

# Evolving the OVB Service Platform Approach to Overcome the AI Experimentation Trap

Heinrich Fritzlär and Andrea Will

OVB Holding AG Germany

## ABSTRACT

OVB, a leading organization in the financial services sector, is proactively leveraging artificial intelligence (AI) to drive innovation and deliver measurable benefits in both customer experience and operational efficiency. During this AI-driven transformation, the company encountered the “AI Experimentation Trap” (Furr and Shipolov, 2025, Huang et al., 2025), the difficulty of converting promising AI prototypes into scalable, compliant, and value-generating solutions. In an era of increasing technological densification, many organizations face similar challenges, compounded by phenomena such as “Shadow AI” and the “Governance Drift Zone” (Silic et al., 2025). To address these challenges, and particularly to embed AI effectively into OVB’s core processes while maintaining customer relevance, the organization adopts a Service-Dominant (S-D) mindset, treating services as the central structuring paradigm. Complementing this approach, OVB employs Service Dominant Architecture (SDA), (Spohrer et al., 2022) as enterprise architecture and organizing logic for both its process design and its technical implementation as core platform. This architectural approach enables the seamless integration of AI-enabled services into the broader business ecosystem. The research picked the Translational Service Research and Design Methodology (TSRDM) (Warg et al., 2025) to systematically generate, translate, and apply knowledge that bridges the persistent gap between scientific advances in AI and their practical, value-creating implementation. In this way the work also contributes to the development of the unifying service language of TSRDM.

**Keywords:** Experimentation trap, Governance drift, Service dominant architecture, Translational service research and design methodology

## INTRODUCTION

The ambition of OVB’s AI First strategy is that every OVB advisor is supported by a personal AI agent, built on OVB data, standards and culture, to stay relevant in a market where competitors and new players massively invest in AI solutions. Core elements of OVB’s AI strategy are: embedding AI in the core business, building a data and platform foundation, establishing a strong AI-governance, and systematically upskilling people and leadership. AI First is defined as a strategic orientation, not as a pure technology project.

The implementation strategy relies on Service Dominant Architecture (SDA) as enterprise architecture and solutions like OVB.OS, a group-wide AI-Operating-System with a central Data Orchestration Platform and a

Webservice Suite to ensure openness, speed, reusability and data sovereignty of AI services. A group-wide AI governance framework defines minimum standards for responsible AI, covering data protection, model transparency, data provenance, robustness, fairness, and security, with the principle “enable, don’t block”.

At this point OVB faces challenges like turning existing use case based “speedboats”, such as the Allfinanz-Tutor for §34d/f exam preparation and the AI-based Sales Training, into robust, durable solutions that work reliably in day-to-day operations across the group or to consolidate and align AI, AI literacy, data, automation in a coherent steering model (Benlian and Pinski, 2025). As long as initiatives are “islands”, OVB faces friction in reusing components, managing risks consistently and achieving economies of scale. This phenomenon of scattered pilots that are not integrated with existing operations and not tied to clear customer or business value, so that results stay local and do not translate into P&L impact is known as “AI Experimentation Trap” (Furr and Shipolov, 2025).

The motivation and purpose of this work is to build knowledge about the phenomenon of the “AI Experimentation Trap” and to elaborate, design, and engineer approaches to overcome this trap and its obstacles, including an analysis of whether the chosen Service Dominant Architecture (SDA) appropriately fits as the enterprise architecture and supports these approaches.

## RESEARCH METHODOLOGY

Translational Service Research and Design Methodology (TSRDM) (Warg et al., 2025) as systematic approach to bridge persistent gaps between scientific discoveries, e.g. AI, and their implementation is chosen as the research design (De Vaus, 2001; Creswell and Creswell, 2017). Through its eight-step process, TSRDM structures linkages, transitions, and frictions between the disciplines of service research, service design, and service engineering. To address these, a translational space of “services” (Gummesson, 1995), architectures, and patterns is developed based on the unifying service language that turn theoretical knowledge into reusable, connectable design prototypes, accelerating innovation cycles.

The subsequent chapters elaborate distinct steps of TSRDM in a coherent translational trajectory. The AI experimentation trap and related organizational challenges deepen the initial problem framing and context analysis. The knowledge base chapter operationalizes the evidence synthesis and conceptual modeling step. Developing translational services and aligning them to Service Dominant Architecture instantiate design and prototyping steps, while the OVB service platform and findings sections address implementation reflection and iterative learning regarding scalability, governance, and value cocreation outcomes.

## THE AI EXPERIMENTATION TRAP

The AI experimentation trap (Furr and Shipolov, 2025; Huang et al., 2025) is a strategic failure mode in which organizations rely on widespread

AI experiments without anchoring them in clear customer or business outcomes, digital foundations, or scale-up pathways, so that most initiatives remain isolated pilots that consume resources, fragment effort, and ultimately erode confidence in AI's value. It mirrors earlier "digital transformation" mistakes where a "let 10,000 flowers bloom" approach produced numerous local successes but very few enterprise-level gains. Recent surveys demonstrate that the vast majority of AI initiatives fail to take off and companies are seemingly stuck in proof-of-concept purgatory (Haefner et al., 2023).

Key characteristics and obstacles are:

- Experiments are scattered across functions and proceed without a unifying roadmap, creating coordination *frictions, lags of integration with existing processes and political resistance* when later attempts are made to prioritize, standardize, or decommission local AI initiatives.
- Because "success" is assessed only at the pilot level rather than at P&L or service-system level, *local wins do not translate into enterprise learning*, which obstructs strategic alignment and slows decisions on which AI initiatives deserve scale-up.
- Every "successful" local AI solution adds new interfaces, data paths, and dependencies, *increasing technical and organizational complexity* that makes subsequent experiments materially slower, more expensive, and operationally fragile.
- Accumulating complexity effects (technical debt, opaque interdependencies, longer and more fragile experiments) directly obstruct robust model testing and validation, raising model risk and *creating friction in demonstrating effective governance, compliance, and internal control* to regulators and auditors.
- Rapid, decentralized AI experiments that repeatedly reuse data and features across systems generate *opaque data flows*; this obscurity hampers data lineage, purpose limitation, and lawful-basis management, creating persistent compliance friction with GDPR (General Data Protection Regulation) requirements and undermining auditability.
- The proliferation of AI tools in organizations has given rise to "Shadow AI", the unsanctioned use of AI. The "Governance Drift Zone" represents a state where formal organizational policies exist but lack real-world traction, falling behind the rapid, informal adoption of AI tools by employees (Silic et al., 2025).

As a result, the organization delays realizing AI's real potential to redesign core services and value creation.

## **ORGANIZATIONAL CHALLENGES RESULTING FROM THE AI EXPERIMENTATION TRAP**

The AI experimentation trap creates a set of intertwined organizational challenges: it locks AI in costly pilots, prevents scaling, and erodes governance

frameworks and thus trust in AI's strategic value. The challenges manifest themselves across technology and data, organizational skills and capabilities, as well as governance, risk, and scaling dimensions.

From a technology and data perspective, many experiments run on fragile, local tech stacks and fragmented datasets so they cannot be industrialized. Without shared, reliable data infrastructures and interoperable, scalable system architectures, AI remains a promising but essentially non-scalable prototype layer. Legacy systems, inconsistent data quality, unclear access rights, and technology-driven complexity reinforce this problem: every new AI use case has to fight the same integration and data issues instead of building on a common platform and re-usable services.

On the organizational skills and capabilities level, AI initiatives frequently sit in isolated innovation labs with weak line-of-business ownership and no dedicated "ninja team" with mandate, budget, and cross-functional reach to drive solutions from pilot to scale. At the same time, limited AI and data literacy among leaders and staff fuels both hype and fear: unrealistic expectations and misuse of tools collide with resistance to changing workflows and decision rights, because people do not clearly see where AI genuinely adds value and where it needs to be limited.

Finally, in terms of governance, risk, and scaling, unclear rules around model ownership, accountability for AI-supported decisions, compliance, and handling of bias or errors mean that legal and risk functions tend to default to caution, slowing or blocking deployment into production. Most pilots were never designed with enterprise-grade security, robustness, and deep process integration in mind, so the step from "lab demo" to an integrated, enterprise-ready solution reveals substantial new risks and integration work, leaving a persistent gap between experimentation and real, organization-wide business impact.

## **KNOWLEDGE BASE**

To build objective-related knowledge, we develop a knowledge base as a foundational element for deriving design principles that guide the engineering of platform capabilities essential to overcoming the AI experimentation trap. This effort centers on value creation and value fulfillment, while delineating services from products (Grönroos, 2008; Vargo and Lusch, 2004b). Consistent with the TSRDM framework (Warg et al., 2025), we adopt "services" as the structuring paradigm and "service" as the "overarching grammar" defined as application of resources for the benefit of another actor (Vargo and Lusch, 2004a).

We draw on Service Dominant Architecture (SDA) (Warg et al., 2016) to elaborate design principles for OVB core platform. Analogous to architecture's dual role as plan and outcome e.g. in construction, SDA facilitates a structure that functions as both plan and realization for the processes recursively organized. This approach facilitates the systematic assignment of services to SDA and orchestrates the process of value cocreation including connecting partners. SDA is examined through conceptual and applied lenses (Warg,

2020). Conceptually, it offers design patterns akin to a construction plan, comprising five interconnected systems. In practice, this pattern manifests as a tangible configuration that coordinates mutual value creation processes. As a structural enabler, SDA allows entities, such as organizations, to evolve roles and systems that, through implementation and reciprocal service exchange, form dynamic value cocreation configurations, ultimately constituting service systems. Finally, SDA as structure is plan and result of the services and processes it recursively organizes (Warg et al., 2015; Warg et al., 2016; Warg, 2020; Spohrer et al., 2022). At the core of SDA are five systems that become Service Systems by being applied by a responsible actor (entity). Service Systems represent open arrangements that (1) enhance another system's state by sharing or applying resources (yielding perceived value for the recipient) and (2) advance their own state by integrating external resources (deriving self-perceived value from interactions). Economic exchanges within this paradigm rely on voluntary, reciprocal value cocreation, wherein all participating systems engage willingly and experience mutual improvement (Spohrer et al., 2008; Demirkan and Spohrer, 2016). SDA consists of the following five (service) systems:

- System of Interaction: supports real-time interaction among customers, providers, and other relevant stakeholders and in doing so the realization of “value in use” by “acting” on AI services.
- System of Participation: supports the integration of external capabilities, solutions, and services into the value creation process.
- System of Operant Resources: supports the development, orchestration, and deployment of strategically relevant services; as a prerequisite for their combination and bundling into new value propositions building value and value creation processes.
- System of Data: supports the construction of data-based customer understanding and tailored event generation derived from interactions.
- System of Institutions: is complemented by institutional rules that coordinate actors and govern access to capabilities, thereby enabling or constraining their use.

## **DEVELOPING TRANSLATIONAL SERVICES FOR LINKING AI TO OVB CORE**

In this chapter, the Translational Service Research and Design Methodology (TSRDM) is applied to move from the previously elaborated obstacles of the AI experimentation trap to concrete service-based interventions. Building on the knowledge base, translational services are elaborated that directly target scattered pilots, opaque data flows, weak process integration, and fragmented governance, thereby creating reusable service patterns that systematically link AI initiatives to OVB's core value creation processes.

**Table 1:** Translational services for overcoming the AI experimentation trap.

Obstacles of the AI Experimentation Trap	Translational Services	SDA System
Obstacle “no customer value” (outcomes)	- “interaction services” for customer engagement and value in use	System of Interaction
Obstacle “opaque data flows”	- Data integration services, e.g. xml, uml, rdf - Data interchange services, e.g. web services, WAP, J2EE, .Net - Data and algorithmic governance services that define shared vocabularies, data products, lineage models. Data orchestration and measurement services that implement event schemas and interoperability KPIs	System of Data
Obstacle “connecting and integrating partner capabilities”	- Partner onboarding and participation services that provide standardized API contracts, test suites, and legal/operational checklists for external AI/data providers and internal “speedboat” teams. - interconnection services, e.g. HTTP, S/MIME	System of Participation
Obstacle “no interlink with organizations capabilities and processes - AI capability gaps”	- “Operant services” on the SDA enable the connection of the different human and non-human actors and their services - SDA service systems” support mutual value creation - service engineering and bundling services for turning AI prototypes into platform capabilities	System of Operant Resources
Obstacle “Governance Drift Zone”; “Shadow AI” - coordination of actors and resources; e.g. platform access	- access and participation services (who can join the platform: Web services, security services) - API terms - compliance rules and role-based access to platform capabilities and resources	System of Institutions

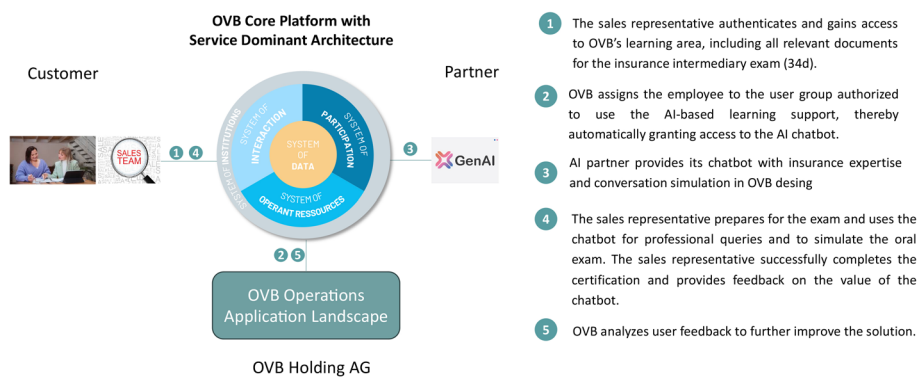
## OVB SERVICE PLATFORM FOR OVERCOMING THE AI TRAP

The paper section on the OVB service platform analyzes how assigning translational services to the five systems of Service Dominant Architecture (SDA) provides a structural backbone for overcoming the AI experimentation trap. Building on a Service-Dominant mindset, the OVB core platform is described as an SDA-based configuration in which interaction, participation, operant resources, data, and institutions become distinct service systems that jointly orchestrate value cocreation between human and non-human actors. Translational services derived from TSRDM are mapped onto these systems to connect AI pilots with OVB’s core processes, ensuring that customer-facing interaction services, partner participation services, operant AI capabilities, data and algorithmic services, and institutional access and governance services form a coherent service hierarchy rather than isolated “speedboats”.

The OVB core platform is positioned as an agile, emergent architecture (Bradley, 2018, Waterman et al., 2015) that grows from concrete AI and automation use cases, rather than from a fully specified Big Architecture Up Front plan. Drawing on agile emergent-architecture concepts, OVB deliberately limits upfront decisions to a small set of irreversible choices such as the selection of partners like e.g. muffintech, and then iteratively extends the “architectural runway” as new use cases are prepared for scaling into the core platform.

Each AI initiative is treated as a use case: once it proves customer and business value, a focused architecture effort lifts it from local “speedboat” to reusable platform capability, by assigning the services to OVB’s enterprise architecture SDA. This approach aligns with the concept of “platformizing,” which allows firms to reuse capabilities in different contexts to fully ramp up AI across the organization (Haefner et al., 2023). At OVB, this is operationalized through the AI Master Capability Program (MCP), a systematic logic to transform isolated MVPs, such as the “Allfinanz-Tutor” or “AI Sales Training,” into resilient, horizontal capabilities.

In this perspective, Service Dominant Architecture (SDA) remains the organizing logic, intentionally evolving the five SDA systems along a prioritized roadmap. For example, a promising advisory AI assistant first runs as a local interaction service; when scaling signals appear or integration with other processes is needed, the relevant SDA systems are extended e.g. the System of Data (shared features, lineage, monitoring), the System of Institutions (roles, policies, accountability), or the System of Participation (standardized onboarding for additional channels and partners) just enough services are build to support safe group-wide reuse. This “last responsible moment” planning keeps the platform closely aligned with validated value-in-use, while avoiding accidental architectures and to ensure controlled enablement and to prevent governance drift.



**Figure 1:** OVB core platform with service dominant architecture (process perspective).

## FINDINGS

The study confirms that OVB’s challenge is not a lack of promising AI pilots, but the difficulty of translating these local successes into robust and reusable

platform services with measurable business impact. Fragmented speedboats, heterogeneous tech stacks, and opaque data flows keep AI solutions trapped in experimentation, while weak line-of-business ownership, uneven AI literacy, and cautious governance further slow the transition from demo to enterprise-ready service.

The findings show that combining a Service-Dominant mindset, Service Dominant Architecture, and TSRDM with a use-case-based agile emergent-architecture approach, provides a viable escape path from this trap. OVB advances its core platform by selecting high-leverage AI and automation use cases, proving customer and business value, and then deliberately “lifting” these use cases into the SDA-based enterprise architecture. In line with emergent-architecture principles, upfront design is limited to a small set of irreversible decisions - e.g. the “tech-stack” -, while technical implementation for each new use case is extended at the last responsible moment, balancing risk, early value, and reusability. In line with the elaborated knowledge base, OVB overcomes the “AI Experimentation Trap” and associated risks like ‘Shadow AI’ or the “Governance Drift Zone”. This use-case-driven evolution yields several concrete effects: AI initiatives that reach scaling are systematically converted into reusable services assigned to SDA and thus to platform capabilities, reducing redundancy and enabling faster onboarding of subsequent use cases. At the same time, low-leverage experiments are deliberately kept at the edge, which prevents accidental implementations and unmanaged complexity from undermining the enterprise architecture and the core platform. The result is a gradually emerging portfolio of reusable SDA core services that reflects adaptation and learning rather than top-down blueprints, while still satisfying regulatory, security, and interoperability requirements.

## **CONCLUSION**

The paper concludes that overcoming OVB’s AI Experimentation Trap depends on treating AI initiatives as services-rendering use cases that deliberately cultivate an agile, emergent enterprise architecture and core platform grounded in Service Dominant Architecture (SDA). By doing so, the study directly addresses its initial motivation to build knowledge about the AI Experimentation Trap and to elaborate, design, and engineer approaches that overcome its obstacles. Rather than pursuing a comprehensive Big Up-Front Architecture or tolerating fragmented, incidental implementations arising from isolated pilots, OVB advances its enterprise architecture incrementally, guided by the most valuable and reusable AI use cases and by last-responsible-moment implementation decisions that are structurally anchored in SDA. In line with the elaborated knowledge base, this use-case-driven evolution demonstrates how SDA, as the chosen enterprise architecture, can be assessed and further developed as a fitting organizing logic for scaling AI, mitigating risks such as Shadow AI and the Governance Drift Zone, and systematically converting local experiments into reusable platform services.

## REFERENCES

- Benlian, A. & Pinski, M. 2025. The AI literacy development canvas: Assessing and building AI literacy in organizations. *Business Horizons*.
- Bradley, C. 2018. An Introduction to Agile Emergent Architecture - Always Intentional. *The Scrum Crazy Blog* [Online]. [Accessed 2018, September, 28th].
- Creswell, J. W. & Creswell, J. D. 2017. *Research design: Qualitative, quantitative, and mixed methods approaches*, Sage publications.
- De Vaus, D. 2001. Research design in social research. *Research design in social research*, 1–296.
- Demirkan, H. & Spohrer, J. C. 2016. Emerging service orientations and transformations (SOT). *Information Systems Frontiers*, 18, 407–411.
- Furr, N. & Shipolov, A. 2025. Beware of the AI Experimentation Trap. *Harvard Business Review*, August 29, 2025, 7.
- Grönroos, C. 2008. Service logic revisited: who creates value? And who co-creates? *European Business Review*, 20, 298–314.
- Gummesson, E. 1995. Relationship Marketing: its Role in the Service Economy. In: GLYNN, W. J. & BARNES, J. G. (eds.) *Understanding Services Management*. John Wiley & Sons Ltd.
- Haefner, N., Parida, V., Gassmann, O. & Wincent, J. 2023. Implementing and scaling artificial intelligence: A review, framework, and research agenda. *Technological Forecasting and Social Change*, 197, 122878.
- Huang, Y., Martin, S. & Qin, Z. 2025. The Trap of Complexity in Experimentation. Available at SSRN 5126086.
- Silic, M., Silic, D. & Kind-Trüller, K. 2025. From Shadow It to Shadow AI—Threats, Risks and Opportunities for Organizations. *Strategic Change*.
- Spohrer, Vargo, S. L., Caswell, N. & Maglio, P. P. The service system is the basic abstraction of service science. *Proceedings of the 41st Annual Hawaii International Conference on System Sciences (HICSS 2008)*, 2008 2008. IEEE, 104–104.
- Spohrer, J., Maglio, P. P., Vargo, S. L. & Warg, M. 2022. *Service in the AI era: Science, logic, and architecture perspectives*, Business Expert Press.
- Vargo, S. L. & Lusch, R. F. 2004a. Evolving to a New Dominant Logic for Marketing. *Journal of Marketing*, 68, 1–17.
- Vargo, S. L. & Lusch, R. F. 2004b. The four service marketing myths: remnants of a goods-based, manufacturing model. *Journal of service research*, 6, 324–335.
- Warg, M. 2020. Architecture and Its Multifaceted Roles in Enabling Value Cocreation in the Context of Human-Centered Service Design. *International Conference on Applied Human Factors and Ergonomics*, 2020. Springer, 79–85.
- Warg, M., Spohrer, J., Böhmman, T., Leitner, C., Carrubbo, L., Gruhn, V., Neuhüttler, J., Weiss, P. & Burmeister, M.-A. 2025. Translational Service Research And Design Methodology (TSRDM) - A new Service Research and Design Methodology in exploration [Online]. Available: <https://www.ifsd.hamburg/translational-service-research-and-design-methodology-tsrdm/> [Accessed].
- Warg, M., Weiß, P. & Engel, R. 2015. *Service Dominant Architecture*. FH Wedel, University of Applied Science.
- Warg, M., Weiß, P., Engel, R. & Zolnowski, A. 2016. Service Dominant Architecture based on S-D logic for Mastering Digital Transformation: The Case of an Insurance Company. In: RUSSO-SPENA, T. & MELE, C. (eds.) *26th Annual RESER Conference*. Naples, Italy.
- Waterman, M., Noble, J. & Allan, G. How much up-front? A grounded theory of agile architecture. *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering*, 2015. IEEE, 347–357.