


Observability? It's a Data Analytics Problem, You Fool!

Dotan Horovits

@horovits








Dotan Horovits

Principal Developer Advocate



     @horovits

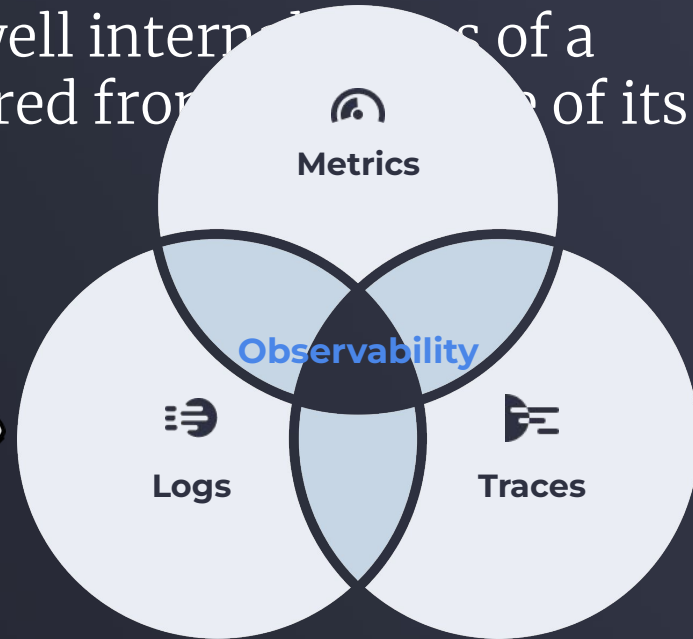


Observability [n.] (Control Theory)

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

Observability [n.] (Control Theory)

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



Observability [n.] (Software, DevOps)

the new definition!

The capability to allow a human to ask and answer questions about the system.

Observability [n.] (Software, DevOps)

the new definition!

The **capability** to allow a human to ask and answer questions about the system.



A property of the system

Observability [n.] (Software, DevOps)

the new definition!

The capability to allow a human to **ask and answer** questions about the system.

A data analytics problem





Observability

So what is observability really about?

Logs

Metrics

Tracing




observability is about

- ... collecting data of different sources, formats and types
- ... structured standardized data
- ... enriching and correlating data
- ... unified query and visualization
- ... data volume & data:noise ratio

Observability is about

**Collecting data of different sources,
formats and types**

It's about collecting data of different sources, formats and types



It's about collecting data of different sources, formats and types



🌸 Alolita Sharma 🌸 @alolita · Aug 3, 2021

O11y needs MELT.

M=Metrics; E=Events; L=Logs; T=Traces

Read the OTEP on Real-time user monitoring (RUM)
[github.com/open-telemetry...](https://github.com/open-telemetry) Please provide your f

[#OpenTelemetry](#) [#Observability](#) [#OpenSource](#)



Yuri Shkuro @YuriShkuro · Sep 19, 2022


It's time to stop talking about three pillars of **observability** and start talking about ... (drumbeat) ... six pillars!



medium.com

TEMPLE: Six Pillars of Observability

It's about collecting data of different sources, formats and types

 **Dotan Horovits** · You
Technology Evangelist | CN
1yr · Edited · 🌐

The #OpenTelemetry commur
Profiling. Here's the vision draf
Check out OTEP-212 on GitHul
<https://lnkd.in/d9TRbPmK>


Are you new to Continuous Pro
<http://ow.ly/Z2AZ50KNeOp>

Cloud Native Computing Four

open-telemetry/oteps

#212 Propose OpenTelemet Vision

7 comments 36 reviews

 Rperry2174 · August 30, 2022

 **Dotan Horovits** 🇮🇱 🏳️‍🌈
Technology Evangelist | t
6mo · 🌐


👉 Going beyond logs, metrics an
Continuous Profiling is out! 🚀
Great work to all the OTel Profiling

New issue

Introduces P Model #237


Open

petethepig wants to m
`open-telemetry:main`
`petethepig:profilin`

 **Dotan Horovits** #CNCFAmbassador 🇮🇱 🏳️‍🌈 @horovits · Feb 29

A major milestone for adding Continuous Profiling to #OpenTelemetry:
The extension proposal (OTEP) outlining the data model has been merged 🎉

I need to look into that deeper @OpenObserv with the SIG members -
@felixge @rperry_ joining me?
#opensource #devops #observability

 **open-telemetry / oteps** Public

<> Code Issues 20 Pull requests 16 Actions

New issue

Introduces Profiling Data Model v2 #239

Merged

jsuereth merged 37 commits into `open-telemetry:main` from
`petethepig:profiling-pproftextended` last week

OpenTelemetry and 3 others

2 13 36 5K



Continuous Profiling:
A New Observability
Signal

It's about collecting data of different sources, formats and types

Charting New Territory: OpenTelemetry Embraces Profiling



7 MARCH 2024

11AM PT / 20:00 CET



DOTAN HOROVITS
Principal Developer
Advocate, Logz.io



RYAN PERRY
Principal Product
Manager, Grafana Labs

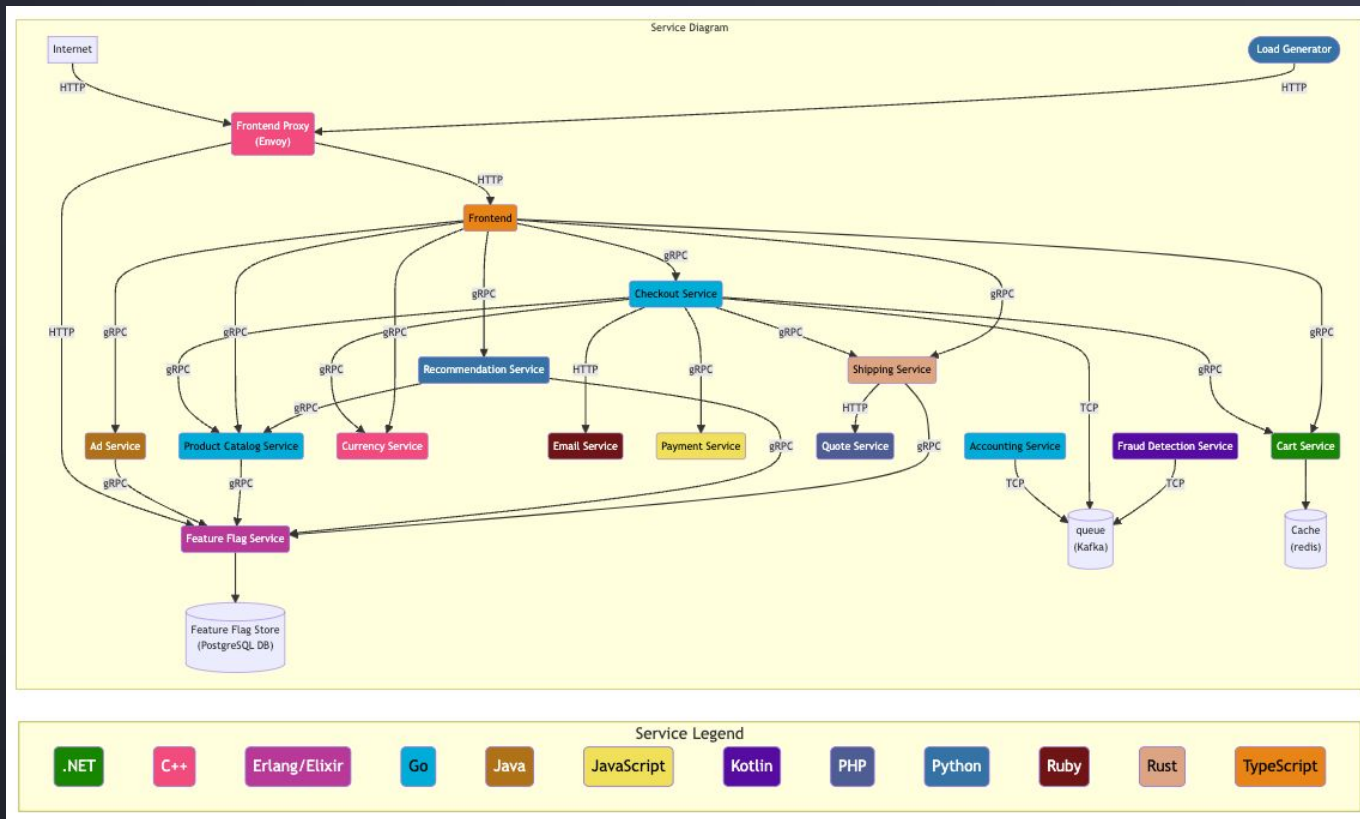


HOSTING

FELIX GEISENDÖRFER
Senior Staff
Engineer, Datadog

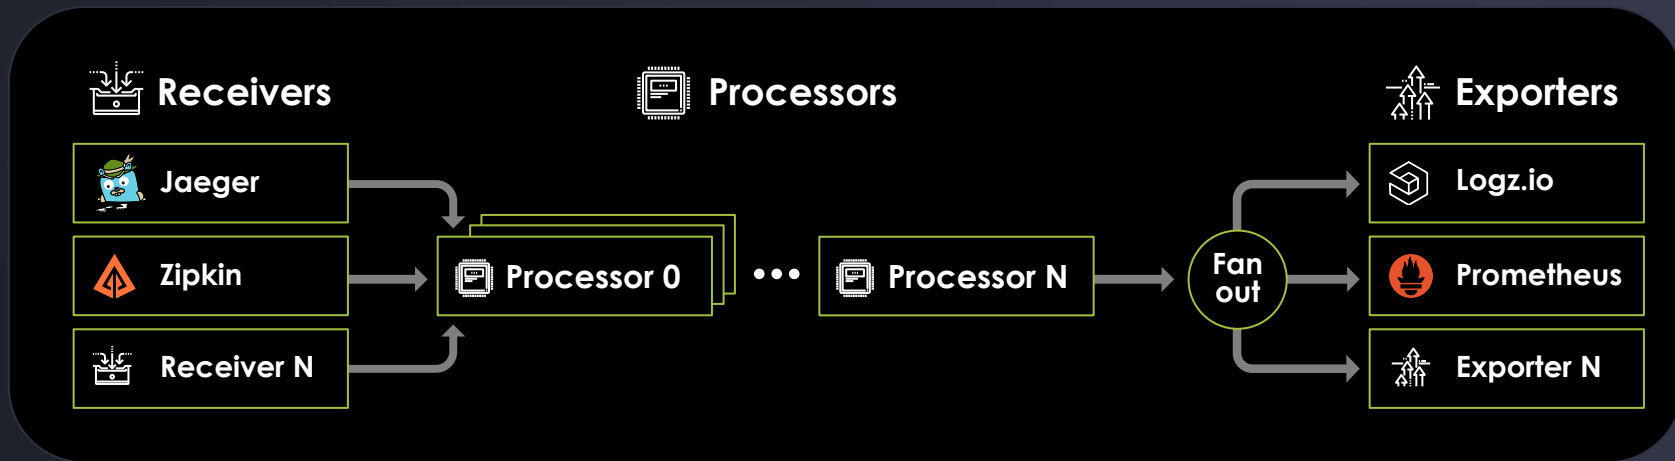


It's about collecting data of different sources, formats and types



It's about collecting data of different sources, formats and types

OpenTelemetry
Guide



Observability is about

Structured standardized data

It's about structured standardized data

```
2023-05-16 11:00:18 ERROR [carts,/] :  
[-nio-80-exec-1] o.a.c.c.C.[.[./].[dispatcherServlet] :  
Servlet.service() for servlet [dispatcherServlet] in context  
with path [/] threw exception [Request processing failed;  
nested exception is  
org.springframework.data.mongodb.UnsupportedMongoCommandException  
: Query failed with error code 352 and error message 'Unsup-  
'Unsupported OP_QUERY command: The client driver may  
require an upgrade. For more details see  
https://www.mongodb.org/core/legacy-opcode-removal on  
server call: s-db:27017; nested exception is  
com.mongodb.MongoQueryException: Query failed with error code  
352 and error message 'Unsupported OP_QUERY command. The  
client driver may require an upgrade. For more details see  
https://www.mongodb.org/core/legacy-opcode-removal' on  
server call: s-db:27017] with root cause
```

It's about structured standardized data

```
{
  "log_level": "INFO",
  "type": "checkoutservice",
  "message": "payment went through (transaction_id: 437d3582be51)",
  "docker": {
    "container_id": "bfc6ddc97255511"
  },
  "env_id": "Astronomy-Shop",
  "timestamp": "2023-08-07T19:17:28.887170752Z"
}
```

It's about structured standardized data



It's about structured standardized data



“Open standards and open specifications have become integral drivers of compatibility, collaboration, and convergence”



Observability is about

Enriching and correlating data

It's about enriching and correlating data

```
{
  "log_level": "INFO",
  "type": "checkoutservice",
  "message": "payment went through (transaction_id: 437d3582be51)",
  "docker": {
    "container_id": "bfc6ddc97255511"
  },
  "env_id": "Astronomy-Shop",
  "timestamp": "2023-08-07T19:17:28.887170752Z",
  "user_id": "dksf70ts",
  "build": "20513097",
}
```



customer-specific?
problematic build?

It's about enriching and correlating data



OPENMETRICS

```
__meta_kubernetes_node_name: The name of the node object.  
__meta_kubernetes_node_provider_id: The cloud provider's name for the  
node object.  
__meta_kubernetes_node_label_<labelname>: Each label from the node  
object.  
__meta_kubernetes_node_labelpresent_<labelname>: true for each label  
from the node object.  
__meta_kubernetes_node_annotation_<annotationname>: Each annotation from  
the node object.  
__meta_kubernetes_node_annotationpresent_<annotationname>: true for each  
annotation from the node object.  
__meta_kubernetes_node_address_<address_type>: The first address for  
each node address type, if it exists.
```

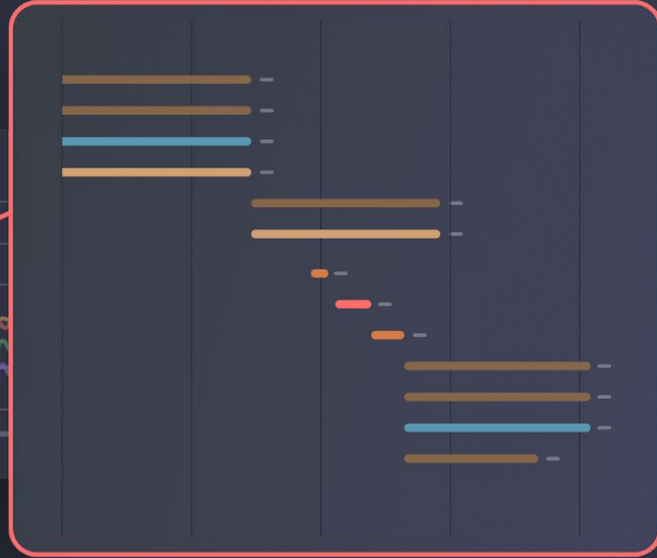
It's about enriching and correlating data



OpenTelemetry

```
k8sattributes:
  filter:
    node_from_env_var: KUBE_NODE_NAME
  extract:
    metadata:
      - k8s.pod.name
      - k8s.pod.uid
      - k8s.deployment.name
      - k8s.namespace.name
      - k8s.node.name
      - k8s.pod.start_time
    labels:
      - tag_name: app.label.component
        key: app.kubernetes.io/component
        from: pod
  pod_association:
    - sources:
      - from: resource_attribute
        name: k8s.pod.ip
    - sources:
      - from: resource_attribute
        name: k8s.pod.uid
    - sources:
      - from: connection
```

It's about enriching and correlating data




Observability is about

Unified query and visualization

It's about unified query and visualization

Logs



Metrics



Tracing



It's about unified query and visualization



OPEN TALKS
OBSERVABILITY



Sponsored by
 logz.io


Streamlining Observability: The Journey Towards Query Language Standardization

The background of the slide features a light blue color with a white bar chart at the bottom. There are several large, stylized 'O' shapes in white and light blue scattered across the background.

Observability is about

Data volume and data:noise ratio

It's about data volume & data:noise ratio



It's about data volume & data:noise ratio

| | | | |
|--------------------------|---------------------------------------|-------------|------------|
| <input type="checkbox"/> | latency_count | 16 labels > | 2067 (5%) |
| <input type="checkbox"/> | kube_pod_status_reason | 19 labels > | 585 (1.4%) |
| <input type="checkbox"/> | kube_pod_status_phase | 19 labels > | 585 (1.4%) |
| <input type="checkbox"/> | container_network_receive_bytes_total | 30 labels > | 537 (1.3%) |

Rollup

Rollups used to aggregate many time-series into few by removing dimension. Choose what dimension to drop and how to aggregate the data.

* Choose account [?](#)
o11y-prod-metrics

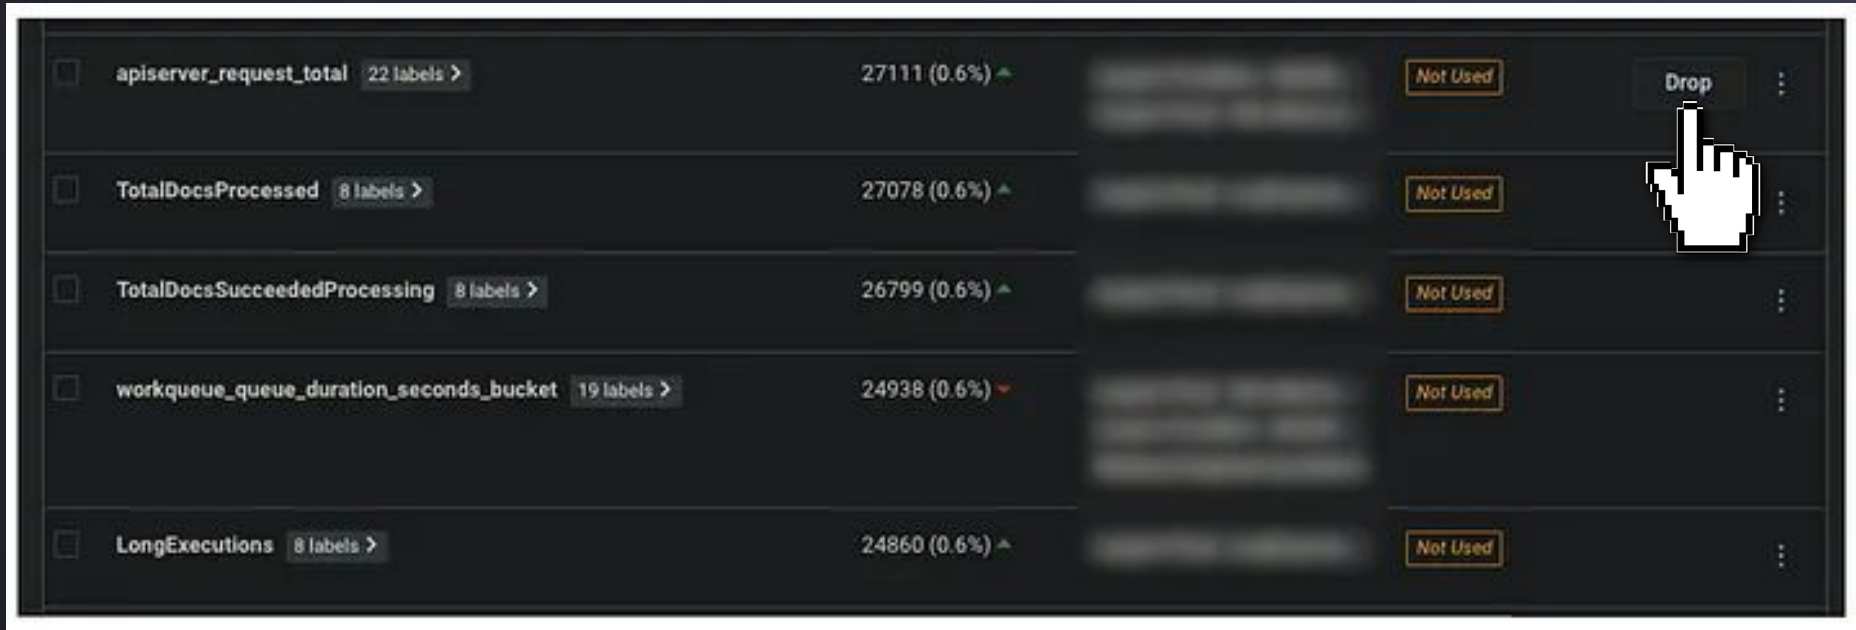
* Choose Metric Name [?](#)
container_network_receive_bytes_total

* Choose dimensions to ignore [?](#)
Select fields to rollup

- beta_kubernetes_io_arch
- beta_kubernetes_io_instance_type
- beta_kubernetes_io_os
- eks_amazonaws_com_capacityType
- eks_amazonaws_com_nodegroup
- eks_amazonaws_com_nodegroup_image
- env_id
- failure_domain_beta_kubernetes_io_region
- failure_domain_beta_kubernetes_io_zone

Source: Logz.io D.

It's about data volume & data:noise ratio



| | | | | |
|--------------------------|---|----------------|----------|------|
| <input type="checkbox"/> | apiserver_request_total 22 labels > | 27111 (0.6%) ▲ | Not Used | Drop |
| <input type="checkbox"/> | TotalDocsProcessed 8 labels > | 27078 (0.6%) ▲ | Not Used | |
| <input type="checkbox"/> | TotalDocsSucceededProcessing 8 labels > | 26799 (0.6%) ▲ | Not Used | |
| <input type="checkbox"/> | workqueue_queue_duration_seconds_bucket 19 labels > | 24938 (0.6%) ▼ | Not Used | |
| <input type="checkbox"/> | LongExecutions 8 labels > | 24860 (0.6%) ▲ | Not Used | |

Source: Logz.io Data Optimization Hub

It's about data volume & data:noise ratio

[Blog](#) / [Best Practices](#)

#Infrastructure Monitoring

#Monitoring

#Observability


Expensive Metrics: Why Your Monitoring Data and Bill Get Out Of Hand

By: Dotan Horovits



Summary

Logs



Metrics



Tracing




Observability is about

Data Analysis


Observability is a Data Analytics Problem

Observability is the capability to allow a human to ask and answer questions about the system.



Observability is a Data Analytics Problem

- ... collecting data of different sources, formats and types
- ... structured standardized data
- ... enriching and correlating data
- ... unified query and visualization
- ... data volume & data:noise ratio



Thank you
&
May the
Open Source
be with you

Dotan Horovits
[@horovits](#)

