How Does Teradata Process Customer Data?

Teradata wants its customers to feel confident about how Teradata processes the data they load into Teradata's Autonomous Knowledge Platforms ("Customer Data"), which might have artificial intelligence components ("AI"). The Autonomous Knowledge Platform ("Platform") may be deployed through Teradata Cloud, as a cloud-based service, or through Teradata Factory or other supported on-premises configurations. Teradata AI Studio is the unified workspace through which customers may access and use analytics, machine learning, model, vector, conversational AI, and agentic capabilities, including Tera and AgentStack. This document explains how Customer Data may be processed across those deployment models and capabilities.

I. What Is Customer Data?

In general, Teradata treats all information received from a customer as confidential. However, Teradata's agreements distinguish between two categories of information: (1) Confidential Information, which broadly covers information disclosed during the parties' business relationship, and (2) Customer Data, a subset of Confidential Information generally defined as all data uploaded by the customer to its Teradata instance. Confidential Information is protected by the confidentiality provisions of the governing agreement, while Customer Data is subject to additional contractual protections applicable to that data.

Although both Confidential Information and Customer Data may contain personal data, Teradata's role differs depending on the information involved. For personal data contained in Confidential Information that is not Customer Data, Teradata generally acts as a controller. For personal data contained in Customer Data, Teradata acts as a processor on behalf of its customers.

In limited circumstances, Teradata may collect limited personal data relating to customer personnel, such as names, email addresses, usernames, authentication credentials, or support-related information, in connection with account administration, authentication, and customer support. This information is Confidential Information, not Customer Data, and is governed accordingly.

When processing personal data collected for these purposes, Teradata provides notice through its Privacy Statement. For example, customer personnel are directed to the Privacy Statement when registering for and using Teradata's support systems, including ServiceNow, and when authenticating through PingID. The Privacy Statement explains how and why Teradata processes such personal information and is available at teradata.com/privacy. Teradata complies with applicable laws and implements appropriate technical and organizational measures to protect personal data processed in systems and operations under Teradata's control. Any transfers of such information are carried out in accordance with applicable privacy laws and Teradata's intragroup processing arrangements.

The remainder of this memorandum focuses on how Teradata processes Customer Data.

teradata.

II. How Does Teradata's Autonomous Knowledge Platform Process Customer Data?

A. Overview of the Autonomous Knowledge Platform

The Teradata Platform may be deployed through Teradata Cloud or Teradata Factory and includes multiple integrated capabilities that customers may choose to enable, depending on their business requirements.

Because the Platform is configurable, the way Customer Data is processed depends on the capabilities enabled by the customer and the customer's selected deployment model. The Platform may store, retrieve, analyze, transform, index, or search Customer Data. It also may generate analysis and outputs using Customer Data.

The customer determines what information is uploaded to, stored in, or otherwise made available through the Platform, and controls who is authorized (including any autonomous operations) to access its environment and the Customer Data contained within it. The Platform processes Customer Data within the customer's selected deployment environment, unless the customer enables a capability or connection that expressly involves processing outside that environment, or an exception described below applies. Except in the limited circumstances described below, Teradata does not access or view the contents of Customer Data as part of day-to-day platform operations.

B. Processing of Customer Data by the Platform

For Teradata Factory and other supported on-premises deployments, the Platform operates within the customer's on-premises or customer-designated environment. Customer Data remains within that environment except in the scenarios described below.

For Teradata Cloud deployments, the Platform is deployed using infrastructure provided by a third-party cloud service provider, or "CSP." Teradata provides the Platform software and related services, while the CSP provides the underlying storage, networking, and computing resources used for the customer's deployment environment. The CSP does not directly access or view Customer Data within the virtual instances Teradata deploys for the customer.

Depending on the CSP selected, Teradata's customers may choose from various geographic regions (including the US East, US West, the EU, etc.) to determine the country or region in which their Customer Data is hosted. Customer Data remains in that selected country/region with two exceptions that both apply to Teradata's legacy VantageCloud Lake product only.

In either deployment model, customers determine what Customer Data is made available to the Platform and control which users, applications, agents, tools, and services are authorized to interact with it. Teradata does not have visibility into or access to Customer Data merely because it is stored in or processed by the Autonomous Knowledge Platform.

C. Platform Components and Processing

The Platform may include AI Studio, Tera, AgentStack, Enterprise MCP, Enterprise Vector Store, Model Services, and infrastructure and compute services. Some Teradata offerings also incorporate third-party technology that is sold or otherwise made available as part of a Teradata product.

1. Platform Component Overview

AI Studio is the unified workspace through which customers build, deploy, manage, and govern analytics, artificial intelligence, machine learning, model, vector, and agent-based solutions. Customer Data remains within the customer's selected environment when AI Studio is used with capabilities deployed in that environment. Customer Data may leave the selected environment only when the customer selects, enables, or configures a capability or workflow (including any autonomous operations) that sends Customer Data to an external model, application, tool, data source, or service.

Tera is an agentic coworker for enterprise data work that allows authorized users to interact with data and analytical systems through natural-language requests. Certain Tera capabilities may use externally hosted models or integrated services; in those circumstances, the information required by the model or service is transmitted outside the customer's selected environment. The specific provider, processing location, and data elements transmitted depend on the applicable deployment and configuration.

AgentStack provides capabilities for building, deploying, operating, monitoring, and governing AI agents. The customer controls the agents it creates or enables, the data and tools made available to those agents, and the permissions and instructions governing their operation. An agent may cause Customer Data to leave the selected environment if the customer authorizes the agent to use an external model, application, tool, data source, or other service.

Enterprise MCP and similar Model Context Protocols provide a governed interface through which authorized agents and applications may use approved Teradata data, resources, and tools. They may also be used to connect an agent or application to an external system or service where the customer configures or authorizes such a connection.

Enterprise Vector Store enables customers to store, index, retrieve, and search vectorized information and unstructured content, including embeddings and associated metadata generated from Customer Data. Customer Data loaded into Enterprise Vector Store is stored and processed within the customer's selected environment.

Model Services supports model development, deployment, inference, evaluation, monitoring, and governance. Where a model is hosted within the customer's selected environment, Customer Data submitted to the model is processed within that environment. Where the customer selects or enables an externally hosted model, Customer Data submitted for inference leaves the customer's selected environment and is transmitted to the external model provider. In some cases, certain selected externally hosted models selected may not be available in the customer’s local deployment region. Where that occurs, Customer Data may be transferred to and processed outside of the customer’s local region. In these cases, to continue using any model-dependent services, the customer’s administrator may need to submit a change request and consent to the cross-region or global processing of Customer Data.

Infrastructure and Compute Services, including Active Compute and Elastic Compute, provide computing resources for workloads operating against Customer Data within the customer's selected environment and do not, by themselves, cause Customer Data to leave that environment.

2. Embedded Third-Party Components

Certain capabilities within the Platform incorporate third-party software that Teradata has licensed and embedded as part of its product. These embedded components are distinct from third-party services that customers may independently configure or connect to the platform using their own

credentials or API keys, which are addressed in Section V below. Because embedded components are incorporated into and distributed as part of Teradata's product, Teradata remains the data processor with respect to any Customer Data that those components process.

Embedded components fall into two categories. The first consists of components that execute entirely within the customer's deployment environment — Customer Data is processed locally and does not leave that environment because of the component's function. The second consists of components whose function involves transmitting Customer Data to the third-party's infrastructure or models for processing; for these, the third-party provider acts as a subprocessor on Teradata's behalf, and Teradata ensures through its agreements with those providers that appropriate data protection obligations apply. The second category is described in Section II.C.2(b) below.

a) Embedded Components That Process Customer Data Within the Customer's Environment

The following embedded components process Customer Data entirely within the customer's selected deployment environment. No Customer Data is transmitted to the third-party provider's infrastructure because of these components' standard operation. Any escalation to the component provider's engineering or support team is limited to configuration, platform, and software issues and does not require or provide the provider with access to Customer Data.

Unstructured.io provides embedded document ingestion and preprocessing capabilities within Enterprise Vector Store. These capabilities enable customers to ingest, parse, chunk, and enrich unstructured content — including documents, PDFs, spreadsheets, and images — and to generate embeddings from that content, within the Customer's deployment environment. Unstructured.io's functions are deployed and executed within Teradata-managed infrastructure in the Customer's environment; outputs are written directly to Enterprise Vector Store within the same environment. Customer Data does not leave the customer's selected environment because of Unstructured.io's standard embedded operation, across Teradata Cloud, on-premises, and air-gapped deployments. Where a customer configures the Unstructured.io integration to use externally hosted model providers by supplying API credentials, the relevant third-party provider is acting as a processor on the Customer's behalf; see Section V below.

Anonomatic provides privacy-enhancing capabilities embedded within the platform, including dynamic data masking, encryption, format-preserving encryption, and tokenization. Where enabled, these capabilities are applied to Customer Data within the customer's deployment environment. Anonomatic's software runs as a containerized service within the customer's environment; Customer Data does not leave the customer's selected environment because of Anonomatic's operation, and Anonomatic does not receive, access, or process Customer Data on its own infrastructure.

NVIDIA AI Components — certain Teradata hardware products, including Teradata Factory, incorporate NVIDIA AI components, including GPU-accelerated model serving infrastructure and pre-packaged inference models, as part of the product. For Teradata Factory and other supported on-premises deployments, these components are deployed and executed within the customer's on-premises environment; Customer Data submitted for inference, embedding generation, or vector search is processed within that environment and does not leave it.

Karini provides agent development, orchestration, evaluation, and governance capabilities. Karini software is deployed within the customer’s deployment environment. Customer Data processing occurs within such environment except to the extent a customer separately configures external

model providers or other third-party services. At that point, the relevant third-party provider is acting as a processor on the customer's behalf; see Section V below.

WisdomAI provides embedded agentic analytics capabilities, including natural-language querying, dashboard generation, contextual analytics, and analytics-agent workflows. WisdomAI processes Customer Data within the customer's selected deployment environment and uses federated access patterns designed to reason across structured and unstructured data without requiring duplication or movement of that data outside the environment.

b) Embedded Components That Involve Processing Outside the Customer's Environment

Pinecone provides vector database capabilities used with Enterprise Vector Store. With this service, customer's embeddings, vectors, associated metadata, and related search requests are transmitted to and processed within Pinecone's hosted service environment for vector storage, indexing, retrieval, and related database operations. These services are used to provide vector storage functionality. Thus, the use of Pinecone will mean that Customer Data may be processed or stored outside of the customer's selected Teradata deployment environment.

c) AI Features

To the extent the customer elects to use any of Teradata's artificial intelligence-related tools, models, and/or features (each, an "AI Feature"), additional details can be found on the Supplemental Notice for Teradata AI Features. The AI Features available to a particular customer, the third-party services used in connection with those AI Features, and the resulting interaction with Customer Data will vary depending on the applicable product, deployment model, customer configuration, and the third-party services selected by the customer.

III. When Could Teradata Personnel Process Customer Data?

Teradata provides customer support, maintenance, managed services, and consulting to support customers' Platform deployments.

Customer support and maintenance personnel generally do not access Customer Data in performing their services. To the extent the customer grants Teradata credentials for diagnosing faults or deploying fixes, patches, and upgrades, Teradata does not use those credentials to access the customer database. Multiple technical and organizational controls prevent the improper use of such credentials, including limits on who receives them, how long they remain valid, how they are shared, as well as the maintenance of access logs. Customers are encouraged to encrypt and/or pseudonymize their Customer Data, which does not prevent Teradata from conducting support and maintenance services.

An exception may arise where the customer support or maintenance team is asked to conduct a query analysis for performance or error diagnosis, or a log or dump analysis. In those cases, the queries or logs analyzed could contain Customer Data, including personal information. Customers are reminded not to load personal information directly into support tickets. Customers may also limit exposure by turning off query logging, or by scrubbing personal information from queries or crash dumps before providing them for analysis. Regardless, Teradata maintains strict security and organizational measures to protect Customer Data in the rare cases its support personnel may encounter it.

Where a customer separately contracts with Teradata for consulting or managed services, those services are typically performed via VPN sessions controlled by the customer and subject to confidentiality obligations. The customer determines the appropriate level of access for those services and may encrypt or pseudonymize Customer Data during the engagement. The same access controls described above apply to any credentials provided in connection with those services. When Teradata personnel access Customer Data remotely in connection with consulting or managed services, such access is typically through Teradata's Global Services Innovation Hubs ("GSIHs"). Customer Data remains stored in the customer's selected location, and the relevant team member accesses the data under the customer's control. Teradata operates all GSIHs under the same data protection principles and technical and organizational measures required by GDPR. Teradata's Autonomous Knowledge Platform and various GSIHs are certified in accordance with ISO 27001.

IV. Summary of Customer Data Processing Activity by Location

The chart below describes Teradata's various processing activities for customers and where they occur currently if the customer has not paid for support to be localized. Teradata may add or remove locations from time to time. Whenever Teradata adds a processing location, it will continue to follow all applicable data privacy laws. Please see Teradata's Transfer Impact Analysis Memorandum for more information regarding transfers subject to the GDPR.

Category of Processing Activity	**Current Location of Activity (as may be updated from time to time)**
CSP Storage and Compute	Region selected by customer
Cloud Operations	No processing for latest version of Teradata Cloud or Cloud on VM; transfer to US only for VantageCloud Lake as set out in Section II.B above
Customer Support & Maintenance for on-premises and Cloud	Australia, Austria, Belgium, Brazil, Canada, Chile, Czech Republic, Egypt, France, Germany, India, Indonesia, Ireland, Italy, Japan, South Korea, Malaysia, Mexico, Netherlands, Pakistan, Poland, Singapore, Spain, Taiwan, UK and the US
Consulting Services for on-premises and Cloud (performed under separate SOW)	Depends on customer location. Usually local country, or GSIHs in Czech Republic, India, Pakistan, and occasionally the US for remotely performed services
Managed Services for on-premises and Cloud	Depends on customer location, with India being the default for remotely performed services

V. Using Teradata Platforms with Third-Party Software

The Platform is designed to be interoperable with third-party services that customers select and configure. Where a customer elects to connect the platform to a third-party service — for example, by providing API credentials for an external model provider, configuring an agent to use an external tool, or enabling a third-party integration through the platform's interfaces — the third-party service provider acts as a processor on the customer's behalf and not as a subprocessor acting on Teradata's behalf. Data

flows associated with those customer-configured connections are not captured in the summary of Teradata's processing activities in Section IV above.

This may include, for example: large language model providers such as Azure OpenAI, AWS Bedrock, OpenAI, Anthropic, Google Vertex AI, or other open-weight models such as those on Hugging Face to which a customer routes inference requests using customer-held API credentials; third-party data sources or APIs that a customer authorizes its agents or workflows to access; and external applications or services connected to the platform through customer-configured integrations.

In these circumstances, the customer is responsible for ensuring that its use of the third-party service — including any transfer of Customer Data to that service — complies with applicable legal obligations and the terms governing the customer's relationship with that provider. Teradata does not control the third-party's processing practices, data retention, or subprocessing arrangements.

VI. Teradata's Response to Data Subject Requests and Government Inquiries

Teradata's obligations related to responding to data subject requests ("DSRs") and government inquiries are typically set out in its agreements with customers. Because Teradata does not have access to Customer Data as part of its day-to-day operations, it does not generally get involved in responding to DSRs on a customer's behalf. Instead, the Autonomous Knowledge Platform allows customers to retrieve, correct, delete, or restrict personal data themselves. In the unlikely event a customer is unable to independently address a DSR through the platform, Teradata will provide reasonable assistance upon the customer's written request or as required by law. If a DSR or other communication regarding the processing of personal data under an agreement is made directly to Teradata, Teradata will promptly inform the customer and advise the data subject to submit their request to the customer directly.

Teradata is also obligated under applicable standard contractual clauses to notify customers if it receives a request from a public authority for government access to customer personal data. Teradata will carefully assess any such request and will provide access only if clearly compelled to do so after a full evaluation. Any public authority must follow applicable legal procedures, and Teradata will refuse any request it determines to be unlawful. Customers will be informed of any such request upon receipt, unless notification is explicitly prohibited by law.

VII. Conclusion

This document is intended to provide customers with information about the approach Teradata has taken to managing the processing and transfer of personal data contained in Customer Data. It is not intended to constitute legal advice and should not be relied on as such. Customers should consider their own circumstances when evaluating their data protection obligations in connection with their use of the Autonomous Knowledge Platform. Please contact your Account Manager if you require assistance in assessing transfer impact or essential equivalence.

Current as of November 2026