

User Guide

Introduction

Most of what you need to know during a working day is already in Aras Innovator. Finding it has meant knowing where to look, knowing how to build the query, or asking the person who does.

The Aras AI Assistant answers the question instead. You ask in plain language and it answers from your own Innovator data, with a link to the record the answer came from.

It can also change data when you ask it to, and it always asks you to approve the change before anything is written. See **Changing Data** on this page.

Who This Guide is For

Anyone who works with Aras Innovator data. No knowledge of the data model or of any query language is assumed. If you can describe what you want in a sentence, you can use the assistant.

Check an Answer Before You Rely on It

The assistant works from your live Innovator data and shows you where every answer came from, which makes checking straightforward. Do check, whenever the answer is going into a decision.

AI-generated output must be reviewed and validated by a qualified person before it is used in any operational, engineering, manufacturing, compliance or business decision. Understanding an Answer describes what to look at, and Terms of Use gives the requirement in full.

Before You Start

The assistant becomes available once your organization's tenant is configured and you have been granted access. If you cannot see it at all, or you can sign in but no agents are listed, contact your Edge AI Administrator.

AI Where You Work

You can use the assistant in two places: inside Aras Innovator, alongside whatever you are working on, or on its own when you are not in Innovator at all.




Starting with InnovatorEdge AI

Go to `edge.aras.cloud/ai` . The home page opens with a greeting that addresses you by name and shows the name of your tenant. A **Start a conversation** section lists the agents published for you. Choose one and the conversation opens.

Most organizations publish the Aras AI Assistant, which answers general questions about Aras Innovator data. Some publish additional agents for particular kinds of work. Where you see more than one, the description under each name tells you what it is for.

If you are an Edge AI Administrator, the home page also shows a **More for you** section with a **Manage agents** card. Selecting **Go to Agent Library** opens the **Agent Library**. The [Administrator Guide](#) covers configuration.





Inside Aras Innovator

Open the assistant panel and it appears beside your work, resizable and dismissible at any time.

Everything you can do in the standalone chat you can do here, and one thing more: the panel knows which record you have open and uses it as context, so you can ask about the item in front of you without naming it. See **Working in Context** on this page.

For a longer session, expand the panel to full screen. Your conversation and your context come with you, and collapsing it returns you to the panel with everything intact.



The screenshot displays the Aras INNOVATOR software interface. The top navigation bar shows the 'Admin Dashboard' and 'Parts' sections. The main content area is titled 'MP2941' and includes a 'BOM' (Bill of Materials) section. The BOM table lists various components and their quantities:

Seq.	Part Number	Rev.	Name	Type	Quantity	created_by_id	State	owned_by_id	Unit	Reference Designator
5	MP2194	B	Thing-O-Matic 2 Front Panel	Component	1	Mike Miller	Released		EA	
10	MP2196	A	Thing-O-Matic 2 Back Panel	Component	1	Mike Miller	Released		EA	
15	MP2199-2	A	New Thing-O-Matic 2 Right Panel	Component	1	Terry Adams	Released		EA	
20	MP2193	C	Thing-O-Matic 2 Top Panel	Component	1	Mike Miller	Released		EA	
25	MP2198	A	Thing-O-Matic 2 Left Panel	Component	1	Mike Miller	Released		EA	
30	MP2195	C	Thing-O-Matic 2 Bottom Panel	Component	1	Mike Miller	Released		EA	

The right sidebar features the 'Aras AI Assistant' chat window, which displays a list of suggested questions and a text input field for user queries.

Which One to Use

- Use the **embedded panel** when your question is about the record in front of you, or when you want to move between answers and records without leaving Innovator.
- Use the **standalone chat** when you are not working in Innovator, or when your question is not about anything you have open.

The one real difference is item context. Only the embedded panel can see what you are looking at.

Switching Agents

Where more than one agent is published, the agent name in the header lets you change to another one.

Switching starts a fresh conversation with the agent you choose, so if your current conversation already has messages you are asked to confirm before it is cleared.

Your First Question

Open the assistant. With nothing in the conversation yet you see its name, a short description of what it does, and a few suggested questions your administrator has set up. Those suggestions are worth reading, because they tell you what this particular agent is good at.



Select one, or type your own. Something similar to:

Show all released Parts changed in the last 30 days

An answer comes back as a list, and each row links to the record in Innovator. That is the whole loop: ask, read, follow the link to the source.

Asking Questions

Ask in the words you would use with a colleague. You do not need to know which item type holds the answer, which properties are involved, or how the data is related. The assistant works that out and tells you what it did.

What You Can Ask

Finding and filtering items. Retrieve items by state, date, owner, attribute, or any combination.

- List all open ECOs assigned to me
- Show released Documents changed in the last 30 days
- Find Parts with no preferred manufacturer

Product structure. Explore assemblies and trace where a component is used.

- Show the BOM for ASM-7700 at rev C
- Where is Part P-4420 used?
- Which assemblies contain capacitor C-100?

Change activity. Review what a change affects, who raised it, and what is in flight.

- What is affected by ECO-2024-0451?
- Who submitted CR-2025-0033?
- Show change requests opened this quarter

Counts and totals. Get a number, or a breakdown.

- How many Parts are in this assembly?
- How many open changes do we have, by type?
- Count Documents per classification



Status and ownership. Check what is waiting on you or on someone else.

- What is assigned to me right now?
- Show items waiting on my approval
- List Parts created by Riley Miller

History. Retrieve what happened to an item, or how it has been revised.

- Show the full history of Part MP0101
- How many revisions does Part MP2954 have, and what state is each in?

Use `full history` or `action history` when you want the record of what was done.

A shorter phrasing such as "history of Part MP0101" may return revision information instead. If that happens, ask again for the action history.

Diagrams and Charts

The assistant can draw what it found as well as list it.

- Diagram the relationships for ASM-7700
- Show me the data model around Part and Document
- Show a pie chart of Parts by classification

Ask for the form you want in your question, or ask afterwards about a result you already have on screen.

Writing a Good Prompt

A prompt that works usually carries three things:

- **the object** you are asking about
- **the action** you want performed on it
- **the scope** or **format** of the answer you want back

Show released Documents changed in the last 30 days



That is action, object, scope. Drop the scope and you get every document. Drop the object and the assistant has to guess what you meant.

You do not need all three every time, and a follow-up question often needs none of them because the conversation already supplies the context. But when an answer is not what you expected, the missing piece is usually one of the three.

- **Name the item type when your question could apply to several.**
 - **Show all open ECOs** is clearer than **show all open changes** because the second could mean several things.
- **Narrow with what you already know.** A date range, a state, an owner or a classification will get you a shorter and more useful answer than a broad question and a long list.
- **Refine rather than restate.** Once you have an answer, build on it with a follow-up instead of writing a new question. See **Working in Context** on this page.
- **Watch what the assistant says it resolved.** If you use a term the system does not recognise, the assistant tells you what it matched, what it could not confirm, and answers with what it was able to establish. That message is the fastest route to a better question, because it tells you exactly which word to change.

Prompts and What They Return

The shape of your question decides the shape of the answer. These examples cover the common ones.

What you ask	What comes back
How many Parts are in 'Released' state?	A single number
List all open ECOs assigned to me	A table of matching ECOs, each row linked to its record
Count Documents per classification	A breakdown, one row per classification
Show a pie chart of Parts by classification	The same breakdown, drawn as a pie chart
Show the BOM for ASM-7700 at rev C	The assembly structure, indented by level
Where is Part P-4420 used?	The assemblies that contain it
Diagram the relationships for ASM-7700	A diagram of the items connected to it
Show the full history of Part MP0101	The changes made to it, with dates and what changed
List Parts created by Riley Miller	Parts filtered by creator. The same works for owned by, modified by and managed by
What is the average number of days for the Design Review activity?	A calculated average, together with how many records it was based on

Asking in Your Own Language



Ask in whichever language you work in. The assistant understands the languages the underlying model supports and answers in the language you used. You can also ask it to reply in a different language at any point in the conversation.

Your PLM data comes back as it is stored. Item names, descriptions and property values are not translated; only the assistant's own explanations are.

Understanding an Answer

An answer is not only the result. It arrives with the record links behind it and with the assistant's own working, and both are there so that you can check it.

Reading the Response

Responses come back as formatted text: paragraphs, lists and tables, with a diagram or a pie chart where you asked for one.

Tables are part of the response rather than a separate grid, so they are for reading rather than for sorting or filtering. If you want the same data ordered differently, or narrowed, ask for it. That is usually quicker than rearranging it yourself would have been.

Where the Answer Came From

Every answer carries links to the Innovator records behind it.

In the embedded panel, following a link opens the record in Innovator and your conversation stays exactly where it was. In the standalone chat, the record opens in a new tab.

Following one link is the fastest check you can make on any answer.

Seeing How the Assistant Worked

Every response carries its reasoning with it, collapsed beneath the answer. Expand it and you see the whole working: the query that was run, what each tool was asked for and what it returned, and the assistant's own description of what it was doing at each stage.

This is not diagnostic detail meant for support. It is the answer's evidence. Open it whenever a result surprises you, because it almost always shows one of two things: that your question was read differently than you meant it, or that the data really does say what the answer says.



Where the agent plans before it acts, the plan appears in a side panel: the steps it intends to take, and the state of each one as it works through them. You can watch it progress and see exactly where it has reached.

The top screenshot shows the Aras AI Assistant interface. The chat history includes a query "What is assigned on me right now?" and a response from the assistant. The response includes a list of tasks: "Call load_skill", "Reasoning", "Call who_am_i", "Reasoning", and "Call odata_schema_fetch". It also shows the arguments for the "Call odata_schema_fetch" task, which are "entity_set_name": "InBasket_Task". The assistant's response continues with a reasoning step, stating "I'm thinking about sorting by due date in ascending order for urgency, but I wonder if that's too complicated—keeping it simple might be better. I should finalize my query with specific fields like id, keyed_name, name, state, status, due_date, start_date, and created_on. I'm considering if I need the instructions field, since it might be lengthy. Including id seems necessary for the table format, even though the user didn't specifically ask for it." The chat input field contains the text "Let me know what you're looking for..." and the character count is 0/500.

The bottom screenshot shows the Aras AI Assistant interface with a more complex query: "which open tasks are assigned to me through an alias identity? And which of them require immediate action?". The assistant's response includes a list of tasks: "Call load_skill", "Call write_todo_list", "Reasoning", "Call who_am_i", "Reasoning", "Call odata_schema_fetch", "Reasoning", and "Call write_todo_list". It also shows a reasoning step: "Thinking... InBasket, with pending tasks waiting for user input while active tasks are already in progress. It seems logical that 'immediate action' means 'Pending' tasks need user attention." The assistant's response continues with a reasoning step: "I might need two separate filters: one for open tasks assigned to me, considering both 'Pending' and 'Active' states, and another for tasks needing immediate action based on due dates or being in the 'Pending' state. I should check if my workflow skill instructions are fully loaded. **Determining immediate action tasks**". The chat input field contains the text "Let me know what you're looking for..." and the character count is 0/500.

The right side of the bottom screenshot shows a plan panel titled "Plan (4)". It lists four steps: 1. "Resolve current user identity" (Done), 2. "Fetch InBasket_Task schema and relevant assignment/state fields" (Done), 3. "Query open tasks assigned through alias identity and identify those needing immediate action" (In progress), and 4. "Query open tasks assigned through alias identity and identify those needing immediate action" (In progress).



When a Result is Bounded

Retrieval is bounded, so a very broad question can come back with part of the answer rather than all of it. When that happens the assistant says so in the answer, and asks you to narrow the request.

Take it at its word. A narrower question does not just avoid the limit, it usually gives you a more useful answer.

Judging an Answer You Will Act On

Three things, and they take about a minute between them.

- **Open the reasoning and read the filters.** Confirm the assistant narrowed the data the way you intended. This catches most misreadings.
- **Follow one link.** Open a single record from the result and confirm it says what the answer says.
- **Read the count.** Where the assistant has calculated across a set, it tells you how many records it used, because records missing the data it needed are left out of the calculation. If that number is lower than you expected, that is the reason, and the answer is describing a smaller set than you had in mind.

Working in Context

The assistant keeps track of two different things: the record you have open, and the conversation you are having. They behave differently and it is worth knowing which is which.

The Record You Have Open

In the embedded panel, the assistant sees the record you are looking at in Innovator. A chip above the input shows which one, so you always know what "this" refers to.

That means you can ask about the item in front of you without naming it.

With ECO-2025-0100 open in Innovator

You: Who submitted this?

Open a different record and the chip follows: your next question uses the new one.

If your question is not about the record on screen, remove the chip and ask without it.

The Conversation So Far



Follow-up questions build on the answer before them, so you do not have to restate anything.

You: Show all open ECOs assigned to me

Twelve ECOs come back.

You: Which of these were created in the last 7 days?

The list narrows to three.

You: Show the affected items for the first one

Each question does less work than the one before it, and gets you further. This is usually faster than writing one long question, and it lets you change direction when an answer is not what you expected.

When the Assistant Asks You Something

If a question could mean several things, the assistant asks before it goes looking. Ask for **the latest change** without saying what kind, and it may ask whether you mean an ECO, an ECR or an ECN.

Answer it and the assistant continues with the narrower question.

Stopping a Response

If an answer is going somewhere you do not need, stop it with the **Stop** button in the input area. The assistant halts, and you can ask something else straight away.

Starting a New Conversation

New Chat clears the current session: the conversation and everything the assistant was holding from it.

It does not change what you have open in Innovator. If you are working in the embedded panel with a record on screen, that record is still your context, because the panel reads it from Innovator rather than from the conversation.

Changing Data

The assistant can change Innovator data, not only read it. Nothing is written without your explicit approval, every time.

What It Can Change



With the right permissions and configuration, the assistant can create a record, update properties on an existing one, move an item to a different lifecycle state, or delete an item.

If your assistant cannot do these things, your organization has not enabled them. Ask your Edge AI Administrator.

Nothing Happens Until You Approve

When the assistant is ready to change something, it stops and shows you an approval card before anything is written.

The card tells you what it intends to do, which item it will act on, and what will be written. For an update it shows the current value alongside the proposed one, so you can see exactly what changes.

The screenshot displays the Aras InnovatorEdge AI Assistant interface. At the top, the header shows 'Aras InnovatorEdge' and a 'Clear chat' button. Below the header, the assistant's name 'Aras AI Assistant' is visible. The main area shows a list of actions: 'Call write_todo_list' (completed) and 'Call odata_item_update' (pending). A large yellow approval card is centered, titled 'Odata Item Update'. It contains a plan step: 'Plan step 2 of 3: Update the Part name to Arm Assembly 20-1000-01'. Below this, it says 'Execute odata_item_update' and shows the entity details: 'entity_set_name: Part', 'id: B03F51ADF40846C8A0E372DCC6992A77', and 'data: { "name": "Arm Assembly 20-1000-01" }'. At the bottom of the card are 'Reject' and 'Approve' buttons. To the right, a 'Plan' sidebar shows three steps: 1. 'Find the latest Part item for AM-1 and verify its current name' (Done), 2. 'Update the Part name to Arm Assembly 20-1000-01' (In progress), and 3. 'Confirm the update result' (Pending). A text input field at the bottom says 'Let me know what you're looking for...' with a character count '0/500'. At the very bottom, a small footer states: 'By using this chat, you agree to Aras Legal Terms and Privacy Policies.'

Approving

Approve and the change is submitted to Innovator under your own identity.

- **If Innovator accepts it**, the change is applied and the assistant confirms it, with the updated record.
- **If Innovator refuses it**, because your permissions do not allow that operation or because a validation rule rejects it, nothing is written and the assistant tells you why.



Rejecting

Reject and nothing is written. The action ends there and the conversation carries on, so you can correct what you asked for and let the assistant propose something different.

What Your Approval Means

Approving authorizes the attempt, not the result.

The change runs under your own Innovator identity, so it can do exactly what you could do by hand and no more. Your permissions, the lifecycle rules and any workflow constraints all still apply.

A change made this way appears in Innovator as your change, in the record's own history.

Important

Read the card before you approve it. It is showing you the change as the assistant understood your request, and understanding is the part worth checking. The permissions behind it are already yours.

Changing Something Back

There is no undo. A change committed through the assistant is a change in Innovator like any other. You can ask for the reverse as a new change, giving the previous value, and it goes through the same approval. The previous value is in the record's history if you no longer have it.

Known Behaviors

A few things behave in ways worth knowing before you meet them.

- **Asking for history may give you revisions.** A short phrasing such as "history of Part MP0101" can return revision and generation information rather than the record of what was done. Ask for the full history or the action history, or just say so in your next message.
- **Some items show an identifier rather than a name in the context chip.** The chip may show an item type and a number instead of a readable name. The context is still correct; only the label is affected.
- **Stopping a response ends that turn completely.** If you stop the assistant mid-answer, the next question in the same conversation starts clean. If the assistant then behaves differently than it did before you stopped it, start a new conversation.

Getting Help



Ask Your Administrator First

Your Edge AI Administrator can resolve most things faster than support can, because they are configuration rather than faults.

Contact them when the assistant is not there at all, when you sign in and no agents are listed, or when the assistant cannot answer about data you know exists. That last one usually means the data has not been made available to the assistant yet, which is something they can change.

Documentation and Support

Documentation and the Aras subscriber portal are both reachable from the help menu.

You can raise a problem in the subscriber portal yourself. Include the agent you were using, the question you asked, what came back, and what you expected instead.

Important

A support case leaves your organization. Before you attach a conversation, an error message or a screenshot, check it for anything covered by your organization's confidentiality or intellectual property obligations, and replace real part numbers, program names and supplier names with placeholders. What support needs is the shape of the problem, not your design data.

Terms of Use

AI-generated output requires human review. 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.

Customer data is not used for model training. 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.

Session data is not retained. The assistant maintains short-term context within an active conversation so that follow-up questions work. That context is not stored or retained after the session ends.

