



SWISSI INSTITUTE FOR AI

Mandate-Bound Agent Exchange

Research consortium and partner invitation

Walter Kurz¹

¹Swissi Institute for AI, kurz@swissi-ai.institute

September 2026

1 Research vision

Companies are starting to use their own AI agents to find counterparties and prepare commercial terms. Each agent acts for a different organisation and interprets messages from interested parties. Mandate-Bound Agent Exchange (MBAX) investigates how signed company authority and purpose-bound disclosure can be enforced throughout such dialogue, with evidence for consequential acts.

The Swissi Institute for AI is assembling a multi-country consortium now for a research proof of principle: a formal policy model, reference implementation and independent evaluation under adversarial multi-turn interactions. We are preparing a European funding application as this work begins.

2 Scientific question and method

Can act-time checks outside the language model preserve authority and disclosure properties while an agent reasons over adversarial content and still negotiates usefully? The project will define its threat model, specify testable properties and compare against strong agent-security and data-space baselines.

The mechanism will combine typed messages, signed mandates, bounded disclosure, revocation, human-approval conditions and session evidence. Tests will measure unauthorised acts, leakage, task utility, interoperability and decision quality. Results will state the tested conditions and observed failure cases.

3 Partners and pilot

We seek security and formal-methods researchers; AI-agent and red-team groups; identity and data-space specialists; human-oversight and business-law researchers; SMEs, industrial adopters and business networks. Beneficiary roles follow specific work packages and capacity. Complementary organisations from the same country are welcome.

The Business AI Endpoint Network offers a planned setting for consented cross-company validation. Participating companies operate their own agents and approve business commitments. The consortium will run a controlled research cohort while development partners recruit across countries towards a wider pilot network.

4 Proposed research work

The work programme will be shaped with participating organisations. Its first stream will define the actors, policies, trust assumptions and measurable properties. A second stream will implement a reference mechanism over existing agent communication protocols, with a conformance suite that independent teams can run. A third stream will perform adversarial evaluation against comparable baselines and report security alongside useful task completion, latency and cost. A fourth stream will study human approval and evidence interpretation in consented company workflows. Results will feed a protocol specification, open research outputs, standards dialogue and a follow-on exploitation plan.

The project brings together formal security, applied AI, distributed systems, human factors and business practice. A partner's contribution should answer a specific research question or supply a practical test setting. Research organisations can lead methods and independent evaluation. SMEs and industrial companies can contribute integration, real operational constraints and pilot scenarios. Business networks can connect a wider range of companies across countries and sectors. Each beneficiary's role, staff and eligible costs will be defined in the final consortium proposal.

The scientific assessment will distinguish established security properties from observed performance. It will test revocation and policy decisions under realistic message timing, examine leakage through free text, and document failures that inform the next design iteration. Public research artefacts will use synthetic or suitably governed data so that other groups can reproduce the conclusions. Companies contributing pilot scenarios will agree the study's consent, confidentiality and publication conditions before participation.

5 Expression of interest

Please send your relevant work, proposed contribution, legal entity and country, named lead, available staff or facilities, and early IP or data-sharing conditions. We welcome discussion of a concrete work package or pilot role.

Contact

Dr. Walter Kurz, MBA, M.Sc.
Swissi Institute for AI

Swissi Holding AG
Baarerstrasse 78
[CHE-189.603.726](tel:+414189603726)
6300 Zug
Switzerland

kurz@swissi-ai.institute