

ICDP Fair Usage and Performance Policy

This policy outlines technical and practical limitations intended to ensure that resources are utilized in a balanced manner while protecting the integrity and performance of the services provided.

Please Note: These terms are subject to change at any time without prior notice. While Treasure Data strives to communicate any changes in a timely manner, we do so on a best-efforts basis. The publication of these terms on this webpage constitutes sufficient notice for such changes to take effect. Your continued use of our services following any updates signifies your acceptance of the new terms.

1. Fair Usage of the ICDP

- Customer shall:
 - Only use the ICDP for Ingestion, Cleansing, Unification, Segmentation, Activation, Journey Orchestration and Analysis of Customer/Consumer Data.
 - Refrain from using resource-inefficient data processing practices as much as possible, including, but not limited to:
 - launching multiple overlapping or parallel workloads instead of running sequentially;
 - maintaining idle sessions;
 - duplicating processing; and
 - repeatedly retrying the same task.

Note: Treasure Data Customer Success and Technical Account Management resources can partner with customers to help identify resource-inefficient data processing practices and suggest optimization steps.

2. Ingest and Storage limits

The following Ingest and Storage Row limitations apply to ICDP usage based on purchased Tiers.

ICDP Tier	Ingest (Monthly)	Storage Rows (Daily)
0 - 1	N/A - order form specifies limits	
2 - 12	2,000 B records	4,000 B rows
13+	10,000 B records	20,000 B rows

- The above limits are calculated by aggregating the Ingest and Storage Rows across all Customer Instances.
 - Total Ingest is measured at the end of each month and compared to the above monthly limits to determine usage.
 - Storage Rows are measured daily. The 4th-highest daily Storage Row total amount in each month is compared to the above limits to determine usage.
- If Customer exceeds any of the above limitations for their purchased ICDP Tier:
 - Notification of the excess will be sent to the e-mail address associated with Customer's Instance (as provided by the Customer);
 - Customer will have until the end of the month in which notice is sent to reduce usage below the limits; and
 - If usage continues to exceed the specified limit at the end of the month in which notice is sent, Customer will be automatically upgraded to the next level per the "Upgrade Pricing & Overages Table" that meets their new usage requirements and shall be charged the incremental increase in price for the remainder of the Subscription Term.

3. Parent Segments

- Max 2,000 Parent Segments per instance.

4. Users

- Max 2,000 Users per instance

5. Production Instances:

- A Production Instance is a live, fully managed environment that Treasure Data makes available to the Customer for processing real-world business transactions and storing production data.
- Max 20 Production Instances per customer.

6. Sandbox Instances:

- We offer three types of Sandbox Instance:
 - Dev
 - QA
 - Full
- Max 5 Sandbox Instances (i.e., Dev, QA, Full) for each Production Instance
- Scope, intended use and performance:
 - Dev: The Dev Sandbox Instance is not intended to be a full replica of production. Data is typically mocked to reflect production data, or a small percentage of production data is used in Dev to test new changes to existing/new workflows. Infrastructure resources provisioned in Dev Sandbox Instances are typically 10% of a Production Instance.
 - QA: The QA Sandbox Instance is not intended to be a full replica of production. While data volume may be similar, workflows and jobs should be ad-hoc and less frequent than in production. Performance of job completion in this instance is not

guaranteed to be equivalent to production. Infrastructure resources provisioned in QA Sandbox Instances are typically 50% of a Production Instance.

- Full: The Full Sandbox Instance is a full replica of production in terms of data volume and workflows. The infrastructure resources provisioned in Full Sandbox Instances are typically 100% of a Production Instance.
- Dev, QA and Full:
 - Activations:
 - None of the Sandbox Instances are intended to run full, live, and/or scheduled activations out of the ICDP. In relation to activations, Sandbox Instances should only be used for testing connectivity to downstream platforms. Full, live, and/or scheduled activations shall only be sent from the Production Instance(s).
 - Conversations:
 - Customers may only use AI Agent Foundry in a Sandbox Instance for isolated test and validation purposes.
 - Customer shall not use AI Agent Foundry in a Sandbox for any production-level, repeated and/or non-ICDP-related purposes.
 - Real Time:
 - Customers may only use Sandbox instances for testing isolated samples of real time data.
 - Customers shall not run production-level volumes and/or frequency of real-time data processing in Sandbox Instances.
 - Customers shall not perform sustained or high-throughput load testing, including ongoing simulations of user activity or latency tracking, in Sandbox Instances.

7. Infrastructure performance

- Treasure Data will provision sufficient infrastructure resources for customers to ingest, cleanse, unify, segment, activate, orchestrate and analyze their profile and behavior data volumes while meeting the following minimum performance standards:
 - Query Success Rate (QSR): 99.9%
 - QSR is the percentage of queries that have completed successfully (including retries) out of the total number of queries submitted (excluding user cancelled or failed due to user errors)
 - Query Execution Time Consistency (QETC) :
 - QETC measures the consistency of query performance. It monitors each distinct query pattern over a rolling-week window, removes obvious outliers, and gauges variability in the remaining run-time samples. Those results are then compared with a historical baseline to spot any drift. By aggregating the share of query patterns that remain within acceptable variability and drift limits, QETC surfaces whether overall query performance is holding steady or starting to deteriorate.
 - These QSR and QETC performance standards are measured every 2 weeks

- Note: Query Queue Time (QQT) is self-managed by the customer by following the terms of this policy, notably the requirement to refrain from using resource-inefficient data processing practices as described in paragraph 1 of this Policy.
- Advance Notice of Significant Usage Changes:
 - Customer shall notify its Treasure Data Customer Success Manager at least five (5) business days before making any planned increase in query volume or complexity (excluding changes in profile/behavior counts) that could affect Treasure Data's ability to meet its published performance SLOs (Query Success Rate and Query Execution Consistency).
 - This advance notice enables Treasure Data to provision the necessary infrastructure to maintain service levels.
 - If Customer fails to provide such notice, Treasure Data may (a) defer non-critical workloads, (b) adjust service delivery, or (c) apply additional fees to cover the incremental resource requirements.