


Kafka-compatible streaming with zero disks, zero interzone fees, and zero ops

WarpStream is the original diskless, Apache Kafka® protocol-compatible data streaming platform. Its stateless Agents run in your cloud account and write straight to object storage, so it's 5–10x cheaper than the alternatives, trivial to operate, and never has access to your data.



10x

cheaper than Kafka

\$0

interzone networking fees

30x

lower storage costs than
cloud disks

100+

customers in production

Why teams switch



5–10x lower cost

Kafka's two biggest cloud bills, interzone networking and triple-replicated disks, go away: Agents write to object storage, and moving data between zones through it is free.



Zero ops

Agents are stateless, so clusters auto-scale on CPU like a web server. No brokers or partitions to rebalance, no capacity planning, and no overprovisioning for peaks.



Zero access BYOC

Your compute, your buckets. Only metadata reaches WarpStream's control plane: no cross-account roles, no VPC peering, and no break-glass access to your data.

In production at over 100 customers, including



character.ai



mixpanel



dish



SUPERWALL

eternal


Secureworks

Backed by IBM, with up to a 99.999% uptime SLA. Runs on AWS, Google Cloud, and Azure, and on any S3-compatible object storage, even on-premise.



Zero is better

WarpStream is one stateless Go binary that speaks the Kafka protocol and keeps everything in object storage: no disks to manage, no brokers to rebalance, and no consensus layer (not even KRaft) to run.



Clients only talk to Agents in their own zone. Agents write to object storage in your account, and send WarpStream Cloud only metadata.



Zero interzone networking

Kafka **\$0.053** → WarpStream **\$0** per GiB

Kafka pays each time data crosses a zone, even with fetch from follower: 60% of the bill at 1 GiB/s and a day of retention. WarpStream moves data between zones through object storage instead, for 3% of that in API fees.



Zero disks

Kafka **\$0.60** → WarpStream **\$0.02** per GiB-month

Kafka keeps three copies on cloud disks it can only half fill. WarpStream keeps none: Agents write straight to object storage, which keeps data durable for you.



Zero ops auto-scaling

Kafka is almost impossible to auto-scale, so most clusters run 2–4x overprovisioned. Stateless Agents scale on CPU like a web server: always right-sized, with no partition rebalancing or capacity planning.



Zero access

Your compute processes the data and your buckets store it; only metadata leaves your environment. WarpStream needs no way to escalate privileges, even in an emergency: **it can't access your data under any circumstances.**

The one trade-off is latency, and it's yours to tune

By default, Produce requests take about 500 ms at the P99, and end-to-end latency is just under a second. Lightning Topics on S3 Express One Zone cut Produce to **33 ms median, 50 ms P99**. Durability never changes: as with Kafka at acks=all, no write is acknowledged until it's in object storage.



5–10x cheaper, and the math is public

WarpStream is the only diskless, Kafka-compatible product that combines usage-based pricing with BYOC, and every price on this page comes from its public calculator.

Pay for what you use

Cluster tier, writes, and storage.
Nothing per core or per Agent, and
no set hourly management fee.

Public unit prices

The only BYOC product that
publishes them, with a calculator for
any workload.

Automatic volume discounts

Write and storage prices step down
on their own as volume grows.

Monthly cost of ownership

Apache Kafka (3 AZ)

100 MiB/s · 7 days



82%
cheaper

Apache Kafka (3 AZ)

1 GiB/s · 1 day · fetch from follower



67%
cheaper

Amazon MSK

1 GiB/s · 7 days · fetch from
follower



87%
cheaper

Amazon MSK Serverless

1 GiB/s · 1 day



89%
cheaper

WarpStream's totals include its licensing fees. Kafka's don't include the engineering hours spent running it.

Where the money goes: 1 GiB/s with 1 day of retention, fetch from follower on

Apache Kafka (3 AZ)

\$57,654/mo



WarpStream

\$18,483/mo



Kafka's interzone networking alone costs more than all of WarpStream.

What customers saved by switching

character.ai **90%**

GOLDSKY **90%**

PIXELFEDERATION* **83%**

eternal **60%**

ShareChat **60%**

SUPERWALL **57%**

Robinhood **45%**

The savings each customer
published in its case study at
warpstream.com.



A complete streaming platform, still zero access

Beyond the Kafka protocol, WarpStream ships what teams usually bolt on, built on the same stateless design and running in your own environment.



Agent Groups

Split one logical cluster into isolated groups: analytics apart from transactions, or one cluster across VPCs, regions, or clouds without VPC peering.



Managed Data Pipelines

Bento stream processing embedded in the Agents (transforms, aggregations, enrichments, and WASM), run from the managed control plane.



Orbit

Copy topics with their offsets, consumer groups, and ACLs from any Kafka-compatible source: for migrations, disaster recovery, and read replicas.



BYOC Schema Registry

API-compatible with Confluent's, embedded in the Agents, with schemas in your bucket and record validation built in.



Tableflow

Iceberg tables from any Kafka topic, with ingestion, compaction, and maintenance automated, queryable from Snowflake, Databricks, BigQuery, and more.



Multi-Region Clusters

Zero data loss (RPO=0) across regions out of the box, with no extra operations, backed by a 99.999% SLA on the Enterprise tier.



Ripcord Mode

Agents keep accepting Produce requests through a control plane outage, rather than halting your pipeline.

Runs anywhere

AWS S3, Google Cloud Storage, Azure Blob Storage, and any S3-compatible object storage, in the cloud or on-premise.



Google Cloud



MINIO

ORACLE

Proven at scale

6.43 PiB

written by one tenant in a single day

44.8 PiB

written by one tenant in a single week

5 GiB/s

streamed by Character.AI

Kafka bytes written by WarpStream's largest production tenants, October 2026; names withheld.

"It looks like Kafka, it talks like Kafka, but it doesn't have all the pain points of Kafka. [...] It just works."

Maor Brill

Principal Engineer, Character.AI

"If you're running Kafka for anything, you're doing it wrong. WarpStream is the way. [...] It's totally elastic. We can increase our volume as much as we want and it's pretty much ops free."

Brian Anglin

Co-Founder and CTO, Superwall

"[...] after 7+ years of running and interacting with all these platforms, I can genuinely say that WarpStream is the easiest platform to run."

Yaroslav Tkachenko

Former Principal Software Engineer, Goldsky



How WarpStream compares

Kafka, and the services built on it, keep its disks, its interzone replication, and its bill. Newer object-storage designs still keep brokers, write-ahead logs, or classic topics.

ALTERNATIVE	WHAT YOU GET	WITH WARPSTREAM
Apache Kafka	Disk-based brokers replicated across zones, overprovisioned for peaks, with partitions to rebalance.	67% cheaper at 1 GiB/s, even with fetch from follower, before counting the engineers who run Kafka.
Amazon MSK	Managed EC2 running disk-based Kafka, at an almost 2x EC2 premium, with no automatic rebalancing or auto-scaling.	48%–87% cheaper at 1 GiB/s, with 1 to 7 days of retention.
Amazon MSK Serverless	Capped at 200 MiB/s of ingress and 3,000 client connections, with a replication factor of 2.	89% cheaper at 1 GiB/s, where MSK Serverless costs 9.2x as much.
Redpanda	Kafka rewritten in C++, with the same disks and interzone fees plus licensing, no public pricing for its flagship products, and BYOC that can escalate to root.	Stateless Agents, public list prices, and zero cross-account privileges.
AutoMQ	A Kafka fork that keeps brokers and the JVM, recovers a failed broker by replaying its write-ahead log on EBS, and leaves KRaft and the control plane to you.	No brokers and no JVM. A failed Agent is simply replaced, and metadata is fully managed.
Aiven diskless	Diskless topics opted into one at a time, beside classic ones on brokers you still manage, with no transactions or compaction on diskless topics.	Every topic is diskless, with transactions and compacted topics fully supported.

See what WarpStream would cost for your workload



Talk to our team for a tailored TCO comparison, or start a cluster in minutes. New accounts get \$400 in free credits.

Talk to us

warpstream.com/contact-us

Start free

console.warpstream.com/signup

Run your numbers

warpstream.com/pricing

Read the full Executive Summary

warpstream.com/executive-summary

About the prices. Every price in this brief comes from WarpStream's public [pricing calculator](#), last refreshed October 19, 2026: AWS us-east-1 list prices, 4x compression, 3x consumer fan-out, 4,096 partitions, and the calculator's default Dev cluster tier. Apache Kafka® is a registered trademark of the Apache Software Foundation.