

# Brücke

---

## Sovereign Decision Infrastructure for Financial Institutions

*“When the decision-maker leaves, the decision still has to be defended.”*

### TARGET AUDIENCE

Chief Risk Officers, Chief AI Officers, Heads of Audit, Heads of Model Risk

### BEACHHEAD

Banking and financial institutions — entering through risk-management expertise and mentor access

### CORE VALUE

Convert senior risk judgment into a human-attested, queryable asset for audit, analyst development and safer AI

*Brücke is not an API company. It is building the Attested Logic Graph — a system of record for an institution's proprietary conditional reasoning.*

## THE PROBLEM

# The Two-Front War

*Banks are losing context at the top and judgment capacity at the bottom*

1

## The demographic drain

Experienced risk officers retire, change roles or leave. The loss is not headcount — it's the conditional reasoning behind decisions: "If X happens, we do Y — unless Z is true — because this precedent changed how we interpret the risk."

2

## The manual archaeology tax

When ECB, BaFin, internal audit or model-risk teams revisit a historic decision, teams reconstruct context from fragmented files, emails and the memories of people who may not have been in the room.

3

## The apprenticeship bottleneck

Junior analysts developed judgment by working alongside seniors. As that access declines, they turn to generic AI that gives plausible answers without showing whether they match the bank's own precedent.

# 30%

of Germany's workforce — 13.3M people — will have passed statutory retirement age by 2040.

# €720K/yr

40 requests/yr × 120 senior hours × €150/hr — illustrative management estimate, not yet validated customer evidence.

*The problem is two-sided: preserve senior judgment and use it to train the next generation.*

## WHY NOW

# Why Brücke, Why Now?

*Three converging waves — demographics, regulation and shadow AI — make this the inflection point*

1

## Demographic wave

*The post-2008 cohort is exiting*

- 13.3M people — 30% of the workforce — pass retirement age by 2040
- Workers 55+ rose from 20.7% (2015) to 27.0% (2025)
- The generation that built post-2008 risk frameworks hits peak retirement 2025–2030

*Their unwritten conditional logic is leaving exactly as macro volatility returns.*

2

## Regulatory wave

*Hard deadlines, not aspirations*

- DORA live since 17 Jan 2025 — ICT risk mapping & third-party governance audited now
- EU AI Act Art. 50 (transparency) from 2 Aug 2026
- Art. 14 (human oversight) bars high-stakes AI without attestation and provenance tracing

*Brücke is architected to generate exactly these artefacts as a by-product of use.*

3

## Shadow AI wave

*Cognitive offloading, happening today*

- Junior analysts already use unapproved generic AI on complex credit/risk questions
- Creates unmapped compliance breaches and erodes independent judgment
- 88% of agentic AI pilots fail to reach production for lack of grounding & governance

*Ungrounded AI is a compliance and risk gap, not just a quality problem.*

*Brücke sits at the intersection of these three waves: preserve judgment, satisfy regulators, and make AI defensible.*

# The Three-Layer Sovereign Risk Infrastructure

*We are not a chatbot. We are building the sovereign operating system for institutional risk.*

1

## The Expert Copilot — Frictionless Input

The workflow layer for the Senior Risk Officer. It assists their daily tasks — checking precedent, drafting memos — and passively extracts their conditional logic and risk boundaries in the background.



2

## The Attested Logic Graph — The Sovereign Spine

The system of record. Extracted logic passes through a Hard Gate where the named expert reviews and digitally signs it — creating a durable, machine-readable repository of the bank's proprietary risk appetite. Unsigned AI guesses cannot enter.



3

## The Application Surface — Deployment

*The spine powers targeted applications across the bank:*

### The Provenance Portal

For CROs and auditors to instantly trace a decision's history to its signed human author (EU AI Act & DORA compliance).

### The Junior Risk Mentor

The talent-acceleration engine that uses the Spine to teach institutional logic to new analysts.

### The Grounding API

The read-only truth layer that keeps the bank's other AI models compliant.

# The Defensible Moat

*Why generic LLMs and horizontal enterprise software cannot replicate this*

1

## The “Data Gravity” flywheel

*Applications feed the infrastructure*

- The Junior Risk Mentor flags a “Knowledge Gap” when it cannot confidently answer a question
- This automatically prioritises the queue for the Senior Risk Officer to review
- The talent intake gap guarantees continuous feeding and updating of the graph

*The more juniors use it, the smarter the graph gets — a localised network effect.*

2

## Standing Assumption Register

*Detects when logic expires, not just stores it*

- Generic AI archives are static. Brücke maps dependencies between decisions
- If the ECB changes a core macro regulation, Brücke flags every historical decision built on that legacy logic
- We do not just store history — we detect when it expires

*Anti-drift by design: signed decisions carry their own expiry conditions.*

3

## High switching costs

*System of record for AI Act compliance*

- Once internal AI tools ground themselves against the validated graph for EU AI Act (Art. 50) transparency duties...
- ...ripping Brücke out requires re-validating the bank's entire AI compliance posture
- The customer still owns and can export its graph — the cost is revalidation, not data hostage-taking

*The moat is operational dependence, not lock-in.*

# The Junior Risk Mentor

*Turning institutional history into a judgment curriculum*

## The Credit Risk Gym

*The Mentor doesn't reveal the answer immediately. It presents a case from the bank's validated decision graph and asks:*

- 1 What would you investigate first?
- 2 Which risk factors matter?
- 3 What alternatives would you reject?
- 4 What evidence would change the decision?
- 5 How did the senior expert handle the historical case?

*The system then reveals the expert's signed reasoning, conditions and exceptions.*

## Live Cognitive Onramp

Brücke

For experienced juniors, Brücke structures the live problem, surfaces applicable precedent, flags where it may not apply, and escalates when coverage is insufficient.

### MEASURABLE HYPOTHESIS

~~9 mo~~



4-6 mo

Target reduction in time-to-operational-independence for junior risk analysts, subject to pilot validation.

*Do not present "9 months to 6 weeks" as an established outcome until measured in a finance pilot.*

# Why Existing Tools Fail in Banking

*The gap is not retrieval — it is validated reasoning*

| Tool category           | Does well                          | Why it fails for risk judgment                                 | Brücke's difference                          |
|-------------------------|------------------------------------|----------------------------------------------------------------|----------------------------------------------|
| Copilot / generic AI    | Drafts & retrieves documents       | Plausible summaries from an incomplete archive; no attestation | Signed decision nodes with provenance        |
| SharePoint / Confluence | Stores documents & procedures      | Assumes expert reasoning is already documented                 | Elicits reasoning from workflow & cases      |
| Enterprise GRC          | Logs risks, controls, remediation  | Manual entry; not built to capture tacit judgment              | Converts reasoning into a structured graph   |
| Transcription + RAG     | Finds passages in transcripts      | Shows what was said, not when it applies                       | Extracts conditions, exceptions, rejects     |
| LMS / e-learning        | Delivers generic training          | Disconnected from the bank's own cases                         | Training from human-attested precedent       |
| Brücke                  | Capture → Attestation → Deployment | —                                                              | Human-attested logic, provenance, escalation |

*Positioning line: Copilot can summarise what a risk officer said. Brücke records what the bank's judgment means, when it applies and who authorised it.*

# Designed for European Banking Realities

*Built backward from trust, governance and institutional control*



## Works Council & Employee Protection

- No individual performance scoring
- No disciplinary use
- Per-node expert review, correction & withdrawal rights
- No ambient surveillance
- Query activity anonymised by design

*§87 BetrVG co-determination is a deployment workstream, not a footnote.*



## DORA & AI Governance Posture

- Provenance for served outputs
- Human validation and escalation
- Capture, validation and query records
- Data-minimised, consent-bounded processing
- Customer-controlled deployment model

*"AI Act ready" means architected toward obligations — classification and counsel opinion pending.*



## Sovereign Substrate

- Single-tenant graph per institution
- On-prem, private-cloud or customer-controlled
- Customer-owned graph, export on termination
- No customer content trains a shared model
- Private model option for isolated environments

# The Self-Correcting Validation Timeline

*We do not ask for scale before we prove the unit economics. The pilot runs on strict, falsifiable gates.*

## GATE 1 · MONTH 1

### The Frictionless Capture Gate

#### FALSIFICATION TEST

Can we extract actionable, conditional risk logic from a Senior Risk Officer using the Expert Copilot without them treating it as a burdensome administrative tax?

#### METRIC

≥ 4 validated conditional nodes produced per expert hour

## GATE 2 · MONTH 3

### The Provenance Falsification Gate

#### FALSIFICATION TEST

We sit with the CRO and ask 3 complex regulatory questions from their last historical audit.

#### METRIC

Accurate, human-attested walkthrough >50% faster than the manual baseline, traced to the named author

## GATE 3 · MONTH 6

### The Decision-Drift Gate

#### FALSIFICATION TEST

Prove the system is an active DORA governance tool, not a static archive.

#### METRIC

Use the Standing Assumption Register to flag legacy logic that no longer fits current macro conditions or updated ECB guidance

# Priced Against Risk Exposure

*Acquisition costs offset by the consulting wedge*

## The revenue model (per site)

The Wedge — Implementation / Consulting

**€15,000**

upfront, funded via the CRO's discretionary “Audit Remediation & Knowledge Capture” budget — the delivery mechanism for the Capture Protocol, not a discretionary add-on.

The SaaS Platform — ARR

**€24,000 / yr**

Anchored against €720K+ in manual audit archaeology and avoided SREP capital add-ons.

**Total first-year value: €39,000 (before expansion)**

## Target unit economics (mature cohort)

**75%**

Gross margin at scale (single-tenant hosting & inference)

**≈€132K**

Customer LTV — 5-year horizon, conservative, excludes optionality

**≈€26.7K**

CAC, fully loaded sales & marketing

**5.0×**

LTV : CAC ratio

**≈ 12 mo**

CAC payback target — the €15K wedge materially offsets acquisition cost

*Investor framing: these are target mature-cohort metrics, not current performance. The pilot is designed to falsify the key inputs — capture yield, sales cycle, expansion.*

# The Path to Breakeven

*A disciplined path to cash-flow positive, built on one highly scrutinised assumption*

YEAR 1

10

active deployments

*Founder-led sales & design partnerships*

YEAR 3

~63

active deployments (~€1.4–1.5M ARR)

*Series A trigger*

YEAR 5

268

active sites + expansion

*€7.6M revenue — EBITDA breakeven*

## INVESTOR DISCIPLINE — THE “SOFTTEST INPUT” IN OUR MODEL

*We do not hide our highest-risk assumption; we state it.*

- To hit Year-5 breakeven, a ramped Account Executive must close ~14 deployments per year
- In enterprise banking, clearing IT Security, DORA checks and CISO approvals is a 6–9 month cycle
- The sensitivity: if AEs close 12 deals/year instead of 14, the plan does not breakeven in 5 years
- Measuring and compressing the sales cycle via the Consulting Wedge is the primary operational focus of the Seed round

Brücke

*A useful lever on this exact risk: every step-up in per-site ARR reduces the deployment count — and therefore the AE quota — needed to hit the same Year-5 revenue target.*

2026

# An €18.2B market, entered through a €150M beachhead

**€18.2B TAM**

Global enterprise risk, governance and compliance software.

**€4.5B SAM**

AI-enabled risk and fraud detection across banking, fintech and insurtech.

**€150M SOM**

Mid-tier banks and fast-growing fintechs in Europe and North America.

## Horizontal expansion vectors

**NOW · Banking & insurance**  
Driven by regulatory pressure, stress testing and model governance.

**NEXT · Manufacturing & industrial**  
Driven by retiring technicians, undocumented judgment and downtime cost.

**LATER · Energy & utilities**  
Driven by safety-critical continuity and scarce grid expertise.

**FUTURE · Aerospace & healthcare**  
Driven by aircraft-on-ground cost, MDR obligations and clinical continuity.

# Brücke Captures What Generic AI Cannot Obtain

---

*Brücke becomes the system of record for a bank's human-attested conditional risk logic — structured, attributable, continuously updated and embedded in the workflows that defend decisions and train the people who inherit them.*

The moat is not the API. The moat is not compliance alone.

**The moat is the compounding, proprietary Attested Logic Graph.**

- ◆ Expert participation
- ◆ Domain ontology
- ◆ Provenance & validation history
- ◆ Junior-analyst workflow adoption
- ◆ Regulatory & Works Council trust
- ◆ Decision-drift & coverage data
- ◆ Reference accounts in a trust-sensitive market

# The other side of the bridge: the talent intake gap

*The thesis: the future is human expertise amplified by AI. The same capture engine that preserves outgoing wisdom also solves the incoming candidate matching problem.*

**Layer 1 · Geopolitical & macro** Trade, regulatory and labour dynamics that shape the desk.

**Layer 2 · Industry & competitive** Sector structure and institutional positioning.


**Layer 3 · Company intelligence** Public-record strategy and leadership alignment.

**Layer 4 · Role reality (expert-authored)** The undocumented daily pressures of the job, extracted directly from the Company Brain. Exists nowhere else.

**Layer 5 · Interview simulator** Calibrated to the real scenarios the candidate will actually face.

# Future Product Surface: Role-Based Micro-Assistants

The same Attested Logic Graph powers targeted, role-specific copilots once the spine is proven.



These assistants are expansion products, not part of the seed scope — built only after the spine is proven and adopted.

Investor discipline: seed phase = spine + Provenance Portal + Junior Risk Mentor only. No engineering hours on micro-assistants before the core is validated in multiple banks.

# The 90-day sovereign pilot

01

## 14-day single-tenant deployment

On-premise or private cloud, entirely inside your security perimeter and through your standard third-party review rather than around it.

02

## Dual risk committee ingestion

Passive capture across two active committee streams — Credit Risk and ALCO, say — measured on attested coverage of decisions taken.

03

## Simulated regulatory dry run

Stress-test the Provenance Portal against past audit findings to prove 10-minute verification depth.

04

## Retrospective test

We answer three real information requests from your last twelve months out of the graph. You compare against what they actually cost you.

*“Let's ensure that when your CRO retires, their judgment remains in the bank — not just in their inbox.”*

# A deferral is not a reprieve. It is the window.

-  **2 Aug 2026** AI Act Article 50 transparency obligations apply. These were not deferred.
-  **2 Dec 2026** Article 50(2) marking applies to systems already on the market; new Article 5 prohibitions take effect.
-  **2 Dec 2027** High-risk obligations for standalone Annex III systems apply — deferred from August 2026 by the Digital Omnibus on AI, Regulation (EU) 2026/1744, in force 27 July 2026.
-  **2 Aug 2028** High-risk obligations for AI embedded in regulated products under Annex I.
-  **Continuous** DORA ICT third-party governance and the EBA SREP cycle. Neither has a deadline because neither ever stops.

**The window:** the evidence base has to be deployed, agreed with the works council and populated with real decisions before December 2027. Institutions that start once the harmonised standards land will be buying under deadline pressure, from whoever is already installed.

# Sources & Sourcing Discipline

Radical honesty: separating published data from our pilot hypotheses

PUBLISHED SOURCE

## Demographics & Labour Market

- **Destatis, 2025 Mikrozensus**  
13.3M people (30.0% of the workforce) pass statutory retirement age by 2040
- **IW Köln, 2026**  
Germany's working-age population shrinks by 4.3M by 2036

PUBLISHED SOURCE

## AI Adoption, Governance & Cognitive Offloading

- **MIT Sloan Management Review, 2026**  
“How AI Helps the Best and Hurts the Rest” — outcomes diverge by baseline user expertise
- **McKinsey & Company, 2026**  
“Five principles for designing brain-powered organizations” — cognitive offloading & skill atrophy
- **IDC & Gartner, 2026**  
88% of agentic AI pilots fail to reach production; 40% of AI projects cancelled by 2027

OUR ASSUMPTION — VALIDATE OR FALSIFY IN PILOT 1

## Internal Baselines & Management Estimates

|                                |                                                                                             |
|--------------------------------|---------------------------------------------------------------------------------------------|
| €720K “Manual Archaeology” tax | 40 requests/yr × 120 senior hours × €150/hr                                                 |
| Time-to-autonomy               | Targeted compression from 9 months toward 4–6 months for junior analysts                    |
| Unit economics                 | €120K ACV, 35% incident-prevention rate, 70% gross margin — tested against the 90-day pilot |