

Frequently Asked Questions

Setup and Prerequisites

Do I need Edge API to use Edge AI?

Yes. Edge API is a prerequisite. All data interactions between agents and Aras Innovator flow through configured Edge API connections. Edge API must be active and connected to your Innovator instance before enabling Edge AI.

Can Edge AI connect directly to the Aras Innovator database?

No. All queries are executed through Edge API endpoints under the user's own session and permissions.

Which Aras Innovator versions are supported?

Edge AI is compatible with all Aras Innovator versions supported by Edge API. See the [Edge API documentation](#) for the supported version matrix.

Is Edge AI available on-premises?

Edge AI is delivered as a cloud service. Your Aras Innovator instance can remain on-premises or in a hybrid deployment, and Edge AI reaches it through the Edge API layer. You can also point Edge AI at your own model provider, including a self-hosted, OpenAI-compatible endpoint, so that inference runs on infrastructure you control. On-premises deployment of Edge AI itself is not part of version 2.0.

Do I need to install anything in Innovator?

Yes. The Aras Innovator AI Assistant Package must be installed on your Innovator instance to enable the embedded assistant panel. This package is not included in the Edge AI service and must be applied independently.

Does the embedded assistant consume a license?

Yes. Opening the assistant from the Innovator command bar consumes an Aras.AIAssistant license.

The agent returns a '401 Unauthorized' error message



The wrong Innovator user is used, or the user is misconfigured. Check the following:

- The signed-in user must exist in Aras Innovator. Task Agents authenticate using the identity of the signed-in user. API Key authentication is not supported. The email address registered in Aras CIAM must correspond to an active Aras Innovator account.
- The CIAM-Innovator user mapping must be correct. Check the user mapping rule in Authentication flow.
- The Innovator user must meet the following conditions: `is_enabled` is set to `true` , and the user has the relevant permission on the items the agent accesses.
- Authorization is specific to the API schema instance. The endpoint URL and the audience value must correspond to the same schema.
- The Innovator instance must have a valid trusted domain certificate for the Innovator OAuth server. End User authentication cannot verify the requesting user's identity without it. See [Connect your Aras Instances](#) .

For the full Edge API context, see "My API call returns a '401 Unauthorized' error message" in the [Edge Builder FAQ](#) .

Roles and Access

What roles does Edge AI use?

Two. Edge AI User can chat with published agents and has no access to configuration. Edge AI Administrator has full control over agents, skills, tools, Access Points, and model configuration within the tenant. Where a user holds more than one role, access is merged.

What happens if a user has no Edge AI role?

They cannot reach the chat or the Agent Management Platform. They see an access denied message with instructions to contact an administrator. Assigning roles is a required onboarding step, not an optional one.

Is access controlled per agent or per resource?

Access is role-based at application level. Per-resource permissions, ownership-scoped access, and a separate agent builder role are not part of version 2.0.

Agents and Configuration



What is the Aras AI Assistant?

The default agent that ships with every tenant, pre-configured by Aras. An administrator connects it to an Access Point, selects the tools it should use, saves, and publishes it, after which end users can reach it. Administrators can extend it further with skills, MCP tools, and sub-agents, and can change the model it uses. Its name, agent type, solution instructions, and platform instructions are locked, and it cannot be deleted.

Can I change the model the Aras AI Assistant uses?

Yes. An Edge AI Administrator can assign a different LLM configuration and deployment. Once you override the model configuration, the tenant no longer receives Aras updates to it. Tenants that keep the default continue to receive them.

How is the Aras AI Assistant configured out of the box?

As a Deep Agent with High reasoning effort, with all Aras-managed skills selected and agent-to-agent delegation enabled with no sub-agents pre-selected.

Why can I not publish the Aras AI Assistant?

Publishing requires at least one Access Point in the tenant. Add one in the Agent Management Platform, then open the assistant, select that Access Point and the tools you want it to use, and save. You can then publish it.

Does adding an Access Point publish my agents?

No. Publishing is always explicit. An Access Point makes publishing possible, but it is not attached to any agent until you select it there, and no agent becomes available to end users until you publish it.

What is the difference between a template and an agent?

A template defines the core capability and behavior model. An agent is a configured instance of a template: it has a name, instructions, and specific Edge API connections bound to it. The Natural Language Analytics (NLA) template is the starting point for all structured PLM data querying agents.

Can I create additional agents beyond the Aras AI Assistant?

Yes. Administrators can create custom agents from templates. Each agent can have different solution instructions, tools, skills, and Edge API connections. Two agent types are available: Deep Agent and



ReAct Agent.

Can I edit the skills and models that Aras provides?

No. Aras-managed skills, tools, and model configurations are shared to your tenant as read-only and are marked as Aras-managed in the interface. You can disable an Aras skill on your own agents, and a skill you disable stays disabled through later updates. You can also add your own skills alongside them.

How do I receive updates to Aras-managed content?

Automatically, at agent load time for skills, and by data migration for new tools, announced in the release notes. Only the latest version is available; there is no version pinning or rollback. An update never interrupts a conversation in progress, and the next conversation uses the updated configuration. Agents you built yourself do not receive Aras updates, even if they started from the same template.

Can multiple agents share the same Edge API connection?

Yes. Edge API connections are tenant-scoped and reusable. A single connection can be linked to any number of agents.

Can I remove an Access Point that an agent is using?

No. An Access Point cannot be removed while it is attached to an agent. Repoint or reconfigure each agent that uses it first, then remove the Access Point.

Can I use my own AI model provider?

Yes. Administrators can register external model providers from the Models screen in the Agent Management Platform. Supported providers are Azure OpenAI and Microsoft AI Foundry, OpenAI, Anthropic, and Google. OpenAI-compatible endpoints, such as Ollama or vLLM, can also be registered using the OpenAI provider type with a custom endpoint.

What happens if I update an Edge API after assigning it to an agent?

Nothing is needed from you. Agents read the Edge API schema at runtime, so a newly exposed item type, property, or action is available on the next request.

Can I restrict what an agent is allowed to do?



Yes. The scope of data and actions available to an agent is determined by the Edge API connections bound to it and the governance policies defined in those APIs. Agents can only access data and perform operations explicitly exposed through their configured connections.

Can I test an agent before publishing it?

Yes. Use the Simulation Chat, accessible from the agent configuration or Agent Overview. It is visible only to administrators. Each response includes expandable reasoning steps showing what the agent queried and why. Only published agents are visible to end users.

What happens if an Edge API connection used by a published agent becomes unreachable?

Connections are exercised when they are used, so the agent stays published and the failure appears during a conversation: the affected tool returns a connection error and the assistant reports that it could not retrieve the data. Correct the endpoint or credentials on the Access Point in the Agent Management Platform. Where the Access Point uses API Key authentication, you can test the connection and see the result.

Data Changes

Can agents update, create, or delete data in Aras Innovator?

Yes. In version 2.0, agents can create records, update properties, change lifecycle state, and delete items. Every operation that changes data requires explicit human approval before it is executed.

Does the assistant ever change data without asking me?

No. Human approval is a platform invariant, not a configurable option. The agent can reason, plan, and propose freely, but it pauses before any operation that writes to Innovator and waits for explicit approval. There is no setting that enables automatic commit.

Where does the approval happen?

Inline in the conversation. The assistant presents an approval card in the chat stream showing the intended action, the target item, and the data payload, with approve and reject controls. There is no separate screen, queue, or application to switch to.

What happens if I reject a proposed change?



Nothing is written and the action ends cleanly. The conversation continues, so you can correct the instruction and let the assistant propose a revised action.

Does approval bypass my Innovator permissions or workflow rules?

No. Approval authorizes the attempt, not the outcome. The write is executed through Edge API under your own identity, so Innovator permissions, lifecycle rules, and workflow constraints still apply. If Innovator rejects the operation, no data is changed and the assistant reports the reason.

When the assistant plans several steps, do I approve each one?

Each operation that changes data is confirmed individually before it is committed. Read and analysis steps do not require approval. Review the full proposed plan before approving the first step, because approving one step does not authorize the rest.

Is there a record of what the assistant did?

Read operations are logged with item type, item ID, user, and timestamp. Changes committed through the assistant appear in Innovator as changes made by you, since the assistant acts under your identity and not a service account.

Data Access and Security

Can a user see data they are not authorized to see in Innovator?

No. All queries execute under the user's own Innovator session via Edge API. The same permission model that governs access in Innovator applies to every query made through the assistant. This applies to search and relationship traversal as well: items you cannot see are excluded from Bill of Materials (BOM) and where-used results rather than hidden from the display.

Does the agent use a service account to reach my data?

No. Every Edge API call made by an agent carries the signed-in user's own authentication token. Attempts to reach items outside your permission scope return a governed error.

Is my data used to train AI models?

No. Aras does not use customer data, user prompts, inputs, or AI interaction data submitted through InnovatorEdge services to train, retrain, or improve any general-purpose or shared AI or machine learning models. See the Aras InnovatorEdge [Product Specific Terms](#) for full details.



How long does the assistant retain conversation data?

The assistant maintains short-term session context within an active conversation to support follow-up questions. This context is not retained after the session ends and is separate from configuration and audit records, which are persisted by design.

Are my agents and connections visible to other tenants?

No. Agents, Access Points, MCP servers, and model configurations you create are scoped to your tenant. Aras-managed resources are shared into your tenant as read-only and cannot be modified or removed by you.

How are credentials stored?

Credentials for Access Points and external MCP servers are held in a managed secret store rather than in the application database, and are never returned through the interface or the API once saved.

Using the Assistant

What types of questions can the assistant answer?

The assistant handles structured PLM data queries: listing and filtering items, BOM structure and where-used, change analysis, version and generation history, counts and aggregations, and status and ownership queries. It maps plain-language questions to the Innovator schema and retrieves results through configured Edge API connections.

Where can I use the assistant?

In two places, with the same capabilities in both. Embedded as a side panel inside Aras Innovator, which is resizable and can be opened alongside your work, and as a standalone chat reached from edge.aras.cloud/ai. Clicking **Full Screen** from the embedded panel opens the standalone view in a new browser tab.

Does the assistant know which item I have open in Innovator?

Yes. When the panel is open in Innovator, the item on the active tab is offered as context and shown as a chip above the input. You can remove the chip if you want to ask an unrelated question. If no item is open, the assistant works normally with no item context.

How do I get from an answer back to the data?



Every result links to its Innovator record. Clicking a link in the embedded panel navigates Innovator to that item.

What if the assistant misunderstands my question?

The assistant asks a short clarifying question when intent is ambiguous before executing a query. If the result is not what you expected, rephrase the question with more detail. Include the item type, property name, or lifecycle state you are interested in.

Can I refine results without starting a new query?

Yes. The assistant retains session context, so you can follow up with refinements such as "filter to items updated this week" or "sort by lifecycle state" without repeating the original question.

Can the assistant visualize results?

Yes. It can diagram the relationships between items and the structure of the data model, and it can produce a pie chart for a simple breakdown. Ask for the form you want in your question, or ask for it afterwards about a result already on screen.

How do I start over?

Click **New Chat**. This clears the session context, so the assistant no longer refers to anything from the previous conversation.

How do I switch to a different agent?

Use the agent switcher in the header. Switching agents starts a fresh conversation, and you are asked to confirm if the current conversation already has messages.

What are the suggested questions on an empty conversation?

Conversation starters, defined by your administrator. Selecting one places the text in the input so you can review and edit it before running it.

Can I stop a response mid-generation?

Yes. Click **Stop**. The assistant halts its current execution. You can then ask a new question or refine your previous one.

Can I export results from the assistant?



Yes. Tables can be exported to Excel and opened in the Innovator Search Grid. Charts can be exported as images.

Can the agent query items created or modified by a specific person?

Yes. The assistant correctly handles queries that filter items by the person who created, modified, owns, or manages a record. For example: "List all Parts created by Mike Miller" or "Parts modified by Jane Smith."

Can the agent show the history of a specific item?

Yes. Ask the agent to retrieve the change history of an item using natural language. For example: "Show me the history of Part MP0101." The agent returns a list of historical changes for that item.

Do AI-generated answers require human review?

Yes. AI-generated outputs must be reviewed and validated by a qualified person before use in any operational, engineering, manufacturing, compliance, or business decision. See the Aras InnovatorEdge Product Specific Terms for full details.

External Tools

Can agents use tools outside Aras Innovator?

Yes. Edge AI acts as an MCP client, so an agent can be connected to an external MCP server and can discover and call the tools that server exposes.

Who controls which tools an agent can use?

Tools are granted per agent by an Edge AI Administrator. An agent can only call tools it has been granted, and changes to its tool grants take effect on its next run.

What authentication do external MCP servers support?

API key, or no authentication for a public server. An external server is reached with the credential configured on it rather than with the identity of the person using the agent. End-user identity is propagated to the Aras-managed tool server, which is where Innovator permissions apply.

Can one agent call another?



Yes. Agent-to-agent delegation lets an agent route part of a request to a specialist agent you nominate as a sub-agent. Where a sub-agent needs approval for a data change, that approval is surfaced to you in your own conversation.

Limitations and Known Behaviors

Are all ItemTypes supported?

The assistant works against your Innovator schema without hardcoded item definitions, so in principle any ItemType exposed through Edge API can be queried. In practice, ItemTypes belonging to some Aras applications, including Variant Management, Technical Documentation, and Simulation Management, may not be fully queryable where Edge API schema exposure for those applications is incomplete. Validate coverage for these applications before relying on them.

Are there limits on how much data one answer can cover?

Yes. Retrieval is bounded per call and large aggregations may be truncated. When a result is capped or truncated, the assistant says so and asks you to narrow the request.

Why does the assistant sometimes report how many records it used?

For calculations over a set, the assistant states how many records were included, because records with incomplete data are excluded from the calculation. Treat that count as part of the answer.

Will the same question always give the same answer?

Within a conversation, identical questions against unchanged data return consistent results. Wording may vary between runs even when the underlying result does not, so compare the data rather than the phrasing.

What happens if no records match?

The assistant tells you that no matching items were found and states the filters it applied, so you can see whether it interpreted the question as you intended.

Does the assistant filter unsafe or inappropriate requests?

Yes. Content safety filtering is applied to conversations. A blocked message is not processed.

Regions and Data Residency



Where does Edge AI run?

Edge AI runs in the United States, Europe, and Japan. The service balances across those regions automatically, choosing where your tenant runs and shifting capacity as the service scales, so performance and availability stay steady without you managing infrastructure.

Can I have processing stay in a specific region?

Yes. Where your organization requires processing to remain within a specific region, that is configured for your tenant on request. Contact your Aras account representative.

Can I keep inference inside my own environment?

Yes. Register your own model provider and point it at an endpoint you choose. Prompts and responses are then processed within that environment rather than by the Aras-provided model.

Languages

Can I ask questions in a language other than English?

Yes. The assistant understands the languages supported by the underlying model, and by default it answers in the language you asked in. You can also ask it to reply in a different language at any point in the conversation.

Can an administrator set the response language?

Yes. Response language can be steered through the agent's instructions, for example to standardize on one language across a team regardless of how the question was phrased.

Are PLM data values translated?

No. Item names, descriptions, and property values are returned as they are stored in Innovator. The assistant translates its own explanations and summaries, not your data.

Getting Help and Support

Where do I find the product documentation?

Click **Documentation** in the help menu in the global header. This opens the Edge AI section of docs.aras.com in a new tab.



How do I contact support?

Click **Contact Support** in the help menu. This opens the Aras subscriber portal.

What details should I include in a support case?

The About panel shows the product name and release version, each individually copyable. The Profile menu shows your tenant name, with a copy action that places the tenant ID on your clipboard. Include the release version and tenant ID, along with the agent name and the question you asked.

Should I sanitize what I send to support?

Yes. A support case leaves your organization, so review transcripts, error messages, and screenshots before you attach them, and sanitize or anonymize anything covered by your intellectual property or confidentiality obligations. Real part numbers, program names, supplier names, and specifications are rarely needed to diagnose a problem. What support needs is the shape of the issue: the item type, the property, the action attempted, and the response received.

