

Observability Cannot Continue *As It Is.*

What one of the largest financial institutions in the world told us last week,
and why it is true for every team here.

Ari Recht

CO-FOUNDER & CEO



“ Observability will be *unaffordable*.”



“I am sitting on a product that everybody says they need. And it will be unaffordable for the business, because the business is not going to pay for it. If we do nothing, we are screwed economically, and we are screwed as a negotiating position.”

Code and deploys grow with AI.

Budgets grow with inflation.

CODE WRITTEN

10×

When writing code is no longer the constraint, the business will solve far more problems with code.

DEPLOYS

3–5×

Models chain exploits now. You patch aggressively just to keep up with the CVEs. CVEs.

OBSERVABILITY SPEND

1×

Nobody pays 10× for tools priced per gigabyte. gigabyte. They are not even going to pay 1× 1× the people.

Every legacy pricing model, ingestion, hosts, indexed terabytes, is linear in exactly the thing that is about to go exponential.

Legacy tools were designed for people.

People will read 10% of the data.

AGENTS AND AUTOMATION

~90%

Every P3. Every root-cause investigation. Every “did my deploy break anything?” Three minutes to answer is fine. That is machine clock, not human clock.

HUMANS

~10%

P0 and P1. A region is down.
Dashboards, war rooms, swarming. Still
Still real, no longer the whole product.
product.

Dashboards, indexed search and per-seat UX are the 10%. We are paying 100% for them.

Our reflex is to collect everything.

More data was never the answer.

01 IT BANKRUPTS YOU

77 billion spans a month

One agent workflow is ~300 spans, not ~10.
Keep it all and the bill, not the outage, is what takes you down.

02 IT DROWNS THE READER

“Great, you have it all. Now what?”

Agents are extraordinary at needle in a haystack. Nobody gets better by making the the haystack bigger.

03 IT IS THE WRONG DATA ANYWAY

Volume without shape

A million log lines written for a human to grep
grep are not the same thing as one trace that
that shows the whole request, end to end.
end.

The instinct that built the last decade of observability, “index it all, decide later,” is the instinct that ends it.

The answer is not more data. *It is better data.*

COMPLETE

The whole request, not a sample of sample of it

One end-to-end trace through the prompt, the API, the database and back carries more truth than a terabyte of logs. Including the services nobody instrumented.

IN CONTEXT

What changed, what it touched, what it did

The deploy three hours ago. The queue that that backed up. The exception in the stack. Telemetry an agent can act on, not telemetry a person has to interpret.

AT THE RIGHT MOMENT

Depth when it matters, silence when it doesn't

Most of the time, almost nothing. When something breaks, everything. The trace that that holds the failure is the only one you ever ever needed.

Quality is a property of what you capture, and of when you decide to capture it.

What if the instrument could *change its mind*?

01

What if instrumentation wasn't a decision decision you made at deploy time?

02

What if depth followed the incident, instead of the other way around?

03

What if the pipeline knew the difference difference between peacetime and wartime, without anyone clicking a button? button?

Observability is the foundation now. Build it like one.

Ari Recht · THANK YOU

