


Security, Privacy, and Trust Overview

AI Assist in FME Form



Contents

- 1. Trust & Transparency.....3
- 2. Feature Overview.....4
 - 2.1 What is AI Assist?.....4
 - 2.2 Key Capabilities.....5
- 3. Value for Organizations.....5
- 4. Our Approach to AI.....5
- 5. Data Handling and Protection.....6
 - 5.1 What Data is Sent.....6
 - 5.2 What Data is Stored.....7
 - 5.3 Where Data is Stored.....7
 - 5.4 Data Training.....7
 - 5.5 Subprocessors.....8
- 6. AI Controls & Configuration.....9
 - 6.1 Disable AI Assist Context Sharing from Within the Workspace (In Product).....9
 - 6.2 Disable AI Assist (In Product).....9
 - 6.3 Disable AI Assist (Machine-Level).....9
 - 6.4 Block Communication to the FME AI Service (Network-Level Control).....10
- 7. Transparency and Compliance.....10
- 8. FAQ.....11
- 9. Resources.....12



1. Trust & Transparency

At Safe Software, we are committed to empowering organizations to securely and effectively solve their data integration challenges. With AI Assist in FME Form, we bring artificial intelligence into the authoring experience in a way that is transparent, secure, and privacy-conscious, aligned with globally recognized best practices and compliance frameworks.

We adhere to three foundational principles:

- **Security:** Our secure-by-design philosophy enforces technical and administrative safeguards to ensure AI features are rigorously tested, responsibly deployed, and monitored for ongoing risk mitigation.
- **Transparency:** We equip organizations with clear, accessible insights into how AI Assist operates, the nature of data processed, and how decisions are made.
- **Privacy:** We align with GDPR, Canada's Personal Information Protection & Electronic Documents Act (PIPEDA), and other applicable laws, and apply privacy-by-design principles, including data minimization and governed retention.

This ensures that AI Assist is not only effective but also trustworthy, controllable, and compliant.



AI Assist brings contextual intelligence into the FME Form authoring experience



2. Feature Overview

2.1 What is AI Assist?

AI Assist is an end-user, in-product opt-in feature, built into FME Form to enhance productivity by offering contextual suggestions and guidance directly in the authoring environment. The AI Service leverages large language models provided through Microsoft Azure OpenAI Service and AWS Bedrock (Anthropic Claude family of models).

Note: AI Assist is available in FME Form. It is not currently available in Esri's Data Interoperability (Data Interop) extension for ArcGIS.

AI Assist requires user consent after login ([FME Account](#)), via a pop-up dialogue, before any data is transmitted to the AI service.



Opt-in User Consent Pop-Up Window

2.2 Key Capabilities

- **Contextual Suggestions:** Offers in-context code generation within select transformers (e.g., Regex, SQL, and Python).
- **Conversational Guidance:** A conversational interface that helps users understand, debug, and improve their workflows by referencing relevant Workspace Context. AI Assist enables users to retain and access their chat history across sessions, allowing them to revisit past conversations. By tying chat data to individual users, it improves problem-solving continuity and the overall AI Assist experience.

3. Value for Organizations

AI Assist offers organizations a responsible and effective way to streamline their data integration workflows by integrating intelligent, context-aware suggestions directly into the user interface. It enhances productivity while maintaining robust governance controls.

AI Assist helps organizations:

- Accelerate workflow development by suggesting code, expressions, next action (reader, writer, transformer), and parameters based on Workspace Context.
- Reduce onboarding and training costs by providing natural-language explanations and smart guidance directly within the interface.
- Maintain operational efficiency by enabling faster discovery of FME features without relying on external support or documentation.
- Ensure safe usage by providing full user control over feature access, Workspace Context sharing, and data retention.

Backed by Safe Software's trusted content (documentation, knowledge base, and community resources), AI Assist empowers teams to innovate with confidence.

4. Our Approach to AI

At Safe Software, we view AI as a collaborative assistant to augment human decision-making, not as an autonomous agent or authority. We encourage users to treat AI suggestions as helpful inputs, not definitive answers. Using AI without critical oversight can be risky, which is why FME is designed to keep users in control of their work and data.

AI Assist does not replace human expertise. It serves as a supportive tool to accelerate learning and enhance workflow quality while maintaining decision-making autonomy in the hands of the user.

Safe Software is ISO/IEC 27001:2022 certified and SOC 2 Type II compliant. Our development practices follow internationally recognized standards for information security and operational controls, reflecting our commitment to secure, reliable, and privacy-conscious software design.

We test all AI-powered features thoroughly to ensure they are accurate and aligned with our product philosophy. AI Assist is part of our internal User Acceptance Testing. We trial, evaluate, and perform quality assurance tests before using new services and models. For more details, please see our [Responsible AI Policy](#).

5. Data Handling and Protection

AI Assist is developed in alignment with Safe Software's secure software development principles. This includes our "Shift Left" security approach, addressing security early and throughout the development lifecycle, and a commitment to continuous improvement as security practices and industry expectations evolve.

In addition to shifting security considerations earlier into development, Safe Software applies data minimization principles so that AI Assist only retains the minimum information necessary to ensure functionality, evaluation, and traceability. Data handling is purpose-limited, transparent, and fully documented to support compliance with organizational policies and applicable regulatory standards.

5.1 What Data is Sent

AI Assist does not send customer data or actual data values to the service. Instead, it may use metadata associated with the user's workspace to provide more relevant responses. Everything we send other than your prompt is called "Workspace Context" data. This includes elements such as bookmarks, annotations, and metadata about the data.

The following information is sent to the FME AI Service:

- For contextual code suggestions
 - **Regex:** Only the user's current prompt is sent to the FME AI Service.
 - **Python:** The user's current prompt and the Python code visible in the Python AI Assist window are sent to the FME AI Service.
 - **SQL:** The user's current prompt is sent to the FME AI Service. There is a user-controlled checkbox to "Send database schema to AI", which, when enabled, includes: database type (e.g., PostgreSQL, OracleDB), table names, column names, and column data types. No actual database values are ever sent. For more details, see [FME Form Documentation - AI Assist](#).
- For conversational guidance
 - Depending on the user's selected operation or prompt, the FME AI Service may reference the following contextual information from the workspace:
 - Canvas settings, including user locale
 - Canvas connections, including how transformers, readers, and writers are linked
 - Canvas content, including annotations, bookmarks, transformers, and format tables/layers/schemas
 - Parameters associated with transformers, readers, and writers, which may include, but are not limited to:
 - Names of databases or web connections
 - Dataset or supporting file paths
 - Database table and field names (e.g. column names in a CSV file)
 - Web URLs, API endpoints, AWS S3 identifiers (e.g. bucket name and IDs, queue names, etc.)
 - FTP server connection details, folders, and filenames

- SQL queries (e.g. column names and values), Python scripts, and embedded expressions
- System Commands (e.g., from SystemCaller)
- Sensitive information may be present if users do not follow best practices and hardcode values in their workspace, such as database hostnames, ports, database names, usernames, or other personally identifiable information (PII).
- To disable sharing Workspace Context, please see the [User Controls and Configuration](#) section.

Important: Password-type fields are never in plaintext and never transmitted to the FME AI Service. However, any credentials exposed in customers' scripts or non-password fields may be included in the context if visible.

Users are responsible for following best practices to avoid exposing sensitive content within workspace scripts or comments.

5.2 What Data is Stored

In alignment with the data minimization privacy principle, AI Assist retains only the minimum information necessary to ensure functionality, evaluation, and traceability.

User Metadata

- A unique user ID associated with the user's FME Account (an internal identifier that can be used to correlate to the user's email and information within the FME Account for traceability).
- The date when the FME AI Service was first accessed by the user (first login date with FME Account).

Conversation History

- AI Chat interactions
- Workspace Context, as outlined above in the ['What Data is Sent'](#) section.

5.3 Where Data is Stored

The FME AI Server and its supporting databases are hosted on AWS in the United States. Conversation history and associated Workspace Context are stored in the FME AI Service Application Database and are retained until explicitly deleted by the user. Users can manage this data directly through the FME Form interface. For full data removal requests, please contact [privacy at safe.com](mailto:privacy@safe.com). Usage statistics related to contextual code suggestions are stored in a separate database. This dataset is fully anonymized, not linked to any FME Account ID, and currently has no defined retention period.

Additional information is sent to third-party providers for the sole purpose of providing the AI Assist Service. Refer to the table in section 5.5 for more details.

5.4 Data Training

Customer data is not used to train the underlying LLM, regardless of which model provider is used. Only minimal metadata and conversation history are stored in the FME AI Service Application database. LangSmith is used internally to evaluate prompt and response quality only, not for LLM training.

AI Assist currently supports models hosted on both AWS Bedrock and Azure OpenAI. Handling of prompt/completion data at the model-hosting layer differs by provider:

- **AWS Bedrock (Anthropic Claude family of models):** No data retained.
- **Azure OpenAI:** Microsoft operates an automated abuse-monitoring system on prompts and completions to detect Code of Conduct violations. In most cases, this review is fully automated and no data is stored. In a limited subset of cases where automated review can't reach a confident determination, a sample of the flagged content may be temporarily stored for human review by authorized Microsoft personnel. In no case is this data used to train any model. For more detail, see Microsoft's documentation on [Azure OpenAI abuse monitoring](#).

5.5 Subprocessors

Subprocessor	Purpose	Data Processed	Data Retention	Location
Microsoft / Azure OpenAI	Language model provider for AI Assist functionality	AI prompts and outputs; Workspace Context	Not retained for automated review. For the limited subset of content flagged for human abuse-monitoring review, retained temporarily (up to 30 days) by Microsoft.	Global*
LangSmith	Tracking and improving AI interactions	AI prompts and outputs; Workspace Context; user metadata	14 days by default / 400 days if user feedback is provided or if flagged for analysis	United States
AWS	Language model provider for AI Assist functionality; hosts FME AI Server & related databases	AI prompts and outputs; Workspace Context; user metadata; usage statistics	Until deleted by user in product or upon request for full data deletion. No data retained by AWS Bedrock (Anthropic Claude family of models).	Global*
Auth0	FME Account authentication	IP address; FME Account credentials	Until FME Account is deleted by the user	United States
Sentry	Monitors FME AI Service for error processing	IP address; User ID; error-related prompts	90 days	United States; EU
DataDog	Logging & monitoring of FME AI Service	IP address	30 days	United States

* "Global" does not include sanctioned countries.

6. AI Controls & Configuration

6.1 How to Disable AI Assist Context Sharing from Within the Workspace (In Product)

AI Assist context sharing can be enabled or disabled at the workspace level. This setting controls whether workspace content is shared as context for AI Assist conversational guidance and contextual suggestions.

Context sharing is configured per workspace. Changing the context level:

- Applies only to the currently open workspace
- Does not affect other workspaces
- Does not modify global workspace defaults

When a workspace is opened, the last saved context level for that workspace is used. If opening an older workspace saved prior to 2026.1, then the context level specified in the workspace defaults will be used.

You can change a workspace's context level in one of the following ways:

1. **From the Navigator:** In FME Workbench, go to **Navigator > Workspace Parameters > AI Assist**. Set **Enable Context in the AI chat** to **No**.
2. **From AI Assist Chat:** In FME Workbench, open **AI Assist Chat > Parameters Icon  > AI Context for this workspace** and turn the toggle off.

Set the desired context level for newly created and older workspaces in **Workspace Defaults**. For more details, see [FME Form Documentation - AI Assist](#).

6.2 How to Disable AI Assist (In Product)

To disable AI Assist entirely, in FME Workbench, go to **FME Options > Workbench > AI Assist Options** and uncheck “**Enable AI Assist**”.

6.3 How to Disable AI Assist (Machine-Level)

Organizations may disable AI Assist via a machine-level administrative control. When set, the machine-level disable applies to all users on that device and cannot be changed by end users from within the product UI.

When the machine-level disable is set, AI settings will be hidden from the user. The only visible location is in **FME Options (Workbench)**, where the **Enable AI Assist** setting is disabled (unchecked and not editable). These UI controls cannot be used to re-enable AI Assist on that machine; changing the state requires modifying the system-level setting (which requires administrative access).

Windows

A machine administrator can disable AI Assist by creating a **Disable AI Assist** key in the HKLM (Local Machine) registry. This registry key disables AI Assist for the entire machine and is not exposed in FME Options. Non-administrator users cannot re-enable AI Assist from the product UI or from the registry.

In Registry Editor, navigate to **Computer\HKEY_LOCAL_MACHINE\SOFTWARE\Safe Software Inc.\FME Workbench**. Create a new key **"Disable AI Assist"** and set the value to **"true"**. Restart FME Workbench for the change to take effect.

macOS

A machine administrator can disable AI Assist using a system-scope preference file (plist). Administrator privileges are required to create or modify this setting.

```
sudo defaults write "/Library/Preferences/com.safe-software-inc.FME Workbench" "Disable AI Assist" -bool YES
```

Restart FME Workbench for the change to take effect.

Linux

A machine administrator (root) can disable AI Assist using a system-scope configuration file. Root privileges are required to create or modify this setting.

```
sudo mkdir -p "/etc/xdg/Safe Software Inc."
echo -e "[General]\nDisable AI Assist=true" | sudo tee "/etc/xdg/Safe Software Inc./FME Workbench.conf" > /dev/null
```

Restart FME Workbench for the change to take effect.

6.4 How to Block Communication to the FME AI Service (Network-Level Control)

If your organization wants to prevent the usage of AI Assist globally, you can block all network communication with the FME AI Service by blocking the domain **fme-ai-service.safe.com** in your firewall settings. Please note that by using this method, FME Form UI elements will remain visible, but no data will be transmitted.

7. Transparency and Compliance

Safe Software is dedicated to empowering users through AI without compromising on security or trust. AI Assist in FME Form is designed with transparency, control, and compliance at its core. Following the Principle of Least Privilege and the Need-to-Know principle, access to the FME AI Service Application Database, Usage Statistics Database, and LangSmith is restricted, with personnel granted only the minimum level of access required to perform their job functions. AI Assist includes visual progress reporting to improve transparency and user confidence. This aligns with Safe Software's ongoing commitment to explainability and transparency in AI-powered features. As this feature continues to evolve, we welcome feedback to ensure it meets the needs and expectations of our global user community. We continuously monitor all applicable laws and may adapt our product functionality from time to time to ensure compliance.

We follow the principles of data minimization, transparency in functionality, user consent and control, and secure handling of sensitive metadata.

8. FAQ

Is AI Assist enabled by default?

Yes, but it requires explicit user consent through an in-product, opt-in dialogue.

Can AI Assist be fully disabled?

Yes. For details, see the [User Controls & Configuration](#) section.

Is any data used to train the AI model?

No. Customer data is not used for model training. Only limited metadata is stored for troubleshooting and service improvement.

Where is the server hosted?

AWS in `us-east-1`.

Where is the model hosted?

Microsoft Azure OpenAI Service.

Is PII sent to the model?

No PII is sent unless included by the user in prompts or workspace content (e.g., annotations, scripts, or parameters not marked as password-type). Users are encouraged to follow best practices and avoid embedding sensitive or personal information in scripts, annotations, or other workspace content.

Is user data stored?

Yes, minimal metadata and conversation history are stored in the FME AI Service Application Database and can be deleted by the user. LangSmith retains conversation history for 14 days by default, or up to 400 days if feedback is provided or the session is internally flagged.

Additionally, Microsoft's abuse-monitoring system may temporarily store a sample of prompts and completions for human review if automated review can't reach a confident determination. Note, this data is never used to train any model. AWS-hosted models do not retain data.

For details, see the [What Data is Stored](#) section.

Where is the data stored?

It is stored in databases secured on AWS and LangSmith. No data is retained in the AI model or AWS Bedrock. Microsoft Azure OpenAI does not use customer prompts or completions to train any model; however, in a limited subset of cases where Microsoft's automated abuse-monitoring review can't reach a confident determination, a sample of flagged content may be temporarily stored for human review by authorized Microsoft personnel.

Can users control or delete their data?

Yes. Users can delete individual conversations and all chat history from within the application. For complete data removal requests, please contact privacy@safe.com.

Is my FME workflow data sent to Safe or OpenAI?

No actual datasets are sent; only metadata embedded in the workspace may be included.

Who owns the output generated by AI Features?

You retain all intellectual property rights to the output generated by AI Features. Safe Software does not claim ownership of any AI-generated outputs.

Does Safe Software transfer my data across borders?

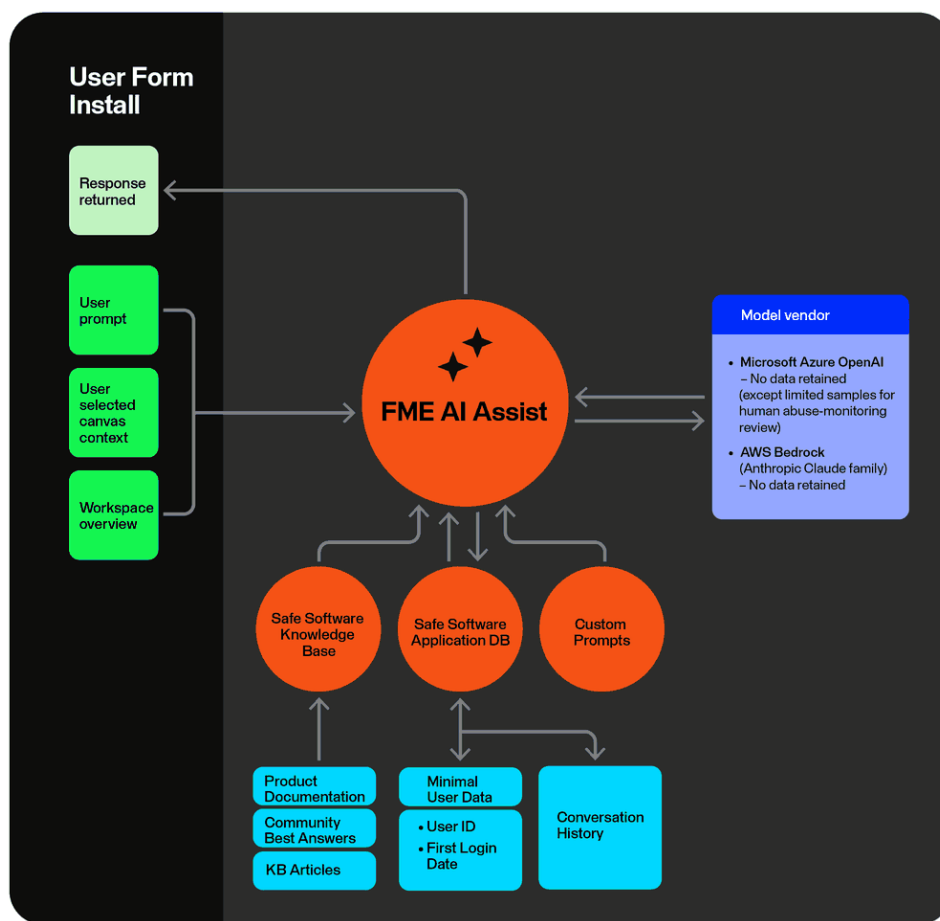
The FME AI Service is hosted in AWS (us-east-1) and LangSmith in the Google Cloud Platform (United States). Depending on the location of our servers and third-party services, your data may be transferred across borders. Cross-border transfers are conducted in accordance with applicable data protection laws and are supported by contractual and technical safeguards, including data processing agreements with our service providers.

Where can I send questions or feedback about AI Assist?

Leave a comment in the [AI Assist in FME FAQ](#) Knowledge Base article or [submit a ticket with support here](#).

9. Resources

- [AI Terms of Use](#)
- [AI Assist in FME Form FAQ](#)
- [AI Assist Documentation](#)
- [Responsible AI Policy](#)



Safe's AI Service flow — what is sent, what is stored, and where