Product Specific Terms

Product	Ibexa Agentic Marketing Platform
Product description	A platform that enables the Customer to build, manage, and run AI Agents (configured automated actors that accept input and produce a result). It connects LLM (Large Language Model), and external tools coming from MCP servers.
Services provision	SaaS licence
Version	2026 Sep

OBLIGATIONS OF THE CUSTOMER 3
USAGE AND LIMITATIONS 3
AGENTS CONFIGURATIONS AND THIRD-PARTY TOOLS 4
CUSTOMER DATA UPON CONTRACT TERMINATION 5
ARTIFICIAL INTELLIGENCE 5
SECURITY AND SAFEGUARDS 5
SERVICE LEVEL AGREEMENT 6
SUPPORT 6
SERVICE CREDITS 7
DATA PROCESSING AGREEMENT 8
DATA PROCESSING ACTIVITIES 8
LIST OF SUB-PROCESSORS 10
SECURITY MEASURES 15

OBLIGATIONS OF THE CUSTOMER

USAGE AND LIMITATIONS

Resources usage

The Customer is granted a monthly AI credit allowance. An "AI credit" represents a standardized unit of consumption corresponding to actions or operations performed within the Product (such as Agent execution or automated workflows). Unused AI credits do not roll over to next month. The Customer may purchase additional AI credit packs at any time.

The Customer is responsible for maintaining a sufficient AI credit balance to avoid service interruption. Thresholds and notification rules for low credit alerts are specified in the Product's technical documentation and user interface. In the event of zero remaining AI credits or non-payment, the Provider reserves the right to suspend or block access to the Product.

Fair use limits

To ensure performance, maintain high system reliability, and protect the infrastructure against security risks or abnormal usage, technical fair use limits apply.

To prepare for and mitigate risks associated with suspicious actions, phishing attempts, or security threats, the Provider reserves the right to suspend or block access for any of the Customer's Users immediately.

Access suspension, Customer's Users blocking, or feature deactivation may be executed without prior notice if the Customer's activity threatens system stability, violates acceptable use policies, or poses a security threat to the Provider or third parties.

LLM provider and Model Context Protocol (MCP) server

The Provider has the right and technical capability to disable or turn off an LLM provider or an MCP server for a specific Customer or globally for all Customers. The

Customer will be notified without undue delay, and any suspension will be limited in duration and scope to what is reasonably necessary.

Agent monitoring

To maintain system performance, safety, and reliability, the Provider reserves the right to immediately stop, suspend, block, or delete any Agent without prior notice or liability if the Agent:

- is used against the law, or generates, processes, or shares unlawful, harmful, or prohibited content;
- influences the system in a negative way, exhibits abnormal behavior, technical anomalies, takes advantage of the system, or exhibits excessive resource usage; or
- causes or threatens to cause damages or service disruption to the Provider, other customers, or third parties.

The Customer is fully responsible for its Agents' activities and outputs, and agrees to indemnify and hold harmless the Provider against any damages, liabilities, or third-party claims resulting from the Agent's misuse or abnormal behavior.

AGENTS CONFIGURATIONS AND THIRD-PARTY TOOLS

Autonomous Agent execution

The Customer acknowledges that the Product utilizes autonomous Agents capable of executing multi-step workflows, automated tasks, and actions via external Model Context Protocol (MCP) servers and third-party tools. The Customer is solely responsible for configuring, defining, and monitoring the instructions, prompts, workflow rules, and operational scope provided to such Agents, as well as managing access to any connected external systems. The Provider shall not be liable for any unauthorized, unintended, or erroneous execution of actions performed by Agents resulting from the Customer's configuration, instructions, or connected MCP tools within the Customer's environment or third-party systems.

Third-party AI models and integrations

The Product relies on and enables access to third-party AI models (such as Large Language Models), cloud infrastructure, or external MCP servers ("Third-Party Providers").

These AI models may be provided either directly by the Provider or brought by the Customer using its own credentials or infrastructure.

The Customer agrees to comply with all applicable terms of use and acceptable use policies of such Third-Party Providers. The Provider does not control, and accepts no responsibility or liability for, the availability, continuity, performance, policy updates, content moderation, or operational changes of such Third-Party Providers. Any temporary unavailability, rate limiting, or modification of the Product resulting from Third-Party Providers shall not constitute a breach of the Contract by the Provider.

CUSTOMER DATA UPON CONTRACT TERMINATION

Upon termination or expiration of this Agreement, the Provider reserves the right to permanently delete all Customer Data following the effective date of termination. Upon the Customer's request, the Provider will, prior to such deletion, make available the Customer's exportable data that is technically feasible to export in a structured, commonly used and machine-readable format.

ARTIFICIAL INTELLIGENCE

SECURITY AND SAFEGUARDS

The latest and most detailed information on the security measures applicable to the AI Features is available as part of our Technical and Organisational Measures.

DESCRIPTION OF AI FEATURES

The latest and most detailed information on the AI Features included in the Product is available at our Trust Center.

SERVICE LEVEL AGREEMENT

SUPPORT

Channels

The primary Support channel is an email to agentic@support.ibexa.ai.

Hours

Standard Support is available during regular Office Hours (Monday to Friday, 09:00–18:00 UK & France local time), including public holidays.

SERVICE LEVELS

Service uptime

Component	
Agent Runtime	99.0%
Agent Administration	95.0%

Agent Runtime: The underlying infrastructure layer that executes, manages, and secures AI agents in production.

Agent Administration: The back-office environment where agents are built, configured, and managed.

First Response Time

First Response Time is measured within Office Hours.

Priority	Estimated response time

P1	2h
P2	4h
P3	8h
P4	Best effort

Mitigation time

Mitigation times in this table are objectives. Depending on the volume and priority of requests submitted to Support, response and resolution times may vary.
Mitigation times are measured within Office Hours.

Priority	Mitigation time objective
P1	8h
P2	72h
P3	1080h
P4	Best effort

SERVICE CREDITS

Only applicable for Agent Runtime

Uptime	Credit
< 99.0%	2%
< 95.0%	5%

DATA PROCESSING AGREEMENT

DATA PROCESSING ACTIVITIES

Item	Description
Purpose and nature of processing	Nature of processing: The Processor provides a cloud platform that orchestrates autonomous AI Agents. The platform ingests Large Language Model ("LLM"), API configurations, processes natural language inputs, and uses the Model Context Protocol ("MCP") to route data payloads between the LLMs and external tools. The platform acts as a transmission channel for data provided by the Customer to the Agents, processing it strictly under the Customer's operational instructions. Core operations include transient and persisted data processing, real-time tool calling, session state management, and logging agent execution metadata. Purpose of processing: The Processor will process personal data solely to deliver, maintain, and optimize the Product. When the Controller or its users provide personal data to the Agents, the Processor handles such data exclusively in accordance with the Controller's documented instructions. This includes executing real-time automated Agents, interacting with third-party LLMs, exchanging data with MCP connected tools, and troubleshooting execution logs as directed by the Controller. The Processor will not use any AI Inputs to train, retrain, or fine-tune its own or any third-party LLMs.
Categories of personal data	• User profile & account data: Names, email addresses, usernames, and authentication credentials.

	• Unstructured Agent inputs & outputs (Controlled by the Customer): Any personal data voluntarily included by the Controller or its end-users within prompts, system instructions, documents, or context payloads fed into the agents and connected MCP tools.
Categories of data subjects	Authorized users: Employees, Partners, contractors, administrators, and consultants of the Controller who access and configure the platform. Third parties (Controlled by the Customer): Any individuals whose personal data is contained within the data sources, databases, or systems connected to the platform via MCP tools or included in user prompts.
Duration of processing	The duration of processing and retention of personal data varies depending on the specific use case: • AI Assistant conversation history is retained for 30 days; • Files in the Knowledge Base, instructions to the agents or reports shall be manually deleted by the Customer; • Personal data deletion of interactions with agents, third-party LLMs or MCPs must be requested on a case-by-case basis via the Support. Overall retention shall not exceed the Term of this Agreement, subject to prompt deletion or anonymization upon termination unless longer retention is required by applicable law.

LIST OF SUB-PROCESSORS

Sub-processors engaged for hosting infrastructure, servers and databases:

OVH

Item	Description
About	https://www.ovh.com/
Registered offices	2 rue Kellermann, 59100 Roubaix, France
Processing purpose	host infrastructure, servers and databases
Processing location	OVH datacenters are located in Europe (3 datacenters in France, one in Germany and one in Poland)
Security measures	• https://www.ovhcloud.com/en/datacenters-ovhcloud/ • OVH is certified SOC 1 type II (SSAE 16 and ISAE 3402), SOC 2 type II and ISO 27001 • https://www.ovhcloud.com/en-gb/enterprise/certification-conformity/iso-27001-27017-27018/

Zendesk

Item	Description
About	https://www.zendesk.com/
Registered offices	Zendesk France SAS 34 Avenue de l'Opéra, 75002 Paris, France
Contact	euprivacy@zendesk.com
Processing purpose	Online helpdesk ticketing service
Processing location	Zendesk datacenters are located in Europe (Frankfurt Germany)
Security measures	• https://www.zendesk.com/trust-center/ • Zendesk is certified ISO 27001, 27701, SOC 2

Optional sub-processors engaged for providing LLM feature depend on LLM providers selected and configured by the Customer:

OPENAI

Item	Description
About	https://openai.com/
Registered offices	3180 18th St, San Francisco, CA 94110, USA
Processing purpose	Artificial intelligence sub-processor: LLM model processing, natural language understanding, text and data generation via API
Processing location	United States
Security measures	• https://openai.com/security-and-privacy/ • OpenAI is certified ISO/IEC 27001:2022, ISO/IEC 27017, ISO/IEC 27018, ISO/IEC 27701, ISO/IEC 42001, SOC 2 Type II, SOC 3, and compliant with CCPA/GDPR through standard DPA and Standard Contractual Clauses (SCCs)

ANTHROPIC

Item	Description
About	https://www.anthropic.com/
Registered offices	548 Market Street, PMB 15301, San Francisco, CA 94104, USA
Processing purpose	Artificial intelligence sub-processor: LLM model processing (Claude models), natural language understanding, text and data analysis via API
Processing location	United States

Security measures	• https://trust.anthropic.com/ • Anthropic is certified ISO/IEC 27001, SOC 2 Type II, and compliant with CCPA/GDPR through standard DPA and Standard Contractual Clauses (SCCs)

GOOGLE

Item	Description
About	https://cloud.google.com/
Registered offices	Gordon House, Barrow Street, Dublin 4, Ireland (Google Ireland Limited for EU) / 1600 Amphitheatre Parkway, Mountain View, CA 94043, USA
Processing purpose	Cloud infrastructure and Artificial intelligence sub-processor: LLM model processing (Gemini models), hosting, text, and multimodal data processing via Vertex AI / Google Cloud Platform
Processing location	United States
Security measures	• https://cloud.google.com/security • Certified ISO/IEC 27001, ISO/IEC 27017, ISO/IEC 27018, ISO/IEC 27701, SOC 1, SOC 2 Type II, SOC 3 • Compliant with GDPR through EU-U.S. Data Privacy Framework (DPF) active certification

MICROSOFT

Item	Description
About	https://azure.microsoft.com/
Registered offices	One Microsoft Place, South County Business Park, Leopardstown, Dublin 18, Ireland (Microsoft Ireland Operations Limited) / One Microsoft Way, Redmond, WA 98052, USA
Processing purpose	Cloud infrastructure, databases, enterprise software, and Artificial intelligence sub-processor: Azure OpenAI Service and LLM deployment/processing
Processing location	United States
Security measures	• https://www.microsoft.com/en-us/trust-center • Certified ISO/IEC 27001, ISO/IEC 27017, ISO/IEC 27018, ISO/IEC 27701, ISO 22301, SOC 1, SOC 2 Type II, SOC 3, C5 • Compliant with GDPR through EU-U.S. Data Privacy Framework (DPF)

MISTRAL AI

Item	Description
About	https://mistral.ai/
Registered offices	15 rue des Halles, 75001 Paris, France
Processing purpose	Artificial intelligence sub-processor: LLM model processing, natural language understanding, text generation, and hosting via La Plateforme API
Processing location	European Union
Security measures	• https://trust.mistral.ai/ • Certified ISO/IEC 27001, SOC 2 Type II • Fully compliant with EU GDPR with data residency within the EU

SECURITY MEASURES

For the latest version of our Technical and Organizational Measures, please visit our Trust Center.