PRACTICAL GUIDE – NOVEMBER 2026 EDITION

AI IN BUSINESS

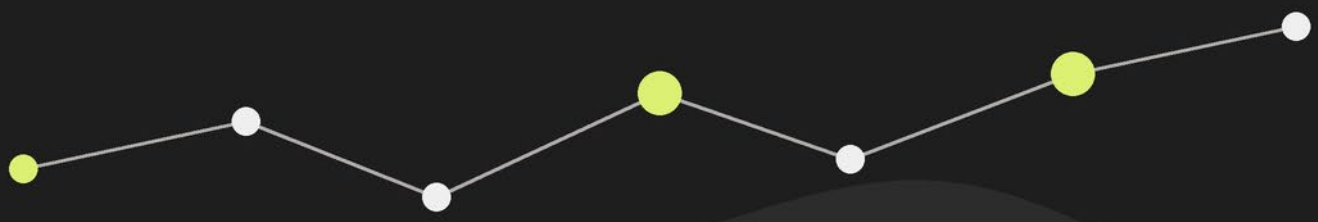
Turning AI into value, without losing control

The practical guide to useful, secure and controlled adoption

November 2026 edition • Switzerland

BUSINESS VALUE · DATA · SECURITY · GOVERNANCE

awsmtech.ch | +41 (0) 22 552 60 70 | mail@awsmtech.ch



NOTE TO THE READER

This white paper is intended for executives, business leaders, IT managers, security managers and compliance functions in Swiss organisations. It offers a decision-making framework: it is neither legal advice, certification of compliance nor a promise of performance.

Artificial intelligence is evolving rapidly. Features, contracts, processing locations, subprocessors, models and obligations must be checked again when selecting or deploying a solution. The examples are deliberately simplified and must be adapted to each organisation's context, sector and risk appetite.

> **THE AWSMTECHVIEW**
> The challenge is no longer simply to provide access to AI. It is to organise its use around data, processes, security, human accountability and measurable value.

How to use this guide

1. Start with the executive summary and the ten leadership decisions.
2. Use the data matrix and checklist as operational tools.
3. Apply the AWSMTECH 6D method to a real, limited and measurable use case.
4. Have the relevant legal, contractual and sector-specific matters validated before going live.

About this edition

This version includes a stronger Swiss perspective, a clear distinction between compliance, cyber security and governance, in-depth coverage of Shadow AI, a chapter on Microsoft 365 Copilot and enterprise AI, a dedicated framework for agents, a six-stage deployment method and a 90-day roadmap.

EXECUTIVE SUMMARY

AI is no longer a future project. It is already being used to write, search, summarise, translate, analyse, prepare for meetings, produce code and automate tasks. This spread creates quick gains, but also a gap between actual use and the company's ability to govern it.

01 — Value
Start with a business problem, not a tool. Define an observable outcome before choosing a technology.

02 — Data
Know what information enters, leaves, is retained and becomes accessible to AI.

03 — Security
Protect identities, permissions, devices, connectors, logs and the actions available to agents.

04 — Governance
Assign an owner, define permitted uses and provide for human validation, measurement and shutdown.

The ten leadership decisions

1. Appoint a business sponsor and an operational AI lead.
2. Map existing uses, including personal accounts and unapproved tools.
3. Define a simple policy: permitted, permitted subject to conditions, prohibited.
4. Classify data and link each category to the permitted AI environments.
5. Select use cases according to value, feasibility, risk and effort.
6. Clean up identities, access rights and sharing practices before introducing connected AI.
7. Separate assistance, recommendation and autonomous action; increase controls as the ability to act grows.
8. Measure the baseline, then quality, time, cost and risk after deployment.
9. Train employees in uses, limitations and their responsibility to verify results.
10. Scale up only those uses whose net value and risk control have been demonstrated.

THE QUESTION TO ASK NOW
Do you know where AI is used in your company, with what data, for what value and under whose control?

CONTENTS

01 From experimentation to control

02 Understanding AI without unnecessary jargon

03 Finding value: use cases and priorities

04 Shadow AI: making the invisible visible

05 Data: deciding what you can entrust to AI

06 Enterprise AI and Microsoft 365 Copilot

07 Architecture: choosing the right level of integration

08 Cyber security: protecting the whole system

09 The Swiss framework, FADP and European scope

10 Governance: rules, roles and evidence

11 AI agents: when software can act

12 The AWSMTECH 6D method

13 Roadmap for the first 90 days

14 ROI, adoption and change management

15 Digital sustainability and the 2027–2030 horizon

16 Checklist, charter and practical tools

17 Regain control before accelerating

CHAPTER 01

FROM EXPERIMENTATION TO CONTROL

AI becomes a leadership issue when it affects data, decisions and processes.

Artificial intelligence is no longer reserved for technical teams or large companies. It is already present in office software, customer relations, security tools, business platforms and browser-based services. For a business leader, the question is therefore no longer whether to take an interest, but how to use it in a useful, controlled and responsible way.

Behind the apparent simplicity of a conversation with an assistant lie fundamental decisions: what data can be entrusted to it? Where is it processed? What permissions does it use? Who checks the answer? What happens if the system takes action? How is the full cost calculated? And how can you ensure that the time freed up genuinely creates value?

Four inseparable pillars

Pillar	Leadership question	Warning sign
Business value	What outcome should improve, and how can it be measured?	Purchasing licences without a use case or baseline.
Data	What information is used, with what permissions and quality?	Historical oversharing or unreliable sources.
Security	How can access, processing, connectors and actions be protected?	An agent or extension has more permissions than necessary.
Governance	Who decides, checks, documents and takes responsibility for errors?	Nobody can stop the use or explain its purpose.

From a tool to an enterprise capability

A standalone chatbot produces answers. An enterprise capability connects AI to a process, authorised data, identities, rules, indicators and human accountability. The difference is not just about the model: it is about integration and operation.

KEY TAKEAWAY FOR LEADERS

AI accelerates what already exists. In a clear, controlled process, it can create an advantage. In a confused process, it accelerates inconsistencies.

CHAPTER 02

UNDERSTANDING AI WITHOUT UNNECESSARY JARGON

A powerful system remains probabilistic: a fluent answer guarantees neither truth nor suitability for the context.

What generative AI is

A language model calculates a likely response from instructions and context. It does not reason like a person, does not automatically know the company's internal policies and cannot always distinguish a fact from plausible wording. Its performance depends on the model's training, the context provided, the tools it can access and the controls surrounding its use.

Five limitations to bear in mind

Limitation	Consequence	Organisational response
Factual error	A credible answer may be wrong or out of date.	Require sources, define checks and retain human validation.
Incomplete context	The system is unaware of a business exception or local rule.	Provide a controlled document base and test exceptions.
Bias	The result may reproduce imbalances in the data or process.	Evaluate across groups and scenarios; document sensitive decisions.
Variability	Two similar phrasings may produce different results.	Use templates, rules, evaluations and acceptance thresholds.
Confidentiality	Data submitted may leave the expected perimeter.	Limit data, approve tools and govern connectors.

The system matters more than the model

In professional use, quality rarely comes from a model alone. It results from a combination of user identity, interface, instructions, reference documents, business rules, deterministic calculations, filters, logging, supervision and a review procedure. Replacing a model without reviewing this whole system may do little to improve the result.

RAG, rules and calculations

A retrieval-augmented generation (RAG) architecture selects authorised documents before asking the model to respond. It can improve relevance and traceability, provided that sources are up to date, properly classified and accessible to the right people. Critical calculations, pricing rules and eligibility conditions should remain deterministic wherever possible: AI explains or assists, but a controlled rule decides.

DESIGN PRINCIPLE

Do not ask an instruction to compensate for weak architecture. Permissions, rules, tests and supervision are stronger controls than a simple "be careful".

CHAPTER 03

FINDING VALUE: USE CASES AND PRIORITIES

The right project is not the most spectacular. It is the one that improves real work with manageable risk.

Four types of value

Accelerate
Search, synthesis, preparation, first-draft writing, classification and extraction.

Improve
Consistency checks, document coverage, standardisation and service quality.

Augment
Contextual assistance, scenario analysis and access to expertise when needed.

Transform
End-to-end process redesign, with orchestration, rules and controlled actions.

Examples for sensitive professions

Sector	Lower-exposure use case	Use case requiring strong controls
Law firms	Summarising public sources, preparing outlines, searching a validated knowledge base.	Analysing client files, professional secrecy, generating legal submissions or decisions.
Fiduciary firms	Preparing document checklists, explaining procedures, classification.	Analysing salaries, personalised tax advice, reconciliations and automatic accounting entries.
Private clinics	Internal documentation, preparing materials, administrative assistance.	Health data, triage, clinical recommendations or patient communication.
Private banks	Monitoring public information, preparing for meetings, document synthesis.	Profiling, credit, advice, transactions or decisions affecting a client.
Notaries	Letter templates, finding approved clauses, completeness checks.	Deeds, identities, inheritance matters, automatic transfers to business systems.
NGOs / institutions	Translation, knowledge capture, summaries of public reports.	Beneficiary data, sensitive locations, allocation decisions.

The priority matrix

Assess each idea against four dimensions: value, feasibility, risk and effort. A quick win combines visible value, available data, limited authority to act and simple validation. A major project warrants a longer pilot. An appealing idea without reliable data or an owner, or with disproportionate risk, should be deferred or abandoned.

AN ALTERNATIVE MAY BE BETTER
Always compare AI with a rule, a form, a process improvement, a search engine, conventional automation or better training. The simplest solution is often the most robust.

CHAPTER 04

SHADOW AI: MAKING THE INVISIBLE VISIBLE

The main risk is not always the AI the company deploys, but the AI its employees already use without a common framework.

A company's first AI project is not necessarily launched by management. It may begin when an employee pastes a client email into a public service to improve its wording, uploads a spreadsheet for analysis, activates a meeting assistant, installs a browser extension or connects an agent to their mailbox.

These uses can deliver immediate value. But when they fall outside the inventory, contracts, access rules and support arrangements, they form Shadow AI: a layer of tools, accounts and data flows that the organisation can neither manage nor stop.

Six questions for mapping AI use

1. Which AI tools and features are actually used, including those in software already licensed?
2. Do users sign in with a work account or a personal account?
3. What data is entered, uploaded, recorded, transcribed or synchronised?
4. Which assistants, extensions, connectors and agents can access email, files or applications?
5. Which contractual terms, reuse options and retention policies apply?
6. Which uses influence a decision, an external communication or an action in the information system?

A ban is not enough

A blanket ban can drive use towards personal accounts, reduce visibility and hinder legitimate uses. A more effective response combines an approved environment, a short policy, practical alternatives, targeted training and detection mechanisms. The aim is to make the right behaviour easier than bypassing the rules.

KEY MESSAGE

Before prohibiting or permitting AI, you need to know what exists. Mapping Shadow AI is the starting point for credible governance.

Minimum deliverable

Item	Expected content
Tool register	Name, supplier, version/offering, owners, user groups.
Data flows	Data types, sources, destinations, retention, subprocessors.
Risk level	Sensitivity, scope, automation, impact on people and operations.
Decision	Permitted, permitted subject to conditions, pilot, suspended, prohibited.
Action plan	Owner, corrective measure, date and evidence of closure.

CHAPTER 05

DATA: DECIDING WHAT YOU CAN ENTRUST TO AI

A useful policy turns information classification into a practical decision for the user.

Three questions before every use

1. How sensitive is the information: public, internal, confidential, personal, subject to secrecy obligations or critical?
2. Are the tool and account approved for this data category and purpose?
3. Is the amount of information submitted strictly necessary to achieve the intended result?

Operational matrix

Category	Examples	Default rule	Minimum controls
Public	Website, brochures, published offers.	May be used in an approved tool.	Check copyright, accuracy and brand communications.
Internal	Procedures, minutes, non-sensitive plans.	Work account and approved environment.	Identity, logging, retention and result validation.
Confidential	Contracts, prices, strategy, client information.	Prior assessment of the use case and tool.	Minimisation, permissions, encryption, contractual clauses, tests.
Personal / sensitive	HR, health, finance, third-party data.	Only where the purpose, legal basis and safeguards have been validated.	DPIA where risk is high, transparency, human control, strict limits.
Secret / critical	Professional secrets, trade secrets, credentials, keys.	Specifically approved environment; otherwise prohibited.	Isolation, least privilege, supervision, traceability and incident procedure.

Minimise before submitting

The best data to protect is often data the system does not need to receive. Remove unnecessary identifiers, reduce the document to the relevant passage, replace real examples with synthetic cases and avoid attaching an entire file to obtain an answer about a single clause.

Quality and permissions

Connecting AI to a document base does not make that base reliable. The company must define authoritative sources, remove duplicates, identify obsolete versions, correct oversharing and assign an owner. Search AI makes historical permissions more visible; it does not correct them on its own.

A SIMPLE RULE

AI should never receive more data or more permissions than necessary for the task, for no longer than the task requires.

CHAPTER 06

ENTERPRISE AI AND MICROSOFT 365 COPILOT

The difference is not just about the model: identity, permissions, data, contracts and administration determine the level of control.

Four levels to distinguish

Level	Description	Point to watch
Public service	Online tool, sometimes accessed with a personal account.	Terms of use, data reuse, lack of central administration.
Work account	Contractual offering administered for a team or organisation.	Configuration, retention, administrative roles and permitted uses.
Integrated Copilot	Assistant connected to applications and accessible data in Microsoft 365.	The quality of permissions, sharing and classification becomes crucial.
Business agent	System connected to internal sources, tools or workflows.	Action permissions, connectors, third-party data, logging and revocation.

What Microsoft 365 Copilot changes

Microsoft states that Copilot uses models, Microsoft Graph and Microsoft 365 applications, and only presents users with content to which they already have at least read access. Prompts, responses and data retrieved through Microsoft Graph are not used to train Copilot's foundation models. These safeguards do not remove the company's responsibility to correct existing access rights, decide which features to enable and govern interaction retention.

Before deployment

1. Check accounts, multifactor authentication, privileged roles and devices.
2. Review SharePoint sites, Teams teams, anonymous links, groups and external sharing.
3. Define sensitive information and apply suitable labels, data loss prevention rules and retention policies.
4. Choose a pilot group, use cases, success criteria and support arrangements.
5. Check integrated agents and applications, their permissions, terms and sources.
6. Train users to cite, verify, minimise and report problematic results.

POINT TO WATCH

AI does not necessarily create a data governance problem. It may reveal, amplify or accelerate a permissions problem that already existed.

Do not confuse protection with relevance

An enterprise environment may offer strong security and privacy commitments while still producing an inaccurate answer. Data protection addresses the processing framework; business validation addresses quality and decision-making. Both controls are necessary.

CHAPTER 07

ARCHITECTURE: CHOOSING THE RIGHT LEVEL OF INTEGRATION

From a standalone chatbot to a connected agent, usefulness increases with access to context — and the need for control increases with it.

Three deployment approaches

Approach	When to use it	Benefits	Constraints
Managed service	Standard office tasks, rapid start-up.	Administration, updates, integration and supplier support.	Contractual dependency, functional scope, residency and subprocessors to verify.
Cloud platform / API	Bespoke business use cases, integration with several systems.	Control over the interface, rules, sources and evaluations.	Architecture, security, operations, variable costs and skills.
Dedicated deployment	Data or constraints that warrant stronger isolation and control.	Greater control over hosting and certain data flows.	Cost, performance, maintenance, updates, supervision and scarce skills.

Reference architecture



Controls must follow the entire flow: authentication, authorisation, data filtering, document retrieval, model calls, tool execution, response generation and retention of audit trails. A review limited to the model supplier overlooks most operational risk.

Seven architecture questions

1. What outcome and service level are expected?
2. What data enters, leaves, is retained or reused?
3. Does the system suggest, recommend or act?
4. What volumes, latency, availability and costs must be guaranteed?
5. Who operates, supervises, updates and responds to incidents?
6. How can the model, supplier or hosting be changed?
7. How will quality be tested after each system change?

REVERSIBILITY
Avoid locking the process into a single model. Separate data, rules, orchestration and the model supplier wherever possible.

CHAPTER 08

CYBER SECURITY: PROTECTING THE WHOLE SYSTEM

Security starts before the first prompt and continues after the response, including any actions triggered.

Main threats

Threat	Example	Priority control
Data leakage	Client document copied into an unauthorised tool.	Approved tools, minimisation, DLP, training and monitoring.
Oversharing	AI retrieves a file to which the user had excessively broad access.	Permissions review, SharePoint/Teams governance, labels.
Prompt injection	An instruction hidden in a document diverts the agent.	Isolate untrusted content, limit tools, validate actions.
Account compromise	An attacker uses a connected assistant with the victim's permissions.	Phishing-resistant MFA, conditional access, detection and revocation.
Supply chain	A connector or subprocessor handles unexpected data.	Due diligence, contractual clauses, inventory, integration controls.
Cost or looping	An agent repeats calls or executes an action at scale.	Quotas, limits, budgets, delays, kill switch.

Compliant does not automatically mean secure

Compliance determines the framework within which data and processing may be used. Cyber security protects identities, systems and information against unauthorised access, modification or loss. Governance defines uses, roles and responsibilities. A tool can meet contractual commitments yet be poorly configured; a robust architecture can support inappropriate processing. All three dimensions must be addressed together.

Basic controls

- Inventory of systems, models, agents, extensions, sources and subprocessors.
- Work identity, strong authentication, least privilege and regular reviews.
- Information classification, data loss prevention and risk-based encryption.
- Separate development, test and production environments; synthetic data where possible.
- Usable logs, alerts, secrets management, backups and response procedures.
- Security evaluations covering injection, exfiltration, malicious content and tool misuse.

DEFENCE IN DEPTH
No content filter can replace authentication, permissions, isolation, limits on actions and supervision.

CHAPTER 09

THE SWISS FRAMEWORK, FADP AND EUROPEAN SCOPE

For a Swiss company, the starting point is the FADP; other frameworks may also apply depending on the activities, individuals and markets concerned.

The Swiss principle

The Federal Data Protection and Information Commissioner (FDPIC) reiterates that the FADP applies directly to data processing involving AI. The organisation must examine, in particular, transparency, purpose, proportionality, security, individual rights, automated decisions and, where risk is high, the need for a data protection impact assessment.

Compliance questions to document

1. What is the purpose, and why is AI necessary or proportionate?
2. Which categories of data and individuals are involved?
3. Who acts as controller, processor or subprocessor?
4. Where is the data processed, and what safeguards govern transfers?
5. What information is provided to individuals, and how can they exercise their rights?
6. Does the result produce an automated decision or significantly influence an individual?
7. Which high risks require an impact assessment, consultation or enhanced measures?
8. What retention period, traceability and incident procedure are planned?

When the GDPR or AI Act may matter

The GDPR may apply to a Swiss company depending on its establishments, its offering of goods or services and its monitoring of individuals in the European Union. The AI Act follows a risk-based approach and may affect operators established outside the Union where the regulation's territorial conditions are met. It applies progressively: certain prohibitions and AI literacy obligations have applied since 2025; the regulation became broadly applicable in August 2026, while some obligations relating to high-risk systems have later deadlines. A specific legal assessment remains necessary.

Sector-specific responsibilities

Professional secrecy, employment law, healthcare obligations, financial regulation, retention rules, contractual requirements and commitments to funders or beneficiaries may impose additional limits. The absence of sector-specific AI legislation does not remove existing obligations.

GUIDING PRINCIPLE

Using AI does not remove any of the company's existing responsibilities for its data, decisions and professional obligations.

CHAPTER 10

GOVERNANCE: RULES, ROLES AND EVIDENCE

Lightweight but explicit governance is better than an ambitious policy that cannot be applied.

The decision-making model

Role	Responsibility
Leadership / sponsor	Set the ambition, risk appetite, priorities and resources.
Business owner	Define the outcome, process, exceptions and validation.
IT / architecture	Choose the integration, operate the service, ensure reversibility and performance.
Security	Assess threats, identities, permissions, connectors, logs and incidents.
Data protection / legal	Determine roles, purposes, safeguards, contracts, rights and impacts.
HR / training	Define skills, support, usage rules and effects on work.
Users	Follow the framework, verify results, report errors and incidents.

The AI register

Every significant use must have a record: purpose, owner, users, data, suppliers, model or service, sources, connectors, autonomy level, controls, metrics, risks, contractual basis, review dates and continuation decision. The register is not a static inventory: it supports audits, model changes and incident response.

Three levels of human control

Level	Application	Requirement
Human in the loop	The decision or action awaits human validation.	Appropriate for external communications, sensitive data and actions with an impact.
Human on the loop	The system acts within limits and a human supervises.	Thresholds, alerts, sampling, ability to stop and reviews.
Human out of the loop	The system acts without operational validation.	Reserved for low-impact, reversible, thoroughly tested and capped actions.

Evidence to retain

- Authorisation decision and named owner.
- Assessment of data, contracts, risks and compensating controls.
- Test datasets, results, thresholds and exception handling.
- System versions, changes, incidents and corrective actions.
- Training, communication and acceptance of rules by the groups concerned.

PROPORTIONATE GOVERNANCE

The more a use case processes sensitive data, influences an individual or can act within the system, the greater the requirements for evidence and control.

CHAPTER 11

AI AGENTS: WHEN SOFTWARE CAN ACT

An incorrect answer is content to correct; an incorrect action can change a system, commit the company or affect a third party.

An assistant provides a suggestion. An agent can use tools, query several sources, chain steps together and perform an action: create a ticket, update a CRM, send a message, produce a document, trigger a workflow or call an API. The essential change is not the claimed autonomy, but the ability to act with real permissions.

The minimum scope to define

1. Authorised objective and success criteria.
2. Accessible sources and the trust level of each source.
3. Available tools and actions, with parameters and limits.
4. Identity used by the agent and associated privileges.
5. Mandatory human validations and escalation thresholds.
6. Logs, alerts, retention and ability to reconstruct events.
7. Quotas, budgets, operating hours, environments and shutdown mechanism.
8. Revocation, correction and rollback procedure.

Classifying actions

Class	Examples	Policy
Read	Search, summarise, compare.	Access according to the user's permissions; cited sources and verification.
Prepare	Pre-populate, draft, propose an update.	Human validation before publication or writing data.
Reversible write	Create a ticket, classify a document, update a field.	Limited scope, log, test, threshold and undo capability.
Binding action	Send, approve, pay, delete, change a permission.	Strong validation, separation of duties; often prohibited by default.

Specific tests

- Malicious instructions in documents or pages consulted.
- Conflict between system instructions, user instructions and external content.
- Tool failure, partial response, duplicate, loop or resumption after interruption.
- Attempt to access a source or action outside the authorised scope.
- Silent degradation after a change to the model, connector or data.

GOLDEN RULE

The greater AI's ability to act, the more essential least privilege, traceability, reversibility and human validation become.

CHAPTER 12

THE AWSMTECH 6D METHOD

Six stages to move from fragmented use to a useful, secure and scalable capability.

1. Discover
Inventory uses, expectations, tools and Shadow AI.

2. Diagnose
Assess processes, data, identities, security and governance.

3. Define
Prioritise use cases and design controls.

4. Deploy
Pilot in a real but controlled setting.

5. Demonstrate
Measure value, quality, adoption, cost and risk.

6. Develop
Scale up, extend, correct or stop.

1 — Discover

Short interviews, observation of work and a technical inventory reveal pain points, duplicate data entry, time-consuming searches, parallel documents, existing uses and team expectations. The deliverable is a map of opportunities and risks, not a product list.

2 — Diagnose

The assessment covers process maturity, data quality and ownership, identities, permissions, architecture, obligations, skills and operational capability. An unstable or unmeasured process is simplified first; AI is introduced afterwards if it adds net value.

3 — Define

Use cases are ranked by value, feasibility, risk and effort. For each one, AWSMTECH documents the purpose, autonomy level, data, target architecture, controls, success criteria and exit decision. Technology selection follows the business need.

4 — Deploy

The pilot uses representative cases, includes exceptions and involves designated users. Access is restricted, logs are enabled, and support and incident procedures are tested. A pilot is not a demonstration: it must reproduce real constraints without exposing the company to uncontrolled risk.

5 — Demonstrate

The outcome is compared against a baseline: time, lead time, quality, correction rate, satisfaction, errors, cost and incidents. Failures are documented. Benefits are calculated after integration, support and control, not solely from the apparent speed of generation.

6 — Develop

Proven uses are scaled up with an owner, budget, supervision, reviews and a reversibility plan. Underperforming uses are corrected or stopped. Development also means training, adjusting policy and progressively enriching sources and controls.

THE 6D LOGIC
Discover → Diagnose → Define → Deploy → Demonstrate → Develop. A decision-making loop, not a race to deployment.

CHAPTER 13

ROADMAP FOR THE FIRST 90 DAYS

An initial short cycle should deliver visibility, a workable framework and proof of value — not a complete transformation.

Days 1 to 30 — Understand and frame

Action	Deliverable
Appoint a sponsor, lead and working group.	Mandate, roles and decision-making cadence.
Map tools, uses and data.	Initial register of Shadow AI and integrations.
Define an interim policy.	Rules: "permitted / subject to conditions / prohibited".
Identify 5 to 10 opportunities.	Use case records with value and risks.
Check the foundations.	State of identities, permissions, sharing and devices.

Days 31 to 60 — Prioritise and prepare

Action	Deliverable
Score use cases by value, feasibility, risk and effort.	Prioritised portfolio and reasoned decisions.
Choose one to three pilots.	Scope, users, data, architecture and controls.
Measure the baseline.	Time, quality, costs, errors and satisfaction before the change.
Train pilot participants.	Rules, methods, scenarios and reporting procedure.
Test security and operations.	Logs, access, support, shutdown and incidents.

Days 61 to 90 — Test and decide

Action	Deliverable
Run on real cases.	Results, exceptions, feedback and incidents.
Compare against the baseline.	Value, cost, quality and risk dashboard.
Correct or narrow the scope.	New version of the process and controls.
Decide.	Scale up, extend, reposition or stop.
Update governance.	Register, policy, training and quarterly plan.

SUCCESS CRITERION AT 90 DAYS

The company knows what is being used, has set understandable rules and can objectively demonstrate the value — or lack of value — of at least one use case.

CHAPTER 14

ROI, ADOPTION AND CHANGE MANAGEMENT

The number of licences measures spending. Value is measured in real work, after checking quality and risk.

Measure more than usage

A process previously taking 45 minutes and reduced to 15 minutes theoretically saves 30 minutes per occurrence. Repeated 20 times a month by five people, it represents a potential 50 hours a month. But this calculation is only useful if quality is maintained, corrections do not consume the gain, the time freed up is reinvested and the full cost remains below the value created.

Measurement dashboard

Dimension	Possible indicators
Productivity	Time per case, lead time, volume, queues, capacity.
Quality	Error rate, completeness, consistency, rework, client satisfaction.
Adoption	Active users, useful frequency, variety of cases, autonomy.
Risk	Incidents, exposed data, unverified outputs, reversed actions.
Economics	Licences, integration, operations, support, training, exit cost.
People	Cognitive load, satisfaction, skills, quality of oversight.

A simple formula

MONTHLY NET VALUE

(hours genuinely freed up × value of a useful hour) + quality gains + avoided risks − full monthly cost.

Quality gains and avoided risks must be estimated cautiously. Assumptions are documented and reviewed after the pilot. A transparent presentation distinguishes theoretical capacity, observed gains and value actually captured.

Managing change

- Start with tedious, repetitive tasks, without promising that work will disappear.
- Involve business experts in tests, rules and quality criteria.
- Train people in verification and escalation, not just prompt writing.
- Create a community of practice and approved examples for each business function.
- Acknowledge failures and retire use cases that shift the workload rather than reduce it.
- Monitor effects on roles, learning, accountability and client relationships.

A SIGN OF MATURITY

A mature organisation can stop a popular but unprofitable use case and invest in a less conspicuous one that improves a process sustainably.

CHAPTER 15

DIGITAL SUSTAINABILITY AND THE 2027–2030 HORIZON

AI should be sized for usefulness: the largest model, longest context and most autonomous agent are not always the best choices.

Principles of digital sustainability

1. Launch only uses linked to a measurable business benefit.
2. Choose the smallest adequate model and limit the length of contexts and responses.
3. Cache, share resources and avoid unlimited agent loops.
4. Reduce transferred data and the frequency of unnecessary calls.
5. Consider infrastructure, lifespan, location and electricity mix in decisions.
6. Measure impact per useful unit: case processed, error avoided, hour genuinely freed up.

Three scenarios to prepare for

Scenario	What changes	Useful preparation now
AI integrated everywhere	Assistance features become standard in software.	Govern enabled features, permissions and adoption.
Specialised agents	Agents execute limited, controlled sequences.	Modularise processes, define actions and strengthen logs.
Demand for evidence	Clients, regulators and leaders demand more documentation.	Maintain a register; retain tests, decisions, versions and named owners.

What will change more slowly

Accountability, data quality, business understanding, client trust, identity security and the need for measurement will remain crucial. Models will evolve rapidly; organisational foundations will progress at the company's own pace. That is why the most enduring investment often focuses on data, processes, permissions and skills.

DO NOT BET ON A PREDICTION

A prepared company tests, measures, learns and changes solutions when results justify it. It does not need to predict 2030 precisely.

CHAPTER 16

CHECKLIST, CHARTER AND PRACTICAL TOOLS

Simple tools to turn principles into verifiable decisions.

Pre-production checklist

✓	Area	Criterion
☐	Purpose	The expected outcome and business owner are defined.
☐	Alternative	A non-AI solution has been compared.
☐	Data	Sources, categories, quality, permissions and retention are documented.
☐	Suppliers	Contracts, subprocessors, processing locations and exit arrangements have been examined.
☐	Security	Identity, least privilege, secrets, logs, tests and incident arrangements are ready.
☐	Quality	Test datasets, thresholds, sources and exception handling are defined.
☐	People	Validation, supervision, review mechanisms and responsibilities are explicit.
☐	Agents	Actions, quotas, validations, revocation and rollback have been tested.
☐	Compliance	The FADP, sector-specific requirements and any EU scope have been assessed.
☐	Value	The baseline, indicators, full cost and decision date are set.
☐	Adoption	Training, support, communication and a feedback mechanism are in place.
☐	Reversibility	Export, replacement, archiving and shutdown are practicable.

AI charter — essential clauses

A concise charter should be understood in less than ten minutes. It specifies:

- the authorised tools and accounts;
- the data categories that are prohibited or require validation;
- the obligation to minimise, verify and cite sources where relevant;
- the prohibition on presenting an AI output as validated without checking it;
- the rules for external communications, decisions and automated actions;
- the channel for reporting errors, leaks, problematic content and unknown uses;
- the consequences of non-compliance and the exception process.

RECOMMENDED WORDING

AI can help prepare, search, summarise and propose. It does not transfer responsibility away from the user or the company. Every use must comply with information classification, approved tools and the required level of validation.

Use case record

Field	Information to provide
Problem / outcome	Process concerned, pain point, indicator and target.
Users	User group, skills, volume and context.
Data	Sources, categories, quality, permissions, retention.
Solution	Service, model, architecture, tools, agent, supplier.
Controls	Rules, tests, validation, logs, security, incidents.
Decision	Pilot, date, budget, owner, criteria for continuing or stopping.

GLOSSARY

Term	Definition
AI agent	System capable of pursuing an objective across several steps and using tools to act.
DPIA	Data protection impact assessment, required in certain high-risk cases.
API	Interface enabling applications to exchange requests and results.
Hallucination	A false or invented answer, plausibly worded by a model.
Generative AI	System producing new content: text, image, sound, video, code or a combination of formats.
LLM	Large language model trained to process and generate text.
Least privilege	Principle of granting only the permissions required, for only as long as necessary.
Prompt	Instruction, question or context submitted to AI.
Prompt injection	Malicious instruction, sometimes hidden in content, intended to divert a system.
RAG	Method that retrieves documents before providing them to the model to support an answer.
Shadow AI	Use of AI tools or features without sufficient visibility, validation or governance.
Traceability	Ability to reconstruct a system's data, versions, decisions and actions.

REFERENCES AND METHODOLOGY

The references below prioritise official sources and practical frameworks. Supplier content and regulatory timetables evolve and must be checked again before any decision. Last editorial review: 11 November 2026.

1. FDPIC — AI and data protection — Direct application of the FADP, transparency and high risks.

2. FDPIC / NCSC — Generative AI at work — Awareness of data protection and security risks in the workplace.

3. European Commission — AI Act — Risk-based framework and implementation timetable.

4. Microsoft Learn — Data, Privacy and Security for Microsoft Copilot — Organisational data, permissions, training, retention and agents.

5. Microsoft Learn — How data is protected and audited — Labels, encryption, SharePoint, OneDrive, auditing and retention.

6. NIST — AI Risk Management Framework 1.0 — AI risk management framework.

7. OWASP — GenAI Security Project — Security risks of applications based on language models.

8. International Labour Organization — Generative AI and Jobs — Task exposure and the transformation of work.

9. International Energy Agency — Energy and AI — Energy, data centres and AI uses.

Editorial methodology

The content brings together principles of governance, security, data protection and change management. It distinguishes suppliers' published commitments from the controls that remain the client's responsibility. The examples are generic and do not guarantee an identical result in another organisation.

CHAPTER 17

REGAIN CONTROL BEFORE ACCELERATING

The aim is not to deploy more AI. It is to use AI where it creates value, under appropriate control.

Your company probably already uses artificial intelligence. The real question is: do you know where it is used, with what data, for what value and under whose control?

AI can accelerate access to information, improve productivity and transform some processes. But its value does not come solely from the chosen model. It depends on data quality, a secure environment, relevant use cases, employee engagement and the governance put in place.

THE AWSMTECHAPPROACH
Map. Prioritise. Secure. Experiment. Measure. Scale up.

From managed IT to Managed Intelligence

AWSMTECH supports sensitive professions in French-speaking Switzerland with managed IT, cyber security, cloud, workplace and outsourced IT leadership services. Managed Intelligence extends this mission: turning technologies, data and artificial intelligence into sustainable, controlled business capabilities.

Support options

Engagement	Expected outcome
AI & Shadow AI assessment	Inventory, risk map, priority rules and use case portfolio.
Microsoft 365 Copilot readiness	Review of identities, permissions, sharing, data, security and pilot.
Prioritisation workshop	Value / feasibility / risk / effort matrix and launch decisions.
6D pilot	Design, controlled deployment, ROI measurement and decision on scaling up.
Governance & security	Charter, register, controls, agents, logging and incident arrangements.

Let's discuss your first use case

TRUSTED MANAGED IT FOR SENSITIVE PROFESSIONS

awsmtech.ch | +41 (0) 22 552 60 70 | mail@awsmtech.ch

Geneva • Martigny

© 2026 AWSMTECH (Switzerland) LTD. All rights reserved.