Case Study for Parallel Sessions

Project ARGUS: Governing AI in the State of IOTA

BACKGROUND

The State of IOTA Revenue Authority (SIRA) is widely regarded as a leader in digital transformation. Facing increasing workloads, reduced budgets, growing taxpayer expectations and pressure to improve efficiency, SIRA has invested heavily in artificial intelligence (AI) to support both operational processes and strategic decision-making.

Over several years, AI has become embedded throughout the organisation through a combination of:

- Internally developed AI: internally developed solutions, including knowledge assistants, audit-support tools and document analysis systems.
- Licensed AI: approved enterprise services such as Microsoft 365 Copilot and other centrally governed AI platforms.

The benefits have been substantial. AI is now used to analyse legislation, prepare audit briefs, draft correspondence, review large datasets and support case management. Productivity has improved significantly, and management considers AI essential to the future of tax administration.

At the same time, adoption has evolved faster than governance. SIRA has an 'AI-First' policy, and new or redistributed tasks are given to AI whenever feasible.

Developers increasingly rely on AI coding assistants and often test new open-source models before formal assessments can be completed. Business units experiment with innovative AI solutions to address operational needs. Employees have embraced AI-enabled wearable devices, including smart glasses and voice assistants, allowing them to access information, receive summaries and interact with AI systems while commuting or travelling.

Meanwhile, suppliers continue to introduce new AI capabilities. Existing contracts are amended to include AI functionality, new sub-processors are engaged, and software updates quietly enable additional AI features that are often undetected by the security and legal departments. Some new services rely on infrastructure and processing arrangements that differ from those originally approved. For example, they process data outside of the "EU data fence", in direct conflict with requirements set by the legal department. Many of the changes to the contracts are not discovered until after the technology has been implemented.

As a result, SIRA is struggling to manage:

- Unapproved or embedded AI: AI capabilities introduced through updates to existing software products and services.
- Shadow AI: external AI tools used by employees, consultants and contractors outside established governance processes.
- Agentic AI: emerging systems capable of gathering information, interacting with multiple systems and performing actions with increasing autonomy.

SIRA remains undeterred and is preparing the next phase of its digital transformation programme. After all, AI is the future.

Using advanced analytics, machine learning and agentic AI, the organisation plans to improve its ability to identify audit targets and allocate resources more efficiently. Several pilot projects are already exploring how AI can assist in selecting potential control objects, combining information from multiple sources to identify cases that may warrant further review. Senior management views these capabilities as critical to the future of risk-based tax administration and expects their use to expand significantly over the coming years.

A recent internal review raised a troubling question:

Can SIRA confidently identify every AI system processing taxpayer information, where data is processed, and which suppliers are involved?

No single department could provide a complete answer. Responsibility remained unclear. IT pointed to Legal. Legal pointed to Security. Security referred the issue to Corporate Governance. Everyone pointed to the DPO.

THE INCIDENT

To further modernise operations, SIRA launches **Project ARGUS**, an AI-powered assistant for auditors, investigators and case handlers.

ARGUS combines internally developed AI, enterprise AI services, supplier-provided functionality and agentic workflows in a single platform. It can retrieve information from multiple systems, analyse documents, draft reports, recommend actions and initiate approved workflows without human intervention.

The pilot exceeds expectations and is rapidly expanded across the organisation.

Six months after deployment, an experienced auditor receives an AI-generated briefing prepared by ARGUS. The report appears accurate and highly relevant.

The information in the report is confidential and contains details on several taxpayers suspected of tax avoidance. It also appears to contain information obtained from third parties, including banks, through legal processes. The ‘case manager’ that signed the legal documents compelling disclosure from the banks, Peder Ås, does not exist in any of SIRA’s employment records.

Some of the information can be traced to "oversharing" in the organisation's M365/SharePoint workspaces. This is a known problem, but it has been difficult to move all documents that should have more limited access out of areas that are accessible to the entire organisation. It is unknown how other information was obtained.

A personal data breach is reported, and an investigation is launched internally. The national data protection authority has also started an audit.

The Director General therefore tasks three expert groups with developing recommendations for the Executive Board.

Parallel Session 5A: Strategy, Governance and Compliance

Questions for discussion

1. What are the most significant governance and accountability failures revealed by the ARGUS incident?
2. How should SIRA govern and oversee different forms of AI across the organisation, including internally developed AI, licensed AI, unapproved or embedded AI, shadow AI and agentic AI?
3. What governance measures should the Executive Board implement to maintain visibility and control over an increasingly complex AI ecosystem?

Parallel Session 5B: Incident Readiness, Response and Collaboration

Questions for discussion

1. What are the main security and privacy risks and incidents associated with the ARGUS project and SIRA's use of AI more generally?
2. How should SIRA strengthen its ability to detect, investigate and respond to AI-related security and data protection incidents?
3. What roles should suppliers, regulators and other stakeholders play before, during and after an AI-related incident?

Parallel Session 5C: People, Technology and Secure Innovation

Questions for discussion

1. How can SIRA continue to realise the benefits of AI while maintaining appropriate safeguards and human oversight?
2. What risks are presented by shadow AI, AI-enabled wearables, developer experimentation and agentic AI, and how should they be managed?
3. What technical, organisational and workforce capabilities are needed to support secure and responsible AI adoption at scale?