

AI Assistant

Generated on October 14, 2026

DataOrchester's AI assistant works inside the Site, next to what you are building. It has three chats: one beside the Program screen's editor, one beside the dashboard editor, and the general assistant in a tab of its own. They read your configuration, the documents you attach and DataOrchester's documentation. The two editor chats propose changes that you review on a card; the general one explains, diagnoses and suggests a plan. This guide covers what each does, how to turn the assistant off and on, how to work with it, what it sends and what it never sends, what it costs, and what to do when it is off.

Important

The assistant proposes; you decide. Nothing it proposes changes your configuration until you apply its card, and nothing reaches the disk until you press **Save**, through the same checks as a change made by hand. It never saves, never starts or stops the engine, and never writes a value to the plant.

The three assistants

Assistant	Where	Ask it to	What it may change
Program	Beside the editor of the Program screen, in the diagram and in the table view	Create, fix and review modules, variables and formulas: <i>"create the Modbus master for this PLC from the attached register list", "why is <code>FIT_501_ALARM</code> never true?", "add the daily average of <code>LIT_301</code>", "review this program and fix its load warnings"</i>	Variables and modules, through change cards. Peers and schemas are read-only to it
Dashboard	Beside the dashboard editor	Design and fix a view: <i>"draw an overview of the dosing station from this P&ID", "make the alarm frame red when any pump fails", "align the pump symbols"</i>	The dashboard open in that editor, through change cards
General	The Assistant tab, opened from the menu entry Assistant	Explain and diagnose: <i>"what does this plant do?", "explain how <code>FIT_501_ALARM</code> is computed", "why has pump 2 stopped, and how do I get it back?", "what would <code>FIT_501</code> read if the PLC sent 20000?"</i>	Nothing: it suggests action plans, and hands configuration changes to the editor chats

All three read the whole configuration, not only the screen they sit on: variables, modules, peers, schemas, dashboards and panels, with the unsaved edits of your open Program tab, and the instance's licence and engine state. So the dashboard chat can bind a gauge to a variable it found in a module's table, and the program chat knows which dashboards and loggers use a variable before it changes it. They also:

- look things up in DataOrchester's documentation, always;
- search the public web - PLC manuals, protocol specifications, standards - only when **Web search** is on;
- read live values, module health and messages and recent console errors only when **Live values** is on.

Before an editor chat proposes a change, it checks its own draft: it validates the whole change, and on the Program screen it also simulates new and changed formulas on test values, in a sandbox that touches nothing. Whatever the assistant did, the Site validates the change again before it shows the card.

On by default

The assistant is on on every instance that has a valid licence and can reach the Portal: nobody needs to turn it on. A user with the Configuration role decides the rest, for the whole instance:

1. Choose **Settings** in the chat's header, beside an editor or in the **Assistant** tab.
2. Read the disclosure: what is sent, to whom, and the AI providers that serve the assistant at present.
3. Choose the two switches. Both stay off unless you tick them:
 - **Web search**: the assistant may search the public web through OpenRouter and its search engine, and Amtiri's Portal fetches the pages it cites;
 - **Live values**: the assistant may read the current values of variables and the health of modules.
4. Choose **Save. Turn off**, in the same window, turns the assistant off for everyone; once it is off, the same window offers **Turn on the assistant**.

What was chosen is kept with who chose it and when, in `entities/.assistant`: a hidden file that no configuration export carries, so it never travels to another installation. An instance where nobody changed anything has no such file, and its assistant is on. Turning it on and off is written to `logs/security.log`.

All three chats need the Configuration role. Every message, attachment, applied card and action-plan button is checked for it on the server.

When the assistant is off

The assistant shows only when all of the following hold, checked in this order. Beside an editor, the chat stays folded when one fails: choose **Assistant** on the editor's toolbar to open it and read the reason. The Assistant tab shows the reason in place of the chat.

Reason	The chat says	What to do
The server forbids it	<i>The AI assistant is switched off on this server.</i>	The operator set <code>application.assistant.enabled</code> to <code>false</code> : see the assistant's keys
Your role	<i>The AI assistant needs the configuration role.</i>	Ask a security administrator for the Configuration role
No licence	<i>This instance has no valid licence, so the AI assistant is off.</i>	Put a licence on the instance: Licence Activation . The two hours an unlicensed instance runs after each start do not count
An offline licence key	<i>This instance runs on an offline licence key, so the AI assistant is off.</i>	An air-gapped instance never calls the Portal, so it has no assistant
Turned off	<i>The AI assistant was turned off for this instance.</i>	A Configuration user turns it on again in Settings , as above
No Portal	<i>This instance cannot reach the Portal, so the AI assistant is off.</i>	The instance must reach <code>dataorchester.com</code> over HTTPS, as it does for its licence
Not in the licence	<i>This organization's licence does not include the AI assistant.</i>	The edition of the organization's licence must include it: see Pricing

No licence, or no connection to the Portal, means no assistant. The instance asks the Portal whether the assistant is available and keeps the answer for ten minutes, or for one minute when the Portal could not be reached, so a licence changed on the Portal reaches the chats within ten minutes and a connection that returns within one.

Working with a chat

Type in the box at the bottom of the chat: **Enter** sends, **Shift+Enter** starts a new line. The assistant answers in the language of your latest message, and shows a line for each tool it uses while it reads the configuration, the documentation or your attachments. **Stop** ends the answer being written.

New conversation starts afresh. A conversation is never stored: closing the tab or starting a new conversation drops it, with its attachments.

Attachments

Attach files adds documents to the conversation: a register list, a P&ID, a manual, a photo, a spreadsheet, a screenshot of an old HMI.

Kind	Files	Largest file
PDF	<code>.pdf</code> , up to 1,000 pages: the text of each page, and each page as a picture for drawings	32 MiB
Images	<code>.png</code> , <code>.jpg</code> , <code>.jpeg</code> , <code>.gif</code> , <code>.bmp</code> , <code>.tif</code> , <code>.tiff</code>	32 MiB
Spreadsheets	<code>.xlsx</code> , <code>.xls</code> : one table per sheet, up to 100,000 rows and 200 columns	8 MiB
Word documents	<code>.docx</code> : headings, paragraphs and tables	8 MiB
Text	<code>.csv</code> , <code>.txt</code> , <code>.md</code> , <code>.json</code>	8 MiB
SVG	<code>.svg</code>	8 MiB

- Any other type - WebP, HEIC, `.doc` , `.ppt` and the rest - is refused, with the list of the types above.
- A conversation holds at most 10 attachments, and 128 MiB together unless the operator changed `application.assistant.attachment-max-bytes` . The ✕ on an attachment's chip removes it.
- Each chip shows the file's pages and about how many tokens reading it whole would take.
- Attachments are held in memory with the conversation, never written to the instance's disk, and dropped when it ends. A page leaves the instance only when the assistant reads it or looks at it.

Page ranges. The assistant reads a long document a few pages at a time, in ranges such as `3` , `1-5` or `1-3, 8, 10-12` , never all of it at once. Tell it where to look - *"the analog registers are on pages 34 to 36"* - and it reads only those, which is faster and costs less.

Documents are data. What a document says is never an instruction to the assistant: when a page tells it to *"ignore your rules"*, it says in one line that the page holds an instruction it did not follow, and goes on.

Change cards

An editor chat proposes each change on a card:

Part	Shows
Header	The title; how many records or components it creates, changes and deletes; Writes to the plant when the change adds or alters something that writes to the plant
Summary	What changes, and what to check
Changes	Field by field: what is added, what changes from what to what, and what is removed
Checks	The self-checks run on exactly this draft - Validated , Simulated - or <i>Not self-checked</i> when none ran
Sources	The attachment pages, documentation sections and web pages the change relied on. An attachment page opens in a new tab; a web page opens only when a search cited it or you typed its address, and its site is always shown beside its title
Assumptions	What the assistant assumed, and what to check before saving and at commissioning

Then choose:

- **Apply**: the change goes into the editor as one step. The tab has unsaved changes, exactly as after a change made by hand, and **Undo** takes the whole change back at once.
- **Apply and continue**: the same, and the assistant goes on with the next step of its work.
- **Discard**: the change is dropped. Write a comment in the box under the buttons first, and the assistant reads it.

Checked again when you apply. The assistant built the card on the records as it read them. If one of them has changed since - you edited it, applied another card, or saved it in its dialog - the card is refused with *Not applied: what it was built on has changed since the assistant read it*, and the list of what changed. Nothing is applied: ask the assistant to read again, and it rebuilds the change on what is there now.

An applied change is then yours: review it in the editor, **Save** it, and commission it like any other change. Each applied card is written to `logs/security.log` with who applied it, the conversation, the card and the records it touched.

Action plans

The general assistant changes nothing. For a problem you describe, it reads the evidence - module health and messages, values with their time, the configuration, the console - and answers with an **action plan**: the problem, the diagnosis with the evidence it rests on, and the steps in order, checks before changes and changes before restarts.

Step	What it is	Its button
On site	Something to check or do outside the Site: power, cables, network, the PLC's own program, a sensor	-
Look	A screen to look at: the Program editor, a dashboard or its editor, the console, the licence	Open
Operate	Something you do to the engine in the Site: start or stop it, enable a module and save	Open
Change	A configuration change, written as a request for the Program chat or a dashboard's chat	Open in editor
Check	What to read afterwards to confirm the fix	Check again

- Each step names its impact - changes the configuration, stops or restarts the engine, writes to the plant - and how to verify it. Tick a step once it is done.
- **Open in editor** opens the editor, and its chat starts with the step's request as your first message. Its answer is a card you review like any other.
- **Check again** asks the assistant to read the evidence again: it confirms the fix or revises the plan.
- No button runs anything. The assistant is instructed never to suggest bypassing, disabling or silencing an interlock, a trip, a safety function or an alarm as a fix; if you ask for one, it explains the risk and hands the change to the Program chat with that warning. Steps on energised or moving equipment say so, and your site's own procedures come first.

Without **Live values**, a diagnosis rests on the configuration alone, and the plan starts with the checks you can read on screen.

What is sent, and what never is

The assistant reaches the AI through Amtiri's Portal. The instance talks only to `dataorchester.com`, over HTTPS and signed with its own install key, as it does for its licence; it never holds an AI provider's key and never talks to a provider. The Portal checks the licence and the allowance, and passes each request to OpenRouter, which routes it to the AI provider serving that assistant. The disclosure names those providers, and every request is routed only to providers that do not collect the data they are sent.

Sent, through the Portal	Never sent
What you type, and the conversation so far	Secrets: a peer's secret, passwords and every other secret setting reach the assistant as «redacted», and so does a credential written into any text - the user and password in an address, a <code>password=</code> or a <code>token:</code> . A change that copies «redacted» back is refused. <code>env:</code> references stay references and are never resolved
The records the assistant's tools return: variables, modules, peers without their secret, schemas, dashboards and panels, role ids and names, the instance's licence and engine state	Users, role members, identity providers, licence files and the instance's identity
The pages and sheets of your attachments that it reads or looks at	The rest of your attachments
Live values, module health and messages and recent console errors, only with Live values on	Logged history, and the unsaved work of other sessions
Documentation searches, to Amtiri's Portal only	The install id: OpenRouter knows the instance only by a code derived from it that changes every day
The web searches it writes, only with Web search on, to OpenRouter and its search engine; the addresses of the pages it cites, which the Portal fetches	Codes, names, addresses, hosts, plant or customer names and document contents in a web search: the assistant is instructed to search with generic technical terms only

- **The instance** keeps nothing of a conversation once it ends. The security log records the assistant turned on and off and each applied card, never a message or a document.
- **Amtiri's Portal** keeps the tokens and the cost of each call, never what was said.
- **A web page** is fetched by the Portal only when a search cited its address or you typed it, from a public address only, and nothing of it is kept.

Cost and fair use

The assistant costs nothing extra: it comes with every licensed instance whose edition includes it - see [Pricing](#).

- Each organization has a monthly fair-use limit, shared by all its installations and counted against the calendar month (UTC). Once it is reached, the assistant answers *This month's AI allowance is used up.* until the first day of the next month (UTC).
- The chat's header shows what is left this month, after each answer.
- An attachment's chip estimates the tokens of reading it whole.

The usage page. In the Portal, an administrator of the organization chooses **AI assistant usage** on the organization's page: the month's limit, what is used and what is left, and the calls by day and by assistant.

To stay well within the limit, point the assistant at the pages to read, start a new conversation for a new subject - every answer reads the whole conversation again - and attach only the documents the question needs.

Troubleshooting

The chat says	What it means, and what to do
<i>This month's AI allowance is used up.</i>	The organization's allowance is spent until the first day of next month (UTC). The usage page shows where it went
<i>Too many requests. Wait a minute and try again.</i>	The Portal limits how often an installation, and an organization, may ask; wait, then send again
<i>The AI provider is busy. Try again in a moment.</i>	The provider refused the call for now; send again
<i>The AI service is unavailable at the moment.</i>	The Portal cannot serve the assistant just now; try again later
<i>The Portal cannot be reached.</i>	The instance lost its connection to <code>dataorchester.com</code> during the answer
<i>The conversation is too large. Start a new one.</i>	Choose New conversation , and attach only what the next question needs
<i>The answer was cut off before its end.</i>	Ask again
<i>The assistant stopped after using every step a turn allows. Ask it to continue.</i>	A turn has a limit of model calls; ask it to continue
<i>Not applied: what it was built on has changed since the assistant read it.</i>	Ask the assistant to read the records again and rebuild the change
<i>Not self-checked on a card</i>	The assistant proposed without checking its draft. Review it with extra care, or discard it and ask for a checked one
<i>... was not attached.</i>	The message says why: a type that is not supported, a file too large, or a conversation that already holds 10 attachments or its byte limit
No chat beside the editor	Choose Assistant on the editor's toolbar to open it and read the reason: see When the assistant is off

Next steps

- [Licence Activation](#)
- [Configuration Reference](#)
- [Pricing](#)