SW4E SOFTWARE ENGINEERING ECOSYSTEM

DIMECC

SW4E ECOSYSTEM

AI SANDBOX

Secure, Transparent and Collaborative AI Experimentation

2025

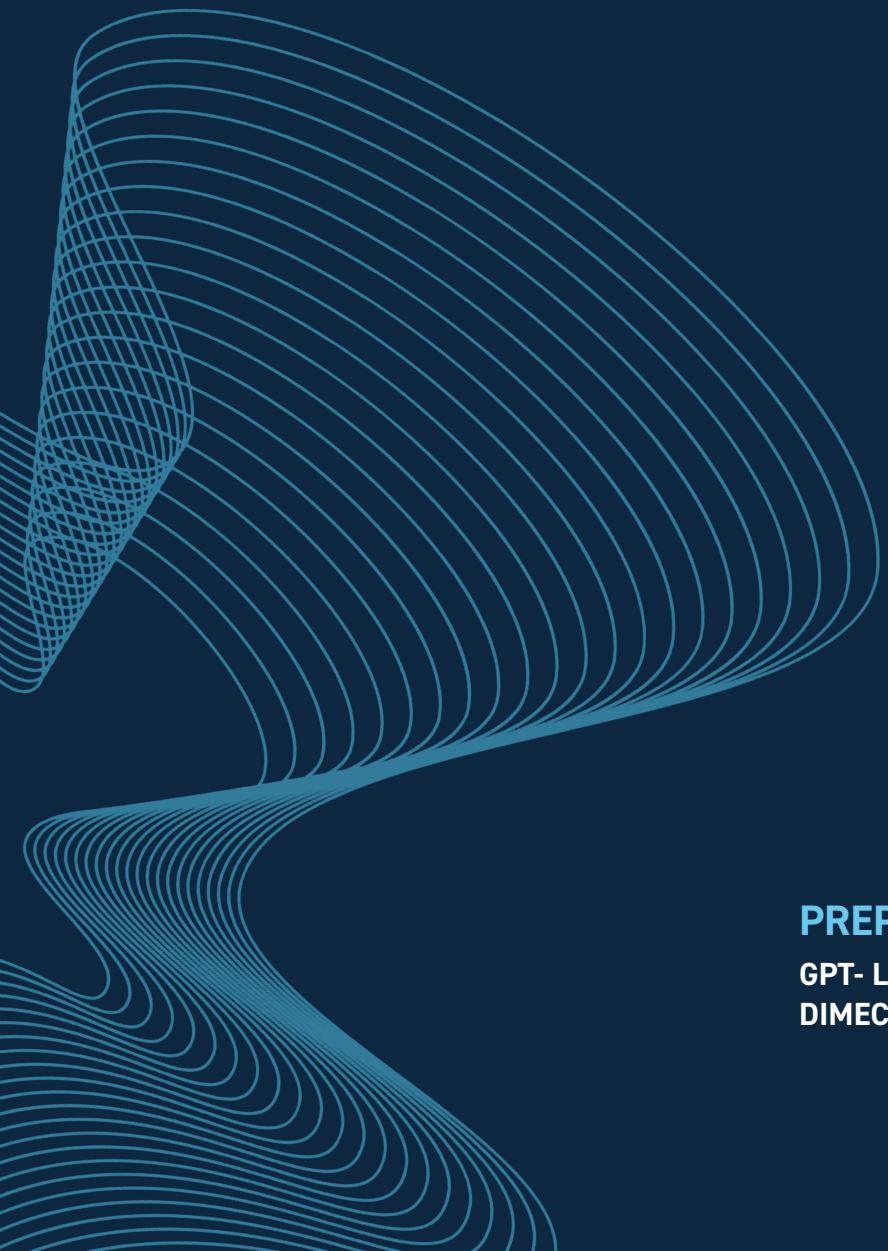

PREPARED BY

GPT- Lab, Tampere University
DIMECC and the SW4E ecosystem

DIMECC PUBLICATIONS
11/2025

SW4E SOFTWARE ENGINEERING ECOSYSTEM

Images - Canva Pro
Graphic Design & layout - Vilma Saari
Publisher - DIMECC Oy

Finland's AI future will be built through 600× faster experimentation.

The envisioned AI Sandbox is a secure, industry-ready environment where Finnish companies, universities, and public organisations can design, test, and deploy AI-powered software systems at unprecedented speed without compromising compliance, data protection, or intellectual property.

As the EU AI Act reshapes how intelligent software must be built, verified, and maintained, the Sandbox gives teams a trusted place to move from concept to compliant AI-powered products. It turns governance into automation, documentation into evidence, and regulation into a competitive advantage.

Developed by Tampere University's GPT Lab together with the DIMECC-led SW4E ecosystem, the Sandbox vision and prototype merges Finland's strengths in software engineering, AI, and digital trust.

Its vision: **to make Finland the European lighthouse for AI-native software engineering and AI-powered software systems.** A place where companies can achieve the culture of 600× experimentation speed and lead the transition to intelligent, trustworthy AI-powered software.

This document explains **why Finland needs a shared and trusted sandbox environment for AI development, what the AI Sandbox MVP would deliver, and how organisations could use it.**

Prepared by SW4E Ecosystem & GPT lab, Tampere University 2025

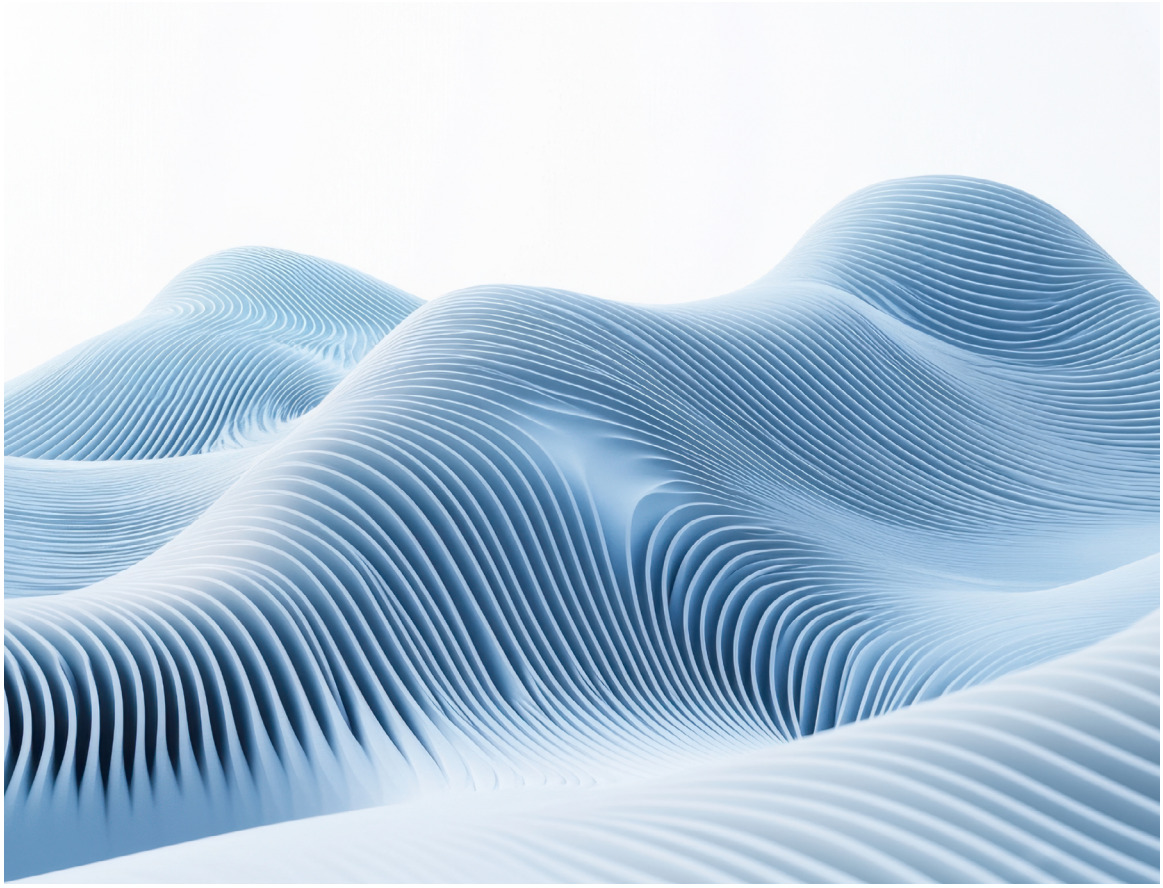

CONTENTS

EXECUTIVE SUMMARY

Artificial intelligence and software-driven innovation are transforming industries faster than most organisations can adapt. New European regulations, including the EU AI Act, have made trust, transparency, and compliance essential preconditions for competitiveness. As organisations race to harness AI, they must also prove that their systems are safe, lawful, and aligned with European values.

The vision for the AI Sandbox is to provide a platform for a software development culture of 600 x faster experimentation

The AI Sandbox provides a solution to this challenge. It will offer a secure, compliance-ready environment where companies, universities, and public organisations can develop, test, and benchmark AI-powered software systems together without compromising data protection or intellectual property.

It is the first Finnish model for shared AI experimentation under full governance transparency — combining technical infrastructure with ethical and legal assurance. Developed within the DIMECC-run SW4E ecosystem, the Sandbox brings together Finland's leading expertise in software engineering, AI, and digital governance. The concept and prototype level technical implementation has been carried out by Tampere University's GPT Lab, a leading AI and software engineering research unit specialising in trustworthy AI and secure computing environments. The result is a functioning prototype deployed on CSC's Rahti Cloud and verified through industry–academia collaboration.

Key Outcomes

- **A shared resource** for Finnish companies and research institutions to collaborate towards the vision "600 times experimentation culture"
- **Proven concept:** A multi-tenant prototype demonstrating how AI innovation can operate safely, lawfully, and collaboratively.
- **Governance framework:** A neutral operating model aligned with the EU AI Act, GDPR, and international standards such as ISO/IEC 27001 and 23894.
- **Technology stack:** Modular, scalable architecture that supports AI benchmarking, compliance auditing, and secure collaboration.
- **Neutral operation:** The Sandbox has been established through the joint coordination of DIMECC and Tampere University's GPT Lab, ensuring that the environment operates under transparent, open, and vendor-neutral principles.

What the envisioned and defined MVP Delivers

The MVP will be a working environment where participants can log in to their organisational tenant, manage projects, access approved data, run AI experiments, and generate audit-ready compliance evidence — all through a unified dashboard.

It includes **25+ AI services** such as benchmarking, fine-tuning, cost estimation, explainability, and model vulnerability scanning. Built-in compliance and data governance controls make every experiment traceable, auditable, and repeatable.

The already existing prototype demonstrates that responsible AI development is not a limitation, but a competitive advantage. **It proves that AI can be built and validated in an open, collaborative, and lawful environment — one that strengthens both Finland's and Europe's innovation sovereignty.**

The Road Ahead

With the MVP definition now available, the focus shifts to real-world piloting between industry and academia. The goal is to grow the Sandbox from a research-driven prototype into a national platform for trusted AI experimentation, engaging Finnish companies, research organisations, and public bodies in hands-on use cases.

The AI Sandbox is ready to serve as Finland's foundation for 600x experimentation culture in responsible, compliant, and collaborative development of AI-powered software systems — enabling organisations to innovate confidently in a rapidly evolving regulatory landscape.

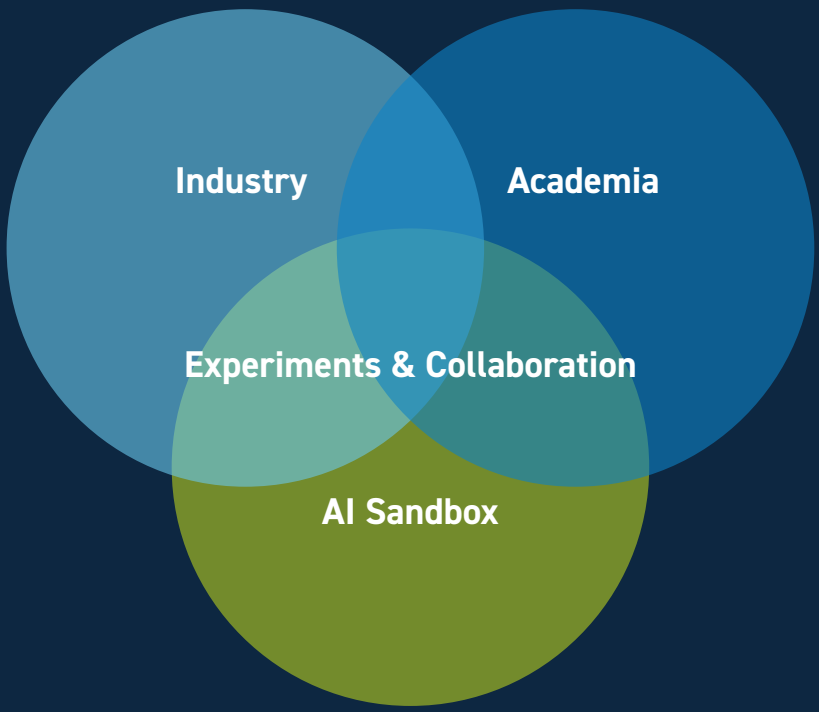

1. BACKGROUND AND THE GLOBAL OUTLOOK

The Global and European Context

AI is revolutionizing software development and software itself is being redefined.

This change requires a new kind of AI-native software engineering where design, implementation, testing, and compliance teams work continuously and transparently with an unforeseen speed.

Artificial intelligence has entered a decisive stage in global technology and policy.
The EU Artificial Intelligence Act (2024) and related frameworks such as NIS2 are transforming the way AI systems must be designed, validated, and governed.

Across Europe, competitiveness increasingly depends on the ability to operate within these new regulatory conditions while still accelerating innovation.

Unlike global technology giants that rely on closed, proprietary ecosystems, Europe is building its own model of innovation: one based on trust, transparency, and cross-sector collaboration. This approach is particularly vital for AI in safety-critical and data-sensitive domains—healthcare, manufacturing, mobility, and public services—where ethical and legal accountability are central to public trust. Responsible innovation is becoming Europe's competitive edge.

The organisations that can combine technical excellence with legal and ethical integrity will define the next generation of trusted AI.

Challenges Facing Industry and Academia

Despite rapid technological progress, most organisations still encounter structural barriers that slow down responsible AI adoption:

- **Data and Compliance Barriers:** Access to relevant data is limited by sensitivity, localisation rules, and high compliance costs.
- **Fragmented Collaboration:** Research and industry often work in isolation, hindering technology transfer and shared learning.
- **Lack of Trusted Environments:** Few infrastructures offer both security and openness for collaborative AI experimentation.
- **Regulatory Uncertainty:** New obligations under the AI Act require risk-based monitoring and documentation that many current testbeds cannot provide.

These barriers prevent Europe from turning AI research into scalable, market-ready solutions. Promising concepts remain locked in prototypes because there are too few trusted, auditable environments where organisations can safely test and validate them together and share expertise.

Finland and the Nordic Opportunity

The Nordic region—and Finland in particular—has all the ingredients to lead Europe's transition toward trusted and human-centric AI ecosystems.
Finland combines advanced digital infrastructure, world-class research capability, and a culture of open collaboration across business, academia, and the public sector. Through initiatives such as DIMECC, Business Finland, and national research and innovation programmes, Finland has demonstrated how ecosystem-based cooperation accelerates technological progress.

The AI Sandbox builds on this foundation to create a neutral, secure, and collaborative environment where:

- **Companies can test and validate AI solutions before deployment and learn from others.**
- **Researchers can move innovations rapidly into industrial contexts.**
- **Public organisations can observe, evaluate, and guide responsible AI development in practice.**

Finland's combination of trust, technological excellence, and cooperation culture positions it as the ideal testbed for Europe's model of responsible AI innovation.

Turning Point for European Innovation

The convergence of new regulation, mature infrastructure, and ecosystem readiness creates a unique moment for action.
The AI Sandbox has emerged at this intersection, and will provide a practical and transparent response to Europe's need for compliant, collaborative AI environments.
It serves as both a national demonstration platform and a scalable blueprint for how trusted digital innovation can operate in practice.
By aligning policy, research, and technology, the Sandbox lays the foundation for the next generation of AI development—technically advanced, ethically grounded, and governed by design.

2. VISION & STRATEGIC VALUE

A Shared Environment for Trusted AI Innovation

The AI Sandbox is a secure, compliance-ready collaboration environment that unites industry, academia, and public organisations in the responsible development of AI and software technologies. It acts as a neutral collaboration layer where organisations can experiment, validate, and benchmark AI solutions under verifiable legal and ethical safeguards.

- Companies can prototype and benchmark AI models without exposing proprietary data.
- Universities can move research directly into applied contexts.
- Public organisations can assess and validate AI systems before adoption.

Together, these capabilities make the AI Sandbox a unique bridge between innovation and implementation, ensuring that Finland's AI breakthroughs are both advanced and responsible.

Strategic Alignment with European and Nordic Priorities

The defined Sandbox aligns with both Finland's national AI strategy and Europe's regulatory and funding priorities, including the EU AI Act, the Data Governance Act, and the Digital Europe Programme.
It advances three strategic priorities:

1. Trust and Accountability

Embedding ethical, legal, and compliance requirements directly into technical workflows.
Providing verifiable audit trails, documentation, and explainability to meet AI Act requirements.

2. Interoperability and Collaboration

Enabling secure data and model exchange between organisations through a federated architecture. Supporting cross-institutional projects while preserving data sovereignty.

3. Sovereignty and Competitiveness

Strengthening Europe's ability to innovate within its own legal and ethical frameworks.
Positioning Finland as a leader in open, standards-based AI infrastructures.

The Sandbox offers a concrete path toward Europe's vision of responsible, innovation-driven digital growth.

Thematic Pillars: Technology, Trust, Transparency

The concept for the Sandbox is built around three interconnected pillars that guide its design and operation:

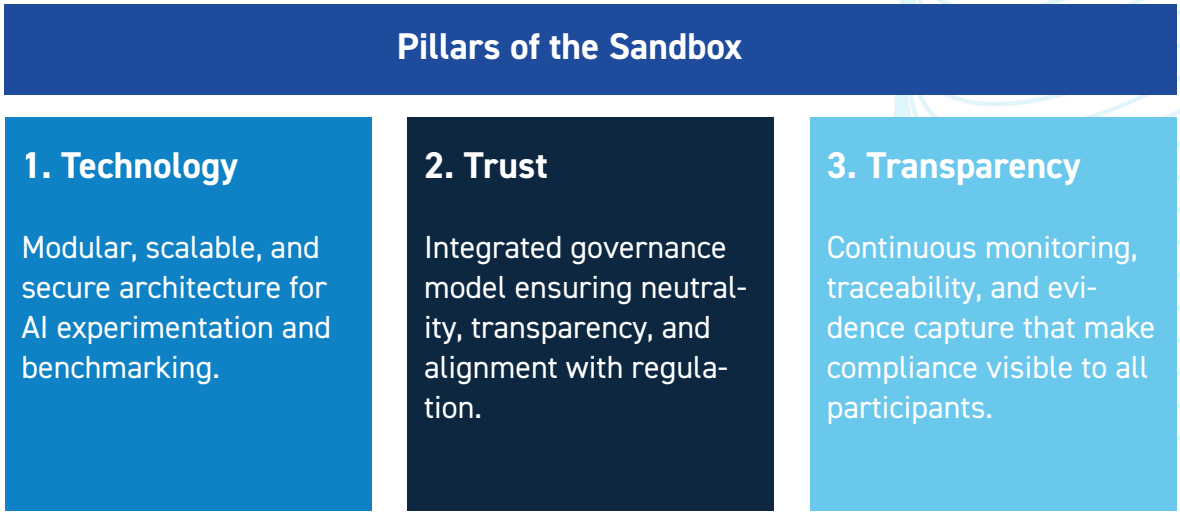

Each project carried out within the Sandbox strengthens these pillars, contributing to Finland's reputation as a hub for responsible, high-impact AI development.

Value for Finland

For Finland, the AI Sandbox would establish as a strategic national capability that enhances competitiveness, supports world-class research, and accelerates responsible AI adoption.

Key outcomes and benefits

- Enabler for cross learning between industry and academia and different industrial sectors
- Faster R&D cycles and experimentation
- Shared infrastructure lowering cost and compliance effort.
- Neutral governance fostering long-term collaboration.
- Built-in compliance shortening time-to-market

The AI Sandbox shows that regulatory responsibility and innovation speed can reinforce rather than limit each other.

SW4E's Response — Building the Trusted AI Environment

The AI Sandbox is designed as a strategic instrument, not just a technical platform. Its purpose is to remove structural barriers to collaboration, ensuring that

- AI development remains compliant by design,
- innovation occurs within a governed and transparent framework, and
- results are reproducible, auditable, and scalable.

The AI Sandbox **Minimum Viable Product (MVP)** demonstrates end-to-end workflows: from organisation onboarding and data governance to experiment execution and compliance monitoring.

The prototype and feedback collected proves that a trusted AI environment can be implemented and maintained with shared resources and transparent rules.

Looking ahead, **the focus is on expanding real-world pilots** between academia and industry, using the MVP as a foundation for applied projects and capability development.

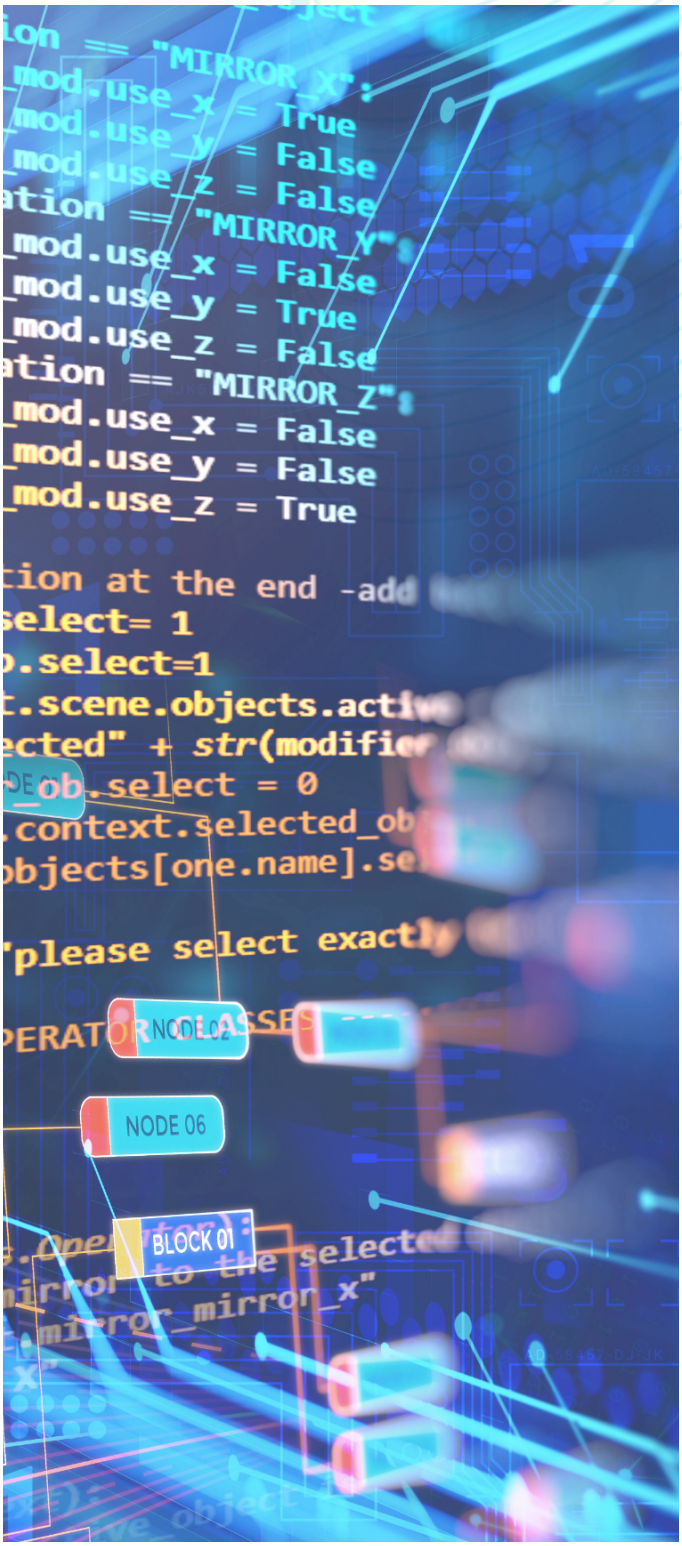

3. CAPABILITIES & DEMONSTRATED RESULTS

From Concept to Working Prototype

The AI Sandbox has evolved from concept to reality through a phased development process, culminating in a fully functional prototype before the Minimum Viable Product (MVP) deployed on CSC's Rahti Cloud in 2025. The prototype validates both the technical feasibility and the governance concept of the Sandbox, demonstrating how secure, transparent, and collaborative AI experimentation can operate in practice.
The initiative was set up through the coordination of DIMECC and Tampere University's GPT Lab, with active input from Bittium, Q4US, Solita, and other SW4E partners.
Together, they defined requirements, tested functionality, and verified the Sandbox's value for different types of organisations.

How it works in practice

- Each organisation receives its own secure tenant with federated sign-on.
- Users access a dashboard that lists projects, data assets, and compliance status.
- The Service Catalogue lets participants launch AI tasks such as benchmarking, fine-tuning, or explainability analysis.
- A Data Workspace connects to both open and controlled data sources and provides automated preprocessing and anonymisation.
- Monitoring and cost-tracking tools visualise resource use, compute time, and emissions estimates.
- A built-in Audit and Evidence panel records lineage, approvals, and reports ready for inspection or funding documentation.

The prototype demonstrates that a trusted AI experimentation environment can combine openness and security without compromising speed or usability.

Core Capabilities of the AI Sandbox

The defined MVP includes more than **25 AI-related services, grouped into six functional areas.**
These turn collected stakeholder requirements into ready-to-use modules for experimentation and compliance.
These modules transform the Sandbox from a research testbed into a production-ready platform where innovation and compliance coexist. The following tables summarise all AI-related services currently deployed in the prototype environment.

Benchmarking & Evaluation	
Model Benchmarking Service	Compare AI model performance against standard benchmarks to assess competitiveness.
Requirement Templates Generation Service	Turn unstructured project or regulatory requirements into structured, compliance-ready templates.
Cost Estimation Service	Use AI analysis to estimate project effort, time, and cost for AI development proposals.
AI Technique Assessment Service	Recommend suitable AI/ML techniques based on data type, task, and user needs.
Results Dashboard and Reporting Service	Generate interactive dashboards and reports for result comparison and documentation.

Data & Compliance Services	
Data Preprocessing Service	Automate the cleaning, transformation, and validation of datasets before experimentation.
Data Residency Control Service	Manage data location and storage jurisdiction to meet EU/EEA compliance requirements
Data Anonymization Service	Apply anonymization techniques (e.g. k-anonymity, l-diversity) to protect sensitive data.
Compliance Auditing Service	Automate GDPR and EU AI Act compliance checks and generate supporting evidence.
Legal Assistant Service	Provide conversational guidance for GDPR, AI Act, and privacy impact documentation.

Controlled Experiment Execution	
Controlled Experiment Execution Service	Run reproducible AI experiments securely with monitored GPU resource allocation.
Experiment Tracking Service	Track, version, and compare experiment parameters, datasets, and outputs

Security & Assurance	
AI Model Vulnerability Scanning Service	Identify model or dataset vulnerabilities, bias, and policy breaches.
AI Security Scanning Service	Perform full-stack vulnerability scanning across code, infrastructure, and runtime.

Infrastructure & Operations	
GPU Resource Allocation Service	Allocate scalable GPU/CPU clusters for AI model training and inference.
Secure Data Storage Service	Provide encrypted, access-controlled data storage compliant with EU regulations.

Core AI Services	
LLM Playground Service	Safe environment for testing large language models across various tasks and data types.
Retrieval-Augmented Generation (RAG) Service	Deliver document-based Q&A and summarization with source citation tracking.
Fine-Tuning Service	Enable domain-specific model training and optimization through hyperparameter tuning.
Text Analysis Service	Offer sentiment analysis, classification, and entity recognition for textual data.
Computer Vision Service	Support image classification and object detection for visual data applications.
Speech Processing Service	Convert speech-to-text, text-to-speech, and analyze voice data with multilingual support.
Recommendation Service	Provide personalized and privacy-preserving recommendation algorithms.
Time Series Analysis Service	Perform forecasting and trend detection on sequential or streaming data.
Anomaly Detection Service	Detect outliers and irregular patterns in structured and unstructured datasets.
AI-Native Anomaly Detection Service	Advanced anomaly detection using LLM and AI-native integration for complex data streams.

Data Management & Sources

The AI Sandbox integrates open and controlled data sources to support experimentation, benchmarking, and model validation.
Examples of data sources:

Source	Domain	Purpose / Use Case	Access
Avoindata.fi	Public services, municipalities	Datasets on government services, infrastructure, demographics	Open API / downloads
Statistics Finland	Demographics, economy	Benchmarking, socio-economic analysis	Open / bulk downloads
THL Sotkanet	Healthcare, welfare	Health and wellbeing indicators for compliant AI use cases	Open API
Findata	Health & social data	Controlled individual-level registers for authorised research	Permit-based
Finnish Social Science Data Archive (FSD)	Social research	Survey-based data for public-sector and societal AI cases	Open or restricted
EU Open Data Portal	Cross-domain	Multilingual European datasets for cross-border pilots	Open API
ELIXIR Finland / CSC Resources	Life sciences	Bioinformatics data for AI model validation	Research agreement

Data are accessed through APIs or secure downloads, processed via automated pipelines for cleaning and anonymisation, and then linked to AI workflows for training, fine-tuning, or evaluation.

Each dataset is classified as Public, Organisational, or Sensitive, with matching controls and full audit logs.

Demonstrated Results and Impact

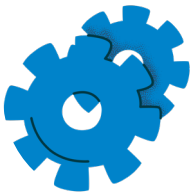

Technical Validation

Fully deployed and prototype operational on CSC Rahti Cloud.
Multi-tenant access, GPU management, and real-time monitoring tested successfully.

Governance Proof

Compliance verified against GDPR and EU AI Act principles.
Functional audit system linking policy requirements to technical evidence.

Stakeholder Engagement

Needs and use cases defined with Bittium, Q4US, Solita, and GPT Lab research teams.

Research-to-Industry Translation

Academic results on benchmarking and compliance integrated into operational tools.
Real-world demonstrations show how regulatory requirements can be embedded in daily workflows.

Key Outcome:

The AI Sandbox has moved beyond a vision to a concrete concept to become a functioning national infrastructure for compliant AI experimentation.
It enables organisations of any size to innovate confidently while meeting European standards for security, fairness, and accountability.

4. GOVERNANCE & TRUST FRAMEWORK

Building Trust

The AI Sandbox is not only a technical platform but also a governed environment designed to ensure that every experiment, dataset, and model operates under clear rules, traceable responsibilities, and verifiable oversight.
Its governance framework is grounded in three principles:

- **Neutrality**: No single organisation owns or controls the AI Sandbox; all participants operate on equal terms.
- **Transparency**: Governance decisions, policies, and updates are openly shared with members.
- **Accountability**: Each participant has defined responsibilities and audit trails supporting regulatory compliance.

By embedding these principles into both the operational and technical design, the Sandbox turns trust from a value statement into a measurable, enforceable practice.
Trust is not assumed — it is built through structure, transparency, and evidence.

Neutral Operator Commitment

The AI Sandbox prototype has been established through the joint coordination of DIMECC and Tampere University's GPT Lab, ensuring that it operates under transparent, open, and vendor-neutral principles. This approach guarantees impartial governance during the pilot phase while enabling future transitions to broader, community-led management.

The neutral operator model allows all participants — companies, universities, and public organisations — to collaborate on equal terms in a secure and governed environment.

Data Governance and Compliance

Data governance is central to trust in AI. Data governance ensures lawful, transparent, and ethical handling of information at every stage.

Key measures

- Data classification: Public, Organisational, Sensitive.
- Lawful basis tracking and consent management.
- EU/EEA residency enforcement for sensitive data.
- Immutable audit logs and automated anonymisation.

Outcome: Every dataset can be shown compliant and ethically handled before, during, and after use.

Security and Assurance Framework

Security underpins trust.
The Sandbox employs a security-by-design approach aligned with ENISA and ISO/IEC standards.
Key security components

- **Identity and Access Management:** Role-based permissions, multi-factor authentication, and federated SSO.
- **Encryption**: TLS 1.3 for data in transit and AES-256 for data at rest with EU-based key management.
- **Network Isolation:** Segregated tenants and private subnets for high-sensitivity projects.
- **Monitoring and Alerts:** Continuous anomaly detection and automated notifications.
- **Incident Response:** Defined playbooks and 72-hour GDPR breach reporting readiness.

These measures ensure verifiable operational integrity across all participants and projects.

Intellectual Property and Licensing

Intellectual property management follows transparent and open principles:

- Background IP: Remains the property of its creator.
- Foreground IP: Results from Sandbox projects belong to the project consortium under predefined terms.
- Infrastructure IP: Core platform components and governance templates are released under permissive open-source licences (Apache-2.0 / MIT) to maintain neutrality and accessibility.

This approach guarantees that shared innovation remains equitable and that all participants benefit from the collective development.

Summary

The governance model transforms trust into an operational principle.
By combining neutrality, transparency, and security, the Sandbox provides a replicable framework for responsible AI collaboration in Finland.

5. STAKEHOLDER VALUE PROPOSITION

The AI Sandbox delivers measurable value for all parts of Finland's innovation ecosystem. It provides a secure and collaborative environment that accelerates responsible AI adoption. Each stakeholder gains a different but complementary advantage from participating.

Value for Companies

For companies, the AI Sandbox is a low-risk gateway to trusted AI development. It removes the uncertainty of compliance and data protection while enabling faster experimentation.

Key benefits

- Shorter innovation cycles. Ready-made frameworks cut project setup time.
- Reduced compliance burden. GDPR- and AI Act-aligned processes simplify audits and documentation.
- Safe collaboration. Enables joint R&D with universities and partners without IP or confidentiality conflicts.
- Access to national expertise. Connects directly to Finland's leading AI researchers and policy experts.
- Reputation for responsibility. Demonstrates ethical and regulatory commitment to customers and funders.

Outcome: Companies can accelerate their AI-powered software development and reduce their time-to-market with faster experimentation and development cycles. Companies can prototype confidently, knowing every step is secure, auditable, and regulation-ready.

Value for Academia and Research Organisations

For universities and research institutes, the Sandbox bridges the gap between research and real-world use. It enables reproducible experiments under controlled, legally sound conditions.

Key benefits

- Applied research testbed. A governed environment for joint academic–industrial projects
- Technology transfer. Moves validated results efficiently into business applications.
- Data access. Provides compliant pipelines to open and controlled datasets.
- Visibility and impact. Positions research outcomes within national and European innovation programmes.
- Funding readiness. Aligns with priorities of Horizon Europe, Digital Europe, and Business Finland.

Outcome: Researchers can demonstrate tangible societal and industrial impact while meeting the highest ethical standards. Collaborative AI Sandbox enables smoother co-creation and experimentation with research project partners.

Value for Public Organisations and Policymakers

For the public sector, the AI Sandbox functions as a trusted observation and testing environment. It allows policymakers to evaluate the behaviour of AI systems before large-scale adoption.

Key benefits

- Evidence-based decision-making. Transparent logs and audits provide real-world insight into AI impacts.
- Policy experimentation. Supports pilot projects exploring AI Act implementation and risk assessment.
- Assurance and accountability. Strengthens public trust in digital services.
- Cross-sector coordination. Creates a meeting point for government, academia, and business to align objectives.

Outcome: Public actors gain practical experience in AI governance while guiding national digital policy.

Value for Investors and Strategic Partners

For investors and funders, the Sandbox proves that responsible AI can be commercially viable and verifiable.

Key benefits

- Lower investment risk. Verified governance reduces exposure to ethical or legal issues.
- Scalable model. Infrastructure can expand across domains and user groups.
- Portfolio insight. Performance and compliance data support informed funding decisions.
- ESG alignment. Investing in transparent, fair AI supports sustainability and accountability goals.

Outcome: Investors can back innovation with confidence, knowing it is grounded in verified trust.

Collective Value – The Ecosystem Effect

The true strength of the AI Sandbox lies in its network effect.
When companies, universities, and public organisations collaborate under one governed framework, their combined capability grows exponentially.

Shared outcomes

- Accelerated R&D. Cross-sector projects form faster through shared tools and standards.
- Knowledge exchange. Expertise flows freely between research and industry.
- Policy learning. Regulators observe innovation in real time and adapt rules accordingly.
- Scalable innovation. Successful pilots can be replicated across sectors nationwide.

Collective impact: The Sandbox transforms collaboration into measurable innovation capacity, strengthening Finland's position as a trusted leader in AI.

6. OPERATING MODEL

The Operating Model

The Sandbox follows a tiered governance model separating strategic oversight from operational management. This ensures flexibility for experimentation while maintaining consistent ethical and legal standards.

Main governance layers

- **Steering Group:** Represents the SW4E ecosystem, providing strategic direction and policy alignment.
- **Operator & Technical Steward:** Manage the infrastructure, implement governance policies, and coordinate updates.
- **Participants (Companies, Universities, Public Bodies):** Conduct experiments, manage their own datasets and projects, and comply with shared governance principles.
- **Advisory Experts:** Offer domain-specific guidance on regulation, ethics, and sustainability when needed.

Proposed Roles and Responsibilities

Each role within the Sandbox has a clear scope of responsibility to maintain accountability and compliance.

Role	Key Responsibilities
Steering Group	Approves governance updates, ensures alignment with EU/national frameworks, and mediates policy decisions.
Operator & Technical Steward	Maintain neutrality, implement security controls, publish transparency reports, and coordinate infrastructure.
Organisation Administrator	Manage users, data, and local compliance within their organisational tenant.
Project Lead / Research Admin	Oversee experiments, validate datasets, and maintain traceability for audits.
Compliance & Security Officers	Review adherence to GDPR, AI Act, and internal data-handling policies.

This structure ensures every decision and data flow has an accountable owner, creating a verifiable chain of trust from governance to experiment level.

Operational Principles

The AI Sandbox would act as a neutral and shared innovation infrastructure.
Its operational approach combines transparency, inclusiveness, and long-term sustainability to maintain trust and scalability.

Three guiding principles define its management model:

1. **Neutrality**
   Governance and coordination are performed by trusted, non-commercial partners to ensure fairness and openness.
2. **Transparency**
   All governance updates, performance reports, and audit results are accessible to participants.
3. **Sustainability**
   The operating model supports continuous use, further development, and financial resilience beyond pilot funding.

This structure allows the Sandbox to remain an independent, reliable environment that serves both research and industry over time.

Governance and Funding Structure

Strategic direction is provided by a multi-stakeholder steering group, supported by technical and policy experts who maintain the governance framework and ensure compliance.
The funding concept follows a hybrid approach that keeps participation accessible while sustaining operations.

Funding sources

- Public innovation programmes (e.g. Business Finland, Horizon Europe).
- Membership or project participation fees.
- Targeted development grants for new features or capacity growth.

The goal is to sustain a national asset — not a commercial product.

Partnership and Growth Mechanisms

Growth is achieved through collaboration rather than ownership.
Partnerships ensure that the Sandbox evolves with user needs and remains relevant for Finland's AI ecosystem.

Engagement paths

- **R&D collaboration:** joint industry–academia projects using Sandbox infrastructure.
- **Innovation pilots:** short, focused experiments validating AI methods or governance concepts.
- **Public sector use:** training, policy pilots, and testing of responsible-AI procedures.

These activities generate both technical and strategic insight, strengthening the case for continued funding and national impact.

Summary

The AI Sandbox demonstrates that responsible AI innovation can be maintained sustainably.
Its neutral governance, transparent operation, and collaborative funding approach ensure that growth never compromises trust.
Together with its partners, the Sandbox now serves as a long-term foundation for compliant and collaborative AI experimentation in Finland.

7. GROWTH & NEXT STEPS

Roadmap & Next Steps

With the prototype now completed and validated, the focus shifts from concept development to further development, real-world use and national scaling.

The coming years will turn the AI Sandbox from a research and vision-driven prototype into an active environment serving Finnish companies, universities, and public organisations in applied AI experimentation and achieve the 600x experimentation speed from 2025.

The near-term focus is piloting and national scaling — deepening impact through real use cases and demonstrating long-term value.

2025

Launch of the MVP version and initiation of the first pilot project.

2026-2027

Launch new pilot projects between industry and academia. Expand use to 10–15 active organisations across different sectors. Refine governance and services based on pilot feedback, and broaden participation through training and outreach.

2028-

Secure continuous funding for operations and establish the Sandbox as a permanent element of Finland's innovation infra-structure.

Collaboration Opportunities

The AI Sandbox grows through practical collaboration, connecting partners who share an interest in responsible and transparent AI innovation.

Upcoming opportunities
Industry–Academia Pilots: Collaborative projects using the Sandbox for AI validation, benchmarking, and data-driven research.
Funding Initiatives: Joint applications for development and operational funding to ensure sustainable growth.
Training and Awareness: Hands-on workshops introducing responsible AI design, auditing, and risk management practices.
Community Building: Continued engagement through SW4E ecosystem activities and national RDI networks.

Each of these efforts strengthens the AI Sandbox as a functional environment, evolving with Finland's digital landscape and ensuring that responsible innovation becomes part of everyday AI practice.

Summary

The prototype has proven the AI Sandbox concept and technical viability.

The next step is to scale its use, expand participation, and secure funding for continuous operation. By linking academic research with industrial application under a shared, governed framework, the AI Sandbox will serve as a cornerstone for Finland's AI ecosystem — enabling innovation that is not only advanced, but also accountable and future-proof.

8. THE PROTOTYPE

AI Sandbox landing page

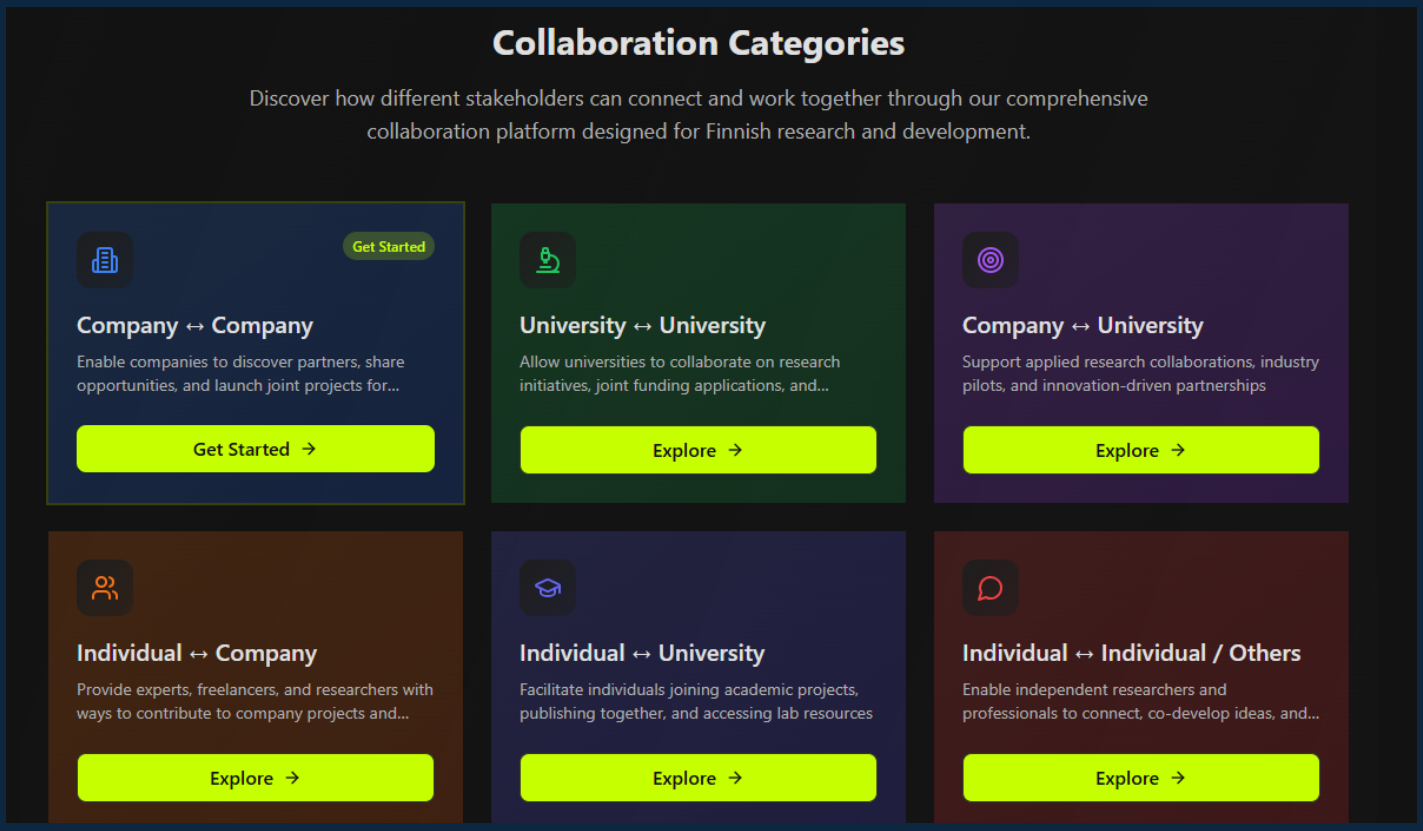

Collaboration opportunities

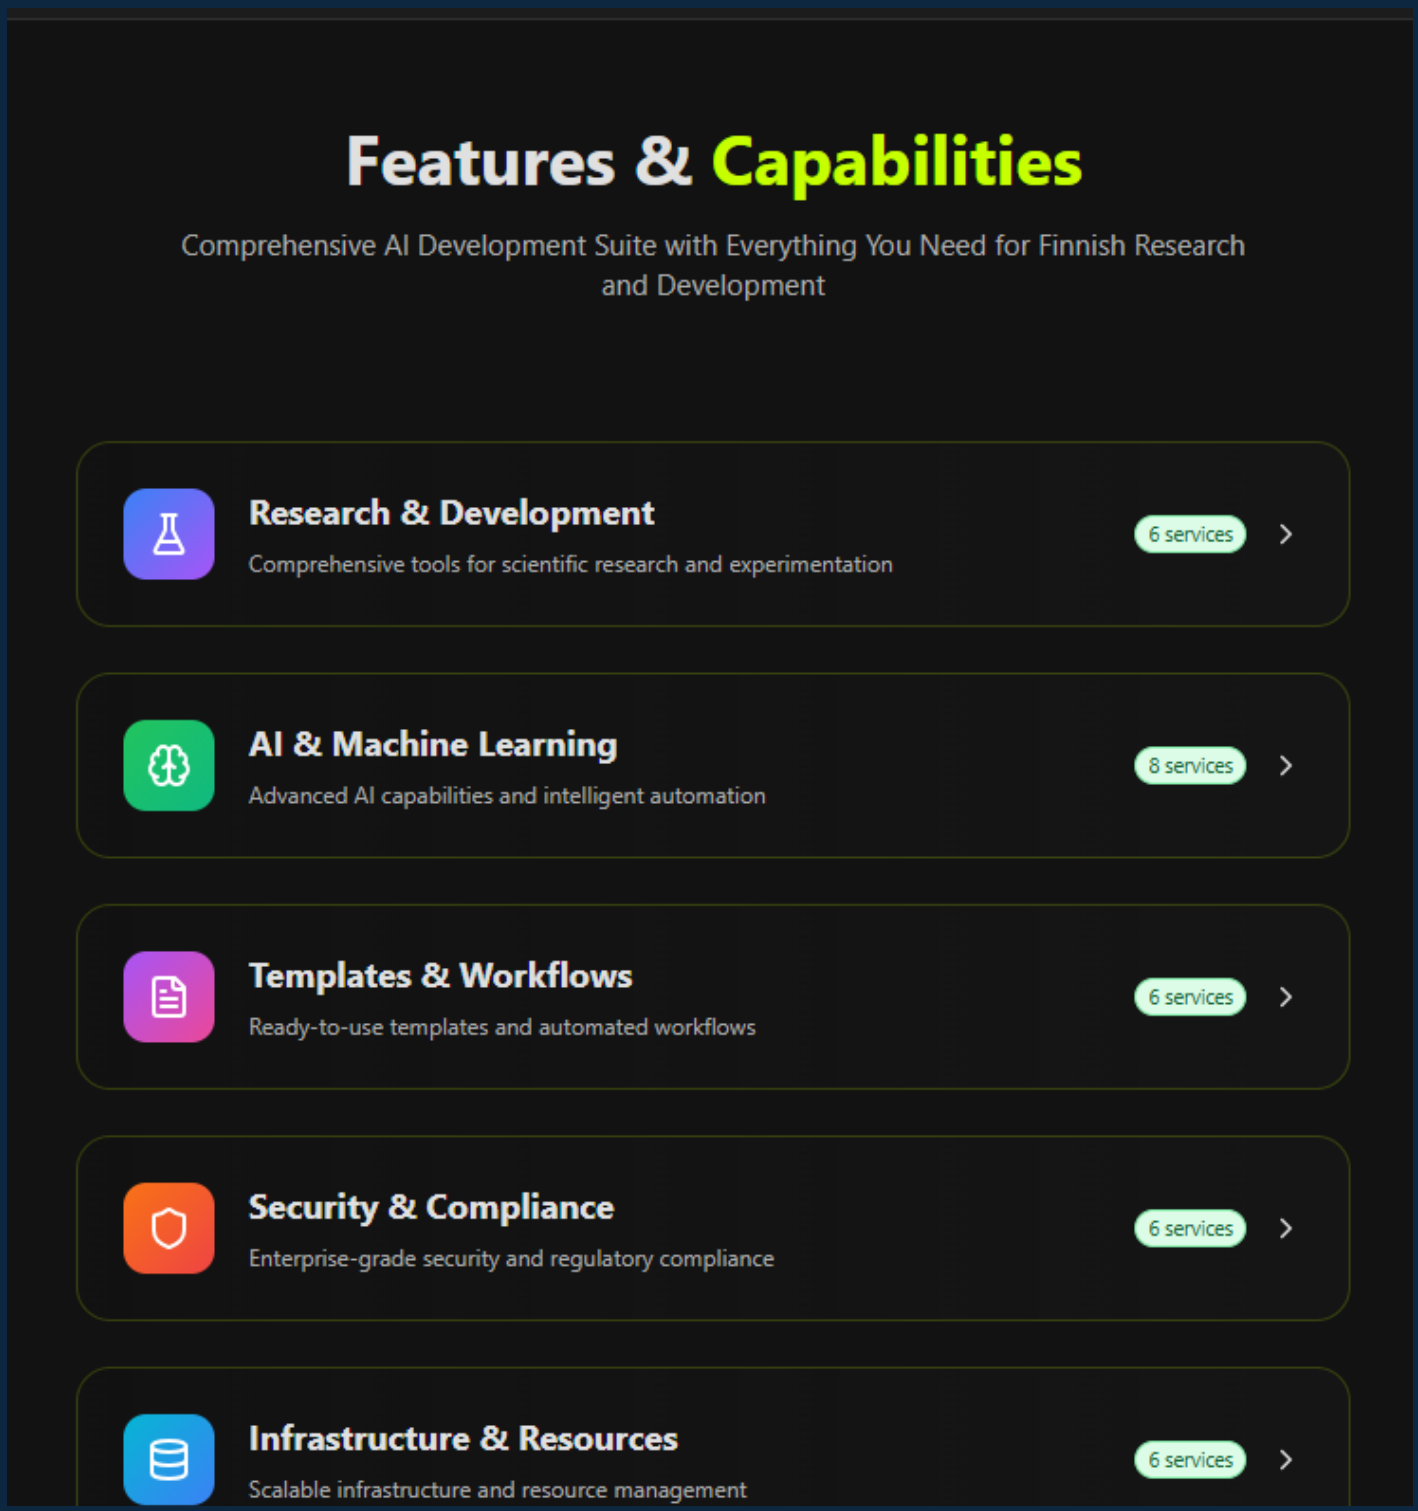

List of features and capabilities

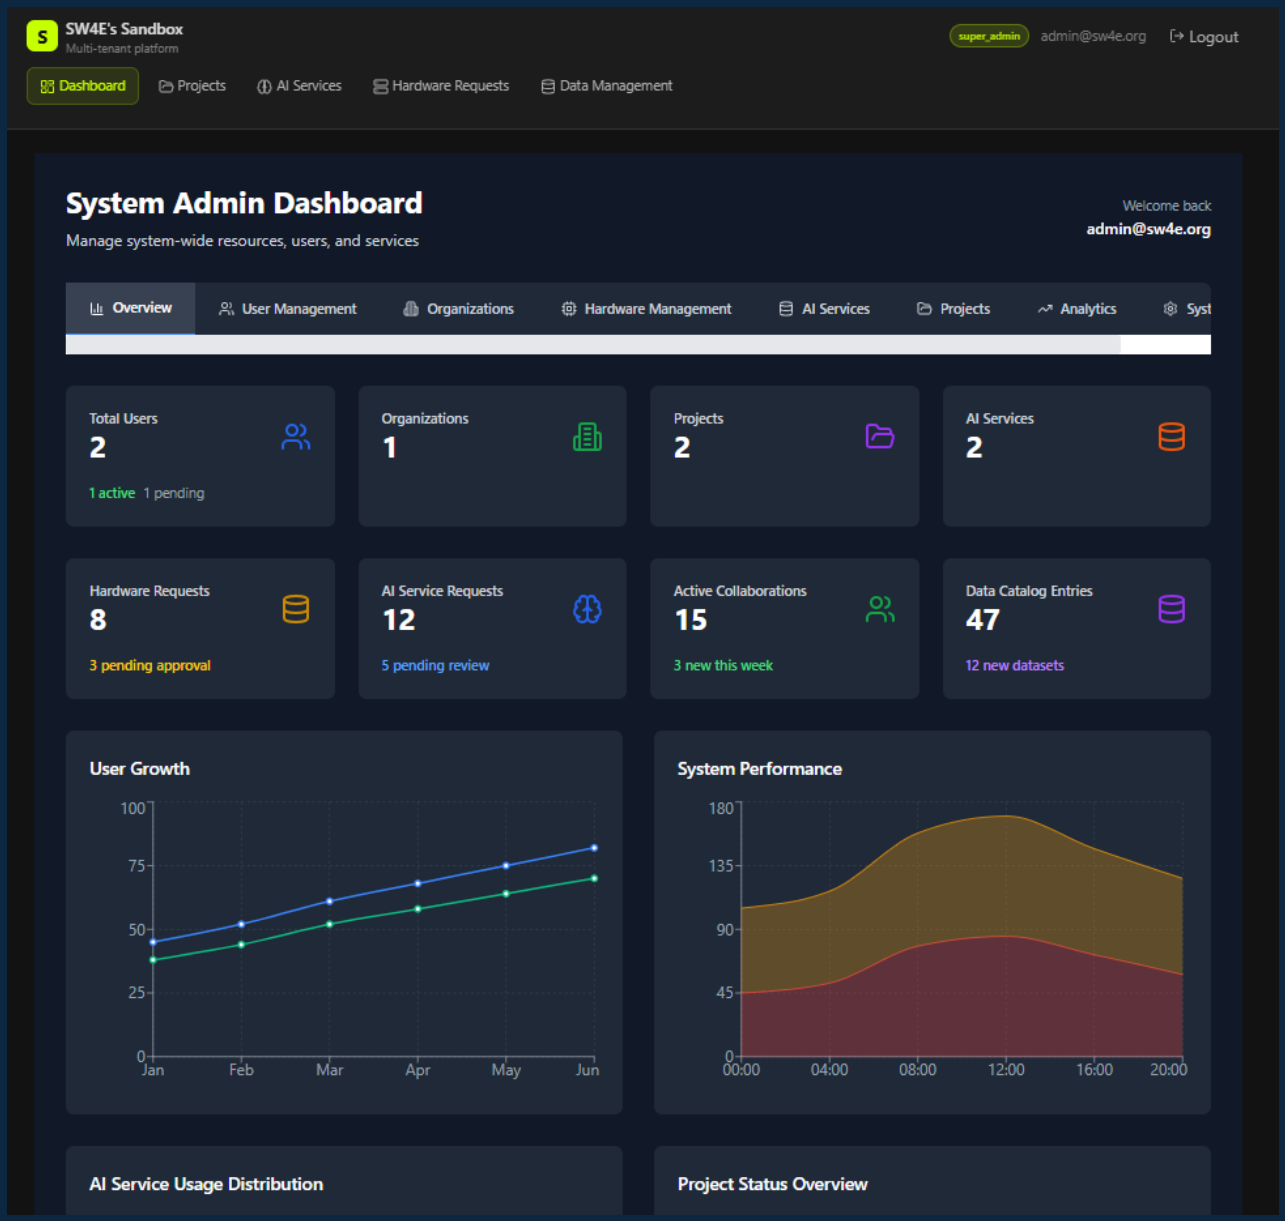

Admin dashboard view

9. FROM PROTOTYPE TO PARTNERSHIP

The AI Sandbox has proven that trusted and transparent AI collaboration is not only possible but practical.

Its prototype demonstrates how research and industry can work together to test, benchmark, and validate AI systems under shared governance and full compliance visibility.
The next phase is to translate this success into lasting partnerships.
With the core environment in place, the focus now shifts toward building an active user community, broadening real-world applications, and securing the resources needed for sustained growth.

How to Join

Organisations can participate through three simple pathways:

1. **Member Access**
   Join the SW4E ecosystem to use Sandbox resources, attend working groups, and access national networks.
2. **Pilot Project Collaboration**
   Propose a joint R&D or benchmarking project to explore AI solutions under trusted governance.
3. **Strategic Partnership**
   Co-invest in the further development and operation of the Sandbox as a national capability.

Participation is open to companies, universities, and public organisations that share the goal of advancing responsible and high-impact AI innovation.

Closing message

The AI Sandbox is more than a project — **it is an evolving platform for Finland's AI future.**
By combining technical excellence, governance clarity, and a shared vision of trust, it enables organisations to innovate boldly and responsibly in an increasingly complex digital world, and to achieve faster experimentation and development cycles than before possible.

dimecc.com

sw4e.fi