

A Gate-Based Decision-Support Web Application for SME AI Adoption: From an Evidence-Based Governance Framework to Implementable Tooling

Thomas Joswig*

White Paper · Version 1.0 · July 1, 2026
München, Germany

Abstract

Small and medium-sized enterprises (SMEs) increasingly must decide *how* to host AI—cloud, on-premises, hybrid, or AI-as-a-Service—under tight resources and rising audit, risk, and compliance expectations. Existing standards, risk frameworks, and cloud playbooks help a firm *govern* or *assess* AI, but none offers an ordered, SME-feasible pathway for the hosting *decision* itself. This white paper turns an empirically grounded doctoral framework¹ into a practical decision-support application—*ToJo Navigator*, whose first module, *AI Governance* (German: *KI Governance*), addresses the AI *hosting* decision as its first use case—that settles *feasibility and governance before cost* and produces an auditable *Decision Record*. The underlying evidence—expert interviews across seven German sectors and a structured practitioner survey—shows that feasibility and capability, not cost, drive the choice.

*Thomas Joswig is a doctoral researcher at Signum Magnum College, Malta, and founder of *ToJo Technologies*. This white paper is based on his dissertation *Governing AI Infrastructure Decisions in SMEs: A Gate-Based Framework Grounded in the German Mittelstand*, which draws on semi-structured expert interviews across seven German sectors and a structured practitioner survey. Contact: thomas.joswig@tojo-technologies.de.

¹Based on Joswig (2026).

Executive summary

- **The problem.** SMEs must decide how to host AI (cloud, on-premises, hybrid, or AI-as-a-Service) under tight resources while satisfying audit, risk, and compliance expectations. Existing standards, risk frameworks, and cloud playbooks do not provide an ordered, SME-feasible decision pathway.
- **The evidence.** Expert interviews across seven sectors and a structured practitioner survey show that *feasibility and capability*, not cost, drive the choice: Technology (30.6 %) and Organisation (21.4 %) carry more than half the decision weight; cost enters only third.
- **The solution.** A three-gate decision tool: Gate 1 screens which options are *feasible* (TOE: technology, organisation, environment), Gate 2 confirms they are *governable and auditable* (COBIT, an IT-governance standard), and Gate 3 *ranks* the survivors (an extended Balanced Scorecard, BSC, scored with the AHP pairwise-weighting method, incl. TCO/breakeven).
- **What you get.** Every run produces an auditable *Decision Record* (with an evidence vault, SOP pack, and KPI plan) from *ToJo Navigator* (its AI Governance module) — a decision that is *cheaper to get right* (no spend on options that cannot work), *easy to defend* to regulators and the board, and *light to run* (four lightweight routines, not a governance department).
- **Why it is credible.** The criteria, weights, and gate sequence were scientifically derived in a doctoral dissertation, not asserted by a vendor.

The bottom line for the CEO

AI hosting is now a board-level decision: get it wrong and you pay twice (rework, lock-in), risk a compliance finding, or watch the project stall. This framework flips the usual order — it settles *can we do it* and *can we govern it before what does it cost*, so money is never committed to an option that cannot legally or technically work. The result is a decision you can make in a single working session, sign off with confidence, and *evidence on demand* to your board, your customers, or a regulator — backed by a documented record, not a gut call. It is built for a lean firm: no governance department, just four lightweight routines with clear ownership. In one line: **a faster, cheaper-to-get-right, auditable AI hosting decision.**

Who should read this. This white paper is written for SME owners and managing directors weighing an AI hosting decision; IT and data leads responsible for feasibility, security, and operations; compliance and risk officers who must evidence governance; and consultants or managed-service providers advising SMEs. Technical implementation detail is kept to the minimum needed to commission or build the tool.

What it does for you. *For the CEO or owner:* a faster, defensible decision — reached in a single session — that avoids the expensive wrong turns (rework, vendor lock-in, a failed audit) which quietly sink AI projects, keeps AI spending tied to business goals, and leaves a signed record that protects the business if the board, a customer, or a regulator ever asks why the choice was made. *For the CTO or IT lead:* a feasibility screen that rules out integration dead-ends early, with vendor-exit and incident rollback built in, so you are not locked in or left without a fallback. *For the compliance, risk, or security officer:* an auditable evidence trail showing how and why each decision was made, mapped to GDPR, the EU AI Act, and NIS2 — prepared before anyone asks.²

²This tool provides decision *support*, not legal advice. It helps a firm evidence how and why a hosting decision was made, but it does not constitute, and cannot replace, a data-protection impact assessment (DPIA), an EU AI

1 Problem and Motivation

SMEs often struggle to operationalise AI governance and infrastructure decisions under resource constraints; surveys report that few SMEs reach deep digital integration, with cost, skills, and capacity the dominant barriers.³ The hosting choice reaches well beyond procurement into strategy, operating capacity, regulatory exposure, and governance. Standard cost-benefit tools handle this poorly, and existing guidance typically falls into three categories: (i) standards and management-system requirements for AI governance (e.g., ISO/IEC 42001), (ii) risk frameworks that structure outcomes and functions (e.g., NIST AI RMF), and (iii) cloud adoption frameworks that emphasise cloud transformation capabilities (e.g., AWS CAF and Microsoft CAF). These resources are valuable but frequently leave a practical gap for SMEs: *a stepwise, evidence-centred decision pathway that gates infrastructure options (cloud/on-prem/hybrid/AIaaS) before cost optimisation, and that produces auditable deliverables.*⁴

Why now. The decision is no longer one an SME can defer. The EU AI Act is phasing into application, NIS2 is being transposed across member states, and customers and insurers increasingly ask suppliers to evidence how their AI is governed. At the same time AI adoption is accelerating—a 2026 Bitkom study reports that around 41 % of companies in Germany now use AI, with SMEs catching up⁵—so firms are committing to hosting arrangements faster than their governance can catch up.

The cost of getting it wrong. A poorly sequenced hosting decision is expensive in ways that rarely show up in the initial price comparison. Choosing on cost first can mean migrating off an option that later fails a data-residency or security requirement (rework and double spend); locking into a provider with no tested exit (loss of leverage and continuity risk); or discovering, during an audit or incident, that the firm cannot show how a decision was made (compliance exposure and reputational damage). Many AI initiatives also stall not because the technology fails but because feasibility and ownership were never settled. Clearing feasibility and governance *before* cost is precisely what avoids these traps.

2 The Underlying Study: Evidence Base

Unlike vendor playbooks or consultancy checklists, the framework in this paper was *scientifically derived* within a doctoral dissertation⁶ and is empirically grounded rather than asserted: its criteria, their relative weights, and the gate sequence are outcomes of a documented, replicable research process, not design opinions. The evidence base has two strands. First, **semi-structured expert interviews across seven German sectors**, which surfaced the factors that actually drive AI hosting decisions and were distilled into eight decision criteria. Second, a **structured practitioner survey** that ranked those criteria by importance. The criteria were cross-mapped to the established COBIT and Balanced Scorecard frameworks, and the weights were checked for robustness; the full methodology is documented in the dissertation. What matters here is the result, not the machinery behind it: the decision turns out to be governed by feasibility and governance, not by cost.

Act conformity assessment, or professional legal and data-protection counsel, and it does not guarantee that any audit or regulatory review will be passed. Regulatory applicability (e.g. EU AI Act risk class, NIS2 scope) depends on the individual firm and use case.

³OECD 2025.

⁴ISO/IEC 2023; National Institute of Standards and Technology 2023; Amazon Web Services 2026; Microsoft 2025.

⁵Bitkom 2026b.

⁶Joswig 2026.

2.1 Empirical priority weights

Table 1 reports the aggregated AHP priority weights for the eight criteria. Two results drive the tool design. First, Technology and Organisation together account for more than half the total weight (51.9 %): infrastructure decisions are dominated by *feasibility* and *capability*, not by economics. Second, Economic Efficiency ranks only third (16.8 %), and Compliance/Security varies narrowly across sectors (10.6 %–11.7 %); firms treat regulatory readiness as a *baseline* to be met rather than as a competitive differentiator. This inverts the cost-first sequence assumed by most practitioner-oriented IT investment guidance and motivates a design in which cost optimisation is deferred until after feasibility and governance have been cleared. A third pattern is the *risk paradox*: Risk ranks last (3.1 %) not because it is unimportant—data breaches, model failure and vendor lock-in are highly salient in the interviews—but because it functions as a *non-negotiable constraint* rather than a tradeable factor. The framework therefore enforces risk as a non-compensatory precondition at Gate 2 (managed-risk objectives and the algorithmic-risk SOP) instead of as a weight at Gate 3, where its low share would otherwise let a cheap, capable but risky option win on points.⁷

Table 1: Empirical AHP priority weights for the eight SME decision criteria (aggregated across the practitioner survey).

Criterion	Weight (%)	Rank / tier
Technology	30.6	1 — dominant (feasibility)
Organisation	21.4	2 — capability
Economic Efficiency	16.8	3 — cost/TCO filter
Compliance/Security	11.0	4 — baseline
Market/Customer	7.4	5
Innovation	5.9	6
Personnel	3.9	7
Risk	3.1	8
Total	100.0	

How to read this table: a low weight does *not* mean a criterion is unimportant. Risk and Compliance/Security carry small weights precisely because the framework treats them as hard pass/fail preconditions at Gate 2 (an option that fails is excluded outright), rather than as factors to be traded off against cost — see the risk paradox below.

3 Contribution

This paper contributes a tool-ready specification of a **gate-based decision-support workflow** that:

- (a) applies **TOE-scoped feasibility screening** to determine an admissible option set;
- (b) enforces a **COBIT-anchored governance and evidence minimum** before optimisation;
- (c) operationalises strategic alignment using an **extended BSC KPI plan** weighted by the empirical AHP priorities; and

⁷This is a deliberate separation of two roles, not a double count. Table 1 reports the *elicited importance* of all eight criteria as surveyed. Risk and Compliance/Security *additionally* act as non-compensatory *floors* at Gate 2: an option that fails the risk or compliance minimum is excluded there and never reaches Gate-3 scoring, so a low Gate-3 weight can never let an unsafe or non-compliant option “buy back” admissibility. Among the options that clear those floors, all eight criteria then contribute to the compensatory ranking at their surveyed weights. The floor (pass/fail at Gate 2) and the weight (degree among survivors at Gate 3) are two distinct functions of the same criterion, not the same test applied twice.

- (d) evaluates economics (TCO/breakeven) and multi-criteria aggregation *within* the admissible set.

The workflow is operationalised as *ToJo Navigator*—a web application whose first module, *AI Governance*, addresses the hosting decision—which outputs an *auditable Decision Record*, *evidence vault*, *SOP pack*, and *implementation backlog*. The approach is designed to be SME-first: minimum viable governance (MVG) with clear ownership, lightweight evidence, and optional deep integrations.

4 From Framework to Tool: The Three-Gate Logic

The doctoral framework is a three-gate decision artefact in which the first two gates are *non-compensatory* (a hard constraint cannot be bought back by economic advantage) and the third is *compensatory* (trade-offs are weighed within the surviving set). Figure 1 shows the sequence; ToJo Navigator implements it as a series of guided steps, mapped in Table 2.




Figure 1: The Gate-Based SME AI Governance Framework. Gates 1 and 2 are non-compensatory screens that filter the option set; Gate 3 applies compensatory multi-criteria evaluation (with economics as a measurement instrument) within the survivors. Adapted from the underlying dissertation.

Table 2: Mapping from the conceptual three-gate framework to the ToJo Navigator application steps. The intake is a preliminary setup step, not a gate.

Conceptual gate	Lens	Application step
(preliminary)	Decision context	Intake: context setup
Gate 1 (non-compensatory)	Adapted TOE	Feasibility screening
Gate 2 (non-compensatory)	Right-sized COBIT 2019	Governance & evidence minimum
Gate 3 (compensatory)	Extended BSC + AHP	Evaluation: KPI plan, then scoring & aggregation (TCO/breakeven)

A **Gate** is an evidence-based checkpoint where named owners review required artefacts against pre-defined criteria and decide **Go**, **No-Go**, or **Defer with actions**. Gates 1 and 2 filter the option set {Cloud, On-Prem, Hybrid, AIaaS}; Gate 3 ranks and measures the survivors. Economic evaluation is a measurement instrument at Gate 3, not a primary selector.

5 Related Landscape: Standards and Commercial Tooling

5.1 Standards and frameworks

ISO/IEC 42001 specifies requirements for establishing, implementing, maintaining, and continually improving an AI management system.⁸ The NIST AI RMF provides a voluntary, technology- and sector-agnostic framework for managing AI risks.⁹ COBIT 2019 provides governance and management objectives for information and enterprise technology.¹⁰

5.2 The current SME solution landscape (2026 market scan)

A scan of the 2026 market shows a healthy ecosystem of capable, well-established tools that cluster into four families.¹¹ The comparison that follows is by *purpose*, not quality: each family is strong at the job it was built for, and the point is simply that none is built for the hosting *decision* itself:

- **AI governance / GRC platforms** — inventory, policy, and lifecycle oversight of AI systems an organisation has *already* deployed, including model monitoring, evaluation, and reporting. These are powerful, well-established platforms aimed primarily at enterprise and mid-market organisations, with pricing scaled accordingly (industry buyer’s guides indicate enterprise tiers in the six figures annually and mid-market tiers from the tens of thousands upward), which can make the lightest SME tier a different fit.¹²
- **Compliance-automation platforms** — certification readiness for standards such as ISO/IEC 42001, the EU AI Act, SOC 2, and ISO 27001, with pre-built controls and evidence collection. These are SME-accessible but are designed to prove *conformance* rather than to choose a hosting option.
- **AI readiness / maturity assessments** — diagnostic questionnaires from public bodies and platform providers that score preparedness across pillars such as strategy, data, infrastructure, and governance. These are SME-friendly and produce a useful readiness *score*, which is a different output from an auditable decision among hosting options.
- **Cloud adoption frameworks** — capability and migration guidance offered by the major cloud providers, which presupposes the cloud as the destination.

The proposed tool occupies a different position: a *decision-first* layer that selects and justifies a hosting option and can then hand off to the governance, compliance, and observability platforms above.

⁸ISO/IEC 2023.

⁹National Institute of Standards and Technology 2023; National Institute of Standards and Technology 2025.

¹⁰ISACA 2018; ISACA 2026.

¹¹Market scan compiled from public vendor documentation and industry buyer’s guides (accessed 2026); illustrative and non-binding, not an exhaustive market survey. Specific products and vendors are deliberately not named; the categorisations reflect each family’s typical positioning from public sources as of 2026, not an assessment of any individual product’s current capabilities, and may change.

¹²Modulos 2026.

5.3 Where current tools leave a gap

Mapping these tool families against the framework’s five distinguishing features (Table 3) shows that each family addresses one slice, governing, certifying, or assessing, but none is designed to perform the SME hosting *decision* itself: none sequences feasibility and governance ahead of economics, gates the hosting options for admissibility, and emits an auditable decision record weighted by empirical priorities. The table compares design *purpose*, not capability or quality: an entry of N simply marks a feature outside that family’s intended remit, and the families remain the natural place to govern, certify, or monitor AI once the hosting option has been chosen.

Table 3: Comparison by design *purpose* (not capability or quality): how current tool families relate to the framework’s five distinguishing features (2026 market scan; to the author’s knowledge). Y = yes, P = partial, N = outside that family’s intended purpose.

Tool family	(a)	(b)	(c)	(d)	(e)
AI governance / GRC platforms	N	N	P	P	N
Compliance-automation platforms	N	N	Y	N	N
AI readiness / maturity assessments	N	N	Y	N	N
Cloud adoption frameworks	N	N	N	N	N
Gate-based tool (this paper)	Y	Y	Y	Y	Y

(a) non-compensatory sequencing feasibility→governance→economics; (b) admissibility gating of hosting options (cloud/on-prem/hybrid/AIaaS); (c) right-sized for SMEs; (d) auditable hosting decision record; (e) empirically derived multi-criteria weights.

The market gap. Today’s tools help an SME *govern* AI it has already adopted, *certify* compliance, or *assess* readiness; they are simply not intended to help it *decide* how to host AI in the first place. The gate-based tool fills exactly that white space and is explicitly complementary to, not a replacement for, the platforms above, to which it can hand off once a hosting option is chosen. This reading is independently corroborated by a 2026 Bitkom industry guide on practical cloud sovereignty, which structures the same risk categories—organisational, technological, skills, economic with tested exit options, and political-regulatory—and explicitly anchors sovereignty as a governance criterion, yet stops at a *sovereignty-degree risk assessment* rather than an ordered hosting decision across cloud/on-prem/hybrid/AIaaS.¹³

5.4 Limitations of established frameworks for the SME hosting decision

Each established resource was designed for a purpose other than an ordered, SME-feasible *infrastructure hosting decision*, and the doctoral study identified a consistent set of gaps. Standard cost-benefit tools reduce the choice to a technical or financial trade-off, while enterprise governance frameworks import logic that SMEs cannot sustain. Table 4 summarises the shortcomings that the gate-based framework was designed to address; the analysis is grounded in the dissertation’s literature review and contrast with prior models.¹⁴

¹³Bitkom 2026a.

¹⁴National Institute of Standards and Technology 2023; ISO/IEC 2023; ISACA 2018; Amazon Web Services 2026; Microsoft 2025.

Table 4: Shortcomings of established frameworks for SME AI hosting decisions, as identified in the underlying doctoral study.

Resource	What it provides	Limitation for the SME hosting decision
Cloud adoption frameworks (AWS CAF, MS CAF)	Cloud transformation capabilities and migration guidance	Assume cloud as the destination; do not compare or gate hosting options (on-prem/hybrid/AIaaS); large-enterprise orientation
NIST AI RMF	Voluntary risk functions and outcomes	Sector- and technology-agnostic; provides functions but no ordered infrastructure decision pathway; not SME-specific
ISO/IEC 42001	AI management-system requirements	Continuous-improvement overhead is hard for resource-constrained SMEs to sustain; specifies <i>what</i> to manage, not <i>how</i> to decide hosting
COBIT 2019	Enterprise governance and management objectives	Designed for mature governance structures; illustrated mainly through large-enterprise cases; no SME right-sizing and no multi-criteria evaluation
TOE adoption studies	Adoption determinants (technology/organisation/environment)	Treat all determinants as compensatory weighted inputs; no non-compensatory feasibility filter and no governance backbone
Standard Balanced Scorecard	Strategic performance measurement	Compensatory across all perspectives; does not scope feasibility constraints or operationalise AI-specific compliance evidence
Cost-first / TCO practitioner guidance	Breakeven and total-cost-of-ownership analysis	Assumes a cost-first sequence, which the empirical weights contradict (Technology and Organisation dominate; economics is a third-stage filter)

The recurring pattern is that these resources are either *compensatory* (every weakness can be offset by a strength) or *single-lens* (technical, financial, or governance alone), and none sequences feasibility and governance as hard preconditions before economic comparison. This is the gap the gate-based framework closes.

6 The Implementation: ToJo Navigator and the AI Governance Module

The framework is delivered as *ToJo Navigator*, a guided web application developed by ToJo Technologies that walks an SME through a short intake and the three gates, assembling the supporting evidence as it goes. ToJo Navigator is designed as a *platform: AI Governance* (German: *KI Governance*) is its first module, and the same gate-and-evidence architecture is intended to carry further governance modules—for example AI Act conformity, data protection, NIS2, and information security—as an SME’s needs broaden. This paper concerns the AI Governance module and, within it, the AI *hosting* decision as its first use case. A working instance has been built and is running in pilot operation. It is built around five principles:

- **Auditability:** every decision ends in a record with linked evidence and a named sign-off.
- **Explainability:** the tool shows in plain language why each option is eligible, conditional, or excluded.
- **SME feasibility:** governance is kept to the lightest set that remains defensible.
- **Repeatability:** the same process is reused across use cases and revisited when circumstances change.
- **Fit with existing systems:** it connects to the tools an SME already uses rather than replacing them.

What every run produces (outputs at a glance)

- **Decision record** — the chosen hosting option with a plain-language rationale and a named sign-off.
- **Eligibility matrix** — each option (cloud / on-prem / hybrid / AIaaS) marked eligible, conditional, or excluded, with reasons.
- **Evidence vault** — the documents and logs that make the decision auditable, mapped to GDPR, the EU AI Act, and NIS2.
- **SOP pack** — four lightweight routines (algorithmic risk assessment; fairness & transparency monitoring; data & model change control; incident & rollback) with named owners.
- **KPI plan** — a short scorecard linking the chosen option to business goals and a review cadence.

6.1 What each gate evaluates

The application opens with a short **intake** that captures the decision context: the use case and business function, the deployment scope, the data involved (in particular personal or sensitive data), any hard constraints such as EU-only data residency, and basic non-functional needs such as criticality and human oversight. This produces a use-case profile and flags cases that need heightened governance attention. In particular, where personal or sensitive data is processed, the intake flags that a data-protection impact assessment (DPIA) under Art. 35 GDPR is likely required and carries that obligation forward into Gate 2.

Gate 1 — Feasibility. Gate 1 asks whether each hosting option is realistic for this particular SME, reflecting the two dominant criteria, Technology (30.6 %) and Organisation (21.4 %). It weighs technological readiness (data quality and accessibility, integration complexity, and where data may be processed), organisational capability (in-house IT and AI skills, the ability to run and monitor a system, and budget preference), the readiness of people to adopt (leadership trust, a sense of control, and capacity to absorb change), and environmental constraints (regulatory intensity, the evidence customers and partners demand, and vendor restrictions). The result is an *eligibility matrix* that marks each of {Cloud, On-Premises, Hybrid, AIaaS} as eligible, conditional, or excluded, each with a plain reason. For example, if no external data processing is permitted, public cloud and AIaaS fall away; if on-premises is required but the SME cannot operate it, that option is only conditionally feasible with managed-service support. Data sovereignty and residency therefore act as *non-compensatory* admissibility constraints at this gate—consistent with industry guidance that treats sovereignty and tested exit options as first-order risks to be settled up front rather than as later trade-offs.¹⁵

Gate 2 — Governance and evidence. Every option that survives Gate 1 must then be *governable*: someone can own it, and its safe operation can be shown on paper. Where the intake flagged personal or sensitive data, completing a data-protection impact assessment (DPIA) under Art. 35 GDPR is a required Gate-2 artefact for any option that processes such data; it is distinct from, and not satisfied by, a data-processing agreement (DPA) with a provider. This is also where *risk is enforced as a hard precondition* rather than traded off later: because Risk carries only a small AHP weight (3.1 %), leaving it to Gate 3 would let a cheap, capable but risky option win on points, so risk and security are screened here, non-compensatorily, not weighted at Gate 3. Gate 2 assigns clear ownership (who is accountable, responsible, consulted, and informed) and puts in place a small set of lightweight standard operating procedures (SOPs) — the four routines SMEs most often lack:

- **Algorithmic risk assessment** — check a model for bias, security weaknesses and robustness *before* it is trusted.
- **Fairness and transparency monitoring** — keep watching that it stays fair and explainable once it is live.
- **Data and model change control** — stay in control when the data or the model changes, so an update cannot silently break things.
- **Incident management and rollback** — a plan for when it goes wrong: how to notice, who reacts, how to switch it off.

Each SOP names an owner and a trigger and keeps a minimal evidence trail, and each is tied to the obligations it satisfies (GDPR documentation, EU AI Act, and NIS2 where relevant).¹⁶ If an option cannot be given an owner, or the required evidence cannot be produced, it is removed or deferred with a clear action list — however attractive its economics. (These four SOPs are the AI-specific reading of COBIT objectives; the controls are kept deliberately light so an SME can actually sustain them.)

¹⁵Bitkom 2026a.

¹⁶ISACA 2018; ISACA 2026.

Table 5: The four Gate-2 SOPs at a glance, preceded by the data-protection impact assessment (DPIA) that a personal-data use case triggers. Each SOP is a small routine with a trigger, an owner, a minimal evidence trail, and the obligation it helps satisfy. (‘Owner’ denotes accountability, not extra headcount; the DPIA is a regulatory precondition that feeds the SOPs, not a fifth routine.)

SOP	Trigger	Typical owner	Evidence kept	Maps to
DPIA (precondition)	Personal/sensitive data flagged at intake	Business owner (DPO advises)	DPIA record, data-flow map, mitigation measures	GDPR Art. 35
1. Algorithmic risk assessment	Before go-live; major model change	IT / technical lead	Risk record (bias, robustness, security)	EU AI Act (where high-risk); ISO/IEC 42001
2. Fairness & transparency monitoring	Ongoing, on a set cadence	Technical / business owner	Monitoring and fairness-audit reports	EU AI Act; GDPR
3. Data & model change control	Any data, feature or model change	IT lead	Change log, model-version register, rollback procedure	EU AI Act; ISO/IEC 42001
4. Incident management & rollback	Incident or performance degradation	IT lead + business owner	Incident playbook, post-incident actions	GDPR Art. 33; NIS2; EU AI Act

Gate 3 — Evaluation. Only options that have cleared feasibility and governance reach Gate 3, where they are compared. The tool sets a measurement plan across the extended Balanced Scorecard perspectives (financial, customer and market, internal process, learning, plus risk/compliance and technology where relevant) and ranks the surviving options using the empirically derived AHP weights (Table 1; options reach this stage only after clearing the Gate-2 risk and compliance *floors*, so the ranking weighs admissible survivors rather than letting a cheap option trade governance away), with total cost of ownership and breakeven entering here as one input among several. A sensitivity view shows how robust the ranking is to changed assumptions. Economics informs the choice at this stage; it never overrides feasibility or governance.

What the SME receives. Each run produces an auditable decision record with rationale and sign-off, an evidence vault with ownership and expiry dates, a standard-operating-procedure pack, a KPI plan with a review cadence, and an implementation backlog of follow-up actions.

6.2 What an SME gains

For a business owner the value is practical, not academic:

- **Fewer expensive wrong turns.** Because feasibility and governance are settled *before* cost, the firm does not invest in a hosting option that cannot legally or technically work — the most costly AI mistakes are caught early, on paper, not after the spend.
- **Supports an auditable decision trail.** Every decision ends in a documented record with linked evidence, so when GDPR, EU AI Act or NIS2 scrutiny arrives the firm can show *how* and *why* it decided — without a fire drill.
- **Faster, defensible decisions.** A structured, one-session process replaces open-ended debate, and the recommendation carries evidential weight because it rests on a peer-reviewable study, not a vendor’s opinion.
- **Governance it can actually sustain.** Controls are right-sized for an SME — four lightweight SOPs with clear ownership, not an enterprise governance department.

- **Less lock-in, fewer surprises.** Vendor-exit terms and an incident-and-rollback routine are built in, so the firm keeps its options open and has a plan when something breaks.
- **Reused, not repeated.** The same process, evidence vault and KPIs serve the next AI decision too, so governance compounds over time instead of starting from scratch.
- **Strategic alignment.** A short scorecard ties the chosen option to business goals and a review cadence, keeping AI spending connected to outcomes.

7 Worked Example

To make the gate logic concrete, consider a fictional but representative case: *Beispiel GmbH*, a German mechanical-engineering SME (180 employees) that wants a GenAI assistant to draft customer-service replies and search internal technical documentation. The use case touches GDPR personal data (customer correspondence) and proprietary engineering know-how.

Intake. Deployment scope: department-level, internal users. Data: GDPR personal + IP/secret. Data residency: EU-only; no training on company data permitted. Criticality: medium; explainability: moderate; human-in-the-loop: yes. The intake raises a “high governance attention” flag (personal + sensitive IP) and flags a required data-protection impact assessment (DPIA, Art. 35 GDPR).

Gate 1 — TOE feasibility. Data is moderately centralised; internal ML capability is low; operations (patching/monitoring) capability is limited; budget preference is OPEX; EU-only residency is a hard constraint. The feasibility screen produces the eligibility matrix in Table 6.

Table 6: Gate 1 eligibility matrix for the worked example (Beispiel GmbH).

Hosting option	Status	Rationale
Public cloud (non-EU region)	Excluded	Violates EU-only residency; non-EU external processing not permitted
AIaaS (EU region, no-train guarantee)	Conditional	Meets residency <i>if</i> a DPA and a no-prompt-training guarantee are obtained (resolved at Gate 2)
Hybrid (private endpoint + EU cloud)	Eligible	Meets residency; leverages existing IT; OPEX-compatible
On-premises	Conditional	Feasible only with a managed-service provider, given low internal operations capability

Gate 2 — COBIT governance & evidence. Named owners are assigned (business owner: head of service; technical owner: IT lead; compliance consulted: external DPO). A DPIA under Art. 35 GDPR is initiated for the surviving options, given the personal-data and IP exposure. For the *AIaaS* option, the required DPA and subprocessor transparency cannot be confirmed within the decision window, so *AIaaS* is *deferred with actions*. *Hybrid* and *on-premises (with MSP)* clear Gate 2 with a minimal SOP pack (algorithmic risk assessment, data/model change control, and incident-and-rollback, plus vendor-exit terms).

Gate 3 — extended BSC / AHP evaluation. The two survivors are scored on the eight criteria using the empirical weights (Table 1). Hybrid leads on Technology and Organisation (lower operational burden, faster time-to-value) and is competitive on Economic Efficiency once the MSP

cost of on-premises is included via TCO/breakeven. A sensitivity panel (stricter residency, higher volume) does not change the ranking. **Recommendation: Hybrid**, with AIaaS retained as a fallback if the DPA gap is later closed.

Sample output: decision record (extract)

Use case	GenAI service-reply + document-search assistant
Admissible set	Hybrid (eligible); On-premises+MSP (conditional)
Excluded / deferred	Public cloud non-EU (excluded); AIaaS (deferred — DPA gap)
Selected option	Hybrid (private endpoint + EU cloud)
Key rationale	EU-only residency met; lowest operational burden; passes governance minimum
Evidence attached	DPIA (Art. 35 GDPR), DPA (provider), data-flow diagram, risk-assessment SOP, incident-and-rollback SOP
KPIs (review quarterly)	Reply turnaround; deflection rate; data-residency conformance; cost per 1k requests
Owners / sign-off	Business: Head of Service · Tech: IT Lead · Compliance: external DPO

8 Differentiation and Positioning

The approach is positioned as a *decision-first* layer that complements standards and platforms:

1. **Sequencing:** explicit non-compensatory gates from feasibility to governance evidence before economics, grounded in empirical priority weights.
2. **Infrastructure focus:** admissibility of hosting options under SME constraints.
3. **SME-first operationalisation:** right-sized governance and lightweight evidence.
4. **Traceability:** a clear line from criteria to evidence and decision rationale.

This directly addresses the shortcomings catalogued in Table 4: it differentiates the tool from cloud adoption frameworks oriented toward cloud transformation capabilities¹⁷ and from risk frameworks that provide functions/outcomes but do not prescribe an ordered infrastructure decision pathway.¹⁸ Crucially, the decision logic and weights are not vendor assertions but the result of a peer-reviewable doctoral study, which gives the tool's recommendations defensible evidential standing in governance reviews and audits.

9 Frequently Asked Questions

- **Isn't this just COBIT or ISO 42001 under a new name?** No. Those standards say *what* good governance looks like; they do not tell an SME in *what order* to decide or which hosting option to pick. This framework supplies the sequence (feasibility, then governance, then cost) and produces an auditable hosting decision. COBIT is used inside Gate 2 only, and in a right-sized form.
- **Is this too heavy for a small company?** No. Governance is deliberately kept to the lightest defensible set: four short SOPs, each with one named owner and a thin evidence trail, run once per decision. There is no enterprise governance department to staff.

¹⁷Amazon Web Services 2026; Microsoft 2025.

¹⁸National Institute of Standards and Technology 2023.

- **Your weights come from German SMEs — do they apply to us?** The *logic* (feasibility and governance before cost) is general and not Germany-specific, and the EU regulatory layer (GDPR, EU AI Act, NIS2) is common across member states. The *weights* are a defensible default; the tool lets you re-tune them for your sector or firm, with the empirical values as the starting point.
- **We have no AI-governance staff. Can we still use it?** Yes. 'Owner' means accountability, not a new hire; the CEO can hold several roles and legal or data-protection advice can be external. The SOPs are designed for firms without a governance function.
- **Why not just follow our cloud provider's recommendation?** A provider optimises for its own platform. This method is vendor-neutral: it screens all options (cloud, on-premises, hybrid, AIaaS) on the same terms and leaves you an evidence record you own — useful precisely when you later renegotiate or switch.
- **Is this a tool or a method?** Both. It is a research-derived decision *method*, delivered as a guided web *tool*. You can apply the gate logic on paper today, or run *ToJo Navigator* for the full audit trail.
- **What happens when the regulations change?** The gate logic stays the same; only the Gate-2 evidence checklist is updated. The framework separates the durable structure (the gates and criteria) from the volatile detail (specific regulatory clauses).
- **How quickly do we see value?** A single hosting decision can be completed in one structured session, producing the decision record immediately; deeper governance (evidence vault, reusable controls) is layered on over time.

10 Conclusion

This paper specified an implementable gate-based decision-support web application for SME AI adoption, derived from an empirically grounded doctoral framework. The evidence base — expert interviews across seven sectors and a structured practitioner survey — shows that infrastructure choices are governed by feasibility and governance rather than by cost. By sequencing TOE feasibility, COBIT-anchored governance evidence, and extended BSC measurement weighted by the empirical AHP priorities, with transparent economic evaluation entering only within the admissible set, the tool produces auditable outputs while remaining feasible for SMEs. The specification has since been implemented: a working web application—*ToJo Navigator*, with *AI Governance* as its first module—is now running in pilot operation. For the decision-maker the promise is plain: an AI hosting decision that is faster to make, cheaper to get right, and that you can evidence the day it is questioned.

Scope and limitations. Two honest caveats. First, the priority weights were derived from German SMEs across seven sectors; they are evidence-based *defaults*, not universal constants, and are intended to be reviewed—and, where justified, re-weighted—for a different sector or jurisdiction. Second, the framework is empirically grounded but has not yet been *prospectively* validated in a live SME hosting decision; a working pilot instance is now in operation, which makes the natural next step—a piloted, longitudinal validation—concrete, but operation is not yet validation. The market scan above is an illustrative 2026 snapshot, not an exhaustive survey. These limits do not weaken the core finding (feasibility and governance precede cost); they mark where further evidence would strengthen it. Finally, the tool provides decision *support*, not legal advice: its outputs support an auditable decision trail but do not replace a DPIA, an EU AI Act conformity assessment, or professional legal and data-protection counsel, and do not guarantee that any audit or regulatory review will be passed. Whether a given obligation applies (e.g. the

EU AI Act risk class of the system, or whether the firm falls within NIS2 scope) depends on the individual firm and use case and should be confirmed with qualified advisers.

Glossary

Term	Definition
AHP	Analytic Hierarchy Process: a pairwise-comparison method for deriving criterion weights.
AIaaS	AI-as-a-Service: AI capability consumed as a managed external service.
BSC	Balanced Scorecard: a multi-perspective strategic performance framework.
COBIT	Control Objectives for Information and Related Technologies: an IT governance framework (ISACA).
DPA	Data Processing Agreement: a GDPR contract governing third-party data processing.
MVG	Minimum Viable Governance: the lightest governance set that keeps a decision auditable.
SOP	Standard Operating Procedure.
TCO	Total Cost of Ownership.
TOE	Technology–Organisation–Environment: a technology-adoption framework.

References

- Amazon Web Services (2026). *AWS Cloud Adoption Framework (AWS CAF)*. URL: <https://aws.amazon.com/cloud-adoption-framework/> (visited on 01/05/2026).
- Bitkom (2026a). *Cloud-Souveränität praktisch umsetzen*. Leitfaden. Bitkom e. V. URL: <https://www.bitkom.org/sites/main/files/2026-05/bitkom-leitfaden-cloud-souveraenitaet-praktisch-umsetzen.pdf> (visited on 06/04/2026).
- (2026b). *Künstliche Intelligenz in Deutschland*. Studienbericht. Bitkom e. V. URL: <https://www.bitkom.org/sites/main/files/2026-02/bitkom-studienbericht-ki.pdf> (visited on 06/05/2026).
- ISACA (2018). *COBIT 2019 Framework: Governance and Management Objectives*. URL: <https://netmarket.oss.aliyuncs.com/df5c71cb-f91a-4bf8-85a6-991e1c2c0a3e.pdf> (visited on 01/05/2026).
- (2026). *COBIT Resources*. URL: <https://www.isaca.org/resources/cobit> (visited on 01/05/2026).
- ISO/IEC (2023). *ISO/IEC 42001:2023 — Artificial intelligence management systems*. URL: <https://www.iso.org/standard/42001> (visited on 01/05/2026).
- Joswig, Thomas (2026). “Governing AI Infrastructure Decisions in SMEs: A Gate-Based Framework Grounded in the German Mittelstand”. PhD dissertation. Malta: Signum Magnum College.
- Microsoft (2025). *Microsoft Cloud Adoption Framework for Azure: Overview*. URL: <https://learn.microsoft.com/en-us/azure/cloud-adoption-framework/overview> (visited on 01/05/2026).
- Modulos (2026). *AI governance tools: the 2026 enterprise buyer’s guide*. URL: <https://www.modulos.ai/best-ai-governance-platforms/> (visited on 06/02/2026).
- National Institute of Standards and Technology (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. NIST AI 100-1. NIST. URL: <https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf> (visited on 01/05/2026).
- (2025). *AI Risk Management Framework (AI RMF)*. URL: <https://www.nist.gov/itl/ai-risk-management-framework> (visited on 01/05/2026).
- OECD (2025). *AI adoption by small and medium-sized enterprises*. Tech. rep. OECD Publishing. URL: https://www.oecd.org/content/dam/oecd/en/publications/reports/2025/12/ai-adoption-by-small-and-medium-sized-enterprises_9c48eae6/426399c1-en.pdf (visited on 06/02/2026).