PromethistAI User’s Documentation

PromethistAI a.s.

Version 1.0

Table of Contents

Platform Introduction

Chapter 1. Introduction

Welcome to the **PromethistAI Documentation** – the central resource for understanding, building, and scaling relational experiences within the **PromethistAI Platform**. This guide shows you how organizations can leverage **Relational Agents** to transform the way they engage with customers, partners, and teams.

It is designed to bring together **relational intelligence** and **business process execution** in one seamless system, enabling agents that are both **empathetic** and **effective**. The platform empowers enterprises to create agents that don't just answer questions but **build relationships, influence behavior, and drive business outcomes**.

Chapter 2. Purpose of This Manual

This documentation exists to help you:

- Understand the **strategic value** of Relational Agents and how they fit into your business workflows.
- Learn how to **design, configure, and deploy Relational Agents** with clear business goals and relational nuance.
- Connect enterprise assets – from **knowledge bases** to **integrations** – into live experiences that reflect your brand, internal processes, and overall CX strategy.
- Track and measure impact using **analytics** that connect agent interactions to business KPIs.

By following this manual, your teams can move from experimentation to enterprise deployment with confidence, clarity, and measurable value.

Chapter 3. Who This Guide Is For

PromethistAI serves multiple stakeholders across the enterprise. This guide is written for:

- **Business & Strategy Leaders** – to understand how relational intelligence can improve customer satisfaction, reduce churn, and increase conversions.
- **Product & Program Managers** – to translate business objectives into agent capabilities, use cases, and measurable outcomes.
- **Designers & Conversation Architects** – to shape agent personalities, tone, and workflows that resonate with users.
- **Engineers & Integrators** – to securely connect external tools, APIs, and systems via integrations.
- **Administrators & IT Owners** – to configure accounts, projects, security, and compliance.
- **Analysts & Operations Teams** – to evaluate performance, governance, and agent-driven business impact.

If you are part of an organization deploying relational technology into your customer experience ecosystem, this documentation is for you. PromethistAI allows you to move beyond static knowledge bots toward **relationship-driven digital partners**. It's designed to evolve alongside customer needs. Your feedback shapes our roadmap. If you encounter gaps or challenges, reach out through your enterprise support channel or use the feedback form embedded in the docs.

Chapter 4. Interacting with the Platform

PromethistAI offers two ways to work with the platform:

- **Studio** – Standard interface for detailed configuration, asset management, and administrative control
- **AI Assistant** – Create agents and perform tasks through a chat window

You can switch between modes at any time depending on your workflow. To explore AI Assistant, see AI Assistant.

Chapter 5. Glossary of Terms

The glossary provides **canonical definitions** of the most important concepts in the PromethistAI Platform. It ensures teams across roles—admins, designers, engineers, and business leaders—share a common understanding of the platform's core building blocks.

5.1. Accounts & Structure

PromethistAI Platform	The administration console and central hub for managing accounts, projects, agents, knowledge bases, integrations, and governance.
PromethistAI Account	An organisational container governing projects, billing, and members. Accounts define ownership and access boundaries across the platform.
PromethistAI Project	A workspace within an account. Projects hold agents, assets, knowledge bases, interactive content, and integrations.

5.2. Clients & Agents

PromethistAI Clients	The applications through which end-users interact with Relational Agents: the iOS and Android apps, the browser-based Web Client, embedded widgets, and kiosk installations.
Relational Agent	A digital agent built on PromethistAI that combines **relational intelligence** (empathy, adaptability, natural conversation) with **business-specific customisations** (identity, objectives, processes and guardrails).
Relational Intelligence Engine	The runtime engine that balances business process execution with relational intelligence, orchestrating how the agents respond in real time.
Avatar	The 3D visual representation of a Relational Agent in the PromethistAI Clients.
Environment	The background scene behind an agent's avatar in the PromethistAI Clients.
Voices	Audio assets used for speech interaction, chosen from multilingual presets or TTS models.
Agent URL	Unique URL address allowing reference to a particular agent (optionally with a specific revision number). Enables distribution via channels such as websites, emails, and messages. Can also be encoded as a QR code for quick access in the PromethistAI client app.

5.3. Assets & Interactions

Knowledge Bases	Repositories of content (e.g. documents, FAQs, product data, troubleshooting guides, web sources). Knowledge bases provide agents with factual grounding.

Integrations	Connections to external servers and APIs using the **Model Context Protocol (MCP)**. These integrations extend agents with enterprise capabilities.
Interactive Content	Structured inline experiences in conversations (forms, images, buttons, web views) that let agents guide users beyond text or voice.
Business Process	The structured rules and workflows that govern how an agent operates, ensuring consistency, compliance, and alignment with enterprise objectives.

5.4. Usage & Billing

Conversational Minute	The billing unit in PromethistAI. Charges accrue based on the number of minutes agents interact with the users.
Billing Model	All usage is calculated at the account level. **Conversational Minutes** aggregate across projects and agents to determine overall billing.

NOTE This glossary focuses on the core terms that appear most often in the platform UI, configuration, and deployment workflows.

Chapter 6. Onboarding & Account Management

Account-level onboarding covers setting up your PromethistAI account, choosing a processing region, inviting members, and managing billing.

For step-by-step instructions, see:

- Account Management
- Managing Members
- Billing and Best Practices

Chapter 7. What are Relational Agents?

7.1. Relational Agent

A Relational Agent is the central entity within the PromethistAI Platform. Each agent is powered by two complementary dimensions:

Relational Intelligence	The default intelligence engine that powers empathetic conversation, adaptability, and user-centric dialogue. This layer provides the psychological and behavioral capabilities that make the agent feel natural, trustworthy, and responsive.
Custom Configuration	The workflows, guardrails, structured rules, and the explicit objectives that are specific for the particular task the agent is intended to execute. This ensures the agent always acts in alignment with organizational standards and compliance requirements.

These two dimensions together enable the Relational Agent to both perform the specific job they are intended for and go beyond simple task fulfilment when needed. They can execute business processes while building meaningful relationships with users.

[Relational Agent Configuration] | *relational-agent-configuration:relational-agent-configutarion.png*

Chapter 8. Creating & Configuring Agents

8.1. Creating a New Relational Agent

8.1.1. From Scratch vs. Templates

When you start creating a new agent, you can choose one of the pre-filled templates or start from scratch.

Templates	A good choice for newcomers. Templates come pre-configured and ready to use. Each template has its own **Template Variables** — a set of fields you fill in to tailor it to your context.
From Scratch	If you don't find a template that fits your use case, you can build an agent from scratch.

Template Variables

When you create an agent from a template, you fill in its **Template Variables** to tailor it to your context.

You can use **@Reference** in any Template Variable to reference a project asset — such as a Knowledge Base, Interactive Content, Integration, or Skill — directly inside the field. Click **Reference** (or type `@`) and pick the asset you want to insert. See Assets Referencing for the asset types you can reference.

8.1.2. Agent Type

After choosing Template or From Scratch, pick an **Agent Type**:

Relational Agent	Leverages advanced voice capabilities and our unique parallel thinking processing — perfectly fitting complex use cases and deep interactions.
Lightweight Agent	Uses efficient voices with a smaller reasoning footprint — ensuring significantly cheaper token spend while staying reliable.

WARNING	Agent Type cannot be changed after the agent is saved. If you need a different type, you'll have to create a new agent.

Lightweight agents share most of the same configuration as Relational agents, but a few sections do not apply — for example, **Business Process Steps** is not available on Lightweight agents.

8.2. Agent Evaluation

Optionally attach an evaluation to the agent. Use the search field to find and select an existing evaluation, or create one in Evaluations.

Buttons in this section let you jump to evaluation management without leaving agent setup — open

the **Evaluations** page to browse what already exists, or start a new evaluation directly from here.

8.3. Agentic Purpose

NOTE

What is Agentic Purpose for?

This is where you tell the agent what it's supposed to do, how it should behave, and what path to follow in a conversation. Think of it as writing a job description and a script — combined.

It has three parts:

- **Objective** — the goal. What is this agent trying to achieve? What should the user walk away with?
- **Guardrails** — the rules. What topics or behaviours are off-limits? What should the agent always or never do?
- **Business Process Steps** — the flow. The ideal step-by-step path through the conversation, from opening to close.

The agent uses all three together. It will follow the steps, stay within the guardrails, and keep the objective in mind throughout — adapting its tone and approach based on how the conversation unfolds.

TIP

All text fields here support rich text formatting, headings, lists, bold, and more. Well-structured instructions are easier for the agent to follow.

[Rich text formatting toolbar] | *markdown.png*

8.3.1. Objective

Define the primary and subsidiary objectives.

- Primary objectives are broad and overarching
- Sub-objectives break down how the agent achieves the primary goal

Example:

- Your primary objective is to inspire the client to love our brand.
- Your job consists in making the clients familiar with your brand's intentions and the most recent product offer.
- In order to do your job well, you should inquire into the client's needs and present the brand in a fresh way.
- Avoid pushing sales too hard if the client is hesitant. Focus on personalized engagement.

8.3.2. Guardrails

Rules and boundaries agents must obey.

- Cover sensitive topics requiring expertise (e.g., psychological treatment, financial advisory)
- Ensure non-standard customer requests are safely redirected back to the business process path

WARNING	When two or more languages are selected in **Allowed Languages**, do not add a guardrail that restricts the agent to speak only one language. This conflicts with the multilingual setup and can cause problems with speech recognition.

8.3.3. Business Process Steps

NOTE	**Business Process Steps** is available on Relational agents only — it does not appear in the configuration for Lightweight agents.

This is the "golden path" of the ideal interaction. Write it step-by-step.

Step 1	Greet the client and explain the advantages of your offer.
Step 2	Learn more about the client’s needs.
Step 3	If the client seems interested, propose the best-fit product. If not, inspire interest with unexpected product features.
Step 4	Direct the client to more information and take a memorable farewell.

8.3.4. Assets Referencing

You can reference project assets by name using @mentions in **Objective**, **Guardrails**, and **Business Process Steps**.

Knowledge Bases	Ground the agent’s responses in specific content.
Interactive Content	Trigger structured interactions like forms, images, or web views.
Integrations	Call external services during or after the conversation.
Skills	Define what is being trained. In the **Objective**, reference Skills to establish the skill catalog for the session. In the **Business Process Steps**, reference a skill to select what is being practised and route to the right scenario. Only Skills available in your project can be referenced.

- **Global assets** – Available automatically. Mentioning them explicitly is optional but improves performance.

Step 5	At the end of the interaction, ask the user how they were satisfied and display the multimodal `@product_feedback_form`

8.3.5. Relational Settings

Determine how strictly the Relational Agent follows the business process vs. focusing on relationship-building.

Task Focus	Strict execution of process. *Example*: Mortgage procedure with no deviation allowed.
Relational Focus	Conversation-driven, lighter process enforcement. *Example*: Brand loyalty and empathy-driven conversations.
Balance	Combines task and relational approaches. *Example*: Informing clients about new offers while still executing specific sub-goals.

8.3.6. Remembers users (memory control)

What it is

A simple toggle that controls whether a Relational Agent should recall previous conversations with the same user.

When to use it

Disable for one-off, shared, or public interfaces (e.g., kiosks) where each conversation should be treated as new. Enable for relationship-based use cases (e.g., account management, support follow-up) where continuity matters.

How to configure

Go to **Agent configuration** → **Agentic Purpose** → **Remembers users**.

- **On (enabled)** – The agent remembers and can refer to past interactions
- **Off (disabled)** – Every conversation is treated as a new interaction

NOTE	Agent-specific setting. For concepts, policies, and governance, see Memory & Relational Intelligence.

8.4. Agent's Identity

8.4.1. Avatar

Select the look of your Relational Agent. The **Avatar Source** determines how the avatar is produced:

On-device rendered	The Promethist engine renders the avatar locally from the available presets. No external service required.
LiveAvatar stream	Displays a live, photorealistic video feed from a HeyGen LiveAvatar instead of rendering locally. Requires a **Live Avatar ID**.

When you select **LiveAvatar stream**, enter the **Live Avatar ID** of the avatar from your HeyGen Live Avatar account — for example, `a2ccb2cf-778d-4f22-ab45-5a4b9e100ed5`.

LiveAvatar stream also needs a **Live Avatar API Key**, set in Project Settings or Account Settings. When a key is set at both levels, the project-level key takes precedence.

8.4.2. Environment

Select the environment in which the Relational Agent appears. The environment sets the visual background for the interaction.

8.4.3. Camera Mode

Control how the camera behaves during a conversation.

Dynamic	The camera moves during the interaction for a more cinematic feel.
Static	The camera stays fixed, similar to a video call.

8.4.4. Voice

Select the voice you want your agent to speak with. Each voice has its own set of supported languages, which determines the options available in **Allowed Languages**.

8.4.5. Allowed Languages

By default, all languages supported by the selected voice are enabled. If you don't need all of them, remove the ones you don't want.

[Voice and Allowed Languages] | *voice.png*

> Your voice supports *American English*, *Czech*, and *German*, but you only want the agent to speak American English and Czech. Remove German from **Allowed Languages**.

When a conversation starts, the platform matches the agent's allowed languages against the **Language** setting on the user's device (app) to decide which language the agent speaks:

- The **primary language** is the highest-priority match between the two lists.
- Any remaining matches become additional languages the agent can also speak.

> The agent's allowed languages are *American English* and *Czech*, and the user's app is set to *Czech, English*:
>
> - Primary language → **Czech**
> - Additional language → **American English**

8.4.6. Character & Backstory

Describe the personal traits of your digital employee.

Elements can include:

- Character traits, style of expression, and backstory
- Previous experiences, likes/dislikes, behaviors in specific situations
- Personal views on topics

NOTE	A well-written character and backstory aligned with the agent's role makes interactions feel natural and memorable.

8.5. Revisions

Every time you save changes to your agent's configuration, a new revision is automatically created. This allows you to track changes and revert to previous versions if needed.

For details on managing revision states and reverting, see Testing & Publishing.

Chapter 9. Our Memory

Our Memory transforms interactions with Relational Agents by mapping key details from conversations into a dynamic Knowledge Graph. This system uses Relational Intelligence, our underlying system for handling context in interactions. It identifies interests, goals, and facts, then connects them to form patterns, making responses more relevant over time. With each message, the visual graph updates incrementally, accumulating context to support agent performance and adaptability.

9.1. How It Works

Our Memory builds its understanding layer by layer as the chat grows. Here's a straightforward look at how it all comes together, powered by Relational Intelligence to create real, lasting connections:

9.1.1. Short-Term Memory

Every time the chat with the agent starts, Our Memory scans it for key details using smart language analysis. The AI breaks the input into categories and only keeps reliable information, enabling immediate, context-aware responses that lay the foundation for trust-building interactions.

9.1.2. Session Memory

At the end of a chat session (e.g., when the app closes), Our Memory reviews the whole conversation. It filters out low-quality notes, removes duplicates, and creates a quick summary like:

> "This session covered travel plans and work stress; user seemed excited about Italy."

This keeps things fresh and relevant for the next session, enhancing continuity and loyalty through seamless recall.

9.1.3. User Memory

Across all sessions, Our Memory builds an overall profile, a living snapshot of interests, goals, and patterns that evolves with every chat. It combines summaries into a master view, like an evolving biography, powered by Relational Intelligence to deepen connections over time.

9.1.4. The Knowledge Graph: Connecting the Dots

Behind the magic is a visual map called the Knowledge Graph. It's like a mind map where details are colorful nodes (dots) linked by relationships (arrows). This graph, at the heart of our Relational Intelligence, helps the AI "think" holistically, uncovering insights that build authentic connections at scale.

- Yellow: Goals (e.g., "Plan a trip" or "Achieve work balance").
- Green: Topics (e.g., "Travel" or "Finance").
- Blue: Evaluations (deeper patterns, like "User engages thoughtfully" or "Prefers concise replies").
- Pink: User Facts (straight facts, like "Lives in Italy" or "Believes in slow travel").

The graph grows with connections, like "User's interest in film → Relates to brand partnerships → Connects to career goals," enabling our Relational Intelligence to engineer personalized pathways for trust and engagement.

[Knowledge Graph example] | *memory:graph.png*

9.2. How Memory Influences Conversations

Memory isn't just storage, it's what turns conversations into meaningful connections, building genuine relationships by linking the dots in the Knowledge Graph.

- Personalization: Remembers the style. If once it was mentioned liking long explanations, future responses get longer, tailored every time.
- Continuity: Picks up where it was left off. "Last time we talked Italy, have you booked that slow-travel spot yet?"
- Smarter Suggestions: Spots patterns through graph connections, offering tailored ideas that bridges goals, interests, and current topics for more relevant advice.
- Deeper Engagement: Uses linked insights to ask better questions, deriving from related beliefs and interests to spark meaningful dialogue.
- Relationship Building: The graph's connections create holistic understanding, building trust and loyalty, like evolving from casual chat to a true relational partner.
- Privacy-First: Data is collected and stored securely within the ProemthistAI system. The feature can be enabled or disabled at the agent level using the anonymous toggle, without losing the core interaction quality.

9.3. User Data Export

User data export is currently available under our enterprise license under request.

Chapter 10. Testing & Publishing

10.1. Test Your Relational Agent

Once saved, click **Open in new Tab** to try your agent in the standalone Web Client.

Click the arrow next to that button for **Copy URL**, which copies the full link, and **Copy ref**, which copies the agent reference on its own.

A QR code is shown alongside the agent. Scanning it opens the agent in the PromethistAI client.

You need the agent's reference when constructing Web Client URLs manually. Use **Copy ref**, or read it from the **Reference** column of the agent list.

10.2. Revisions

NOTE

What are Revisions for?

Every time you save your agent's configuration, a new revision is created automatically. This gives you a full history of changes — so you can test a new version without touching the live one, roll back if something goes wrong, and always know exactly what your users are experiencing.

Each revision moves through a lifecycle:

- **Draft** — work in progress. Only Drafts can be edited.
- **Preview** — ready to test. Share the preview URL with your team before going live.
- **Published** — the live version your users see. Only one can exist at a time.
- **Archived** — retired versions, kept for history.

[Revision lifecycle] | *project-management:revision.png*

10.2.1. Revision States

Draft

The initial state when you create an agent. Remains Draft until you move it to Preview.

Only Drafts can be edited. Once a revision moves to Preview, Published, or Archived, its configuration is permanently locked.

When you make changes to a locked revision and save, a new Draft is automatically created. You're not editing the locked revision, you're creating a new one based on it.

Preview

Used for testing. Anyone with the preview URL can test the agent, including non-logged users.

Multiple Preview revisions can exist at the same time.

Published

The live, user-facing version. The agent URL (without a specific revision number) always opens the Published revision.

Only one Published revision can exist at a time. When you publish a new revision, the previous Published revision is automatically archived.

Archived

A retired revision that is no longer available to the public.

Which URL opens which revision, and who can use it

A URL without a revision number always opens the agent's **Published** revision, for example `https://eu.promethist.ai/a/bianka_ma931`. Adding a revision number opens that specific revision instead, for example `https://eu.promethist.ai/a/antonin_a07fk.6`.

NOTE

A **Published** link works for anyone.

A **Preview** link also works for anyone you send it to, even if they have no access to the project. If the project requires people to sign in, they will need to sign in first. This makes Preview the right state for leaving an agent available after a demo.

A **Draft** or **Archived** link only opens for people who have a role on the project, so neither is suitable for sharing outside your team.

10.2.2. Reverting to a Previous Version

If you need to restore your agent to an earlier configuration:

1. Locate the revision you want to restore.
2. Click the **Revert** button next to that revision.

IMPORTANT Reverting archives all revisions created after the one you're reverting to.

Chapter 11. Memory

Memory powers continuity in Relational Agents — letting them recall details from past conversations and apply that context to new ones.

For the agent-level memory toggle (the **Remembers users** setting), see Remembers users (memory control) in **Creating & Configuring Agents**.

For the broader concept and how Memory builds a Knowledge Graph from interactions, see Memory & Relational Intelligence.

Chapter 12. Relational Agent Configuration

Relational Agent configuration defines how an agent behaves — its purpose, identity, voice, and the assets it can use during a conversation.

For the full configuration reference, see Creating & Configuring Agents.

Chapter 13. Distribution Channels

Distribution Channels covers how end users access and interact with Relational Agents built on the PromethistAI Platform. This section explains the available channels, their capabilities, and how to configure them for different deployment scenarios.

13.1. What the Client Is

The PromethistAI client is the end-user application for experiencing Relational Agents created in the PromethistAI Platform. The same client runs on phones and in a browser, so an agent behaves the same way wherever people meet it.

Each agent is delivered as a multimodal experience blending:

- A locally rendered 3D visual scene
- Real-time voice interaction
- A chat interface
- Contextual interactive elements, such as forms, buttons, image pickers, and embedded web pages

A reliable internet connection is required because voice and reasoning are streamed; the 3D scene renders on the device for responsiveness and smooth animation.

13.2. Available Channels

Mobile apps	The iOS and Android apps, published in the Apple App Store and Google Play Store under the name PromethistAI.
Web Client	The same client in a browser, opened from a link with nothing to install.

Out of the box, the clients connect directly to your SaaS platform configuration. This provides end-to-end functionality without requiring any technical expertise:

- Administrators configure agents in the Platform
- End users reach them through whichever channel suits them
- Branding, processes, and updates flow automatically

For enterprise customers, there is an additional option:

- Relational Agents can be integrated into existing enterprise applications via the PromethistAI SDKs
- SDKs are available exclusively to enterprise clients and allow deeper embedding of relational intelligence into custom digital touchpoints

This dual approach supports quick out-of-the-box adoption alongside highly customized enterprise distribution.

13.3. Session Results in Administration Console

Conversation outcomes propagate back to the administration console.

This allows teams to:

- Review sessions
- Analyze engagement
- Track KPIs and business objectives
- Review multimodal inputs (e.g., forms, survey responses) under the project’s data-governance settings

Chapter 14. Authentication & Projects

The PromethistAI client organizes agents into **projects**. A project is the workspace an agent was published from. When you add an agent by scanning its QR code or opening its link, that agent's project joins your Projects list.

You can talk to a project's public agents without signing in at all. Signing in unlocks agents your organization has restricted to authenticated users, links conversations to you, and enables Insights.

In the mobile app you sign in to each project separately. In the Web Client one account covers all of them.

14.1. Projects

14.1.1. The Projects Screen

Open the Projects screen using the project logo on the home screen.

It lists every project you have added. Each card shows the project logo, the project name, and the organization that owns it.

To remove a project, use the on its card. You can add it back later.

14.1.2. Adding a Project

There are three ways to add a project:

Open Project from the Platform

In the PromethistAI Platform, open the project and click **Open Project**. It opens the project in the Web Client and adds it to your Projects list. This is the quickest route if you already have the Platform open.

Scan a QR code or open a shared link

Adding an agent this way also adds the project that agent belongs to.

Add it by reference

Tap the + button on the Projects screen and paste the project's URL, or its reference on its own.

14.1.3. Signing In

Sign in from the Projects screen using your credentials. Sign-in opens in a browser.

Projects differ in what they expect. Some let you use their agents without signing in at all, some offer sign-in when you open them and let you skip it, and some require you to sign in before you can go any further.

NOTE	In the Web Client you are signed in with the same account you use for the PromethistAI Platform, so opening an agent or project from the Platform carries

that account across.

14.1.4. Managing an Account

In the mobile app, tap the account on a signed-in project card to open the account panel. It shows your name and email address, along with:

Switch account
: Sign in to that project with a different account.

Log out
: Sign out of that project. Other projects keep their own sign-in state.

Request data deletion
: Initiates a user data deletion request, covering your profile and identifiable interaction data.

In the Web Client, **Log out** is on the Projects panel and signs you out of everything at once.

14.2. Home Screen

The home screen is built around the project you have open. Depending on the project, it shows either:

- The project's **Featured agent**, with its description, a section describing what you can do with it, and a button to start talking
- The project name and description, when no featured agent is set

An **Agents** section lists the project's agents. Tap the star on a card to keep one handy. The QR button opens a scanner for adding an agent by code.

Which agents you see depends on your access to the project. Signed out, you see only **Published** agents. If you have any access to the project you also see **Preview** agents, and only as an Owner or Editor do you also see **Draft** agents.

The project logo reflects the project's branding, configured in the PromethistAI Platform. See Project Administration.

14.3. Anonymous Mode

Anonymous Mode lets you talk to an agent without linking the conversation to your identity.

Turn it on under **Settings** → **Conversation** → **Anonymous mode**. It applies to every agent and stays on until you turn it off. You can use features normally; conversations simply aren't associated with your account.

An anonymous identity is tied to the device you are using, so anonymous conversations do not carry over to another device.

If an agent requires sign-in, you must authenticate to use it.

Chapter 15. Mobile Applications

The PromethistAI Mobile Clients are available on the Apple App Store and Google Play Store under the name PromethistAI. These applications are the main way end-users access and interact with Relational Agents built on the PromethistAI Platform.

They are designed to be intuitive, support modern devices, and provide feature-rich experiences such as QR enrollment, agent sharing, and interactive content.

The apps share their interface and settings with the Web Client, so an agent behaves the same way in either.

15.1. Supported Devices Matrix

PromethistAI classifies supported devices into three categories to clarify performance expectations and feature availability.

Category	iOS	Android	Implications
Fully Supported	• iPhone (12+) • iPad Pro (2021+) • iPad Air (2022+) • iPad Mini (2021+) • iPad (9th gen+) • Macs with Apple Silicon (M1+)	• Google Pixel 7 or newer • Samsung Galaxy S22 or newer • Samsung Z Fold5 or newer • Samsung Z Flip4 or newer • OnePlus 12 or newer • OnePlus Nord 4 or newer • Xiaomi 12 or newer • Motorola Edge 60 Pro or newer • Motorola Razr 50 Ultra or newer	Full experience enabled: • Smooth rendering, HDR graphics • Low-latency audio

Category	iOS	Android	Implications
Best Effort (Performance Tier)	• iPhone 11–12 • iPad Pro (2018–2020) • iPad Air (2020) • iPad Mini (2019) • iPad (7th–8th gen)	Other Android devices meeting mid-range criteria: • 4–7 GB RAM • Android 10–12 • Released 2020–2021 (e.g., Pixel 4a/5/5a; Galaxy A52–A54; Oppo Reno 6–8; Xiaomi Mi 10/11; older mid-range and flagships)	Core features work with reduced performance: • Graphics downscaled • Slower transitions
Unsupported / Disqualified	• iPhone X and older • iPhone SE (all gens) • iPhone 8 and earlier • iPad Pro (pre-2018) • iPad Air (pre-2020) • iPad Mini (pre-2019) • iPad (6th gen and older) • Intel Macs (all models, pre-2020)	Devices failing baseline: • < 4 GB RAM • Android 9 or lower • Released 2019 or earlier • Android Go devices • Huawei (no Play Store) • Low-end OEMs	Installation may be possible, but not supported: • Instability likely • Features may fail

15.2. Home Screen

The home screen is built around the project you have open. It shows either the project's **Featured agent**, with a section describing what you can do with it and a button to start talking, or the project name and description when no featured agent is set. An **Agents** section lists the project's agents. See Home Screen for which agents appear.

The project logo opens your Projects, the gear opens Settings, and the QR button opens a scanner for adding an agent by code.

15.3. Agent Detail

Opening an agent shows its visual, name and description, a button to start talking, and three tabs:

Description

What the agent is for.

Insights

Scores from any user-facing evaluations the agent has enabled. See Insights.

History

Your conversation history with this agent.

All of this content is configured in the PromethistAI Platform by the agent creator.

15.4. During a Conversation

The conversation controls switch between the voice and chat views, and open Settings. The microphone button is how you speak when **Conversation mode** is set to Push to talk.

- The 3D presence runs locally
- Voice interaction streams in real time
- The chat view shows a timestamped transcript, and lets you type instead of speaking
- Multimodal steps such as forms, confirmations, image tiles and embedded web pages display inline

15.5. Insights

When an agent has at least one user-facing evaluation enabled, the **Insights** tab on the agent screen shows a **Summary** of the signed-in user's own conversations with that agent. Each enabled insight type appears as a card containing:

- The insight name and a short description of what it measures
- **AVG**, **MIN**, and **MAX** scores across the user's sessions with the agent
- A **Distribution** chart that visualizes how the user's scores are spread

Insights are private to each user — they reflect only the conversations the signed-in user has had with that agent, never another user's data.

NOTE	Insights require the user to be signed in and are configured per-agent through the **User Facing** destination on an evaluation. Learn more about Evaluations.

15.6. Sharing

Agents can be shared via QR code or direct link.

Recipients without the app are redirected to the store page and, after installation, taken directly to the shared agent.

15.7. Connectivity & Performance

- The voice pipeline requires a stable internet connection

- 3D rendering happens on the device, so visuals stay fluid on congested networks

How the agent listens is your choice rather than the device's. Set it under **Settings** → **Conversation** → **Conversation mode**. See Conversation mode.

Chapter 16. Settings & Troubleshooting

16.1. In-App Settings

Open Settings from the home screen, or from the conversation controls while talking to an agent.

These options let you tune the interface, the way the agent listens, and your privacy.

16.1.1. Application

Language

Sets the interaction language used when talking to agents. It does not change the interface language, which follows your device's language setting. This setting is matched against the agent's **Allowed Languages** to determine which language the agent speaks. See Allowed Languages.

UI scale

Zooms the interface in or out. Choose **80%**, **90%**, **100%**, **110%**, or **120%**. The default is 100%. Larger values make text and controls easier to read; smaller values fit more on screen.

16.1.2. Visual

Visual quality

Controls 3D fidelity: materials, lighting, and rendering resolution. Choose **High**, **Medium**, or **Low**. Lower values improve stability and battery life on less powerful devices.

16.1.3. Conversation

Conversation mode

Controls how the agent listens during a conversation. Choose one of three modes:

- **Natural conversation**: hands-free. You can speak at any time, including interrupting the agent while it is still talking
- **Take turns**: the agent finishes what it is saying before it listens to you. If you speak while it is talking, it carries on and does not pick up what you said
- **Push to talk**: the agent hears you only while you press and hold the microphone button. Useful in noisy environments, or when you want precise control over when the agent is listening

Anonymous mode

Starts conversations without linking them to your identity. You can use features normally; the conversation simply isn't associated with your account. The setting applies to every agent and stays on until you turn it off.

16.1.4. About

- **Application version**: provide this when contacting Support
- **Terms & conditions**: opens the current terms
- **Privacy policy**: opens the current privacy policy
- **Data processing notice**: opens the data processing notice

16.1.5. Support

Report an issue

Sends a problem report to PromethistAI.

16.2. Troubleshooting

16.2.1. Quick Checks

High latency?

Use Wi-Fi or ensure a strong cellular signal.

Rendering performance issues?

Lower **Visual quality**; close background apps.

Failed to interrupt the agent?

Set **Conversation mode** to **Push to talk** so the agent only listens while you hold the button.

Agent talks over you?

Set **Conversation mode** to **Take turns**.

Wrong interaction language?

Update **Language** under Application.

A project opens with no agents?

If the project tells you there are no agents available and to contact the project administrator, your account isn't permitted to use that project's agents. Ask whoever manages the project for access.

Can't see the agents you expect?

Signed out you see only Published agents, so sign in to the project. In the mobile app each project is signed in to separately. See Projects.

Chapter 17. Project Management

Projects are designed to mirror real business goals. A customer support project may include multiple agents (one for FAQs, another for warranty claims, a third for escalation), all sharing the same knowledge bases and integrations. A sales enablement project might bring together demo agents, a retail catalog, and integrations into CRM systems. By structuring work into projects, organizations ensure that everything connected to one initiative — agents, assets, branding, and analytics — is aligned and working toward the same purpose.

Each project defines four things:

- The Relational Agents it contains.
- The shared assets those agents can access, including Knowledge Bases, Integrations, and Interactive Content.
- The branding and visual style of the experience as it appears to end users in PromethistAI clients.
- The analytical scope measuring engagement, usage, and business outcomes at the project level.

This structure makes projects both cohesive (all elements tied to a single business goal) and scalable (assets can be reused across agents without duplication).

Chapter 18. Agents & Assets

18.1. Overview

Agents in a project share a common set of assets: Knowledge Bases, Integrations, and Interactive Content. Whether an agent can access an asset depends on how it is configured: Global assets are available to all agents automatically; non-Global assets must be explicitly referenced. See Assets Visibility for details.

18.2. Knowledge Bases

NOTE

What are Knowledge Bases for?

Knowledge bases are what your agent knows. Add your content and the agent can reference it in conversation, accurately and consistently.

- **Customer Support**: FAQ repositories, policy handbooks, escalation flows.
- **Sales and Retail**: Product catalogs, promotional guides, price lists.
- **Technical Support**: Troubleshooting manuals, diagnostic checklists, configuration notes.
- **Onboarding and Training**: HR guidelines, compliance materials, training decks.

Knowledge bases are shared repositories of information that Relational Agents use to ground their responses in reliable facts. They function as structured content libraries within each project, enabling agents to answer questions accurately, perform complex reasoning, and remain consistent across conversations.

18.2.1. How Knowledge Bases Work in PromethistAI

Knowledge bases are project-level assets:

- Every agent in a project inherits access to the same set of knowledge bases.
- Updates or new uploads become immediately available to all agents in that project.
- Knowledge bases are stored and parsed centrally, so administrators maintain one canonical version.
- **Renaming:** You can change a knowledge base's display name at any time. This does not affect its content or project-level availability. If an agent references a knowledge base (e.g., `@old_name`), it will be updated with the new reference `@new_name`.

When an agent receives a query, it combines relational intelligence with structured retrieval from the knowledge base. This allows it to adaptively pull facts, mirror user context, and deliver an answer that is both empathetic and precise.

If multiple knowledge bases are available, the agent will attempt to select the most relevant source. For best results, you can instruct the agent explicitly to use a particular knowledge base for specific

topics.

18.2.2. Supported Formats

PromethistAI supports a wide range of enterprise file formats:

Text and Documents	PDF, TXT, DOC, DOCX, HTML, HTM
Data Files	CSV, XLSX
Presentations	PPTX
Images	JPG, JPEG, PNG, TIFF, TIF

Maximum file size is 10 MB. When uploaded, each file is parsed so its contents are searchable in conversation.

18.2.3. Best Practices

To get the most from knowledge bases:

- Be explicit: Instruct an agent which knowledge base to consult for a given scenario. Example: "When answering questions about Company X, always use the Company_data knowledge base."
- Keep content structured: Use clear headings, consistent formatting, and concise entries. Parsed documents with logical sections produce better retrieval quality.
- Segment by purpose: Create separate knowledge bases for distinct domains (e.g. **Support_FAQ**, **Retail_Catalog**, **Compliance_Guides**).
- Update regularly: Outdated documents stay live until removed. Review knowledge bases periodically to ensure accuracy.

18.2.4. Managing Knowledge Bases

Adding a Knowledge Base

To add a new knowledge base:

1. Click **Knowledge Management** in the left sidebar.
2. Select + **Create** in the top right corner.
3. In the **Create Knowledge** modal, choose a **Knowledge Type**: **File** or **Link**.
4. Fill in the fields for the type you chose (see below).
5. (Optional) Toggle **Global** on to make this knowledge base available to every agent in the project automatically. See Assets Visibility.
6. Click **Add Knowledge** to save.

File fields:

- **Knowledge Name** (required): display name for the knowledge base.
- **Description** (optional): short summary of what the file contains.

- **Upload File** (required): drag and drop a file into the upload area or click it to browse. Maximum size 10 MB. See Supported Formats for accepted extensions.

Link fields:

- **Knowledge Name** (required): display name for the knowledge base.
- **Webpage URL** (required): the address of the page to ingest.
- **Type** (required): controls how the URL is crawled:
  - **Static URL**: a fixed URL that is crawled once to collect information. The content does not change and no additional links are followed.
  - **Dynamic URL**: a URL that is crawled every time information is needed. The content may change, and the crawler can follow links multiple levels deep to discover new data.
- **Description** (required): describe what the page contains so the agent knows when to consult it.

Once added, the content is parsed automatically and made available to agents in the project. If **Global** is off, agents must explicitly reference the knowledge base to use it.

Viewing

Knowledge bases are accessible from the **Knowledge Management** section in the left sidebar of the administration console.

- To preview content, select **View**.
- To retire a knowledge base, use the **Archive** option (bin icon). Archived knowledge bases are removed from agent access but remain in history for audit.

18.3. Integrations

NOTE

What are Integrations for?

Integrations connect your agent to external systems, so it can pull in data before a conversation starts, or push data out when one ends.

- **Before Action**: the agent fetches relevant data from an external source (a CRM, a database, another AI model) and uses it as context. For example, loading a customer's profile so the agent already knows who it's talking to.
- **During Action**: the agent calls external tools or services in real time via MCP servers. For example, checking live inventory, looking up an order status, or running a calculation mid-conversation.
- **After Action**: the agent sends a summary or extracted data to an external service. For example, logging the outcome to a CRM, posting a summary to Slack, or triggering a follow-up email.

Integrations can be connected to platforms like Zapier, Make, N8N, or any custom service your team maintains.

18.3.1. Integration Types

When creating a new integration, the app first asks you about the type. Currently, only the **custom type** is available.

18.3.2. Custom Integrations

There are three sub-types of custom integrations, each tied to a different point in the conversation:

- **Before Action**: a webhook called before the conversation starts. Useful for session initialization (e.g., loading customer profile or context).
- **During Action (MCP)**: an MCP server the agent can call during the conversation. The agent decides when to invoke it (e.g., live tool calls or data lookups).
- **After Action**: a webhook called after the conversation ends. Useful for post-conversation processing (e.g., logging summaries, updating a CRM).

All these integrations are technology agnostic: both webhooks and MCP servers can be implemented in many ways, including various low-code/no-code platforms like N8N, Zapier, Make, and Workato. In the platform, you simply link your external services and tell the agent how it can use them.

Before Action

A **Before Action** is a webhook triggered before the first speech of the conversation. It is typically used to fetch data into the context: customer profiles, relevant information from other systems, user preferences, predictions from other AI models, and more.

Configuration Fields

Description	This text is added into the context of the language model before the actual output of the webhook. It should explain what kind of information the agent is getting from this source and/or how it should be used.
URL	URL of the webhook to be called before the conversation starts. The endpoint should return text or structured output the agent can use as context at the start (e.g., user profile, last session summary).
Method	HTTP method to be used with the webhook. Currently only POST is supported.
Headers	Key-value pairs that should be sent as HTTP headers. Typically used for authentication headers.
Timeout	How long (in milliseconds) the engine should wait for the webhook to finish.
Retries	How many times the engine will try in case of failure or timeout before giving up. If your action is not idempotent (it has side-effects that should not be repeated), you may want to disable retries (value = 0).
Default Value	If the webhook does not work (after all retries the service still fails or times out), this value is added into the context instead of the real response. This can inform the model that contextual data could not be retrieved, or it can be empty if the model does not need to know.

Request Format

The webhook receives information in this format:

```
{
  "userId": "<user_id>",
  "deviceId": "<device_id>",
  "agentRef": "<agent_ref>",
  "sessionId": "<session_id>",
  "attributes": {
    "name": "<name_from_identity_provider>",
    "email": "<email_from_identity_provider>"
  }
}
```

NOTE	The `name` and `email` attributes are only present if the client has authenticated against an identity provider. This is optional for public projects but can be enforced.

After Action

An **After Action** is a webhook triggered after the conversation ends. It is used to inform other services about the outcomes of the conversation and trigger various reactions. You can create a summary and send it via email, extract information and post it into a Slack channel, or update your CRM.

Configuration Fields

URL	URL of the webhook to be called after the conversation ends. The endpoint should accept the session data and can return status or confirmation; it's mainly for logging or updates, not for enriching the live conversation.
Method	HTTP method to be used with the webhook. Currently only POST is supported.
Headers	Key-value pairs that should be sent as HTTP headers. Typically used for authentication headers.
Timeout	How long (in milliseconds) the engine should wait for the webhook to finish.
Retries	How many times the engine will try in case of failure or timeout before giving up. If your action is not idempotent (it has side-effects that should not be repeated), you may want to disable retries (value = 0).

Request Format

Your webhook will receive information in this format:

```
{
  "sessionId": "<session_id>",
  "contentRef": "<agent_ref>",
  "userAgent": "<client_version>",
```

```
  "userId": "<user_id>",
  "username": "<username>",
  "transcript": "<transcript>",
  "attributes": {
    "name": "<name_from_identity_provider>",
    "email": "<email_from_identity_provider>"
  }
}
```

NOTE	The `name` and `email` attributes are only present if the client has authenticated against an identity provider. This is optional for public projects but can be enforced.

During Action (MCP)

A **During Action** uses **MCP (Model Context Protocol)**, a standard mechanism that lets agents interact with other services and tools during the conversation. While knowledge bases provide static content and Before/After Actions trigger at fixed conversation points, an MCP-based During Action provides dynamic, self-described tools that the agent (or more precisely, the Large Language Model) can choose to call mid-conversation.

MCP integrations are managed at the project level. Once connected, they are available to every Relational Agent inside that project. This makes them shared assets: one integration can power many agents without duplicate configuration.

When invoked in conversation, the MCP integration acts as an extension of the agent. The agent decides when to call it, passes the required context, and integrates the response back into the dialogue. This ensures that the technical exchange with the external system remains invisible to the end user. They experience a continuous, natural interaction, while the agent silently handles the back-end integration. MCP integrations are typically used to:

- Connect agents to internal ticketing or CRM systems.
- Allow secure calls to enterprise APIs.
- Integrate specialised data or toolchains directly into conversational flows.

Configuration Fields

Name	The name of your integration.
URL	The endpoint URL of your MCP server.
Auth Header Name	The name of the authentication header if required by your MCP server.
Auth Header Value	The value of the authentication header if required by your MCP server.

Connections use the **Streamable HTTP** transport, the current MCP standard.

Selecting Tools

Once you've filled in the configuration, click **Try Connection** and select which tools from the linked MCP server should be added into the context.

TIP Typically, you only want to add the tools that are actually needed into the context. Too many unnecessary tools could confuse your agent.

18.4. Interactive Content

NOTE

What is Interactive Content for?

Sometimes conversation alone isn't enough. Interactive content lets your agent show or collect structured content mid-conversation, without the user ever leaving the session.

- A support agent shows a troubleshooting form instead of asking the user to describe their issue step by step.
- A retail agent displays product images so the user can pick what they're looking for.
- A training agent drops a knowledge check inline to test what was just covered.

You configure these interactions once at the project level, and any agent in the project can use them. When the agent decides the moment is right, it triggers the interaction. The app renders it, captures the user's response, and hands it back to the agent to continue the conversation.

18.4.1. Supported Interaction Types

PromethistAI supports a growing set of multimodal assets, including:

Choices	Selectable tiles that let users choose between multiple options.
Images	Enable relational agents to display images during the conversation.
Input Fields	Structured forms for capturing details such as account numbers, contact information, or survey responses.
Webpages	Embedded views that let users interact with existing enterprise web tools without leaving the conversation.
Handover	Transfer the conversation to another relational agent, ending the current session and starting a new one with the selected agent.
Video	Enable relational agents to play video content during the conversation.

18.4.2. Formatting the Description

When you create **Choice** or **Input Field** interactive content, the **Description** field is a rich-text (Markdown) editor. Use the toolbar above the field to format the text with:

- Headings (H1, H2, H3)

- **Bold**, *italic*, and strikethrough
- Inline code
- Bulleted and numbered lists
- Block quotes

18.4.3. Handover Interaction

The Handover interaction allows agents to transfer conversations to other relational agents. When triggered, the current session ends and a new conversation begins with the target agent.

This is useful for:

- Escalating conversations to specialist agents
- Routing users based on their needs or intent
- Transferring to agents with specific expertise or capabilities

Configuration

When creating a Handover interaction, configure the following fields:

Title	The display name for this handover interaction (e.g., "Talk to relational agent")
Tool Description	Instructions for the agent about when and how to use this handover. This description helps the agent understand the purpose of the transfer (e.g., "End the session, switch to relational agent and start the conversation")
Agent	Choose the target agent for the handover: • **Automatic** - Let the system select the most appropriate agent based on context • **Specific Agent** - Select a particular published agent. Only published agents appear in the picker; drafts and unpublished revisions are not selectable. The handover always points to the agent's latest published revision, so you don't need to reconfigure it when the target agent is republished.
Show as detail	Toggle on to show the target agent's details to the user before the handover starts, rather than switching immediately. The user sees a preview of the agent they're about to be handed over to.
Global	Toggle on to make this handover available to all agents in the project. Toggle off to limit access to agents that explicitly reference it.

18.5. Evaluations

NOTE	**What are Evaluations for?** After each conversation, an evaluation can automatically score what happened, so you don't have to listen to every recording or read every transcript to know how

things are going.

Example: Sales training. Your agent runs practice calls with sales reps. After each call, an evaluation scores how confident the rep sounded, whether they handled objections well, and what they should work on next. You see the results in Analytics without doing anything manually.

Example: Customer support. Your agent handles support requests. An evaluation checks whether the customer's issue was actually resolved and how much effort it took them. If a particular topic keeps coming up unresolved, you'll see it.

Evaluations run in two modes. Pick one or both:

- **Real-time Steering**: runs during the conversation and adjusts the agent's behaviour on the spot. For example, if the conversation is going off-track, the agent can course-correct without waiting for the session to end.
- **Post-Conversation Analysis**: runs after the session ends. Good for reviewing performance over time.

When using **Post-Conversation Analysis**, results can go to three places. Choose any combination:

- **Agent Memory & Context**: fed back into the next session so the agent remembers what happened and picks up where it left off.
- **Analytics & Visualization**: visible to admins in the Analytics Suite, or surfaced to end users as insights in the app.
- **3rd Party Systems**: sent to an external service via webhook for CRM updates, notifications, or custom workflows.

18.5.1. Creating an Evaluation

Navigate to **Performance** → **Evaluations** in the side menu and click **+ Create**. Fill in the required fields:

TIP	As you configure the evaluation, a live preview panel appears on the right side of the page. It shows the output structure your evaluation will produce, including the JSON schema and contextual notes for each component you've enabled. Use it to verify your setup before saving.

Evaluation Name	A unique name to identify the evaluation.
Description	A short description of what the evaluation measures.

18.5.2. Evaluation Purpose

The **Evaluation Purpose** determines when and how the evaluation runs. Both modes can be selected at the same time.

Real-time Steering	Injects context into the active agent during the conversation.
Post-Conversation Analysis	Runs after the session ends for analytics, feedback, and coaching.

Evaluation Prompt

This is where you define what the evaluation should look for and how it should interpret the results.

Real-time Steering: Look Back Window

When **Real-time Steering** is selected, a **Look back window** field appears. This controls how much conversation history the evaluation considers. Two options are available:

From start to now	Includes the entire conversation up to the current point. This is the default.
4 messages back to now	Includes only the last 4 messages. Use this to focus the evaluation on recent exchanges rather than the full history.

Featured Evaluation

Turn on **Enable featured output** to add a single, top-level result for the whole evaluation, on top of the individual insights. Configure it with these fields:

Result type	Choose **Pass/Fail** for a verdict or **Number** for a numeric score.
Featured prompt	Describe how to compute the featured result from the individual insights. This is a dedicated field, separate from the main **Evaluation Prompt**.

The remaining fields depend on the **Result type**:

- **Pass/Fail**: set a **Pass label** and a **Fail label** for the two outcomes.
- **Number**: set a **Max value** for the top of the scale.

Output Data Elements

Output data elements, also called **insights**, are the individual data points the evaluation extracts and returns. Use the buttons at the bottom of the section to build them:

- **Add group**: adds a group that several insights can share (see Grouping Insights).
- **Add insight**: adds a single insight.

Each insight has a **Group** dropdown. Leave it on **Standalone** to score the insight on its own, or select a group to score it together with the others in that group.

Each insight also has a **Type** that determines which fields are available:

Text

Variable name	The name used to reference this output.
Extraction prompt	Describe what the evaluation should extract.
Description	Describe what this variable represents.

Number

Variable name	The name used to reference this output.
Extraction prompt	Describe what the evaluation should extract.
Min value	The minimum value on the scale.
Max value	The maximum value on the scale.
Description	Describe what this variable represents.

List

Variable name	The name used to reference this output.
Extraction prompt	Describe what the evaluation should extract.
Limit to specific items	Optionally restrict the output to a defined set of values.
Description	Describe what this variable represents.

Pass / Fail

Variable name	The name used to reference this output.
Extraction prompt	Describe what the evaluation should extract.
Pass label	The label shown when the result is a pass.
Fail label	The label shown when the result is a fail.
Description	Describe what this variable represents.

Grouping Insights

By default, each insight is scored on its own. Grouping lets you bundle related insights so they're scored together as one coherent judgment, keeping linked outputs consistent with each other.

- **Standalone**: the insight is scored on its own.
- **Group**: the insights in the group are scored together, sharing the same reading of the conversation.

Group insights that belong together, for example a score and the rationale behind it. **Keep standalone** anything independent that doesn't need shared context, such as a detected scenario type or the language used.

To create a group, click **Add group**, give it a **Group name** and an optional **Group description**, then click **Add insight to group**. To keep an insight on its own, leave its **Group** dropdown set to **Standalone**.

NOTE	Evaluations created before grouping was introduced show a **This evaluation uses the legacy format** message. They still run and score as before; set up groups to move an evaluation to the new format.

Formatting Output as Markdown

You can instruct an evaluation to return its output as **Markdown**. The result is then rendered with formatting (headings, bold text, lists, and emoji status marks) in **Session Insights**, instead of as plain text.

To do this, describe the structure you want in the output's **Extraction prompt**. For example:

```
OUTPUT FORMAT — respond in clean Markdown only, in exactly this structure and order:

**Score:** `X / 6`

**Result:** ✅ **PASS**  (use ❌ **FAIL** instead if the score is below 3)

### Criteria
- ✅ **Register match** — one or two sentences of reasoning with a short quoted snippet
from the transcript.
- ✅ **Visible willingness to act** — reasoning with a short quoted snippet.
- (…one bullet per criterion, in the order listed above, each beginning with ✅ for
pass or ❌ for fail…)

Replace the example marks with the real verdicts. Do not add any preamble, closing
remarks, tables, or text outside this structure.
```

Adapt the structure to whatever your evaluation needs.

Agent Memory & Context

Available when Post-Conversation Analysis is selected.

Session agent	Updates context for the specific agent used in this session.
Featured agent	Updates the main Project Lead agent with new insights.

Analytics & Visualization

Available when Post-Conversation Analysis is selected.

User Facing	Visualises results as insights to the end-user in the app.
Admin Facing	Logs data to the Analytics Suite for comprehensive reporting.

3rd Party Systems (Webhook Integration)

Available when Post-Conversation Analysis is selected.

Enter a webhook URL to forward evaluation results to an external system automatically.

18.6. Assets Visibility (Project-Wide)

Project assets can be configured as **Global** or **Non-Global**, giving you control over which agents can access them.

18.6.1. Global Assets

Assets marked as **Global** are available to all agents in the project automatically, without requiring explicit references.

- All agents can choose to use Global assets, even if not mentioned in their purpose, business process, or guardrails
- Includes Knowledge Bases, MCP Integrations and Interactive Content
- Best for shared resources that all agents should have access to

18.6.2. Non-Global Assets

Assets that are **not** marked as Global are available only to agents that explicitly reference them using @mentions (e.g., `@Support_FAQ`).

- Agent must include the asset reference in their Custom Configuration (purpose, business process, or guardrails)
- Best for specialized resources that should be isolated to specific agents
- Provides fine-grained control over asset access per agent within a project

18.6.3. Configuring Asset Visibility

To mark an asset as Global:

1. Navigate to the asset (Knowledge Bases, MCP Integrations or Interactive Content)
2. Open an existing asset or create a new one
3. Toggle the **Global** setting
4. Save changes

NOTE

You can optionally **@mention** a specific asset (e.g., `@Support_FAQ`) in an agent's Custom Configuration to:

- Guide the agent to prefer it
- Avoid confusion with similar names

@mentions are **never required** they're just for precision.

Chapter 19. Project Administration

Collaboration in PromethistAI happens inside projects. Project membership is governed by roles inherited from the account, ensuring consistent permissions across the platform.

19.1. Viewing Project Members

Each project has a **Members** tab that shows the full list of people who can access it, including:

- Name and email
- Their account role
- When their access was last updated

It's simply a quick way to see "who can see and touch this project."

NOTE	To add someone new to the account or change their role, go to **Account → Members**. See Managing Members for instructions.

19.2. Editing Roles

Roles control what members can see and do inside a project:

Owner	Full control, including the ability to manage billing, invite members, and assign roles.
Editor	Can view and edit all core resources (agents, knowledge bases, analytics, interactive content), and create or delete projects.
Viewer	Read-only access to project content and analytics.

Project roles also affect which agents members see in the client. Everyone sees the project's **Published** agents. Members of the project also see **Preview** agents, and Owners and Editors additionally see **Draft** agents. See Home Screen.

To change a role:

1. Select **Edit Role** and choose a new role.
2. Click **Edit** to confirm.

19.3. Revoking Access

Users can be removed at any time:

1. Select **Revoke access**
2. Confirm the action

Pending invitations can also be revoked in the same way.

19.4. Project Settings

Projects also define how agents appear to end users in the PromethistAI clients. Administrators can customize branding so that the visual presentation of agents reflects the organization's identity, keeping the experience consistent with the enterprise brand even though it is delivered through the PromethistAI application.

Under **Project** → **Settings** you can set:

Project Name and Description
: How the project is identified.

Featured Agent
: The project's default main agent. It displays prominently at the top of the home screen.

Logo
: Shown in the top left of the client when the project is active.

Primary and Secondary Color
: Applied in the client near the call to action, with a live preview beside the color fields.

LiveAvatar API Key
: See below.

Click **Save Changes** to apply.

19.4.1. Live Avatar API Key

The **Live Avatar API Key** is the key used for the **LiveAvatar stream** Avatar Source. Paste the key from your Live Avatar (HeyGen) account in the project's Main Information settings.

You can set a Live Avatar API Key at both the **project** level and the account level.

NOTE	If only one key is set, that key is used. If both are set, the **project-level key takes precedence** over the account-level key.

19.5. SSO and Identities

For enterprises requiring secure distribution, projects can be tied to **Single Sign-On (SSO)**. When SSO is enabled, user identities and entitlements flow directly into the PromethistAI clients. This means that only authorized employees, partners, or customers see the agents they are entitled to access, and access can be managed centrally through the organization's identity provider.

NOTE	SSO configurations are managed at account level, under **Account** → **Identity Providers**. Use **Custom Integration** to connect your organization's own SSO. SSO and identity features are available as part of the enterprise offering and can be enabled by contacting the PromethistAI Sales team.

19.6. Relational Agents Listing

The **Agents Listing** shows all Relational Agents within the Project. From here you can:

- Edit existing Relational Agents.
- Clone an existing agent to reuse its configuration (see Cloning Agents).
- Archive those no longer active (see Archiving Agents).
- Create new Relational Agents with + **Create** (see Creating a New Relational Agent for the full process).

The agent set as the project's **Featured Agent** is marked with a **Featured** tag in the listing, so you can tell at a glance which agent is the project default. Set the Featured Agent in Project Settings.

A Project often contains multiple Relational Agents working toward the same goal. For instance:

- In a **support Project**, one Relational Agent may handle FAQs, another processes warranty claims, and another escalates complex cases.
- In a **training Project**, one Relational Agent might welcome new employees, another deliver compliance lessons, and another manage leadership development.

The listing provides visibility and control so that all these Relational Agents stay aligned.

19.6.1. Cloning Agents

To reuse an existing agent's configuration as the starting point for a new one, use the **Clone** action in the **Actions** column. Cloning creates a copy you can rename and adjust, so you don't have to rebuild a similar agent from scratch.

19.6.2. Archiving Agents

To remove an unused agent from the active Agents list, use the **Archive** icon in the **Actions** column. Archived agents are hidden from the default view but remain in the project's history with an **Archived** status badge.

Toggle the **Show archived** checkbox at the top of the Agents list to reveal archived agents alongside active ones.

WARNING	Archiving is permanent. Once an agent is archived, it cannot be restored or returned to active status. Make sure you no longer need the agent before archiving it.

Chapter 20. Analytics Suite

NOTE

What is the Analytics Suite for?

The Analytics Suite is where you see how your agents are actually performing: not just whether conversations are happening, but whether they're going well. It combines usage statistics with qualitative signals, so you can evaluate both the efficiency of conversations and the outcomes they generate.

You can view data across the whole project or drill down to a specific agent or user.

20.1. Filters

The filter bar at the top of the Analytics Suite controls the data shown in every card below it.

From / To	Set a custom start and end point for the data, including the time of day. Pick the day and time in the picker, or type the value straight into the field in the `Aug 18, 2026 @ 13:58` format. Times use your local timezone.
Time Period	Choose a preset range, from **Last 15 minutes** up to **All time**, instead of setting the dates by hand. Adjusting **From** or **To** switches this to **Custom**.
Select Agent	View data for all agents, or narrow it down to specific agents and versions. See Agent and Version Selection.
Filter by User	View data across all users, or filter by a specific one. Use the search field to find a user.
Export data	Download the currently filtered data as **Excel (.xlsx)**, **CSV (.zip)**, or **JSON (.json)**.

20.1.1. Agent and Version Selection

Select Agent opens the **Agents version selection** panel: agents on the left, the versions of the selected agent on the right. The number of selected versions is shown in the top-right corner of the panel, and each agent in the list shows how many of its versions are selected.

By default **All Agents** is selected, which includes every version of every agent.

NOTE

To analyse specific versions, first click **All Agents** to clear the selection. Then choose an agent on the left and toggle the versions you want on the right. Otherwise, everything stays selected.

Use **All drafts** to select every draft version of the chosen agent at once, and **Clear drafts** to remove them again. Use the **Search agents** field to find an agent in a long list.

The button itself reflects the current selection: **All Agents**, a single agent version, or the number of selected agent versions.

20.2. Core Metrics

- **Conversations** – visualises the conversational throughput of all sessions by length (short, medium, long) with lifetime totals.
- **Unique Users** – the number of distinct users in the selected period, with a chart of unique users over time and a lifetime total.
- **Total Conversation Length** – cumulative duration of all conversations.
- **Average Conversation Length** – average duration per conversation.
- **Average Turns per Session** – average number of interaction turns within a session.
- **Relational Metrics** – qualitative measures such as interaction style matching, relationship building, user engagement, empowerment, and task resolution.
- **Business Metrics** – purpose fulfillment rate and business process adherence.
- **Evaluation Metrics** – results from evaluations run across conversations. Filter by evaluation type and view all evaluations or select a specific one.
- **Interactive Content Metrics** – counts of structured actions such as forms filled, images displayed, or web views opened.
- **Conversations** – a table listing individual conversations and a summary of the interaction.

Cards that show a green or red percentage chip compare the selected period with the period of the same length immediately before it. Hover over the chip to see the exact comparison range.

These metrics allow product teams to continuously improve agents: reducing drop-off, shortening resolution times, and optimising relational quality. They also provide evidence for business leaders on how relational intelligence contributes directly to KPIs like customer satisfaction, efficiency, and conversion.

Chapter 21. Platform Administration

21.1. Getting Started

This section explains how to get started with Promethist, from your very first login to managing accounts, members, and billing. It is designed to provide both step-by-step instructions and the broader context you need to understand how accounts fit into the overall platform structure.

21.2. Signing in with Google or Microsoft

PromethistAI supports secure, password-free authentication via **Google** and **Microsoft** using OAuth. This provides strong identity management without requiring users to create and maintain new credentials.

To sign in:

1. Go to the PromethistAI login page.
2. Select **Continue with Google** or **Continue with Microsoft**.
3. Approve the OAuth prompt to authorize access.
4. You will be redirected into the PromethistAI administration console.

NOTE	PromethistAI requires a Google or Microsoft account. If your organization uses Google Workspace or Microsoft Entra ID, your access will be tied to that enterprise identity.

21.3. Browser Support

We officially support **Google Chrome (latest stable release)** for the PromethistAI administration console.

Other modern browsers (e.g., Safari, Firefox, Microsoft Edge) are supported on a **best-effort** basis: the console may work, but some features might be limited or behave differently.

NOTE	If you encounter a browser-specific issue outside Chrome, try reproducing in **Chrome (latest)** before filing a ticket.

21.3.1. Accepting an Invitation

Most users will join PromethistAI by being invited into an existing account.

- Invitations are sent by email or link.
- After accepting the invitation and signing in with **Google or Microsoft**, you'll immediately see the projects and agents available under that account.

TIP If the invitation email doesn’t appear, check your spam or promotions folder. If issues persist, an account Owner can re-send the invite.

Chapter 22. Account Management

Accounts are the **foundation** of PromethistAI. They represent your organization within the platform and define the boundaries for billing, member access, and project ownership.

Think of the structure as:

- **Accounts** – top-level container. Owns billing, members, and projects.
- **Projects** – workspaces that hold agents, assets, and integrations.
- **Agents** – the relational experiences delivered to end-users.

This hierarchy ensures that every agent is tied to a project, every project belongs to an account, and billing is consolidated at the account level.

22.1. Creating an Account

If you are the first person from your organization to log in:

1. You'll be prompted to create a new account.
2. Provide:
    - **Account name** – usually your organization's name.
    - **Billing email** – the address where invoices and usage alerts will be sent.
    - Additional required information (contact details, region, etc.).

Once set up, this account becomes your **primary organizational container**, and you can begin adding projects and inviting team members.

22.2. Accessing Accounts

It's common for consultants, agencies, or enterprises with multiple divisions to have access to more than one account.

[Account Switcher] | *onboarding-account-management:account-selection.png*

- Use the **Account Switcher** in the administration console to move between accounts.
- The switcher shows all accounts you belong to and your role in each.
- Your role determines what you can see and do in that account (see Roles & Permissions.)

Example: A design consultant may be an **Owner** in their agency's account, but only a **Viewer** in a client's account.

22.2.1. Opening a Project from a Client Profile

From a client profile, use the **Open project** button to jump straight into the project. It redirects you to the Web Client, where you can open the agent you need.

NOTE	The **Open project** button only appears once the project has at least one published agent. If no agent has been published yet, the button is not shown.

22.3. Account Settings

Account Settings is where account-level options are managed. Settings apply across the whole account, so they affect every project under it. Use **Save Changes** to apply your edits.

22.3.1. Display full conversation data

Account **Owners** see a **Display full conversation data** toggle in Account Settings. It controls who can see **Conversations** in the Analytics Suite across every project under the account.

- The toggle is **enabled by default**.
- When **enabled**, anyone with access can see the Conversations in the Analytics Suite of any project under the account.
- When **disabled**, no one can see the Conversations.

22.4. Selecting Platform Region (US/EU)

You can choose which backend cluster the **administration console** connects to. Use the region switcher in the top navigation (globe icon) to toggle between **EU** and **US**.

22.4.1. How to switch

1. Click the globe pill in the top bar (e.g., **EU**).
2. Choose **US** or **EU** from the popover.
3. The console reloads to that region's workspace and shows the content stored there.

- The platform region defines where your console resources live (projects, agents, knowledge bases, assets, analytics). Switching regions does **not** move data; it only changes which region you're viewing.
- Your selection is remembered for your next visit (per browser).

Chapter 23. Managing Members

Collaboration in PromethistAI happens at the account level. Account membership is governed by roles that define what each user can see and do across projects and agents.

23.1. Inviting Members

Owners and Editors can expand collaboration by inviting new members:

1. Navigate to **Account** → **Members**
2. Click **Invite User**
3. Enter the person's email address
4. Assign a role (Owner, Editor, Viewer)
5. Send the invitation

The invitee will receive an email and join the account after signing in with Google or Microsoft.

TIP If the invitation email doesn't appear, ask them to check spam or promotions. If issues persist, an account Owner can re-send the invite.

23.2. Roles & Permissions

NOTE

What are Roles for?

Roles control what each member can see and do in the account. Assign them based on what someone actually needs — not more.

- **Owner** — for whoever is responsible for the account. They have full control including billing and member management. Usually one or two people.
- **Editor** — for people actively building and maintaining agents: content creators, trainers, IT. They can create and edit everything except account-level settings and billing.
- **Viewer** — for stakeholders who need visibility but shouldn't touch anything: managers, analysts, clients. Read-only access across the account.

When in doubt, start with Viewer and upgrade if they need more.

IMPORTANT **Best Practice:** Assign the lowest necessary role to preserve governance and security.

23.3. Editing Roles

To change a member's role:

1. Navigate to **Account** → **Members**
2. Select **Edit Role** and choose a new role
3. Click **Edit** to confirm

23.4. Revoking Access

Users can be removed at any time:

1. Navigate to **Account** → **Members**
2. Select **Revoke access**
3. Confirm the action

Pending invitations can also be revoked in the same way.

Chapter 24. Plans, Pricing & Billing

PromethistAI offers three subscription plans to suit different needs, from individual exploration to enterprise-scale deployments.

[Pricing overview with Free, Growth, and Enterprise cards] | *platform-administration:pricing.png*

24.1. Overview of Subscription Plans

- **Free** – Perfect for individuals or small teams trying out the most advanced relational agents.
- **Growth** – Designed for growing teams that want to quickly deploy digital humans into real-world workflows.
- **Enterprise** – Built for companies deploying adaptive, scalable, enterprise-ready relational AI in complex workflows.

24.2. Plan Comparison

Feature	Free	Growth	Enterprise
Platform Access	Unlimited	Unlimited	Unlimited
Active Conversation Minutes	60 minutes included	Flexible (per-seat or usage-based with volume discounts)	Custom
Cost for Additional Minutes	$0.5 per minute	Included in plan pricing	Custom
Support	Best-Effort	Priority Support	24/7 Support & Dedicated Onboarding
Security & Authentication	Basic	SSO & Basic Security Controls	Advanced Security & Compliance
Deployment Speed	–	Production-Ready, Deployed in Hours	Custom Workflows
Customizations	–	–	Bring-Your-Own LLMs Custom Avatars, Voices & Workflows
Pricing Model	Free + pay-per-use overage	Flexible per-seat or usage-based	Custom Enterprise Pricing

24.3. Selecting or Upgrading Your Plan

To select the Growth plan:

1. Click **Select** on the Growth card.
2. Choose your preferred billing period (Monthly or Yearly).

3. Select the license parameter (Minute-based or Seat-based).
4. Adjust the monthly volume slider to match your expected usage.
5. Review the calculated price and any savings.
6. Click **Select This Plan** and provide payment information.

To select the Enterprise plan:

1. Click **Select** on the Enterprise card.
2. You will be directed to the Enterprise customization screen.
3. Click **Contact Us** to reach the sales team for tailored licensing and pricing.

Higher-plan features become available immediately after completing the process (payment for Growth, or sales confirmation for Enterprise).

24.4. Billing & Usage

- Billing is applied at the **account level**, regardless of the number of projects or agents.
- Charges are based on **Conversational Minutes** – the total active interaction time between agents and end users.
- All payments are processed securely via **Stripe**.
- The billing dashboard in the Administration Console provides real-time usage, historical invoices, and projected costs.
- Some advanced features (SSO, premium analytics, custom integrations, Bring-Your-Own LLMs, etc.) are available only on Growth or Enterprise plans.

24.5. Best Practices

- Start with the Free plan to evaluate capabilities and estimate usage.
- Use the Growth calculator to find the most cost-effective volume and billing options.
- Monitor usage reports regularly to forecast spend and avoid unexpected overage.
- Use separate accounts when strict billing separation or compliance isolation is required across divisions.
- Keep your billing email actively monitored for invoices and important notices.

By choosing the right plan and managing usage effectively, you can scale PromethistAI confidently—from your first agent to enterprise-wide deployments.

Chapter 25. Platform Security & Compliance

PromethistAI is designed with security, compliance, and privacy at its core. This document describes the security principles, operational safeguards, and compliance measures that govern the platform, the distribution channels, and the Relational Intelligence Engine.

25.1. Platform Overview

The PromethistAI platform operates as a regionalized, multi-tenant SaaS environment, hosted on Microsoft Azure under an enterprise license. The platform provides account management, project administration, and integration into the Relational Intelligence Engine. All customer instances are logically isolated, and access to administrative functions is restricted to authenticated and authorized users through secure identity providers such as Microsoft and Google.

- All communication between platform components is encrypted in transit using TLS 1.2 or higher, and all stored information is encrypted at rest using AES-256.
- Encryption keys are managed by Azure Key Vault and rotated regularly.
- Access to sensitive operations is further secured by role-based access control and auditable activity logging.

25.1.1. Distribution Channels

Relational agents created and managed on the platform are accessed through the PromethistAI clients, which run on iOS and Android devices and in a browser. These applications serve purely as a distribution channel and do not persist sensitive data locally, aside from secure authentication information. All processing continues to occur within the selected deployment region, ensuring both compliance with data residency requirements and reduced latency for end users.

The distribution layer supports kiosk mode and additional controlled pathways while continuing to respect regional boundaries. Regardless of the mode of use, the design ensures that user data remains within the locality selected at the time of account provisioning.

25.1.2. Relational Intelligence Engine

At the core of the solution lies the **Relational Intelligence Engine** responsible for managing the interaction and data handling with the end user. This component runs within the same Azure regional environment as the platform itself though separated in architectural sense. It processes conversational data securely, using a decoupled microservices architecture with strong inter-service authentication and isolation.

Internal teams at PromethistAI do not have general access to customer conversations; operational access is strictly controlled and allowed only in exceptional circumstances, subject to least-privilege enforcement and audit logging.

25.1.3. Regional Deployment and Data Residency

When creating an account, customers can choose whether their data is deployed and hosted in the

European Union or in North America. This regional deployment ensures that data never leaves the selected geography, meeting strict residency requirements. It also provides performance benefits by minimizing latency and ensuring reliable access to relational agents.

The regional choice applies consistently across the platform, the Relational Intelligence Engine, and the distribution channels. As PromethistAI expands, additional hosting regions will be made available, enabling greater flexibility for global organizations.

25.1.4. Identity and Access Management

PromethistAI employs a layered approach to identity and access management (IAM), ensuring that both platform administrators and end users authenticate securely and operate with the least privileges necessary. Our IAM framework governs how accounts are created, how users are authenticated, and how permissions are enforced across the platform and the client applications.

Authentication

By default, PromethistAI supports federated authentication through trusted identity providers such as Google and Microsoft. This eliminates the need for password management and reduces risks associated with weak or reused credentials.

For organizations with more advanced requirements, PromethistAI offers **Single Sign-On (SSO)** through integration with enterprise identity providers. SSO is powered by **Keycloak** and supports a wide range of federation standards including **SAML 2.0, OpenID Connect**, and **OAuth 2.0.** With SSO enabled, employees, partners, or customers can access the PromethistAI platform and assigned relational agents using the same credentials they use internally, simplifying access and strengthening compliance with enterprise security policies.

SSO is available as part of the **Enterprise Tier license**, and configuration can be customized per organization in coordination with our technical team.

Authorization and Roles

Authorization within the platform is governed by **role-based access control (RBAC)**. Roles define the scope of actions available to each user and are consistently applied across accounts and projects. By default, the following roles exist at the account level:

- Owner – Full administrative rights, including billing, member management, and project governance.
- Editor – Operational control over projects, agents, and assets, without access to sensitive account-level functions such as billing.
- Viewer – Read-only access for oversight and review purposes.

Within projects, these same roles cascade, ensuring consistent governance across the entire hierarchy. This uniformity reduces complexity while maintaining clear boundaries of responsibility.

Logging and Auditability

PromethistAI maintains baseline operational logs for authentication events, role assignments, and administrative actions. These logs are used internally to support security monitoring and incident response.

For enterprise customers requiring deeper visibility, **enhanced audit logging** is available on demand. In this mode, detailed records of authentication attempts, role changes, and privileged actions can be exposed to account Owners for compliance and governance purposes.

Client Access and Identities

On the client side, the PromethistAI clients respect the same IAM model. Agents that require authentication enforce identity checks before interaction begins. When SSO is enabled, entitlements flow directly from the enterprise identity provider into the client applications, ensuring that only authorized users can access designated agents.

This design allows organizations to securely distribute agents to employees, partners, or selected customer groups, while maintaining centralized control over who can see and interact with each digital agent.

25.1.5. Operations and Infrastructure

PromethistAI's infrastructure is deployed exclusively on **Microsoft Azure**, leveraging Azure's enterprise-grade compliance with ISO 27001, SOC 2, GDPR, and HIPAA frameworks. All deployments are automated and defined through **Infrastructure-as-Code (IaC)**, ensuring consistency, repeatability, and auditability across environments.

Data Protection and Access Controls

Data protection is embedded into every layer of the platform. All communications are encrypted in transit using TLS 1.2+ and all databases and object storage are encrypted at rest using AES-256. Keys are centrally managed in **Azure Key Vault**, with strict rotation policies and access logging.

No PromethistAI personnel have routine access to customer conversational content. Limited access is technically possible only for a small, highly secured subset of our **operations and DevOps team**, and only in exceptional cases such as resolving a customer support request or fulfilling a compliance obligation. Such access is time-bound, governed by least-privilege policies, and fully logged.

Customers do not interact with raw infrastructure data directly. Instead, they may access relevant information through **dedicated APIs** (for example, APIs exposing memory or account-level resources) or request access through formal support channels. This approach protects the integrity of operational systems while giving customers the transparency they need for governance and compliance – and this feature is available upon request as a part of our Enterprise Tier.

Operational Safeguards

PromethistAI maintains a robust security operations program that includes continuous vulnerability scanning, automated patch management, and 24/7 monitoring across all critical

systems. Automated alerting is backed by documented incident response playbooks.

The platform is protected against denial-of-service attacks through **Azure Front Door**, while regional redundancy ensures that services remain resilient even in the event of localized infrastructure failures.

Business Continuity and Disaster Recovery

PromethistAI is designed for resilience and continuity. Each regional deployment includes redundancy across availability zones, as well as daily encrypted backups. Restoration procedures are regularly tested to validate data integrity and recovery reliability.

The platform's **Recovery Point Objective (RPO)** is 24 hours, while the **Recovery Time Objective (RTO)** for critical services is under four hours. These measures ensure that customer data and services remain available even under adverse conditions.

25.1.6. Compliance and Certifications

PromethistAI is certified under ISO/IEC 27001 (inherited from PromethistAI governing company certifications), demonstrating our commitment to international information security standards. The platform is fully compliant with GDPR, ensuring customer rights such as consent, data portability, and the right to erasure. SOC 2 Type II certification is underway, and HIPAA compliance support is available for customers operating in regulated industries. Regular third-party penetration tests validate the ongoing effectiveness of our security controls.

25.1.7. Privacy & Governance Policy

This section explains how PromethistAI governs data on the platform, how responsibilities are shared with customers, and what rules apply to relational agent behavior. Technical controls (encryption, hosting, IAM, operations) are described in Security & Compliance and are referenced here where relevant.

Purpose and Scope

PromethistAI enables organizations to build and operate relational agents. With that capability comes responsibility for how data is processed and how agents behave. This policy covers data handled within PromethistAI accounts, including platform administration, the Relational Intelligence Engine, and distribution via the PromethistAI clients.

Where this policy references security mechanisms (encryption, monitoring, incident response), the authoritative details live in Security & Compliance.

Roles and Responsibilities under Privacy Law

For Customer Content (e.g., knowledge bases, agent configuration, Memories), the customer is the controller and PromethistAI acts as the processor. For platform operations data (e.g., account admin details, billing, core service telemetry), PromethistAI acts as an independent controller to operate and secure the service.

A Data Processing Addendum (DPA) and the current list of sub-processors are available upon

request.

Data Categories We Handle

Customer Content (processor role)

Content you upload or configure (documents, knowledge bases, multimodal assets, integrations metadata), agent configuration (identity, guardrails, processes), and any persisted artifacts produced by the platform (e.g., Memories) or created via your calls to dedicated APIs.

Account & Operational Data (controller role)

Account details (organization, emails, roles), billing/usage data, and minimal service telemetry required to operate, secure, and improve reliability of the platform and its clients.

> PromethistAI clients act as a distribution layer and do not persist sensitive content locally beyond secure session tokens required for authentication; 3D scenes render on-device.

Data Residency

At onboarding, each account is bound to a hosting region (currently EU or North America). Customer Content is processed and stored exclusively within the selected region. This region binding applies to the administration platform, the Relational Intelligence Engine, and distribution via PromethistAI clients.

Cross-region migration occurs only upon explicit customer request and under a controlled migration plan.

Collection and Use of Data

- To provide the service. Route and render agent interactions, manage projects/assets, and operate integrations.
- To secure and maintain the service. Authenticate users, enforce entitlements, detect abuse, and diagnose incidents.
- To support customers. Investigate issues when you ask us to, or when required for compliance; see “Access by PromethistAI Personnel” below.

We do not sell Customer Content or use it to train unrelated third-party models.

Retention and Deletion

PromethistAI retains conversational interactions and associated artifacts by default. This ensures that relational agents can maintain continuity across sessions, support compliance and audit requirements, and provide customers with the ability to review historical data where needed.

All stored content is:

- **Encrypted at rest** using AES-256 and contained within the region selected at onboarding (e.g.,

EU or North America).

- **Protected by strict access controls**, with access limited to a highly secured subset of PromethistAI operations staff, and only under approved, time-bound, and logged conditions (such as a customer support request or compliance requirement).

Customers may request data exports or deletions through formal support channels. When such requests are processed:

- **Active systems** are updated immediately.
- **Backups** are purged according to PromethistAI's standard backup lifecycle, ensuring compliance while preserving system integrity.

Unless otherwise agreed in writing, retention follows PromethistAI's default platform policy. This model ensures data availability for operational continuity and compliance, while giving customers governance rights over export and deletion on demand.

Access by PromethistAI Personnel

PromethistAI personnel do not have routine access to Customer Content. Access is technically possible only for a highly secured subset of our **Operations and DevOps team**, and only under specific circumstances such as incident response, compliance requirements, or when a customer explicitly requests support. **All such access is time-bound, requires prior approval, follows least-privilege principles, and is fully logged.**

No other PromethistAI teams have access to Customer Content by default. In cases where a customer explicitly grants access — for example, by inviting PromethistAI staff into a dedicated account or project for implementation, testing, or support purposes — access is limited to the scope defined by the customer and governed by the same security controls.

Customer Access to Data and Audit Information

Customers do not access raw infrastructure storage or logs. You interact with your own data through available platform interfaces. Where audit information is exposed, it is provided in a structured, customer-safe form. Additional audit extracts can be provided upon request, subject to security review.

Shared Responsibility & Customer Controls

PromethistAI provides the secure, compliant foundation; customers decide how it is used.

You control:

- Region selection at account creation (applies platform-wide, including engine and distribution).
- Identity and access: choosing Google, Microsoft, or enterprise SSO; assigning roles and entitlements within accounts and projects.
- Agent configuration: knowledge bases, multimodal assets, integrations, guardrails, and the agent's purpose/process.
- Retention and deletion of persisted artifacts via formal requests.

- User notices/consents where your organization is the controller (e.g., informing end users they are interacting with an automated system).

PromethistAI controls:

- Hosting security, encryption, regional processing, secure operations, vetted sub-processors, monitoring and incident response, and execution of controller instructions consistent with the DPA.

Agent Content Responsibility

Relational agents are configured and governed by the customer. PromethistAI does not pre-moderate or continuously monitor agent behavior.

You (the customer) are solely responsible for:

- The content and responses your agents produce, including accuracy, legality, and appropriateness.
- Ensuring agent behavior complies with applicable laws, regulations, industry rules, and your internal policies.
- Implementing and maintaining guardrails to prevent unlawful, offensive, discriminatory, or otherwise prohibited output.
- Obtaining any required consents or notices from end users.

PromethistAI provides technical controls and may suspend or restrict access where we detect abuse or violations of platform policies, but responsibility for agent output rests with the customer.

Data Subject Rights (GDPR and analogs)

When you are the controller, PromethistAI (as processor) will assist you in fulfilling data subject requests—access, rectification, deletion, portability, objection, and restriction—through available APIs and formal support channels. Requests sent directly to PromethistAI that relate to your Customer Content are routed to you as the controller.

Chapter 26. PromethistAI Release Notes

A changelog of the latest updates for PromethistAI.

26.1. September 2, 2026

26.1.1. Agent visibility depends on your access (New)

- Signed out, you see a project's **Published** agents. With any access to the project you also see **Preview** agents, and Owners and Editors also see **Draft** agents.
- A **Preview** link works for anyone you send it to, so it is the right state for leaving an agent available after a demo. **Draft** and **Archived** links only open for people with a role on the project.

Learn more about Home Screen and Revisions.

26.1.2. Conversation mode replaces Push-to-talk and Turn Taking (Improvements)

- The two separate settings are now one **Conversation mode** setting with three options: **Natural conversation**, **Take turns**, and **Push to talk**.

Learn more about Conversation mode.

26.1.3. Projects (New)

- Agents are now grouped into projects. The home screen is built around the project you have open, and you can add or remove projects from the Projects screen.

Learn more about Projects.

26.1.4. UI scale and Report an issue (New)

- **UI scale** zooms the interface between 80% and 120%.
- **Report an issue** sends a problem report to PromethistAI.

Learn more about Settings & Troubleshooting.

26.2. August 20, 2026

26.2.1. Filter analytics by agent and version (New)

- **Select Agent** in the Analytics Suite now opens an **Agents version selection** panel, so you can narrow the data down to specific agents and specific versions of those agents.

Learn more about Agent and Version Selection.

26.2.2. More time periods and time-of-day filtering (Improvements)

- **Time Period** now offers presets ranging from **Last 15 minutes** up to **All time**.
- **From** and **To** now include a time of day, so you can narrow a range down to a specific part of a day.

Learn more about Filters.

26.2.3. Pass / Fail evaluation outputs (Improvements)

- The **Boolean** output data element type is now called **Pass / Fail**, and its two labels are now **Pass label** and **Fail label**.

Learn more about Output Data Elements.

26.3. July 30, 2026

26.3.1. Grouping Insights in Evaluations (New)

- You can now bundle related insights into a **group** so they're scored together as one coherent judgment, or keep an insight **Standalone** to score it on its own.

Learn more about Grouping Insights.

26.3.2. Featured Evaluation (New)

- **Enable featured output** adds a single high-level result (**Pass/Fail** or a **Number**) that summarizes the whole evaluation.

Learn more about Featured Evaluation.

26.3.3. Clone an agent (New)

- You can now clone an agent from the Agents list to reuse its configuration as the starting point for a new one.

Learn more about Cloning Agents.

26.3.4. Featured Agent tag in the Agents list (New)

- The project's **Featured Agent** now shows a **Featured** tag in the Agents list.

Learn more about Relational Agents Listing.

26.3.5. Manage evaluations from agent setup (New)

- When configuring an agent, buttons in the **Agent Evaluation** section let you open the **Evaluations** page or create a new evaluation without leaving agent setup.

Learn more about Agent Evaluation.

26.3.6. MCP transport standardized on Streamable HTTP (Improvements)

- **During Action (MCP)** integrations now always use the **Streamable HTTP** transport, the current MCP standard.

Learn more about During Action (MCP).

26.3.7. Open a project from a client profile (New)

- A new **Open project** button on a client profile takes you straight to the project in the Web Client.

Learn more about Opening a Project from a Client Profile.

26.4. June 18, 2026

26.4.1. Markdown-formatted evaluation output (New)

- You can now instruct an evaluation to return its output as **Markdown**. The result is rendered with formatting (headings, bold text, lists, and emoji status marks) in **Session Insights**, instead of as plain text.

Learn more about Formatting Output as Markdown.

26.4.2. @Reference in Template Variables (New)

- When you create an agent from a template, you can now use **@Reference** in any **Template Variable** to reference a project asset (such as a Knowledge Base, Interactive Content, Integration, or Skill) directly inside the field.

Learn more about Template Variables.

26.4.3. Formatted Interactive Content descriptions (New)

- When creating **Choice** or **Input Field** interactive content, the **Description** field is now a rich-text (Markdown) editor. Format the text with headings, bold, italic, strikethrough, inline code, bulleted and numbered lists, and block quotes.

Learn more about Formatting the Description.

26.4.4. Display full conversation data (New)

- Account **Owners** now have a **Display full conversation data** toggle in Account Settings. It controls who can see **Conversations** in the Analytics Suite across every project under the account. The toggle is enabled by default. When disabled, no one can see the Conversations.

Learn more about Display full conversation data.

26.5. June 4, 2026

26.5.1. Avatar Source (New)

- You can now choose how an agent's avatar is rendered with the new **Avatar Source** setting: **On-device rendered** (rendered locally by the Promethist engine) or **LiveAvatar stream** (a live video feed from a HeyGen LiveAvatar). LiveAvatar stream needs a **Live Avatar ID** from your HeyGen account.

Learn more about Avatar Source.

26.5.2. Stop AI Assistant Mid-Response (New)

- You can now stop AI Assistant while it's generating a response: a new **Stop** button cuts it off instead of making you wait for it to finish.

Learn more about AI Assistant.

26.5.3. Handover Preview (New)

- You can now show a preview of the target agent before a handover starts, using the new **Show as detail** toggle on Handover interactions.

Learn more about Handover Interaction.

26.6. May 21, 2026

26.6.1. Web Client – Sign-in & Home Screen (New)

- You can now sign in to the Web Client. The home screen shows your **Featured agent**, **Recent agents**, and other entry points.

Learn more about Sign-in & Home Screen.

26.6.2. Push-to-Talk on Web Client (New)

- Push-to-talk mode is now available in the Web Client and is on by default. The client listens only while the button is held, which is useful in noisy environments or when you want precise control over when the agent is listening. You can toggle it off in Settings.

Learn more about Push-to-Talk.

26.6.3. Insights on Web Client (New)

- Signed-in users can now see Insights on the Web Client. When an agent has user-facing evaluations enabled, the Summary view appears directly on the agent screen with score cards for each insight type, including average, minimum, and maximum values along with a distribution chart.

Learn more about Insights on Web Client.

26.6.4. AI Assistant (Improvements)

- Refreshed chat UI: previous messages are no longer rendered on screen, so each session starts clean. AI Assistant still retains the context of the conversation so you can keep referring back to earlier topics.
- The UI is now temporarily blocked while AI Assistant is applying changes, preventing

conflicting actions until the update finishes.

Learn more about AI Assistant.

26.6.5. Archiving Agents (New)

- You can now archive agents from the Agents list to keep your project view focused on what's active. Archived agents are hidden by default; toggle **Show archived** at the top of the list to see them again. Archiving is permanent and cannot be undone.

Learn more about Archiving Agents.

26.6.6. Handover Interaction (Improvements)

- The agent picker for Handover now lists only published agents. Drafts and unpublished revisions no longer appear, so handovers always point to a working version.
- When the target agent is republished, the handover automatically picks up the latest published revision. You no longer need to reconfigure the handover after each publish.

Learn more about Handover Interaction.

26.6.7. Lightweight Agents (New)

- When creating an agent, you can now pick an **Agent Type**: **Relational** for deep interactions with parallel thinking, or **Lightweight** for a smaller reasoning footprint and cheaper token spend. The choice is locked once the agent is saved.

Learn more about Agent Type.

26.7. April 30, 2026

26.7.1. Insights (New)

- Signed-in users can now see Insights about their own conversations with an agent. When the agent has user-facing evaluations enabled, an **Insights** button on the agent screen opens a summary with average, minimum, and maximum scores along with a distribution chart for each insight type. Available on iOS and Android.

Learn more about Insights.

26.7.2. Visual & Performance Updates (Improvements)

- We continued refining avatar realism with new wrinkle support in the skin shader, more natural mouth shading, and softer, cleaner skin edges for a warmer, more lifelike appearance.
- We refined depth of field for a more natural focus on the avatar and improved texture quality on the Web Client.
- We polished avatar behavior at session start. Expressions, head positioning, and eye rendering now feel more consistent and natural from the very first moment.
- Visuals now stay consistent when the on-screen keyboard appears.
- We improved audio reliability so playback stays smooth and responsive across longer sessions and on the Web.

26.8. April 16, 2026

26.8.1. Skills Reference in Agent Configuration (New)

- When Skills are available in your project, you can now reference them in your agent's configuration.

26.8.2. Evaluations (New)

- You can now create evaluations to measure how your agent is performing against criteria you define. Results are available in the Analytics Suite.

Learn more about Evaluations.

26.8.3. Soniox Speech Recognition (Engine)

- Enabled Soniox speech recognition to improve the quality of voice transcription.

26.8.4. Video Interactive Content (New)

- Relational Agents can now play video content during a conversation.

Learn more about Interactive Content.

26.8.5. Agent Templates (New)

- When creating a new Relational Agent, you can now choose from a set of pre-configured

templates. Each template includes its own Template Variables, a set of fields you fill in to tailor it to your context.

26.9. March 5, 2026

26.9.1. Conversation Transcript Retrieval (Engine)

- Agents can now retrieve the full transcript of a specific past conversation. This allows them to accurately recall what was discussed and the exact words used in a previous session, not just a memory summary.

26.9.2. Knowledge Base Source Attribution (Engine)

- When an agent retrieves content from a Knowledge Base, it can now identify exactly where that information comes from: the document, the page, the section, and the position on the page (top, middle, or bottom).

26.9.3. Camera Mode (New)

- Agents now have a configurable camera mode. Choose **Dynamic** for a cinematic, moving camera or **Static** for a fixed, video-call-style view.

Learn more about Camera Mode.

26.9.4. Rich Text Formatting in Agent Configuration (New)

- You can now format text across agent configuration fields use headings, bullet lists, bold, and more to write clearer, better-structured instructions.

Learn more about Configuring Agents.

26.9.5. Account Settings (New)

- You can now manage settings at the account level. Update your account name and add a profile image directly from Account Settings.

26.9.6. Visual & Performance Updates (Improvements)

- Improved skin rendering for a more lifelike appearance. Light now interacts with avatar skin

more naturally, passing through and softly glowing beneath the surface. This gives avatars a warmer, more human look.

- The camera has been reworked from the ground up to better follow the tempo of the conversation, frame the avatar more naturally, and feel more natural overall.
- Lip movement now tracks speech more closely, making the avatar feel noticeably more present and natural to talk to.
- Improved adaptive performance reliability on iOS and Android.

26.10. February 12, 2026

26.10.1. Extractor Insights in Analytics (New)

- You can now see the outcomes of your custom extractors directly in the Analytics Suite.

Learn more about Analytics Suite.

26.11. February 5, 2026

26.11.1. Agent Revisions (New)

- Track your agent's configuration history with automatic versioning. Every time you save changes, a new revision is created. You can preview, publish, or revert to any previous revision.

Learn more about Revisions.

26.11.2. Mobile Image Upload (New)

- New controls include taking photos with the camera or choosing images from the gallery on OpenAI agents. This enables more interactive and contextual discussions.

26.11.3. Visual & Performance Updates (Improvements)

- Reduced the frequency of idle animations for a less distracting experience.

26.11.4. Visual Quality Settings (Improvements)

- Visual quality settings have been simplified. Scene resolution is now automatic, and the new "Auto" quality mode adapts to your device's performance, no manual adjustments needed.

26.11.5. Kiosk Mobile Microphone (New)

- Kiosk stations now support using your mobile device as a microphone. Scan the QR code on the kiosk with your phone to use it as the audio input for the session.

Learn more about Kiosk Mobile Microphone.

26.11.6. Extractors (Improvements)

- We made improvements to extractors helping your agents track and optimize specific metrics during conversations.

26.11.7. Handover Multimodal Interaction (New)

- Agents can now transfer conversations to other relational agents using the new Handover multimodal interaction. Configure automatic or specific agent routing to escalate conversations or direct users to specialists.

Learn more about Handover Interaction.

26.12. January 15, 2026

26.12.1. Widget Embedding (New)

- You can now embed Relational Agents directly into your website using the PromethistAI Widget. Add a simple script tag to display a floating assistant that users can interact with.

Learn more about Widget Embedding.

26.12.2. Web Client – Expanded Platform Support (New)

- The Web Client now supports Android devices and Linux desktop browsers, making it accessible on even more platforms without installing a native application.

Learn more about Web Client.

26.12.3. Visual & Performance Updates (Improvements)

- Visual quality now automatically adjusts to your device for smoother operation (Android and iOS only).

- Performance has been improved in multiple areas for a faster, more stable experience.

26.13. December 11, 2025

26.13.1. AI Assistant (New)

- You can now control the platform through a conversational interface. Create agents, navigate between sections, and make configuration changes using natural language instead of clicking through menus.

Learn more about AI Assistant.

26.13.2. Web Client (New)

- Access your Relational Agents directly in your browser without installing any applications. The Web Client is now available for Windows (Chrome, Edge) and iOS 16+ (Safari) with WebGPU support.

Learn more about Web Client.

26.13.3. Integrations (New)

- Connect your agents to external services and data sources using webhooks and MCP servers.

Learn more about Integrations.

26.14. November 27, 2025

26.14.1. Visual & Performance Updates (Improvements)

- We improved avatar visual fidelity with a new realistic hair shader and better lighting, including fixes for hair blurring.
- We refined overall movement and scene behavior with reduced stance-change swinging and a unified, simplified loading sequence across Avatar and Environment.
- We improved compatibility and reliability, including fixes to static camera behavior.

26.14.2. Turn Taking (New)

- We added Turn Taking option in Settings (iOS and Android), when enabled, prevents interruptions.

Learn more about Turn Taking.

26.14.3. Interaction Continuity for Kiosks (New)

- Kiosks can now display a QR code that mobile users can scan to start a personalized conversation

Learn more about Interaction Continuity.

26.15. November 13, 2025

26.15.1. Global Assets (New)

- Project assets (Knowledge Bases, Integrations and Interactive Content) can now be marked as Global to make them accessible to all agents within the same project

Learn more about Assets Visibility

26.15.2. Featured Agent Display (New)

- A primary featured agent now displays prominently at the top of the home screen for signed-in users

---

26.15.3. Push-to-Talk on iOS (New)

- Push-to-talk mode is now available also on iPhone

Learn more about Push-to-Talk.

26.15.4. Audio & Performance Updates (Improvements)

- We improved audio streaming with reduced buffering time for faster, more responsive voice interactions
- We improved scene and avatar loading reliability

26.16. October 30, 2025

26.16.1. Extractors (New)

- Let agents optimize conversations for custom, measurable goals (e.g., brand sentiment, product knowledge, issue detection).

26.16.2. Remembers users toggle (New)

- Choose whether agents recall past conversations for each user

Learn more about Remembers users (memory control).

26.16.3. Filter analytics by agent version (New)

- Compare performance across versions to see what changes worked

Learn more about Filters in the Analytics Suite.

26.16.4. Silence Detection (New)

- The agent detects when you stop speaking and adjusts continuing or offering to end the session

Notes: The silence timeout is configurable per agent.

26.16.5. QR Button Relocation (UI)

- The **QR** button has moved from the **Identity** sidebar to the **main Talk screen** (lower-right corner) for quicker sharing and device pairing

---

26.16.6. Visual & Performance Updates (Improvements)

- We improved conversational responsiveness with faster audio processing
- We made avatar animations smoother across different devices
- Cleaner, smoother background blur with less pixelation

26.17. October 16, 2025

26.17.1. Visual & Performance Updates (Improvements)

- We made major performance improvements with an updated rendering system

---

- We refined avatar posture and movement to look more natural
- We improved audio reliability for smoother conversations

26.17.2. New Planning (Coming Soon, Beta Preview)

- We're putting the finishing touches on a new planning component that will let you opt in per agent for quicker reasoning and more reliable function calls

---

26.17.3. Vision Sensors on Kiosks (New)

- If enabled, kiosk agents can now use the device camera to perceive their surroundings (e.g., detect presence, read signage, or react to context)

Notes: Requires camera permissions and adherence to your organization's privacy settings.

26.17.4. Image Modality Visibility (Fix)

- Addressed cases where images weren't shown to users; display is now consistent

---

26.17.5. Conversation Ending Reliability (Improvements)

- The agent automatically shows a close-session prompt when it detects that you're done with the conversation

---

26.17.6. Mobile Region Switching (New)

- In the PromethistAI iOS and Android apps, you can now switch between US and EU clusters directly in Settings → Region. The default is autodetected based on your location for optimal performance

Learn more about Processing Region.

26.17.7. Cross-Region QR Scanning (New)

- Now you can scan a QR code for an agent hosted in a different region (e.g., EU QR on a US-configured device). A warning will appear, but the session will launch without interruption

Learn more about Cross-Region QR Scanning.

26.18. October 02, 2025

26.18.1. Switch between US and EU clusters (New)

- You can now choose where your agent runs (US or EU) and which backend cluster the Administration Console connects to

Learn more about Switching Clusters.

26.18.2. Rename Knowledge Bases (Improvements)

- Keep your workspace tidy and searchable. You can now rename any Knowledge Base from its menu or within the Knowledge Base settings page

---

26.18.3. Accepted Languages selector (UI & Controls)

- Limit which languages an agent will speak using the new Accepted Languages control

---

26.18.4. Anonymous Mode (New)

- Want a little extra privacy? Start a session without linking it to your identity: toggle **Anonymously** under the agent's **Talk** button (per session)

Learn more about Anonymous Mode.

26.18.5. Agent animations (Improvements)

- Your agent just got smoother: new "thinking" animations, better lip-sync, and steadier resolution scaling for a more natural feel

---

If you have questions or want help enabling any of these features, reach out to PromethistAI Support.