


## FIELD BRIEFING

# Your data has a route

Follow information beyond the model.



## Where does information go when I use AI?

“Local” and “cloud” are useful starting labels, but neither describes an entire information journey. A local application can call a remote model. A cloud model can receive only redacted text while originals stay elsewhere. A local model can still expose information through logs, backups, plugins or exports.

Explanatory schematics and worked examples are illustrative. No live model, observed client results or measured market performance are presented.

# Read the mechanism.

Draw the route for one task. Identify the input, preprocessing, model endpoint, tools, output store, logs and backup locations. Then identify who controls each step. This makes questions about access, retention and deletion concrete enough to investigate.

An invented document-summary workflow illustrates the distinction. A team could send a full customer file to an approved remote service; it could remove direct identifiers before sending; or it could run an appropriate model on controlled local hardware. Each option has different quality, operational and data-handling requirements. Redaction reduces some disclosures but is not automatically anonymisation: a document's contents may identify a person without a name.

Check policies and settings for the exact service and plan. Questions include whether inputs are used for training, which logs are retained, where processing occurs, who can access content, which subprocessors participate and how deletion works. A statement about one product should not be silently applied to another product from the same vendor.

For local operation, ask who administers the machine, where model and application files are obtained, how updates are reviewed, what network access exists and where backups land. Owning the hardware changes some boundaries; it does not remove operational responsibilities.

For hybrid systems, list the crossings explicitly. A document's original text, embeddings, snippets and generated summaries may travel separately. Embeddings are derived representations, not proof that the source information has become non-sensitive.

Start with the minimum authorised information needed for the task. Keep access and secrets outside prompts where possible, and test with synthetic data while the route is being designed. Legal and regulatory obligations depend on the actual processing context; this guide is a decision framework, not legal advice or a claim of compliance.

Technical references: Microsoft · Responsible AI; Microsoft · Small Language Models; OWASP · Sensitive Information Disclosure (2026); Hugging Face · Secure code execution. Links and source scope appear in the final section.

# Try the idea.

## INVENTED TEACHING EXAMPLE

### Three routes, different questions

Local: original → local preprocessing → local model → output, plus backups and logs. Cloud: approved information → remote endpoint → output, with service-specific retention and access. Hybrid: original remains local while selected excerpts or derived representations cross a boundary. No route is automatically certified private.

# Go deeper.

## Map derived copies

Inventory snippets, embeddings, summaries, logs and caches as well as original files. A transformed representation can retain sensitive information, so a claim about the source location is not a complete processing map.

## Verify the exact service

Retention, training use, regional processing and deletion differ by product, plan and configuration. This briefing provides questions rather than unsupported provider assurances.

## Minimise the task input

Supply the minimum authorised information required. Check whether combined innocuous fields identify someone or reveal a confidential event. A local machine still needs access controls and a backup policy.

# Keep the evidence close.

---

## Data boundary

A crossing between systems, owners or access scopes.

---

## Retention

How long information or derived copies remain stored.

---

## Derived representation

A transformed version such as an embedding or summary that can still reveal information.

---

## Anonymisation

A context-dependent standard of preventing identification; removing a name is not enough by itself.

---

## Sources · accessed 4 October 2026

---

[Microsoft · Responsible AI](#)

---

[Microsoft · Small Language Models](#)

---

[OWASP · Sensitive Information Disclosure \(2026\)](#)

---

[Hugging Face · Secure code execution](#)

Technical and security references support the mechanisms, not our fictional results. KAF's worksheets and examples are original editorial recommendations. No current provider pricing, comparative ranking or legal compliance is asserted. Verify the exact release, configuration and policy for your use.

---