Loonar Compass

User guide

version 0.6 · October 2026

Loonar Mariusz Głowacki
ul. Gombrowicza 12, 73-110 Stargard
loonar.eu · mariusz@loonar.eu

Contents

1. What Loonar Compass is
2. Requirements
3. Installation
4. Application window and navigation
5. Settings and language
6. Talking to the assistant
7. Working with the Enterprise Architect model
8. Working with the open EA window
9. Diagram editor
10. Sending diagrams to Miro
11. Documents: Word, Excel, images, Confluence
12. Document templates
13. Assistant memory
14. Approving changes and the audit log
15. Conversations: history, search, background work
16. Data, security and privacy
17. Troubleshooting
18. Reporting bugs and new releases
19. Updating and uninstalling
20. License, trademarks and contact

1. What Loonar Compass is

Loonar Compass is a standalone Windows desktop application: an AI assistant for working with an Enterprise Architect model (Sparx Systems). You write to it in plain language, and it:

- **reads the model** (packages, elements, diagrams, connectors, attributes and operations) and describes what is in it,
- **analyzes the model**: the impact of changing an element on the rest of the model, where operations are used, elements without connectors,
- **changes the model**: creates packages, diagrams, elements and connectors, adds attributes and operations to classes, sets keywords and tagged values, lays out elements on diagrams, moves them, removes an element from a diagram, deletes objects from the model, colors them, adds legends, copies packages and creates baselines of them,
- **draws** diagrams from a description, a photo, a sketch or an attached file,
- **shows diagrams** as images in the chat, just as Enterprise Architect draws them,
- **prepares documents** from the model's content: Word, Excel, PNG images and Confluence pages, also based on templates,
- **remembers notes** between conversations (modeling conventions, preferences, corrections).

The application uses the Claude Haiku 5.5 language model (Anthropic) over the Internet. It connects to the EA model in two ways: it reads it directly (from the PostgreSQL database or from the .qea/.qeax file) and writes to it through Enterprise Architect, using the model user account given in the settings when the model has one. With a model file you have open in EA, the changes go through your EA window; otherwise EA runs in the background.

The interface and the assistant's answers are available in four languages: Polish, English, German and French (chapter 5). Answers and changes are prepared by an AI model and may contain errors: review them before you approve them (chapter 14).

2. Requirements

Prerequisites. The application works only when the following are met:

1. **The model is an EA project file .qea or .qeax, or is in a PostgreSQL database.** Older file formats (.eap, .eapx, .feap) and other databases are not supported; Enterprise Architect converts an .eap/.eapx file to .qea with a project transfer (Settings → Model → Transfer → Full Project Transfer).
2. **A model in a database has user security enabled** (in Enterprise Architect: Settings → Security → Enable Security), in "Require User Lock to Edit" mode. In a model file, security is optional: without it, the assistant needs no EA account. The application logs in to the model as a specific user and locks objects while saving. When security is turned off, the locks are skipped; this is not a supported configuration.
3. **A model with security enabled has a user account for the assistant** (Settings → Security → Users), with a password and edit permissions. The account name can be anything, for example `claude` or `assistant`, but the account must exist. You enter its login and password in the application settings.

What	Details
System	Windows 10 or Windows 11 (64-bit)
Enterprise Architect	Installed on the same computer, with a valid license. The application starts it in the background; while it runs, it takes up one license seat.
EA model	An EA project file .qea or .qeax (you need only its path), or a repository in a PostgreSQL database with user security enabled (you need: the server, port, database name, and the database user and password). A .qea file is meant for one person; .qeax for a small team sharing the file.
Model user account	For a database model, or a model file with security enabled: the login and password of a user defined in the EA model. Ideally a separate account just for the assistant, so that its changes and locks are visible separately and do not clash with yours.
Claude API key	A key from the Anthropic Console.
Internet	An HTTPS connection to api.anthropic.com. The application uses the proxy from the Windows Internet settings or from the HTTPS_PROXY and ALL_PROXY variables (chapter 16).
Confluence (optional)	Confluence Cloud or Data Center, only for publishing pages through the assistant: the Confluence address (https://), a token and the space key (chapter 11). Without them, pages are still created and you paste the code yourself.
Miro (optional)	Only for sending diagrams to a Miro board: an access token of a Miro app (chapter 10).
Mail program (optional)	For reporting bugs and subscribing to news about new releases: a mail program set as the default in Windows, for example Outlook.

Costs. The application itself is free. You pay Anthropic for your use of the Claude model, through the account the API key belongs to.

3. Installation

Run the file `LoonarCompass-Setup-0.6.exe`. The wizard guides you through the steps with the **Next** and **Back** buttons. You do not need administrator rights unless you install the application for all users of the computer.

Windows SmartScreen. The installer is not digitally signed yet, so when you run it for the first time, Windows SmartScreen may show the warning “Windows protected your PC”. Click **More info** and then **Run anyway**. If Smart App Control is turned on, Windows blocks unsigned applications altogether; in that case, wait for the signed version.

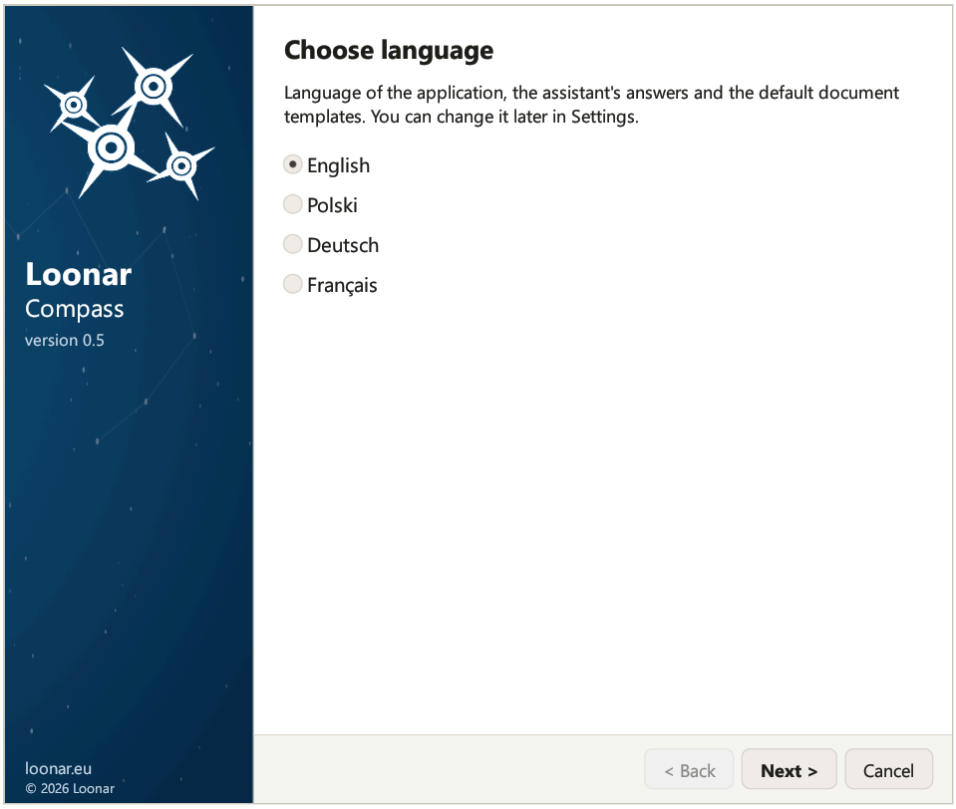

The first page of the wizard: choosing the language

1. **Language.** Polish, English, German or French. For a new installation, English is selected; for an update, the language saved in your settings. The chosen language applies to the wizard, the application, the assistant's answers and the default document templates; you can change it later in the settings (chapter 5).
2. **Welcome.** The prerequisites and a list of things worth preparing before installation.
3. **What's new.** When updating: the changes since the installed version. For a new installation: the changes in the current version.
4. **License agreement.** Select “I accept the terms of the license agreement”.
5. **Destination folder.** The installation type (described below), the program folder (**Browse...**) and the shortcuts: Start menu (application, user guide, uninstall) and Desktop.
6. **Configuration.** Choose “Enter the details now” or “Configure later in the application ("Settings" view)”. The page contains the Claude API key (with a note on what the application sends to Anthropic), the model (a PostgreSQL database: server, port, database name, database user, database password; or an EA file: the path to the .qea/.qeax file, which you can select with the **...** button), Enterprise Architect (the path to `EA.exe`, which you can select with the **...** button, the EA user account and the EA account password) and the approval mode.
7. **Ready to install.** A summary of the installation type, folder, shortcuts and configuration. Click **Install**; the progress bar shows the percentage and the current step.
8. **Installation complete.** With “Launch Loonar Compass” selected, the application opens after you click **Finish**.

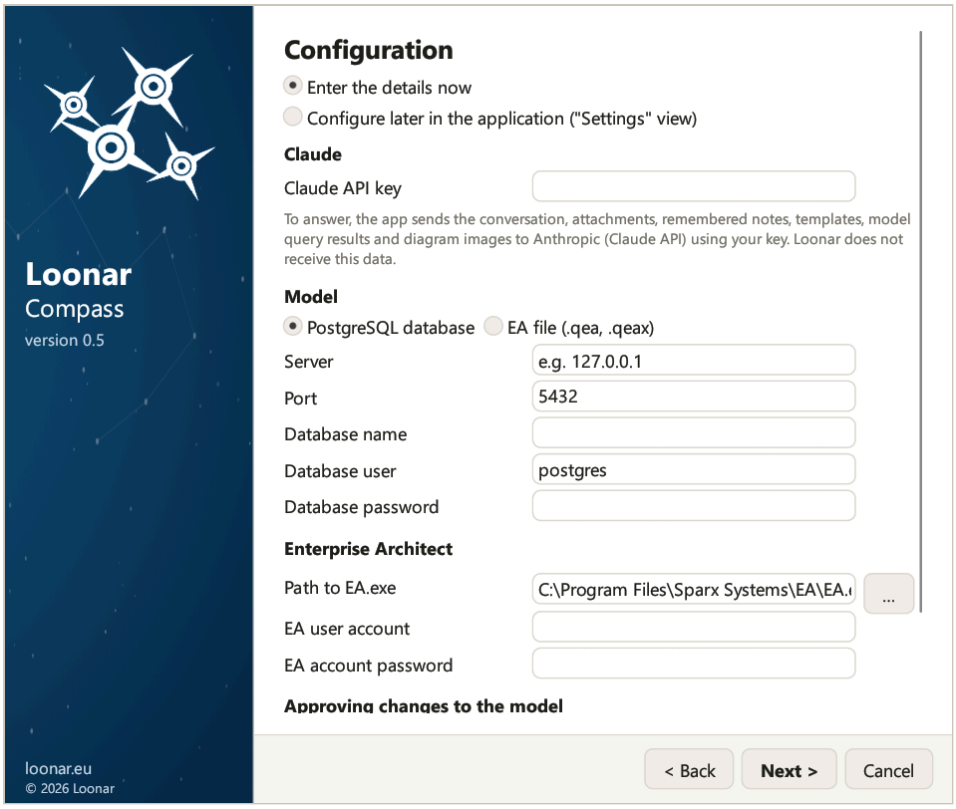

Configuration in the wizard is optional: you can enter the same details later in the “Settings” view

Installation type

On the “Destination folder” page, you choose who the application is installed for:

- **“Just for me (recommended, no administrator rights)”**: the program is installed in `%LOCALAPPDATA%\Programs\Loonar Compass`. Windows does not ask for administrator rights.
- **“For all users of this computer (requires administrator rights)”**: the program is installed in `%ProgramFiles%\Loonar Compass`, where only administrators can change files. Only in this case does Windows ask for administrator rights, after you click **Install**.

Settings, conversations and passwords are always stored for the user who runs the wizard (in `%APPDATA%\Loonar Compass` and in Windows Credential Manager), also when another person confirms the administrator rights and also in an installation for all users.

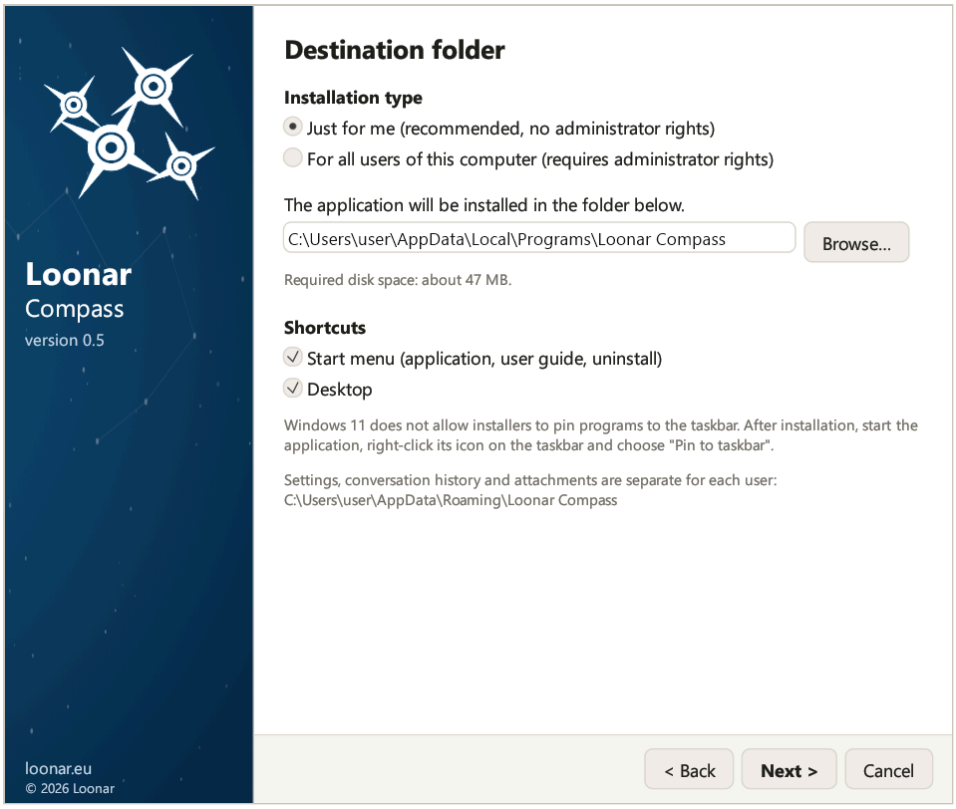

Installation type, program folder and shortcuts

Taskbar. Windows 11 does not allow installers to pin programs to the taskbar. Start the application, right-click its icon on the taskbar and choose “Pin to taskbar”.

Updating from an earlier version

Earlier versions of the application were called Loonar EA Assistant: Loonar Compass takes over their settings, conversations, memory, templates and passwords automatically.

- An earlier version installed just for you is removed during installation (program files, shortcuts and the Windows entry); your data stays.
- If an earlier version is installed for all users, the “Ready to install” page shows it. You can remove it with the **Remove the old version (administrator rights)** button, or later in Windows Settings → Apps.
- On the first page, the wizard preselects the language saved in your settings (settings without a saved language, from versions 0.1 and 0.2, count as Polish), or English when there are no settings. Check this choice before you continue.
- Changing the language in the wizard does not translate the assistant's memory or your templates. To carry them over to a new language, change the language on card 4 in the settings after installation (chapter 5).
- A database password that older versions kept in the settings file is moved to Windows Credential Manager, and files with Polish names in the data folder get English names (for example `audit_log.jsonl`).

4. Application window and navigation

On the left side of the window there is a navy-blue **navigation bar** with icons. When you hover the mouse over an icon, its name appears. Clicking switches the view.

Icon	What it contains
Chat	The list of conversations and the conversation with the assistant (chapters 6 and 15)
Settings	Claude API key, model database, Enterprise Architect, language and approval mode, new releases, Confluence (chapter 5)
Assistant memory	Notes the assistant has remembered between conversations (chapter 13)
Document templates	Document templates that the assistant fills with content from the model (chapter 12)
Model audit log	The assistant's saves to the model, with export to Excel (chapter 14)
Changelog	The list of features and known limitations of each version
Report a bug	A bug report form with images (chapter 18)
Support the project (heart icon, at the bottom of the bar)	Voluntary support for the development of the application, links to GitHub and loonar.eu
About (at the very bottom of the bar)	Version, application owner, links (privacy policy, source code on GitHub, loonar.eu), trademark notice, license text and the licenses of the open-source components (chapter 20)

An orange **"!"** on the Settings, Assistant memory and Document templates icons means that after a language change the translation has not finished yet (chapter 5).

Support the project. Loonar Compass is free, also for commercial use, and it will stay free. If it saves you time, you can support its development with a voluntary contribution: the **Support on suppi.pl** button opens suppi.pl/loonar. The contributions pay, among other things, for a code-signing certificate. Other ways to help: report bugs and ideas (Report a bug), recommend the application to other Enterprise Architect users or star the project on GitHub.

Chat view

- **The list of conversations** on the left is grouped into "Today", "Yesterday", "Last 7 days" and "Earlier". You start a new conversation with the **+** button. The selected conversation is highlighted.
- **The conversation header** shows its title, the **Model: …** button with the model of the conversation (chapter 15) and, while the assistant is working, the "Claude is working" indicator.
- **Messages:** yours are in a beige bubble, the assistant's answers have the application icon. Each message shows the time (older ones also the date).
- **The new conversation screen** shows the logo, a hint about what you can do, and which model the conversation will use, with the **Choose another model…** button.

- **The message card** at the bottom: the text field, the paste button, the model and approval mode, and the round send and stop buttons (chapter 6). Below the card, a reminder says that answers and changes are prepared by an AI model.

5. Settings and language

The **Settings** view has six cards. You save each one separately with the button below the card, and the application immediately checks that the details work. The result appears below the button: green with a description or red with the cause of the error.

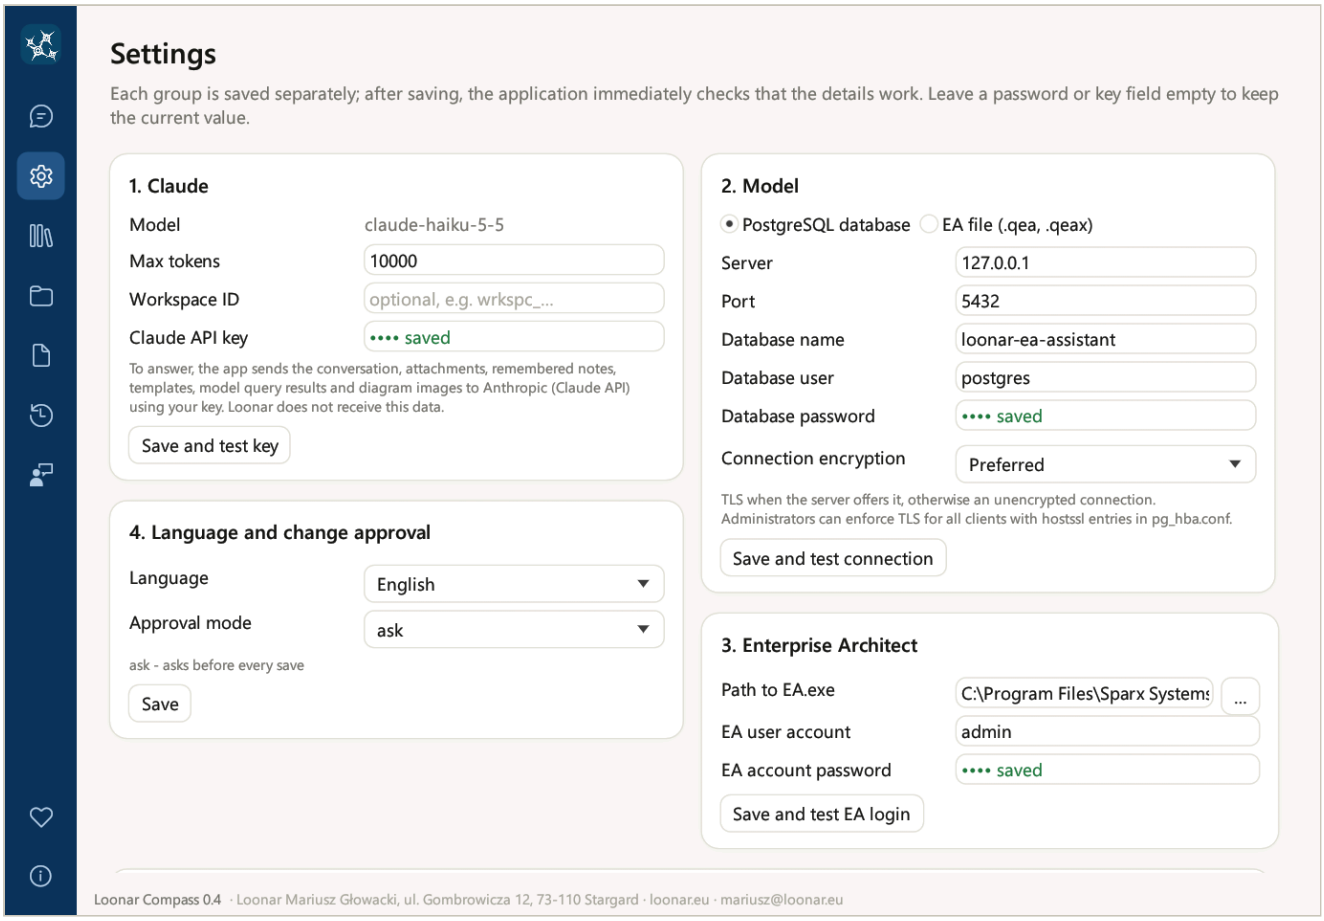

Settings: each card with its own save and test

Card	Fields	What the test checks
1. Claude	Model (Claude Haiku 5.5, read-only), Max tokens (the length limit of a single model answer), Claude API key. A note below the fields says what the application sends to Anthropic (chapter 16).	Sends a short request to the model
2. Model	Choice: **PostgreSQL database** (Server, Port, Database name, Database user, Database password stored in Windows Credential Manager, Connection encryption) or **EA file (.qea, .qeax)** (Model file; the **...** button opens a file selection dialog)	Database: connects, counts packages, elements and diagrams, says whether the connection is encrypted and warns if the database user is a superuser. File: opens it read-only and counts packages, elements and diagrams
3. Enterprise Architect	Path to EA.exe (the **...** button opens a file selection dialog), EA user account, EA account password (for a model file without security, both stay empty)	Opens the model in EA (a model file you have open: in your EA window) and logs in with the given account
4. Language and change approval	Language; Approval mode: ask, plan or auto (chapter 14); after a language change, also the **Translate now** button	Saving only
5. New releases (newsletter)	Email address and consent to news about new releases, with an information note and a link to the privacy policy (chapter 18)	Prepares a message in the mail program
6. Confluence	Confluence address, Email (Cloud only), Token, Space key, Parent page ID (chapter 11). The address must start with https:// (http:// only for a server on the same computer).	Connects to Confluence with the given token and checks the space
7. Miro	Access token of a Miro app with the boards:read and boards:write permissions (chapter 10)	Connects to Miro and counts the boards the token can see

An empty password or key field means "no change". The label **saved** next to a field indicates that the value is already stored. The API key, the passwords and the Confluence token are kept in Windows Credential Manager, never in the settings file. The database password must not contain a semicolon (`;`), because Enterprise Architect cannot pass it in its connection string.

Default values. When there are no settings yet, the server, database name, EA user account and password fields are empty (the Server field shows the gray hint "e.g. 127.0.0.1"), the port is 5432, the database user is `postgres`, the path to EA.exe is `C:\Program Files\Sparx Systems\EA\EA.exe`, the approval mode is ask and Max tokens is 10000. Existing settings are not changed.

Max tokens. If the assistant stops working with a message about a truncated answer, increase this value (for example to 20000). The limit costs nothing by itself: you pay for the text actually generated.

Package locking while saving is always on, which is why it is not in the settings.

Changing the language

The application works in four languages: **Polish, English, German and French**. You choose the language in the installer (English by default) and change it on the **"4. Language and change approval"** card:

1. Choose a language from the **Language** list.
2. Click **Save**. The interface switches to the new language, and the assistant answers in it from the next message.
3. The application carries your content over to the new language: it replaces untouched default templates with the versions in the new language, and **the assistant's memory and your own and modified templates are translated by Claude**.

Only the memory and the templates are translated. **Conversation history is never translated**: conversations stay in the language in which you held them.

> **No connection to Claude.** Translation requires a working API key. If there is no connection, an orange **"!"** appears on the **Settings**, **Assistant memory** and **Document templates** icons, and the tooltip explains that the content has not been translated yet. The translation resumes the next time the application starts, or immediately when you click the **Translate now** button on card 4. Saving other cards does not start it.

Database connection encryption

The **Connection encryption** field on the "2. Model" card (PostgreSQL database) decides whether the assistant reads the model over TLS:

Setting	Effect
Preferred (default)	TLS when the server offers it; otherwise an unencrypted connection. Works with any server.
Required	TLS is required: without it the application does not connect. The server certificate is not verified.
Required, certificate verified	TLS is required, and the server certificate is verified against the Windows certificate store (including company certificates), together with the server name.

- The connection test says whether the connection is encrypted: "The connection is encrypted (TLS)." or "The connection is NOT encrypted: the server does not offer TLS.".
- With certificate verification, the name in the **Server** field must match the name in the server certificate.
- The database administrator can enforce TLS for all clients with `hostssl` entries in the `pg_hba.conf` file.
- The setting applies to the assistant's reading of the model. Enterprise Architect's own connection to the database (when saving) follows EA's settings.

6. Talking to the assistant

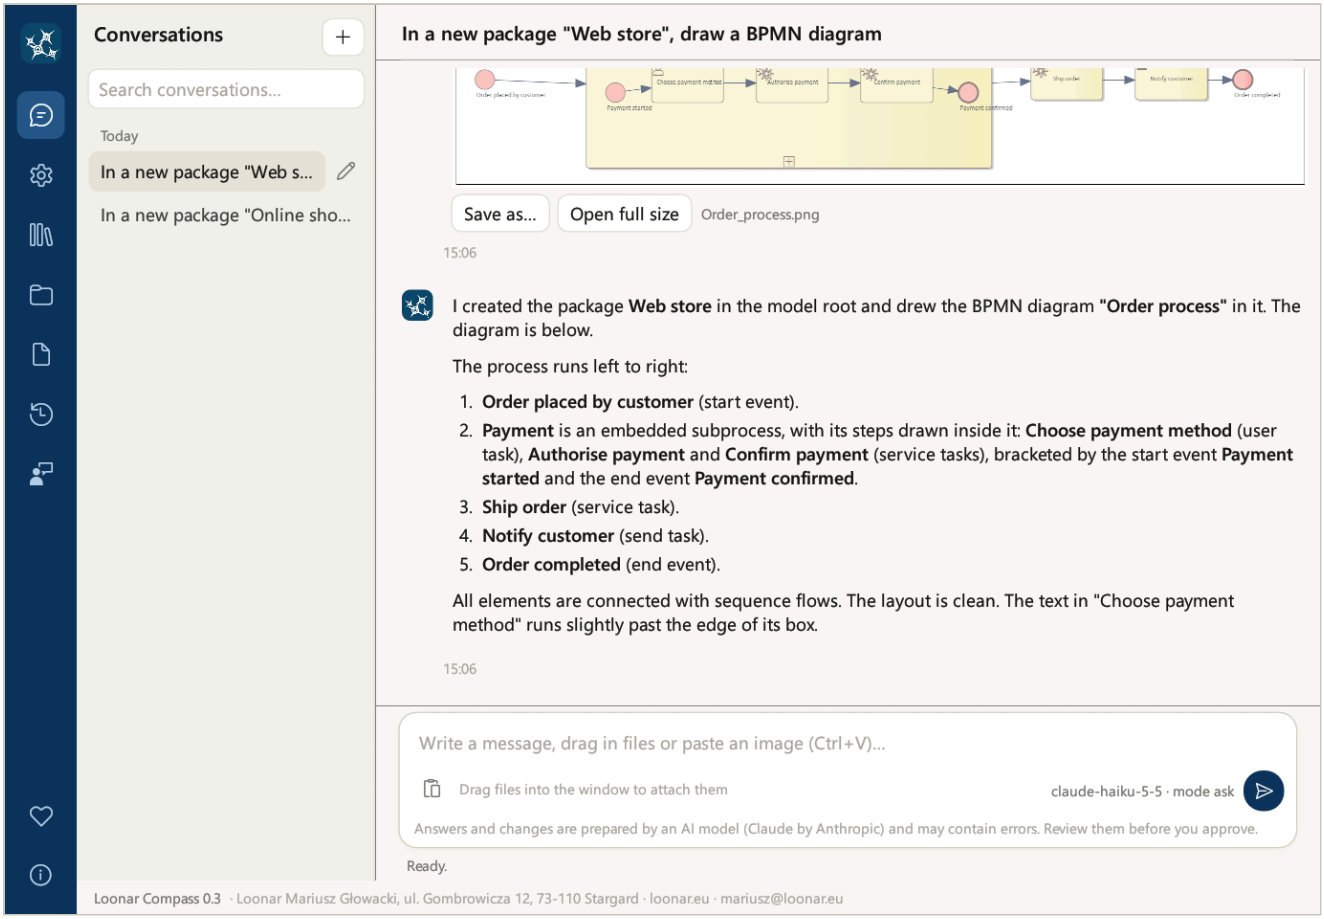

A conversation with a diagram shown as an image in the chat: the list of conversations on the left, the message card at the bottom

Writing

You type a message in the **message card** at the bottom of the window. The card contains the text field, the paste button, the model name and the current approval mode, and round buttons: **Send** and, while the assistant is working, **Stop**. Below the card you see the reminder: “Answers and changes are prepared by an AI model (Claude by Anthropic) and may contain errors. Review them before you approve.”

Key	Action
`Enter`	Sends the message
`Shift` + `Enter`	New line in the message
`↑` in an empty field	Recalls your last message from this conversation; further presses go back to earlier ones, `↓` goes forward again
`Ctrl` + `V`	Pastes text, and when the clipboard contains an image or files, adds them as attachments

The message field grows with the text. Spelling errors are underlined in red using the Windows system dictionary, in the language from the system settings.

Attachments

You can attach up to 5 files to a message, each up to 10 MB. Ways to do it: drag files into the window (from Explorer or from a browser), paste them from the clipboard (`Ctrl` + `V`) or use the paste button in the message card.

File type	What the assistant sees
Images (PNG, JPG)	The image: a photo of a whiteboard, a sketch, a screenshot
Images (GIF, WebP)	The image; in the chat, such an attachment is shown as a file card, not as an image
PDF	The whole document
DOCX	The text of the document
Text files (BPMN, XML, CSV, JSON, code…)	The content of the file; the file must be saved in UTF-8
XLSX, PPTX, DOC and other binary files	Not read: the assistant only learns that the content of the file cannot be passed to it
ZIP	The files inside it, each according to its type: at most 20 files and 30 MB in total; a ZIP inside a ZIP is not unpacked

While the assistant is working

- Below the messages you see the **steps card**: what the assistant is doing right now, how long it takes and a short result of each step (in Polish, German and French the short results are translated, and the technical details stay in English). While the assistant is working, the card is expanded, and when it finishes, it collapses to a single line; clicking the card header expands it again.
- You can **keep writing**: the next message goes into a queue and the assistant takes it into account at the next step, without interrupting its work.
- The **Stop** button stops the work after the current step. A step that is already running (for example saving to EA) is completed.
- An answer can take at most **24 steps**. If the assistant has not finished by then, it stops so as not to run in an endless loop; send a new message (for example “continue”) to carry on.

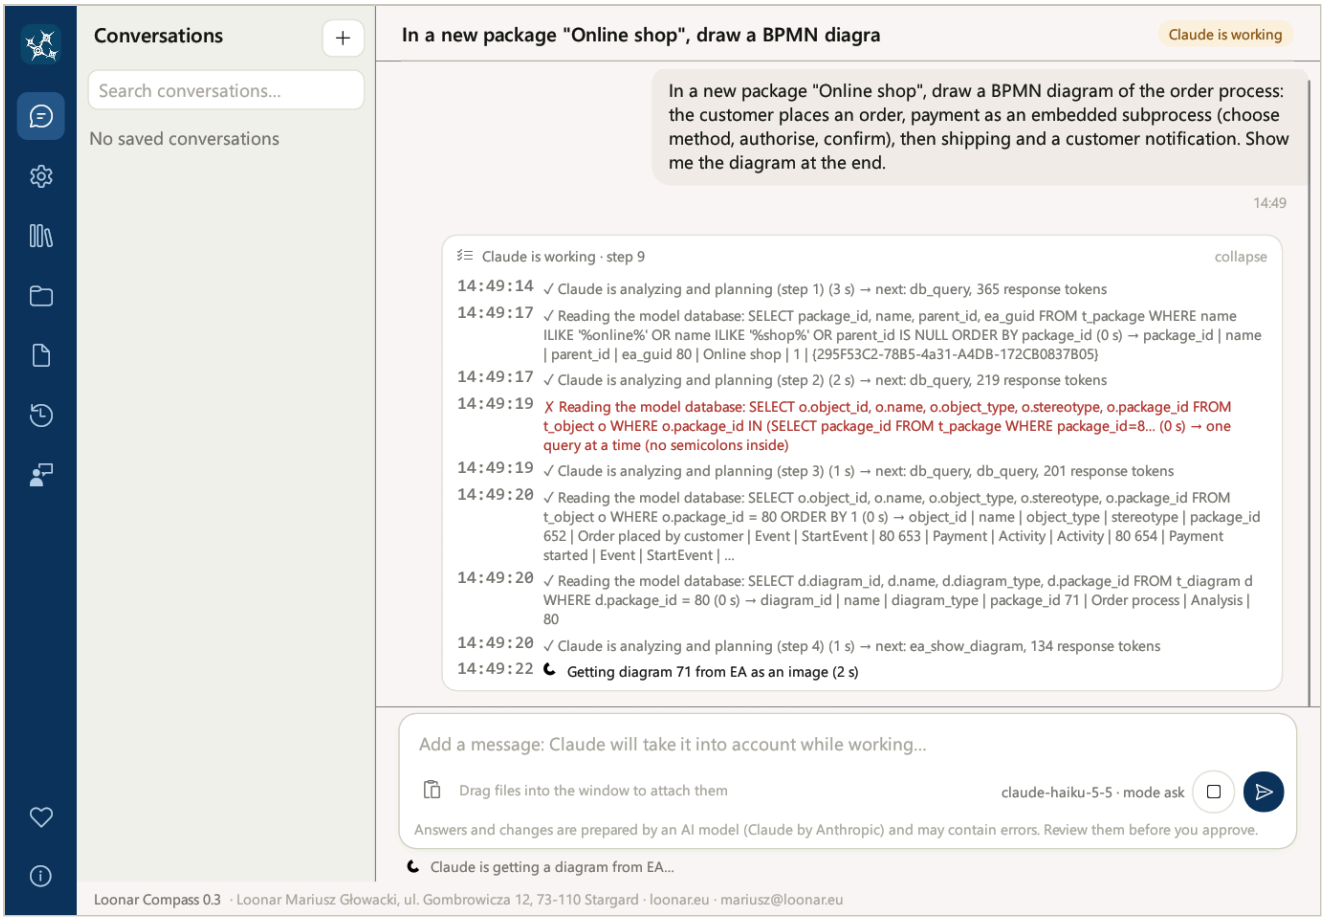

The assistant at work: the list of steps

What can the assistant do? The assistant knows the application's features from the changelog. Ask it: “What can you do in this application?”.

7. Working with the Enterprise Architect model

Write as you would to a colleague who knows EA. Refer to packages and diagrams by name. The assistant refers to model objects by name, package and GUID. Below are examples of requests.

Reading

> *Which packages are in the model and how many diagrams does each of them contain?*

> *Describe what the "Order process" diagram shows and point out elements without connectors.*

> *Show the "Shop use cases" diagram.*

> *List the attributes and operations of the "Order" class together with their parameters.*

The diagram shown appears in the chat as an image. Below it are the **Save as...** and **Open full size** buttons.

Change impact analysis

The assistant goes from the given element along relationships deep into the model (up to 6 levels, in one or both directions, optionally only along selected relationship types). For each affected element it gives the relationship through which the impact arrives, the path from the starting element and the diagrams on which the element appears. You can get the result as a description, an image, a Word document or an Excel workbook.

> *What will change if we change the "Payment handling" component? Check 3 levels in both directions and give me the result in Excel.*

Where operations are used

For a class, component or interface, the assistant finds where its operations (services) are called: on sequence diagrams (messages to the element or its lifeline) and activity diagrams (actions calling the operation), with the package and the caller. It also lists the operations not used on any diagram.

> *Where are the operations of the "Order API" interface used? Which ones are not used anywhere?*

Drawing

> *Create a "Shop" package under the "Model" package and draw a use case diagram in it: the customer browses the catalog, adds a product to the cart, places an order and pays. Add actors and short descriptions of the use cases.*

> *Redraw the process from the attached photo as a BPMN diagram in the "Processes" package.*

> *Extract the requirements from the attached document, create them in the "Requirements" package and show them on a requirements diagram.*

BPMN subprocesses

- **Embedded:** the steps of the subprocess are its child elements and are drawn inside its rectangle on the same diagram.
- **Collapsed:** the steps are on a separate diagram, and double-clicking the subprocess opens that diagram. The assistant can also link a subprocess (or another element) to an existing diagram as a composite diagram.
- The assistant adds new steps to an existing subprocess.

> *In the "Complaint handling" process, replace the "Verification" task with an embedded subprocess with three steps: check documents, check warranty, make a decision.*

Changes to an existing model

> *Move requirements FR01–FR05 to the "Catalog and cart" package.*

> *On the "Order process" diagram, color the tasks performed by the system green and the manual ones yellow.*

> *Remove the cell leading to the "Archive" diagram from the "Start" diagram.*

> *Add a "Start" diagram with link tiles (navigation cells) to all use case diagrams, four per row.*

> *Rename the "Draft" diagram to "Return process" and add a short description to it.*

> *Enlarge the "Warehouse" pool on the "Shipping process" diagram so that all the tasks fit in it.*

The assistant also changes the position and size of an element on a diagram without removing it (connectors stay), removes an element from a diagram while keeping it in the model, and changes the name, description and stereotype of an existing diagram.

Keywords and legends

The assistant sets elements' **keywords** (the Keywords field in EA) and **tagged values**, and creates a **diagram legend with automatic coloring**, as in EA: elements (by keyword, stereotype, status or type) and connectors (for example by stereotype) get a color from the legend, and the legend shows the value, color and label.

> *Mark the requirements with the keyword "MVP" or "later" and add a legend to the diagram that colors them by that keyword.*

Connector lines

- **Line style** (like Line Style in EA): direct, auto routing, custom line, orthogonal square, orthogonal rounded, tree and lateral; for selected connectors or all connectors on the diagram.
- **Connector route:** the assistant routes the line through the given bend points and the points where it meets the element edges, for example following a sketch in a screenshot.
- **Hiding connectors** on a single diagram (and showing them again): the relationship stays in the model and on other diagrams, so you can declutter a dense diagram without deleting relationships.

> *On the "Architecture" diagram, set all connectors to orthogonal and hide the dependencies to the "Login" component.*

Package copy and baseline

- **Package copy** creates a duplicate of the package with all its contents.
- **A baseline** is a saved version of a package in EA. Before deleting or substantially restructuring an existing package, the assistant is instructed to create a baseline, but this is not enforced: for important packages, ask for a baseline explicitly and use the **ask** mode.
- **Restoring** is done in Enterprise Architect: right-click the package, choose **Manage Baselines**, select the baseline and use **Compare/Restore**. The assistant does not restore baselines itself, because restoring through automation makes EA hang.

Good to know

- **The first EA call takes from a few to over ten seconds**, because the application starts Enterprise Architect in the background and opens the model. Subsequent saves and diagram retrievals take less than a second. After 10 minutes without calls, the background EA closes and releases the license. Closing the application also closes its EA session.
- **Locks.** In a model with security enabled, the assistant locks an object under its own account while saving and releases the lock afterwards. It will not change an object locked by another user; the message states who holds the lock on the element (account or group, since when), its package and diagram.
- **Deletion is permanent.** Deleting a package deletes all its contents. With important models, use the **ask** mode (chapter 14), use baselines and back up the database.
- **Partial saves.** When one operation of a save fails, the earlier operations of that save stay in the model. The message says what was saved.
- **Check the result.** The assistant can make mistakes: misread a photo, choose a different element type or lay out a diagram illegibly. Ask it to show the diagram after the changes and to make corrections.
- **Large diagrams** are saved by the assistant in several steps. This is normal.

8. Working with the open EA window

If you have **the same model** open in Enterprise Architect, the assistant works together with your EA window:

- **it knows what you have open and selected**: which diagram is open and which elements are selected on the diagram or in the Project Browser. So you can write “this element” or “the selected tasks”,
- **it opens diagrams** in your EA window and **highlights elements** in the Project Browser,
- **after saving, it refreshes** the changed diagrams and packages in your window by itself, without pressing `F5`.

> *Describe the element I have selected and show which other diagrams it appears on.*

> *Open the “Order process” diagram in my EA and highlight the “Requirements” package in the Project Browser.*

Only when the EA window is not open or shows a different model does the assistant work without it; you will then see the changes after refreshing the package in EA (`F5`). The connection string of your EA window is never sent to the model.

9. Diagram editor

Every EA diagram shown in the chat has an **Edit diagram** button. It opens the diagram in an editor window where you change it with the mouse. The editor works on the model, not on a picture: what you save goes into EA through the same operations the assistant uses, with the approval panel and the change log (chapter 14). The editor needs WebView2, which is part of Windows 10 and 11.

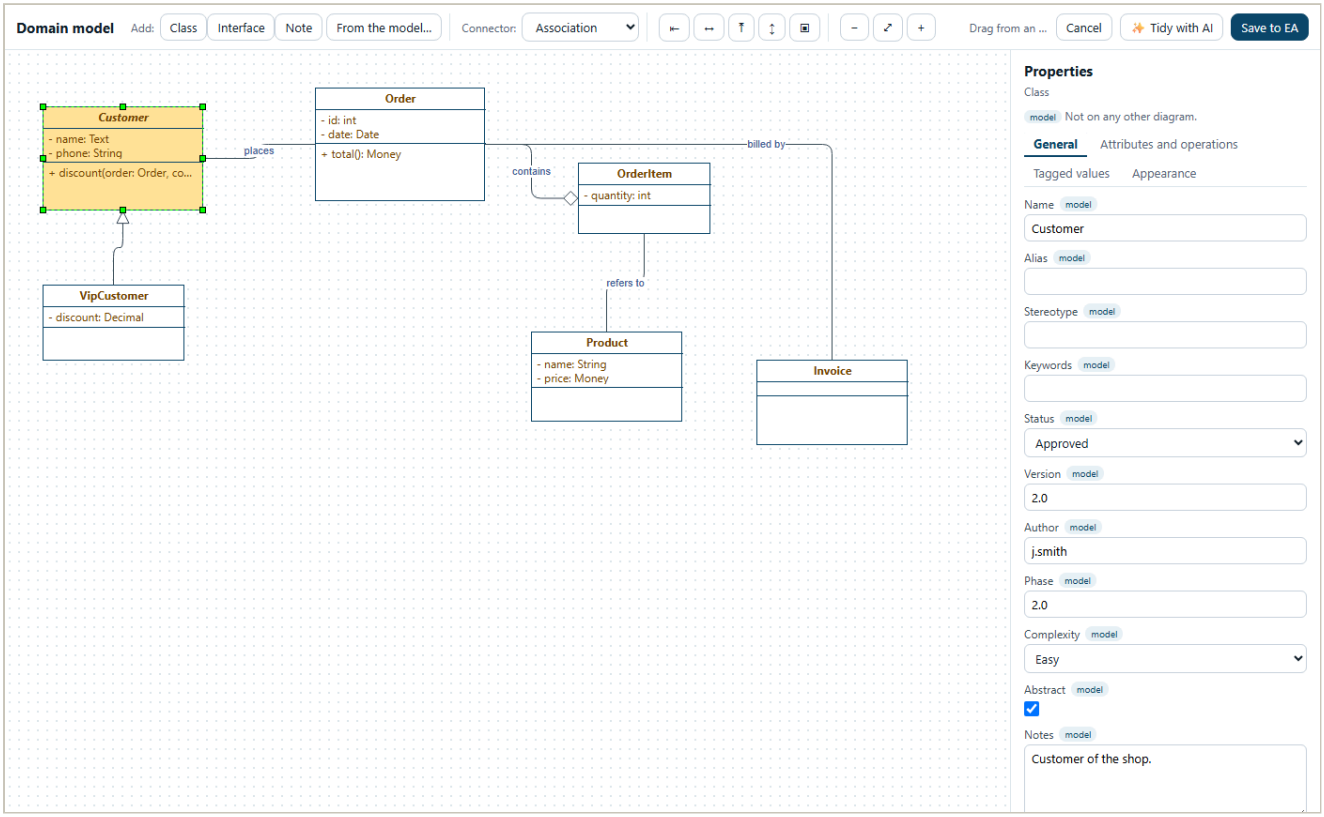

The diagram editor: the toolbar, the diagram and the Properties panel

Drawing

- **Move, resize and align** elements with the mouse and the alignment buttons; `Ctrl` + `Z` and `Ctrl` + `Y` undo and redo.
- **Add**: the buttons depend on the kind of diagram (class, component, use case, activity, BPMN, state machine, deployment, package, sequence). A new element goes to a free spot near the middle of the view; type its name straight away.
- **Connect**: drag from one element to another. The kind of relationship comes from the **Connector** list.
- **From the model...** puts an element that already exists in the model on the diagram (search by name, type or package). Its relationships with the elements of the diagram appear after saving.
- Double-click an element or a relationship to rename it.

Properties

The **Properties** panel shows the selected element or relationship. Each field is marked **model** (the change applies everywhere the object is shown) or **diagram** (only this diagram).

- **General**: name, alias, stereotype, keywords, status, version, author, phase, complexity, abstract and notes; the package, GUID and dates are read-only. For a relationship: name, stereotype, direction, the

multiplicities and roles of both ends, notes.

- **Attributes and operations** (classes, interfaces, components): add, change (visibility, name, type) and delete; a new operation takes its parameters as `name: Type, other: Type`.
- **Tagged values**: change, add and remove; where the profile lists the allowed values (for example BPMN task types), you choose from a list.
- **Appearance**: the fill color on this diagram or the element's default color, and the line style of a relationship on this diagram.

Deleting

The `Delete` key removes what you have added in the editor without asking. For an object that is already in the model, the editor asks: **Remove from this diagram** (it stays in the model and on other diagrams) or **Delete from the model** (permanently, with its relationships, from every diagram). It shows on which other diagrams the element is and how many relationships it has.

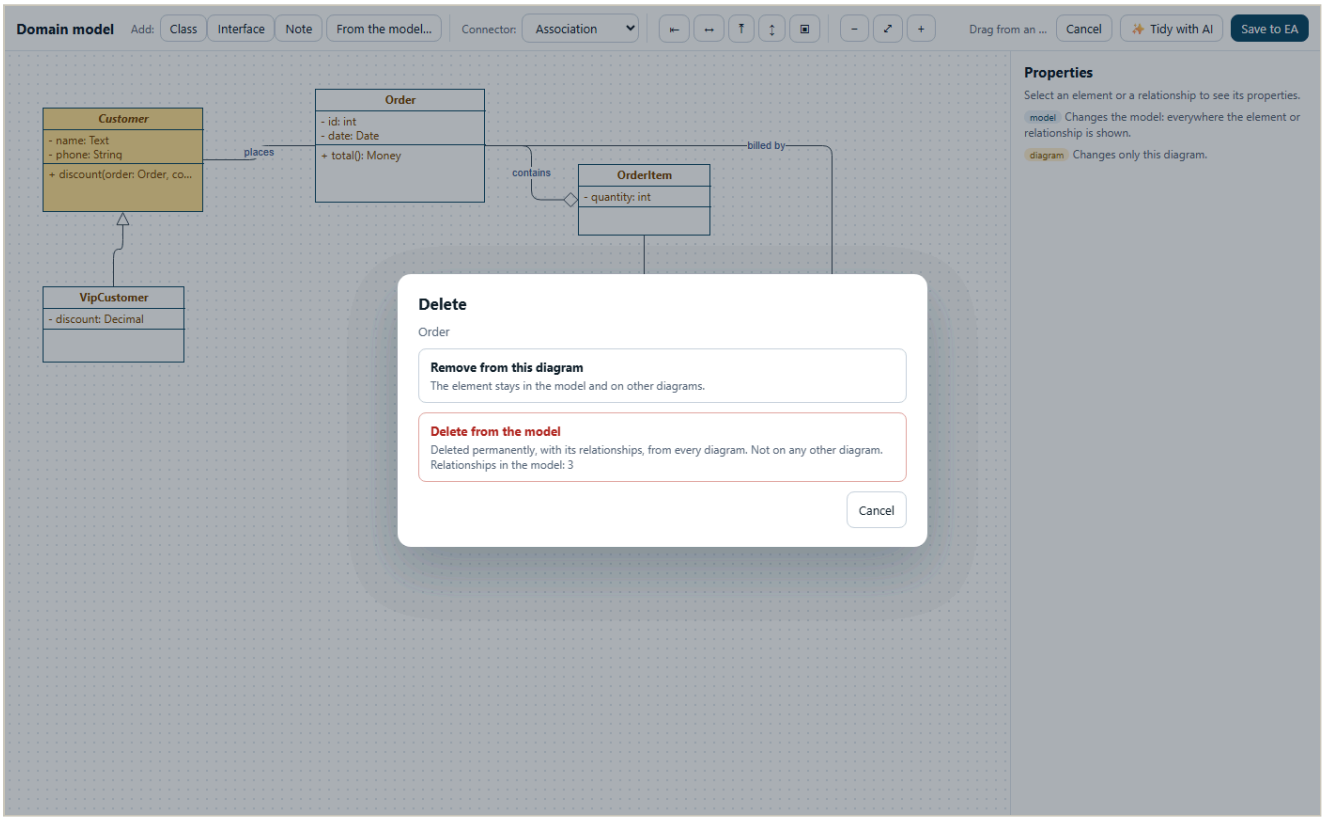

Deleting: from this diagram or from the model, with the other diagrams the element is on

Sequence diagrams

Lifelines have a header and a dashed line; they move only sideways. Messages go from top to bottom in their order, with their numbers, as EA draws them. Change the order of the selected message with ▲ ▼ in the panel or `Alt`+`↑` / `Alt`+`↓`. A new message (drag from one lifeline to another) goes right after the selected one. In the panel you set the kind (synchronous or asynchronous), whether it is a reply, and you can name it after an operation of the receiver.

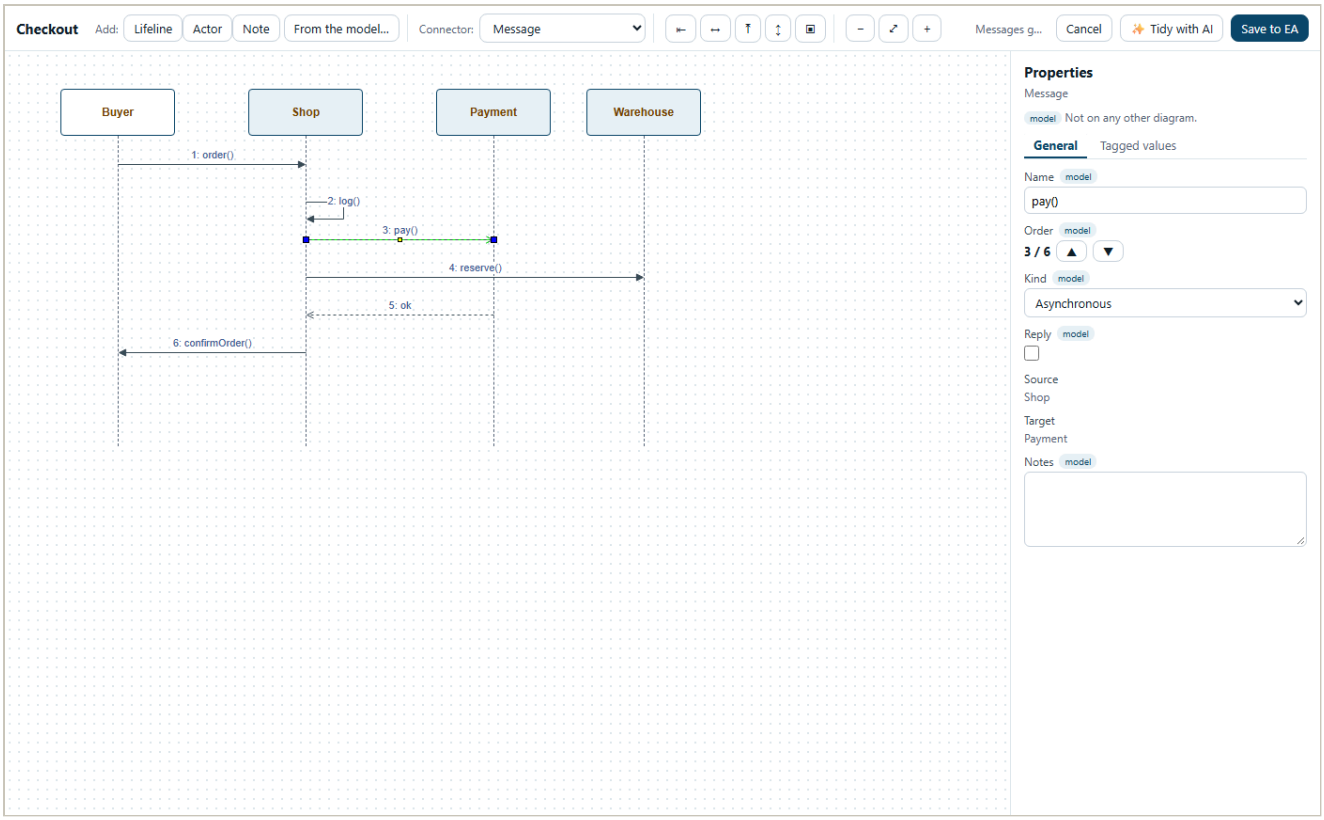

A sequence diagram: numbered messages, the order changed with ▲ ▼

Saving

- **Save to EA** saves the changes and closes the editor; the result appears in the conversation. **Cancel** closes it without saving.
- **Tidy with AI** saves the changes and asks the assistant in the conversation to arrange the diagram so it reads well, without adding, deleting or renaming anything.

10. Sending diagrams to Miro

A diagram from EA can go to a Miro board as editable shapes and connectors in a frame with the diagram's name. Only from EA to Miro: what people change on the board does not come back to the model.

Setup, once:

1. In Miro, open the settings of your profile → **Your apps** → **Create new app**. Miro first asks you to create a **Developer team** (a free sandbox for your apps); accept its terms and create it. In an organization on the Enterprise plan, the organization's admin sets it up.
2. Give the app a name (for example Loonar Compass) and **leave "Expire user authorization token" unchecked**: an expiring token stops working after an hour.
3. Under **Permissions**, select `boards:read` and `boards:write`.
4. Click **Install app and get OAuth token** and choose the organization and team whose boards the application may use (for a first test, the Developer team). Copy the token.
5. Paste it in Settings, card **"7. Miro"**, and click **Save and test Miro connection**. The token is kept in Windows Credential Manager.

Diagrams can go only to boards of the team the app is installed in; to use another organization, install the app there too (the same button).

Sending:

1. Click **Send to Miro** under a diagram in the chat.
2. Choose the board (the one you used last time is offered first; the search field filters by name).
3. The thumbnail shows what is already on the board (grey) and where the diagram will go (a blue frame, red when it would cover something). The application finds a free spot by itself; drag the frame to choose another one.
4. Click **Send to Miro**. When it is done, the conversation gets a message with a link that opens the board at the new frame.

Sequence diagrams go as lifelines with numbered messages in their EA order (replies dashed). Sending the same diagram again puts a new frame next to the earlier one, so nothing changed on the board is overwritten.

11. Documents: Word, Excel, images, Confluence

The assistant prepares documents from the content of the model and the conversation. The finished file appears in the chat as a card with buttons.

Type	What it contains	Buttons in the chat
Word (DOCX)	Headings, paragraphs, lists, tables, diagrams from EA and images drawn earlier in the conversation	**Open**, **Save as...**
Excel (XLSX)	Several sheets, headers with a filter and a frozen first row, numbers as numbers; for example a list of elements, a matrix, the result of an impact analysis	**Open**, **Save as...**
Image (PNG)	A matrix (for example of use cases and requirements), a table, a chart, a schematic. The assistant sees the result and can correct it.	**Save as...**, **Open full size**
Confluence page	The page code for the source editor, and diagrams and images as files	**Copy page code**, **Open images folder**, **Save as...**

> *Prepare a Word document describing the "Shop" package: a list of use cases with descriptions and the use case diagram.*

> *Draw a matrix of use cases and requirements from the "Shop" package as an image.*

Confluence pages

The assistant prepares a page in the Confluence source editor format (storage XHTML): headings, tables, lists, a table of contents, panels (info, note, warning, tip, expand section), statuses, Jira issues, user mentions and code blocks. Diagrams from EA and images are added as attachments. The page code is saved in the attachments folder, and the page images in a folder next to it (`<page>_images`).

Pasting manually (without Confluence settings):

1. Click **Copy page code** on the page card.
2. In Confluence, open the page for editing, open the Source Editor, paste the code and confirm.
3. Click **Open images folder** and add the images to the page as attachments.

Publishing through the assistant. Once you have filled in the **"6. Confluence"** card, the assistant publishes the page via the REST API at your request: it creates a new page under the parent page or updates an existing one (given by ID or title). Images are uploaded as attachments; existing ones get a new version. Diagrams are exported from EA anew on every publication, so republishing refreshes the page after changes in the model. The assistant remembers the pages it has published. The connection uses HTTPS only (http:// is accepted only for a server on the same computer).

Publishing always requires approval, also in **auto** mode. Before the page is sent to Confluence, a panel with the **Approve** and **Reject** buttons appears. It says whether a new page will be created or whether the content of an existing page (with the same title or ID) will be **REPLACED**; the previous version stays in the page history.

Field of the "6. Confluence" card	What to enter
Confluence address	On your organization's own servers (Data Center), for example `https://confluence.example.com`; in the cloud (Cloud) `https://your-site.atlassian.net/wiki`
Email (Cloud only)	For Confluence Cloud: the account's email address. For Data Center, leave it empty.
Token	Data Center: a personal access token (Profile → Personal Access Tokens). Cloud: the account's API token. The token is stored in Windows Credential Manager.
Space key	The space key, for example `ARCH`
Parent page ID	Optional: the ID of the page under which new pages are created

Confluence Cloud: where the details come from.

1. **Token**: it is created on your Atlassian account, not in Confluence. Open `https://id.atlassian.com/manage-profile/security/api-tokens` (Atlassian account → Security → API tokens). Atlassian may first ask you to confirm it is you with a code sent to your email.
2. Click **Create API token** (the plain one, not **Create API token with scopes**), give it a name, for example Loonar Compass, and choose how long it is valid: one week by default, at most one year. Click **Create** and copy the token: Atlassian shows it only once. When it expires, create a new one and paste it in card 6.
3. **Email**: the email address of the Atlassian account the token belongs to.
4. **Address**: your site with `/wiki`, for example `https://your-site.atlassian.net/wiki`.
5. **Space key**: it is in the address of any page of the space, after `/spaces/` (for example `…/wiki/spaces/PM/…` gives `PM`), and under **Space settings**.
6. **Parent page ID** (optional): the number in the address of the page, after `/pages/` (for example `…/pages/65818/…` gives `65818`).

Save the card with the **Save and test Confluence connection** button.

> *Prepare a Confluence page describing the "Complaint handling" process with the diagram and publish it in the space.*

12. Document templates

A template tells the assistant how to build a document: which sections it should have and what to put in them. The assistant fills it with content from the model and prepares a Confluence page or a Word document.

After installation, the following default templates are available:

- **Preliminary feasibility analysis**,
- **BRD - Business requirements**,
- **System change description**,
- **Process description**.

The “Document templates” view

- A list of templates with a short description.
- **Use** inserts into the chat a request to prepare a document based on the template. Complete it (for example with the name of a package or process) and send it.
- **Edit** opens the template file in Notepad.
- **Open templates folder** opens the `templates` folder in Explorer (older installations may still use the `szablony` folder).

Your own templates

Templates are Markdown files with instructions for the assistant. The first line has the form `> Template: description`; the description is shown in the list. You can modify the existing templates, add your own files to the templates folder or ask the assistant to create a template from a sample document. In the ask and plan modes, the assistant asks for your approval before it overwrites a template.

> *Create a template from the attached document: keep the layout of the sections and add what each of them should contain.*

13. Assistant memory

The assistant keeps notes between conversations: modeling conventions, your preferences, corrections and facts about projects. It knows them in every new conversation. A note on the same topic is updated instead of duplicated, and outdated ones are removed.

The assistant saves a note by itself when it learns something worth remembering, and tells you so in one sentence. You can also ask it explicitly, starting with “remember…”:

> *Remember: we number requirements FR01, FR02… and keep them in the “Requirements” package.*

> *Remember that on BPMN diagrams we color system tasks green.*

Notes are shared by all models. When you want a note to apply to one model only, say so, for example “remember this only for this model”: the note is then used only in conversations with that model, and the “Assistant memory” view shows which model it belongs to.

In the **ask** and **plan** modes, saving a note requires your approval, just like a change to the model. Saving a note does not interrupt the assistant's work on the current task. The whole content of the memory is sent to Claude with every request (chapter 16).

The “Assistant memory” view

- A list of notes with dates.
- **Remove** next to a note deletes it.
- **Open memory folder** opens the `memory` folder in Explorer. The notes are Markdown files in it, and the index is in the `MEMORY.md` file.

14. Approving changes and the audit log

You set the mode in the **Settings** view, on card 4. It decides whether and when the application asks for approval before saving to EA. Reading the model and showing diagrams never require approval. The current mode is also visible in the message card.

Mode	When it asks for approval	For whom
ask	Before every save, including saving a note to the memory and overwriting a template	Working on an important model; full control
plan	Once, before the first save in a given assistant answer. Further saves in that answer go ahead without asking, except deleting objects, moving elements and diagrams, and copying packages: these always ask again.	Larger tasks, when the first save shows the direction
auto	Never (exception: publishing to Confluence)	A test model or work that is easy to undo

When approval is needed, a **"Claude wants to save changes to the model"** panel appears above the message field, with a description of the changes (how many packages, diagrams, elements and connectors, and the names of the objects being deleted, moved or copied) and the **Approve** and **Reject** buttons. A rejected save does not change the model; the assistant is informed of the rejection and may suggest something else.

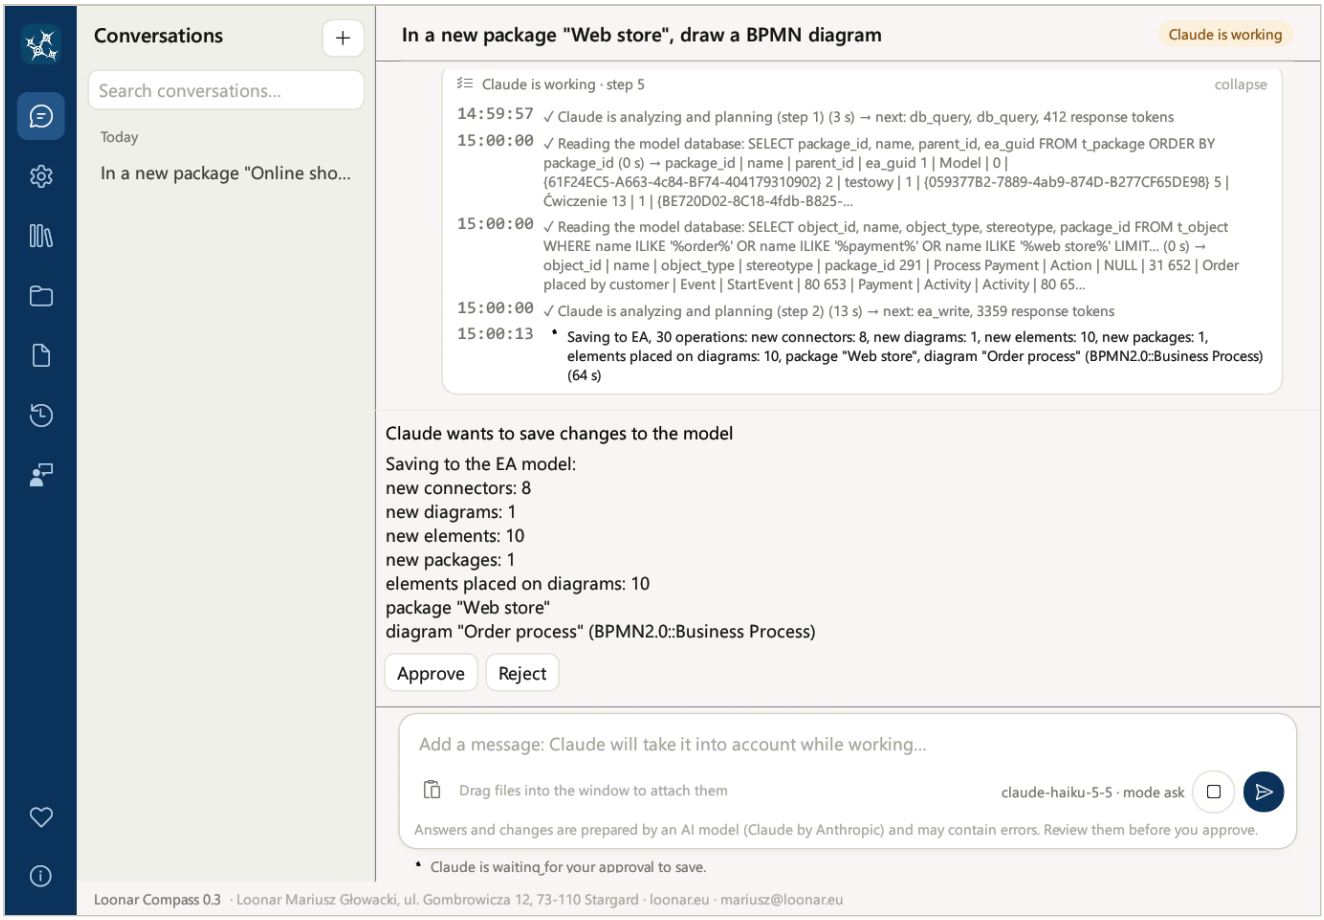

The approval panel for a change in the model

Publishing to Confluence always asks for approval, in every mode, and says when an existing page will be replaced (chapter 11).

Audit log

Every save by the assistant to the model, including failed ones, goes into the audit log (saves you rejected are not recorded, because nothing was saved). The **Model audit log** view shows the latest 200 entries, newest first: time, result, EA account, database and changes. The **Export to Excel** button saves the whole log as an XLSX workbook, for example for reviewing changes as a team, and **Open log file** opens the `audit_log.jsonl` file.

15. Conversations: history, search, background work

- **New conversation:** the **+** button above the list of conversations. Every conversation is saved and comes back after the application is restarted.
- **Groups:** conversations are grouped into “Today”, “Yesterday”, “Last 7 days” and “Earlier”.
- **Search:** the “Search conversations…” field filters the list by title and by message content.
- **Renaming:** click the pencil next to the selected conversation, type a name; `Enter` saves, `Esc` cancels.
- **Model of the conversation:** every conversation works with its own model. A new conversation uses the default model from the settings (the new conversation screen says which one); **Choose another model…** or the **Model: …** button in the header opens the model of the conversation. There you choose **Default from the settings** or **Another model**: a PostgreSQL database (server, port, database

name, user, password, connection encryption) or an EA file (.qea, .qeax), and the EA account. **Test the connection** opens the model and counts packages, elements and diagrams; **Save** sets it for the conversation. Passwords of a connection that differs from the settings are kept in Windows Credential Manager. Changing the model of a conversation that already has messages applies to the next messages; the earlier answers refer to the previous model.

- **EA file dropped on the window:** a .qea or .qeax file dropped on the chat becomes the model of the conversation instead of an attachment; in a conversation with messages, a new conversation starts with that file.
- **Conversations from earlier versions** get a copy of the model from the settings when version 0.5 or later starts for the first time, so they keep working with that model even if the settings change later. Hovering over a conversation in the list shows its model.
- **Several conversations at once:** while the assistant is working, you can start a new conversation or switch to another one. The work already started continues in the background.
- **Deleting:** conversations cannot be deleted in the application. If you need to, close the application and delete the conversation's file in the `history` folder (chapter 16).

Dot next to a conversation	Meaning
● **orange**	The assistant is working in this conversation (or waiting for your approval to save)
● **green**	The assistant has finished working in the background, and you have not looked at the result yet

Saves to EA from different conversations are carried out one after another, not simultaneously. EA in the background keeps one model open at a time: switching between conversations with different models reopens it, which takes from a few to about 25 seconds. When the chat is scrolled up, the circle with an arrow above the message field takes you to the end of the conversation.

16. Data, security and privacy

Where your data is

What	Where
Program	Just for you: `%LOCALAPPDATA%\Programs\Loonar Compass`; for all users: `%ProgramFiles%\Loonar Compass`. The folder also contains the license, the changelog and `THIRD-PARTY-NOTICES.txt`; the user guide is on loonar.eu (Start menu → User guide).
Settings (`app.properties`, without passwords or keys)	The user data folder `%APPDATA%\Loonar Compass`. Each user of the computer has their own.
Conversation history (`history` folder)	
Attachments and the files prepared by the assistant: Word, Excel, PNG and Confluence pages with their image folders `<page>_images` (`attachments` folder)	
Assistant memory (`memory` folder with the `MEMORY.md` index)	
Document templates (`templates` folder)	
Audit log (`audit_log.jsonl`) and the list of published Confluence pages (`confluence_pages.json`)	
Claude API key, EA account password, database password, Confluence token	Windows Credential Manager, entries “loonar-compass” (encrypted by the system). Nothing secret is stored in `app.properties`.

Conversation history is stored as plain files in your Windows profile, protected by your Windows account; the application does not encrypt it.

What is sent to Anthropic

So that the assistant can answer, every request to the Claude service (Anthropic) contains: the content of the conversation, the attached files, the full content of the assistant's memory, the list of templates with their descriptions, and what the assistant has read from the model in that conversation, i.e. the results of its database queries (names and descriptions of packages, elements, diagrams) and images of the diagrams shown. The whole database is not sent, only what the assistant asked for. After a language change, the content of the memory and templates to be translated is also sent to Claude. The data goes to Anthropic with your API key; Loonar does not receive it.

Before you use the application on a company model, check whether your organization's rules and your agreements with customers allow the content of the model to be passed to an external AI service, and which data processing terms apply to your Anthropic account.

The assistant's permissions

- The assistant reads the database in a read-only session, and a model file read-only, with a time limit for each query; server functions that read files or affect other sessions are blocked. All changes go through Enterprise Architect, under the account given in the settings, so the permissions and locks of that account apply to them.
- In Confluence, the assistant works with the permissions of the token from card 6 and publishes only after your approval.
- The assistant has no access to files on disk other than those you attach to it yourself and its own memory and templates folders. Images drawn by the assistant cannot read files from disk or from the network.
- The assistant treats the content of attachments, documents and the model as data, not as instructions.
- The connection string of your open EA window, with its login details, is never sent to the model.

Connections and encryption

- The Claude API and Confluence are used over HTTPS only.
- The connection to the model database supports TLS (card 2, see “Database connection encryption” in chapter 5). Tip: administrators can enforce TLS for all clients with `hostssl` entries in `pg_hba.conf`. Enterprise Architect's own connection to the database follows EA's settings.
- The application uses the proxy from the Windows Internet settings and the HTTPS_PROXY and ALL_PROXY variables. Proxies with a PAC script and NTLM/Kerberos proxy authentication are not supported.
- Temporary Claude API errors (overload 529 or other 5xx errors, rate limit 429) are retried automatically.

17. Troubleshooting

Symptom	What to do
The application opens on the “Settings” view	Settings or secrets are missing. Fill in the cards and save each one with its “Save and test…” button.
“The Claude API rejected the API key (401)”	The key is wrong or has been revoked. Paste the current key from the Anthropic Console on card 1.
“The Claude API returned … with a response that is not from the API”	A proxy or firewall is probably blocking api.anthropic.com. Ask your IT department to allow the connection. The application uses the proxy from the Windows Internet settings or the HTTPS_PROXY / ALL_PROXY variables; PAC scripts and NTLM/Kerberos proxy authentication are not supported.
“The Claude API rate limit was reached (429)” or “The Claude API is overloaded or temporarily unavailable”	The application has already retried the request several times. Wait a moment and send the message again.
File test: “this does not look like an Enterprise Architect model file”	Choose a .qea or .qeax project file created by Enterprise Architect 16 or later. Older .eap/.eapx files must first be converted to .qea in EA (project transfer).
Database test: no connection	Check the server, port, database name, user and password, and whether the PostgreSQL server accepts connections from this computer. The database password must not contain a semicolon (;).
“The connection is NOT encrypted: the server does not offer TLS”, or the connection fails with “Required” or “Required, certificate verified”	The server does not offer TLS: ask the database administrator to enable it. With certificate verification, the server certificate must be trusted in the Windows certificate store (for example through your company's certificate authority) and the server name must match the certificate.
The database test warns that the account is a superuser	Superuser rights are not needed. Ask the administrator for a normal account with read and write rights on the model tables, without superuser rights: the assistant only reads the database, but Enterprise Architect saves changes with the same account. Enter it on card 2.
“EA did not respond within 300 s”	Most often a wrong model user account or password (EA then waits at an invisible login window) or no free license. Correct the details on card 3 and use “Save and test EA login”.
The first save or the first diagram takes a long time	This is Enterprise Architect starting in the background (from a few to over ten seconds). Subsequent calls are fast.
“The model's answer was cut off at the limit of … tokens”	Increase “Max tokens” on card 1 or ask for the task to be done in smaller parts.

“The assistant took 24 steps without finishing”	The answer reached the step limit. Send a new message, for example “continue”, and the assistant carries on from where it stopped.
Changes are not visible in Enterprise Architect	If the same model is open in EA, the assistant refreshes it by itself. Otherwise, refresh the package in the Project Browser (`F5`) or open the diagram again.
“Cannot lock the object”	The object is locked by another model user; the message says who. Release the lock in EA or ask its owner to do so.
Orange “!” on the Settings, Assistant memory and Document templates icons	After a language change, the translation is waiting for a connection to Claude. Check the API key on card 1 and click “Translate now” on card 4, or restart the application.
Publishing to Confluence fails	Check card 6 with the “Save and test Confluence connection” button: the address (it must start with https://), token, space key, and for Cloud also the email address. You can also paste the page manually with the “Copy page code” button.
After a baseline you need to go back to an earlier version of the package	In EA: right-click the package → Manage Baselines → Compare/Restore.
The assistant claims it cannot do something	Ask: “What can you do in this application?”. If the feature is on the list in the “Changelog” view, tell it so explicitly and ask it to try again.
Files cannot be dragged into the window	The application was started as administrator. Close it and start it with the normal shortcut.
“The program hit an unexpected error and has to close”	Your conversations and settings are saved. The details are in the file `%TEMP%\loonar-compass-crash.log` : attach it to a bug report (chapter 18).
“The settings file … contains an error”	The message gives the faulty value. Correct it on the relevant card and save the card again.
Windows SmartScreen warns when the installer starts	The installer is not digitally signed yet. Click “More info” and then “Run anyway”. With Smart App Control turned on, Windows blocks unsigned applications; wait for the signed version.

18. Reporting bugs and new releases

Report a bug

1. Click the **Report a bug** icon in the navigation bar.
2. Enter a subject and a description: what you were doing, what you expected, what happened. The description is required.
3. Add images, for example screenshots: drag them into the window, paste them from the clipboard (`Ctrl` + `V`) or choose them from disk. The added images are shown as thumbnails; together they can take up to 15 MB. Remove personal and confidential data from the screenshots before sending.
4. Optionally, enter the email address to which I should reply. The option “Include technical information (application version, system, model, approval mode, last message; no passwords or keys)” is selected by default; clear it if you do not want to include this information.
5. Click **Send report**. A ready message to the author, with the images as attachments, opens in the Windows mail program. Check it and send it.

The report goes to Loonar Mariusz Głowacki and is used only to answer it and to improve the application; the details are in the privacy policy, linked from the form. The temporary copies of the images are deleted as soon as the mail program has taken them. If the mail program does not accept attachments from other applications, a message without images opens together with the folder containing the images: drag them into the message. Only in this case do the copies stay in that folder.

New releases (newsletter)

On the **“5. New releases (newsletter)”** card in the settings, enter your email address and select “I want to receive news about new releases”: a short email when a new version is released, with a description of the changes and where to download it, without ads. Your choice is stored in the settings, and subscribing and unsubscribing are sent to the author as a message from the mail program. Your address is used only for this news, based on your consent, until you withdraw it: clear the option and save, or write an email. The card links to the privacy policy.

19. Updating and uninstalling

Updating. Close the application and run the installer of the newer version. The wizard shows the “What's new” page with the changes since the installed version, suggests the same installation type and folder, and keeps your settings, conversation history, attachments, memory, templates and audit log. Updating from the earlier version under its former name is described in chapter 3.

Closing the application. Closing the window also closes the Enterprise Architect session started in the background by the application, so the EA license is released immediately.

Uninstalling. Close the application. Use the “Uninstall Loonar Compass” shortcut in the Start menu or the “Loonar Compass” entry in Windows Settings (Apps → Installed apps). The uninstaller removes the program files, the shortcuts and the system entry, showing the progress on a bar. For an installation for all users, Windows asks for administrator rights. Your data stays unless you select:

- “Remove settings, conversations, attachments, memory, templates and the audit log”: deletes the whole user data folder `%APPDATA%\Loonar Compass`,

- “Remove the API key, passwords and tokens from Credential Manager”: removes all the secrets, including the Confluence token and the database password.

In an installation for all users, only your own data is removed; the data of other users stays in their profiles.

20. License, trademarks and contact

License. Loonar Compass is freeware: you may use it free of charge on any number of computers, also for professional and commercial purposes, without time or territorial limits, and pass on unmodified copies free of charge. Modifying the application is prohibited. The license also sets out rules on AI-generated output (you review it before you approve it or rely on it), on data and on open-source components. The Polish version of the license is binding; the English, German and French versions are translations. The full text is in the **About** view and in the installation folder: `LICENSE.txt` in the installation language and, for languages other than Polish, `LICENSE.pl.txt` with the binding Polish version.

Privacy. The application collects no data (no telemetry). How bug reports and newsletter sign-ups are processed is described in the privacy policy; the owner's details are in the legal notice.

Open-source components. The application includes open-source libraries and fonts; each keeps its own license. The list and the license texts are in the `THIRD-PARTY-NOTICES.txt` file in the installation folder; in the **About** view, click **Show licenses of the components**.

Support. The application is free and will stay free. If you want to support its development, use the **Support the project** view or go to suppi.pl/loonar. The source code is at github.com/loonar-eu/loonar-compass.

Application owner: Loonar Mariusz Głowacki, ul. Gombrowicza 12, 73-110 Stargard.
Contact: mariusz@loonar.eu, loonar.eu. The most convenient way to report bugs is the **Report a bug** view.

Loonar Compass 0.6. Sparx Systems and Enterprise Architect are registered trademarks of Sparx Systems Pty Ltd. Claude and Anthropic are trademarks of Anthropic, PBC. Atlassian, Confluence and Jira are trademarks of Atlassian. Microsoft, Windows, Word and Excel are trademarks of the Microsoft group of companies. PostgreSQL is a registered trademark of the PostgreSQL Community Association. Loonar Compass is an independent product and is not affiliated with, endorsed or sponsored by these companies.