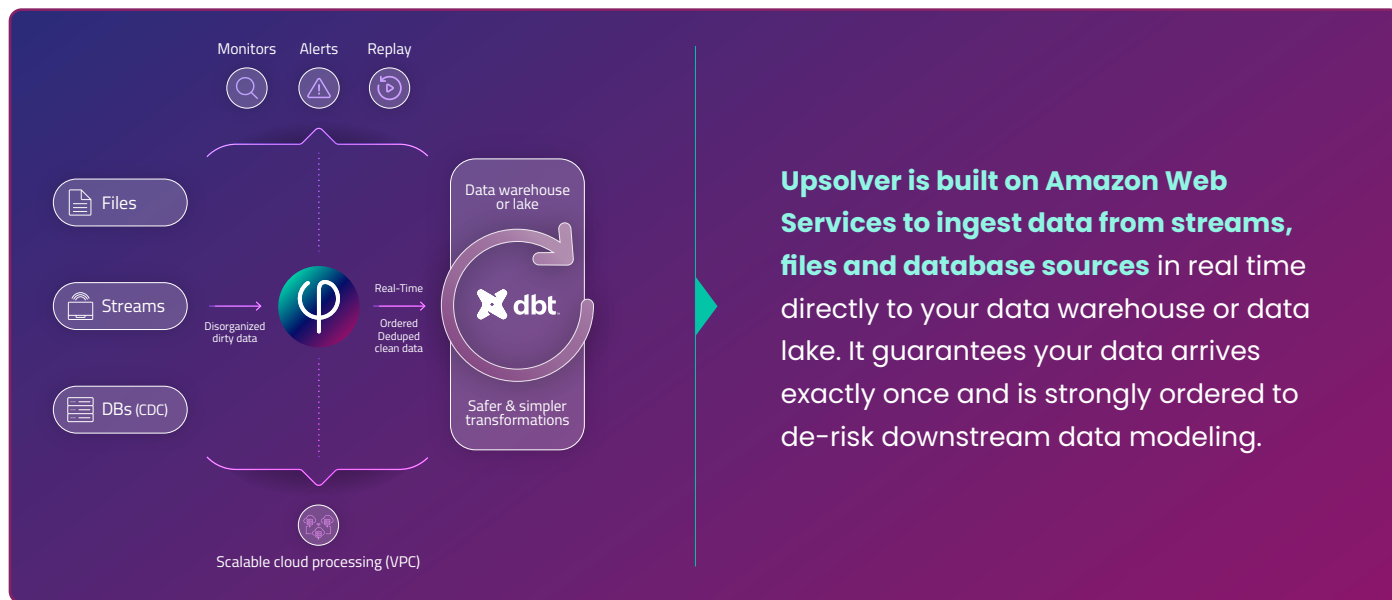
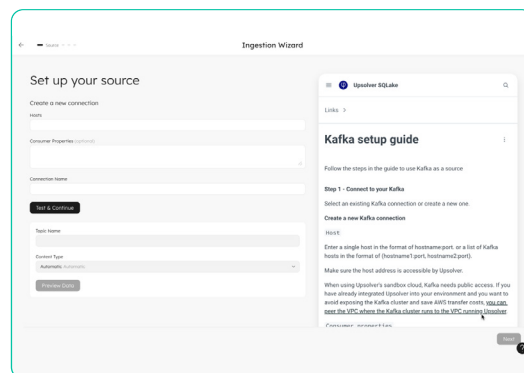


High quality, observable data ingestion

Stop bad data from polluting your data warehouse and lake

- **Reliable** - Guaranteed exactly-once, strongly-ordered data
- **Observable** - Detect and resolve data and schema issues upstream
- **Easy** - No-code ingestion for batch and streaming sources
- **Fresh** - Up-to-the-minute data in your warehouse and lake



Upsolver is built on Amazon Web Services to ingest data from streams, files and database sources in real time directly to your data warehouse or data lake. It guarantees your data arrives exactly once and is strongly ordered to de-risk downstream data modeling.

Loved by enterprises and startups across industries:



Reduce Modeling Risk in Your Data Warehouse

For the data warehouse, Upsolver Big Data Ingestion ensures that your incremental models work as intended, reducing risk of bad data making it into your warehouse, analytics and applications, and avoiding the unnecessary engineering effort of patching issues downstream. By simplifying your models, Upsolver removes obstacles to improving data freshness from days or hours to minutes.

- Derisk data modeling with high quality, observable data
- Eliminate unnecessary engineering effort to fix data quality issues downstream
- Improve freshness by eliminating lookback windows for late-arriving data

Power Transformations in Your Data Lake

For the data lake, Upsolver SQLake enables you to transform raw data into analytics-ready tables. You can perform continuous joins and aggregations at scale and leverage 150+ SQL functions including table functions, data functions, window operations, and transformation functions. SQLake slashes the time it takes to develop high-quality production data lake pipelines by eliminating manual orchestration and automating data management and optimization tasks that take the bulk of an engineer's time.

- Eliminate DAGs — Upsolver automates orchestration
- Automate compaction, vacuuming and metadata management
- Dynamically apply source schema changes downstream with smart SELECT*
- Define data transformations directly on fields with native array support
- Leverage INSERT and MERGE commands for table UPSERTs
- Time travel and instant replay to retroactively fix data issues or apply new logic

Get Started

Start deploying high quality, observable data pipelines today with a [free 30-day trial](#) of Upsolver or purchase Upsolver on the [AWS Marketplace](#). Get started with wizard-based, intuitive Big Data Ingestion, and when you are ready for advanced transformations, leverage dozens of out-of-the-box templates with SQLake.

About Upsolver

Upsolver is obsessed with removing friction from big data pipelines for data warehouses and data lakes. Upsolver Big Data Ingestion is a no-code tool for exactly once, strongly-ordered delivery of files, streams and database events at scale, while Upsolver SQLake provides a SQL-based IDE to perform real-time transformations in your AWS data lake. Headquartered in San Francisco with R&D in Tel Aviv, customers include Cox Automotive, IronSource, and SimilarWeb.