



# Teradata VantageCloud Lake

## Cloud Service Description and AWS Addendum

---

This Cloud Service Description applies to VantageCloud Lake. In addition, an Addendum to this Cloud Service Description exists for each Cloud Service Provider's cloud platform on which VantageCloud Lake operates. If there are any conflicts between the terms of the cloud platform Addendum and the provisions of this Cloud Service Description, then such Addendum-specific terms shall apply and take precedence over the conflicting provisions of this Cloud Service Description.

**August 11, 2025**

## Table of Contents

1.	Teradata VantageCloud Lake .....	3
2.	VantageCloud Lake Packages .....	4
3.	Responsibilities .....	6
4.	Security .....	7
5.	Data Protection .....	11
6.	Data Storage and Management.....	14
7.	Elasticity Features.....	14
8.	Interfaces and Tools .....	14
9.	External Connector .....	15
10.	Analytics.....	15
11.	Education Services .....	16
12.	Support and Operational Services .....	16
13.	Pricing .....	21
14.	Pre-General Availability Offerings.....	22
15.	Additional Services (Sold Separately) .....	22
16.	Availability Service Level Agreement.....	26

## Appendix: Platform-Specific Addendum

## 1. Teradata VantageCloud Lake

Teradata VantageCloud Lake is a cloud native deployment option for the VantageCloud Analytics and Data Platform, delivered as a service, and powered by an enhanced version of the Teradata Analytics Database. The following list defines some of the terms used in the Cloud Service Description and addenda:

- 1.1 Analytics Database is the core engine of the connected multi-cloud data platform for enterprise analytics.
- 1.2 Primary Cluster serves as the hub of Customer Environment, containing an Analytics Database that consists of Banyan Network (BYNET) connected nodes. The Primary Cluster manages query planning and distribution of work across the Lake Environment. The Primary Cluster has access to high-performance persistent Block Storage to support operational queries and also uses Object Storage for data lake analytics.
- 1.3 Compute Cluster, a collection of BYNET-connected compute-only nodes, receives query work from the Primary Cluster for execution. Compute Clusters don't have any persistent storage attached; Compute Clusters use the Object Storage for data analytics. Compute Clusters enable elasticity and an extensible compute VantageCloud Lake architecture. There are three types of Compute Clusters: Standard, Analytic, and Analytic GPU.
  - a) Standard Compute Clusters are suitable for a wide breadth of enterprise decision support application queries and in-database data engineering and analytics.
  - b) Analytic Compute Clusters are suitable for memory-intensive analytic queries utilizing open-source language and analytics to run in an open analytics framework.
  - c) Analytic GPU Compute Clusters add Graphics Processing Unit (GPU) compute to Analytic Compute Clusters. Analytic GPU Compute Clusters are ideal for running Deep Learning and Large Language Model inference tasks and are suitable for running traditional Machine Learning Model training and inference use cases.
- 1.4 Access Module Processor (AMP) is a set of virtual processing units (vprocs) on the Primary Cluster and the Compute Cluster responsible for managing the interactions with the subsystems.
- 1.5 Block Storage is the persistent block-level data storage offered by a Cloud Service Provider and is used by the AMP to store data on the Primary Cluster.
- 1.6 Object Storage is the persistent object-level data storage offered by a Cloud Service Provider and is used by the AMP to store data on the shared storage tier for the Primary Cluster and the Compute Cluster.
- 1.7 Organization refers to the billing entity for a Customer.
- 1.8 Environment refers to a single instance of a VantageCloud Lake system deployment. Each Environment consists of one Primary Cluster with Block Storage, and optional Compute Cluster(s) with Object Storage that Customer may choose to provision.
- 1.9 Customer Data refers to all data uploaded by Customer to the Teradata Cloud Service through the VantageCloud Platform.

## 2. VantageCloud Lake Packages

The following VantageCloud Lake Packages table lists the features and services available in each package.

VantageCloud Lake Packages			
Features And Services	Standard	Lake	Lake+
<b>Compute Resources</b> (see Section 1 for details)			
Primary Clusters	2-9 nodes	2-128 nodes	2-128 nodes
Compute Clusters	No autoscaling	•	•
<b>Availability</b> (see Section 16 for details)			
Availability Service Level Agreement (SLA)	99.9%	99.9%	99.9%
<b>Storage Types and Related Fees</b> (see Section 13 for details)			
Block Storage	•	•	•
Object Storage	•	•	•
Standard Backup Storage	•	•	•
Data Transfer (Ingress and/or Egress)	•	•	•
<b>Elasticity Features</b> (see Section 7 for details)			
Primary Cluster Scale Up/Down		•	•
Block Storage Expansion		•	•
<b>Network Connectivity</b> (see Section 4 and platform-specific addendum for details)			
Network policies to control Environment access by Internet Protocol (IP) addresses	•	•	•
Internet connectivity	•	•	•
Private connectivity to VantageCloud Lake Service		•	•
<b>Security</b> (see Section 4 for details)			
Authentication Methods: Federated Authentication/SSO, Standard Teradata Authentication (TD2), LDAP Authentication	•	•	•
Data at-Rest Encryption	Platform Managed Encryption Keys	Platform & Customer Managed Encryption Keys	Platform & Customer Managed Encryption Keys
Data in-Transit Encryption	TLS 1.2 Encryption	TLS 1.2 Encryption with Customer's Certificate Signing Request (CSR)	TLS 1.2 Encryption with Customer's Certificate Signing Request (CSR)
Supported Bring Your Own License 3 <sup>rd</sup> party security software		•	•
<b>Compliance Certifications / Audit Reports</b> (see Section 4 for details)			
ISO 27001:2022, ISO 27017:2015, PCI DSS 4.0, SOC1 Type 2, SOC2 Type 2	•	•	•
HIPAA Type 2, HITRUST r2, TX-RAMP Level 2		•	•
<b>Data Storage and Management</b> (see Section 6 for details)			
Java Open Table Format (OTF)	•	•	•
Native Object Store (NOS) read/write	•	•	•
<b>Connectors</b> (see Section 9 and platform-specific addendum for details)			
QueryGrid (Teradata connectors)		•	•
QueryGrid (external connectors)		○	○

VantageCloud Lake Packages			
Features And Services	Standard	Lake	Lake+
<b>Analytics</b> (see Section 10 for details)			
ClearScope Analytics	●	●	●
Industry Data Model or Industry Analytic Schema			●
<b>Data Protection Plan</b> (see Section 5 for details)			
Same Environment Full Backup & Full Restore	●	●	●
Same Environment Partial Backup & Partial Restore		●	●
<b>Interfaces and Tools</b> (see Section 8 for details)			
VantageCloud Lake Console	●	●	●
<b>Education Services</b> (see Section 11 for details)			
Premier Education Package (eLearning)	●	●	●
Instructor-led classes	○	○	○
<b>Support Services</b> (see Sections 12 and 15 for details)			
Premier Cloud Support	●	●	●
Priority Service		○	●
Platform Services Management (PSM)		○	○
Critical Platform Services Management (CPSM)		○	○
<b>Operational Services</b> (see Section 12 for details)			
Updates and Maintenance	●	●	●
Service Disruption Recovery	●	●	●
Service Advisory	○	●	●
Backup and Restoration Assessment	○	○	●
Cost Optimization Assessment	○	○	●
Data Management and Workload Optimization Assessment	○	○	●

● Available

○ Optional Paid Add-On

### 3. Responsibilities

Teradata, Customer, and the Cloud Service Provider all have responsibilities for managing Customer's VantageCloud system as summarized in the VantageCloud Lake System Management table.

<b>VantageCloud Lake System Management</b>	
<b>Responsibility</b>	<b>Responsible Party</b>
Hardware	Cloud Service Provider
Data Center / Hosting	Cloud Service Provider
Initial Data Migration	Customer <i>Note: Customer may hire Teradata to perform certain responsibilities under a Teradata Consulting engagement (sold separately).</i>
Environment Availability Monitoring – Operating System (OS) and Teradata software	Teradata
Software Patching and/or Upgrading	Teradata <i>Note: Third-Party Tools patching and/or upgrades require a Teradata Consulting engagement (sold separately).</i>
Cloud Software Support	Teradata
Cloud Platform Support	Cloud Service Provider
Database Administration / Operations <ul style="list-style-type: none"> <li>• Database access security monitoring – managing VantageCloud Lake security roles, passwords, and access rights</li> <li>• Maintaining VantageCloud Lake structures, space, users, and jobs</li> <li>• Monitoring alerts, queries, access locks, and database performance</li> <li>• Analyzing database activity and priority of jobs/queries to identify performance tuning opportunities</li> <li>• Managing consumption and query performance</li> </ul>	Customer  <i>Note: Customer may hire Teradata to perform certain responsibilities under a Teradata Consulting engagement (sold separately).</i>
Operating System (OS) Administration / Operations <ul style="list-style-type: none"> <li>• OS security monitoring</li> <li>• Volume encryption</li> <li>• OS user administration for Teradata personnel</li> </ul>	Teradata
Network Administration / Operations	Teradata – <ul style="list-style-type: none"> <li>• Restrictions / filtering of incoming traffic to VantageCloud Lake Environment</li> </ul> Customer – <ul style="list-style-type: none"> <li>• Configuring Customer's networks for all applications to connect to the VantageCloud Lake Environment on customer-defined ports and protocols.</li> <li>• Configuring Customer's infrastructure components and cloud services including, but not limited to, gateways, endpoints, access profiles, firewalls, DNS &amp; proxy servers.</li> </ul>

VantageCloud Lake System Management	
Responsibility	Responsible Party
Cloud Environment Administration / Operations <ul style="list-style-type: none"> <li>Security monitoring of the VantageCloud Lake Environment (excluding database administration)</li> <li>Cloud Environment access management for Teradata personnel</li> </ul>	Teradata
Backup Archive Restore (BAR)  <i>Note: See details in Section 5 (Data Protection).</i>	Teradata – <ul style="list-style-type: none"> <li>Backup Support and Full Restore requests per defined scope</li> </ul> Customer – <ul style="list-style-type: none"> <li>Sets up, changes, and manages the Backup Plans and execute Partial Restore jobs using self-service capability utilizing VantageCloud Lake Console</li> <li>Defines Backup Lifecycle and Storage Policy Management requirements</li> <li>Resolves issues causing backup warnings and exceptions</li> <li>Submits Requests for a one time Full Restore via Change Request process</li> </ul>
Setting and managing DBC password	Customer
User Management <ul style="list-style-type: none"> <li>Adding and maintaining users in the VantageCloud Lake Console</li> <li>Configuring and managing user logs</li> </ul>	Customer

#### 4. Security

Teradata VantageCloud Lake Security uses a shared responsibility model. Teradata is responsible for security up to the OS layer, whereas Customer is responsible for security at the application, Customer Data, and user access/identity layers. Customer is also responsible for their own user management and data authorization policies. Customer should also note that Teradata’s security practices described in this service description are standard across VantageCloud Lake service and any customizations are not supported.

- 4.1 Customer’s Responsibility for Security Controls. Teradata requires that Customer implement reasonable security protections for the elements of the Teradata Cloud Service operated by Customer. This includes, without limitation, (a) analyzing the Center for Internet Security’s Critical Security Controls and implementing all applicable controls in keeping with information security industry best practices; and (b) implementing Multi-Factor Authentication (MFA) for VantageCloud user credentials controlled by Customer.
- 4.2 Security Standards. Teradata uses independent, industry-recognized auditors to annually audit the VantageCloud service. Compliance certifications and independent audit reports for each Package are listed in Section 2 (VantageCloud Lake Packages).
- 4.3 Additional Audits or Reviews. Any audits or reviews other than those listed under Compliance Certifications / Audit Reports in Section 2 (VantageCloud Lake Packages) will only be allowed if expressly permitted in the Teradata Master Cloud Service Agreement.
- 4.4 Secure Authentication. Teradata recommends the use of Federated Authentication / Single Sign-on (SSO) as the primary authentication method for VantageCloud Lake.

#### 4.5 Authentication Options

- a) Federated Authentication/SSO. Teradata supports Federated Authentication/SSO. This capability is specific to Customers that have an Identity Provider (IdP) and enables VantageCloud Lake users to log on to the VantageCloud Lake system and supported applications with a single set of their corporate credentials and move seamlessly between applications. These applications include Teradata Studio, Viewpoint, and certain third-party applications. Customer's IdP can be integrated via the VantageCloud Lake Console. This Federated Authentication/SSO offers the following features:
  - i. Bring Your Own IdP (BYOIDP) - Customer can bring their own IdP to integrate with the VantageCloud Lake native IdP using the self-service capability on the VantageCloud Lake Console. In this digital authentication approach, user identity is managed by Customer's IdP.
  - ii. Bring Your Own Multi-Factor Authentication (BYOMFA) - VantageCloud Lake customer business users enrolled in MFA with their IdP can engage in a multi-level authentication flow.
  - iii. Multiple Identity Provider Support - VantageCloud Lake's native IdP can be integrated with multiple instances of Customer's IdP. This allows Customers to group specific environments within an organization, restrict user access to those specific groups and help isolate activities of different departments.
  - iv. Federation Standard Protocols - VantageCloud Lake supports both SAML and OpenID Connect (OIDC) protocols. SAML is an XML-based open-standard for exchanging authentication and authorization data between applications. OIDC uses JSON Web Tokens (JWT) obtained through a standard OAuth 2.0 flow.
- b) Standard Teradata Authentication (TD2). VantageCloud Lake supports TD2 mechanism as an authentication method for the database.
- c) Lightweight Directory Access Protocol (LDAP) Authentication. VantageCloud Lake supports LDAP directory services integration over SSL/TLS (LDAPS). LDAP directory services integration for authentication requires a Teradata Consulting engagement for an additional fee. LDAP authentication is supported for direct database access from thick client applications and is extended for scheduled ETL jobs as well.

**Note:** *The VantageCloud Lake Console does not support authentication via LDAP.*

4.6 Network Policies. To control Environment network access, Customer can use network policies to enable/disable IP addresses.

#### 4.7 Encryption

- a) Data-at-rest Encryption. All data-at-rest is encrypted using National Institute of Standards and Technology (NIST) approved encryption standards. Teradata generates separate encryption keys for each tenant and automates the lifecycle management of those keys. The Separation of Duties principle is enforced so only Customer can decrypt their data; Teradata employees do not have rights to read any Customer Data. In addition, data-at-rest is stored in encrypted volumes in the Cloud Service Provider's storage.
- b) Data-in-transit Encryption. Data is encrypted in transit between the VantageCloud Lake Environment and connecting client sessions. Data is also secure from public exposure as it traverses network segments in the Cloud Service Provider's infrastructure when using a platform-specific private network link as a connectivity option. TLS 1.2 is supported by default in VantageCloud Lake. Customer may have the option to install their own CSR certificate. To do so, Customer must open a Change Request for Teradata to assist with the installation.

- 4.8 Internal service accounts. Teradata uses internal service accounts (e.g., TDaaS\_####) to provide many of the services mentioned in this service description. These service accounts automate system events and processes only. Customer does not have access to these internal service accounts.
- 4.9 Vulnerability Management. Teradata performs regular scans of the Environment and code to identify vulnerabilities in the software and operating Environments. Through a combination of static application security analysis and both network and application-level vulnerability assessments, Teradata promptly mitigates and remediates any vulnerabilities detected pursuant to applicable Teradata policies.
- 4.10 Access Control
- a) Customer controls access to their Customer Data. Teradata treats all Customer Data as sensitive and can only access Customer Data when Customer grants specific permissions.
  - b) As part of its Access Protection Policy, Teradata requires its personnel to complete training and to sign security agreements before receiving Environment access.
  - c) For the passwords that Teradata controls, Teradata enforces password complexity, stores and transmits only encrypted password representations, and sets minimum and maximum lifetime restrictions on passwords. Teradata cannot view Customer Data without Customer explicitly granting access—and we never transfer Customer Data between countries unless instructed to do so by Customer or unless required to perform Teradata’s agreed obligations.
  - d) If required for support or troubleshooting access, Customer must provide database and application-level credentials to Teradata personnel via Secure Password Vault or other approved / established mechanism.
- 4.11 Network Security Layers. Teradata employs multiple layers of network security as follows:
- a) At the perimeter, Teradata uses the Cloud Service Provider’s (CSP’s) Application Programming Interface (API) gateway, content delivery network, and distributed denial-of-service protection service to secure inbound traffic.
  - b) Within the perimeter, Teradata employs Virtual Private Clouds (VPCs) to isolate resources from each other.
  - c) Teradata uses network access control lists and security groups to add an additional layer of security controls.
  - d) Teradata continually monitors the network security configuration against the baseline using a combination of in-house security scanning services and CSP tools.
- 4.12 Access Rights and Security Re-Approval Process. Teradata establishes access rights and a security re-approval process to regularly review credentials and access rights for Teradata personnel consisting of:
- a) Managing access to the Environment from remote devices by Teradata via Virtual Private Network (VPN) and Multi-Factor Authentication (MFA). Teradata logs all access to Environments into immutable storage where they are correlated and analyzed.
  - b) Reviewing and approving account management actions
  - c) Monitoring account management operations for unauthorized actions
  - d) Disabling inactive accounts after 90 days
  - e) Disabling VantageCloud Lake accounts issued to a Teradata user after that Teradata user is transferred or terminated

- f) Modifying role-based access when a Teradata user's Environment usage or need-to-know requirement changes

#### 4.13 Environment Monitoring and Telemetry

- a) Teradata requires collection of telemetry data with any selected VantageCloud Lake package.
  - i. The telemetry measurements and data include: ResUsage, DBQL, obfuscated query logs, event logs, and infrastructure resource utilization (e.g., uptime of provisioned resources, storage used, data transfer).
  - ii. Teradata uses this information for support and operational services, which include:
    - a. detecting, troubleshooting, and resolving issues
    - b. gauging, analyzing, and optimizing performance and functionality
    - c. providing service packages, feedback, and recommendations related to the services Customer receives from Teradata
    - d. billing and invoicing Customer
    - e. determining how Customer uses our software, accesses functions, and performs analytics
- b) To deliver the assessments detailed in Section 12 (Support and Operational Services), Teradata requires the collection of additional telemetry data, which include:
  - i. Platform capacity and database utilization metrics
  - ii. Aggregate workload processing metrics
  - iii. Object usage metadata, including database, table, and column names
  - iv. Obfuscated query processing information
- c) The telemetry data that Teradata collects does not require Customer Data but may include database, table, and column names. It is incumbent on Customer to use best practices when configuring their VantageCloud system to not include underlying Customer Data in those names.
- d) Teradata uses reporting features contained within the Cloud Service for analyzing cloud system workload and usage and for determining fees and capacity requirements. Such reports will not require use or access to Customer Data. Customer can monitor consumption through Teradata's standard reporting mechanisms.
- e) If Customer limits or interferes with Teradata collecting telemetry data or using reporting features described above, the possible implications are:
  - i. Teradata may be unable to provide the security, operations, availability, billing, and other services that make up the Teradata Cloud Service as described herein;
  - ii. Customer may experience lower system availability;
  - iii. Teradata may be unable to meet Support SLAs for responsiveness and recovery times;
  - iv. Teradata may be unable to honor support expectations; and/or
  - v. Termination of database access.

## 5. Data Protection

- 5.1 Data Protection Plans (Backup and Recovery). Customer can configure and manage Data Protection Plans using the self-service capability on the VantageCloud Lake Console. A Data Protection Plan describes the schedule, frequency, and retention policy for data backups using Standard Backups.
- “Data Protection Plan” refers to a defined set of rules used to perform Environment backups according to a predetermined schedule and lifecycle (retention) policy.
  - “Backup” refers to the execution of the Data Protection Plan.
  - A “Backup Job” is the term used to define the process of performing the Backup per the Data Protection Plan.
  - A “Restore Point” is one instance of a successful Backup.
- 5.2 Backup Pricing Options. Section 13.2 (Storage and Data Transfer Pricing) of this Cloud Service Description describes the available Backup Pricing Options.
- 5.3 Configuring the Data Protection Plan. The VantageCloud Lake Console enables Customer to configure Data Protection Plans. Using self-service capabilities in the VantageCloud Lake Console, Customer can configure a Data Protection Plan by providing following inputs:
- Data Protection Plan name
  - Schedule (i.e., date, time, frequency)
  - Retention period (i.e., how many Restore Points to keep)
- 5.4 Backup Lifecycle and Storage Policy Management. Using the VantageCloud Lake Console self-service capability, Customer can make changes and modifications to the Data Protection Plans and perform on-demand backups themselves.
- Backup Jobs are run as an automated process and will continue per the requested policy unless modified or cancelled.
  - Backup Jobs can be run daily, weekly and/or monthly.
  - Retained restore points that have reached the end of the defined retention period will be deleted automatically.
  - Backups are encrypted at rest using cloud provider-managed encryption keys. Optionally, Customer may choose to use customer-managed encryption keys.
  - Backups are retained in the same cloud provider region where the VantageCloud Lake instance is located.
  - If a Customer creates a backup policy, Customer is solely responsible for managing that backup policy.
  - Through the VantageCloud Lake Console and/or Support Portal, Customer can perform the BAR services shown in the BAR Services and Availability table below.

BAR Services and Availability		
BAR Services	Availability	
	Support Portal	VantageCloud Lake Console
Enable daily, weekly, or monthly backups	•	•
Disable scheduled Backup Jobs	•	•
View the status of upcoming and previously executed Backup Jobs		•
Sign up for weekly backup status reports by email		•
Enable auto-restart to re-run a failed Backup Job one time		•
Re-run Backup Jobs		•
Cancel in-progress Backup Jobs		•
Change default and custom backup schedules		•
Edit backup retention policies <i>Note: The Cloud Service Description (CSD) Addendum for the specific cloud platform provides the backup retention terms.</i>		•
Configure email alerts indicating the Backup Job's status (i.e., start, failure, completed successfully, completed with exception(s), completed with warning(s), and/or schedule changes)		•
Configure Backup Job priorities		•
View backup storage utilization over time		•
Perform same-Environment Full Restores		•
Perform same-Environment Partial Restores (available for select Packages)		•
Request cross-Environment Full Restores (available for select Packages)	•	

- h) Customer may set completed backups as “do not delete” for an additional cost due to increased storage usage.
- i) Customer Data stored in OFS uses native cloud object store resiliency (i.e., three copies of data are saved in different zones) to protect against accidental data loss. Versioning retains all changed objects for a rolling 30-day period. Customer can request a reversion to a previous version (i.e., a roll-back or restore) within the 30-day window by submitting a Change Request, including the object name, date to revert to, and date/time to perform the restore.

### 5.5 Backup Support

- a) Backup Jobs will initiate a Severity 3 Case Level on job “failure” and will include the status as "Failed." Failure is defined as unable to complete a Backup to its end point.

- b) Teradata is responsible for investigating Backup Job failures and for resolving Backup Job failures that are under Teradata's control. Depending on the job failure error, Customer may be responsible for fixing the root cause of the job failure.
- c) Standard Backup Jobs completed with exceptions and/or warnings should not be considered as failed or unusable. To ensure Backup Jobs can complete within a reasonable timeframe, objects will be skipped (i.e., excluded) when multiple retries are unsuccessful or if the object is unavailable to backup. Those situations will be reported in the VantageCloud Lake Console as an exception or warning. If a particular Backup Job continues to skip objects, Customer should review the job schedule to see if the Backup Job can be rescheduled for a day/time when the affected object(s) will be available.
- d) If needed, Customer can submit a Case in the Support Portal to request more information about any exceptions or warnings.

5.6 Restores. Customer can perform Partial Restores (available with select Packages) or Full Restores to the same Environment via the VantageCloud Lake Console. Cross-Environment Full and Partial Restores are also available with select Packages. See details in the Cloud Service Provider addendum.

- a) For Cross-Environment Full Restores requested from a Change Request:
  - i. A Full Restore is performed manually by Teradata as a one-time operation at the date and time agreed upon in the Change Request.
  - ii. Customer must provide the following information in the Change Request:
    - a. Standard Backup name and date (i.e., point-in-time)
    - b. VantageCloud Lake Environment names to restore to/from
    - c. Restore maintenance window (i.e., date and time the restore should be implemented)
  - iii. Change Requests for Restores will be designated as "Normal" unless "Emergency" S1 situation requirements are met.
- b) For all Restores:
  - i. A suitable Backup Job must be available to restore from the point-in-time that is requested.
  - ii. Restores can only be applied to the same Environment from which the Backup Job was generated.
  - iii. Customer can request a reversion to a previous version (i.e., a roll-back or restore) for data objects in OFS to a previous point in time within the 30-day window via submission of a Change Request including the object name, date to revert to, and time/date to perform the restore.

### 5.7 Additional Customer Responsibilities and Terms

- a) Customer is responsible for scheduling, monitoring, and correcting conflicts (e.g., Data Integration [ETL], data loads, external queries, jobs/flows, queries, etc.).
- b) Customer must keep Teradata QueryService enabled to take advantage of self-service Standard Backup features.
- c) Customer can trigger Backup Jobs using available API calls. Customer is solely responsible for the use of APIs, including the development of scripts, configuration of 3rd party scheduling tools (e.g., Control-M, Cron Jobs), integration with external event management systems, and ensuring there are no conflicts between API-triggered backups and scheduled Standard and/or Snapshot Backups.

- d) Customer is responsible for ensuring that an appropriate Backup Restore Point exists to meet Customer's desired recovery point objective. These Backups will be used by Teradata to restore data and VantageCloud Lake services in case of an Environment outage or hardware / software / OS induced data corruption that has rendered the database unusable and requires full data restoration.

**Note:** Full Environment restores are performed only when needed to recover from an Environment outage or data loss/corruption event.

- e) Customer should not schedule concurrent or overlapping Backup Jobs due to possible performance degradation or prolonged job completion times.
- f) Customer is responsible for ensuring that retained Backups are capable of being restored.

## 6. Data Storage and Management

This section describes the VantageCloud Lake features to access data in Object Storage. See Section 2 (VantageCloud Lake Packages) for availability of features.

- 6.1 Open Table Format (OTF) are open-source standard table formats designed for working with very large datasets in a performant manner. VantageCloud Lake supports Apache Iceberg and Delta Lake OTFs, with REST Open Catalog integration on all major CSPs and for on-premise support of hybrid deployments. VantageCloud Lake supports only Externally Managed OTF (where Customer owns and manages their catalog and data storage).
- 6.2 Native Object Store (NOS) is a Teradata Vantage capability to read from and write to S3-compatible object storage.

## 7. Elasticity Features

This section describes the elasticity features that are available in select VantageCloud Lake packages. See Section 2 (VantageCloud Lake Packages) for availability of features.

VantageCloud Lake Package Elasticity Features and Descriptions	
Features	Descriptions
Primary Cluster Scale Up/Down	<ul style="list-style-type: none"> <li>• Change Primary Cluster instance size by up to 2 sizes up or down from the initial deployment size without affecting persistent block storage via a Service Request in the Support Portal, unless otherwise specified in the platform-specific addendum</li> <li>• Scale Up/Down will increase/decrease the rate of Primary Cluster Unit consumption</li> </ul>
Block Storage Expansion	<ul style="list-style-type: none"> <li>• Expand block storage after initial deployment at the Primary Cluster Environment level via a Service Request in the Support Portal</li> <li>• Block storage on the Primary Cluster cannot be reduced once increased</li> </ul>

## 8. Interfaces and Tools

The VantageCloud Lake Console is a web-based user interface that enables Customer to create and manage the VantageCloud Lake components in a self-service manner.

- 8.1 Access to VantageCloud Lake Console. Teradata will enable access to the VantageCloud Lake Console for Customer's designated Organization Administrator.
  - a) Teradata will create the Organization Administrator username.

- b) Customer is responsible for setting the password at first use.
- 8.2 Environment Provisioning. The Organization Administrator can provision a VantageCloud Lake Environment or add other users as Organization Administrators who can provision an Environment.
- 8.3 Operations. Operational capabilities include, but are not limited to, the following:
- a) Creating and maintaining the VantageCloud Lake Environments
  - b) Creating and maintaining compute groups and compute profiles for workloads
  - c) Query Management with the ability to view running queries and the query history
  - d) Tracking and monitoring of compute and storage consumption
- 8.4 User and Access Management. Customer's Organization Administrator is responsible for the initial onboarding of all VantageCloud Lake users and the ongoing assignment/maintenance of users and roles.
- a) One Organization Administrator is required, although Teradata recommends that Customer consider adding multiple Organization Administrators. As the Organization Administrator, this user can add and manage additional users as designated Organization Administrators.
  - b) Organization Administrators can:
    - i. Provision and manage Environments
    - ii. Manage Organization Administrators, roles, and users
    - iii. SQL Commands using a SQL editor can also be used to manage users and Compute Groups.
- 8.5 Editor in VantageCloud Lake Console. Customer can use the web-based SQL Editor to easily connect to provisioned databases, without installing anything on a local computer, and then submit database queries and run a query on any Compute Group accessible to the user.

## 9. External Connector

- 9.1 Teradata QueryGrid delivers data access, processing, and movement across Teradata Environments. Details for the VantageCloud Lake connectivity specific to each platform is detailed in the platform-specific addendum. QueryGrid is available in select VantageCloud Lake packages, as noted in Section 2.

## 10. Analytics

- 10.1 Teradata ClearScape Analytics™. VantageCloud Lake includes ClearScape Analytics, which consists of comprehensive in-database data manipulation, modeling and scoring functions, open and connected third-party integrations/APIs, as well as features enabling full-scale analytics activation and operationalization. ClearScape Analytics is built on the following analytic functions and frameworks:
- a) In-Database Analytics includes a wide variety of analytic functions for data-preparation, cleaning, transformation, feature engineering, model training, and model scoring.
  - b) Bring Your Own Analytics refers to an extensible framework that enables execution of open-source code/libraries and import of pre-trained models for scoring data in Teradata Vantage. In addition, Teradata provides client analytic libraries for Python (teradataml) and R (tdplyr) languages.

- c) Partner Integration connects third-party partner analytics to the data which may allow parallelizing computations in the Teradata Analytics Database.

10.2 Industry Data Model or Industry Analytic Schema. Teradata's Industry Data Models encapsulate an enterprise view of data designed to include industry-specific business objectives and rules. During the Cloud Service Term, Teradata grants customer a non-exclusive, non-transferable, worldwide license to use either the Industry Data Model or the Industry Analytic Schema. The Teradata Industry Data Models and Analytic Schemas are confidential information of Teradata. Customer may only use the Industry Data Model or Industry Analytic Schema in accordance with the limitations set out in the Cloud Service Agreement.

- a) Customer can choose one of the following Industry Data Models or Industry Analytic Schemas to license:
  - i. Communications Data Model (CDM)
  - ii. Financial Services Data Model (FSDM)
  - iii. Healthcare Data Model (HCDM)
  - iv. Life Sciences Data Model (LSDM) (not available for Industry Analytic Schema)
  - v. Manufacturing Data Model (MFGDM)
  - vi. Media and Entertainment Data Model (MEDM)
  - vii. Retail Data Model (RDM)
  - viii. Transportation and Logistics Data Model (TLDM) (not available for Industry Analytic Schema)
  - ix. Travel and Hospitality Data Model (THDM)
  - x. Utilities Data Model (UDM)
- b) Teradata will provide the chosen Industry Data Model or Industry Analytic Schema by installing it on VantageCloud as a physical data model or schema without customization for a customer's data demographics, access paths, or temporal requirements. Customer can modify the data model or schema and source-to-target mappings or for an additional fee, contract with Teradata Consulting to assist with such modifications.
- c) Additional Industry Data Models and/or Industry Analytic Schemas are available to license for an additional fee.

## 11. Education Services

Teradata provides each VantageCloud Lake Customer a Premier Education Package and offers additional training packages for individual purchase.

- 11.1 Premier Education Package. The Teradata Premier Education Package is a comprehensive learning resource designed to help customers enhance their skills and knowledge in various areas related to Teradata's offerings. The package includes a curated selection of on-demand courses in the categories of Getting Started, Analytics and SQL, ClearScape, Concepts, and Tools & Applications.
- 11.2 Teradata Training Packages. Education packages that include virtual instructor-led training as private or public sessions, and subscriptions to Teradata University (with hands-on labs) are available for individual purchase.

## 12. Support and Operational Services

Teradata provides maintenance and support services for all VantageCloud Lake packages, including defined coverage hours and response times, access to the Support Portal, and other features. Support and Operational Services are provided in English; however, Support Tickets (Cases, Change Requests, and Service Requests) will have language translation capabilities based on Customer settings in the Support Portal.

12.1 Support Portal. VantageCloud Lake Support and Operational Services are accessible through the Support Portal ([support.teradata.com](https://support.teradata.com)), which is available 24 hours a day, 7 days a week, 365 days a year. The Support Portal gives Customer self-service access to features such as Support Tickets, analytics, Customer dashboards, and Knowledge Articles.

12.2 Support Tickets. Customer is responsible for submitting Support Tickets (i.e., Cases, Change Requests, and Service Requests) through the Support Portal. Teradata uses best practices from the Infrastructure Technology Information Library (ITIL). The Support Ticket Types and Submission Options table (below) illustrates these options.

Support Ticket Types and Submission Options	
Support Ticket Types	Support Ticket Submission Options
Case	<ul style="list-style-type: none"> <li>• Support Portal</li> <li>• Telephone (S1, S2 Cases only): 1-877-MY-TDATA, Option 3</li> <li>• Automatic Case Creation: Generated by Teradata</li> </ul> <p><i>Note: Additional global support telephone numbers are available via the Support Portal.</i></p>
Change Request	Support Portal
Service Request	Support Portal

- a) Case is a record to track customer issues, requests, and resolution activity.
- i. When opening a Case in the Support Portal, Customer must select a severity level based on the level of impact (see the Case Severity Levels table below).
  - ii. Teradata will respond to Cases based on the assigned Severity Level.
  - iii. Teradata also provides automated event management to monitor VantageCloud Lake health and to create Cases.
  - iv. In some instances, Cases may be set to a State of “Awaiting Info.” When that happens:
    - a. The Case resolution duration pauses until Customer responds.
    - b. If no response is received within 14 days, the Case closes automatically.
  - v. Once a solution has been provided in a Case, the State is set to “Resolved.” When that happens:
    - a. Customer needs to review the update and either Accept or Reject the solution.
    - b. If Customer does not respond within 14 days, the Case will close automatically.

Case Severity Levels		
Severity	Impact	Description
S1 (Critical)	The mission-critical Production Environment is down, corrupted, or severely degraded, or it is unusable and requires immediate attention to return the Environment to service.	Daily business is being critically impacted, causing revenue/risk exposure. Users are unable to perform primary functions; no workaround exists. Many users cannot access the Environment or login. <b>Note:</b> For all S1 Cases, Teradata requires Customer engagement to provide the appropriate Support Levels. Teradata will default to contact the Primary Teradata Support Contact(s) and the Case submitter (if registered in the Support Portal). If the Case submitter or Primary Teradata Support Contact(s) are not available, Teradata will change the Case from S1 to S2 and respond accordingly.
S2 (High)	The Production Environment is up and operational, but the issue has severe, ongoing, daily impact to operations; a non-mission critical Environment is down and requires expedited engagement and urgent resolution.	If not resolved, daily operations will be impacted. Many users are unable to perform primary functions; no workaround exists which significantly affects ability to sufficiently achieve business objectives. Many users are affected.
S3 (Medium)	The issue interferes with normal work efforts, but work can continue. Environment response / performance is degraded. Non-critical application functionality is not available.	Medium and manageable impact to business, with little revenue/risk impact. Users are unable to perform secondary functions without a sufficient short-term workaround. Several users are affected.
S4 (Low)	A minor issue exists; normal operations can continue. Functionality is impacted, but not down.	Issue has no business impact and low impact to operations. Additional research or information is needed to address a question. Impacts only a few users.
S5 (Planned)	No issue exists; normal operations can continue.	No business impact exists. Teradata uses this severity level for proactive, planned Cases.

- b) Change Request is a request to change something about an Environment. These changes can include the need to add, modify, configure, or even decommission a site or to discontinue use of a service component, application, or other associated elements.
- i. Customer can submit new “Normal” Change Requests and view existing Change Requests in the Support Portal.
  - ii. Teradata can designate the Change Request as one of three types:
    - a. Standard – Low risk, pre-approved change plans that follow a specified and repeatable procedure or work instruction. These changes do not require case-by-case approval.
    - b. Normal – Changes without predefined plans that require both Customer approval and approval from the Teradata Cloud Change Advisory Board (i.e., a formal approval authority whose function is to control changes to the approved VantageCloud Lake architecture).
    - c. Emergency – Unplanned changes necessary for service restoration. These changes require Customer approval and approval from the Teradata Cloud Change Advisory Board. An Emergency Change can only be created in one of the following situations:
      - The Change is necessary to restore service and is recommended by Teradata Services SMEs during an S1 case investigation. The change must be for the same account as the Severity 1 case.

- A critical security vulnerability exists and, if not expeditiously corrected, could cause harm to the Cloud Environment and its Customers.
- iii. Teradata will schedule the action taken in response to a Change Request in advance and will coordinate the date and time with both Customer and the assigned Teradata resources.
- c) Service Request has predefined and specific actions, services, or work activities that happen upon submission of the request. For example, the Service Request, “Create a Contact,” on the Support Portal will allow Customer to add a contact to Customer’s account. Some Service Requests, however, may need to go through an approval process or include tasks that must be completed by Teradata. Customer can submit new Service Requests and view existing Service Requests in the Support Portal.

12.3 Access Management (i.e., Support Role Assignment). Customer is responsible for initial registration to the Support Portal, and then ongoing assignment and maintenance of Support Portal Customer roles as shown in the following table. Premier Cloud Support and Priority Service require the following roles: Customer Admin, Change Control Contact, Event Notification Contact, Remote Access Contact, and Primary Teradata Support Contact. Premier Cloud Support and Priority Service require the roles and functions shown in the Support Portal Customer Roles and Functions table below.

Support Portal Customer Roles and Functions	
Customer Roles	Functions
Customer Admin*	<ul style="list-style-type: none"> <li>• Approves all access requests</li> <li>• Grants site roles to customer contacts</li> <li>• Grants Secure Password Vault responsibilities</li> </ul>
Change Control Contact**	<ul style="list-style-type: none"> <li>• Approves Change Requests</li> <li>• Approves maintenance schedules</li> </ul>
Event Notification Contact**	Serves as the default Customer contact for both Teradata-owned Cases and Cases generated from Environment events (will be notified and/or added to Case watchlists)
Remote Access Contact**	Serves as support contact for Environment credentials
Primary Teradata Support Contact**	Serves as primary customer contact for escalations
Customer User	Creates, tracks, and manages Support Tickets
* Teradata recommends that Customer considers assigning two or more Customer Admins.	
** Assignment required for each Site ID.	

12.4 Support Levels (i.e., Premier Cloud Support or Priority Service) define the types and standards of services to be offered.

- a) Premier Cloud Support (included with all Packages) provides integrated maintenance and support services for all VantageCloud Lake subscriptions, enabling access to the Support Portal, downloadable software, knowledge base searching, communities and forums, and other features. Premier Cloud Support coverage hours, Remote Response, and status update cadence are described in the Premier Cloud Support table below. Remote Response time is measured during Customer contracted hours of remote coverage by the interval between Customer’s initial contact (via electronic receipt of case or phone call) to

Teradata and the first contact (via electronic receipt or phone call) with a Teradata representative.

Premier Cloud Support					
	S1 (Critical)	S2 (High)	S3 (Medium)	S4 (Low)	S5 (Planned)
Case Coverage Hours	24 x 7	9 standard business hours, 5 business days per week <sup>1</sup>			N/A
Remote Response	2 hours	2 hours <sup>1</sup>	Next business day <sup>2</sup>		N/A
Customer Status Updates	Every hour	Every 6 hours <sup>1</sup>	Daily <sup>1</sup>	Weekly	N/A

<sup>1</sup> Same business day: Monday – Friday (time zone of Customer’s address as set out in the Order)  
<sup>2</sup> Next business day example: If Customer opens an S3 or S4 Case after 8 pm Friday, Teradata will respond after 9 am Monday.

- b) Priority Service (included with select Packages or as an optional add-on) enhances Premier Cloud Support by providing Customer with increased VantageCloud Lake Support coverage hours and accelerated response times for certain Case severities as described in the Priority Service table below.

Priority Service					
	S1 (Critical)	S2 (High)	S3 (Medium)	S4 (Low)	S5 (Planned)
Case Coverage Hours	24 x 7		9 standard business hours, 5 business days per week <sup>1</sup>		N/A
Remote Response	30 minutes		2 hours <sup>1</sup>	Next business day <sup>2</sup>	N/A
Customer Status Updates	Every hour	Every 6 hours <sup>1</sup>	Daily <sup>1</sup>	Weekly	N/A

<sup>1</sup> Same business day: Monday – Friday (time zone of Customer’s address as set out in the Order)  
<sup>2</sup> Next business day example: If Customer opens an S3 or S4 Case after 8 pm Friday, Teradata will respond after 9 am Monday.

12.5 Updates and Maintenance (included with all Packages)

- a) Teradata updates Customer Environments on a regular basis, sometimes as frequently as once a week, and publishes a summary of changes that were made available as part of the update. This process allows Teradata to incrementally and regularly deliver service improvements in the form of new features, enhancements, and fixes. The updates also allow Teradata to apply security fixes to keep the Environment secure and compliant. These updates are intended to have limited impact on service to Customer.
- b) As new versions are released, Customer will need to upgrade the Teradata Tools and Utilities (TTU) software on all Customer devices that connect to VantageCloud Lake to ensure compatibility and stability. In addition, Customer will need to update all third-party dependent applications to ensure the VantageCloud Lake update does not cause disruptions due to version incompatibilities.
- c) Changes in Capacity. Any large increases in capacity will require coordination between Customer and Teradata to schedule and implement such an increase.

12.6 Service Disruption Recovery (included with all Packages). In the event of an unplanned, extended service disruption where VantageCloud is deployed in a cloud provider region that has multiple availability zones, Teradata will use reasonable efforts to deploy a new system in a secondary, unimpacted availability zone and restore from an existing backup to a secondary,

unimpacted availability zone as part of our service. Teradata offers disaster recovery service options at an additional fee to reduce downtime during extended service disruptions or long-term outages.

12.7 Service Advisory (included with select Packages or as an optional add-on)

- a) Teradata provides access to Teradata experts and resources to advise and assist Customer with initial Support Portal onboarding, including Access Management and an overview of Support and Operational Services.
- b) Teradata also provides ongoing Support and Operational Services Advisory via a Support Ticket in the Support Portal.

12.8 Backup and Restoration Assessment (included with select Packages or as an optional add-on).

In this assessment service, Teradata will review Customer's Data Protection Plan and help identify opportunities to bolster their data protection strategy if areas of concern are identified. The assessment focuses on identifying structures that are not included in Backup Jobs, structures that are included in multiple Backup Jobs, and structures configured to Backup Jobs that are no longer present on the system. Teradata will also provide a review of Backup Job execution.

12.9 Cost Optimization Assessment (included with select Packages or as an optional add-on). This assessment service helps organizations identify and reduce unnecessary cloud costs by providing insights into component usage and efficiency, which can lead to potential changes and cost-saving strategies. This assessment focuses on the efficiency of applications through either the optimization of applications or the appropriate use or sizing of VantageCloud Lake components. The assessment will be performed at the query level, the application level, the compute group level, and the overall component level.

12.10 Data Management and Workload Modernization Assessment (included with select Packages or as an optional add-on). Since VantageCloud Lake offers multiple options of where data is placed and how it is processed, this assessment helps identify if data is better suited for block storage or object storage based upon the types of workloads. Teradata will include an assessment of overall data placement, loading techniques, and statistics to provide Customer with information and recommendations on potential opportunities to optimize data management and workload performance.

## 13. Pricing

VantageCloud Lake uses a consumption-based pricing model with Units as the compute measurement metric. Units are consumed when VantageCloud Lake compute resources are running and available to perform Customer workloads. Customer can control when compute resources are running or not, allowing for cost flexibility.

13.1 Committed Quantities. Units, Storage and Data Transfer are entitlements that are sold in committed quantities. At the beginning of the Order's term, Customer commits to an initial number of Units, Storage and/or Data Transfer for the Order's term. The entitlements are allocated across one or more 12-month terms that begin with the anniversary of the effective date of the originating Order and must be used within such 12-month term. Any partial period of less than 12 months will be prorated. The commitment(s) can be used by any VantageCloud Lake Environment compute resources set up under the Organization.

- a) Any unused quantities of a committed entitlement will expire at the end of each 12-month term and do not roll over; there is no credit or refund for unused commitments.

- b) During an Order's term, Customer may purchase additional entitlements for any given product feature without being obligated to increase the quantities of their commitments to other parts of the Order. Customer can buy additional committed entitlements for either:
  - i. The balance of that 12-month period and such committed entitlements will expire at the end of that 12-month period.
  - ii. The balance of the Order's term, provided that each remaining full or partial 12-month period of the Order's term is increased by the same number of entitlements, with the partial period prorated accordingly.
- c) Quantities purchased at a committed rate will increase the commitment for subsequent terms of the Order, including any renewals beyond the initial term.
- d) In the absence of a committed quantity, Customer will be invoiced monthly in arrears at the on-demand rate stated in the Order for the calendar month's usage.
- e) At the end of a committed Order's term, unless Customer has turned off the VantageCloud Lake compute Environment resources, use of such Environment is converted to an On Demand Order (on a month-to-month term) at the applicable then-current, on-demand rate.

13.2 Storage and Data Transfer Pricing under this VantageCloud Lake Services Description are listed below as separate billable items.

- a) Block Storage is consumed based on the maximum amount provisioned at any point during each calendar month.
- b) Object Storage is consumed based on the average storage capacity utilized at any point during each calendar month.
- c) Standard Backup Storage is consumed based on the average storage capacity utilized at any point during each calendar month.
- d) Data Transfer (Ingress and/or Egress) fees are based on the actual total ingress and/or egress of data during each calendar month.

13.3 Consumption Monitoring: In the VantageCloud Lake Console, Customer can view compute consumption, expressed as Units, at the Organization level. Drill-down capabilities allow Customer to view Unit consumption at the Primary Cluster and Compute Cluster levels. Storage consumption can also be viewed in the VantageCloud Lake Console. Refer to the platform-specific addendum for details of Primary Cluster and Compute Cluster Unit Consumption.

## 14. Pre-General Availability Offerings

Teradata may make available, and Customer may choose to use, pre-general availability features for VantageCloud Lake Edition that are identified as "Limited Availability," "Early Access," "Private Preview," "Alpha," "Beta," or similar designation in related documentation or materials. Pre-General Availability offerings are not necessarily feature-complete, nor do they necessarily have technical support commitments. Unless otherwise stated by Teradata, Pre-General Availability features are intended for use in test Environments only and should not be used to process personally identifiable data or data subject to legal or regulatory compliance requirements. Customer may provide feedback and suggestions to Teradata about pre-general availability offerings, and Teradata and its affiliates may use any feedback or suggestions provided without restriction and without obligation to Customer. Pre-General Availability offerings may be changed, suspended, or discontinued at any time without prior notice to Customer and are not covered by any Service Level Agreement.

## 15. Additional Services (Sold Separately)

15.1 Platform Services Management (PSM) and Critical Platform Services Management (CPSM) provide proactive account management delivered by a designated SEM along with local language support. PSM and CPSM are available as an optional add-on with select Packages.

- a) Platform Services Management (PSM) serves as the entry tier for SEM services, offering a light-touch approach. The offer comes with a Service Experience Manager (SEM) who will oversee the overall Customer service experience.
- b) Critical Platform Services Management (CPSM) builds upon PSM by providing proactive guidance and reviews, operational success enablement and monitoring, and on-call support.

Features	Platform Services Management (PSM)	Critical Platform Services Management (CPSM)
SEM Support	Designated Resource	Designated Resource
Business and Operational Reviews	Quarterly	Monthly
Support Services Account Management	•	•
Local Language Support*	•	•
Proactive Guidance & Reviews		•
Operational Success Enablement & Monitoring		•
On-Call Support		•
<i>*Supported languages are limited. Customers should contact their account team for more information.</i>		

- c) Definitions of the features for PSM and CPSM are as follows:
  - i. SEM Support is a designated resource that serves as Customer’s primary point of contact, offering strategic guidance, proactive management, and specialized assistance to Customer. The designated SEM will perform services during normal business hours, based on Customer’s region as set out in the Order.
  - ii. Business and Operational Reviews assess optimization strategies and improvement opportunities. The reviews occur monthly for CPSM and quarterly for PSM.
  - iii. Support Services Account Management involves proactive account management of customer accounts by a designated SEM. The SEM collaborates closely with Customer during reviews, ensures smooth onboarding, and advocates for Customer during critical Cases.
  - iv. Local Language Support is an advisory service delivered by the designated SEM. The SEM can speak and interact with a Customer’s team in the supported language(s) to accommodate regional needs, facilitating effective communication and understanding. This service does not include official documentation translations.
  - v. Proactive Guidance & Reviews are monthly collaborative reviews led by the SEM that help Customer achieve meaningful outcomes. The SEM analyzes data to identify areas for improvement and suggest corrective actions.
  - vi. Operational Success Enablement & Monitoring involves the SEM proactively monitoring key performance indicators, offering timely insights, and providing guidance to Customer to help improve operational success.
  - vii. On-Call Support builds upon SEM Support by providing S1 Case support after local business hours through a team of SEMs.

15.2 VantageCloud Services are services (sold separately) provided by Teradata to manage Customer's VantageCloud analytics and data platform(s) ("VantageCloud Ecosystem") and maintain system efficiency and performance. All services are delivered remotely by Teradata consultants. For the provision of these VantageCloud Services, Customer is responsible for security obligations per the shared responsibility model and for controlling access to any Customer Data that Teradata consultants may require. Any specific activities that Customer wants Teradata to perform to meet Customer's security obligations must be mutually agreed on by Teradata and Customer.

- a) List of Services. The following is a list of VantageCloud Services current as of the date of this publication. The complete set of VantageCloud Services with links to details can be found in this [CSD Knowledge Article](#).
  - i. Analytic Pipeline Operations
  - ii. Analytic Platform Operations
  - iii. Database Admin & Operations. This also includes Workload Optimization as an optional add-on service.
  - iv. Database Insights
  - v. Database Optimization
  - vi. Enterprise Feature Store (EFS)
  - vii. Model Operations Center Management
  - viii. SAS Optimization
  - ix. Standalone Backup and Restore (BAR) Operations
  - x. Use Case Solutions
  - xi. DataDNA Operational Services
- b) Services Onboarding and Start Date of Service. Teradata will start work and begin billing for each VantageCloud Service on the start date of the Order. Customer is responsible for tasks outlined in the respective Services addenda that need to be completed before Teradata resources can be fully functional.
- c) Remote Access. Teradata will perform VantageCloud Services (as described in the applicable VantageCloud Service addendum) remotely through the Teradata Services Tools (e.g., Viewpoint and PDCR), and Customer will manage database and application-level credentials for Teradata's consultants. Please refer to Section 4.13 (Environment Monitoring and Telemetry) of this Cloud Service Description.
- d) Shared Services. Teradata will utilize a shared resource pool of Teradata consultants to deliver services. Customer must allow Teradata to meet the minimum connectivity requirements below.
  - i. Customer will provide user logons for personnel in the shared services pool as defined during the Transition Service period. The VantageCloud Services require up to 12 user logons depending on the service(s) selected. The service-specific addendum will provide details on the number of user logons needed for each service.
  - ii. Customer will allow the shared resource pool to use Teradata laptops for logging on to the Vantage Cloud Ecosystem.
- e) Performance Data Collection and Reporting (PDCR) Requirements. The following are the Customer requirements for PDCR if the service-specific addendum states that PDCR is required.
  - i. Data Collection Requirements. Customer will establish PDCR data collection at least three weeks before the start of the VantageCloud Service. Should Customer have

concerns regarding the amount of data storage used for PDCR, Teradata will assist in reviewing and redefining logging protocols to optimize storage efficiency while maintaining the essential performance data collection requirements for the VantageCloud Service. If PDCR is already active and with collection jobs enabled, Customer will verify that the advocated data collection best practices have been activated for at least three weeks prior to the start of the VantageCloud Service.

- ii. DBQL. Customer will set usage logging at the detail level. The only exception is usage that is running strictly as sub-second, known work (tactical applications). In addition to the normally recommended detail logging, StepInfo and EXPLAIN data are required for this Cloud Service.
  - iii. ResUsage. Customer will ensure availability of the latest 90 days of ResUsage data stored in the DBC database or PDCR.
  - iv. Collection and Logging Settings. Customer will set the PDCR Collection at 60 and PDCR Logging at 600 using Viewpoint or Remote Console.
- f) Data Responsibility. With respect to the VantageCloud Services and in accordance with the applicable Data Processing Addendum, Customer is the data Controller and shall remain responsible for determining the purposes and means of the processing of Customer Data. In certain cases, when specifically contracted by Customer to do so, Teradata may process Customer data on behalf of Customer. Customer is responsible for knowing the location of all Customer Data maintained on their VantageCloud Ecosystem and is responsible for implementing Customer-controlled data security measures, including but not limited to encryption—both at rest and in transit. In addition, Customer shall maintain all Teradata system logs to validate access, record retention, and deletion of this Customer Data for regulatory purposes. Customer must enable logging features during transition to accurately log the data processing activities on Teradata VantageCloud.
- g) Teradata Service Tools
- i. Teradata may utilize or develop certain computer programs, data, scripts, documentation, tools, and other materials that Teradata uses to deliver VantageCloud Services (“Teradata Services Tools”). These Teradata Services Tools include but are not limited to the following: Monthly Service Report Tool, Operations Assessment Tool, Workload Management Assessment Tool, Multi-Value Compression Tool, and BAR Assessment Tool. Teradata Services Tools are the confidential and proprietary intellectual property of Teradata and are not licensed or transferred to Customer.
  - ii. These tools may not be copied, transferred, disclosed, or used by anyone other than Teradata without Teradata's prior written consent, and no rights or licenses will be implied. Teradata may install, update, change, and/or remove Teradata Services Tools at its discretion. Teradata will remove the Teradata Services Tools from Customer’s Environment by no later than the last day of Teradata providing the applicable service.

## 16. Availability Service Level Agreement

### 16.1 Definition of Availability and Related Terms

- a) Teradata shall make the Analytics Database of the VantageCloud Service Available, as measured over the course of each calendar month, at the availability service level percentage defined in Section 2 (VantageCloud Lake Packages).
- b) “Available” (and correlative term “Availability”) means the Analytics Database is operating.
- c) Total Minutes in a Calendar Month is the number of days in the calendar month multiplied by 1440 minutes per day.
- d) Unplanned Outage is any period of time during which the Service is not Available but excludes intervals of time that are an Exception. An Unplanned Outage commences on the earlier of (a) when Customer informs Teradata that the Service is not Available or (b) when Teradata otherwise becomes aware that the Service is not Available, excluding intervals when the Service is not Available due to an Exception. Unplanned Outages will continue until the Service is Available.

16.2 Exceptions. When intervals of the Service not being Available are the result of any of the following conditions (“Exceptions”), these intervals will be excluded from the reported calculation during Teradata’s analysis of the Unplanned Availability Percentage for that calendar month.

- a) Customer’s or any of its representatives’ use of the Service in breach of the Agreement;
- b) Internet or other network traffic problems other than problems with networks controlled by Teradata;
- c) an event beyond Teradata’s reasonable control where the cloud provider availability zone or region is impacted;
- d) a Customer-agreed outage when the Service is not Available, including for updates, maintenance, support, and operational services; or
- e) for VantageCloud Systems provided at no charge, such as proof of concept Sites.

### 16.3 Reporting of Service Availability

- a) Within ten (10) business days after the end of each calendar month, Teradata will report on the Support Portal (support.teradata.com) the Unplanned Availability Percentage, which is the proportion of time the Service was Available during the Total Minutes in a calendar month after subtracting Unplanned Outages, as calculated below.
- b) Customer must permit the required telemetry collection included in the VantageCloud Lake subscriptions as set out in Section 4.13 (Environment Monitoring and Telemetry) of this Cloud Service Description to be eligible for this Availability SLA.

16.4 Calculation of Unplanned Availability Percentage. The Unplanned Availability Percentage is calculated as follows:

$$\frac{\text{Total Minutes in a Calendar Month} - \text{Total Minutes of Unplanned Outages in a Calendar Month}}{\text{Total Minutes in a Calendar Month}} \times 100$$

### 16.5 Remedies for Unplanned Service Availability Failures

- a) If the Unplanned Availability Percentage for a Site during a calendar month is less than the Availability SLA percentage, Customer Admin may request a corresponding service credit as set out in the table below based on the affected Site’s fees for the calendar month (the “Service Availability Credit”) by submitting a Service Availability Credit Service Request in the Support Portal. Customer’s Service Availability Credit request must be received by

Teradata by the end of the next calendar month following the report of the unavailability giving rise to the claim. The parties acknowledge and agree that the Service Availability Credits assessed pursuant to this Availability SLA (i) are a reasonable estimate of the compensation for the anticipated or actual damage to Customer that may arise from the corresponding failure to meet the Availability SLA, which would be difficult to accurately estimate; (ii) must be used within twelve months of Teradata issuing the credit; and (iii) may, at Customer's option, be credited or set off against any fees or other charges payable to Teradata under the Agreement or payable to Customer upon demand.

<b>Service Availability Credit Table</b>	
<b>Unplanned Availability Percentage</b>	<b>Service Availability Credit*</b>
Less than Availability SLA% but equal to or greater than 99.0%	10%
Less than 99.0% but equal to or greater than 95.0%	25%
Less than 95.0%	50%
<i>* The total Service Availability Credits owed in each contract year of an Order's Cloud Service Term shall not exceed 10% of such Order's annual Teradata Cloud Service fees.</i>	

- b) If in any three months in a rolling six month period, the Unplanned Availability Percentage for a Site is less than the Availability SLA percentage, then, in addition to the credits set out in Section 16.5.a, Customer may terminate the Cloud Service for that Site on written notice to Teradata with no liability, obligation, or penalty to Customer by reason of such termination, and Teradata will promptly, but in no event more than thirty days after termination of such Site, provide a refund of any pre-paid Customer fees for such Site prorated based on the date of termination.



# Teradata VantageCloud Lake on Amazon Web Services (AWS)

## Cloud Service Description: AWS Addendum

---

This document supplements the Teradata VantageCloud Lake Cloud Service Description for the AWS cloud platform.

**August 11, 2025**

## Table of Contents

1. Teradata VantageCloud Lake on AWS .....	3
2. Supported AWS Connectivity Options .....	3
3. Supported AWS Configurations .....	3
4. Security on AWS .....	4
5. Data Protection on AWS .....	5
6. Elasticity Features on AWS.....	5
7. Interfaces and Tools on AWS .....	5
8. External Connectors on AWS .....	6
9. Analytics on AWS.....	7
10. Applications on AWS.....	7
11. Pricing on AWS .....	8
12. Storage and Data Transfer Pricing on AWS .....	10
13. Supported AWS Regions .....	11

## 1. Teradata VantageCloud Lake on AWS

Provisioning a VantageCloud Lake Environment creates a set of resources in a Teradata-owned AWS account. This document supplements the Teradata VantageCloud Lake Cloud Service Description. In case of overlap or conflict, this addendum supersedes the relevant sections in the Teradata VantageCloud Lake Cloud Service Description.

## 2. Supported AWS Connectivity Options

Connectivity for VantageCloud Lake is described in the Supported AWS Connectivity Options table below.

Supported AWS Connectivity Options	
Connectivity Options	Details
AWS PrivateLink	<ul style="list-style-type: none"> <li>Connect over a private network to VantageCloud Lake</li> <li>AWS PrivateLink is supported by an endpoint (created in same AWS region as VantageCloud Lake) in a Customer Virtual Private Cloud (VPC)</li> <li>One instance of AWS PrivateLink Service is included per VantageCloud Lake Environment</li> </ul>
Public Internet	<ul style="list-style-type: none"> <li>Connect over the internet to VantageCloud Lake</li> <li>Customer must add a source CIDR block to whitelist access in the VantageCloud Lake Console</li> </ul>

## 3. Supported AWS Configurations

VantageCloud Lake offers node and cluster options based on customer infrastructure and performance requirements. Primary Cluster Nodes and Compute Clusters are preconfigured with Teradata software, with the size options as shown in the table below.

Primary Cluster	
Node Size Options	XSmall–Standard to 3XLarge–Standard <sup>12</sup> XSmall–High I/O to 2XLarge–High I/O <sup>1</sup> 4XLarge–High I/O <sup>1</sup> (detailed in Section 11.1)
Node Limit	2–128 nodes
Compute Cluster	
Cluster Size Options	XSmall to 2XLarge (detailed in Section 11.2)

<sup>1</sup>Only available in select AWS Regions

<sup>2</sup>3XLarge-Standard is only available for new deployments after January 2025

## 4. Security on AWS

This section supplements Section 4.7 (Encryption) in the VantageCloud Lake Cloud Service Description.

- 4.1 Customer-Managed Encryption Keys (CMEK). For encryption of data at rest, Customer has the option to use CMEK in addition to the default AWS-managed encryption keys. The scope of CMEK support is limited to single region, symmetric software keys.
- a) CMEK can only be enabled during system provisioning and cannot be enabled for an existing site.
  - b) CMEK is available for use on VantageCloud Elastic Block Storage (Database Disks), Object Storage (OFS), Standard Backups and Snapshot Backups.
    - i. CMEK is not available for all VantageCloud services; for such services, data is secured-at-rest with Cloud Provider managed encryption keys.
    - ii. Multi-Region Keys, HSM based keys, and Asymmetric Keys are not supported.
  - c) Customer is responsible for creation, rotation, and maintenance of CMEK key(s) and key policy using their AWS Key Management Service (KMS)
  - d) If desired, CMEK keys can be rotated by Customer by creating a new key and configuring the correct key policy, or by enabling Automatic Key Rotation for their existing key in AWS KMS.
  - e) Following key rotation, data will be re-encrypted with the new key as follows:
    - i. New key will be used to encrypt all data in Elastic Block Storage as part of the next upgrade. Customer should consult with Teradata to align key rotation with upgrade to minimize downtime.
    - ii. New key will be used to encrypt all new objects on AWS S3 (including updates to existing data). The old key will continue to be used for all existing objects in AWS S3; hence Customer should not disable or destroy any old CMEKs.
    - iii. Over time, all old data in the object store will migrate to the new key as those objects receive updates, or it gets purged due to retention policies.
    - iv. After data re-encryption, older CMEKs should be retained for the same time period as Customer-configured retention period for Standard Backups and Snapshot Backups (VantageCloud Lake does not support re-encryption of backups).
  - f) Customer should not disable or destroy their current or old CMEKs. Old CMEKs are required to access previously encrypted data after key rotation.
    - i. As an exception, Customer may choose to disable their CMEK when they believe their key is compromised.
    - ii. Revoking/Disabling Teradata's access to the CMEKs will cause VantageCloud Lake to cease operations and may result in data integrity issues, corruption, and/or the unrecoverable loss of Customer Data.
    - iii. Destroying/Deleting the CMEKs will result in complete and unrecoverable loss of Customer Data. If keys are deleted, everyone including Customer will permanently lose access to the data.

## 5. Data Protection on AWS

This section supplements Section 5 (Data Protection) of the VantageCloud Lake Cloud Service Description. These features are available with select VantageCloud Lake Packages, as shown in the table below.

- 5.1 Cross-Environment Full/Partial Backup. Customer can request a Full Backup across VantageCloud Lake AWS Environments by submitting a Change Request in the Support Portal. For Partial Backups, Customer can utilize the VantageCloud Lake Console to execute the Partial Backup job across Environments.
- 5.2 Cross-Environment Full/Partial Restore. Customer can request a Full Restore across VantageCloud Lake AWS Environments by submitting a Change Request in the Support Portal. For Partial Restores, Customer can utilize the VantageCloud Lake Console to execute the Partial Restore job across Environments.
- 5.3 Availability of the Cross-Environment Data Protection features by package is shown below.

VantageCloud Lake Packages			
Features And Services	Standard	Lake	Lake+
Cross-Environment Full/Partial Backup		•	•
Cross-Environment Full/Partial Restore		•	•

## 6. Elasticity Features on AWS

This section supplements Section 7 (Elasticity Features) of the VantageCloud Lake Cloud Service Description with AWS-specific information.

VantageCloud Lake Package Elasticity Features and Descriptions	
Features	Descriptions
Primary Cluster Scale Up/Down	<ul style="list-style-type: none"> <li>• XLarge–Standard and 2XLarge–Standard node sizes do not support scale down operation with io2 block storage</li> <li>• 3XLarge–Standard can only be scaled down by 1 size with io2 block storage</li> </ul>

## 7. Interfaces and Tools on AWS

This section supplements Section 8 (Interfaces and Tools) of the VantageCloud Lake Cloud Service Description.

- 7.1 Data Copy in VantageCloud Lake Console. Data Copy is a feature that enables Users via the VantageCloud Lake Console to provision, configure, and run data copy jobs. With this feature, Users can copy data between VantageCloud Lake on AWS and either VantageCloud Enterprise (on any Cloud Service Provider) or VantageCore (on-premises systems). Data Copy uses Teradata QueryGrid as the copy method to perform full or partial table copies. Data Copy is only available for Lake and Lake+ packages.
- 7.2 ask.ai in VantageCloud Lake Console. ask.ai is Teradata's context-aware Artificial Intelligence (AI) application that uses natural language processing and is available for use via the VantageCloud Lake Console. ask.ai leverages VantageCloud Lake APIs and data to provide personalized responses and supports Knowledge Intelligence related prompts. "Knowledge

Intelligence” means that users can ask for help on a variety of topics such as general documentation, questions about what Teradata functions exist in a specified database, and detailed descriptions on a particular function and SQL generation for that function. The availability of ask.ai varies by region; see Section 13 (Supported AWS Regions) for region availability.

- 7.3 ModelOps in VantageCloud Lake Console. ModelOps is a feature available for Lake and Lake+ packages via the VantageCloud Lake Console that allows users to deploy, monitor, and manage machine learning models. It includes a model registry, automated lifecycle management (which covers aspects like data drift and performance monitoring), and job scheduling capabilities for VantageCloud Lake Compute Clusters. Users can utilize Bring Your Own Model (BYOM) and Open Analytics Framework (OpenAF) within this feature.

## 8. External Connectors on AWS

This section supplements Section 9 (External Connector) of the VantageCloud Lake Cloud Service Description.

- 8.1 Teradata QueryGrid delivers data access, processing, and movement across Teradata Environments. The below table for QueryGrid Connectivity for VantageCloud Lake on AWS details the external connectors for the AWS platform.

QueryGrid Connectivity for VantageCloud Lake on AWS	
Connects to	With connector
VantageCloud Lake on AWS	Teradata-to-Teradata <b>Note:</b> To set up and configure the required PrivateLink endpoint on VantageCloud Lake, open a PrivateLink change request on the Support Portal.
VantageCloud Enterprise on AWS	Teradata-to-Teradata <b>Note:</b> To set up and configure the required PrivateLink endpoint on VantageCloud Enterprise, open a PrivateLink change request on the Support Portal.
VantageCore (on-premises)	Teradata-to-Teradata
Amazon EMR	Teradata-to-Apache Hive <sup>1</sup> Teradata-to-Apache Spark <sup>1</sup>
Other JDBC-compatible query engines (such as Amazon Redshift)	Generic JDBC <sup>1</sup> <b>Note:</b> Refer to the “QueryGrid Installation and User Guide” for driver node requirements.

<sup>1</sup>Connectors to non-Vantage systems are an additional fee

- 8.2 VantageCloud Lake on AWS supports bi-directional query initiation with VantageCloud Enterprise on AWS and VantageCore (on-premises) using the Teradata-to-Teradata connector.

## 9. Analytics on AWS

This section supplements Section 10 (Analytics) of the VantageCloud Lake Cloud Service Description.

9.1 Teradata ClearScape Analytics™. In addition to the analytic functions and frameworks of In-Database Analytics, Bring Your Own Analytics, and Partner Integration, ClearScape Analytics on AWS also includes:

- a) Enterprise Vector Store. This feature introduces the VECTOR datatype, vector indexing, and search algorithms to enable Retrieval Augmented Generation (RAG) and other GenAI applications. For pricing details, refer to Section 11.5.

## 10. Applications on AWS

This section describes applications, sold separately, for use with VantageCloud Lake.

10.1 Third-Party Tools Descriptions (Sold Separately). Teradata will host approved and licensed third-party software as described in this section and shown in the Third-Party Tools table below.

- a) Third-Party Tools are available only through a bring-your-own-license (BYOL) model where Customer is required to enter into a separate agreement directly with the third-party vendor to license, maintain, and support the Third-Party Tool for the duration of Customer's VantageCloud Lake subscription.
- b) Customer is responsible for engaging the Third-Party Tool vendor for software maintenance and support.
- c) Customer is responsible for enabling installation of Third-Party Tools via the VantageCloud Lake Console. This enablement grants Teradata permission to install, upgrade or uninstall the software components and perform activities that are necessary for such maintenance and support of the Third-Party Tool software in a VantageCloud Lake environment that require infrastructure or node-level access.
- d) Service Level Availability (SLA) commitments do not apply to the extent that Teradata reasonably determines a failure was caused by the Third-Party Tool software.
- e) Customer is responsible for managing their own Third-Party Tool configurations, including application policies, passwords, and encryption keys.
- f) Implementing one or more of the below mentioned third-party encryption tools might result in an overall performance impact on the VantageCloud Lake system.

Third-Party Tools	
Third-Party Tools	Details
Protegrity	<p><u>Protegrity</u> develops enterprise data security software and provides scalable, end-to-end data security solutions. Protegrity Data Security Platform helps Teradata Customers secure sensitive data and comply with privacy regulations. Protegrity includes two components that pertain to VantageCloud Lake:</p> <ul style="list-style-type: none"> <li>• <u>Protegrity Database Protector</u> provides policy-based data encryption and tokenization capabilities for Analytics Database</li> <li>• <u>Protegrity Enterprise Security Administrator (ESA)</u> provides Customers with centralized, visual administration of data security policies, key management, auditing, and reporting. Protegrity ESA must be installed, hosted, and managed by Customer, outside of the VantageCloud Lake environment managed by Teradata and made accessible through Private Link from VantageCloud Lake.</li> </ul> <p><b>Note:</b> Customer may contract with Teradata Consulting to assist with implementation and support as needed for an additional fee.</p>
SAS Embedded Process (SAS EP)	<ul style="list-style-type: none"> <li>• <u>SAS</u> develops and markets a suite of analytics software, which helps access, manage, analyze, and report on data to aid in decision-making. SAS Embedded Process helps bring to Teradata Customers advanced analytics capabilities through SAS code in Teradata VantageCloud without moving data out of VantageCloud.</li> <li>• <u>SAS EP</u> includes two components that pertain to VantageCloud: <ul style="list-style-type: none"> <li>○ <u>SAS User-Defined Function (UDF) Server components</u> includes SAS EP, SAS EP Support Functions, and the SAS Formats library. Installing or upgrading SAS EP is independent of VantageCloud.</li> <li>○ <u>SAS client-side software</u> provides an interface that Customers can use to initiate function calls and analyze the results.</li> </ul> </li> </ul>

## 11. Pricing on AWS

This section supplements Section 13 (Pricing) of the VantageCloud Lake Cloud Service Description with consumption-based pricing model details.

11.1 Primary Cluster Unit Consumption will be charged per minute, with partial minutes rounded up to the next higher minute. Units are charged for every minute that the Primary Cluster is on and available for running workloads. The Primary Cluster is always running, available to perform Customer workloads, and consuming Units per the table below until deprovisioned. Changes to node size and count will impact Unit consumption. Availability of node sizes varies by AWS region.

Primary Cluster Unit Consumption per Hour			
Unit Rate per hour per Node Count: Min 2, Max 128			
Node Size	Standard Package	Lake Package	Lake+ Package
XSmall–Standard	1.6	2.0	2.4
Small–Standard	3.2	4.0	4.8
Medium–Standard	5.6	7.0	8.4
Large–Standard	8.0	10.0	12.0
XLarge–Standard	10.4	13.0	15.6
2XLarge–Standard	16.0	20.0	24.0

<b>Primary Cluster Unit Consumption per Hour</b>			
<i>Unit Rate per hour per Node Count: Min 2, Max 128</i>			
<b>Node Size</b>	<b>Standard Package</b>	<b>Lake Package</b>	<b>Lake+ Package</b>
3XLarge–Standard	30.4	38.0	45.6
XSmall–High I/O	2.6	3.2	3.8
Small–High I/O	5	6.3	7.6
Medium–High I/O	10.2	12.8	15.4
Large–High I/O	15.3	19.1	22.9
XLarge–High I/O	20.5	25.6	30.7
2XLarge–High I/O	30.6	38.2	45.8
4XLarge–High I/O	40.8	51.0	61.2

11.2 Standard & Analytic Compute Cluster Unit Consumption for Standard and Analytic Compute Clusters will be charged per minute, with partial minutes rounded up to the next higher minute. Compute Cluster(s) are optional, and Customer can provision, schedule, and scale Compute Cluster(s) to perform Customer workloads. Units are charged for every minute that the Compute Cluster(s) are on and available for running workloads. Compute Cluster(s) consume Units per the table below until suspended or deleted. Changes to the cluster size and schedule will impact the Unit consumption as shown in the table below.

<b>Compute Cluster Unit Consumption per Hour</b>			
<b>(Standard and Analytic Compute Clusters)</b>			
<b>Cluster Size</b>	<b>Standard Package</b>	<b>Lake Package</b>	<b>Lake+ Package</b>
XSmall	8	10	12
Small	16	20	24
Medium	32	40	48
Large	64	80	96
XLarge	128	160	192
2XLarge	256	320	384

11.3 Analytic GPU Compute Cluster Unit Consumption will be charged per minute, with partial minutes rounded up to the next higher minute. Analytic GPU Compute Cluster(s) are optional, and the Customer can provision, schedule, and scale Analytic GPU Compute Cluster(s) to perform Artificial Intelligence / Machine Learning (“AI/ML”) workloads. Units are charged for every minute that the Analytic GPU Compute Cluster(s) are on and available for running workloads. Analytic GPU Compute Cluster(s) consume Units per the table below until suspended or deleted. Changes to the cluster size and schedule will impact the Unit consumption as shown in the table below.

<b>Analytic GPU Compute Cluster Unit Consumption per Hour</b>			
<b>Cluster Size</b>	<b>Standard Package</b>	<b>Lake Package</b>	<b>Lake+ Package</b>
XSmall	10	12	14.4
Small	19	24	28.8
Medium	38	48	57.6
Large	77	96	115.2
XLarge	154	192	230.4
2XLarge	307	384	460.8

11.4 Supported AWS Region Availability Zones for Analytic GPU Compute Clusters. Analytic GPU Compute Clusters are supported in VantageCloud Lake AWS Region Availability Zones, subject to Cloud Service Provider availability.

11.5 Enterprise Vector Store Pricing Methodology. For Enterprise Vector Store, charges for model token usage, both from embedding usage and generated by LLMs, will be based on the number of input tokens processed and/or output tokens generated. These charges will be billed monthly in arrears at CSP published rates. The fees will reflect the actual total number of tokens processed during each calendar month.

## **12. Storage and Data Transfer Pricing on AWS**

This section supplements the list of billable items in Section 13.2 (Storage and Data Transfer Pricing) of the VantageCloud Lake Cloud Service Description.

12.1 IOPS (Input/Output Operations per Second) is provisioned based on instance size and quantity. IOPS is an entitlement that is sold in committed quantities; refer to pricing terms in Section 13.1 (Committed Quantities) of the VantageCloud Lake Cloud Service Description.

### 13. Supported AWS Regions

VantageCloud Lake is supported in the AWS Regions shown in the VantageCloud Lake Supported AWS Regions table below.

VantageCloud Lake Supported AWS Regions	
Region Name	Region
<b>North America</b>	
Canada (Central) <sup>1 3 7</sup>	ca-central-1
US East (N. Virginia) <sup>7</sup>	us-east-1
US East (Ohio) <sup>3</sup>	us-east-2
US West (Oregon) <sup>7</sup>	us-west-2
<b>South America</b>	
South America (Sao Paulo) <sup>1 2 3 4</sup>	sa-east-1
<b>Asia Pacific</b>	
Asia Pacific (Tokyo)	ap-northeast-1
Asia Pacific (Osaka) <sup>1 2 3 4 5 7</sup>	ap-northeast-3
Asia Pacific (Mumbai) <sup>1 3 4</sup>	ap-south-1
Asia Pacific (Singapore) <sup>3</sup>	ap-southeast-1
Asia Pacific (Sydney) <sup>1 3</sup>	ap-southeast-2
Asia Pacific (Jakarta) <sup>1 2 3 4 5 6</sup>	ap-southeast-3
<b>Europe</b>	
Europe (Frankfurt) <sup>4</sup>	eu-central-1
Europe (Ireland) <sup>3</sup>	eu-west-1
Europe (London) <sup>1 3 4</sup>	eu-west-2
<b>Africa</b>	
Africa (Cape Town) <sup>1 2 3 5 7</sup>	af-south-1

<sup>1</sup> High I/O Primary Cluster sizes not supported

<sup>2</sup> Lake Primary Cluster node limit is 2-9 nodes

<sup>3</sup> ask.ai not supported

<sup>4</sup> Region may require additional time for activation

<sup>5</sup> 3XL Standard Primary Cluster size not supported

<sup>6</sup> Analytic GPU Compute Clusters not supported in this region

<sup>7</sup> Analytic GPU Compute Clusters not supported in all Availability Zones