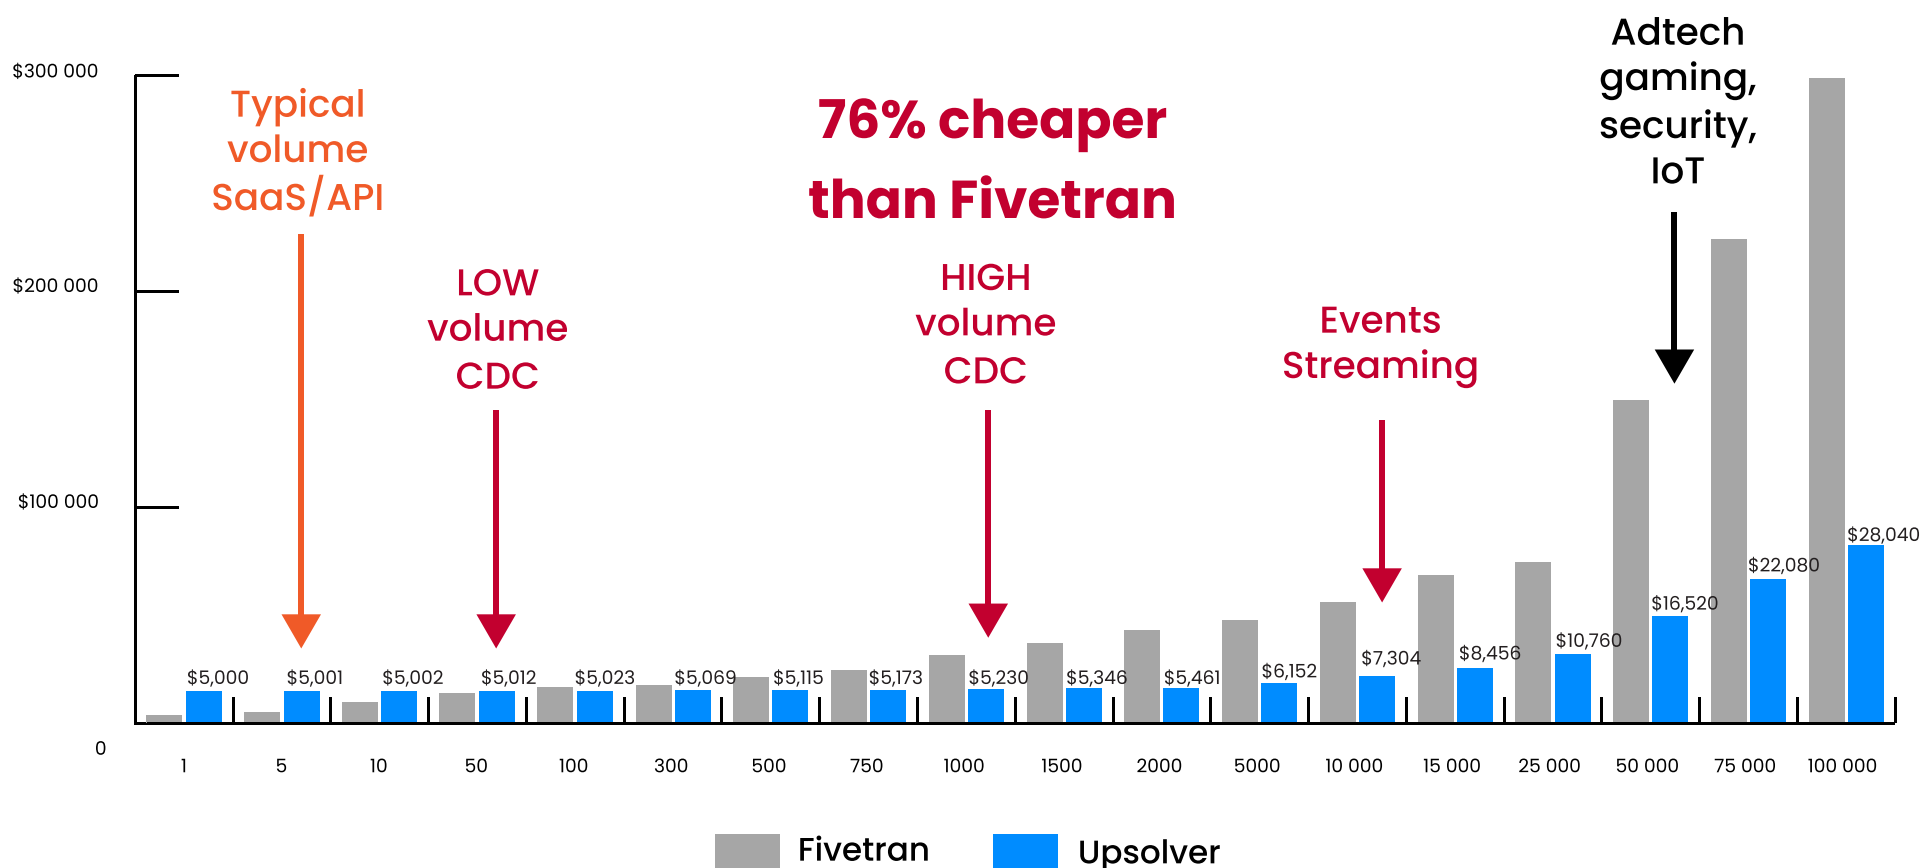


Upsolver vs Fivetran – Pricing Comparison

Most companies end up overpaying for Fivetran. Choose Upsolver when you need to scale.

Upsolver charges by data volume. Fivetran charges by 'active rows'.
 This key difference drives order-of-magnitude savings when data volumes scale



Fivetran

- Charges per monthly active row by complex calculations per source; very difficult to predict costs ahead of time
- Costs spike rapidly with data volume, or when data needs to be resynced frequently

Upsolver

- Charges per terabyte ingested: \$225/mo (standard edition not including volume discounts)
- Predictable pricing

Based on real-world workloads:

Upsolver delivers

76% cost reduction

For continuous CDC and file ingestion workloads, compared to Fivetran

Upsolver delivers

90% cost reduction

For streaming ingestion workloads, compared to Fivetran

Comparing common high-scale use cases

Upsolver vs Fivetran for CDC

Example use case: replicating operational databases into warehouse and lake.

Fivetran

- Fivetran uses full snapshots and incremental loads. This fails to scale and stresses source databases as data volume and frequency of changes increase.
- For higher volumes Fivetran offers a High Volume Agent, which comes with additional cost and infrastructure complexity.

Upsolver

- Upsolver leverages open source Debezium technology and distributed processing to provide a linearly scalable CDC solution without additional software.
- Upsolver's volume-based pricing is predictable as data scales.
- No agents installed or manage.

Bottom line: Fivetran's pricing and infrastructure requirements spike rapidly with CDC volume. Upsolver's CDC scales efficiently at a consistent price per TB.

Upsolver vs Fivetran for Semi-Structured Ingestion from Object Storage

Example use case: Ingesting logs or data dumps from legacy systems (DB2)

Fivetran

- Fivetran's MAR-based pricing is not designed for file-based data movement from object stores.
- Frequent file changes result in over paying for repeated MARs as files need to be fully re-synced

Upsolver

- Upsolver automatically detects file changes and processes updates incrementally.
- Volume-based pricing is predictable as data scales and file changes accelerate.
- You pay only for the data you need, not redundant information.

Bottom line: Fivetran's pricing is designed for changed rows, not files. Incremental processing is inefficient because changed rows are more difficult to detect. Upsolver pricing is designed to scale with your data. Incrementally processing data means you pay less.

Upsolver vs Fivetran for Streaming Ingestion

Example use case: Ingesting application or security events from Kafka or Kinesis.

Fivetran

- Fivetran counts each event's unique offset as a distinct MAR.
- You pay for every single message consumed from a stream, resulting in exorbitant costs.

Upsolver

- Upsolver charges by the volume, in TB, of data consumed from streams.
- Volume-based pricing is predictable regardless of the number of unique events.
- Additional volume discounts are available.

Bottom line: Fivetran's pricing is a non-starter for any streaming workload. Upsolver is a streaming engine built and priced for scale. Upsolver's largest customer transfers over 5MM events/second.

Stop overpaying for an entry level data ingestion tool. Choose a simple to use, highly scalable and cost effective ingestion tool to serve your current and future needs.

Get in touch to discuss your data ingestion needs and discover how much you can save.

Start here: <https://upsolver.com/demo>

Free trial available: <https://www.upsolver.com/signup>