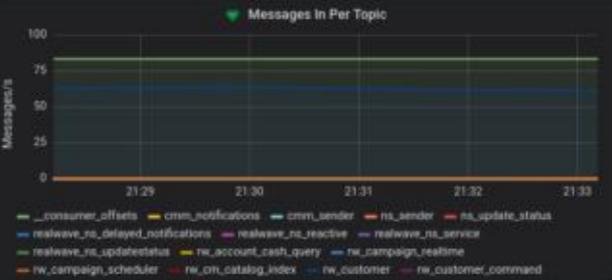
Kafka Metrics

Job kafka_exporter Topic __consumer_offsets ConsumerGroup None

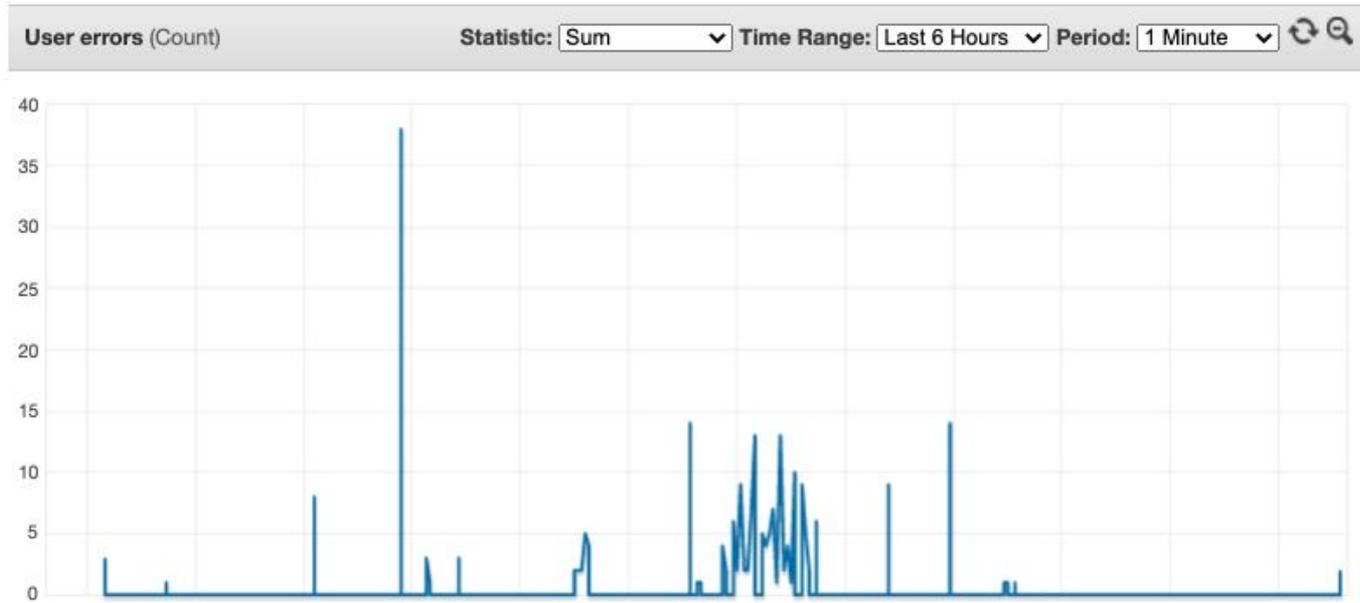


Kafka Total Partitions

888



DynamoDB Metrics



We need more debug data → logs

Search for log entries... (e.g. host.name:host-1)

2018-11-08 20:14:30.309	redis Connecting to MASTER redis-master:6379	
2018-11-08 20:14:50.378	redis MASTER <=> SLAVE sync started	
2018-11-08 20:15:04.000	apache2 192.168.65.3 - "GET /guestbook.php?cmd=set&key=messages&value=Lorem+ipsum+dolor+sit+amet%2C+consectetur+adipiscing,+B%27duis%27+b%27mi%27+b%27a%27+b%27quam%27,+B%27eget%27+b%27ad%27+b%27a%27+ HTTP/1.1" 200 255	09 AM
2018-11-08 20:15:04.000	apache2 192.168.65.3 - "GET / HTTP/1.1" 200 826	
2018-11-08 20:15:13.000	apache2 192.168.65.3 - "GET / HTTP/1.1" 200 826	
2018-11-08 20:15:13.000	apache2 192.168.65.3 - "GET /guestbook.php?cmd=set&key=messages&value=Lorem+ipsum+dolor+sit+amet%2C+consectetur,+B%27eros%27+b%27ve%27+b%27a%27+b%27a%27+b%27a%27+b%27curabitur%27+b%27a%27+b%27odio%27+b%27a HTTP/1.1" 200 255	12 PM
2018-11-08 20:15:16.000	REPLCONF ACK 13902	
2018-11-08 20:15:17.000	apache2 192.168.65.3 - "GET / HTTP/1.1" 200 826	
2018-11-08 20:15:17.000	apache2 192.168.65.3 - "GET /guestbook.php?cmd=set&key=messages&value=Lorem+ipsum,+B%27amet%27+b%27ac%27,+B%27orci%27+b%27ut%27,+B%27pede%27+b%27eu%27+b%27ac%27+b%27quam%27+b%27a%27+b%27a%27+b%27a%27+b%27enim%27+b%27a%27.+ HTTP/1.1" 200 255	
2018-11-08 20:15:19.000	REPLCONF ACK 14037	
2018-11-08 20:15:20.000	apache2 192.168.65.3 - "GET /guestbook.php?cmd=set&key=messages&value=Lorem+ipsum+dolor+sit+amet%2C+consectetur+adipiscing,+B%27arcu%27+b%27ut%27+b%27a%27+b%27a%27+b%27a%27+b%27tactic%27+b%27a%27,+B%27 HTTP/1.1" 200 255	03 PM
2018-11-08 20:15:20.000	apache2 192.168.65.3 - "GET / HTTP/1.1" 200 826	
2018-11-08 20:15:21.000	apache2 192.168.65.3 - "GET /guestbook.php?cmd=set&key=messages&value=Lorem+ipsum+dolor,+B%27quis%27+b%27in%27+b%27a%27+b%27sem%27+b%27sed%27+b%27nibh%27+b%27amet%27+b%27a%27+b%27a%27+b%27nibh%27,+B%27nisi%27+b%27ad%27 HTTP/1.1" 200 255	
2018-11-08 20:15:21.000	apache2 192.168.65.3 - "GET / HTTP/1.1" 200 826	
2018-11-08 20:15:26.000	apache2 192.168.65.3 - "GET / HTTP/1.1" 200 826	
2018-11-08 20:15:26.000	apache2 192.168.65.3 - "GET /guestbook.php?cmd=set&key=messages&value=Lorem+ipsum+dolor+sit+amet,+B%27nunc%27+b%27et%27+b%27egestas%27+b%27a%27+b%27nec%27+b%27sociis%27+b%27purus%27+b%27nec%27,+B%27eget%27+ HTTP/1.1" 200 255	06 PM
2018-11-08 20:15:34.000	apache2 192.168.65.3 - "GET /guestbook.php?cmd=set&key=messages&value=Lorem+ipsum+dolor+sit+amet%2C+consectetur+adipiscing+elit+b%27amet%27+b%27mi%27,+B%27ante%27+b%27in%27+b%27laoreet%27+b%27s HTTP/1.1" 200 255	
2018-11-08 20:15:34.000	apache2 192.168.65.3 - "GET / HTTP/1.1" 200 826	
2018-11-08 20:15:41.000	apache2 192.168.65.3 - "GET /guestbook.php?cmd=set&key=messages&value=Lorem+ipsum+dolor,+B%27pede%27+b%27ad%27,+B%27ante%27+b%27ve%27+b%27a%27+b%27etiam%27,+B%27nunc%27+b%27eu%27+b%27a%27+b%27a%27+b%27a%27+ HTTP/1.1" 200 255	
2018-11-08 20:15:41.000	apache2 192.168.65.3 - "GET /guestbook.php?cmd=set&key=messages&value=Lorem+ipsum+dolor,+B%27quis%27+b%27mi%27+b%27ut%27+b%27id%27+b%27a%27+b%27a%27+b%27nam%27,+B%27eros%27+b%27et%27+b%27a%27+b%27id%27 HTTP/1.1" 200 826	
2018-11-08 20:15:41.000	apache2 192.168.65.3 - "GET / HTTP/1.1" 200 826	
2018-11-08 20:15:44.543	redis Connecting to MASTER redis-master:6379	
2018-11-08 20:15:44.543	redis Timeout connecting to the MASTER...	
2018-11-08 20:15:44.544	redis MASTER <=> SLAVE sync started	
2018-11-08 20:15:50.000	apache2 192.168.65.3 - "GET /guestbook.php?cmd=set&key=messages&value=Lorem+ipsum+dolor+sit+amet%2C+consectetur+adipiscing+elit+b%27quam%27,+B%27a HTTP/1.1" 200 255	
2018-11-08 20:15:50.000	apache2 192.168.65.3 - "GET / HTTP/1.1" 200 826	
2018-11-08 20:15:51.675	redis Timeout connecting to the MASTER...	
2018-11-08 20:15:51.675	redis Connecting to MASTER redis-master:6379	
2018-11-08 20:15:51.676	redis MASTER <=> SLAVE sync started	
2018-11-08 20:15:57.000	apache2 192.168.65.3 - "GET /guestbook.php?cmd=set&key=messages&value=Lorem+ipsum+dolor+sit+amet%2C+consectetur+adipiscing+elit+b%27urna%27,+B%27 HTTP/1.1" 200 255	
2018-11-08 20:15:57.000	apache2 192.168.65.3 - "GET / HTTP/1.1" 200 826	
2018-11-08 20:16:00.000	apache2 192.168.65.3 - "GET /guestbook.php?cmd=set&key=messages&value=Lorem+ipsum+dolor+sit,+B%27amet%27+b%27et%27+b%27eu%27+b%27ve%27+b%27a%27+b%27eros%27+b%27hac%27+b%27a%27+b%27a%27+b%27a%27+b%27tel%27,+B%27 HTTP/1.1" 200 255	
2018-11-08 20:16:00.000	apache2 192.168.65.3 - "GET / HTTP/1.1" 200 826	
2018-11-08 20:16:01.000	apache2 192.168.65.3 - "GET /guestbook.php?cmd=set&key=messages&value=Lorem+ipsum+dolor+sit+amet,+B%27eget%27+b%27ac%27,+B%27elit%27+b%27 HTTP/1.1" 200 255	
2018-11-08 20:16:01.000	apache2 192.168.65.3 - "GET / HTTP/1.1" 200 826	
2018-11-08 20:16:09.000	apache2 192.168.65.3 - "GET / HTTP/1.1" 200 826	
2018-11-08 20:16:09.000	apache2 192.168.65.3 - "GET /guestbook.php?cmd=set&key=messages&value=Lorem+ipsum+dolor+sit+amet%2C+consectetur+adipiscing,+B%27quis%27+b%27 HTTP/1.1" 200 255	
2018-11-08 20:16:11.000	apache2 192.168.65.3 - "GET /guestbook.php?cmd=set&key=messages&value=Lorem+ipsum+dolor+sit,+B%27nunc%27+b%27id%27+b%27a%27,+B%27arcu%27+ HTTP/1.1" 200 255	



Java Logs

```
[http-nio-8080-exec-10] INFO io.jaegertracing.internal.reporters.LoggingReporter - Span reported: 615f47e4c32f589d:4e8220be4a768563:615f47e4c32f589d:1 - placeNewOrder
```

```
[http-nio-8080-exec-10] INFO io.jaegertracing.internal.reporters.LoggingReporter - Span reported: 615f47e4c32f589d:615f47e4c32f589d:0:1 - POST
```

```
[kafka-producer-network-thread | producer-1] INFO io.jaegertracing.internal.reporters.LoggingReporter - Span reported: 615f47e4c32f589d:9b14b78b08321244:4e8220be4a768563:1 - produce
```

```
09:26:16.894 [http-nio-8080-exec-27] INFO com.epsagon.java.rest.OrdersService - placing new order {}
```

```
09:26:16.894 [http-nio-8080-exec-27] INFO c.epsagon.java.kafka.producer.Sender - sending new order='NewOrder{itemId=0, username='9a7ed47bfe21c01387fa3d93d3each', discountCode='XMASSAVE30', quantity=4}' to topic='queuing.retail_site.new_orders'
```

```
09:26:17.242 [http-nio-8080-exec-27] ERROR Missing required parameter in input: "Key"  
Unknown parameter in input: "Item", must be one of: TableName, Key, AttributeUpdates, Expected, ConditionalOperator, ReturnValues, ReturnConsumedCapacity, ReturnItemCollectionMetrics, UpdateExpression, ConditionExpression, ExpressionAttributeNames, ExpressionAttributeValues
```

Things missing?

- How do we correlate between metrics and logs?
- How do we correlate data between different services?
- How do we find the **where** when something goes wrong?

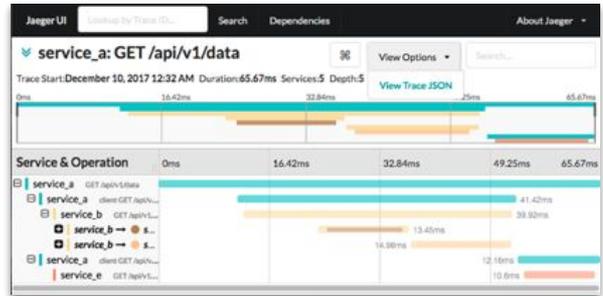
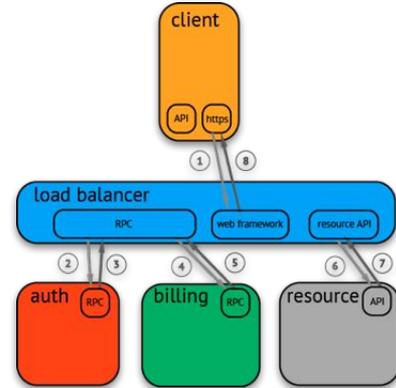


Distributed Tracing

What is Distributed Tracing?

“A **trace** tells the story of a transaction or workflow as it propagates through a distributed system.”

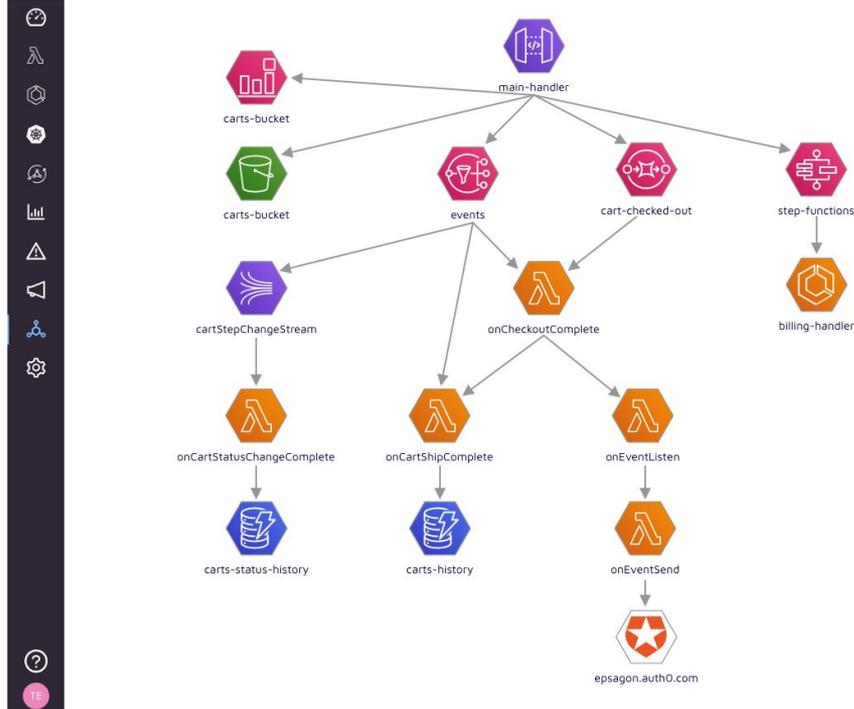
Since distributed tracing connects every request in a transaction, it allows you to know and see what’s happening to every service component and app in production



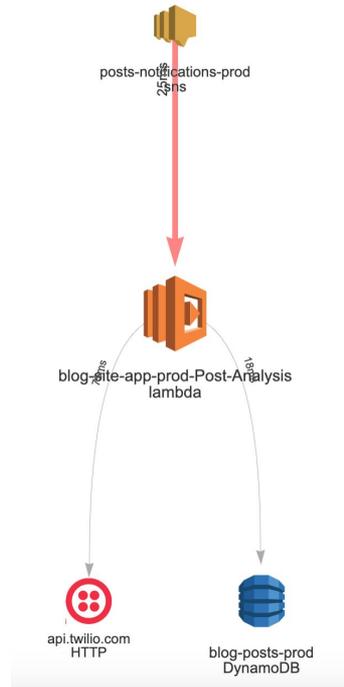


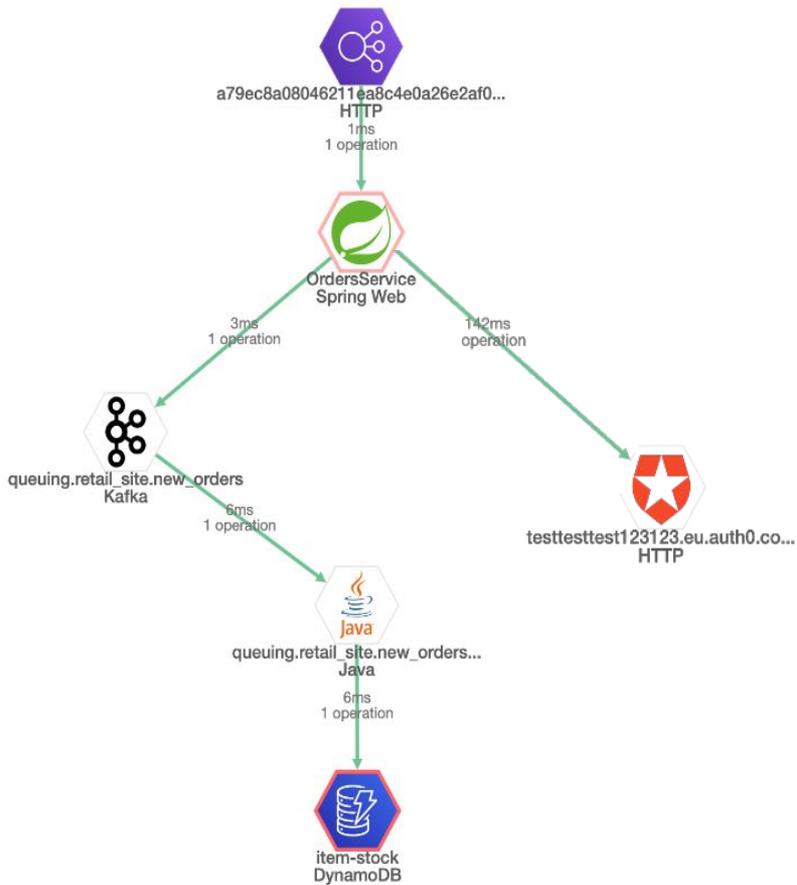
Benefits of Distributed Tracing

Visualize and Understand



Bring Focus to the Problems





UpdateItem | 5.73ms

Sep 14, 2020 8:01:18.326 PM



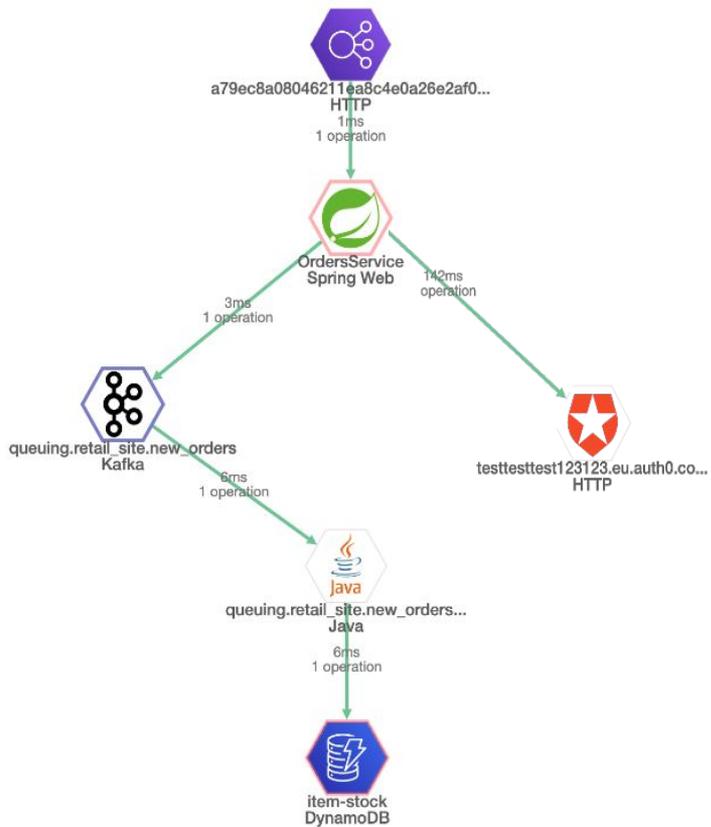
An error occurred (ValidationException) when calling the PutItem operation: One or more parameter values were invalid: Missing the key id in the item

[Collapse](#)

Tags

[Index Tags](#)

component	aws-sdk
error	True
hostname	stock-updater-856884bbd6-9t97s
ip	100.96.3.58
is_k8s	true
k8s_pod_name	stock-updater-856884bbd6-9t97s
aws.agent	aws-sdk
aws.agentVersion	>1.11.0
aws.endpoint	https://dynamodb.us-east-1.amazonaws.com
aws.operation	PutItemRequest
aws.region	us-east-1
aws.service	AmazonDynamoDBv2
env.runtime	opentracing-java
epsagon.version	Java-0.35.4
http.method	POST
http.url	https://dynamodb.us-east-1.amazonaws.com
span.kind	client
aws.dynamodb.table_n...	item-stock



queuing.retail_site.new_orders (1 operation)

Service Map

→ | produce | 2.66ms
Sep 14, 2020 8:01:17.260 PM

Tags

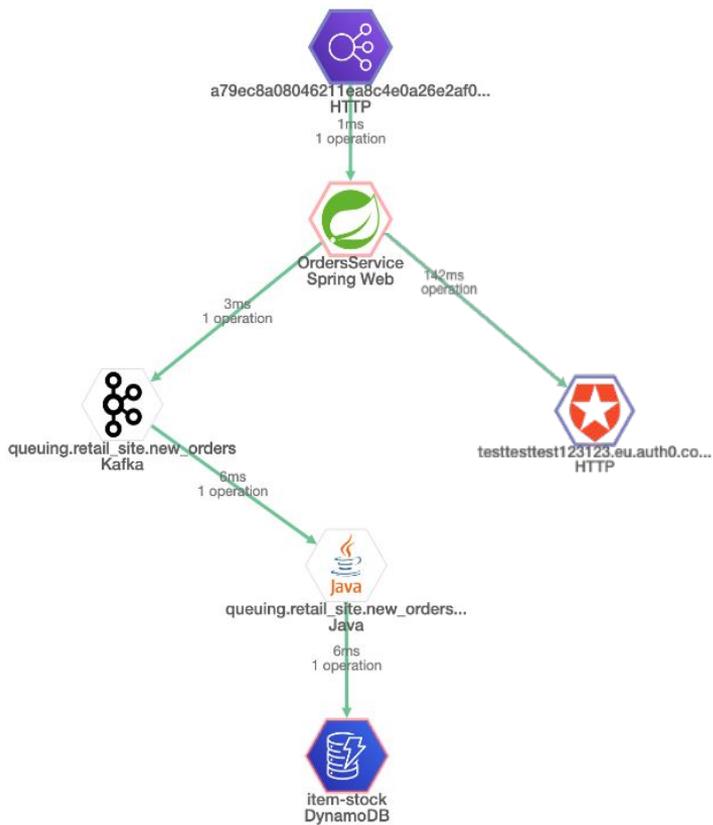
Index Tags

component	kafka-clients-0.11
hostname	orders-service-64cbdfcb5c-b2ghw
ip	100.96.1.35
is_k8s	true
k8s_pod_name	orders-service-64cbdfcb5c-b2ghw
env.runtime	opentracing-java
epsagon.version	Java-0.35.4
kafka.key	null
span.kind	producer

JSON Tags

kafka.value

```
{ 4 items
  "itemId" : 1
  "username" : ""
  "discountCode" : "DEL15"
  "quantity" : 122
}
```



Tags

[Index Tags](#)

http.host	testtesttest123123.eu.auth0.com
http.scheme	https
http.status_code	401
http.request.path	/api/v2/users/auth0%7C5ba1a9227dc7232e1aec4fd0

JSON Tags

http.request.headers >

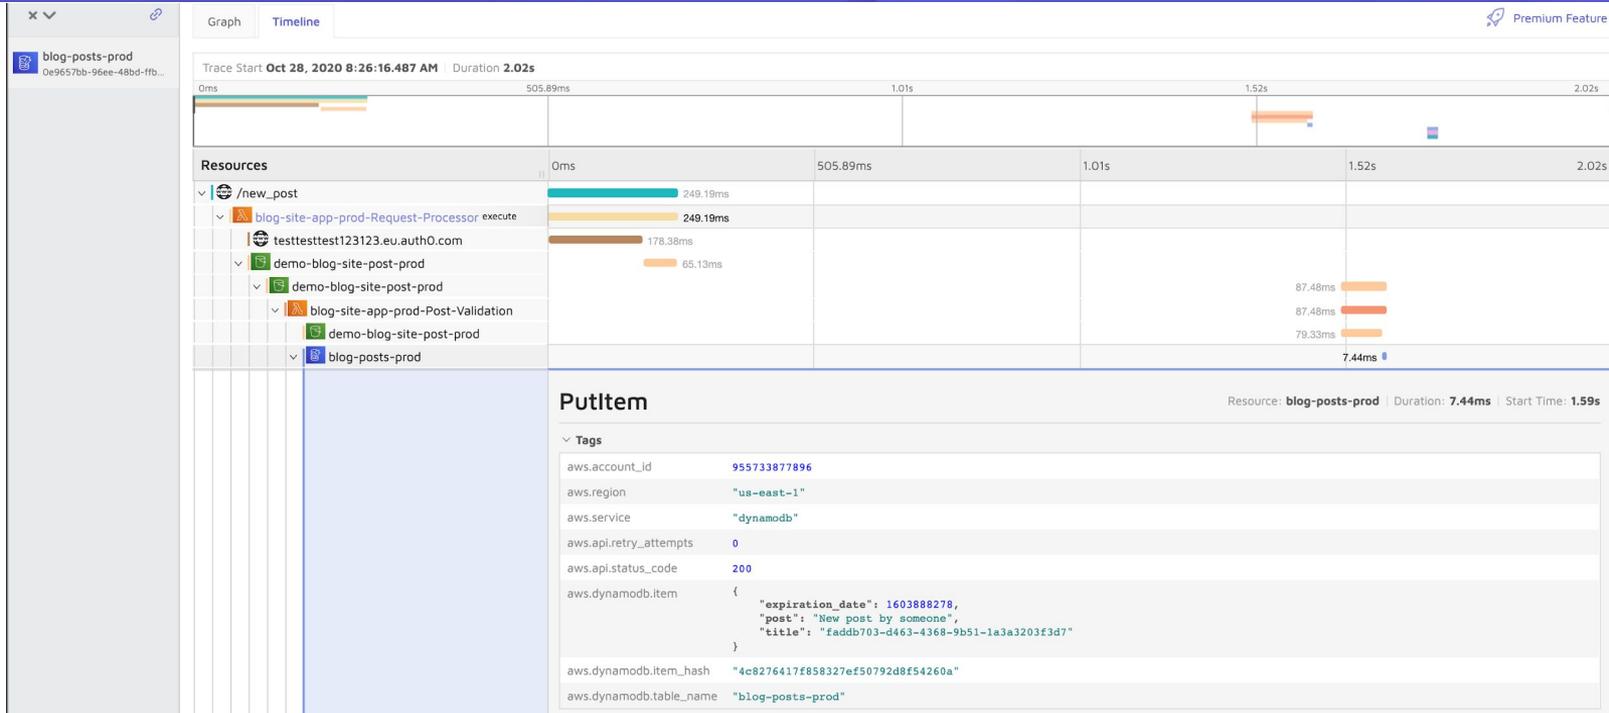
http.response.body >

```
{ 4 items  
  "message": "Expired token received for JSON Web Token validation"  
  "statusCode": 401  
  "error": "Unauthorized"  
  "attributes": { 1 item  
    "error": "Expired token received for JSON Web Token validation"  
  }  
}
```

http.response.headers >

```
{ 21 items  
  "CF-Cache-Status": "DYNAMIC"
```

Where Does Our Code Spend Time?

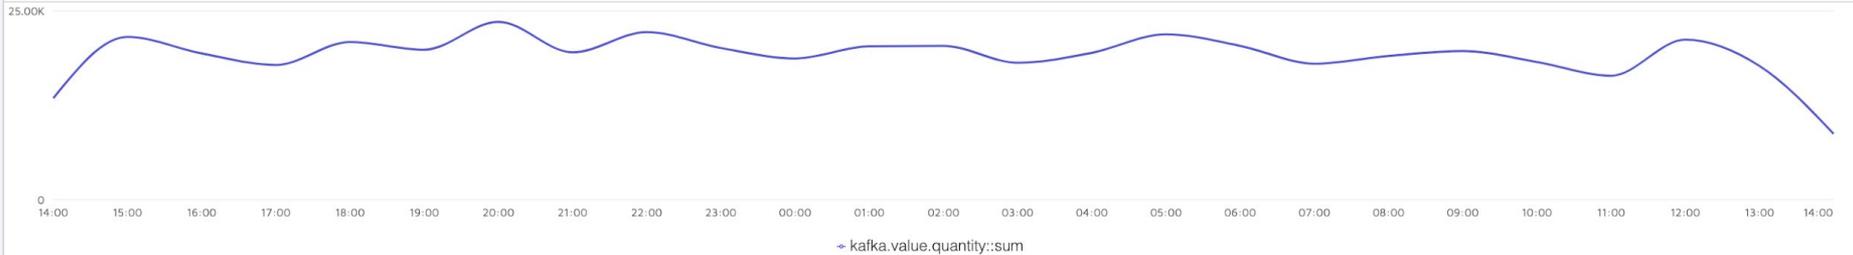


Business Insights

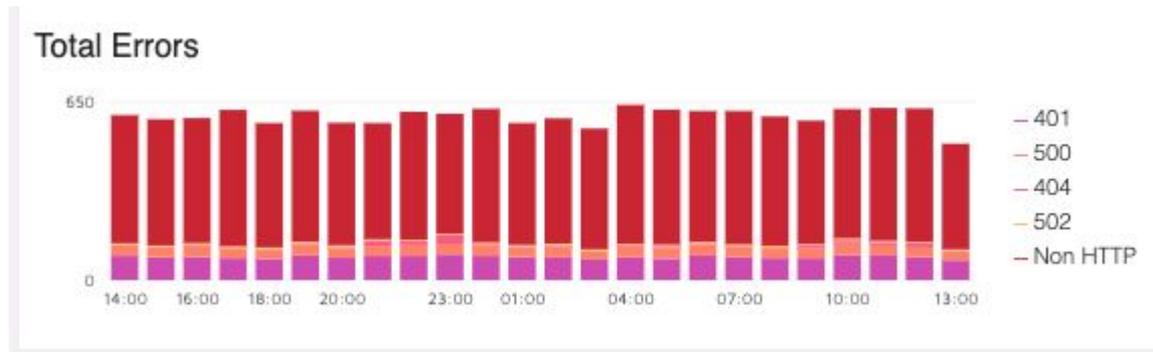
Count of error



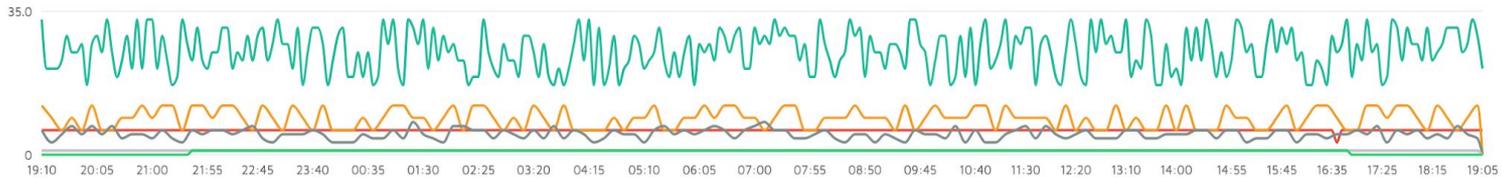
Total Items Ordered



Errors, Categorized



Count of error by exception.type



[Add Chart to Custom Dashboard](#)

- error::ConnectTimeoutException
- error::ConditionalCheckFailedException
- error::ZeroDivisionError
- error::TypeError
- error::JSONDecodeError
- error::ReadTimeoutError

Monitor with Trace-based Metrics and Alerts

The screenshot displays the 'Kubernetes Nodes (4)' monitoring page. The selected node is 'ip-172-20-38-39.ec2.internal'. The main metrics section shows four charts: CPU usage (peaking at 6%), Memory usage (fluctuating between 743.87 MB and 801.09 MB), Network Receive (constant at 1.95 KB), and Network Sent (constant at 1.95 KB). Below the charts is the 'Node Information' table.

Node Name	Cluster	Provider ID	CPU	Memory	Disk	Network
ip-172-20-38-39.ec2.internal	https://api.retail-site-demo.k8s.epsagon.uk	aws:///us-east-1a/i-007383206775fc346	3.693%	776.39 MB	R: 1.85 KB W: 6.93 KB	R: 4.23 KB T: 4.1

Node Information									
Date Created	5 months ago	Internal IP	172.20.38.39	External IP	3.80.55.226	Internal DNS	ip-172-20-38-39.ec2.internal	Hostname	ip-172-20-38-39.ec2.interna
External DNS	ec2-3-80-55-226.compute-1.amazonaws.com								
Machine ID	5207f949b6a24d4ab82b539b0cf4e8df			System UUID	EC2C10FF-2133-C2AD-2A07-BBD16365280F				
Boot ID	Od7f5fec-df67-4a35-bf1c-6d2ddf8a0374			Kernel Version	4.9.0-9-amd64				
Os Image	Debian GNU/Linux 9 (stretch)			Container Runtime Version	docker://18.6.3				

OpenTelemetry Framework, Open-source Tooling

- OpenTelemetry is a framework, not a service!
- Jaeger (Uber) and Zipkin (Twitter)
- Manual tracing requires heavy lifting: instrumentation and maintenance
- Lack visualizations, context, and tracing *through* middleware



JAEGER



Generating Traces with OpenTelemetry

- Instrument every call (AWS-SDK, http, postgres, Spring, Flask, Express, ...)
- Create a **span** for every request and response
- Add **context** to every span
- **Inject** and **Extract** IDs in relevant calls

```
def handle_request(request):
    span = before_request(request, opentracing.global_tracer())
    # store span in some request-local storage using Tracer.scope_manager,
    # using the returned `Scope` as Context Manager to ensure
    # `Span` will be cleared and (in this case) `Span.finish()` be called.
    with tracer.scope_manager.activate(span, True) as scope:
        # actual business logic
        handle_request_for_real(request)

def before_request(request, tracer):
    span_context = tracer.extract(
        format=Format.HTTP_HEADERS,
        carrier=request.headers,
    )
    span = tracer.start_span(
        operation_name=request.operation,
        child_of(span_context))
    span.set_tag('http.url', request.full_url)

    remote_ip = request.remote_ip
    if remote_ip:
        span.set_tag(tags.PEER_HOST_IPV4, remote_ip)

    caller_name = request.caller_name
    if caller_name:
        span.set_tag(tags.PEER_SERVICE, caller_name)

    remote_port = request.remote_port
    if remote_port:
        span.set_tag(tags.PEER_PORT, remote_port)

    return span
```

The Epsagon Solution



Easy to use

5 minute setup, fully automated, no training or maintenance required



Runs anywhere

Kubernetes, ECS, containers, serverless, and more



Correlated

metrics, logs, traces with payload visibility



End-to-end

product, from monitoring to troubleshooting across Ops and Dev

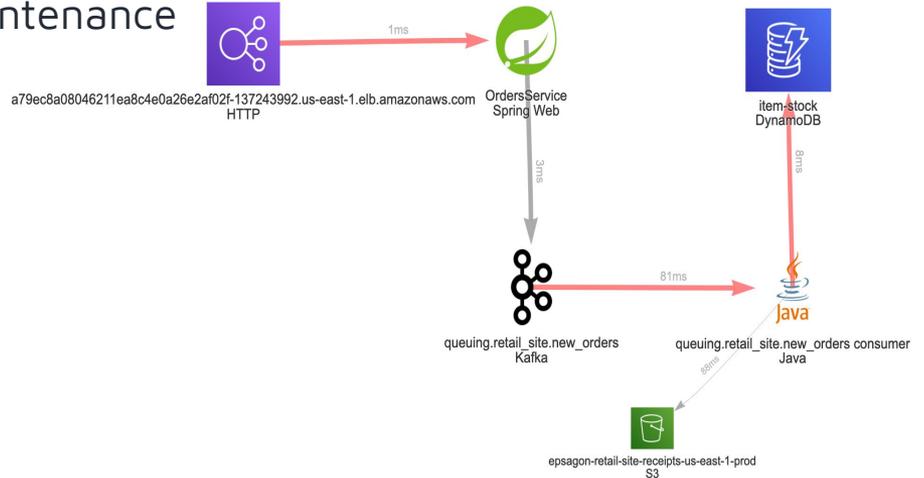


**SOC 2
TYPE II
CERTIFIED**



Best Practices for Observability

- **Automated setup** and minimal maintenance
- Support **any environment** (containers, K8s, cloud, Serverless)
- Connects **every request** in a transaction
- Search and **analyze** your data
- Helps to quickly **pinpoint** problems



The Journey to Observability



- Identify your business goals and architecture model
- Determine your approach: DIY or managed
- Implement observability solutions
- Ensure scalability of observability strategy

Summary

- Distributed applications bring unique benefits and challenges
- Advantages of using Distributed Tracing approach
- Observability is critical to:
 - Keep track of the architecture
 - Detect performance issues and reduce MTTR
 - Reduce Ops, Dev and Opportunity costs

Be **PROACTIVE** not REACTIVE





Thank you!

chinmay@epsagon.com

To learn more and for our special offer visit:
<https://epsagon.com/skilupglobal/>