

STARTER KIT

AI Governance Starter Kit for SMEs

Disclosure rules, approval workflows, prompt and asset registers,
and data boundaries sized for a Swiss SME under FADP -
governance without enterprise bureaucracy.

Governance that fits an SME

Enterprise AI governance assumes a compliance department. An SME has the same exposure - customer data in prompts, unreviewed output reaching customers, staff using unapproved tools - with none of the staff. This kit is a minimum viable policy: four documents and one habit, enough to answer 'what is our AI policy?' credibly in front of a customer or an auditor.

Document one: acceptable use

- Approved tools list - which AI services staff may use for work, and which accounts (company-managed, not personal)
- Data classes - what may never enter a prompt: personal data of customers, credentials, unpublished financials, contract text of third parties
- Output rules - AI output is a draft; external-facing text, code and analysis requires human review before use
- Disclosure - when customers must be told AI contributed (support replies, assessments, generated media)

Document two: the approval workflow

One named person approves: new tools, new use-cases for personal data, and anything AI-generated that leaves the company under the company name. The workflow is deliberately short - a request in writing, a yes/no with conditions, a line in the register. If approval takes more than a day, people route around it.

Document three: the register

A single spreadsheet with three tabs answers most audit questions: tools in use (name, purpose, account owner, data classes allowed), prompts and templates in production use, and AI-assisted deliverables shipped to customers. It takes ten minutes a week to maintain and is the difference between 'we have governance' as a claim and as a fact.

Document four: data boundaries

- Under FADP (the Swiss data protection act), personal data fed to an external AI service is a disclosure to a third party - document the legal basis or do not send it

- Prefer tools with contractual no-training terms for anything involving customer data
- Self-hosted inference is an option for sensitive workloads - it removes the third-party question entirely, at the cost of operating it properly
- Log retention: decide how long prompts and outputs are kept and who can see them

The one habit

Quarterly, thirty minutes: review the register, confirm the tool list still matches reality, and re-issue the acceptable-use page to all staff. Governance that is not rehearsed is decoration.

- ☐ Acceptable-use page written and distributed
- ☐ Named approver for new tools and data uses
- ☐ Register exists: tools, prompts, shipped deliverables
- ☐ Data classes defined; customer data rules explicit
- ☐ Quarterly review scheduled

This document is part of the UHL resources collection. Direct download and related work: <https://uhl-services.ch/en/resources/>

