laserdata

RUST INDIA CONFERENCE · BENGALURU · 2026

Fast Isn't Enough.

Why latency still breaks streaming systems.

Kranti Parisa
Founder & CEO, LaserData · Apache Iggy PPMC

laserdata

REAL-TIME AI, WHEN IT WORKS

A sales call with live transcription.

Every word appears as it is spoken. Action items surface automatically. When it works, the AI is invisible.

END-TO-END < 100 MS · ZERO DROPPED WORDS

laserdata

THEN THIS HAPPENS

One delay.
The conversation breaks.

WORDS DISAPPEAR

A sentence drops. The rep loses context mid-call.

ACTION ITEMS WRONG

AI captures the wrong task from an incomplete sentence.

TRUST COLLAPSES

The rep stops reading the transcript. The AI becomes noise.

laserdata

WHERE SYSTEMS FAIL

The tail is where experience breaks.

MEDIAN	95TH	99TH	99.9TH
P50	P95	P99	P99.9
Unnoticed.	Noticeable.	Painful.	System failure.

Your worst users — the ones most likely to churn — **live at the tail.**

laserdata

THE UNCOMFORTABLE TRUTH

Everyone's building workarounds.

Capacity.

Scale E2E infra artificially.

Cache.

To mask broker latency.

Bypass.

Custom RPCs around the broker.

Patch.

Hand-rolled pub-sub for hot paths.

Not solutions. Symptoms of unpredictability.

laserdata

THE TAIL LATENCY TAX

Complexity as a coping mechanism.

TODAY

Producers · Services

↓

L1 Cache Layer — masks broker lag

↓

Custom RPC Bypass — hot paths

↓

3–5× partitions — unpredictable P99

↓

Consumers · Services

~4× infra cost · P99 unknown

WITH A PREDICTABLE ENGINE

Producers · Services

↓

Streaming engine

predictable P99 — designed, not discovered

↓

Consumers · Services

Right-sized · P99 as a contract

laserdata

WHY NOW

Data is exploding — and going real-time.

394 ZB

GLOBAL DATA BY 2030 — 23% CAGR

~40%

NEEDS REAL-TIME PROCESSING BY 2030

75 B

CONNECTED DEVICES STREAMING EVENTS

laserdata

AI ERA

Not one prediction. Continuous reasoning.

TRADITIONAL ML

Move data to the model.

Batch

Nightly ETL · hours late is fine.

One hop

Request → prediction → done.

AGENTIC AI

Stream events between models.

Streams

Continuous, multi-source, always live.

Many hops

Agents, tools, memory — every turn.

Same building blocks. **Completely different data movement.**

laserdata

AT SCALE

Message streaming is the nervous system.

AI is the **brain.** | Streaming is the **nervous system.**

laserdata

Compute evolved.
Data movement didn't.

Streaming must evolve the way GPUs evolved for AI.

Not band-aids. A fundamental shift.

laserdata

THE INSIGHT

Same fundamentals. New engine.

PREVIOUS GENARATION Designed 2011 · JVM · Throughput-first	IGGY Rust · Designed for predictability
Streams ✓	Streams ✓
Topics ✓	Topics ✓
Partitions ✓	Partitions ✓
Consumer groups ✓	Consumer groups ✓
JVM · GC pauses	**Rust** — no GC, zero-cost
Complex, external consensus	**Built-in consensus**
OS page cache	**Direct I/O** — bypass page cache
Throughput-optimized	**Predictability-first**

laserdata

APACHE SOFTWARE FOUNDATION

The next phase in data streaming.

Engineered for predictability, not throughput alone.

Written in Rust | Low protocol overhead | Minimal data movement | Predictability-first

4,100+
GITHUB STARS

Top 3
RUST PROJECT AT APACHE

PPMC
LEADING FROM THE FRONT

laserdata

PROOF, NOT PROMISES

Benchmarked.
Not estimated.

1M+ msg/s
Throughput
+1 GB/s

1.01 ms
Avg latency
Producer-side, sustained

2.05 ms
P99 latency
The tail, measured

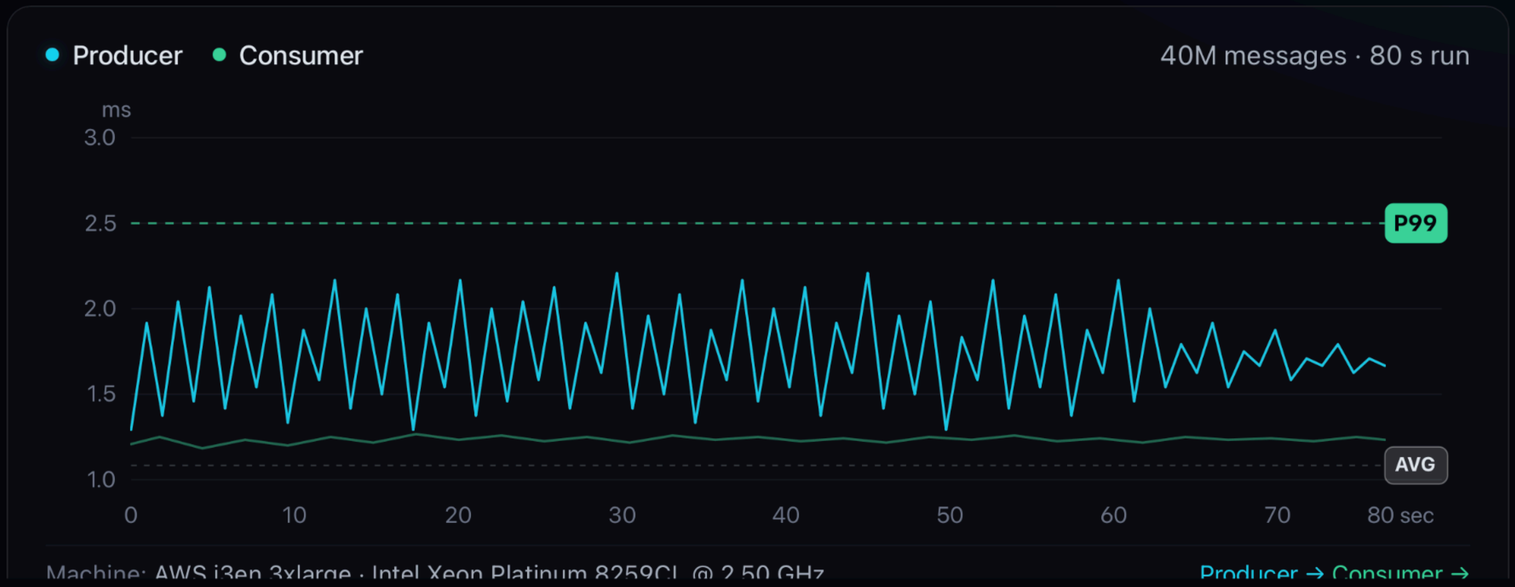

laserdata

LASERDATA

Iggy at internet scale.

Enterprise streaming for the AI era. Built on Apache Iggy.

- **Blazing Fast & Predictable**
  1 ms P99 · 1 GB/s R+W · 1 M+ msgs/sec per node
- **Multi-Cloud**
  AWS · GCP · more coming
- **Dedicated Infrastructure**
  No noisy neighbors · Built-in observability
- **Open Source Core**
  Creators of Apache Iggy · 4 protocols · 6 SDKs
- **Flexible Deployment**
  Managed · BYOC · Private · On-prem
- **Zero-Trust Security**
  Enterprise-grade · RBAC · No inbound, no vendor access

laserdata

Fast is easy.
Predictable is hard.

The compute layer got its GPU moment.
Data infrastructure is next.

If not the Rust community — who?

laserdata

Obsessed about tail latencies?

Come build with us.

Register for the FREE TIER. Join the community.
Let's take this to the next level — together.

Kranti Parisa
Founder & CEO, LaserData · Apache Iggy PPMC

laserdata.com

iggy.apache.org

SCAN TO VISIT

Thank you.