

Appendix: Glossary

Essential terms and their meaning in the context of Aras InnovatorEdge AI.

Term	Definition
Access Point	A registered Edge API connection as it appears in the Agent Management Platform. Access Points are tenant-scoped and reusable across multiple agents. At least one is required before an agent can be published. See also Access Points in the Administrator Guide .
Activity status	The health and availability state of an agent, separate from its publish state. An agent's activity status is active, disabled, error, or archived.. End users can reach an agent only when it is both active and published.
Agent	A configured AI instance that interprets user questions, selects the appropriate tools, and returns results. Agents are created from templates and governed by solution instructions, platform instructions, bound skills and tools, and bound Access Points.
Agent Library	The central screen in the Agent Management Platform, where administrators create, configure, publish, and manage agents.
Agent Management Platform	The management application for InnovatorEdge AI, where administrators configure agents, models, skills, MCP Servers, Access Points, and access. Reached at its own route, separate from the chat application. Available only to Edge AI Administrators.
Agent Overview	A read-only summary available after an agent is first saved. Displays the agent's identity and configuration state.
Agent type	The reasoning approach an agent uses. Two agent types are available: Deep Agent (plan-and-execute) and ReAct Agent (iterative reason-and-act).
Approval card	The inline card presented when the Aras AI assistant proposes a data change. Shows the action type, target item, and data payload. The user approves or rejects before any change is executed.
Aras AI Assistant	The default governed agent that ships with every tenant. Pre-configured by Aras with locked core settings and editable extension points.
Archived	An agent that has been soft-deleted. An archived agent is hidden from end users regardless of its publish state, and can be returned to service with Restore.
Authentication mode	The authentication method used by an Access Point or an MCP Server. End User propagates the signed-in user's identity, so calls execute under that user's permissions and no credential is stored. API Key uses a static credential held in the secret store, and calls execute under the permissions of the user the key belongs to. MCP Servers additionally support no authentication; End User is not selectable for MCP Servers.
Built-in tools	The Innovator data tools Aras provides with the platform, covering item retrieval, search, BOM retrieval, relationship traversal, and lifecycle state queries. Made available to an agent by selecting them on the agent, from an Access Point the agent is bound to.
Centralized Identity Access Management (CIAM)	Centralized Identity Access Management, the Aras identity service that authenticates users and carries their Edge AI role assignment. Edge AI is registered as a single CIAM application.



Context awareness	The embedded panel's ability to detect the Innovator item the user is viewing and use it as implicit context for questions.
Context chip	The indicator above the chat input showing which Innovator record is currently in context. The user can remove it to ask an unrelated question.
Conversation starters	Suggested prompts configured by the administrator that appear on the start screen when no conversation is active. Selecting one places the text in the input for review before it is sent.
Deep Agent	An agent type that plans before it acts: it writes out the steps it intends to take, then works through them, adjusting as it learns. Suited to questions whose answer must be assembled from several retrievals. The Aras AI Assistant is a Deep Agent.
Delete	Permanent removal of an agent. Unlike Archive, a deleted agent cannot be restored.
Draft	The publish state of an agent that end users cannot reach. Draft agents are visible only to Edge AI Administrators, who can test them in Simulation Chat.
Edge AI	The InnovatorEdge service that delivers conversational intelligence and governed action on top of Aras Innovator data. Depends on Edge API for all data access.
Edge AI Administrator	Tenant-level administrator role. Manages agents, models, skills, Access Points, and MCP Servers within one tenant. Can simulate draft agents and converse with published agents.
Edge AI User	End-user role. Can converse with published agents through the embedded panel or the standalone application. Cannot access the Agent Management Platform.
Edge API	The InnovatorEdge service that exposes Aras Innovator data and actions as governed endpoints. All Edge AI data interactions flow through Edge API. Agents read the Edge API schema at runtime, so changes to an API definition take effect on the next request.
Edge API connection	A registered, tenant-scoped link between Edge AI and a configured Edge API endpoint. Stored once and reusable across multiple agents. Displayed in the Agent Management Platform as an Access Point; the two terms refer to the same object.
Embedded Panel	The in-Innovator access mode for the Aras AI Assistant: a resizable conversational panel that operates alongside the Innovator workspace without interrupting the user's active context.
Ephemeral Interaction	The data handling model for LLM interactions within Edge AI. Data passed to the language model for reasoning is not stored, retained, or used for model training after the interaction completes.
Governed agent	An agent whose core configuration is locked by Aras and cannot be changed by the tenant. For the Aras AI Assistant the locked settings are name, agent type, solution instructions, and platform instructions, and the agent cannot be deleted. The LLM configuration is not locked and can be changed by an Edge AI Administrator.
Human-in-the-loop (HITL)	The requirement that a person explicitly approves any operation that changes Innovator data before it is committed. In InnovatorEdge AI this is a platform invariant rather than a configurable option: there is no setting that enables automatic commit. The validation flow is carried by the tools that change data.
LLM configuration	A registered model provider and deployment that an agent uses for reasoning. Aras provides a default configuration to every tenant, and administrators can register their own. Overriding the Aras default stops the tenant from receiving Aras updates to it.



MCP Server	An external tool server that provides additional tools to agents. The platform discovers the tools a server offers and makes them available for binding to agents.
Metadata Refresh	The process of forcing Edge AI to re-read and re-cache an Edge API schema after the underlying API definition has been updated. Triggered by editing and re-saving the connection in Access Points.
Models	The registry of AI model providers available for agent configuration. Model selectors distinguish between Reasoning Models (used for language understanding and response generation) and Embedding Models (used for semantic search and knowledge retrieval).
Platform instructions	Aras-authored instructions that govern platform-level agent behavior. Locked for the Aras AI Assistant.
Polyitem	An Aras Innovator data concept representing an item type that can resolve to multiple concrete item types depending on context. The Aras AI Assistant can query relationships involving Polyitem types.
Published	The publish state of an agent that end users can reach. An administrator publishes an agent explicitly, once it has an Access Point and the tools it needs selected and saved. Only published agents appear to Edge AI Users.
Role-Based Access Control (RBAC)	The Aras Innovator permission model applied to all data queries executed by Edge AI agents. Users can only retrieve data they are already authorized to see in Innovator.
ReAct Agent	An agent type that works in a single loop: it reads the question, takes one action, looks at the result, and decides the next action. Suited to questions that resolve in a few steps.
Reasoning effort	The depth of reasoning an agent applies to a request. The Aras AI Assistant ships at High.
Restore	The action that returns an archived agent to service.
Schema Caching	The platform's local storage of Edge API metadata. Keeps agent responses fast and ensures the agent works from a consistent view of the data structure. Refreshed by editing and re-saving the relevant connection.
Session context	Short-term conversational context retained within a single user session. Enables follow-up questions to refine previous results. Not retained after the session ends.
Simulation Chat	An administrator-only environment for testing an agent against live data before publishing. Each response includes an expandable reasoning trace. Simulation works on draft agents, so an administrator can test an agent before end users can reach it.
Skill	A reusable instruction bundle that teaches an agent domain-specific query patterns, rules, or behaviors. Skills are bound to agents through the agent editor. Aras-managed skills are read-only; an administrator can disable one on their own agents, and it stays disabled through later updates.
Solution instructions	Instructions that define an agent's behavior, scope, and personality. Locked for the Aras AI Assistant.
Standalone application	The out-of-Innovator access mode for the Aras AI assistant, reached at its own URL, separate from the Agent Management Platform. Offers the same conversational capabilities as the embedded panel.
System guardrails	Internal instructions built into every agent that enforce security, data integrity, and platform governance. Custom instructions that conflict with system guardrails are automatically ignored.



Template	A pre-built agent definition that provides a starting configuration (solution instructions, tool bindings and skill bindings) for a new agent. A template is a one-time copy source: an agent created from it does not inherit later changes to the template and does not receive Aras updates. The Aras AI Assistant is not delivered through a template.
Tenant	The organizational account unit within InnovatorEdge. Access Points, agents, models, skills, MCP Servers, and user access are all scoped to a tenant.
Test Connection	A check available on an Access Point that uses API Key authentication, confirming that the endpoint is reachable and the credential is accepted. The platform reports the result.
Tool	A single callable capability available to an agent, such as retrieving an item or traversing relationships. Tools come from Access Points, for Innovator data, or from MCP Servers, for external systems, and are selected per agent. Tools that change data carry the human-in-the-loop approval flow.

