

Distributed Tracing

Fighting Bugs with
Confidence

30.04.2025 Freiburg Devsmeetup, Maksym Prokopov



Agenda

- Intro
- The problem
- How to improve
- Tools Overview 
- 3 Demos 
- Summary

Me

Maksym Prokopov - SRE

- Many years in the software development
- Work at **Billie** - Buy Now Pay Later platform from Berlin
- Love "fiddling" with technology and building pet projects



The problem

How do we fight a bug?



Sentry APP 3:48 PM

replied to [a thread](#)

 [App\Infrastructure\InvoiceButler\InvoiceNotFound](#)

/src/Infrastructure/InvoiceButler/InvoiceButlerClient.php in

App\Infrastructure\InvoiceButler\InvoiceButlerClient::App\Infrastructure\InvoiceButler\{closure}

```
App\Infrastructure\InvoiceButler\InvoiceButlerClient::App\Infrastructure\InvoiceButler\{closure}
```

Events: 4542 State: Ongoing First Seen: 2023-12-01

Resolve

Archive

Select Assignee...



Suspect Commit: [31dbe2](#) by Tomi Atanasoski 2023-06-12

'Om 1900 invoice external code validation is broken (#1875)' [View Pull Request](#)

Project: [paella](#) Alert: [Paella prod exceptions slack](#) Short ID: PAELLA-60D

ELK

Checking logs...

DQL Last 3 days Show dates Refresh

+ Add filter

production-docker-*

Search field names

Filter by type 0

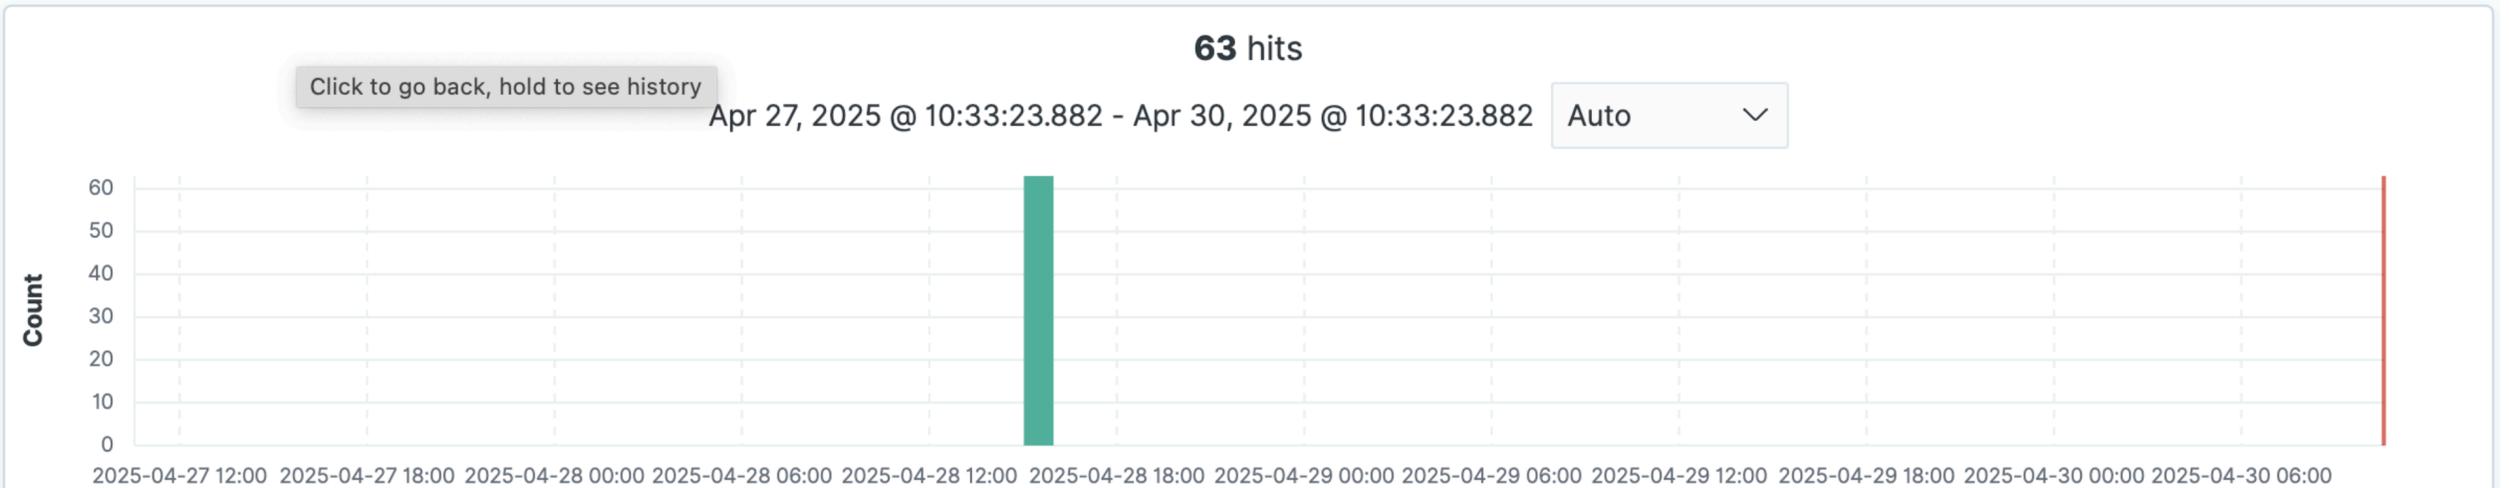
Selected fields

- _source

Available fields

Popular

- @message
- data.context.big_sobaka
- data.context.rid
- data.context.route_name
- data.context.url
- data.level
- data.level_name
- data.message
- service_name
- _id



Time	_source
> Apr 28, 2025 @ 15:56:26.990	<code>data.context.rid: 86ac30e5-4ff5-4122-92be-5c2e42aad04 message: {"message":"Dispatching domain event: App\\Domain\\Order\\Events\\OrderPlaced", "context":{"rid":"86ac30e5-4ff5-4122-92be-5c2e42aad04", "strict_logging_processed":1}, "level":200, "level_name":"INFO", "channel":"app", "datetime":{"date":"2025-04-28 15:56:26.989491", "timezone_type":3, "timezone":"Europe/Berlin"}, "extra": [], "entity.name":"buyerportal", "entity.type":"SERVICE", "entity.guid":"MTYxMjQ2MnxBUE18QVBQTE1DQVRJT058MTYyMjMjY4Nw", "hostname":"bu</code>
> Apr 28, 2025 @ 15:56:26.989	<code>data.context.rid: 86ac30e5-4ff5-4122-92be-5c2e42aad04 message: {"message":"Guzzle service organisation success", "context":{"name":"organisation", "status_code":200, "big_sobaka":{"body":{"uuid":"e18507d1-be9c-4ce0-9992-</code>

More logs Diving deeper...

📄 "86ac30e5-4ff5-4122-92be-5c2e42aad04" 📅 Last 3 days Show dates Refresh

+ Add filter

production-docker-*

🔍 Search field names

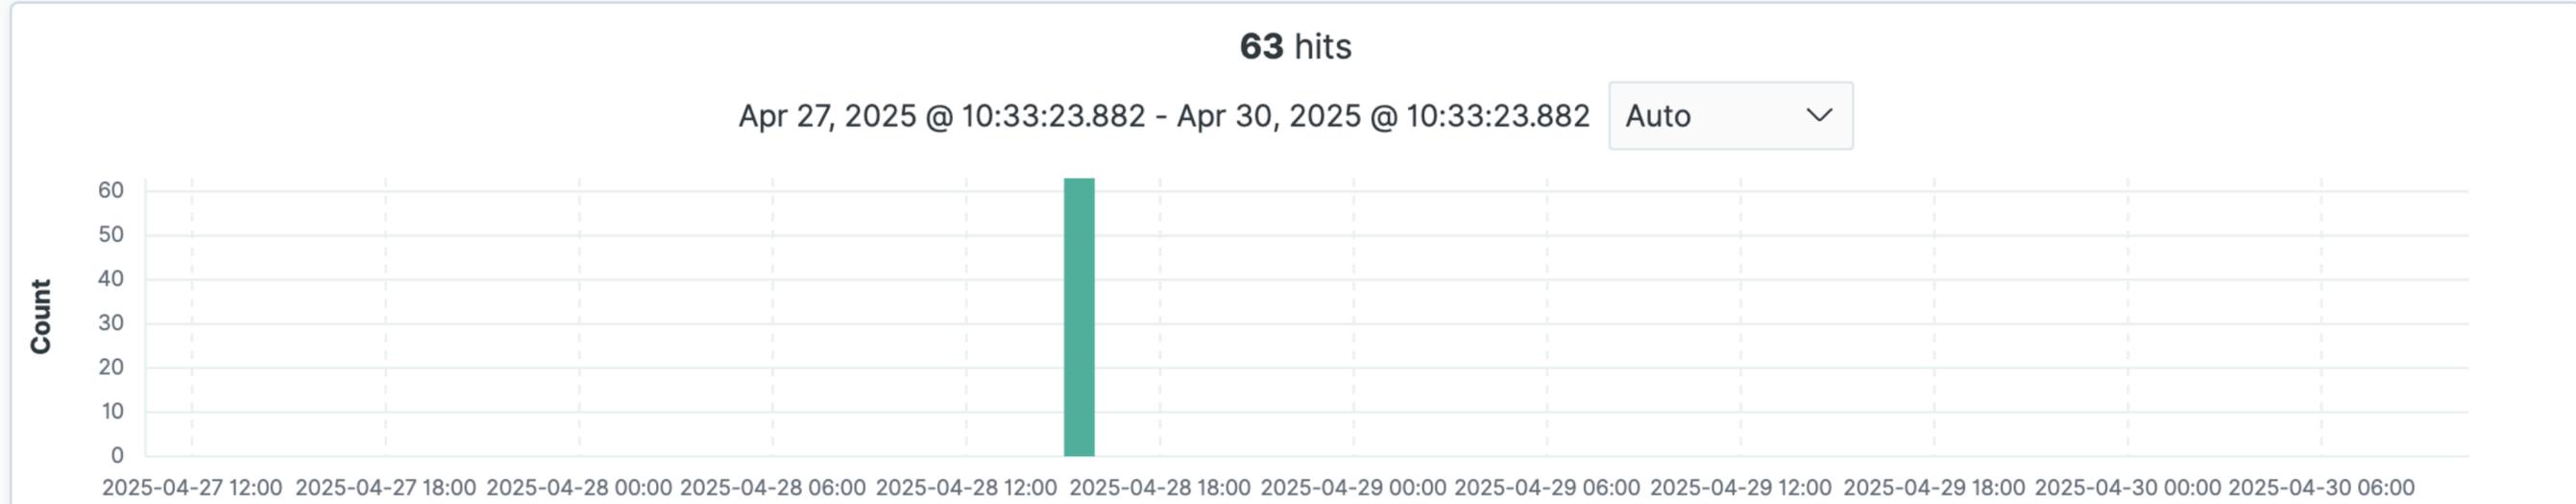
Filter by type 0

Selected fields

- data.message
- service_name

Available fields

- Popular
- @message
 - data.context.big_sobaka
 - data.context.rid
 - data.context.route_name
 - data.context.status_code
 - data.context.url
 - data.level
 - data.level_name
 - instance_name
- _id
 - _index
 - _score
 - _type
 - @id



Time	service_name	data.message
> Apr 28, 2025 @ 15:56:26.990	buyerportal	Dispatching domain event: App\Domain\Order\Events\OrderPlaced
> Apr 28, 2025 @ 15:56:26.989	buyerportal	Guzzle service organisation success
> Apr 28, 2025 @ 15:56:26.971	buyerportal	Guzzle service organisation outgoing request
> Apr 28, 2025 @ 15:56:26.957	buyerportal	Request to http://env-production.haproxy-application.service.production.billie.internal:8004/private/compact/orders/cdba2f52-e126-494d-9495-1b2798326823 took 13165ms
> Apr 28, 2025 @ 15:56:26.957	buyerportal	Guzzle service paella success
> Apr 28, 2025 @ 15:56:26.955	paella-core	Response for private_oa_compact_order_get_details sent
> Apr 28, 2025 @ 15:56:26.953	paella-core	Request to get_invoices took 59ms
> Apr 28, 2025 @ 15:56:26.953	paella-core	Guzzle service invoice_butler success
> Apr 28, 2025 @ 15:56:26.952	invoice-butler	Response for get invoices sent

Demo

Current debugging flow

- Find Request ID
- Filter Logs
- 

The problem?

- Too many logs to filter
- It takes time
- Switching between different systems

- Some requests were slow?
- And when?
- And, more importantly, why?



How to Improve?

Imagine building own solution

- A database with RequestId as an index
 - Let's call it **TraceID**
- Record Every Component separately
 - Component can be any part of your system
 - Let's call it a "span" and assign unique Id as **SpanID**
 - Save start and end time
 - Save parentId
 - Add metadata as **attributes**



Components

- Backend
 - Write/Query
 - Jaeger, Grafana Tempo, commercial
- Client
 - **Instrumentation**
 - **Collector** may help to buffer
 - Emerging Standard - OTEL (Open Telemetry)

Tools Overview

Cloud Only

- NewRelic
- DataDog
- HoneyComb
- and many more...

Tools Overview

Cloud and self-hosted

- Grafana Stack with Grafana Tempo
- Signoz OSS

Demo

Grafana Tempo

Demo

NewRelic

Summary

- We can have better experience!
- Explore solutions supporting OpenTelemetry
- Debugging is much better with distributed tracing

Follow Me

- Blog <https://prokopov.me>
- LinkedIn: <https://www.linkedin.com/in/max-prokopov/>
- Twitter/X: @mprokopov