TDengine Datasheet

TDengine® is an AI-powered data historian that consists of of two products seamlessly integrated: high-performance time-series database TDengine TSDB and industrial data management platform TDengine IDMP. More than just a database, TDengine delivers everything a traditional historian provides and more: high-performance time-series storage, industrial data management, contextualization, analytics, events, visualization, and AI.

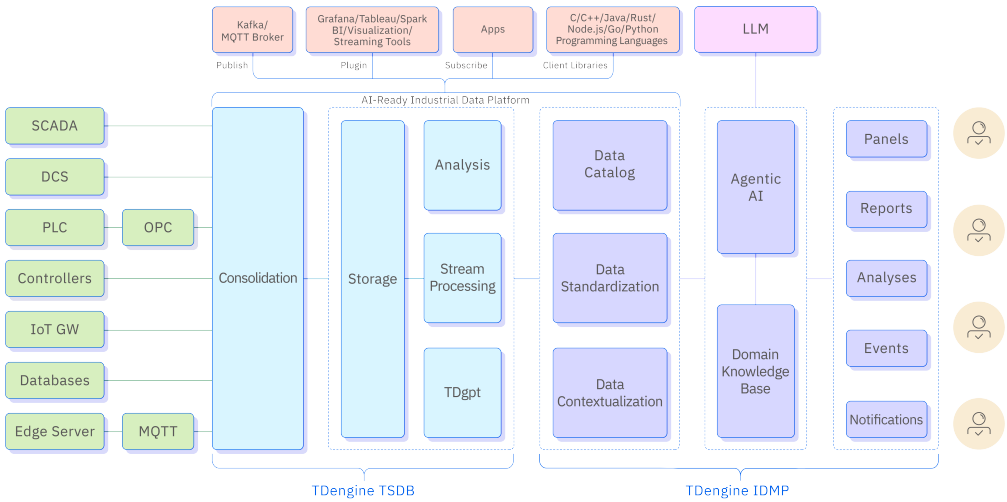

TDengine differentiates itself from other data historians in the following ways:

Proven Industrial-Grade Technology

Built for mission-critical operations, TDengine is trusted by over 1,000 global customers, with the largest production deployment managing more than 50 million tags continuously for over three years.

Radically More Affordable

TDengine delivers enterprise-grade historian capabilities at roughly one-third the cost of traditional systems like PI System with no extra charges for visualization, connectors, or add-ins, and no seat limits.

AI-Powered by Design

TDengine is built on an AI-ready platform supporting AI/ML for forecasting and anomaly detection as well as LLM-driven panel and analysis creation, natural-language operational insights, and assisted exploration of industrial data.

Open Ecosystem, No Lock-In

Designed for interoperability and long-term flexibility, TDengine runs across cloud and on-prem environments, integrates easily with existing data stacks, and keeps your data portable and under your control.

TDengine TSDB

TDengine TSDB plays a critical role in the TDengine historian package, providing high-performance data ingestion, storage, and compression. It is a purpose-built time-series database specifically designed for the characteristics of industrial operational data and includes the following key features:

Efficient Data Ingestion

With a built-in high-performance write engine that features automatic thread scheduling and batch processing, TDengine TSDB delivers high throughput and easily handles high-frequency data ingestion scenarios.

Optimized Query Engine

TDengine TSDB supports standard SQL with query capabilities optimized for time-series scenarios, including built-in time- and tag-based filtering, multi-level aggregation, downsampling, and interpolation.

Data Subscription

TDengine TSDB includes a built-in Kafka-like data subscription mechanism, allowing users to create topics based on databases, super tables, or SQL queries to stream incoming data in real time.

Read Caching

TDengine TSDB features a built-in read cache that stores the latest data for each table, enabling fast current-value queries out of the box and eliminating the need for external systems like Redis.

Stream Processing

TDengine TSDB comes with a built-in stream processing engine that supports SQL-based windowed aggregation and real-time computation, writing results directly to target tables with millisecond-level latency.

Edge–Cloud Synchronization

With automated replication of data between edge and cloud deployments, TDengine TSDB ensures cross-site data can be reliably stored and makes it easy to build distributed systems that combine local responsiveness with centralized insight.

Zero-Code Data Connectors

With built-in connectors for a wide variety of industrial sources — MQTT, Kafka, OPC, PI System, and more — TDengine delivers zero-code data ingestion and ETL in a centralized platform that acts as a single source of truth for your business.

AI Agent for Time-Series Analytics

TDengine TDgpt provides time-series data forecasting and anomaly detection, supporting AI/ML, including time-series foundation models and large language models, as well as traditional statistical algorithms, all in a single SQL statement.

TDengine IDMP

TDengine IDMP provides critical historian features including data modeling and cataloguing, contextualization, and standardization, as well as visualization, analytics, events, and notifications. It integrates with LLMs to bring the power of generative AI to the industrial sector, delivering an industrial AI agent that enables truly independent, real-time decision-making.

Industrial Data Modeling
: IDMP digitizes physical and logical entities through a tree hierarchy. By defining entity types, attributes, and relationships, it unifies fragmented sensor, device, and system data into a semantically rich digital twin model.

Data Contextualization
: IDMP enriches industrial data with business semantics by dynamically linking raw data to specific business entities. It transforms low-value operational data into understandable business objects.

Data Standardization
: IDMP builds a unified governance framework for industrial data by using entity templates to standardize attributes, analytics, and dashboard definitions for equipment and assets.

Industrial Data Visualization
: IDMP automatically identifies application scenarios and business priorities, recommending optimal panels and reports and eliminating the need for industry expertise and manual configuration.

Real-Time Analytics
: IDMP uses AI to dynamically recommend real-time analyses with millisecond-level responsiveness. Even without domain knowledge, you can access real-time insights to support business decisions.

Event Management
: IDMP automatically transforms real-time analysis results into actionable events with defined severity levels, pushing alerts to responsible personnel through an intelligent routing engine with optional confirmation.

Designed for Industrial-Scale Operations

TDengine is designed for high-frequency, large-scale industrial workloads where reliability, performance, and data integrity are critical. It is used across energy, manufacturing, utilities, and other industrial environments to support continuous data ingestion and analysis in production systems.

TDengine's distributed architecture supports high availability, replication, and fault isolation, helping systems remain available in the face of software maintenance, hardware failures, or network disruptions. This makes it suitable for mission-critical deployments that require consistent access to operational data.

TDengine runs on Linux and Windows and supports deployment on-premises, in private clouds, or in public cloud environments, enabling flexible infrastructure choices without sacrificing performance.

Secure and Governed Data Management

TDengine provides role-based access control, secure authentication, and encrypted data transmission to protect industrial data. Auditing and traceability capabilities help record data access and changes, supporting internal governance and compliance requirements in regulated environments.

TDengine complies with SOC 2 Type II as well as ISO 27001 and ISO 27017 standards.

Hybrid-Ready and Scalable Architecture

TDengine supports distributed and hybrid architectures, enabling data to be collected at the edge and synchronized with centralized systems for analysis and long-term storage. Built-in replication and flexible deployment models allow organizations to combine local responsiveness with centralized insight while evolving their data infrastructure over time.

Open and Extensible Ecosystem

TDengine is designed around an open ecosystem, with standard SQL, open APIs, and native support for modern industrial protocols. It integrates easily with existing tools and platforms, supports real-time data streaming, and enables data to be shared or exported across systems, helping organizations avoid vendor lock-in and retain full control over their data.