

# CRIBLCON<sup>24</sup>

POWERED BY  Cribl

## Correlating Traces with Logs for Enhanced Observability

**BHOOPESH**

Senior Software Engineer, Observability, Autodesk





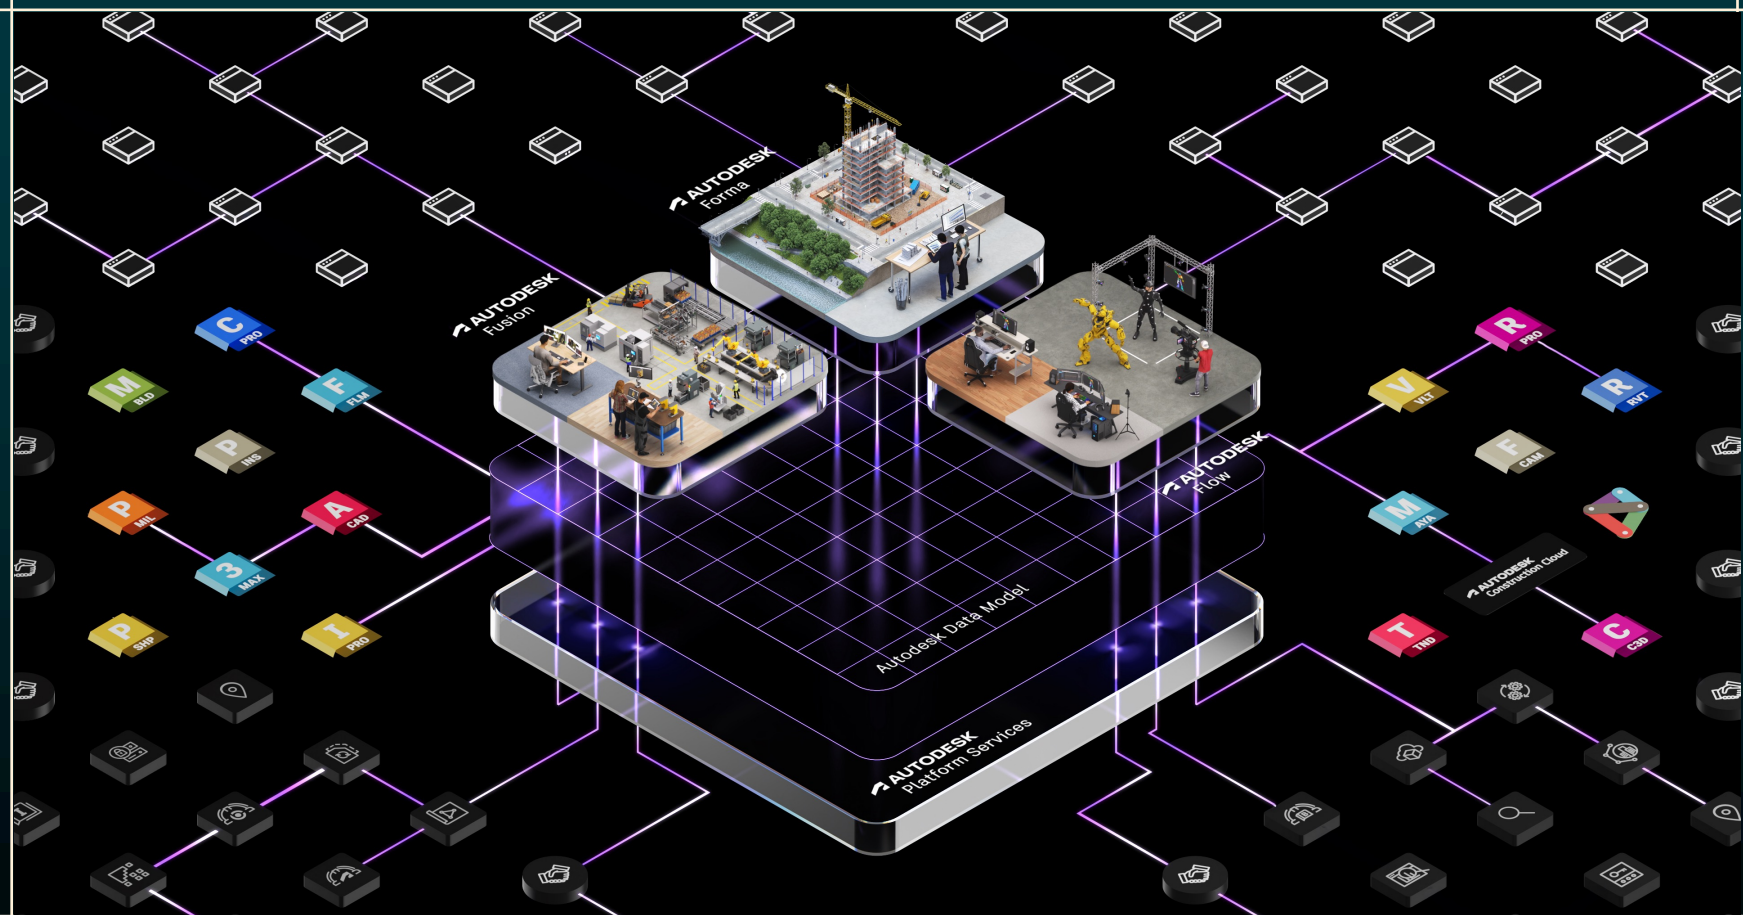
## **BHOOPESH**

Senior Software Engineer,  
Observability, Autodesk

# What does Autodesk Do?

Autodesk is changing how the world is designed and made. Our technology spans architecture, engineering, construction, product design, manufacturing, media, and entertainment, empowering innovators everywhere to solve challenges big and small. From greener buildings to smarter products to more mesmerizing blockbusters, Autodesk software helps our customers to design and make a better world for all.





# Agenda

Here's what to expect:

1

## INTRODUCTION

Setting the stage for enhanced observability in distributed apps

2

## THE CHALLENGE

Identifying issues with isolated observability data and siloed tools

3

## INTEGRATION

Learn the benefits of combining OTel traces with Splunk logs

4

## OUTCOME & HOW

Learn more about the impact and how to implement solution

5

## USE CASE

See how to enhance observability in an e-commerce application

6


## BEST PRACTICES

Sharing lessons learned and best practices for effective observability

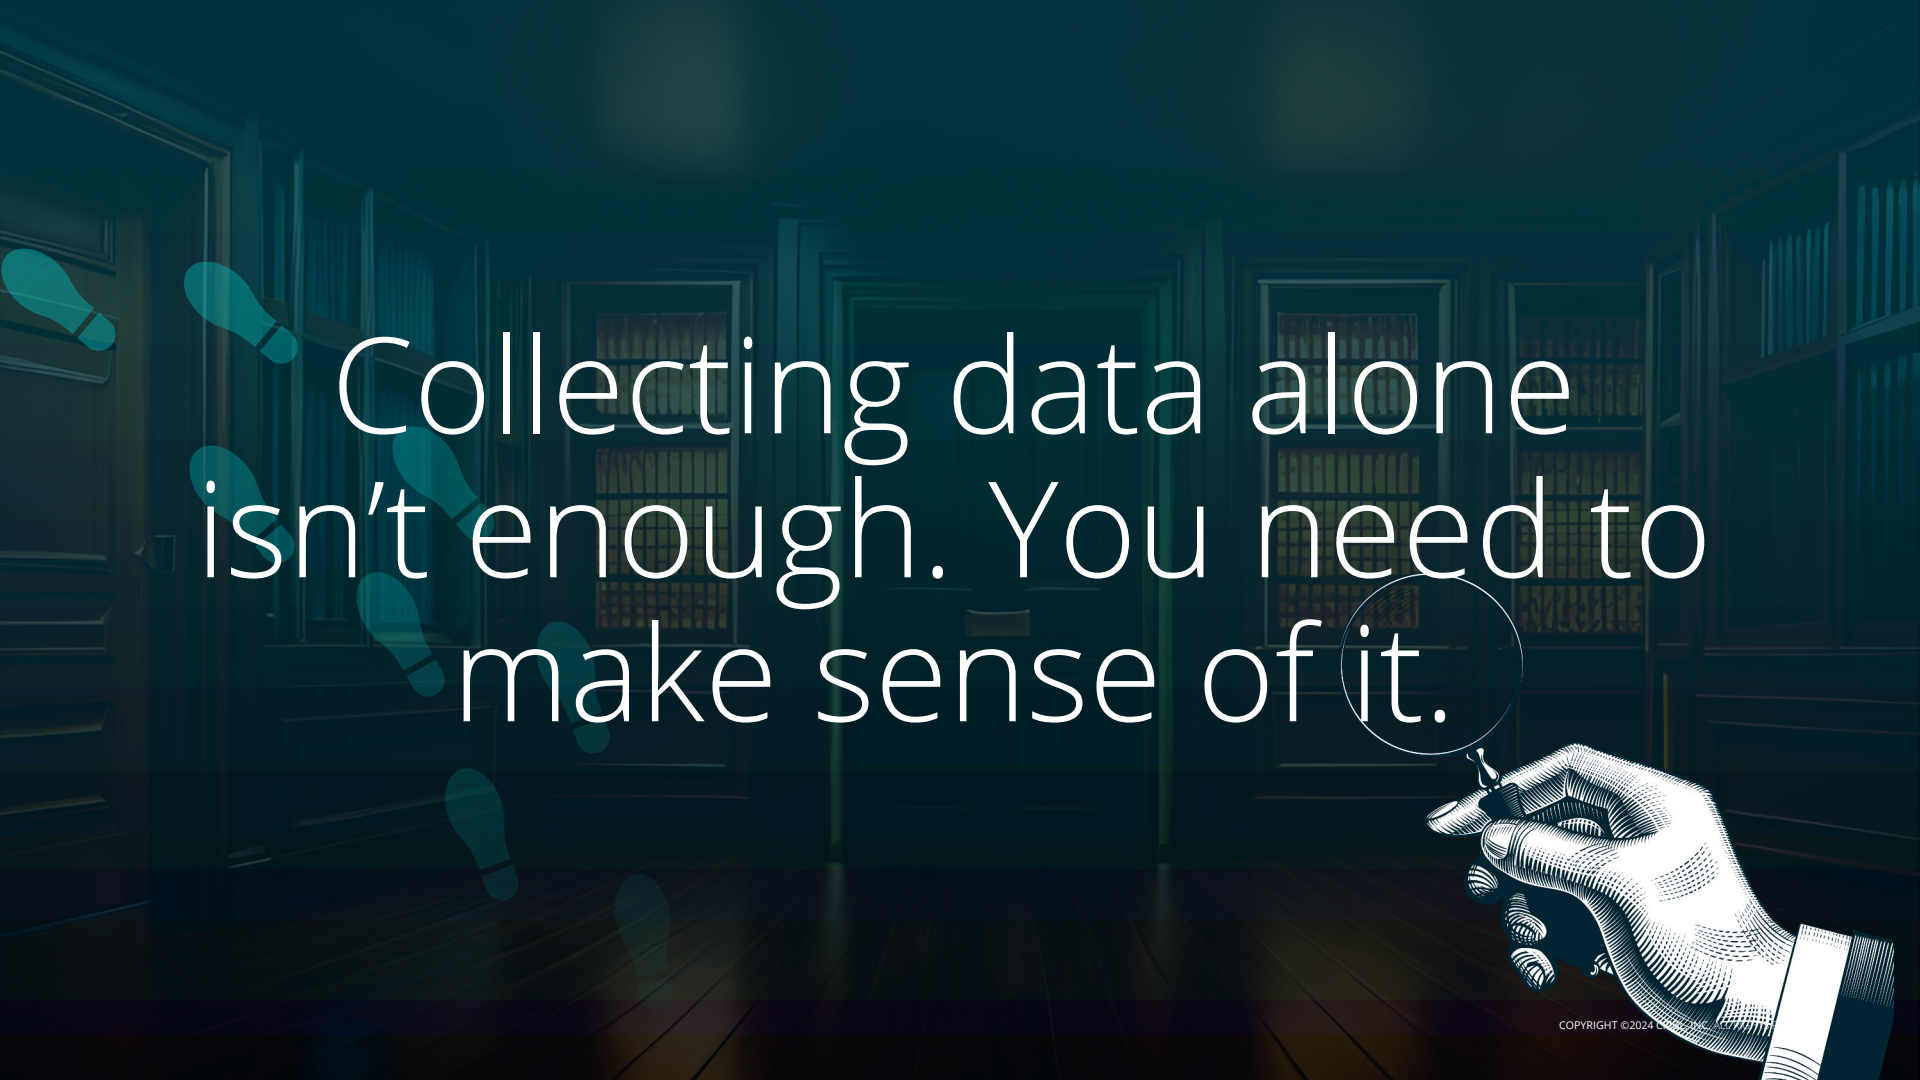
7

## DASHBOARDS

Building a one-stop shop for enhanced system insights







Collecting data alone  
isn't enough. You need to  
make sense of it.

# The Challenge



# The Challenge

- 1** Unifying diverse telemetry data requires tracking unique request identifiers across a distributed system
- 2** Switch between different tools and lose the valuable context pertinent to fixing performance issues
- 3** Manual mapping of service instrumentation to log identifiers is complex and time-consuming
- 4** Maintaining and expanding this system becomes increasingly difficult as new services are built and new languages, frameworks, and libraries are adopted, often resulting in visibility of only a subset of logs for any particular request



# The Challenge

Siloed observability data slows down effective troubleshooting and performance optimization



**Data silos**



**Incomplete  
insights**



**Extended  
downtime**



**Scalability  
issues**



# Impact of Challenges

Reinforcing the need for integrated observability



## **Increased mean time to respond (MTTR)**

Higher resolution times due to fragmented data




## **Higher operational costs**

More resources needed for manual correlation



## **Reduced system reliability**

Impact on uptime and user experience



“We can’t have all  
these things working  
on a separate basis.  
This problem has to  
be addressable.”



# The Outcome and how





# A unified view of system behavior

Bridge silos for a seamless flow of info

## APPLICATION TEAM

## OBSERVABILITY TEAM





# Integrating Traces and Logs



## Quick root cause identification

Faster root cause analysis by correlating events across services



## Service dependency mapping

Visualize how different services interact



## Metrics tracking

Monitor and analyze latency and Endpoint performance across various parts of the system

# Solution

## Collect

Enable  
OTel  
Source



OpenTelemetry

Collect Trace  
data from OTEL  
Collectors

```
exporters:
  debug:
    verbosity: detailed
  otlp/cribl:
    #endpoint: ${CRIBL_GRPC_ENDPOINT}
    endpoint:
    tls:
      insecure_skip_verify: true
    keepalive:
      time: 10s
    retry_on_failure:
      max_elapsed_time: 10s
    sending_queue:
      queue_size: 10000
    compression: gzip
```

Sources > OpenTelemetry  
otel-input-1


Configure Status Charts **Live Data** Logs Notifications

Filter Expression\*

\_\_inputId=="open\_telemetry:otel-input-1"

Fields All None

#	Event
1	<div><div>#_time: 1718540966.167</div><div>#_index: ese_test_jb</div><div>2024-05-24 16:06:06.167</div><div>1606:06.167</div><div>405:30</div><div><div>instrumentation_library_spans:</div><div><div>#:</div><div><div>instrumentation_library:</div><div><div>name: opentelemetry/instrumentation-http</div><div>version: 0.41.2</div><div>schema_url:</div><div>spans: 7 items...</div><div>3 items...</div><div>3 items...</div></div></div></div></div><div>resource:</div><div><div>attributes:</div><div><div>service.name: ace-hub-sync-service</div><div>service.version: 0.0.1</div><div>telemetry.sdk.language: nodejs</div><div>telemetry.sdk.name: opentelemetry</div><div>telemetry.sdk.version: 1.15.2</div><div>dropped_attributes_count: 0</div><div>schema_url:</div></div></div></div>



# Solution

## Reduce

```
# Event
1 X # _time: 1715009299.026
2024-05-06 # index: ese_test_jb
20:58:19.026 # instrumentation_library_spans:
+05:30 # 0:
# instrumentation_library:
# name: Mojo
# version: unknown
# schema_url:
# spans: 100 items...
# resource:
# attributes:
# deployment.environment: prd
# process.command: /home/lem/bundle/jruby/3.1.0/bin/puma
# process.pid: 1
# process.runtime.description: jruby 9.4.6.0 (3.1.4) 2024-02-20 576fab2c51 OpenJDK 64-Bit Server VM 25.402-b08 on 1.8.0_402-b08 +jit [x86_64-linux]
# process.runtime.name: jruby
# process.runtime.version: 3.1.4
# service.name: lem
# service.version: lem:1316
# telemetry.sdk.language: ruby
# telemetry.sdk.name: opentelemetry
# telemetry.sdk.version: 1.2.1
# dropped_attributes_count: 0
# schema_url:
```

```
# spans:
# 0:
# attributes:
# component: http
# deployment.environment: prd
# http.client_ip: 10.40.226.123
# http.method: GET
# http.route: /service/entitlements/v2/users/7PUCM4N95JH3
# http.status_code: 200
# http.url: /service/entitlements/v2/users/7PUCM4N95JH3?includeExpired=false&includeParentAsset=true&source=legacy
# router.params.includeExpired: false
# router.params.includeParentAsset: true
# router.params.reference_id: 7PUCM4N95JH3
# router.params.source: legacy
# dropped_attributes_count: 0
# dropped_events_count: 0
# dropped_links_count: 0
# end_time_unix_nano: 1715009296381282300
# events:
# kind: 2
# links: 0 items...
# name: /service/entitlements/v2/users/:reference_id
# parent_span_id:
# span_id: a0690d522be14af4
# start_time_unix_nano: 1715009296378018000
# status: 2 items...
# trace_id: afc2a962c77043ec432b1a6a450ad1ca
# trace_state:
# 15 items...
# 15 items...
# 15 items...
# 15 items...
# 15 items...
# 15 items...
# 15 items...
# 15 items...
```



# Solution

## Shape

```
1 2024-05-24 19:53:00.590 +05:30
{
  "@_raw": {
    "# process.pid: 1
    "service_name": "lem
    "spans": {
      "attributes": {
        "component": "http
        "http.client_ip": "10.40.226.123
        "http.method": "GET
        "http.route": "/service/entitlements/v2/users/7PUCM4N9S3H3
        "# http.status_code: 200
        "http.url": "/service/entitlements/v2/users/7PUCM4N9S3H3?includeExpired=false&includeParentAsset=true&source=legacy
        "router.params.includeExpired": false
        "router.params.includeParentAsset": true
        "router.params.reference_id": "7PUCM4N9S3H3
        "router.params.source": "legacy
      }
      "# kind: 2
      "name": "/service/entitlements/v2/users/:reference_id
      "parent_span_id":
      "span_id": "a0690d522be14af4
      "status": {
        "# code: 0
        "message":
        "trace_id": "afc2a962c77043ec432b1a6a450ad1ca
        "trace_state":
      }
      "# tms: 3
    }
    "# _time: 1716560580.59
    "cribl_pipe": "pipeline-splunk-ttl_master-v1
    "host_ip": "10.226.2.2"
    "index": "ese_test_jb
    "instrumentation_library_spans": 1-item...
    "# process.pid: 1
    "resource": 2-items...
    "schema_url":
    "source": "open_telemetry:otel-input-1
    "sourcetype": "ttl
  }
}
```



# How

## Shed the Dead Weight

Unroll

true

Filter ⓘ

true

Description ⓘ

Enter a description

Final ⓘ

No

Source Field Expression\* ⓘ

instrumentation\_library\_spans

Destination Field\* ⓘ

instrumentation\_library\_spans

Unroll

true

Filter ⓘ

true

Description ⓘ

Enter a description

Final ⓘ

No

Source Field Expression\* ⓘ

instrumentation\_library\_spans.spans

Destination Field\* ⓘ

instrumentation\_library\_spans.spans

Parser

true

Parser

true

Eval

true

Auto Timestamp

true

Parser

true

Eval

true

add http fields below

Parser

true

spans drop

Eval

true

Rename

true

Serialize

true

final parsed fields removal

Eval

true

FIN





# Injection of trace metadata


## Trace Context Propagation

```
from opentelemetry import trace
    You, 1 second ago • Uncommitted changes
def _log(self, level, msg, *a, **kw):
    kw["extra"] = self._data(kw.get("extra"))
    ctx = trace.get_current_span().get_span_context()
    kw["extra"]["otelTraceId"] = ctx.trace_id
    kw["extra"]["otelSpanId"] = ctx.span_id
    self._logger.log(level, msg, *a, **kw)
```

```
from opentelemetry import trace

def process(self, msg, kwargs):
    kwargs["extra"] = {"data": {k: v for k, v in self._gdata(kwargs)}}
    if self.extra:
        kwargs["extra"]["data"].update(self.extra)
    ctx = trace.get_current_span().get_span_context()
    kwargs["extra"]["otelTraceId"] = ctx.trace_id
    kwargs["extra"]["otelSpanId"] = ctx.span_id    You, 3 weeks ago •
    return msg, kwargs
```


```
class ContextFilter(logging.Filter):
    def filter(self, record):
        record.id = str(uuid4())
        record.utctime = _timestamp(DEFAULT_DATE_FORMAT)
        record.source = "urn:adsk:sfdc:moniker:SFDC-CPQ"
        record.message = record.getMessage()
        ctx = trace.get_current_span().get_span_context()
        record.otelTraceId = ctx.trace_id
        record.otelSpanId = ctx.span_id
        if hasattr(record, "data"):
            record.datacontenttype = "application/json"
        return 1
```



# Correlation begins here

## Route

- Route the traces to corresponding Index/source/sourcetype using eval
- Enrich the Log data with trace metadata i.e. traceIds and spanIds and environment metadata i.e. service.name and service.environment
- This will allow you to connect the dots between a specific operation (traces) and what was happening in your system at that time (logs)
  - To reconstruct the code path taken by reading a trace
  - To derive request or error ratios from any single point in the code path



# Demo



A stylized illustration of a server room. On the left, a door is labeled 'DOWN' in red. A hand in a white suit sleeve holds a magnifying glass over the door. To the right of the door are several server racks filled with equipment. The floor has a checkered pattern.

# Real-world use case

## E-commerce application

- Initial challenges with multiple services and data sources
- Integration of OTel traces with Splunk logs using Cribl Stream

### Results:

- Improved issue resolution
- Enhanced performance monitoring
- Increased system reliability

# Lessons learned / Best practices





# Best practices

How to achieve integrated observability with Cribl



Leverage Cribl  
Stream for routing



Use reduction  
techniques to  
minimize noise



Don't forget to  
parse and enrich

# Visualization Best Practices...

...for better data readability and operational decision-making

## Custom dashboard creation

- Create custom views tailored to specific operational needs or teams
- Useful for environments where conditions change rapidly and up-to-date information is crucial

## Real-time monitoring and alerts

- Set up dashboards that reflect live data flows
- Trigger alerts based on predefined thresholds

# Visualization Best Practices...

...for better data readability and operational decision-making

## Correlating data across sources

- Pull data from multiple sources into a single dashboard
- Combine logs, metrics, and traces in one view to quickly pinpoint issues across the infrastructure

## Sharing and collaboration

- Share dashboards across teams to so all stakeholders can access relevant insights
- Trigger alerts based on predefined thresholds



Thank you!