COLIBO®

Whitepaper · October 2026 · Version 2.5

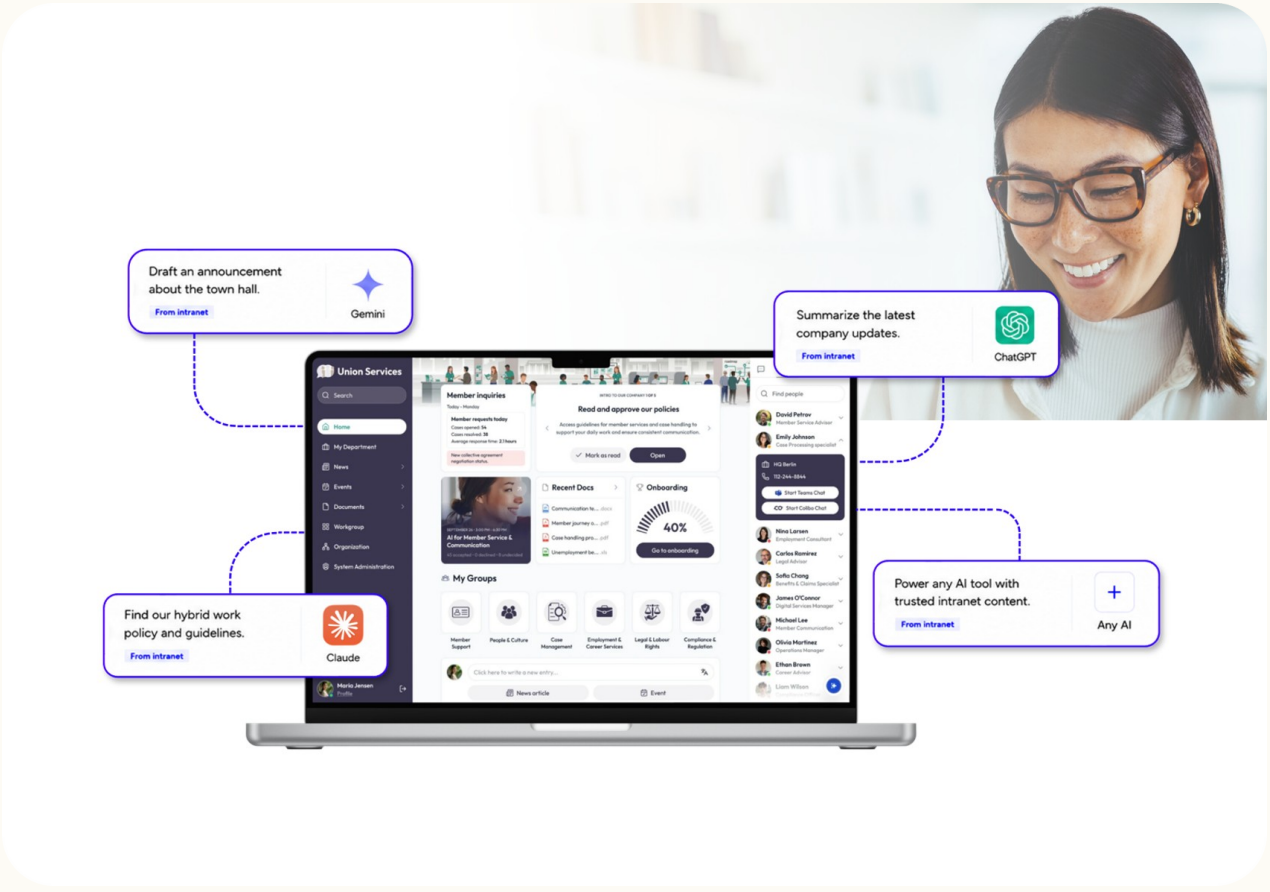
Colibo AI Enabler

Get answers from your intranet, inside Colibo and in your own AI tools

Permission-governed access to intranet content

Flexible choice of AI platform

Cloud or on-premises



About this white paper

Colibo AI Enabler connects your organisation's AI tools with the content in your Colibo intranet. The connection works both ways: employees can draw on intranet knowledge through the AI tool they already use, or get AI assistance directly inside Colibo.

The first five chapters are written mainly for communications, HR and intranet managers. Chapters 6 to 9 cover the choice of AI platform, technology, security and costs, and are written primarily for IT and security managers. Chapter 10 covers getting started and pricing.

Terms used in this white paper

- **Colibo AI Enabler:** The complete product, comprising Colibo MCP Server, Colibo Embedded AI and the Colibo Intelligence Layer.
- **Colibo MCP Server:** Gives external AI tools access to the parts of the Colibo intranet that each employee is permitted to see.
- **Colibo Embedded AI:** AI assistance directly inside Colibo. Employees can ask questions about intranet content and use relevant actions. The function uses your organisation's own AI API account. AI Enabler Administration comes with it.
- **Colibo Intelligence Layer:** Instructions and default settings developed for using AI together with intranet content.
- **Assistant name:** The name shown in the Colibo interface. It can be changed in Colibo Admin and is “AI assistant” by default.

Contents

01 When AI cannot reach the intranet	**02** What is Colibo AI Enabler?	**03** Colibo Embedded AI	**04** Colibo Intelligence Layer	**05** Use Colibo through your existing AI tools
06 Choice of AI platform	**07** Technical architecture	**08** Security and data handling	**09** AI usage and costs	**10** Getting started

01
When AI cannot reach the intranet

Most organisations already hold much of the information their employees look for on the intranet. That includes HR policies, guides, onboarding material, internal news and descriptions of working procedures. The content is often carefully written and kept up to date.

But employees' AI tools do not necessarily have access to it. An employee might ask an AI assistant about the rules for working from home, parental leave or submitting expenses. If the assistant cannot reach the organisation's own documents, the employee risks getting a general answer that does not match the internal rules.

This is an integration challenge. The organisation's knowledge already exists, but it does not automatically become available in the AI tools employees use.

What it means for communications, HR and the intranet

When employees cannot find a reliable answer, they often end up contacting HR or IT, asking a colleague or trying their luck. That takes time and can lead to mistakes or inconsistent answers.

For communications and intranet managers it also means that relevant content is used less than it could be. Good search and clear navigation remain important, but they do not solve the problem on their own if the employee starts the task in an AI tool that has no access to the intranet.

What it means for IT and information security

Many organisations have already adopted AI tools through Microsoft, another platform vendor or separate subscriptions. The tools can help with writing, analysing and summarising, but without access to internal sources they are of limited value for questions about the organisation's own affairs.

Colibo AI Enabler gives existing AI tools access to Colibo content through a structured connection in which the user's permissions continue to apply. The solution can be used without changing AI provider or rebuilding the intranet. How data is handled depends on the component chosen and is covered in chapter 8.

02
What is Colibo AI Enabler?

Colibo AI Enabler connects the Colibo intranet with the organisation's AI tools. The solution builds, among other things, on the Model Context Protocol (MCP), an open standard for how AI tools can communicate with other systems and data sources.

The aim is to make the organisation's existing intranet content usable in AI-assisted work. Colibo provides the connection to the intranet, while the organisation chooses which AI solutions to use.

Bring your own AI

Some digital workplace platforms come with a built-in AI solution, where platform and AI service are tightly bound together. With Colibo AI Enabler, the organisation can instead use an AI platform it has already approved, for example Claude or Microsoft Copilot.

Colibo provides the access to intranet content. The agreement for the AI service itself stays with the chosen provider. This makes it possible to replace or supplement AI tools without changing the foundation of the Colibo integration. The specific responsibility for data and data processing is described in chapter 8.

A short note on MCP

MCP stands for Model Context Protocol. The standard allows an AI tool to request data from a connected system and use it to answer the user's question. In Colibo AI Enabler, it is the MCP server that handles the connection to Colibo and returns the relevant information.

MCP is an open standard and not a Colibo-specific technology. AI platforms that support MCP can therefore in principle be integrated with Colibo MCP Server. See chapter 6 for the supported platforms and setup requirements.

Two components, one product

Colibo MCP Server: Colibo in your AI. **Colibo Embedded AI:** your AI in Colibo.

Colibo MCP Server makes it possible to reach Colibo content through external AI tools. The organisation can, for example, let one team use Claude and another Microsoft Copilot. The AI tools work as they always have, but can retrieve information from Colibo once the integration is set up.

Colibo Embedded AI gives employees an AI assistant directly in the intranet. Colibo sends relevant information to the AI provider the organisation has chosen, and shows the answer together with source references. Embedded AI requires a separate AI API account that belongs to the organisation. AI Enabler Administration comes with Embedded AI, an administration site for content quality, feedback and usage.

The two components can be used separately or together. Each works without the other, and either can be activated later without changing the setup of the other.

Examples of use

With access to Colibo, an AI assistant can, among other things, help an employee find and understand an internal policy, summarise the latest intranet news, get an overview of a team page or draft a text based on the organisation's existing content. The assistant can refer to the pages and documents on which the answer is based.

03
Colibo Embedded AI

Primarily for communications and intranet managers

The intranet holds your organisation's answers, but reaching them still takes a search, a few clicks, and a careful read. Embedded AI shortens that path for whoever opens the intranet, with the same permission-governed access described throughout this document.

Colibo Embedded AI lets employees ask questions about intranet content directly from Colibo. They do not first have to find the right document or open a separate AI tool. The function uses the organisation's own AI API account, and Colibo's usual access permissions apply here too.

What the employee sees

The assistant appears in a panel that can be opened anywhere in the intranet. The name can be adapted in Colibo Admin and is “AI assistant” by default. The employee can type their own question or choose one of the available shortcuts.

There is no need to install anything or create a personal account with the AI provider in order to use the function in Colibo.

Questions and source references

The employee can, for example, ask how a particular procedure works, or which rules apply in a specific situation. Colibo Embedded AI finds relevant material among the pages and documents the employee has access to, and presents an answer with source references. The sources are shown with document type and last-updated date, so that the employee can check the information.

The assistant can also take the page the employee is on as its starting point. Follow-up questions can be asked along the way without starting a new conversation.

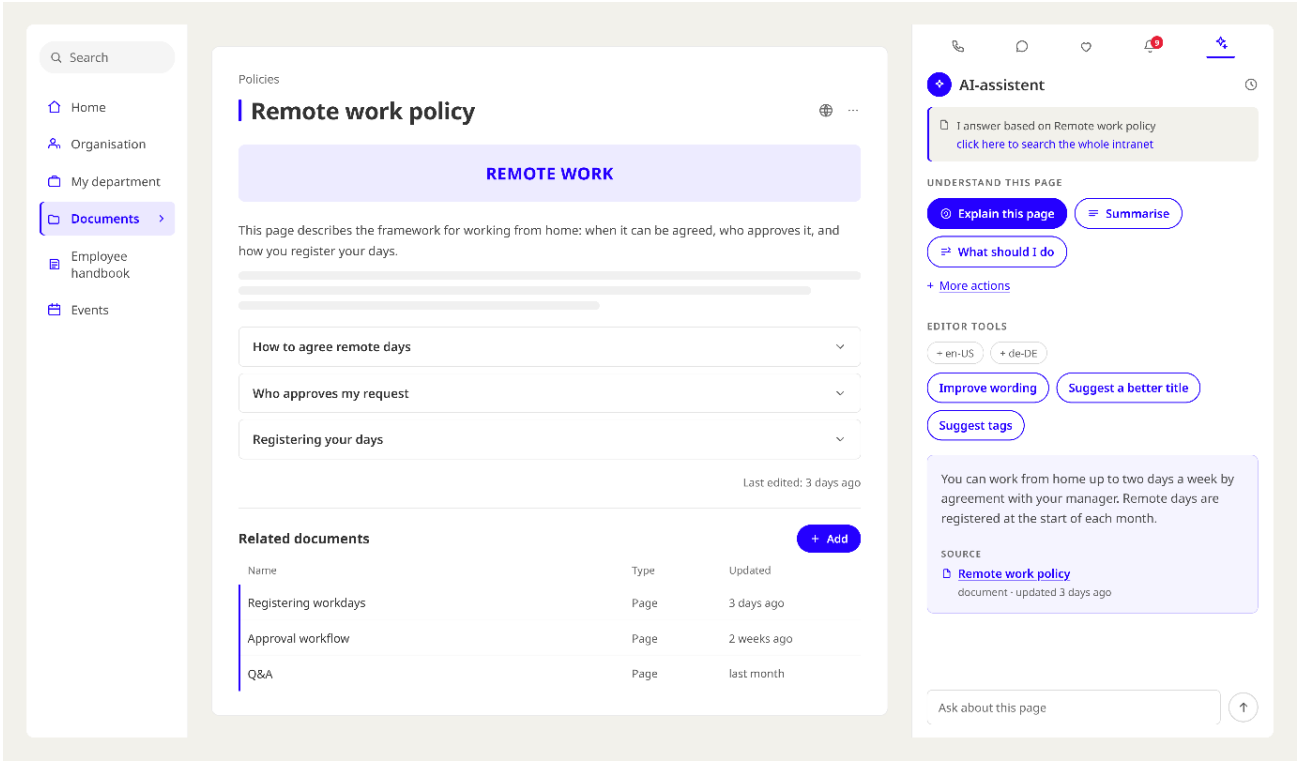
Ordinary intranet search remains relevant when the employee wants to find and read the documents themselves. Embedded AI is a supplement that can help answer a specific question from the sources it finds.

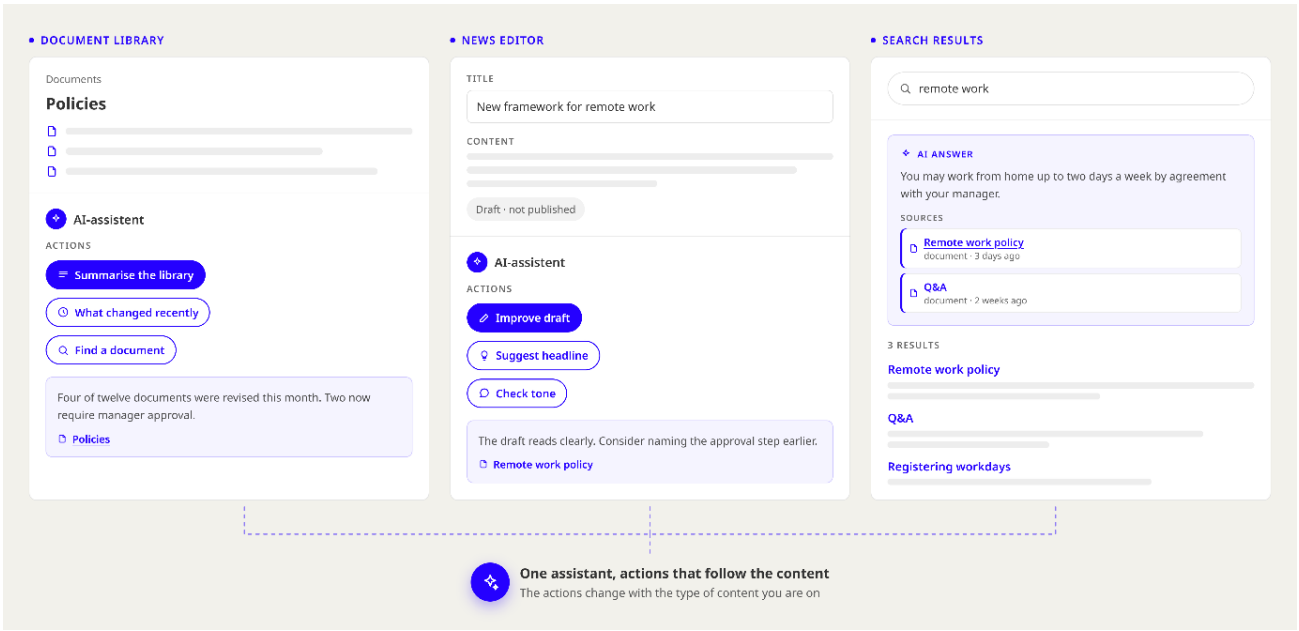
The quality of the answers depends on the quality of the content. If a document is outdated, or if there are several conflicting versions of the same policy, that can affect the answer. For this reason the employee can mark an answer as helpful or not helpful, and indicate, for example, that the information is wrong or outdated. This feedback becomes visible to the content owners in the administration.

Shortcuts and actions

The assistant can offer relevant actions depending on where the employee is in Colibo. On a document page it could be, for example, “Explain this page” or “Summarise the content”. In a news editor, an editor can get help to improve a text, suggest a headline, add tags or translate content.

The actions shown are drawn from a register and can be adapted to the individual organisation, the content type and the user's role. Administrators decide which actions are available to different groups, and can switch them off individually or all at once.





What happens when someone asks a question?

Colibo gathers the relevant intranet content and associated metadata that the employee has access to, and sends the information to the API of the AI provider the organisation has chosen. The answer is then shown in Colibo's assistant panel.

All lookups in Colibo are made with the current employee's permissions. The function can therefore not retrieve content that the employee does not have access to. It is Colibo's API that enforces the permissions, independently of the settings for the AI model.

Requirements for the AI provider

To use Embedded AI, the organisation needs an API account with a supported AI provider and must configure the API key in Colibo Admin. The account is separate from any subscriptions to AI chat tools. A subscription such as Claude Team or ChatGPT Enterprise does not, for example, give access to the API calls that Embedded AI uses.

The organisation creates the API account itself and pays the AI provider directly. The IT department can manage usage, limits and rate limits in the provider's administration environment. Colibo does not resell AI usage. Chapter 9 describes the further options for managing usage in Colibo.

Embedded AI supports Anthropic, OpenAI and Azure OpenAI.

Administration, content quality and usage

AI Enabler Administration comes with Embedded AI and is a separate administration site that gives an overview of content quality, employee feedback and the use of Embedded AI.

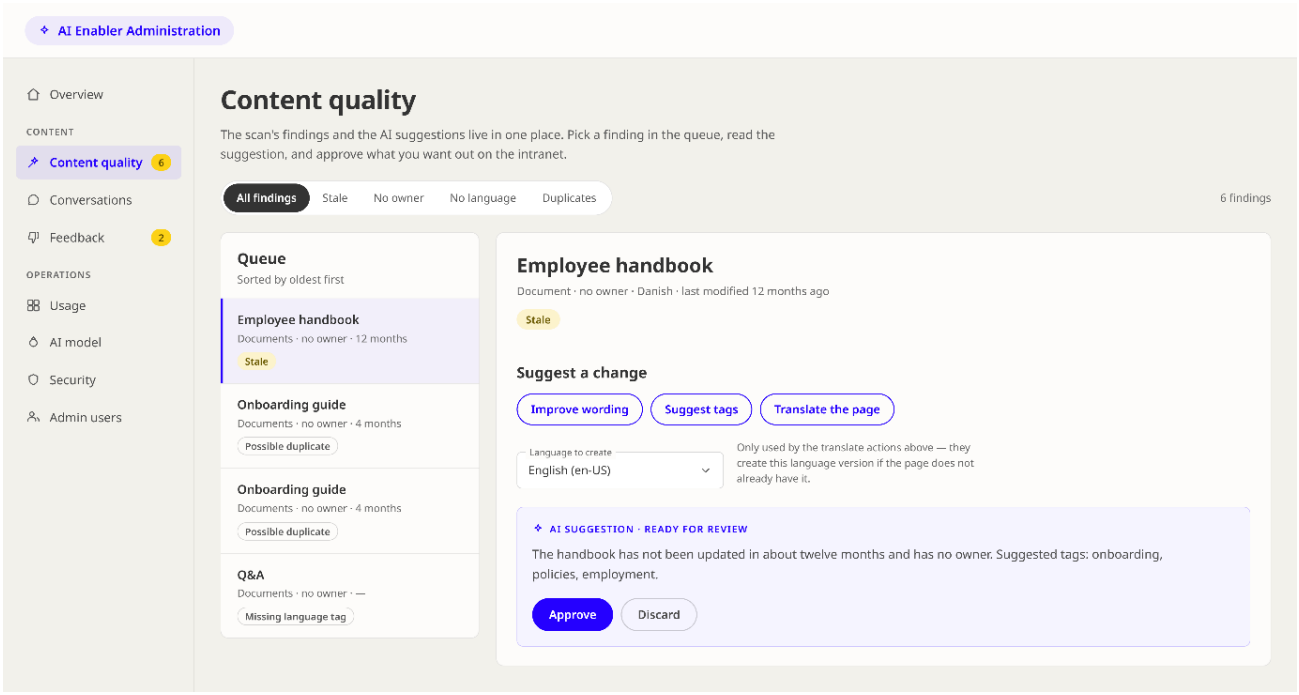
Here administrators can, among other things, find pages that need updating. These can be documents without a responsible owner, duplicates, or translations that do not follow the original. The solution can suggest a correction, which an administrator must review and approve before the content is changed.

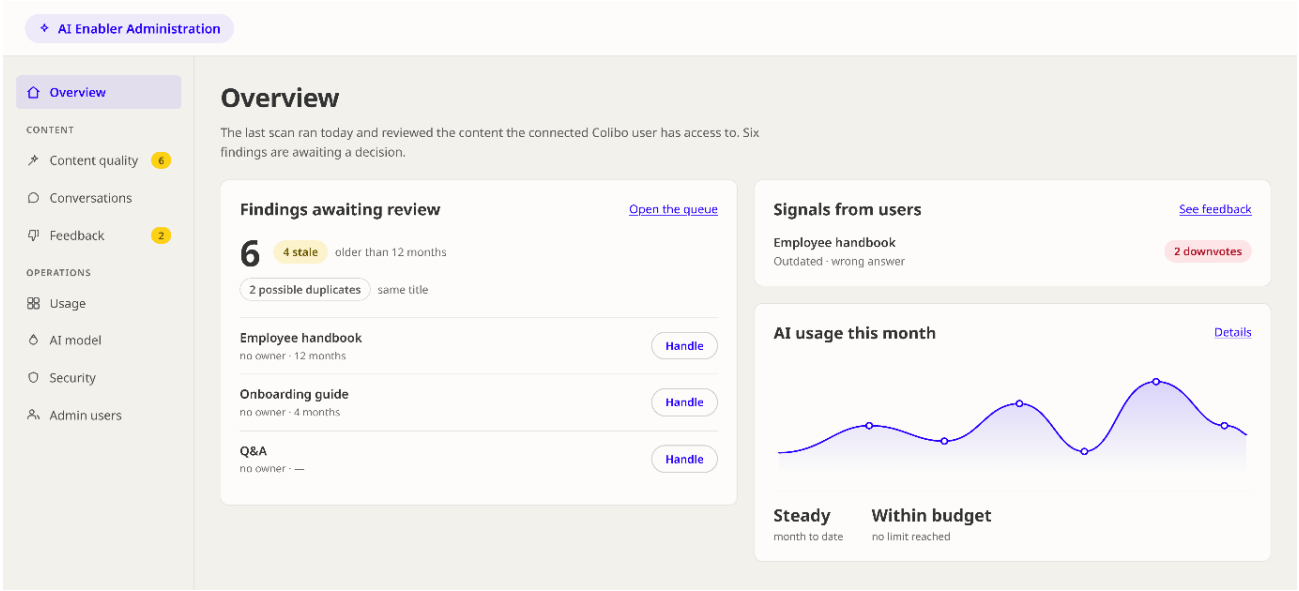
Administration also shows information about usage and costs, as well as feedback on answers. If an employee marks an answer as outdated, for example, it can be linked to the source the answer is based on. That makes it easier for content owners to prioritise maintenance.

Your own instructions for the AI

In AI Enabler Administration, administrators can add their own special instructions for the AI to answer from in Embedded AI. The instructions can apply generally to all answers or be tied to an individual action, for example to shape how a summary is worded. They apply only to Embedded AI and do not affect AI tools that use the MCP server. They supplement the default instructions in the Intelligence Layer and do not change Colibo's access control, which is still enforced in the API. Like all AI instructions, they are not a technical guarantee, but they influence how the model responds.

There is also a combined overview of matters that need attention.





The components can be activated separately

Colibo Embedded AI can be activated without Colibo MCP Server. The organisation does not need to have set up MCP connections to external AI tools first in order to give employees the assistant in the intranet. Equally, Embedded AI is not a prerequisite for the MCP server.

If the other component is activated at a later point, it requires no changes to the existing configuration, users' access permissions or the Intelligence Layer setup.

04
Colibo Intelligence Layer

When an AI assistant gets access to the intranet, there is still a need to specify how it should use the information. That includes, for example, source references, approval before changes, and handling of questions about sensitive HR or legal matters.

The Colibo Intelligence Layer contains default instructions for the AI model and can be adapted to the organisation's needs. The instructions are used with both the MCP server and Embedded AI.

Important: what the instructions can and cannot do

The instructions influence the AI model's answers and suggestions, but are not a technical guarantee that the model will always follow them. Results can vary between AI platforms, model versions and individual questions. This is a significant difference from Colibo's access control, which is enforced technically in the API and does not depend on the model's response.

For actions with significant operational consequences, Colibo continues to recommend that an employee reviews the result.

Sources in the answers

The instructions ask the AI assistant to state the pages or documents it uses as the basis for an answer. This lets the employee read and check the original information.

In Embedded AI the sources are also shown as clickable links in the Colibo interface, with information about the last update. Here the display of source references is a function of the interface and not only an instruction to the AI model.

Approval before changes

The Intelligence Layer contains instructions that the AI assistant must describe proposed changes and ask the user for approval before it, for example, creates, edits, publishes or deletes content or sends an onboarding message. How this is handled in practice depends on the AI platform and integration used.

Use of internal sources

For questions about the organisation's own rules and procedures, the AI assistant is instructed to start from the intranet content rather than general information from the internet. This reduces the risk of a general answer being mistaken for the organisation's current practice.

Questions that should be referred on

Some topics are not suited to being decided by an AI assistant. These include, for example, employment disputes, disciplinary cases, legal questions and confidential pay matters. The Intelligence Layer can instruct the assistant to refer such questions to HR, the legal department or another relevant contact. The organisation can adapt the categories itself.

Adaptation by role

Instructions and tools can be adapted to the user's role. Communications staff can, for example, get access to functions for writing and editing, while administrators can get tools for content analysis. HR staff can get instructions that fit the department's own procedures for referral and escalation.

Default setup and your own instructions

The Intelligence Layer comes with a set of default instructions. The organisation can use them as a starting point, add its own rules or write its own instruction set. In Embedded AI, administrators can add their own instructions in AI Enabler Administration, see chapter 3.

When implemented together with Colibo's team, the review can cover the choice of default behaviour, special internal rules, relevant contact routes and testing before go-live. This is particularly useful in organisations with special requirements for communication, approval or sensitive information.

05
Use Colibo through your existing AI tools

Primarily for communications and intranet managers

With Colibo MCP Server, employees can use their usual AI tool to find information and carry out relevant tasks in Colibo. The examples below show how the solution can be used.

The scenarios build on functions that have been demonstrated during the development and testing of the MCP server. Names, data and results are illustrative and do not come from a specific customer.

Example 1. Questions about HR policies

Situation: An employee works from home and wants to know whether the organisation gives an allowance for ergonomic equipment. The employee is unsure where the policy is kept, and would normally search the intranet or contact HR.

With Colibo MCP Server: The employee asks their AI assistant: “What are our rules for home office equipment?” Through the MCP server, the assistant can retrieve the relevant section of the HR policy in Colibo and give a short explanation with a link to the document.

Possible benefit: The employee finds the current policy more easily, and HR may receive fewer routine enquiries.

Example 2. Overview of intranet content

Situation: An intranet manager wants to find out where guides are missing and which pages have not been updated.

With Colibo MCP Server: The intranet manager can, for example, ask: “What do employees search for without finding a guide?” or “Which pages have not been updated in the last six months?” The MCP server can retrieve relevant search data and document metadata from Colibo. On that basis the assistant can prepare a prioritised list with links to the pages.

Possible benefit: The intranet manager gets a better basis for prioritising clean-up and new content without reviewing every page personally.

Example 3. Onboarding reminders

Situation: An administrator can see that several new employees have not yet started or completed an onboarding programme. Finding and contacting them one by one can be time-consuming.

With Colibo MCP Server: The administrator asks their AI assistant to identify the relevant employees, draft a personal reminder and prepare it for sending through Colibo's notification system. The MCP server retrieves status from the onboarding API. The assistant shows the message to the administrator and asks for approval before it is sent, in line with the instructions in the Intelligence Layer.

Possible benefit: The administration can save time on finding recipients and drafting reminders. The person responsible still reviews the message before it is sent.

Example 4. News based on existing policies

Situation: A communications employee needs to write a news item about new opportunities for flexible working. The item must be consistent with the organisation's hybrid working policy and refer to the current HR guidance.

With Colibo MCP Server: The employee asks the AI assistant for a draft with references to the existing policy. The MCP server retrieves the relevant text from Colibo, and the assistant drafts a proposal for the news item with links. The draft is shown to the employee, who can review and approve it before publication.

Possible benefit: It becomes easier to write content that builds on the organisation's own guidelines, and the editorial review can focus on quality and accuracy.

06
Choice of AI platform

Colibo AI Enabler is set up so that the organisation does not have to choose the same AI provider for every purpose. The platforms used through the MCP server can differ from the provider that may run Embedded AI in Colibo.

With the MCP server, one team can, for example, use Claude and another Microsoft Copilot or ChatGPT, provided that the relevant connections are configured. With Embedded AI, one AI provider is chosen for use inside the intranet itself. The two choices are independent of each other.

That separation can matter for organisations that already have agreements, data processor requirements and internal guidelines for different AI services. The choice of intranet integration does not have to settle at the same time which AI platform the organisation will use in the longer term.

Platforms for the MCP server

Colibo MCP Server is built on the open MCP standard and is designed to work with MCP-compatible AI platforms.

The following platforms are verified and ready for use with the MCP server:

- **Claude (Anthropic):** Can be used with cloud and API installations.
- **Microsoft Copilot:** Relevant for organisations with Microsoft 365 that want access to Colibo content through Copilot.
- **ChatGPT (OpenAI):** Can be used by organisations with an existing OpenAI Enterprise agreement.

Other MCP-compatible systems can also be connected, but may require additional configuration.

The difference between MCP Server and Embedded AI

Colibo MCP Server can connect several AI platforms to the same Colibo installation. If another tool is added, it does not need to affect the existing connections.

Colibo Embedded AI uses one AI provider at a time, because Colibo sends its API calls to the provider specified in Colibo Admin. The provider can be changed later.

An organisation can therefore use Claude and Microsoft Copilot through the MCP server and at the same time choose Azure OpenAI as the provider for Embedded AI.

A note on Microsoft

Microsoft 365 Copilot can be connected to Colibo through MCP. For Embedded AI, Azure OpenAI is used when Colibo needs to call an AI API directly. These are therefore two different integration methods, even though both solutions are delivered within Microsoft's product area. Access depends on the organisation's specific agreements.

If you change AI platform

Because the MCP server uses an open standard, an organisation can switch between supported AI platforms without developing a completely new integration to Colibo. Colibo's access control and the configuration tied to the intranet stay in Colibo.

This gives flexibility if the organisation later wants to change its AI platform or take more tools into use. Connecting a new platform still needs to be configured and tested, of course.

07 Technical architecture

Primarily for IT and security managers

This chapter describes how Colibo MCP Server connects the organisation's AI environment with Colibo. The architecture is also the starting point for the review of data security in the next chapter.

Two separate environments

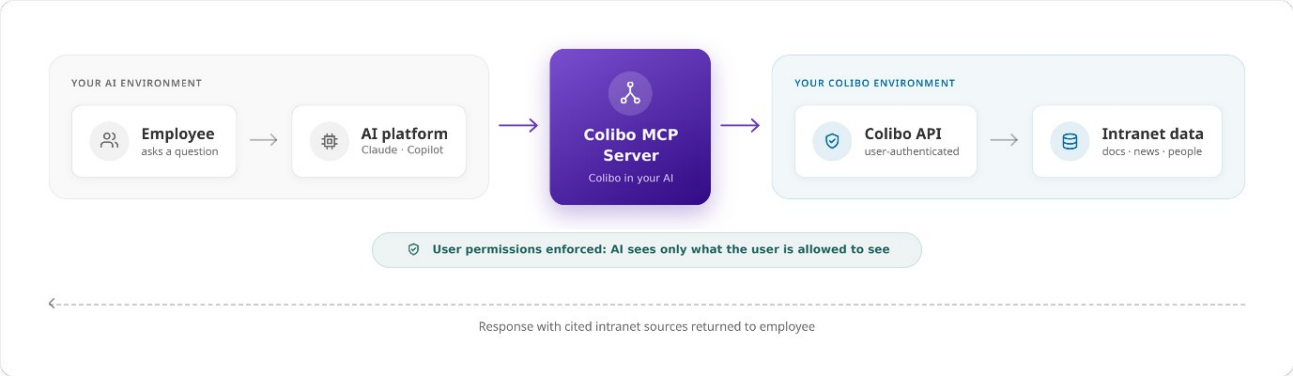
The solution connects two environments, each with its own function.

The AI environment is the platform the employee uses, for example Claude or Microsoft Copilot. This is where the employee types their questions, and where the AI service processes requests and generates answers. The processing is covered by the organisation's agreement with the AI provider.

The Colibo environment contains the intranet's documents, news, employee information and other data. Colibo's API controls access, so that an employee can only retrieve the content they have permissions for.

Colibo MCP Server connects the two environments. When the AI tool requests information from the intranet, the MCP server converts the request into a call to Colibo's API. It retrieves information using that employee's credentials and returns it in a structured format to the AI tool.

The two environments keep their own roles: Colibo controls access to the intranet content, while the AI provider processes the AI request and produces the answer.



User identity and access permissions

All calls from the MCP server to Colibo's API are made with the credentials of the employee who asked the question. The MCP server therefore does not use a shared administrator access that could bypass the existing permission model.

If an employee does not have access to a document in Colibo, their AI assistant cannot retrieve it through the MCP server either. This also applies to access restricted to particular groups or departments.

The permission check sits in Colibo's API and therefore does not depend on how the employee words their question, or on how the AI model responds.

08
Security and data handling

Primarily for IT and security managers

When AI gets access to internal information, it is important to know which data is exchanged, who processes it, and which access permissions apply. This depends, among other things, on whether the organisation uses Colibo MCP Server or Colibo Embedded AI.

Data handling with Colibo MCP Server

With the MCP server, the organisation's AI tool asks for information from Colibo. The server retrieves the content the employee has access to and sends it back to the AI tool. The AI processing of the user's question and the generation of the answer take place in the AI provider's environment, under the organisation's agreement with the provider.

Colibo does not receive, process or store the AI tool's prompts or responses. Colibo delivers the intranet information the AI tool requests, within the user's access permissions.

Data handling with Colibo Embedded AI

With Embedded AI, the request starts inside Colibo. Colibo gathers the relevant intranet content and sends it to the API of the AI provider the organisation has chosen. Colibo is therefore involved in handling the request and not only in delivering data.

Prompts and responses from Embedded AI are stored in anonymised form in Colibo, so that administrators can review them through AI Enabler Administration. Colibo's processing is covered by Colibo's standard data processing agreement, which also covers user-generated content that may contain personal data.

The AI processing itself takes place with the chosen AI provider, under the organisation's own API account and the agreements the organisation has entered into with the provider.

The difference between MCP Server and Embedded AI should form part of the organisation's assessment of data processing and information security. If the organisation's guidelines distinguish between the two models, we recommend involving the data protection function before Embedded AI is activated.

Access permissions follow the employee

Whichever component is used, Colibo AI Enabler retrieves content only through the current employee's Colibo identity. Access permissions for groups, departments and individual documents continue to be enforced through Colibo's API.

If, for example, an employee in the finance department asks for a document from the executive team's private area without having access to it, Colibo does not return the document's content to the AI tool.

Security documentation and standards

Colibo holds the following certifications and assurance reports:

- **ISO 27001:** Colibo's work on information security and management of security risks is certified to the standard.
- **ISAE 3000:** An independent assurance report on information security and data processing practices is available.
- **GDPR:** Colibo processes personal data as a data processor under agreements with its customers and in accordance with the General Data Protection Regulation.

The documentation should of course form part of the organisation's usual supplier and security assessment.

Cloud and on-premises

Colibo Cloud is the standard setup. Here the MCP server is operated by Colibo, and access to intranet data takes place through ordinary API calls over encrypted connections.

On-premises is relevant for organisations that want to run the integration component in their own infrastructure, or where regulations do not allow a cloud-hosted solution. Here the MCP server is placed within the organisation's own operating environment.

09
AI usage and costs

Primarily for IT managers and budget owners

The price of using an AI model depends, among other things, on how much text the model receives and generates. When using Colibo Embedded AI, the organisation has its own account with the AI provider and pays the provider directly.

It is therefore the IT department that sets usage limits, any budget alerts and limits on the number of requests, in the provider's administration environment. These limits apply regardless of the settings that can be made in Colibo.

Embedded AI also contains a number of features designed to limit unnecessarily high token consumption.

How token consumption is limited

Reuse of fixed instructions. Part of the information in an AI request is the same every time, for example instructions and descriptions of the available tools. Embedded AI can send this part in a form that allows the AI provider to reuse it through caching. This depends on whether the chosen provider supports the feature.

Limited conversation history. With follow-up questions, the model needs the relevant part of the conversation so far sent again. Embedded AI limits the amount of history by removing the oldest messages and ending inactive conversations automatically.

A limit on how much content is sent. Each action has a fixed budget for the amount of context sent to the model. If the material is too large, it is shortened in a set order. The employee's question and the page's key metadata are kept. Longer texts are cut down, and the model is told that it sees only an excerpt. If the necessary information cannot fit within the limit, the AI call is not made.

Smaller models for simpler tasks. Background tasks such as naming conversations, categorising content and translating category labels can be carried out with a faster and cheaper model. Where possible, several tasks are combined into one request, and the results are stored so that the calculation is not repeated unnecessarily.

Content gathered in one step. For several actions, Colibo gathers the necessary material itself before a request is sent to the model. This can reduce the number of follow-up lookups and also makes it possible to show the sources the answer is based on.

Settings for administrators

In AI Enabler Administration an administrator can adjust several settings without changing the installation itself:

- **Reasoning level:** Where the chosen AI model supports it, the level can be set globally or for each action. The setting affects the balance between depth of answer, response time and cost.
- **Output limit:** Limits how much text the model may generate per request.
- **Monthly spending cap:** An optional cap based on Colibo's estimate from recorded tokens. When the cap is reached, new requests stop until an administrator steps in. The cap does not replace the limits at the AI provider.
- **Usage overview:** Records token usage for each request and shows the trend in a dashboard.

Protection against unexpected usage

Embedded AI limits the number of requests that each user can send per minute. This reduces, for example, the effect of a faulty program or an automated script. Colibo is also designed not to retrieve the same intranet content several times while preparing a single request.

Costs when using MCP Server

When employees use an external AI platform through the MCP server, it is the AI platform that governs the size and price of the requests. Colibo's ability to influence usage is here limited to the data the MCP server returns.

Recommended setup

Colibo recommends using both the hard limits at the AI provider and the operating limits that can be set in Embedded AI. Together they give better control of usage.

10 Getting started

Prerequisites

What needs to be in place depends on which components the organisation activates:

- **An active Colibo subscription.** AI Enabler is an add-on to existing Colibo subscriptions.
- **For Colibo MCP Server: an approved AI platform with MCP support.** The organisation needs an active subscription to a relevant AI platform.
- **For Colibo MCP Server: IT approval of the connection.** The MCP server creates an authenticated API connection between the AI environment and Colibo. The organisation's IT department must review and approve the setup, especially for on-premises installations.
- **For Colibo Embedded AI: an AI API account.** The function requires an account with a supported AI provider and a corresponding API key in Colibo Admin. Embedded AI does not require the MCP server, and the account is not needed if only the MCP server is used. See chapters 3 and 9.

How the implementation works

Colibo's implementation team helps with the setup. The process is described in three steps:

01 **Technical setup.** The chosen components are set up in the organisation's Colibo environment and tested. For the MCP server this covers the connection to Colibo's API, for Embedded AI the configuration of the AI provider's API key in Colibo Admin.

02 **Adapting the Intelligence Layer.** The team reviews the default instructions and adapts them as needed, for example with special rules, topics to be referred on, and role-based settings. The configuration is tested.

03 **Testing and handover.** The organisation carries out a user acceptance test of the most important scenarios before the solution is handed over to operation.

Pricing

The price of Colibo AI Enabler depends on the size of the organisation and is a combined price for the components the organisation activates: Colibo MCP Server, Colibo Embedded AI, including AI Enabler Administration, or both. All price levels include the standard library for the Intelligence Layer. Setup and onboarding are charged as a one-off fee.

Employees	Price per month
1-499	EUR 200
500-999	EUR 300
1,000-4,999	EUR 400
5,000 or more	Contact Colibo

What is included?

Area	Terms
Operating model	Colibo Cloud or on-premises
Colibo MCP Server and Colibo Embedded AI	Included, separately or together. AI Enabler Administration comes with Embedded AI. Embedded AI requires the organisation's own AI API account, which is paid directly to the AI provider
Intelligence Layer	Standard library included, customisation on request is charged by the hour
AI platforms	One platform included, further platforms can be added by the hour
Setup and onboarding, Colibo Cloud	EUR 2,000 one-off
Setup and onboarding, on-premises	EUR 3,000 one-off

All prices exclude VAT. Contact your Colibo account manager or see colibo.com if you would like a demonstration or want to discuss a specific setup.

About Colibo

Colibo provides a digital workplace platform for internal communication, knowledge sharing and collaboration. The platform is used by private and public organisations in Europe, including municipalities, public institutions and companies.

Colibo is headquartered in Aarhus, is ISO 27001 certified and holds an ISAE 3000 assurance report. The platform can be operated in the cloud or on-premises and integrates with Microsoft 365.

colibo.com · marketplace.colibo.com