

PRODUCT DESCRIPTIONS & INTEGRATIONS

Table of Contents

Agentic Experience Platform	2
A. Overview	2
B. Platform Entitlements	3
C. Add-On Services	5
Treasure AI Suites	6
A. Personalization AI Suite	6
B. Engagement AI Suite	8
C. Creative AI Suite	10
CDP For Marketing	11
A. Applications	11
B. Treasure AI Platform	12
C. Add-On Services	15
D. Pricing Meter	17
Treasure AI Voice	25
Integrations	26
A. Input Integrations	26
B. Output Integrations	29

Agentic Experience Platform

A. Overview

Treasure AI's Agentic Experience Platform (or "AEP") provides an integrated suite of data, predictive models, automations, journeys, and activations designed to identify, score, and to create personalized experiences and target audiences.

The main components of AEP are described below. In addition, AEP comes with basic versions of Treasure Insights, Data Tank and Custom Scripts (each described under "Add-on Services" below); more advanced versions can be purchased separately.

Note: AEP was previously referred to as "Intelligent Customer Data Platform" or "ICDP."

Components	Description
Ingestion and Unification	The AEP allows users to ingest, cleanse, and unify data into customer profiles. See "Input Integrations" below for information on available integrations.
Segmentation and Activation	The AEP allows users to leverage the enriched profile data to create segments and target audiences with activations using integrations. The capabilities include an out of the box dashboard displaying summary statistics such as recent profile count in segment and growth rate. Users also have access to tools to analyze and segment profiles by attributes and select behaviors. See "Output Integrations" below for information on available integrations.
Journey Orchestration	The Journey Orchestration tool includes a drag-and-drop canvas for users to create journeys that are connected to each buying cycle stage across channels and business units. Features like Wait Steps can be used to send profiles in a series of activations with time delays in between. Decision Points make Journeys flexible, providing the ability to send profiles on a personalized journey by splitting the journey path based on profile attributes or behaviors.
AI Agent Foundry	The AI Agent Foundry is a service that enables data professionals to leverage large language models and offers capabilities to create GenAI agents in the Treasure AI ecosystem. The platform includes Policy Based Permission capabilities which enable access controls and data governance, facilitating user access and use of customer data in a privacy-conscious manner.

Components	Description
	AEP includes Audience Agent, which users can prompt to suggest marketing strategies, create segments, and analyze customer profiles. Audience Agent integrates with the Customer's Parent Segment Data within its AEP instance.

B. Platform Entitlements

Components	Features	Description
Data Management		
Data Lakehouse (Ingestion and Storage)	Batch Input Integrations	Provides the capability to use the out of the box bulk data integrations set forth in "Integrations" below to enable Customer to ingest data into Treasure AI's AEP.
	Streaming Input Integrations	Provides the capability to use the out of the box streaming data integrations set forth in "Integrations" below for streaming ingestion of data for Real Time (RT) use cases.
	Distributed analytical database	Provides access to Treasure AI's distributed columnar data storage system.
Data Engineering	SQL Query Engines	Provides access to Trino and Hive big data query engines for data analysis. Trino query engine is ideal for interactive query performance while Hive is used for high-volume batch processing.
	Python Custom Scripts	Provides the capability to run Python scripts in isolated containers from within Treasure Workflow.
	Treasure Workflow Data Orchestration	Provides the capability to manage data processing jobs, by defining the order in which processing tasks run, scheduling the runs, and monitoring the task status.
	Data Landing Area (DLA)	Data Landing Area (DLA) enables an intermediate SFTP storage for users to push their data from on-premises or other places.
ID Resolution and Data Enrichment	ID Unification	Provides the capability to create an ID graph using Treasure AI's out of the box deterministic algorithm.

Components	Features	Description
	Server Side Cookie	Server Side Cookie (SSC) enables users to identify and track visitors on a website. SSC allows for the setting of first-party cookies on the server side of a website domain instead of the client side.
Insight		
Business Analytics	Analytics/BI Output Integrations	Provides the capability to use out of the box integrations to enable Customer to connect data from the AEP to the analytics and business intelligence tools, as set forth in “Integrations” below.
Data Science (ML/AI)	Treasure Machine Learning Library	Provides Hivemall, Treasure Machine Learning library that runs on Apache Hive.
	Predictive Scoring (Self-serve Predictive ML)	Provides machine learning based methods to generate unmeasured values from known values. The new sets of values can be used in predictions of potential future customer outcomes.
Security Management		
Permissions	Access Control	Provides global access control for workflows and segmentations, and user-level access control for databases.
Authentication	Local User	Provides the capability to create user accounts with ID and password that are local to the AEP for the users to log in to the AEP.
	Single-Sign-On (SSO) at User Level	Allows users to use Google SSO and SAML 2.0 for login.
	IP Whitelisting	Enables the administrator to set up an IP whitelist to allow user login or user API Keys originated within certain IP ranges.
Audit Logs	Premium Audit Logs	Provides the capability to view and access audit logs which capture operations within the AEP.

C. Add-On Services

AEP comes with basic versions of each of the services described below. More advanced versions and/or increased quantities can be purchased separately.

Add-On Services	
Product	Description
Data Tank 2.0	Provides a hosted data mart for analytics, containing a subset of the data from the data lake. Typically used for business intelligence (BI), it offers easy access from external BI tools like Tableau.
Custom Scripts	Provides the capability to run Python scripts in isolated containers from within Treasure Workflow.
Treasure Insights	Provides data visualization (dashboard) capability through various charting options based on the data set, to facilitate deeper analysis and insights.

Treasure AI Suites

A. Personalization AI Suite

Personalization AI Suite		
Component	Feature	Description
Personalization Studio	Personalization Studio	Provides advanced web personalization capabilities by coupling enhanced profile data and real-time capabilities with a content rendering solution. Users can target audiences to receive personalized experiences, set specific rules based content targeting, and define the website areas for optimal content delivery.
Real Time Trigger Activation	Real Time Trigger Activation Engine	Provides a Decision Engine to create Real Time (RT) Profiles (known or unknown) via ID Stitching and make decisions based on Incoming Events, utilizing data stored in the RT Decision Engine. This data can include Attributes and Segment Membership from the AEP, plus RT Attributes including RT derived attributes. The RT Trigger Activation Engine stores this information, makes decisions and makes it available via the Trigger Activation Integration. One Trigger Activation Engine is required per Parent Segment and cannot be shared with Personalization.
	Trigger Ingest API	Provides an API Service to capture upstream events from multiple sources and route them to the Real Time Decision Engine and Core AEP.
	Trigger Activation Integration	Provides an integration service to initiate downstream actions using streaming connections, based on decisions made by the Real Time Decision Engine.
	Trigger Canvas	Provides a drag-and-drop canvas for users to create trigger activation logic based on data in the Real Time Trigger Activation Engine, which enables entry criteria, wait steps, activations, exit criteria, decision points and merges.
AI Signals	AI Signals	Provides access to the automated machine learning capabilities from Treasure Workflows. Using AI Signals, the processes to build models are automated.
	Next Best Product	The Next Best Product (NBP) solution is an AI-powered recommendation engine that generates personalized product suggestions. This enables companies to strategically recommend items, content and services that will resonate most with the customer. The output data that are associated with each profile can be used for immediate actionability within Treasure AI's ecosystem like in Audience Studio's Segment Editor or through Customer

Personalization AI Suite		
Component	Feature	Description
		Journey Orchestration and Personalization. Marketers are able to inspect and visually validate the predicted NBP data using intuitive data visualization dashboards. The personalized product suggestions are refreshed on a daily basis, ensuring that recommendations remain current and relevant to each customer's evolving preferences and behavior.
	RFM Model	The RFM solution is an analysis model that uses Recency, Frequency, and Monetary metrics to segment customers. It identifies high-value and at-risk groups based on recent activity, purchase frequency, and spending. Integrated within Treasure AI's ecosystem, it enables targeted marketing via tools like Audience Studio and Customer Journey Orchestration, supported by intuitive dashboards for data validation.

B. Engagement AI Suite

Engagement AI Suite		
Component	Feature	Description
Engage Studio	Engage Studio	Provides an intuitive, no-code message creation tool designed to help create, automate, and optimize personalized email and mobile push messages with AI Agent. Users can integrate email and mobile push notification (Google Firebase Cloud Messaging) with Treasure AI AEP to build targeted campaigns, execute personalized communications, and track performance.
Real-Time Trigger Activation	Real Time Trigger Activation Engine	Provides a Decision Engine to create Real Time (RT) Profiles (known or unknown) via ID Stitching and to make decisions based on Incoming Events, utilizing data stored in the RT Decision Engine. This data can include Attributes and Segment Membership from the AEP, plus RT Attributes including RT derived attributes. The RT Trigger Activation Engine stores this information, makes decisions and makes it available via the Trigger Activation Integration. One Trigger Activation Engine is required per Parent Segment and cannot be shared with Personalization.
	Trigger Ingest API	Provides an API Service to capture upstream events from multiple sources and route them to the Real Time Decision Engine and core AEP.
	Trigger Activation Integration	Provides an integration service to initiate downstream actions using streaming connections, based on decisions made by the Real Time Decision Engine.
	Trigger Canvas	Provides a drag-and-drop canvas for users to create trigger activation logic based on data in the Real Time Trigger Activation Engine, which enables entry criteria, wait steps, activations, exit criteria, decision points and merges.
AI Signals	AI Signals	Provides access to the automated machine learning capabilities from Treasure Workflows. Using AI Signals, the processes to build models are automated.

Engagement AI Suite		
Component	Feature	Description
	Next Best Product	The Next Best Product (NBP) solution is an AI-powered recommendation engine that generates personalized product suggestions. This enables companies to strategically recommend items, content and services that will resonate most with the customer. The output data that are associated with each profile can be used for immediate actionability within Treasure AI's ecosystem like in Audience Studio's Segment Editor or through Customer Journey Orchestration and Personalization. Marketers are able to inspect and visually validate the predicted NBP data using intuitive data visualization dashboards. The personalized product suggestions are refreshed on a daily basis, ensuring that recommendations remain current and relevant to each customer's evolving preferences and behavior.
	RFM Model	The RFM solution is an analysis model that uses Recency, Frequency, and Monetary metrics to segment customers. It identifies high-value and at-risk groups based on recent activity, purchase frequency, and spending. Integrated within Treasure AI's ecosystem, it enables targeted marketing via tools like Audience Studio and Customer Journey Orchestration, supported by intuitive dashboards for data validation.

C. Creative AI Suite

Creative AI Suite		
Component	Feature	Description
Creative AI Suite	Creative AI Suite	Provides a set of tools that leverage AI agents to automate various aspects of the creative content production.
	Email Copy	Provides an AI-powered agent to generate engaging email subject lines and body copy tailored to campaign goals and brand voice.
	Image Generation	Creates innovative and commercially safe marketing images from simple text descriptions.
	Image Remix Agent	Offers a suite of tools to modify existing images, including the ability to change styles and automatically remove backgrounds.
	Brand Guideline Agent	Automates the review of text content for compliance with pre-loaded brand and legal guidelines.
AI Signals	AI Signals	Provides access to the automated machine learning capabilities from Treasure Workflows. Using AI Signals, the processes to build models are automated.
	Next Best Product	The Next Best Product (NBP) solution is an AI-powered recommendation engine that generates personalized product suggestions. This enables companies to strategically recommend items, content and services that will resonate most with the customer. The output data that are associated with each profile can be used for immediate actionability within Treasure AI's ecosystem like in Audience Studio's Segment Editor or through Customer Journey Orchestration and Personalization. Marketers are able to inspect and visually validate the predicted NBP data using intuitive data visualization dashboards. The personalized product suggestions are refreshed on a daily basis, ensuring that recommendations remain current and relevant to each customer's evolving preferences and behavior.
	RFM Model	The RFM solution is an analysis model that uses Recency, Frequency, and Monetary metrics to segment customers. It identifies high-value and at-risk groups based on recent activity, purchase frequency, and spending. Integrated within Treasure AI's ecosystem, it enables targeted marketing via tools like Audience Studio and Customer Journey Orchestration, supported by intuitive dashboards for data validation.

CDP FOR MARKETING

(Application Pricing Model)

As an alternative to its Agentic Experience Platform (AEP) and AI Suites, Treasure AI offers CDP for Marketing as a suite of “*a la carte*” applications, with access to additional platform features.

A. Applications

The table below describes the features the applications offered with CDP for Marketing. Of these applications, Martech and Adtech Engagement are provided by default with every purchase of CDP for Marketing. Each of the other applications (Web Personalization, Journey Orchestration and Engage Studio) is offered separately with its own pricing.

CDP FOR MARKETING		
App	Features	Description
Martech and Adtech Engagement	Activations editor	Provides the capability to create activations from Audience Studio to export data to Treasure AI output connectors (as set forth in “ Integrations ” below).
	Output integrations	Provides the capability to use the out of the box output connectors as set forth in “ Integrations ” below).
Web Personalization	Profile API - Web	Provides an API service that returns profile-level segments and attribute data for web personalization.
	Real Time Segments	Provides the capability to return segment data from the Profile API - Web that is designed to be processed within seconds.
Journey Orchestration	Journey Canvas	Provides a drag-and-drop canvas for users to create journeys that are connected to each buying cycle stage across channels and business units.
	Wait Steps	Allows marketers to send profiles in a series of activations with time delays in between.
	Decision Points	Allows marketers to send profiles on a personalized journey by splitting the journey path based on profile attributes or behaviors.
Engage Studio	Engage Studio	Engage Studio is an intuitive, no-code message creation tool designed to help marketers create, automate, and optimize personalized email messages with AI Agent.

		With Engage Studio, users can seamlessly integrate email with Treasure AI CDP to build targeted campaigns, execute personalized communications, and track performance—all from a single interface. Key features include Content Editor, Email performance logging, Custom Template, Email Tracking, SSL & Domains.
--	--	--

B. Treasure AI Platform

Customers who purchase CDP for Marketing receive access to Treasure AI's platform, including the features described below. Platform access is offered at two tiers, Standard and Advanced. Features annotated with an "[A]" are available with Advanced Tier only.

Platform Feature	Description
Data Management	
Data Lakehouse (Ingestion and Storage)	
Batch input integrations	Provides the capability to use the out of the box bulk data connectors (as set forth in " Integrations " below) to enable Customer to ingest data into Treasure AI's CDP.
Streaming input integrations	Provides the capability to use the out of the box streaming data connectors (as set forth in " Integrations " below) for streaming ingestion of data for real-time use cases.
Distributed analytical database	Provides access to Treasure AI's distributed columnar data storage system.
Data Engineering	
SQL query engines (Presto, Hive)	Provides access to two big data query engines (Presto and Hive) for data analysis. Presto query engine is ideal for interactive query performance while Hive is used for high-volume batch processing.
Python custom scripts	Provides the capability to run Python scripts in isolated containers from within Treasure Workflow.
Treasure Workflow data orchestration	Provides the capability to manage data processing jobs, by defining the order in which processing tasks run, scheduling the runs, and monitoring the task status.
Data Landing Area (DLA) [A]	Data Landing Area (DLA) enables an intermediate SFTP storage for users to push their data from on-premises or other places.
ID Resolution and Data Enrichment	
ID Unification	Provides the capability to create an ID graph using Treasure AI's out of the box deterministic and probabilistic (fuzzy matching) algorithm.
Server Side Cookie	Server Side Cookie (SSC) enables users to identify and track visitors on your website. SSC allows Treasure AI to set first-party cookies on the server side behalf of your website domain instead of the client side. In consultation with Treasure AI Support, Customer may

Platform Feature	Description
	configure your DNS setting based on CDP requirements and data collection needs in order to initialize SSC. SSC is available with unlimited SSC access and unlimited SSC website domains.
Insights	
Business Analytics	
Analytics/BI output integrations	Enables Customer to connect data from the CDP to the analytics/BI tools described below under “Integrations”.
Data Science (ML/AI)	
Treasure machine learning library	Provides Hivemall, Treasure Machine Learning (ML) library that runs on Apache Hive.
Content affinity ML	Provides the capability to create profile-level content affinity attributes by analyzing a profile's behavior history and determining contextual categories the profile has expressed affinity towards.
Predictive scoring (Self-serve predictive ML)	Provides machine learning based methods to generate unmeasured values from known values. The new sets of values can be used in predictions of potential future customer outcomes.
Audience Studio Foundation / Engagement	
Segmentation	
Batch segmentation	Provides the capability to create rule-based batch segments and activate the attribute segment data to external tools for activities and to schedule activations on a recurring basis.
Segment view	Provides an out of the box dashboard displaying summary statistics such as recent profile count in segment and growth rate.
Actionable Profiles	
Customer profile view	Provides summary view of attributes and select behaviors associated with a profile accessible from the Audience Studio console.
Security Management	
Permissions	
Access control for databases, workflows, and segmentations	Provides global access control for workflows and segmentations, and user-level access control for databases.
Granular access control for workflow and segmentation [A]	Provides policy-based access control for workflows and workflow projects, parent segment, segments, and segment folders.
Policy-based permissions for database [A]	Provides policy-based access control for databases.
Authentication permissions [A]	Provides policy-based access control for Authentications.
Authentication	

Platform Feature	Description
Local user	Provides the capability to create user accounts with id and password that are local to the CDP for the users to log in to the CDP.
Google SSO at user level	Allows users to use Google SSO for login.
IP whitelisting	Enables the administrator to set up IP whitelist to allow user login originated within certain IP ranges.
SAML-based identity federation [A] (Azure, Okta, Ping Identity)	Allows integration with SAML-based identity providers (IdPs) for user login.
Audit Logs	
Standard audit logs	Provides the capability to view and access audit logs which capture basic operations within the CDP (last 24 hours).
Premium audit logs	Provides the capability to view and access granular audit logs which provide a detailed audit trail of the activities that occur in an account in the CDP.
Data Governance	
Data classification and column-level access control [A]	Provides the capability to classify data at the column-level with custom tags and column-level access control based on tags.
PII hiding in Audience Studio [A]	Provides the capability to label personally identifiable information (PII) attributes of a parent segment and restrict the view and use of PII attributes in Audience Studio.

The additional platform features described below may be purchased as add-on items.

Paid Add-On Platform Features	
Add-On	Description
Data Tank	A managed database service that customers can use as a hosted data mart for analytics alongside the main Treasure AI data lake. Data Tank allows Customer to prepare summary data for analytics from their raw tables and push it into their Data Tank 2.0 instance to share with a broader audience that may not have or need direct data lake access. Data Tank 2.0 offers easier access from external BI tools like Tableau, etc. and comes with a separately purchased high availability (HA) option to support automatic failover by having another instance that can take over operations if the primary instance becomes unavailable.
AutoML	AutoML provides access to the automated machine learning capabilities from Treasure Workflows. Using AutoML, the processes to build models is automated resulting in reduced time to deploy ML models.
AutoML Solution Notebooks	Solution notebooks provide out-of-the-box ML models, built using Treasure AutoML, executable from Treasure Workflows

C. Add-On Services

Under the Application Pricing Model, the following additional application-based Services may also be purchased separately by Customer:

AI Services		
Application	Features	Description
AI Agent Foundry	AI Agent Foundry	The AI Agent Foundry is a fully managed service that enables data professionals to easily leverage powerful foundational large language models and offers a wide set of capabilities to create GenAI agents in the Treasure AI ecosystem. This service helps organizations and professionals take advantage of generative AI without significant machine learning or infrastructure expertise. The platform includes Policy Based Permission capabilities which enable access controls and data governance, facilitating user access and use of customer data in a privacy-conscious manner. Customer and application data are not used to train any base models or shared with any third parties.
Audience Agent	CDP Data Integration	Audience Agent seamlessly integrates with the Customer's Parent Segment Data within its CDP instance and does not do any model fine-tuning or training on customer data, consistent with privacy-by-design best practices. PII-hiding is currently not supported for this application.
	Natural Language UX	Users can interact with Audience Agent in Audience Studio using natural language text with the same ease as instant messaging a colleague.
	Attributes and Behavior Analytics	Leveraging a RAG (Retrieval Augmented Generation) architecture, Audience Agent is able to find, interpret, and analyze attributes and behaviors and subsequently display these data using interactive and modifiable visualizations.
	Marketing Strategy and Execution Brainstorming	Audience Agent can suggest marketing opportunities and strategies based on customer data, develop step-by-step execution plans, and brainstorm additional

		modifications or enhancements to previously developed ideas.
	Segment Creation	Leveraging all of Audience Agent's capabilities, users can easily propose and create segments from the Audience Agent interface in Audience Studio, and immediately take action by using them throughout the Treasure AI ecosystem.

Real Time		
Application	Features	Description
Personalization (Real Time 2.0)	Real Time Personalization Engine	Provides a Decision Engine to personalize downstream systems, based on data in Real Time Profiles (known or unknown). Real Time Profiles can be existing or created via ID Stitching and can consist of stored Attributes and Segment Membership from the Core CDP, plus Real Time Attributes including Real Time derived attributes. The Real Time Personalization Engine stores this information, makes decisions and makes it available within 100 milliseconds via the Personalization API. One Personalization Engine is required per Parent Segment and cannot be shared with Trigger Activations.
	Personalization API	Provides an API Service for server-side integration which ingests an upstream event and returns a payload of the latest profile information including Real Time Attributes and Core DCP Attributes including Segment Membership which are stored in the Personalization Engine. Results are returned within 100 milliseconds at p95 (<i>requires Enterprise Tier Support SLA</i>).
	Personalization Canvas	Provides a drag-and-drop canvas for users to create website personalization logic based on data in the RT Decision Engine.
Trigger Activation (Real Time 2.0)	Real Time Trigger Activation Engine	Provides a Decision Engine to create Real Time Profiles (known or unknown) via ID Stitching and make decisions on them based on incoming streaming Trigger Ingest events, utilizing data stored in the Real Time Decision Engine. This data can include Attributes and Segment Membership from the Core CDP, plus Real Time Attributes including Real Time

		derived attributes. The Real Time Trigger Activation Engine stores this information, makes decisions and makes it available within ~3 minutes via the Trigger Activation Integration. One Trigger Activation Engine is required per Parent Segment and cannot be shared with Personalization.
	Trigger Ingest API	Provides an API Service to capture upstream events from multiple sources and route them to the Real Time Decision Engine and Core CDP.
	Trigger Activation Integration	Provides an Integration Service to initiate downstream actions using streaming connections, based on decisions made by the Real Time Decision Engine. End-to-end system response time is ~3 minutes.
	Trigger Canvas	Provides a drag-and-drop canvas for users to create trigger activation logic based on data in the Real Time Trigger Activation Engine, which enables entry criteria, wait steps, activations, exit criteria, decision points and merges.

D. Pricing Meter

The table below defines some of the key pricing meters (i.e., metrics) under the Application pricing model and explains how they are measured.

Pricing Meter	Period of time for measurement	Definition
App: Martech and Adtech Engagement (Default)		
Activations (Activated Profile)	Monthly	A Profile in the Treasure AI platform that is exported to and activated via Audience Studio, Data Workbench or any other Treasure AI application, to an external system, such as without limitations, an email service provider or ad network.
App: Engage Studio (Add-on)		
Emails Sent	Monthly	Emails Sent offers the sending of emails sent by a customer using Engage Studio within a given month. It reflects the scalability of the service and provides visibility into the customer's usage, helping ensure that the platform aligns with their communication needs.

Pricing Meter	Period of time for measurement	Definition
LINE OA Messages Sent	Monthly	LINE OA Messages Sent offers the sending of LINE OA messages sent by a customer using Agentic Engage within a given month. It reflects the scalability of the service and provides visibility into the customer's usage, helping ensure that the platform aligns with their communication needs.
Mobile Push Messages Sent	Monthly	Mobile Push Messages Sent offers the sending of Mobile Push messages sent by a customer using Agentic Engage within a given month. It reflects the scalability of the service and provides visibility into the customer's usage, helping ensure that the platform aligns with their communication needs.
App: Web Personalization (Add-on)		
Profile API - Web	Monthly	A call to the Profile API - Web with a user identifier, returning attribute(s), a segment or set of segments that the individual belongs to. For the purposes of reporting and compliance, the number of Profile API - Web calls for a given month will be the total number of Profile API - Web calls per token recorded in the CDP for the given month.
Platform		
Profile	Monthly	A Profile is the result of ID unification and contains a collection of data about a known or unknown individual's behaviors and attributes, which may include personally identifiable information such as an email address, phone number or digital ID such as a cookie or mobile advertising ID. <i>100 B2C Known Profile = 1 B2B Known profile 1 B2C Known Profile = 20 Unknown Profile</i> <i>Unless stated otherwise, any reference to Known Profiles in the Order Form will mean B2C Known Profiles.</i> <i>For the purposes of reporting and compliance, the number of Profiles for a given month will be the maximum number of Profiles held in databases in the CDP.</i>
Known Profile, Consumer	Monthly	(a.k.a., B2C Known Profile) This type of Profile consists of behavioral or attribute data that contains personally identifiable information such as an email address or phone number and is used when Customer is using the Service for its business-to-consumer enterprise.
Unknown Profile	Monthly	This type of Profile consists of behavioral or attribute data that does not include personally identifiable information, such as an email address or phone number.

Pricing Meter	Period of time for measurement	Definition
Known Profile, Business	Monthly	(a.k.a., B2B Known Profile) This type of Profile consists of behavioral or attribute data that contains personally identifiable information such as an email address or phone number and is used when Customer is using the Service for its business-to-business enterprise.
Activations (Activated Profile)	Monthly	A Profile in the Treasure AI platform that is exported to and activated via Audience Studio, Data Workbench or any other Treasure AI application, to an external system, such as without limitations, an email service provider or ad network.
Compute Units	Monthly	Compute Unit is the basic unit of compute (vCPU and memory) needed to run Customer's queries concurrently, by considering the various stages of the query execution and resources needed to execute the full query plan. Compute Units are how Treasure AI allocates and measures consumption of Customer's monthly resources for compute tasks (e.g., SQL queries for profile stitching, scanning some volume of data from an input table, batch queries). Consumption relates to Customer's users' requirements for instantaneous computing headroom or elapsed time for batch queries (how much computing can be done simultaneously) and total computing use (how much can be done over the course of a month). Compute Units are credits bought for processing data. Purchased Compute Units can be used for either Presto, best used for interactive SQL queries, or Hadoop, used for Hive batch SQL queries as well as some other functions in the CDP. <i>1 Compute Unit = 1 Presto Unit = 2 Hive (Hadoop) Units.</i> <i>Each Presto Compute Unit comes with 10GB per query memory and 15,000 Presto Split-Hours. Each Hadoop Unit comes with 2 cores. Systems limitations exist for the Presto Unit and the Hive Unit and are as described in the Treasure AI Documentation.</i> In consultation with Treasure AI Support, Customer may allocate Compute Unit resources to Presto or Hive based on CDP requirements and performance needs. For example, a customer who has purchased 80 Compute Units may allocate those units to 40 Presto Units and 80 Hive Units. More details about how Compute Units, Presto Units and Hive Units are calculated and parameters that impact them may be found in Treasure AI's Documentation at https://docs.treasure.ai/.

Pricing Meter	Period of time for measurement	Definition
Ingest Rows	Monthly	The number of rows ingested by or on behalf of Customer to the CDP from all channels, such as, without limitations, web, mobile, ERP, CRM, or external databases. For example, a CSV file containing 10,898 rows of data will increase the monthly measured Ingest Rows by 10,898 rows. Each row assumes an average compressed row size of 75 bytes. <i>For the purposes of reporting and compliance, the Ingest Rows for a given month will be the maximum number of rows stored in the CDP on any given day within the month and the sum of all those daily values during the given month.</i>
Custom Scripts	Monthly	Custom Scripts enable users to run Python scripts from within workflows in secure isolated containers and provide for greater flexibility of custom logic for building data pipelines, transformations and running production ready ML models. Treasure AI offers 4 levels of Custom Scripts, which correspond to the resource configuration and total number of allocated runtime hours for the scripts: Level 1 (offered as default with every purchase of CDP for Marketing): • 1 vCPU • 8GB Memory 20GB Disk Space • 100 hours of runtime for custom scripts Level 2: • 4 vCPU • 30GB Memory 20GB Disk Space • 500 hours of runtime for custom scripts Level 3: • 8 vCPU • 60GB Memory 20GB Disk Space • 500 hours of runtime for custom scripts Level 4: • 16 vCPU • 120GB Memory 20GB Disk Space • 500 hours of runtime for custom scripts <i>Note: Customer has option to purchase an add-on of 500 hours of runtime for custom scripts under each of the levels stated above at an additional fee.</i>

Pricing Meter	Period of time for measurement	Definition
Add-On		
Treasure Insights - Data Modeler	Monthly	A pricing meter used in measuring the consumption of Treasure Insights, a Data Modeler can perform the following functions: (i) select and push data from the CDP to Treasure Insights; (ii) create and maintain data models; (iii) share data models with other Data Modelers within the CDP; (iv) share dashboards that the Data Modeler has edit permissions for; and (v) have capabilities belonging to a Viewer.
Treasure Insights - Viewer	Monthly	A pricing meter used in measuring the consumption of Treasure Insights, a Viewer can perform the following functions: (i) access dashboards that are shared with the Viewer; (ii) filter data in the dashboards; and (iii) download dashboards in an Adobe PDF format.
Treasure Insights-Infrastructure Block	Monthly	(a.k.a. "TI Block") A pricing meter used in measuring the consumption of Treasure Insights, an Infrastructure Block is 5GB of storage on the CDP.
Data Tank 2.0	Subscription Term	A managed database service that customers can use as a hosted data mart for analytics alongside the main Treasure AI data lake. Data Tank 2.0 allows Customer to prepare summary data for analytics from their raw tables and push it into their Data Tank 2.0 instance to share with a broader audience that may not have or need direct data lake access. Data Tank 2.0 offers easier access from external BI tools like Tableau, etc. and comes with a separately purchased high availability (HA) option to support automatic failover by having another instance that can take over operations if the primary instance becomes unavailable. Data Tank 2.0 is available in the following configurations in a single instance: • Small: 16GB Memory, 80GB Storage, 60 R/W* per second • Medium: 32GB Memory, 320GB Storage, 240 R/W per second Data Tank 2.0 is available in the following configuration in the high availability option: • Small: 16GB Memory, 80GB Storage with HA (2 availability zones), 60 R/W per second • Medium: 32GB Memory, 320GB Storage HA (2 availability zones), 240 R/W per second • Large: 128GB Memory, 1024GB Storage with HA (2 availability zones), 960 R/W per second • Extra Large: 256GB Memory, 4096 GB Storage with HA (2 availability zones), 3072 R/W per second If more read/write per second storage is needed, Customer may purchase an add-on option as follows: • Tier 1: (40GB Storage, 30 R/W per second), only available with Small • Tier 2: (160GB Storage, 120 R/W per second), only available with Medium

Pricing Meter	Period of time for measurement	Definition
		- Tier 3: (512GB Storage, 480 R/W per second), only available with Large - Tier 4: (2048GB, 1536), only available with Extra Large <i>Note: R/W means Read/Write</i>
AutoML	Subscription Term	Automates machine learning (ML) process steps such as data processing, feature engineering, model selection and hyper-parameter tuning to reduce time to deploy ML models. Treasure AutoML provides users access to the AutoML capabilities from Treasure Workflows. AutoML is available in the following options: AutoML Tier1: - Up to 256GB RAM - Maximum runtime hours: 500 Unit Hours 4 AutoML tasks can be run concurrently 5 TD users can run AutoML tasks AutoML Tier2: - Up to 384GB - Maximum runtime hours: 1000 Unit Hours 8 AutoML tasks can be run concurrently 15 TD users can run AutoML tasks AutoML Tier3: - Up to 512GB RAM - Maximum runtime hours: 2000 Unit Hours 16 AutoML tasks can be run concurrently - Unlimited TD users can run AutoML tasks In addition to basic AutoML, solution notebooks provide out-of-the-box ML models, built using Treasure AutoML, executable from Treasure Workflows. Solution notebooks are separately purchased from Treasure AutoML. Option to purchase an additional bundle of 2 Solution Notebooks at an additional fee.
Real-Time		
Real Time Profiles (B2C)	Monthly	A Real Time Profile is a collection of Core CDP attributes, Real Time Attributes and Behaviors and RT Decisions related to a single individual held in the Real Time Decision Engine. Each Real Time Profile may have up to a combined two hundred (200) Core CDP Attributes, Real Time Attributes, Real Time Behaviors and Real Time Decisions, with a maximum size of 500 bytes per profile. Each Real Time Profile is counted as one (1) Real Time Profile, regardless of if it is a Known Real Time Profile (has Personally Identifiable Information) or an Unknown Real Time Profile (no PII).

Pricing Meter	Period of time for measurement	Definition
		For the purposes of reporting and compliance, the number of Real Time Profiles for a given month will be the maximum number of Real Time Profiles held in the Real Time Decision Engine at any point during the month. To accommodate profile retention over time, this is maximum number of Real Time Profiles stored.
Incoming Trigger Event	Monthly Average	An incoming streaming event that is ingested into the Real Time Trigger Engine. The Monthly Average number of events is the sum of 12 months of events, divided by 12 starting at the contract effective date. Overage occurs once 12x monthly average is reached during any 12-month period. For reporting and compliance purposes, the number of Incoming Trigger Events is the number of Incoming Trigger Events recorded in the Real Time Trigger Engine(s) for a given month. This number includes Incoming Trigger Events in any Environment (Production, QA, Development, etc.)
Outgoing Trigger Activation	Monthly Average	An outgoing streaming event that is generated by the Real Time Trigger Engine, which can result in a call to downstream streaming connections for an action to be taken. Response time from Incoming Trigger Event to Outgoing Trigger Activation is approximately three (3) minutes. The Monthly Average number of events is the sum of 12 months of events, divided by 12 starting at the contract effective date. Overage occurs once 12x monthly average is reached during any 12-month period. For reporting and compliance purposes, the number of Outgoing Trigger Events is the number of Incoming Trigger Events recorded in the Real Time Decision Engine(s) for a given month. This number includes Outgoing Trigger Activations in any Environment (Production, QA, Development, etc.)
Personalization Calls (Monthly Average Users)	Monthly	A call to the Personalization API which using a real-time user identifier, returns a payload of up to 20 RT Attributes and Core CDP attributes which are stored in the Real Time Decision Engine. Note: The 20 attribute limit is necessary for the system SLA to be met. The Monthly Average number of events is the sum of 12 months of events, divided by 12 starting at the contract effective date. Overage occurs once 12x monthly average is reached during any 12-month period. For the purposes of reporting and compliance, the number of Personalization Calls for a given month will be the total number of Personalization Calls recorded in the Real Time Decision Engine(s) for a given month. This number includes Personalization Calls in any Environment (Production, QA, Development, etc.)

Pricing Meter	Period of time for measurement	Definition
Treasure AI Foundry		
Conversations	Monthly	A Conversation is the basic metric of usage of Treasure AI Foundry. A session that consists of 5 back-and-forth's will be metered as a single Conversation using Treasure AI's base Large Language Model, Claude 4.5 Sonnet. Shorter sessions will be metered as a fraction of a Conversation. For longer sessions, the metering rate (i.e., the ratio of back-and-forth's to Conversations) will decrease below 5.0 if and to the extent the number of back-and-forth's exceeds 5 (for example, a session with 10 back-and-forth's would be metered as 2.87 Conversations). Conversations are metered on all Production Instances.

The Monthly count is defined as: sum of all daily counts in the given month.

References to terms can be found in the Treasure AI Documentation at <https://docs.treasure.ai/>

Treasure AI Voice

Treasure AI Voice leverages generative AI technology to enable authorized users to capture, transcribe, and summarize audio recordings (for example, from in-person meetings or conference calls) using compatible voice capture devices and mobile applications. The solution consists of the following features, capabilities and components.

Component	Description
Voice Capture Device	The voice capture device is the hardware that enables Treasure AI Voice. Manufactured by Plaud LLC, the devices allow users to capture audio from conversations, meetings, calls, and other interactions, and to send the recordings to their compatible mobile devices. Plaud voice capture devices are covered by a 12-month manufacturer's warranty.
Treasure AI Voice Application	Allows the user to pair a compatible mobile device with their Voice Capture Device. Through the application, the user can initiate, view and manage recordings and upload captured audio to the Treasure AI Voice hosted service for processing. The mobile application is available for free download through select app stores.
Treasure AI Voice Hosted Service	Leverages artificial intelligence to perform cloud-based speech-to-text transcription on uploaded audio (including timestamps and speaker identification), and to generate summaries (with information such as topics, speakers, decisions, action items, etc.) and various metadata. Provides short-term storage of uploaded audio, transcripts, summaries and other outputs, subject to capacity, retention, and deletion settings configured by Customer and Treasure AI's standard technical limits and backup practices.
Web-Based Administrator Console	A web-based admin console allows designated administrator(s) to activate and resource an account for Customer's organization; create, delete and manage user accounts; enroll Devices into the organization's account and assign Devices to users; monitor usage metrics; and configure other settings within Treasure AI Voice.
Security Management Permissions	Provides global and user-level access control for Treasure AI Voice.

Integrations

As of the date of this document, Treasure AI makes generally available to its customers input and/or output integrations with the third-party services listed below, subject to the terms and conditions of Customer's service agreement with Treasure AI ("Agreement") and the following:

- Customer must separately contract with the applicable third-party provider of each listed application for the right to use such application and to import/export data to/from such application as contemplated by the Agreement.
- Because access to each third-party application is not entirely within Treasure AI's control, Treasure AI cannot guarantee that any listed application will continue to be available for interconnection with Treasure AI's platform throughout Customer's subscription term.

In addition to the listed integrations, Customer may also input data through the following mechanisms:

- *For batch input:* (1) manual upload (CSV, TSV, JSON); (2) thin client to push data to the service (e.g., Embulk, Database JDBC/ODBC Drivers); and/or (3) use of Data Landing Area; and
- *For streaming input:* (1) server-side cookies; and (2) server-side agent (e.g., streaming log data).

A. Input Integrations

Type	Category	Integration
Connectors (Batch)	<i>Advertising</i>	• theTradeDesk • Google My Business • Google Audience Partner API Status via SFTP • Google Ad Manager via Audience Partner API • Google Ads • Facebook (Lead Ads) • Eyeota • Brandwatch
	<i>Business Intelligence</i>	• Google Sheets
	<i>Cloud Storage and File Storage</i>	• Amazon S3 • FTP • Box • SFTP • OneDrive • Microsoft Azure • Microsoft Azure Data Lake • Karta • Google Drive • Google Cloud Storage

Type	Category	Integration
	<i>CRM</i>	• Bombora • Gigya (SAP Customer Data Cloud) • Intercom • Microsoft Dynamics • OneTrust • Salesforce • Sansan
	<i>Customer Engagement Platform</i>	• Airship
	<i>Customer Support</i>	• Jira • Zendesk
	<i>Database</i>	• Amazon Redshift • Google Big Query • Microsoft SQL Server • Mongo DB • MySQL • Neighborly (Not on UI) • PostgreSQL • Snowflake
	<i>DMP</i>	• Salesforce DMP Krux • Repro
	<i>Ecommerce</i>	• Magento • Salesforce Commerce Cloud • Shopify • Vtex
	<i>Email Marketing</i>	• Cvent • DotDigital • HubSpot • Iterable • Mailchimp • Mailpublisher Smart • Marketo • Oracle Eloqua • Salesforce Marketing Cloud
	<i>ERP</i>	• Netsuite
	<i>Payment</i>	• Stripe • Zuora
	<i>Servers</i>	• Riak CS
	<i>Social Media</i>	• Facebook (Ads Insights) • Facebook (Lead Ads) • Facebook (Page Insights) • Instagram User and Media

Type	Category	Integration
		• Sprinklr • Twitter Tweet Insights • YouTube
	<i>Web/Mobile Analytics Services</i>	• Adobe Analytics • Amplitude • App Annie • Apptopia • Dynalyst • Google Analytics • Google Search Analytics • Mixpanel
SDKs	<i>Web SDK</i>	• JavaScript SDK •
	<i>Mobile App SDK</i>	• Android • iOS
	<i>Gaming SDK</i>	• Unity SDK (mobile and PC) • Unreal Engine SDK (mobile and PC)
	<i>Client SDKs / APIs / Webhooks</i>	• Python • Ruby • Java
	<i>In-Store SDK</i>	• In-Store SDK
IOT	<i>IoT streaming</i>	• Apple TV (iOS) • Arduino (Fluentd) • C / C++ SDK (Fluent Bit) • Mbed OS (Fluentd) • RaspberryPi (Fluentd)
Connectors (Streaming)	<i>Martech streaming connector</i>	• Salesforce Marketing Cloud Streaming Import Integration • Azure Event Grid
	<i>IaaS/PaaS streaming connector</i>	• Amazon Kinesis

B. Output Integrations

Category	Integration
<i>Business Intelligence</i>	• Birst (Presto JDBC) • Chartio (Presto connection) (TD Export) • Domo (Presto connection) • JasperReports (Presto JDBC) (TD Export) • JasperSoft (Presto JDBC) • Looker (JDBC postgres - Data Tanks) (TD Export) • Metabase (Presto connection) (TD Export) • Microsoft Excel (Presto ODBC) • Microsoft Power BI (plugin) (TD Export) • Mode Analytics (Presto connection) (TD Export) • Pentaho (Presto JDBC) (TD Export) • QlikView, QlikSense (Presto ODBC) (TD Export) • Redash (Presto connection) • Slemma (Presto connection) • Tableau (plugin) (TD Export)
<i>Object Storage Systems [Cloud Storage and File Storage]</i>	• Amazon S3 • FTP • SFTP • Microsoft Azure • Korte • Google Cloud Storage
<i>Databases</i>	• Amazon Redshift • Elastic Cloud (Output only) • Google Big Query • Microsoft SQL Server • Mongo DB • MySQL • Neighborly (Not on UI) Input Only • PostgreSQL • Snowflake
<i>DMP</i>	• Adobe Audience Manager • Audience One • Beeswax (Liveramp) • Facebook (Direct Onboarding - Liveramp) • Facebook (Liveramp) • Google Display & Video 360 (Liveramp) • IntimateMerger • Liveramp • Oracle BlueKai • Twitter (Liveramp) • Verizon Media • Yahoo! DMP

Category	Integration
<i>Ecommerce</i>	• Shopify Metafield
<i>Email Marketing</i>	• Cuenote • Loopass • Oracle Responsys • Sailthru ESP • Sendgrid • Webcas
<i>Marketing Automation</i>	• Oracle Netsuite Bronto • Salesforce Pardot
<i>Social Media</i>	• Facebook (Conversion API) • Facebook (Custom Audiences) • Facebook Offline Conversion • Instagram Custom Audiences • Line • LinkedIn • Snapchat • Twitter Tailored Audiences
<i>Advertising</i>	• Amazon DSP (plugin) • Braze (plugin) • Criteo (plugin) • Firebase Cloud Messaging (plugin) • Google DoubleClick Bid Manager (Now: Display & Video 360) (plugin) • Nend (plugin) • Supership (plugin) • Reddit Custom Audience Integration¹ • Yelp Conversions API (attribution services)²

¹ Use of this integration is further subject to the “Reddit Ads API Terms and Conditions” (available at <https://business.reddithelp.com/s/article/Reddit-Ads-API-Terms>).

² Use of this integration is further subject to the “Yelp Attribution Agreement” (available at https://terms.yelp.com/attribution/20231206_en_us/).