

99.99% of Your Traces Are (Probably) Trash

Paige Cruz, Chronosphere



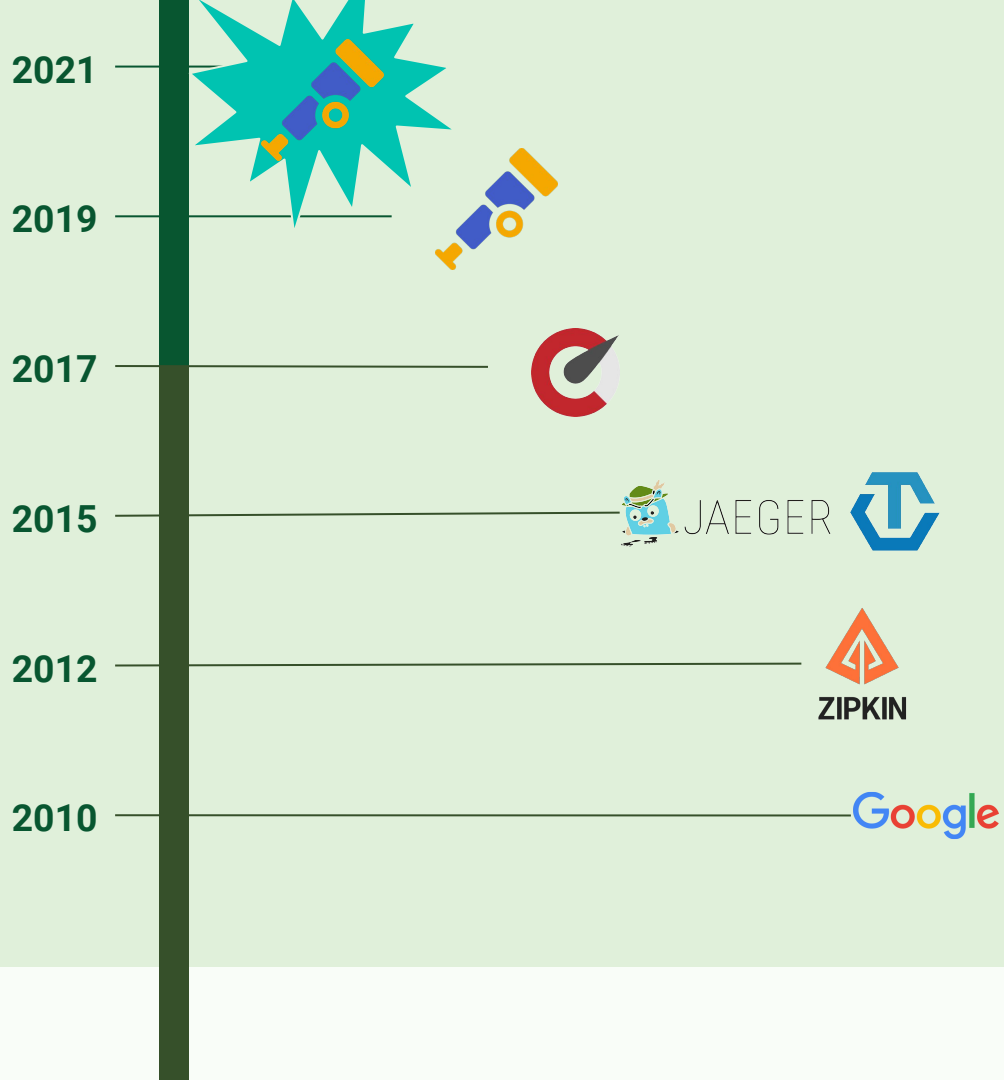
01



Tracing Today



Brief History of Tracing



Tracing circa 2012-2019



No Standard

Tracing circa 2012-2019



Immature

Tracing circa 2012-2019



Steep learning curve

Tracing circa 2019-



No Standard

Tracing circa 2019-



No Standard



Immature

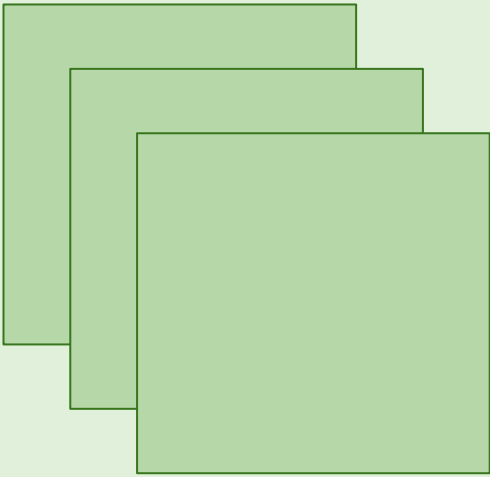
Tracing circa 2019-



Steep learning curve



The True Cost of Tracing




Applications

- Generate spans
- Send spans



Collector*

- Ingest spans
- Process spans
- Export spans



Storage

Tracing Cost Factors




- Throughput
- # Traced Apps
- Metadata



In practice, sampling rates can be as low as

0.01%



Dapper paper from Google

02



Sampling





**What traces
should we
keep?**



**What traces
should we
toss?**



**What makes an
“interesting”
trace?**

What constitutes an “interesting” trace?

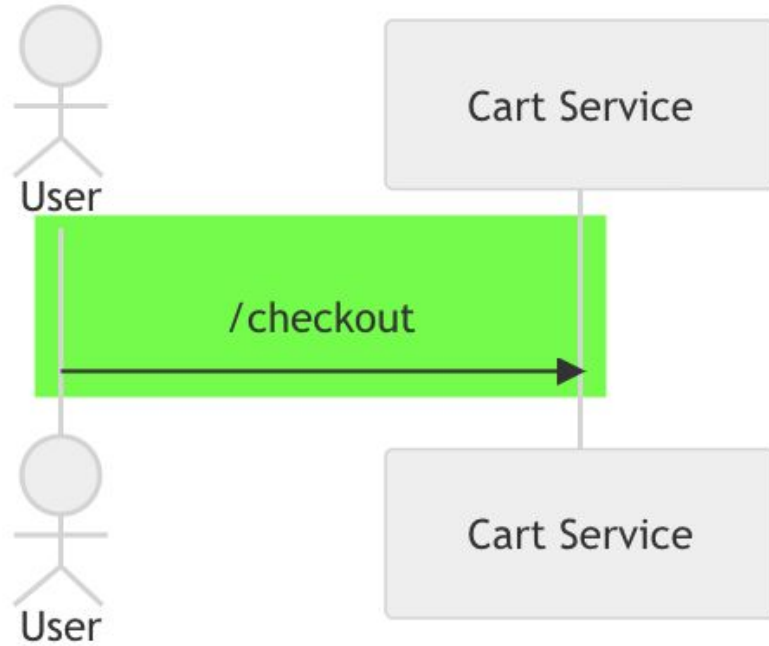
- Infrequent Request Types
- Anomalous Events
- Edge Case
 - Error
 - Tail latency



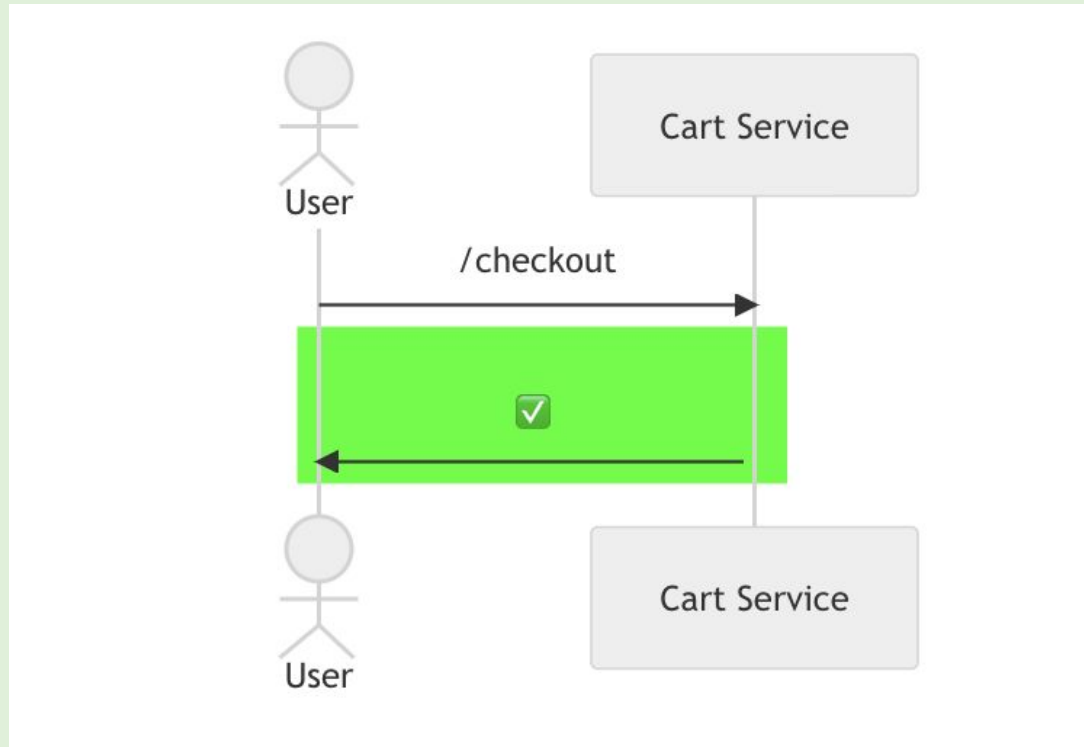
i MUST get the
latest drop from
Purrada



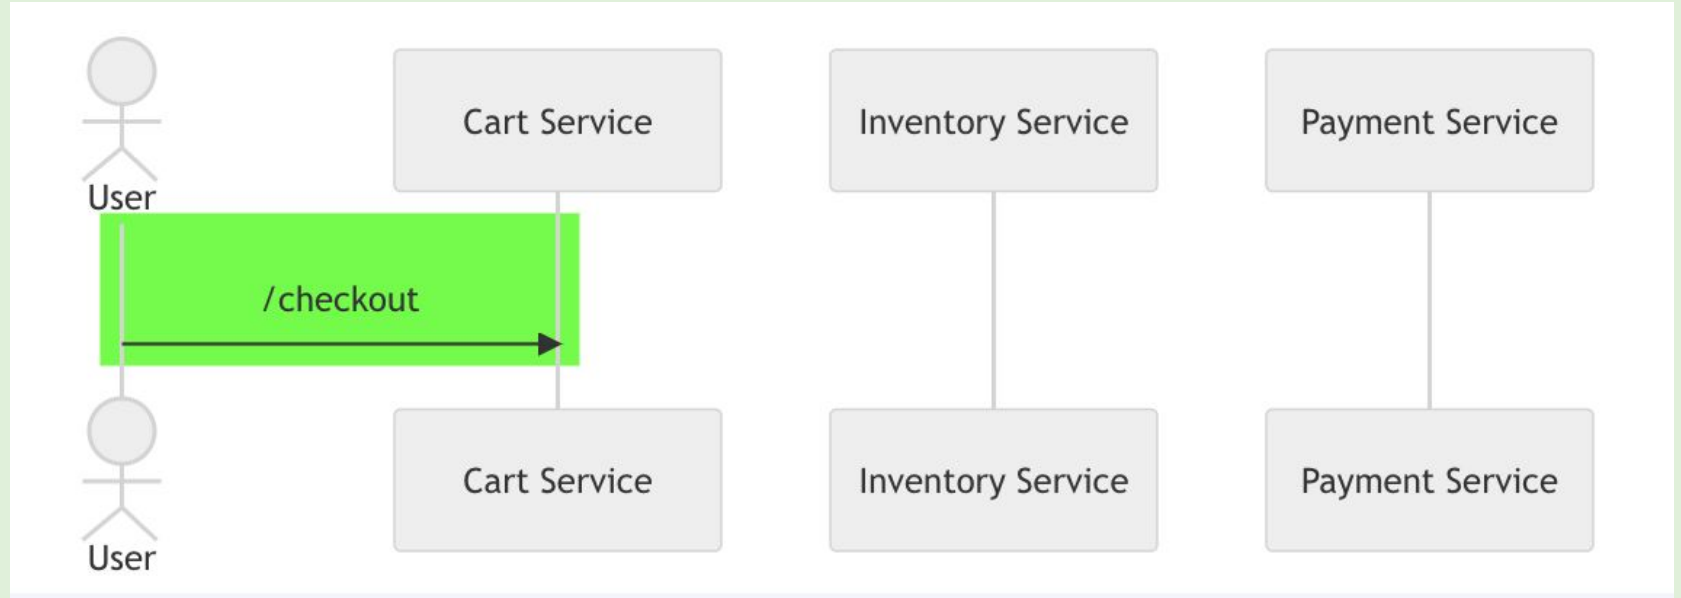
User



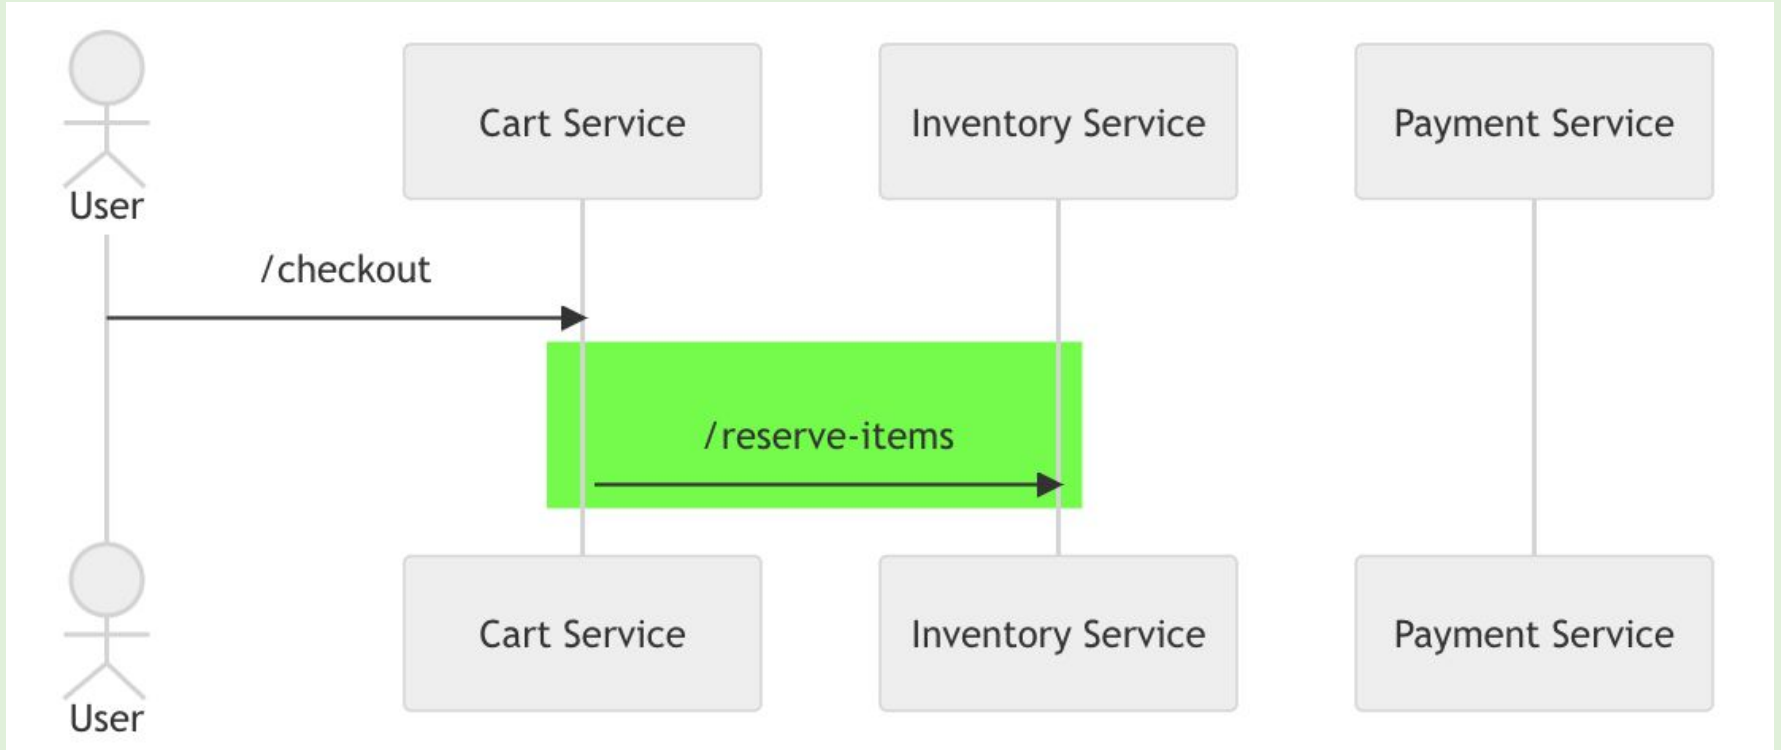
User



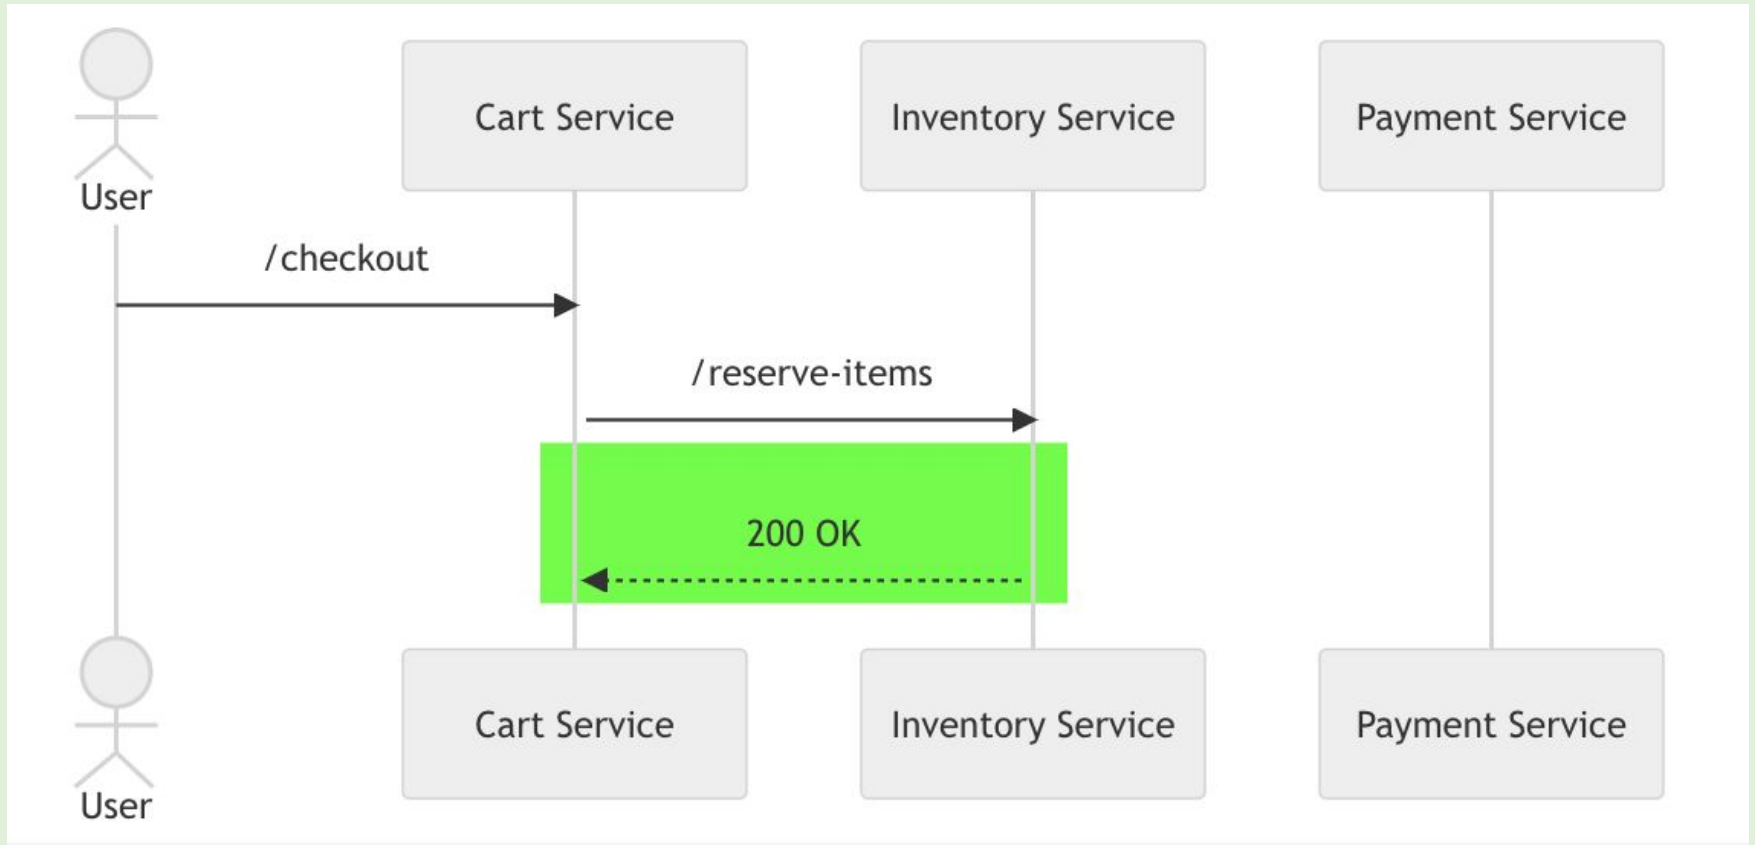
System



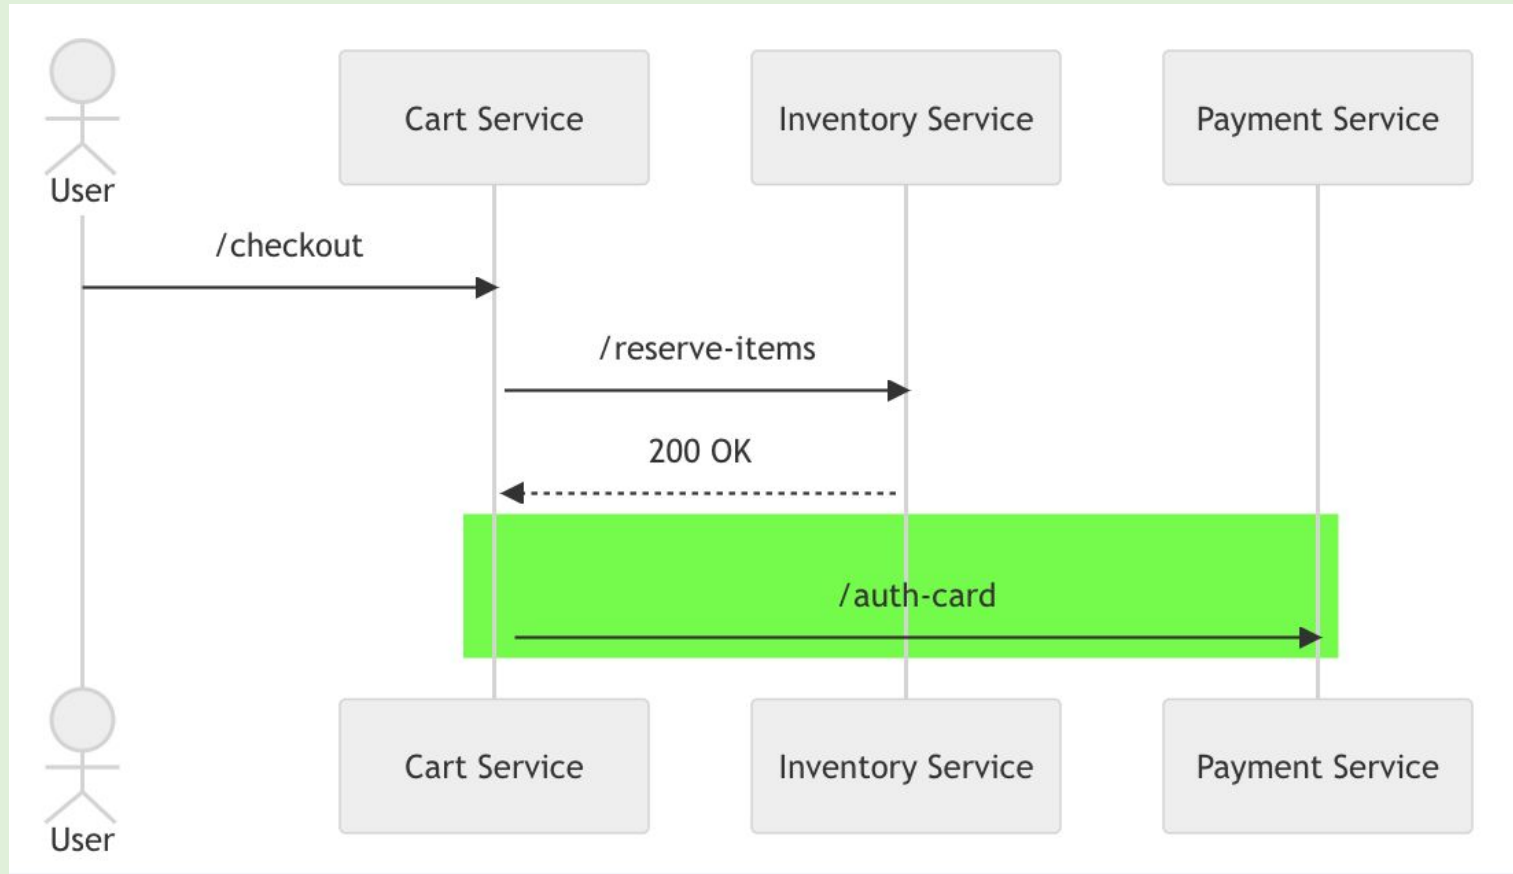
System



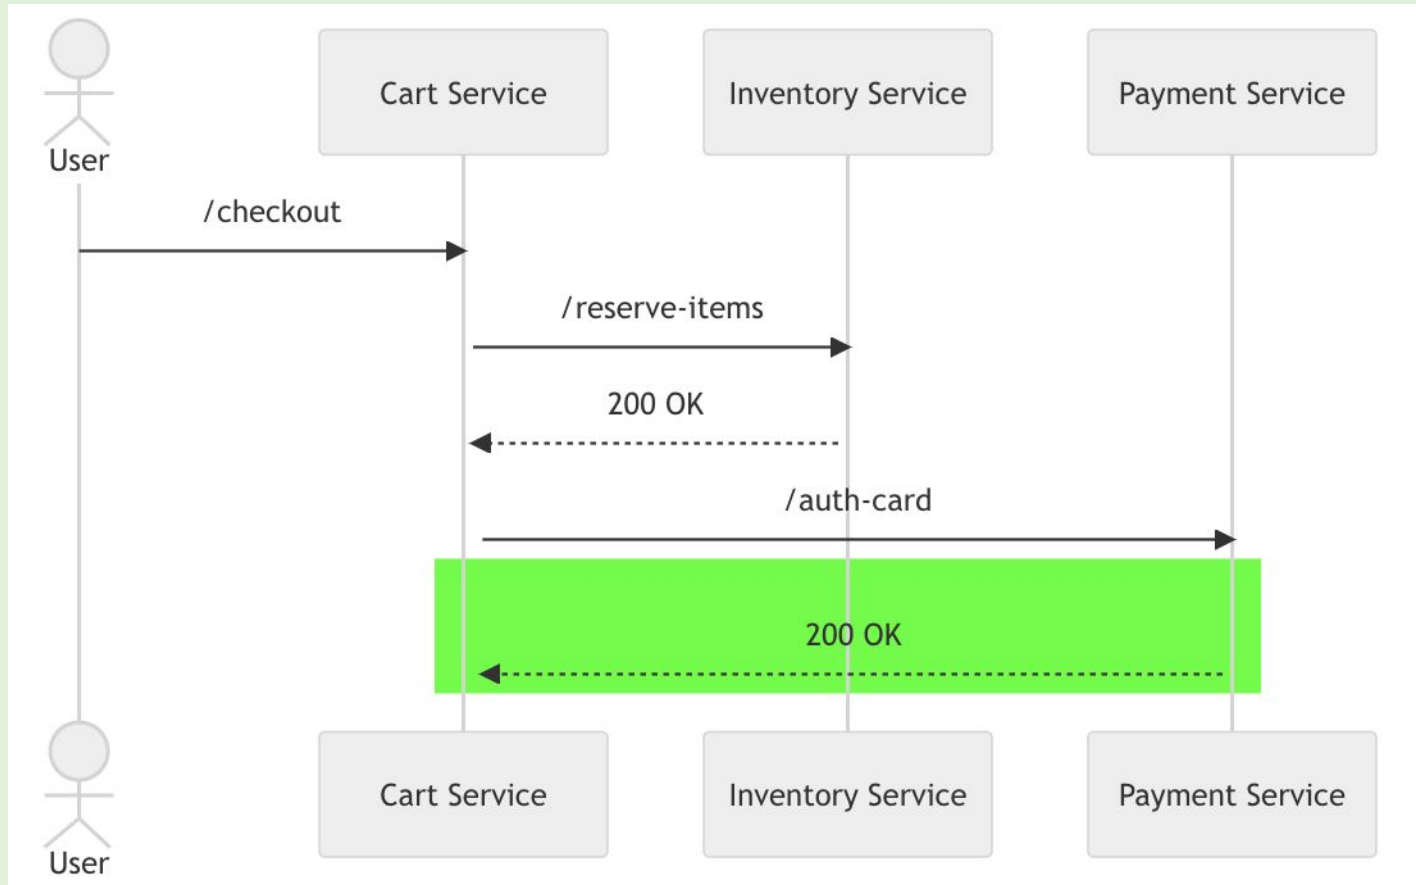
System



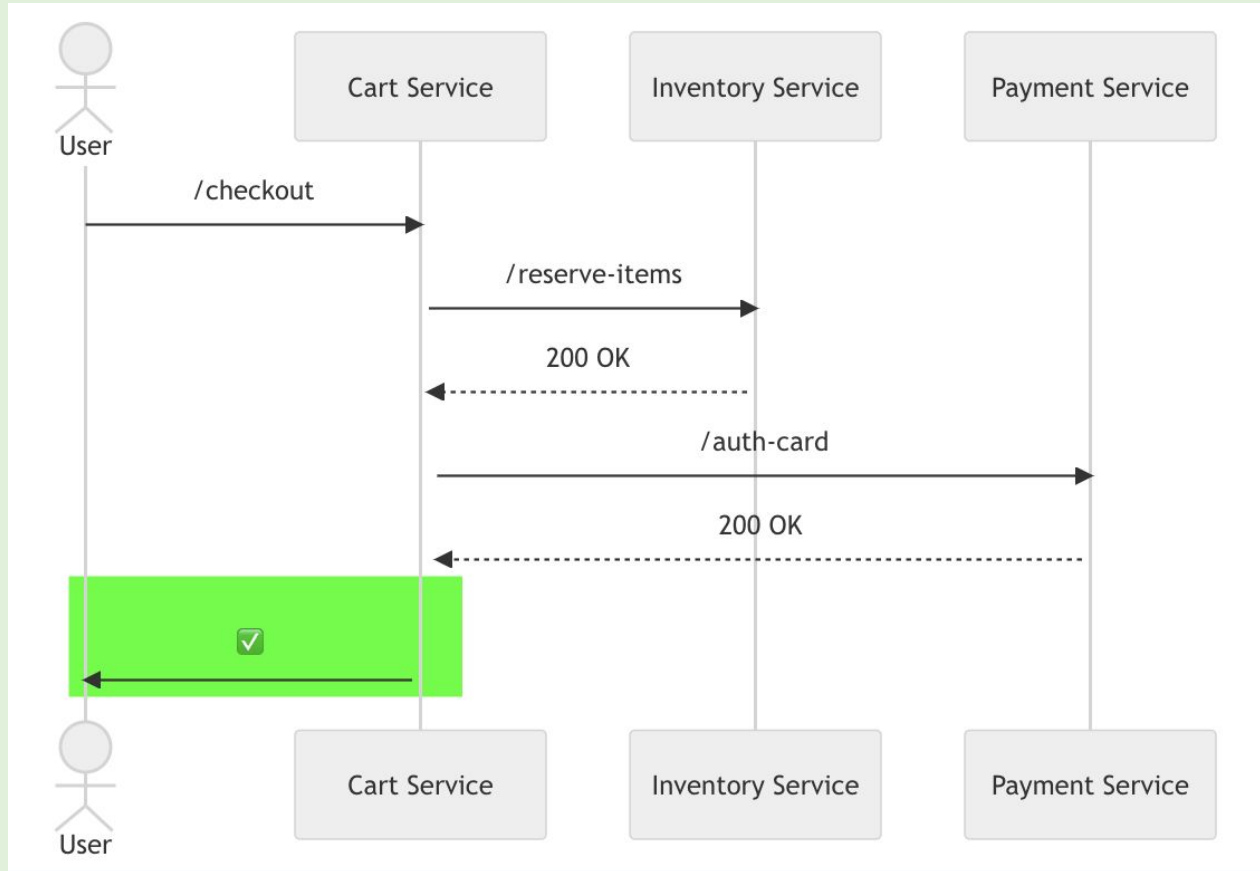
System



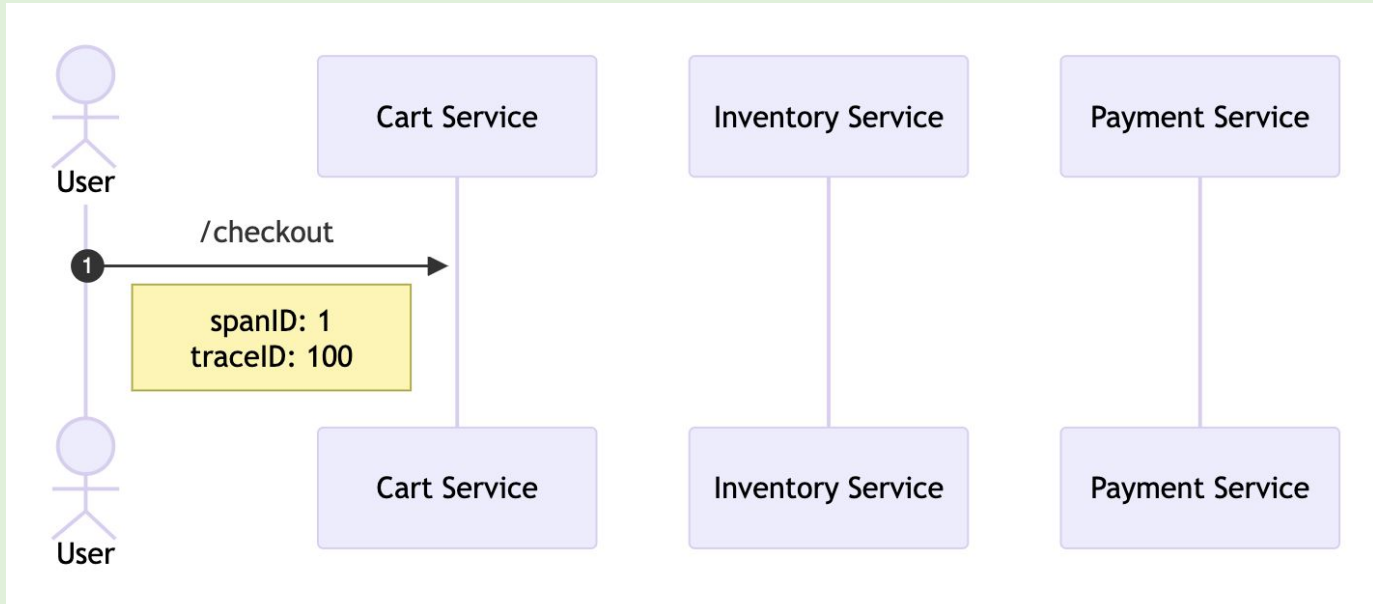
System

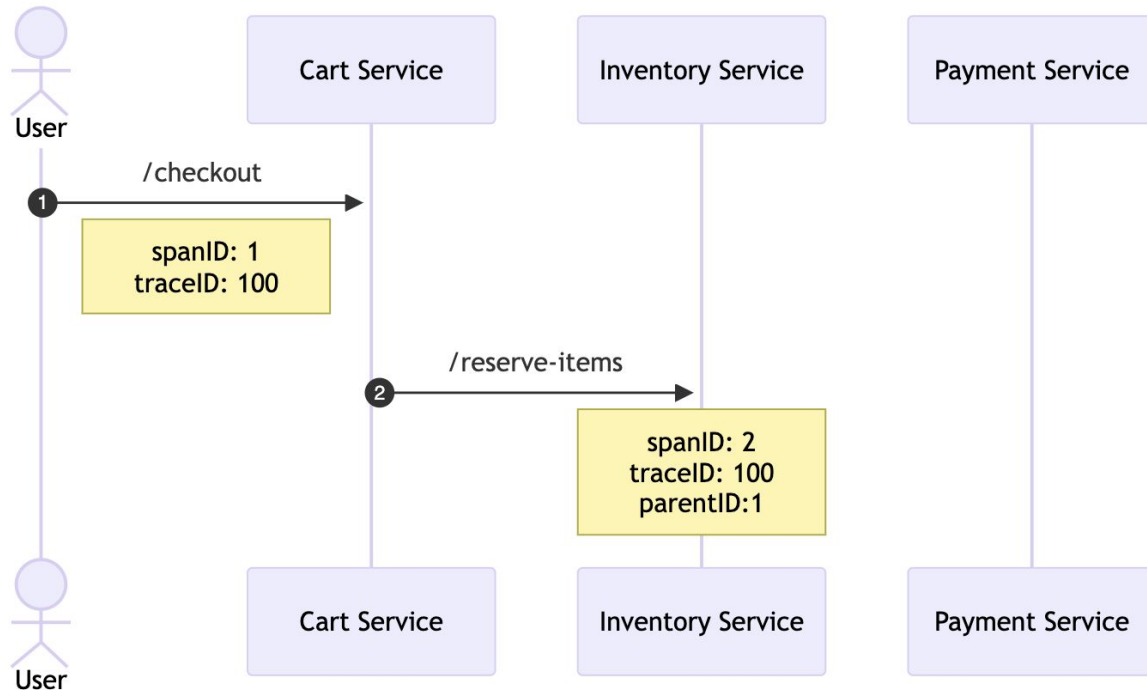


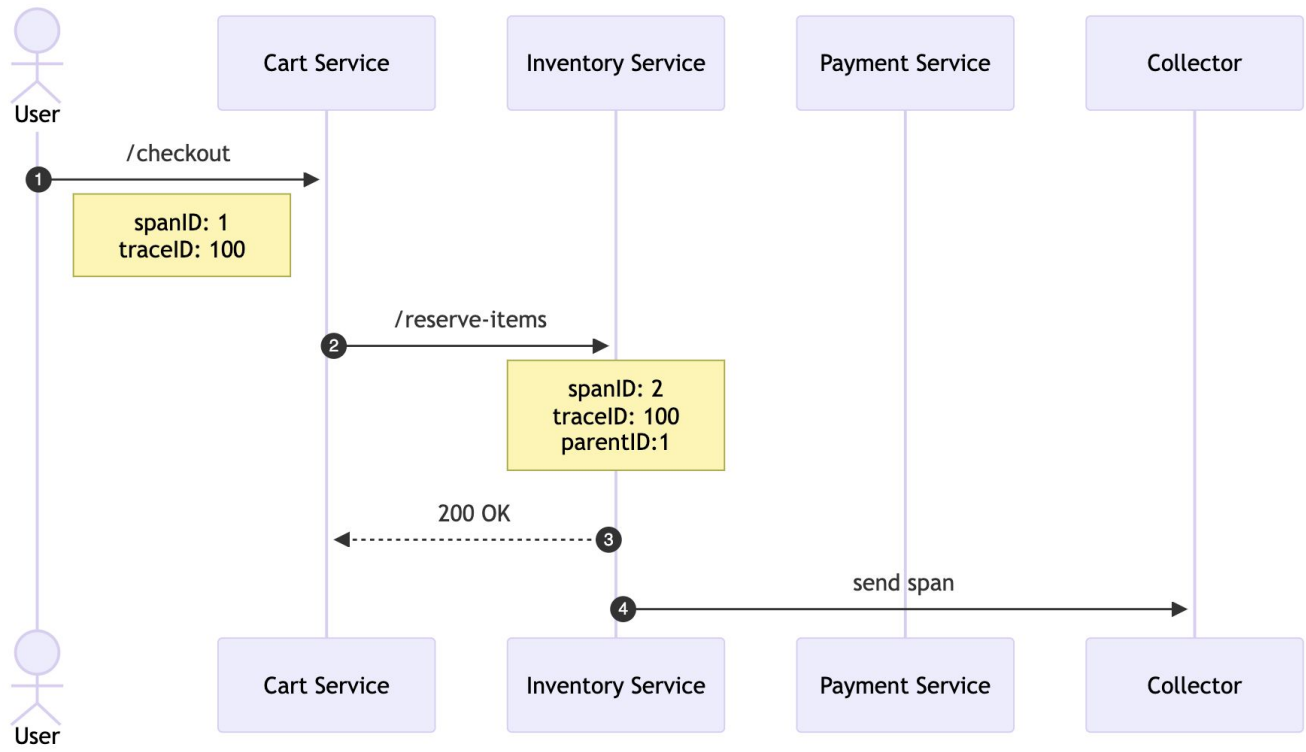
System

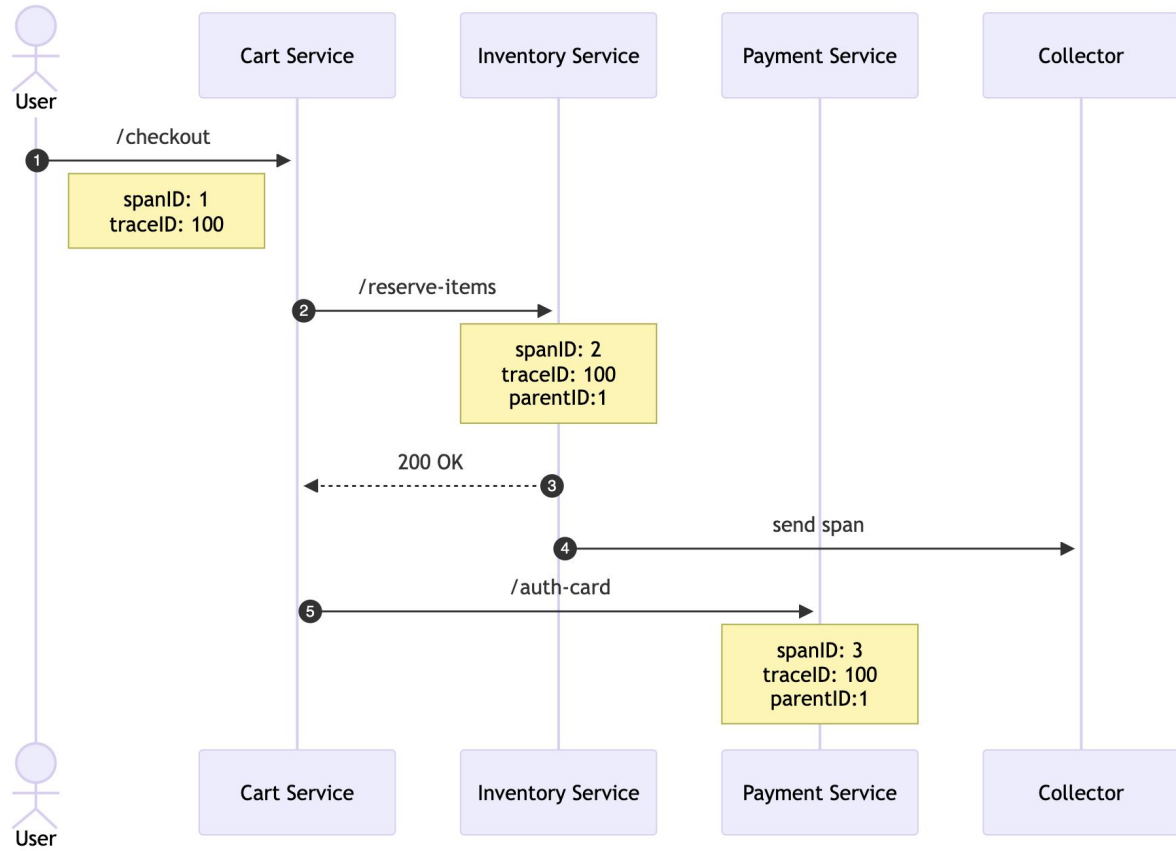


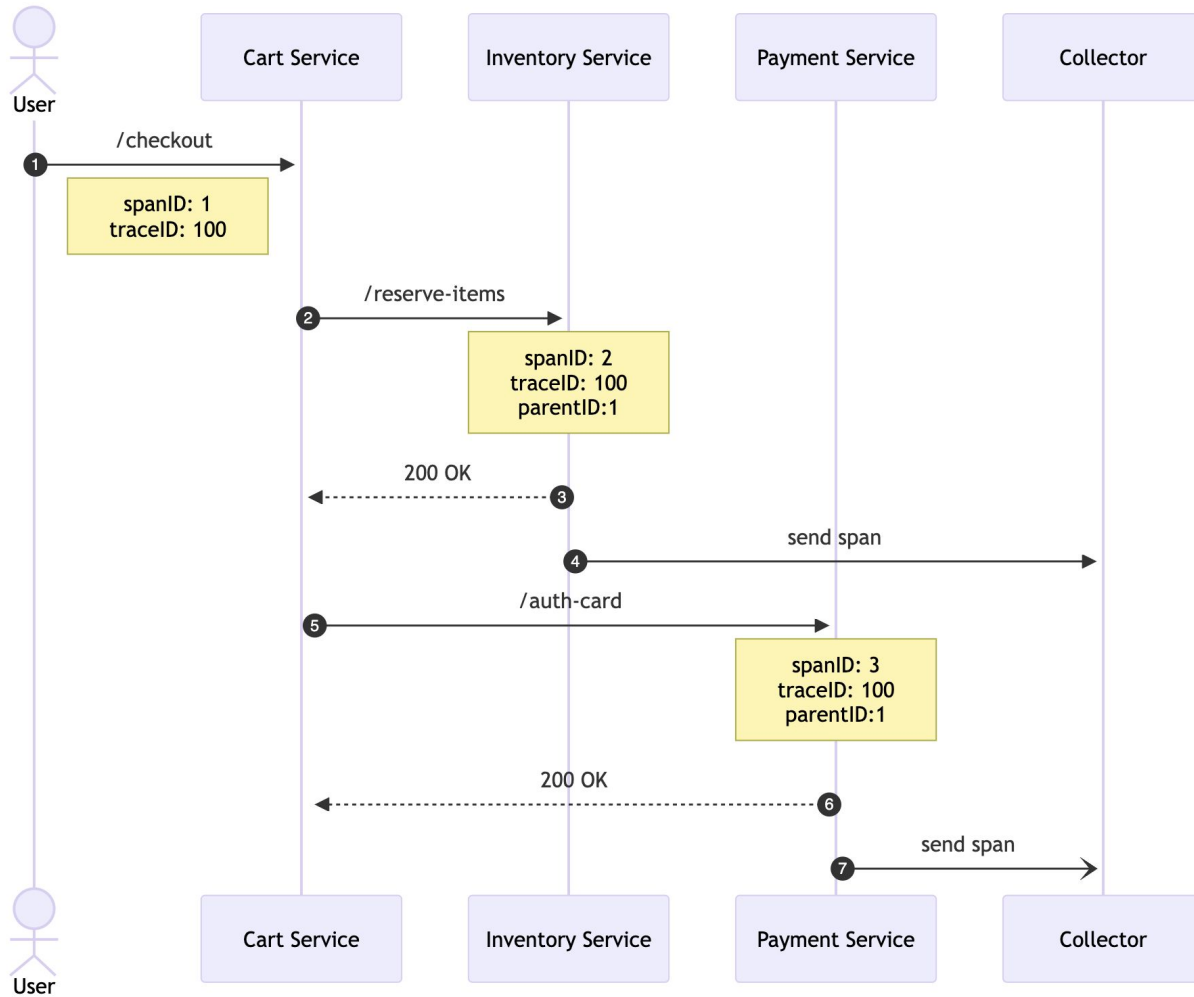
How Spans Become A Trace

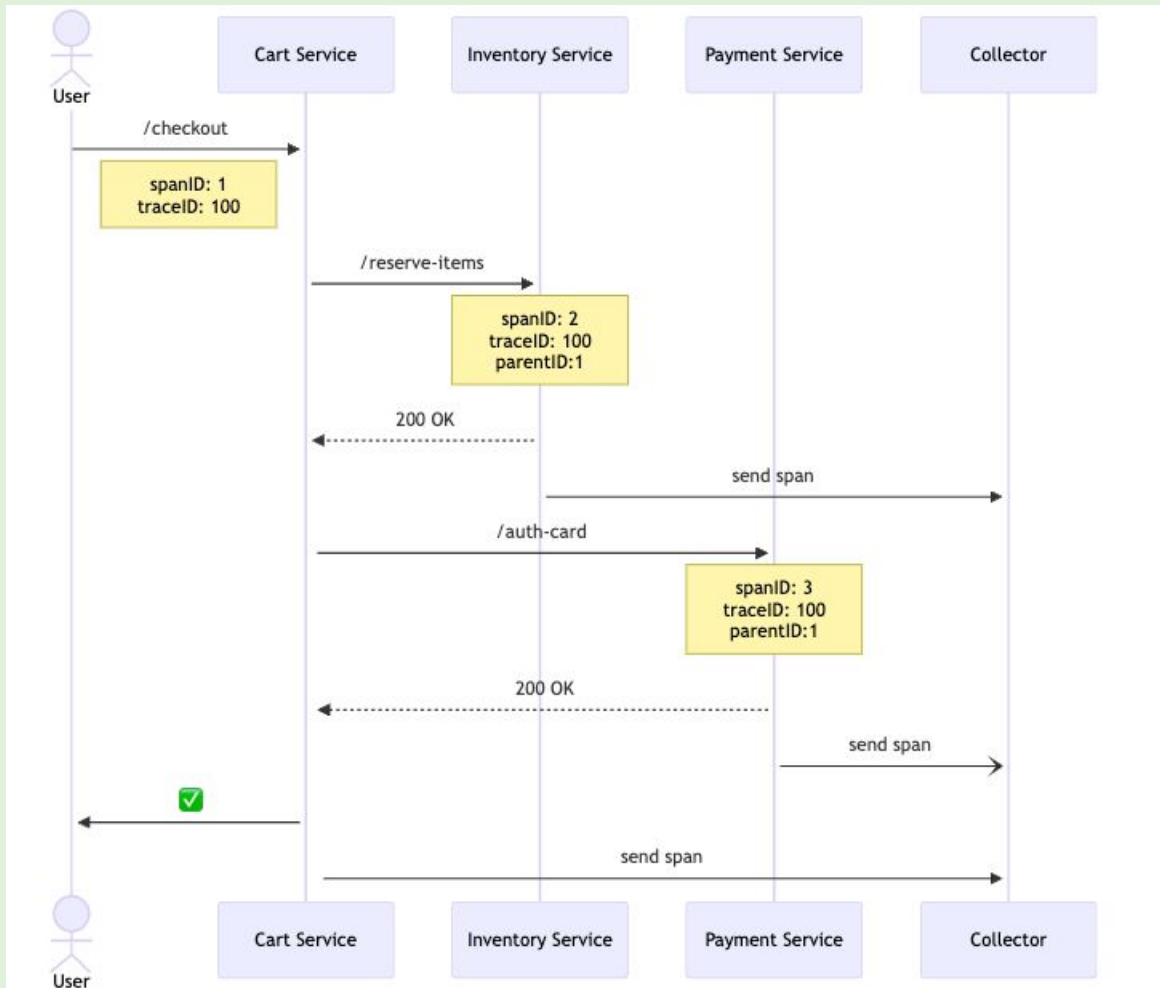













Trace Waterfall



03




**Heads or Tails?
Sampling Strategies**

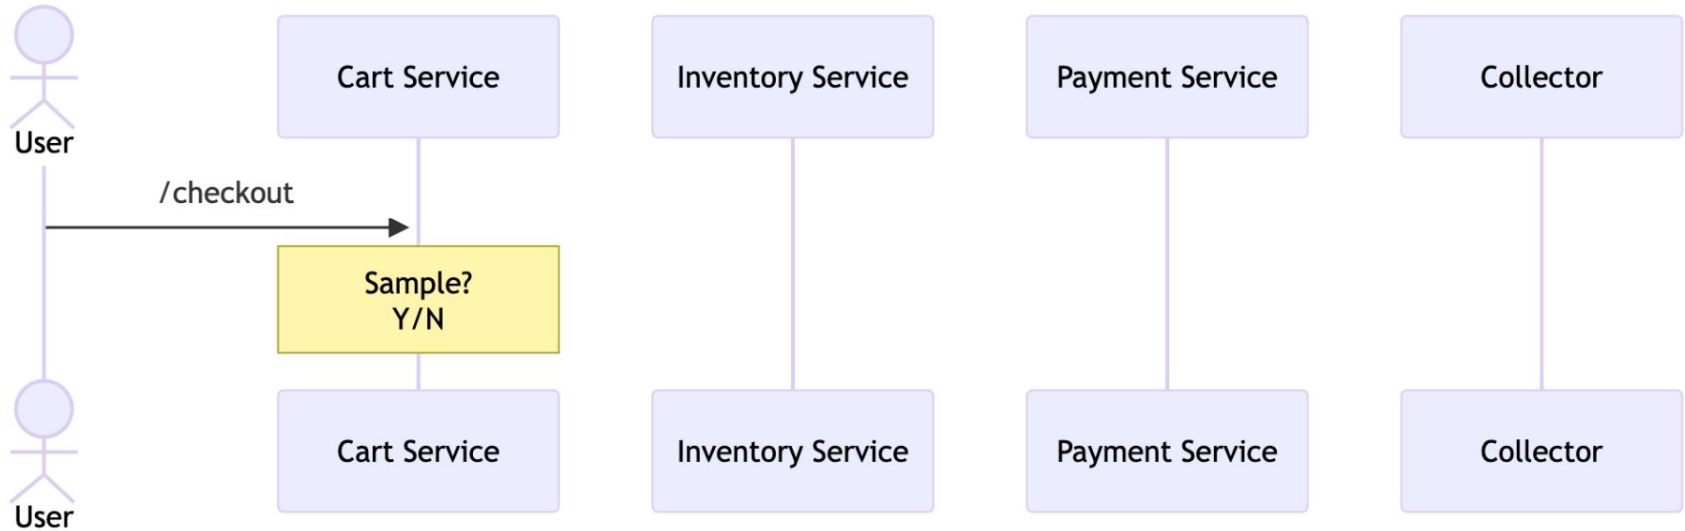
Constant Sampling



Head Based Sampling



Head Based Sampling



Head Based: Probabilistic



“Sample 1% of traces”

Configured in the SDK

```
OTEL_TRACES_SAMPLER=parentbased_traceidratio
```

```
OTEL_TRACES_SAMPLER_ARG="0.01"
```



Head Based: Rate Limit

“Sample up to 10 traces per second”

Configured in the SDK & Jaeger Collector*


`OTEL_TRACES_SAMPLER=parentbased_jaeger_remote`

`OTEL_TRACES_SAMPLER_ARG=`

endpoint=your-endpoint-here

pollingIntervalMs=5000

initialSamplingRate=0.1




Head Based: Adaptive

“Sample 1 trace per second per endpoint”

Configured in the SDK & Jaeger Collector

OTEL_TRACES_SAMPLER=parentbased_jaeger_remote

OTEL_TRACES_SAMPLER_ARG=
endpoint=<http://localhost:14250>
pollingIntervalMs=5000
initialSamplingRate=0.25



Head Sampling Trade Offs

Benefits

- Efficient for high throughput systems
- Limits data generation
- Easy to reason about
- Simple set up



Head Sampling Trade Offs




Drawbacks

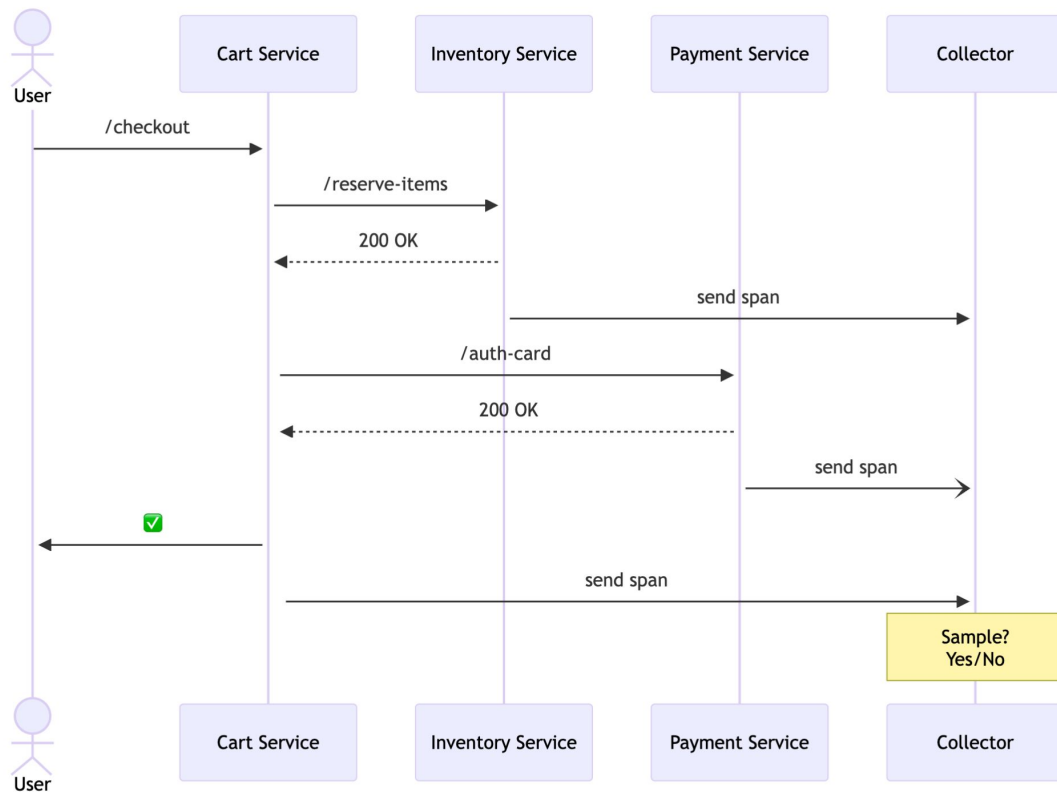
- Can sacrifice “interesting” or edge case traces




Tail Based Sampling



Tail Based Sampling



Tail Based Sampling*



Latency



Status Code




Numeric



String




Error



Span Count



TraceState




Composite



Prob.




Rate Limit



Tail Based: Latency

“Sample all traces to /checkout where latency > 2m”

Configured in the *TailSamplingProcessor* in OTel Collector



AND policy

- string_attribute **service.name** = **cart**
- string_attribute with regex for **http.route** = **[/checkout/.+]**
- latency with **threshold_ms**: 120000



Tail Based: Error

“Sample all traces on /checkout that have errors”




AND policy

- string_attribute **service.name** = **cart**
- string_attribute with regex **http.route** = **[/checkout/.+]**
- status_code {**status_codes**: **[ERROR]}**



Tail Based: Attributes

“Sample all traces on /checkout where cart is > \$1000”



AND policy

- string_attribute **service.name = cart**
- string_attribute with regex **http.route = [/checkout/.+]**
- numeric_attribute
 - **key: cart.total**
 - **min_value: 1000**

Tail Sampling Trade Offs

Benefits

- Focuses on capturing interesting traces
- Flexible configuration



Tail Sampling Trade Offs

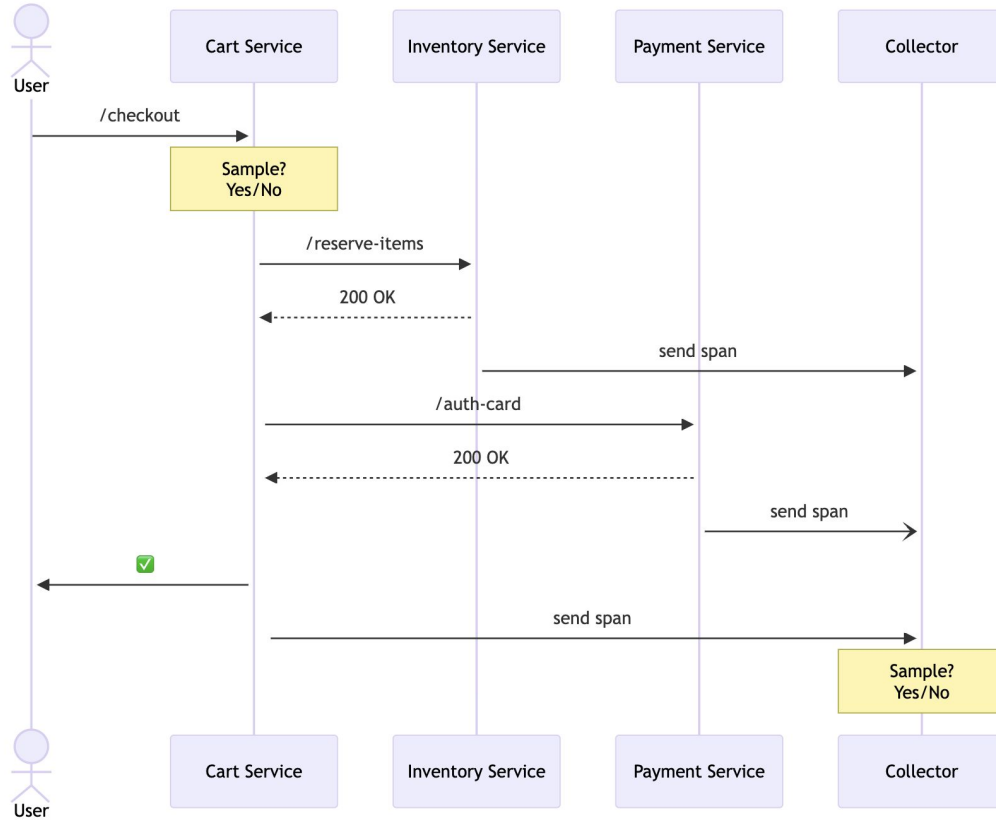


Drawbacks

- Inherently higher cost (computational/financial)
- Variable/unpredictable, fluctuates with throughput, span size, trace size
- Introduce latency to make sampling decision



Head & Tail Combo





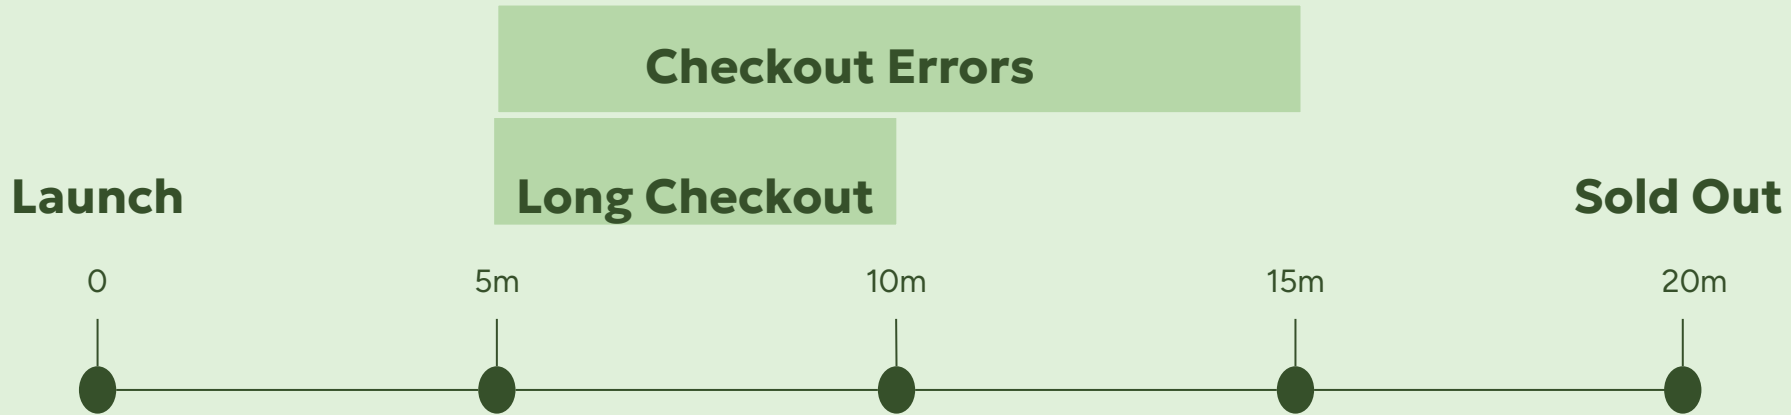
Example Time!




HALP!

- Yesterday's drop sold out in 20m
- Cat-lebrity complained about a long checkout time
- Confirmed reports of failed payments

Timeline



Constant Sampling - volume




/checkout requests 600,000

traces generated 600,000

traces stored 600,000

Constant Sampling - “interestingness”





success	590,500
---------	---------

error	9,000
-------	-------

high latency	1,500
-----------------	-------

Constant Sampling



Head Sampling 1% - volume




/checkout requests	600,000
--------------------	---------

traces generated	6,000
------------------	-------

traces stored	6,000
---------------	-------



Head Sampling 1% - “interestingness”





success	5,895
---------	-------

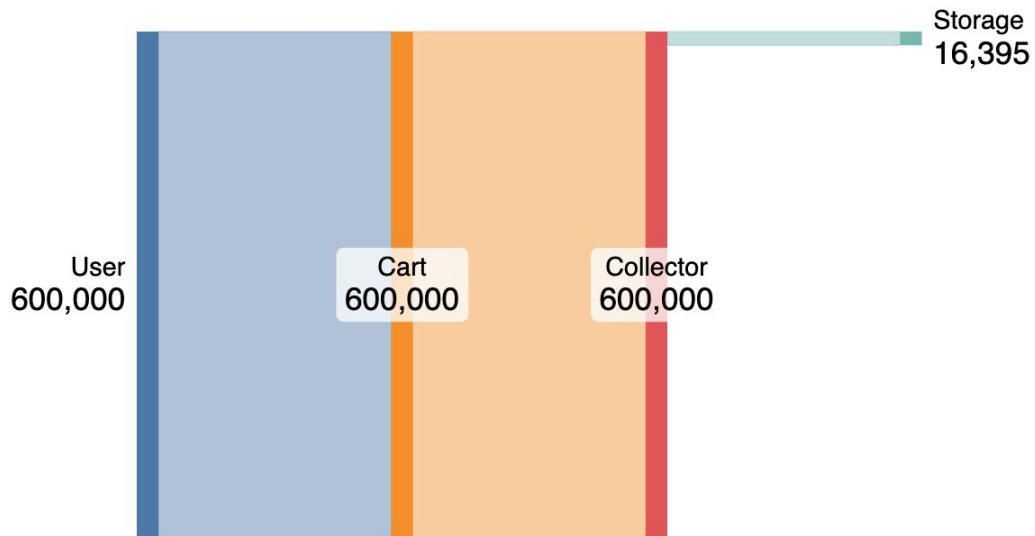
error	90
-------	----

high latency	15
-----------------	----

Head Sampling 1%



Tail Sampling - volume




/checkout requests 600,000

traces generated 600,000

traces stored 16,395

Tail Sampling - “interestingness”


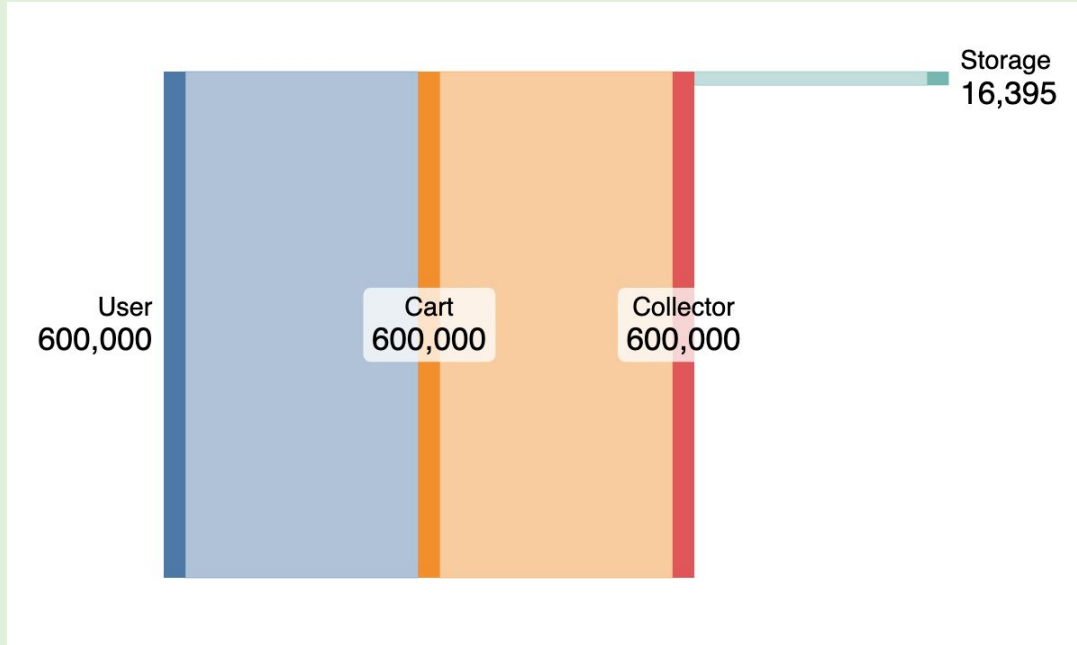


success	5,895
---------	-------

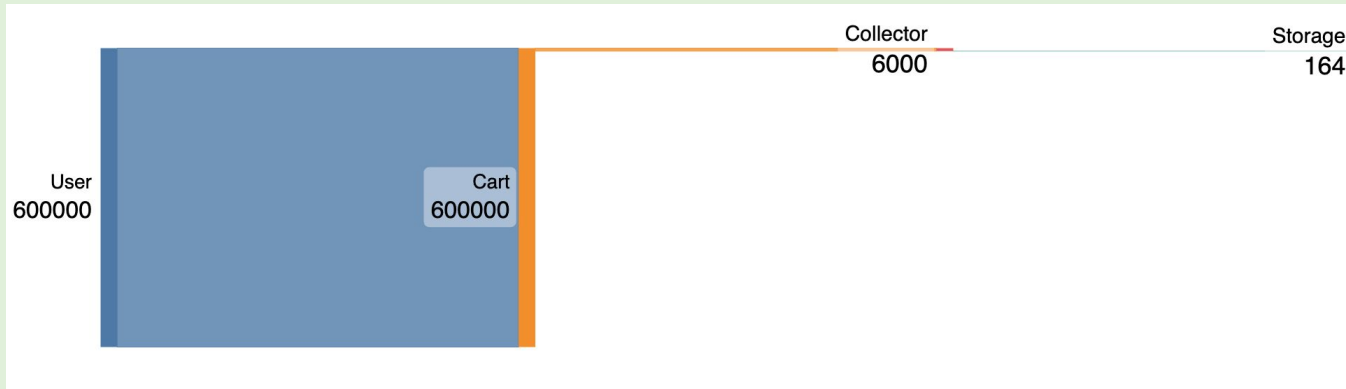
error	9,000
-------	-------

high latency	1,500
-----------------	-------

Tail Sampling 1% success, 100% error + lat



Head & Tail Sampling Combo - volume




/checkout requests 600,000

traces generated 6,000

traces stored **164**

Head & Tail Sampling Combo - “interestingness”





success 59

error 90


high
latency 15

Head & Tail Sampling Combo




Interestingness Comparison

No Sample




600,000 traces

Tail




16,405 traces

Head @ 1%



6,000 traces

Head & Tail




164 traces

04



“Gotchas”

Set It & Forget It Sampling



Siloing Telemetry

Traces < > Logs


- Inject traceIDs & spanIDs into logs

Metrics < > Traces

- Exemplars



Sampling as *only* Shaping Strategy



05



**Tomorrow's
Tracing**



Head


vs.

Tail




[Tracing] made queries like

**find me requests where
the user was logged in**




[Tracing] made queries like

**find me requests where
the user was logged in
the request took more than 2s**




[Tracing] made queries like

**find me requests where
the user was logged in
the request took more than 2s
only certain databases were used**




[Tracing] made queries like

**find me requests where
the user was logged in
the request took more than 2s
only certain databases were used
a transaction was held open for > 500ms**



[Tracing] made queries like

**find me requests where
the user was logged in
the request took more than 2s
only certain databases were used
a transaction was held open for > 500ms**



possible



paigerduty

.com



@hachyderm.io



@chronosphere.io



CREDITS: This presentation template was created by [Slidesgo](#), and includes icons by [Flaticon](#), and infographics & images by [Freepik](#)