



Self Service Data Streams

The Internet of Things (IoT), along with web and mobile applications, present organizations with rich sets of real-time data. This data can yield enormous business value when effectively analyzed. Insights from this data enables businesses to better understand their customers' behavior, improve product offerings, and gain operational efficiencies. Capturing data from real-time data streams enables businesses to take action faster than ever before, leading to improved operational business outcomes.

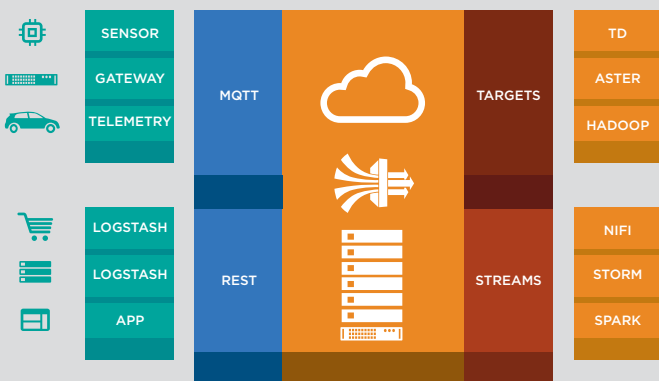
The challenge is reliably capturing scalable streaming data. In addition, both programmers and operations staff need easy self-service ways to set up and manage new data streams. Many organizations have realized benefits from building their own streaming data infrastructure. Later they realize that scalability—and adding new streams and targets—introduces development and maintenance hurdles. Leading data-driven organizations are moving beyond expensive, error-prone manual coding to self-service streaming products based on open source and modern software engineering.

Teradata Listener enables data administrators, developers, and data scientists across the organization to easily configure multiple data streams for reliable ingestion and writing to various data stores within the analytic ecosystem.

Benefits

- Easy-to-use, self-service delivery of data streams
- Ingestion and distribution decoupled for fast and flexible data streaming
- Accelerates time-to-business insights by reducing manual coding efforts
- Removes IT complexity, maintenance, and costs of custom-built systems
- Guaranteed, reliable data-ingestion to ensure zero data loss

Multiple Sources and Targets



Teradata Listener™

Teradata Listener™ is a self-service solution for ingesting and distributing fast-moving data streams throughout the analytic ecosystem. It's designed to be the primary ingestion framework for real-time data streams. Teradata Listener is designed to:

- Provide a self-service solution for programmers that minimizes the complexity of building and supporting data streams
- Operate as a central service to ingest and write hundreds of streams
- Reliably deliver data without loss
- Provide low latency ingestion for near real-time applications
- Run as a software-only solution, in the data center or in the cloud

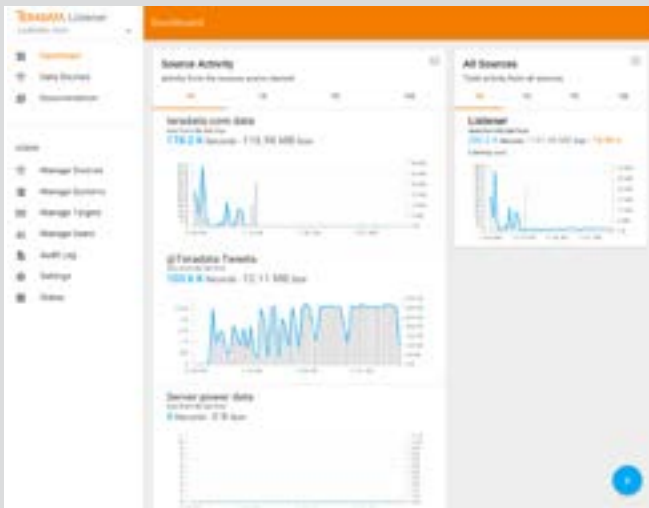
Ingest High-Volume, Continuous Data Streams

Sources of streaming data abound from web logs, email, sensors, social media, machine data, and others. Teradata Listener unifies the big data ingest process by collecting high volume data streams continuously from a variety of sources, and persisting them to one or more data stores.

Listener has capabilities to write to a variety of targets, such as the Teradata™ Integrated Data Warehouse, Hadoop, or complex event processors. Additionally, Listener works with both structured and unstructured data. Teradata Listener scales out horizontally to meet the growing demands of the business.

Teradata Listener's architecture enables the decoupling of incoming data streams with the outgoing distribution processes. It buffers the distribution output when the target systems are full, offline, or extra busy. Listener will resume data distribution later when the target system allows it—all without any manual intervention.

Listener Dashboard



Self Service

Teradata Listener simplifies data ingestion through a self-service dashboard, which can be accessed by developers, administrators, and data scientists. The intuitive dashboard makes configuring data sources and targets a simple ten-minute task, eliminating the need for custom programming. Users can easily add, remove, or manage sources and targets to create streaming data pipelines. With Teradata

Listener, there are no IT work requests to deploy a streaming pipeline—and no waiting for a programming team to build and test another interface in a home-grown streaming tool.

Listener's ingest services are accessible by popular interfaces:

- RESTful ingest: a universally accepted protocol for modern applications
- MQTT: the standard protocol for M2M and IoT connected devices where bandwidth or battery power are at a premium

Additionally, APIs provide flexibility for developers to configure and manage the data flowing through Listener. These APIs can be used to create custom visualizations, or integrate with data transformation engines.

Data Intelligence

Listener continuously monitors incoming data streams, gathering critical information and displaying it on the dashboards. Dashboard metrics help users understand what is coming in to, and going out of, Listener. Users can quickly discover when a stream is paused or when a target stops accepting data. This helps administrators understand throughput, and source or target system interruptions.

Enterprise Grade with High Reliability

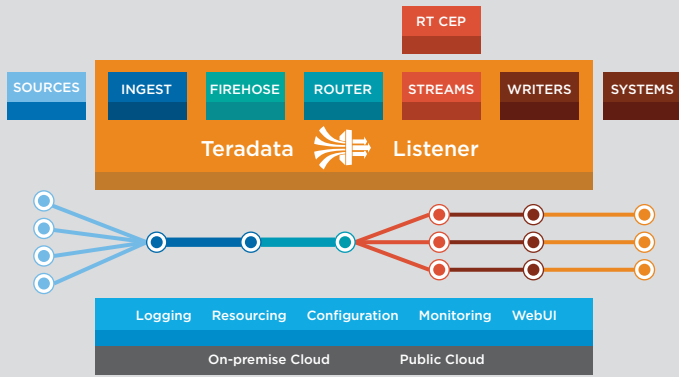
Teradata Listener is built to reliably deliver data without loss. When the target systems are full, and cannot accept any more data, its intelligent buffering mechanism persists data and delivers it later when the target becomes available. This ensures that no data is lost. Similarly, surges in source data are managed by the elastic ingestion micro-services.

Listener's built-in failover architecture deals with server failures automatically. If any service on Listener goes down for any reason, it is automatically re-started on the server, ensuring continuous availability of all services. It maintains redundant copies of the data throughout the data ingestion and distribution process in case the primary copy is lost.

Deployment Flexibility: On Premises or Cloud

As a software-only solution, Teradata Listener can run on premises commodity hardware or in a public cloud, such as Amazon Web Services (AWS).

Listener Data Flow



Modern Software Engineering

Teradata Listener combines the latest innovations and engineering from both open source software (OSS) and Teradata expertise. Listener is built leveraging proven OSS projects along with modern software engineering, taking full advantage of micro-service and container technologies. Built upon Docker® and Kubernetes® micro-services, Listener follows the Teradata Everywhere design goal ensuring it runs in popular environments.

Business Use Cases

Teradata Listener accelerates deployment of streaming data use cases within enterprises, such as click-stream analytics, recommendation engines, IoT sensor analytics, email compliance, network analytics, supply chain tracking, and cyber security. Any business process that needs real-time data delivery is applicable to Listener.

Listener Use Cases

Website Activity	IoT / Sensors	Email Compliance	Sales Spikes	Track Logistics
Track Social Media	Customer Satisfaction	Reservation Brokers	Ratings	Security Data

24/7 Support from Teradata

As a fully supported product, Teradata Listener enjoys an array of customer support advantages, including:

- Industry-certified regional and global Customer Care Centers
- Experienced service representatives, available 24/7
- End-to-end support
- Secure remote connectivity options to pinpoint problems and react rapidly

Why Teradata?

Teradata is the world's largest company focused on analytic data solutions through integrated data warehousing, big data analytics, and Hadoop. Only Teradata gives organizations the advantage to transform data across the organization into actionable insights—empowering business leaders to think boldly and act decisively to make the best decisions.

For More Information

To learn more Teradata and how we can help you with big data streaming, please contact your local Teradata representative, or visit teradata.com/listener.

10000 Innovation Drive, Dayton, OH 45342 Teradata.com

Listener is a trademark and Teradata and the Teradata logo are registered trademarks of Teradata Corporation and/or its affiliates in the U.S. and worldwide. Teradata continually improves products as new technologies and components become available. Teradata, therefore, reserves the right to change specifications without prior notice. All features, functions, and operations described herein may not be marketed in all parts of the world. Consult your Teradata representative or Teradata.com for more information.

Copyright © 2017 by Teradata Corporation All Rights Reserved. Produced in U.S.A.

09.17 EB9158



TERADATA.