
A MINIMUM VIABLE
AI
GOVERNANCE
(MV-AIG)
BY
PIKOM
OVERSIGHT COMMITTEE

Published by:

PIKOM
E1, Empire Damansara,
No. 2, Jalan PJU 8/8 A, Damansara Perdana
47820 Petaling Jaya, Selangor
T : +(603) 7622 0079
E : info@pikom.org.my
W : www.pikom.org.my

Release date: June 2026

Disclaimer:

All information furnished in this publication is provided strictly on an 'as is' and 'as available' basis and is so provided for your information and reference only. With this caution, kindly be informed that this release is not presented to address the circumstances of any particular individual or entity. As such, PIKOM including their sponsors, partners and associates, whether named or unnamed, do not warrant the accuracy or adequacy of the data and findings. Moreover, all parties concerned explicitly disclaim any liability for errors or omissions or inaccuracies pertaining to the contents of this publication. Therefore, the use of data and findings presented in this publication is solely at the user's risk. PIKOM shall in no event be liable for damages, loss or expense including without limitation, direct, incidental, special or consequential damage or economic loss arising from or in connection with the data and / or findings published in this series. However, professional advice can be sought from the producers of this publication. Disclosure – the document was prepared with assistance of Microsoft Co-Pilot.

Copyright

A MINIMUM VIABLE AI GOVERNANCE (MV-AIG)

Table of Contents

Message from the PIKOM Chairman

"Embracing Responsible Innovation"

Adj. Practice Prof. Alex Liew

As we navigate the transformative landscape of 2026, Artificial Intelligence has transitioned from a futuristic concept to a routine operational engine within our industries. While this shift unlocks unprecedented productivity, it also brings a new set of responsibilities that we must address with both urgency and pragmatism.

The **Minimum Viable AI Governance (MV-AIG)** framework, spearheaded by our Oversight Committee, represents a landmark shift in how we approach technology oversight. We recognize that for many organizations—especially our vital SMEs—complex compliance can feel like a barrier to innovation. This model changes that narrative by proving that effective oversight does not require massive infrastructure; it requires clear accountability and a proactive culture.

Why MV-AIG Matters Now?

- **Risk Mitigation**:
  Most AI-related issues arise from well-intentioned staff using tools without standardized guidance; this framework provides that necessary guardrail.
- **Practicality over Complexity:**
  By focusing on high-impact scenarios and providing streamlined templates, we ensure that ethical AI use becomes a seamless part of daily workflows.
- **Scalable Foundation:**
  Whether you are a micro-enterprise or a large corporation, this model serves as a credible starting point to build toward more sophisticated governance in the future.

I want to extend my gratitude to the **PIKOM Oversight Committee (2026)**, and specifically **Dr. Dzaharudin Mansor**, for their visionary work in initiating this study and managing the sandbox implementation.

We encourage every member organization to take full cognizance of this report. By adopting the MV-AIG model, we can move together toward a future of confident, governed, and truly responsible AI adoption.

Message from Chair of PIKOM Oversight Committee (2026)

Woon Tai Hai

The **PIKOM Oversight Committee** serves as a critical governance mechanism designed to ensure transparency, accountability, and strategic alignment within the association. By reporting directly to the **PIKOM Council**, it acts as a "check and balance" entity for the organization's operations and financial integrity.

Our mandate is to ensure that while we embrace the rapid transformative power of Digitalization including Artificial Intelligence, we do so with an unwavering commitment to integrity and strategic discipline. Hence supported by independent expertise, this inaugural report was initiated and published by the Committee.

Minimum Viable AI Governance (MV-AIG)

AI has shifted from an experimental novelty to a routine operational tool, deeply embedded in activities like data analysis and meeting transcription. While this integration drives productivity, it introduces risks that many organizations—particularly smaller ones—are only beginning to address. Most AI-related harm stems not from malice, but from well-intentioned staff using familiar tools without adequate guidance.

The Need for Proportional Governance

For SMEs and micro-enterprises, where a single security breach can be existential, structured governance is no longer optional. Responsible AI practices must become an immediate operational foundation rather than a future goal.

The **Minimum Viable AI Governance (MV-AIG)** model is a practical, lightweight framework designed for organizations with limited resources. Recognizing that most entities are AI **users** rather than developers, the model prioritizes simplicity and real-world applicability.

Key Framework Components

The MV-AIG distils complex ethical principles—fairness, transparency, and accountability—into actionable daily workflows. It simplifies compliance through:

- **Targeted Scenarios**: Focusing on high-impact areas like document processing.
- **Practical Tools**: Providing usage inventories, risk tiering, and basic incident protocols.
- **Streamlined Policy**: Consolidating requirements into templates that integrate into existing routines.

National Alignment and Impact

This work builds on **PIKOM's** national leadership in AI ethics, specifically the 2024 and 2025 policy papers. It proves that effective oversight does not require large teams or elaborate infrastructure; it requires clear accountability and a culture of awareness.

The MV-AIG model serves as a credible entry point for small organizations and a foundational building block for larger ones to scale toward more sophisticated frameworks. By adopting this model, organizations in Malaysia and beyond can move from hesitation to confident, governed AI adoption.

Finally, we would also like to thank PIKOM Oversight Committee (2026) for this initiative and publication in particular **Dr Dzaharudin Mansor** who not only initiated this study but is responsible for the sand box implementation. I would also like to extend my appreciations to Ms Nurul Asyiqin Nasir, PIKOM Head of Strategic Relations & Communications for developing the layout and design of the publication, for ease of reference and reading. We strongly encourage all organizations to take cognizance of this report content and do direct any queries to PIKOM.

"Effective oversight does not require large teams or elaborate infrastructure; it requires clear accountability and a culture of awareness including in an AI governance's deployment"

PIKOM OVERSIGHT COMMITTEE:

WOON TAI HAI (CHAIR)

DANNY LEE

SHAIFUBAHRIM SALEH

DR DZAHARUDIN MANSOR

Executive Summary

Lead Researcher

Dr Dzaharudin Mansor

The Immediate Challenge

As AI seamlessly integrates into daily operations, the risks of ungoverned usage have shifted from theoretical to immediate. Recent high-profile data leaks—from proprietary code exposed in chatbots to sensitive clinical meetings recorded via personal accounts—highlight a critical reality: AI incidents are rarely malicious, but rather the result of convenience and a lack of clear boundaries. For small and micro-organizations, where a single oversight can cause disproportionate reputational or legal harm, implementing a "do nothing" approach is no longer sustainable.

The Solution - MV-AIG

The **Minimum Viable AI Governance (MV-AIG)** model offers a pragmatic, lightweight alternative to heavy corporate bureaucracy. Specifically designed for organizations acting as AI users with a low-risk appetite, it focuses on the two most common, high-risk activities: document analysis and meeting transcription. Built on the tenets of proportionality and simplification, MV-AIG translates high-level standards—such as **PIKOM's AIEG** and **MOSTI's AIGE**—into actionable staff commitments that prioritize fairness, privacy, and transparency.

Core Operational Mechanisms

The "Safe Path" Inventory:
A single, definitive list of approved AI tools. If a tool is not on this list, staff cannot use it with organizational data.

Three-Question Procurement Gate:
Before adopting any new tool, the organization must verify: Where does the data go? Who owns the output? What is the human fallback?

Risk Impact Classification:
AI tools are classified into Low, Medium, or High impact by assessing factors such as decision influence and reversibility, human oversight, data sensitivity, affected population and vulnerability, regulatory considerations, and autonomy level, with each tier determining the required level of governance, risk assessment, and approval.

Strategic Value

MV-AIG **translates high-level national standards**—specifically *MOSTI's AIGE* and *PIKOM's AIEG*—**into immediate office habits.** It demonstrates that responsible AI adoption does not require an exhaustive compliance department; it requires clarity, culture, and a small number of well-designed mechanisms that empower teams to innovate with confidence.

"Responsible AI adoption does not require an exhaustive framework; it requires clarity, culture, and a small number of well-designed mechanisms."

1. INTRODUCTION

"AI governance is no longer a luxury for large enterprises; for small organizations, where a single mistake can have a disproportionate impact, it is a functional necessity."

Artificial intelligence is now embedded in everyday work, and even small organizations rely on it to boost productivity—often without realizing the real risks when its use goes ungoverned. Recent incidents[iv] show that well-intentioned shortcuts can quickly lead to legal, operational, and reputational harm, making even minimal AI governance essential for organizations where a single mistake can have outsized impact.

1.1 Shadow AI

Ungoverned AI usage poses immediate and significant risks to organizations of all sizes, regardless of whether they perceive their operations as too small or their use cases as trivial. The rise of Shadow AI, the unauthorized use of AI tools by employees, has created a new frontier for business, legal, and reputational exposure. When staff utilize pervasive AI tools without formal guidance, even routine tasks like document summarization or meeting transcription can inadvertently lead to the leakage of confidential information or non-compliance with data privacy laws.

1.2 Strategic Context: The AI Governance Imperative

The regulatory landscape is shifting rapidly. With emerging guidelines and legislation relevant to AI, such as the Malaysian National Guidelines on AI Governance and Ethics (AIGE) [i], Malaysian Cybersecurity Act 2025[ii] (CSA 2025) and new regulations associated with the Personal Data Protection Act 2010[iii], organizations that fail to implement structured oversight risk legal consequences and a profound loss of stakeholder trust. Unlike traditional IT systems, AI risks present unique challenges:

- **Impact Velocity:** Risks can escalate from detection to full-scale impact much faster than expected.
- **Feedback Potential:** Behavioural and technical feedback loops can amplify minor errors, making them difficult to contain once they enter a production environment.

Real-world evidence[iv] suggests that AI-related harm rarely stems from malicious intent. Instead, it typically arises from everyday habits—choosing convenience over security or relying on personal accounts for professional tasks. Without clear rules, these small actions can snowball into major breaches.

While specific industry incidents involving leaked source code and unauthorized clinical recordings illustrate these dangers vividly, they underscore a central truth: AI governance is no longer a luxury for

large enterprises. For small organizations, where a single mistake can have a disproportionate impact, a minimal governance framework is a functional necessity.

1.3 AI Risk versus Cybersecurity Risks

AI risks and cybersecurity risks both affect AI-driven digital systems, but they arise from different sources and call for different approaches. Cybersecurity is focused on protecting systems, networks, and data from threats such as malware, phishing, ransomware, and breaches. AI risks, by contrast, stem from how AI systems are designed, trained, and behave, including issues like data poisoning, model manipulation, bias, hallucinations, and opaque or unintended decision-making. In essence, cybersecurity defends systems from attack, while AI risk addresses whether the system itself can be trusted. The boundaries, however, are not rigid: AI depends on secure data, infrastructure, and access controls, while also introducing new vulnerabilities and amplifying traditional cyber threats. As a result, the two areas overlap significantly, with MV-AIG focusing on managing AI-specific risks while integrating into established governance frameworks, including cybersecurity governance.

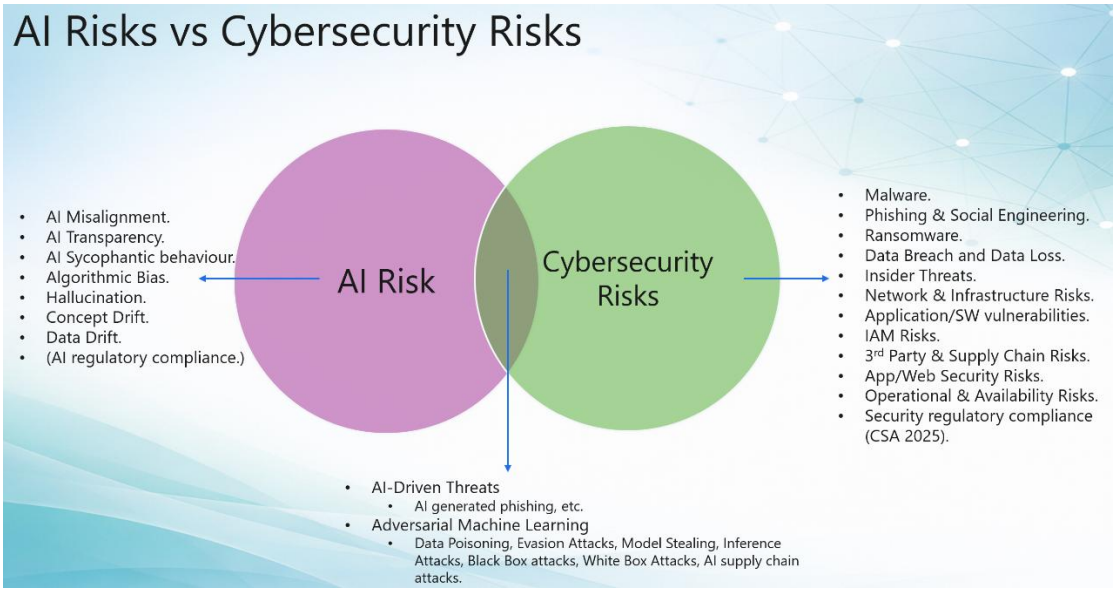

1.4 Scope, Applicability, and Exclusions

This framework is primarily designed for organizations acting as AI Users rather than developers. It is applicable to all employees, contractors, and third-party partners who utilize AI-powered tools for business purposes.

The focus of this "Minimum Viable" approach is intentionally narrow, prioritizing high-frequency activities such as document analysis and meeting transcription. Specialized AI development, high-stakes autonomous systems, and highly regulated sector-specific AI deployments may require oversight beyond the scope of this baseline framework.


The Need for AI Governance
Why Oversight of AI is Essential for All Organizations
Shadow AI Risks
Data Leaks
Compliance Violations
Reputational Damage
Regulatory Pressure
New AI Laws
Cybersecurity Regulations
Privacy Concerns
Key Challenges of AI
Rapid Impact
Fast Escalation of Risks
Feedback Loops
Amplified Errors & Bias
AI Governance is Essential
Minimize Legal, Operational & Reputational Risks
Essential for Organizations of All Sizes
Clear Oversight is Critical to Safe AI Use

2. THE MV-AIG FRAMEWORK: TENETS AND ASSUMPTIONS

"Creating just enough governance structure to keep the use of AI tools safe and responsible without slowing people down."

The Minimum Viable AI Governance (MV-AIG) described here is about creating just enough governance structure to keep the use of basic and widely used AI chatbot and AI transcription tools safe, responsible, and aligned with organizational goals without slowing people down. These two classes of AI use cases were selected because most organizations use these. The goal is not bureaucracy, but clarity and confidence in using these AI powered products. MV-AIG is meant for small and micro-organizations. However, we believe that larger organizations without AI Governance in place but wants to start with a relevant but minimalist approach before embarking on a more comprehensive one, may benefit from MV-AIG.

MV-AIG is grounded in a set of guiding tenets and assumptions that shape a practical, streamlined governance model. The strategies, methods, and resources presented here draw directly from an active, real-world implementation of AI governance within a small organization, combined with the pragmatic, dynamic approach advocated by AI Career Pro[v] —an approach intentionally designed to address real-world AI risks rather than theoretical ones. This foundation is further strengthened by the author's hands-on experience in AI governance roles at Microsoft and AWS.

2.1 Core Tenets for Small & Micro-Organizations

Tenets[vi] are a set of believes adopted by a team or organization that clarify and align people on what is important and what, by extension, isn't. Tenets help to accelerate decision making especially when there are differing and often conflicting points of views or interpretation of data that can lead to unnecessary decision bottlenecks. Tenets are being used here to help the AI Governance team to make faster, clearer, and more consistent decisions by giving everyone a shared understanding of what truly matters.

The tenets adopted for the development of MV-AIG are as follows:

1. **Proportionality — Governance that fits your size**	AI governance should match the scale and complexity of the organization. Small teams don't need enterprise-level frameworks. They need simple, clear rules that protect them without overwhelming them.
2. **Start Small — Focus on the highest-impact areas first**	Instead of trying to govern every possible AI scenario, begin with the tools and use cases

	that matter most—typically document analysis and meeting transcription.
3. **Practicality Over Perfection**	Governance must work in real life. Policies that are too complex or idealistic will be ignored. Practicality ensures adoption.
4. **Integrate, Don't Add Burden**	AI governance should plug into existing processes—information security, procurement, HR policies—rather than creating parallel structures.
5. **Simplify — Make the right thing the easy thing**	Good governance removes friction. When safe AI use is simpler than unsafe use, compliance becomes natural.
6. **Avoid Bureaucracy**	Small organizations cannot afford heavy administrative processes. Governance must be lightweight, fast, and easy to maintain.
7. **Single-Threaded Ownership (STO)**	Clear ownership prevents confusion. One person—usually a senior leader—acts as the anchor for AI governance, supported by others as needed.
8. **Just Culture — Reward, Coach, Sanction**	People should feel safe reporting mistakes or uncertainties. A just culture encourages learning, not blame, while still enforcing accountability when needed.

Sanctions are reserved for repeated, wilful violations of the "Prohibited Use Cases," rather than honest mistakes. Organisations must practice the culture of carrot and stick to adherence of in the prohibition Use cases!

Ref: 8. Just Culture – Reward, Coach, Sanction

2.2 Defining the "AI User" vs "AI Developer" Mindset

Assumptions help to scope the initial focus of AI Governance to what is important and immediately relevant for the organisation in line with "Start Small" and "Simplify" tenets. Without defining the assumptions, the policy will be too generic, hard to interpret and difficult to apply in practice. If the small to micro-organization begins developing AI models, expands use cases, or handles higher-risk data, the governance framework must be reviewed and strengthened. Several assumptions that shape the scope and design of an initial MV-AIG are as follows:

1. **The organization is primarily a user of AI, not a developer**: This means governance focuses on safe use of third-party tools rather than model development or training.

2. **AI use is currently limited widely used high-value tasks**: Most organizations, including small and micro-organizations, today use AI for:
   - Document analysis and preparation
   - Meeting transcription

   This restricted scope makes governance easier to implement and reused for other similar organizations.

3. **The organization has a low-risk appetite**: For the limited uses cases mentioned above, a conservative approach is appropriate—especially when handling sensitive information or public-facing content.

4. **Principles are aligned with established AI ethics frameworks**: The governance model builds on principles that are the organization are prepared to adopt. However, well established and globally recognized principles can be used as a starting point. Examples are those defined in the "PIKOM AI Ethics & Governance 2025" (PIKOM AIEG) publication[vii] and AIGE[i]. They include globally recognized principles such as fairness, accountability, transparency, and privacy.

3. MINIMUM VIABLE AI GOVERNANCE POLICY

" Consolidating complex requirements into a single, accessible 'Streamlined' AI Governance Policy that integrates into existing routines."

Policy documents in governance act as formal, written instruments that define an organization's or government's rules, principles, and strategic intentions. They provide authoritative guidance, ensure consistent decision-making, ensure compliance with legal requirements, and manage risks by clarifying expectations and roles. It is recommended that AI Governance for a large organisation requires three different policy document each tailored for different stakeholders namely AI Governance Policy, AI Risk Management Policy and AI Use Policy.

However, for small to micro-organizations and simple implementation that MV-AIG is designed for, these three policies are simplified and merged into ONE document, the "Streamlined" AI Governance Policy[viii] (moving forward this will be simply referred to as the MV-AIG AI Policy). Some of the key contents of the policy are as follows:

1. **Guiding Principles -** Outlines practical principles the organization has decided to adopt. This is detailed in section 3.2. Principles & Commitments This is to guide responsible AI use.
2. **Accountability -** Assigns responsibilities across committees, leaders, system owners, mechanism owners, and all employees. This is detailed in the sub section Accountability.
3. **Risk Screening and Classification -** Establishes mandatory classification of all AI systems into Low, Medium, or High Impact with corresponding requirements.
4. **AI System Inventory -** Requires maintaining a central inventory of all AI systems, including embedded features and supporting documentation. This is developed in the operationalization phase as described in section 7.2 Phase 2: Operationalization (Core Mechanisms).
5. **Acceptable Use Guidelines -** Defines permitted, prohibited, and cautionary AI uses to ensure safe and compliant adoption. This is further elaborated in section 3.1 Acceptable Use Guidelines.
6. **Procurement and Vendor Management -** Sets approval, assessment, evidence, contract, and oversight requirements for acquiring AI systems from vendors. This is to be established in the operationalization phase as described in section 7.2 Phase 2: Operationalization (Core Mechanisms).
7. **Human Oversight and Control -** Ensures humans retain meaningful oversight, especially for consequential decisions, and understand how to intervene. This is to be established in the operationalization phase as described in section 7.2 Phase 2: Operationalization (Core Mechanisms).

8. **Reporting and Support -** Provides channels and expectations for reporting AI incidents, seeking guidance, and offering feedback. This is to be established in the operationalization phase as mentioned in section 7.2 Phase 2: Operationalization (Core Mechanisms).
9. **Exceptions -** Describes how exception requests are submitted, evaluated, and time-limited. This is to be established in the operationalization phase as described in section 7.2 Phase 2: Operationalization (Core Mechanisms).
10. **Policy Review -** Establishes annual review requirements and approval processes for substantive revisions. This is to be initiated as part of maintenance and continuous improvement briefly mentioned in section 7.3 Phase 3: Maintenance & Continuous Improvement .

3.1 Acceptable Use Guidelines

The following guidelines outline what is permitted, what is prohibited, and what requires extra care when using AI tools. These guidelines apply alongside the principles in section 3.2. Principles & Commitments The Acceptable Use Guidelines here are derived from Streamlined AI Governance templates from AI Career Pro[ix].

Permitted Uses

AI tools may be used to enhance staff productivity if they are used in line with the principles in the Policy. Staff are encouraged to use AI to improve efficiency by automating repetitive tasks, summarising documents, drafting routine content, and generating insights, as well as to support creative brainstorming by producing first-draft ideas, text, code, or designs that will later be refined by humans. AI may also assist with data analysis and decision support by highlighting trends or patterns, provided the results are validated and human judgment is applied. For customer communication, AI can help research issues or draft responses, but staff remain fully responsible for accuracy and may only use tools approved for customer data. AI may also be used for learning and skill development, using only dummy or public data for practice. Small pilot projects and experiments are encouraged when conducted with test data and with a manager's awareness, but any move to real or customer-facing use requires approval from the AI Governance Lead. Overall, staff are encouraged to use approved AI tools in their work, following the organisation's authorised list.

Prohibited Uses

AI must never be used in ways that put confidential information, security, or individuals at risk. This means staff must not enter confidential, proprietary, or personal data into AI tools that are not approved, and anything typed into an external AI service should be treated as if it could become public. AI may not be used to bypass security controls, access unauthorised information, or engage in illegal or fraudulent activity. It must also not be allowed to make important decisions about people—such as those involving employment, credit, or safety—without proper human review. AI must not be used to create harmful, deceptive, or misleading content, including impersonation, deepfakes, harassment, or phishing. Staff must avoid generating content that infringes copyrights or other intellectual property rights and should seek advice if they are unsure about the source of AI-generated material. AI must not be used in ways that violate laws, regulations, or internal policies, and AI involvement must not be hidden when transparency is required. If an AI tool produces biased, inaccurate, or harmful outputs, staff must stop using it for that purpose and report the issue rather than ignoring or concealing it. Any violation of these

prohibitions may result in disciplinary action, and staff should seek guidance whenever they are uncertain about whether a use case is permitted.

Use with Extra Caution

Some uses of AI are allowed but require extra care because they can introduce risks if not handled properly. Extra caution is needed when external or unapproved AI tools are used, especially free or consumer versions, and only non-sensitive, publicly available information should be entered unless the tool has been formally approved; even then, outputs must be checked for accuracy. Any AI-generated content must be thoroughly reviewed and edited to ensure it is correct, appropriate, and aligned with the organisation's standards, and it should be labelled as AI-assisted when transparency is required. When AI tools are integrated into existing workflows, they should be tested on a small scale first, with attention to data handling, security, and compliance, and IT should be consulted if shared systems or data may be affected. If personal devices or personal AI accounts are used for work, the same rules apply—personal and work data must not be mixed, and organisation-approved tools are preferred. Additional scrutiny is also required when AI is used in regulated areas, with sensitive customer information, or in high-stakes situations, and staff should seek guidance from compliance or legal teams whenever they are unsure.

3.2. Principles & Commitments

With the tenets and assumptions in place, the next step is defining the principles that guide responsible AI adoption — and the practical commitments that make those principles real. MV-AIG doesn't require complex frameworks. Instead, it focuses on a small set of clear, actionable principles that help staff use AI confidently and safely. It is important that the principles and commitments are tailored for the organization, as well as agreed by the organization leadership including the CEO.

For illustration for this report, the principles listed here are based on the principles defined in PIKOM AI Governance and Ethics (AIGE)[vii] that an organization has decided to adopt. Each principle needs to be translated to commitments that an organization pledge to practice that should be co-created and agreed by organization leadership lead. For example, each of the seven core principles in PIKOM AIGE can be paired with a simple commitment may be as follows:

- **Fairness**

AI should support work without introducing bias, misrepresentation, or inappropriate language. Even simple tasks like summarizing documents can unintentionally distort meaning or tone.

Commitment:

> Staff review AI-generated content—summaries, drafts, transcripts—to ensure it is accurate, respectful, and free from biased or misleading language.

- **Reliability, Safety & Control**

AI outputs are not final answers—they are starting points. Treating them as drafts ensures that errors, hallucinations, or misinterpretations do not slip into official work.

Commitment:

> Important facts, figures, and interpretations must be verified before use. Staff should fact-check AI-generated figures against original source documents and meeting transcripts are checked for accuracy before distribution. If an AI tool behaves unpredictably, staff stop using it and report the issue.

- **Privacy & Security**

AI tools process information, and that information must be protected. Small organizations are especially vulnerable to accidental data leakage through unapproved or consumer-grade AI tools.

Commitment:

> Only approved AI platforms may be used for documents or meeting recordings. Sensitive or personal data must never be uploaded into unapproved tools. Staff follow existing data protection and information security policies when using AI.

- **Inclusiveness**

AI should empower everyone—not just the tech-savvy. When tools are accessible and guidance is available, the entire organization benefits.

Commitment:

> All relevant staff receive basic guidance or training. AI-generated documents and transcripts must be clear and accessible to their intended audience. Anyone unsure about using an AI tool can request support.

- **Transparency**

Being open about AI use builds trust internally and externally. Transparency also reinforces accountability and helps others understand how content was produced.

Commitment:

> When sharing drafts or summaries created with AI, staff indicate that AI assisted in generating the content. They can explain, in simple terms, what the AI did and what humans edited. Final documents always reflect human review and approval.

- **Accountability**

AI does not replace responsibility. Humans remain fully accountable for all work produced with AI assistance.

Commitment:

> Staff review and approve all AI-generated content before using it. AI tools are not used to make final decisions that affect people or the business. It is possible to invoke another AI agent to check and validate the authority of the reported ones in addition to the human in the loop process. Any issues or concerns with AI outputs are reported promptly.

- **Pursuit of Human Benefit**

AI should enhance work, not replace judgment or create harm. The goal is to improve quality, reduce manual effort, and support better communication.

Commitment:

> AI is used to improve productivity and document quality. Staff ensure that AI-generated summaries or documents accurately reflect the original content. Any potential negative impacts on staff, customers, or partners are raised for review.

4. AI GOVERNANCE INTEGRATION: HARMONIZATION WITH EXISTING POLICIES

"Making the 'safe path the easy path' by building on what already exists rather than introducing parallel bureaucracies."

For a Minimum Viable AI Governance (MV-AIG) model to be effective, it must move beyond policy and become part of the organization's daily operations. This chapter outlines the clear roles required for accountability and the practical steps to embed AI governance into the existing organizational culture.

4.1 Roles: From Governance Leads to AI Users

Clear ownership prevents confusion and ensures that governance is not a "side project" but a core business function.

- **AI Governance Lead:** A senior staff member who anchors the program. They oversee policy rollout, maintain the AI inventory, coordinate training, and serve as the primary point of accountability.
- **AI Governance Committe**e: A small group of senior leaders who provide high-level oversight. They review and approve key documents, evaluate high-risk use cases, and ensure alignment with organizational values.
- **AI System Owners:** Individuals responsible for specific AI tools. Their duties include approving tool usage, ensuring secure configurations, and monitoring for misuse.
- **Mechanism Owners:** Assigned to maintain specific processes such as the AI inventory, incident handling, or policy update cycles.
- **AI Users:** All staff, contractors, and partners using AI tools. They are responsible for following the Acceptable Use Guidelines and reviewing AI-assisted content for accuracy and fairness.

In micro-enterprises, the Single Threaded Owner (STO), AI Governance Lead, and System Owner may be the same individual to avoid the impression that a large committee is required. STO is a function or role not a title or designation

4.2 Cultural Integration: Embedding AI into Daily Workflows

The goal of MV-AIG is to make the "safe path the easy path". This is achieved by building on what already exists rather than introducing parallel bureaucracies.

- **Build on Existing Policies:** AI governance should plug into established Information Security, Procurement, HR, and Privacy policies.
- **Use Existing Governance Bodies:** Wherever possible, use existing management meetings or oversight committees to review AI risks rather than forming entirely new entities.
- **Align with Operational Cycles:** AI policy reviews and risk screenings should be timed with existing annual budget or security review cycles to ensure consistency without adding administrative burden.
- **Inventory Management:** Maintain a simple, centralized AI Inventory as the "single source of truth" for all authorized tools, owners, and risk classifications.

5. RISK MANAGEMENT & OPERATIONAL MECHANISMS

"Every AI tool must be checked and classified based on impact to understand the safeguards needed before it becomes part of our work."

To keep everyone safe and to make sure we use AI responsibly, every AI tool used must be checked and classified before it becomes part of our work. For AI, the risk screening is based on impact[x]. This helps us understand the level of risk involved and what safeguards we need to put in place. Some AI tools are simple and low-risk, like tools that help draft documents. Others can affect people's rights, finances, or access to services. Some simple low-risk tools can become high risk when used to process sensitive or personal information. Because of this, we classify each AI system that involves the type of tool and in relevant cases, the classification of data that the AI tools process. The clear classification will allow the organization to determine how carefully it needs to be managed.

5.1 Impact Assessment: Low, Medium, and High-Risk Tiers

All AI systems must be classified before deployment. AI systems are grouped into three levels based on how much impact they could have.

Low Impact

Low impact tools are those that help with simple tasks and don't process sensitive, personal or data bounded by regulations. They don't make important decisions and have very little potential to cause harm.

Examples:

- Drafting tools for non-sensitive information.
- Document summarisation using non-sensitive information.
- AI transcription tool used for non-sensitive meetings and discussions.

What's required:

- Register the tool
- Use it responsibly

Medium Impact

These tools handle more sensitive situations. They may use personal or sensitive data, or their outputs may influence decisions that affect many people.

Examples:

- AI tools that process personal data with explicit consent by the data subject.
- AI used for large-scale marketing.
- AI transcription tool used for sensitive meetings and discussions.
- AI tool being used to draft internal sensitive information.

What's required:

- A documented risk assessment
- Approval from the AI Governance Lead

High Impact

These tools can significantly affect regulatory compliance, people's rights, finances, employment, or access to services. They may also involve highly sensitive information.

Examples:

- Use of AI tools using information protected under the Official Secrets Act 2010 (OSA).
- AI tools being used to process personal data in a manner that is non-compliant with the Personal Data Protection Act 2010 (PDPA).
- AI that influences major decisions that impact lives.

What's required:

- Comprehensive risk assessment
- Executive-level approval
- Strong human oversight and ongoing monitoring

THE 60-SECOND RISK IMPACT ASSESSMENT CHECKLIST Before adding a tool to the **AI Inventory**, ask these questions:	
1. **Compliance:** Does this tool process government OSA data or other sensitive data without explicit consent (e.g., personal data such as NRICs, phone numbers, financial information, or copyrighted corporate materials like documents or code)?	
2. **Impact on Transparency & Rights:** Is there a risk that the AI could generate deceptive, biased, or harmful content that reaches the public or customers?	
3. **Decision Authority:** Will the AI's output be used to make a final decision about a person's employment, credit, or legal rights?	
4. **Processing Sensitive Corporate Information:** Is the AI tool used to process sensitive corporate information that, if exposed to a third party, will adversely impact the organization.	
IMPACT LEVEL	
If YES to 1 or 2	***HIGH IMPACT***
If YES to 3 or 4, NO to 1 or 2	*MEDIUM IMPACT*
If NO to all four	*LOW IMPACT*

5.2 Prohibited Use Cases

The use of AI that violates applicable law or presents unacceptable risks should not be allowed. Some uses of AI are simply too risky. We will not use AI that:

- breaks the law
- makes important decisions without human oversight
- violates privacy or data protection rules
- creates unacceptable risks to people or the organisation

6. CASE STUDIES: REAL-WORLD APPLICATIONS & INCIDENTS

"AI incidents are rarely malicious, but rather the result of convenience and a lack of clear boundaries."

In the landscape of AI adoption, the divide between intended use and actual outcome is often bridged by oversight gaps. Even when organizations implement baseline controls, a lack of comprehensive monitoring can lead to systemic failures. This chapter analyses high-profile incidents to illustrate how "convenient shortcuts" and "familiar tools" can escalate into significant regulatory and security breaches.

6.1 Analysis of Global AI Incidents

The following cases highlight how a lack of AI governance or the presence of incomplete monitoring, can expose sensitive organizational data.

a) *Proprietary Data Leak when using Generative AI: Incomplete Monitoring*[xi]

The Samsung case serves as a primary example of how sensitive information can be exposed when monitoring is insufficient.

- **The Incident:** Employees utilized generative AI tools to assist with coding and meeting notes, inadvertently feeding proprietary source code and internal meeting data into public AI models.
- **The Governance Gap:** While the organization had existing data security protocols, they were not adequately adapted to the specific "input-output" risks associated with public AI interfaces.
- **Outcome:** Intellectual property was absorbed into the AI's training set, representing a loss of corporate assets and a failure in oversight.

b) *Hospital privacy breach when using transcription tool: Absence of AI Governance*[xii]

The Otter.ai incident demonstrates the danger of allowing unmanaged personal AI tools into professional environments.

- **The Incident:** A personal AI transcription account was used to record confidential hospital meetings without the knowledge or consent of the participants.
- **The Governance Gap:** There was a total absence of AI governance, allowing a familiar tool to be deployed in a high-stakes, regulated environment (healthcare) without a risk assessment.
- **Outcome:** The unauthorized recording of sensitive medical discussions led to regulatory breaches and triggered mandatory reporting requirements.

6.2 Lessons from AI Incidents

These incidents provide critical insights into how AI risks manifest within small and large organizations alike. Here is a non-comprehensive list of the lessons that can be learned:

• **Non-Intentional Harm:** AI incidents rarely stem from malicious intent; instead, they are the result of employees taking shortcuts or assuming that "nothing bad will happen" when using familiar tools.

• **The Velocity of Risk:** What begins as a simple oversight—such as a single employee using an unapproved app for convenience, can escalate rapidly into a mandatory reporting event.

• **The Need for Integrated Oversight:** Organizations must move beyond "set and forget" policies. Continuous feedback loops and active monitoring are essential to identify when people rely on familiar tools in ways that bypass official security measures

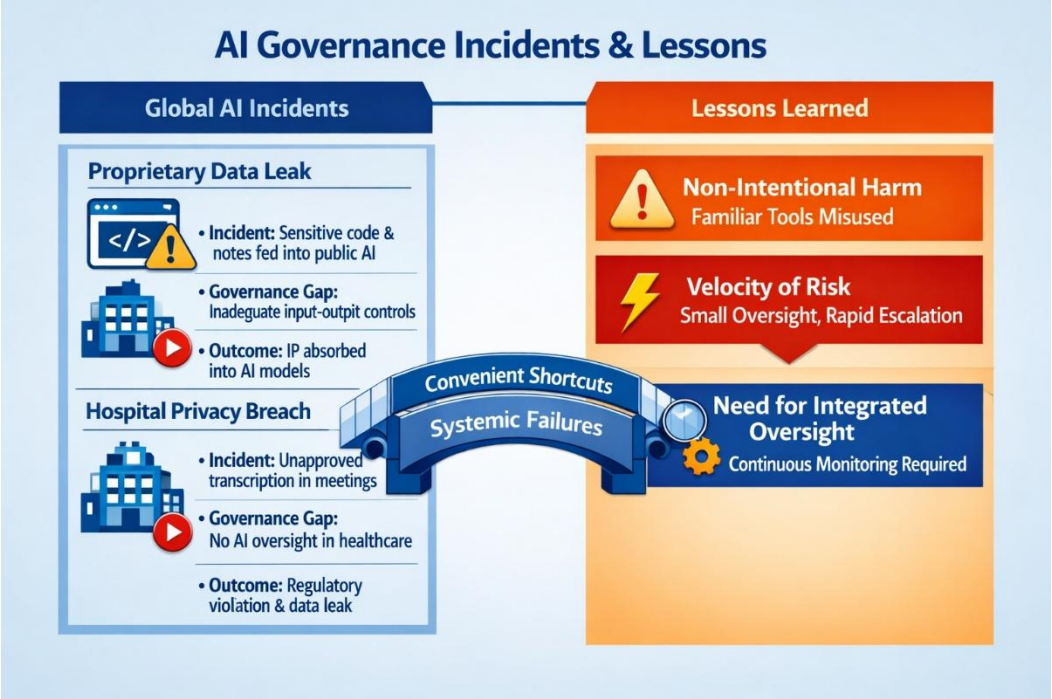

7. IMPLEMENTATION ROADMAP

"A 'start small, grow as needed' philosophy that ensures the organization matures without being overwhelmed."

The MV-AIG roadmap follows a "start small, grow as needed" philosophy. It integrates AI oversight into existing workflows to ensure adoption is safe without being bureaucratic. Implementation follows a pragmatic, three-phased approach to ensure the organization matures without being overwhelmed.

7.1 Phase 1: Establishment (Setting the Groundwork)

The goal is to secure formal authority and define the structure for AI oversight.

Key Activities:

1. Leadership Alignment: Present to the board/senior management to secure top-level buy-in.
2. Charter & Lead: Formalize a signed AI Governance Charter and appoint an AI Governance Lead (Single Threaded Owner) to drive decision-making.
3. Pro-Tem Committee: Form a temporary small group of senior leaders to oversee initial setup.
4. Policy Development: Draft the core AI Governance Policy and an Acceptable Use Guide (AI Addendum).
5. Initial Inventory: Create a basic spreadsheet of current AI tools with preliminary risk classifications.

Outputs: Signed Charter, appointed Lead and Committee, finalized Policy/Addendum, and a draft AI Inventory.

7.2 Phase 2: Operationalization (Core Mechanisms)

This phase transforms documentation into day-to-day action through lightweight, high-impact processes.

Key Mechanisms:

1. Inventory & Procurement: Activate a register to track tool owners and use cases, integrating it with the procurement process.

2. Impact Assessments: Implement risk assessments for any workloads classified as Medium Impact or higher.
3. Monitoring & Oversight: Identify "Shadow AI" and establish human-in-the-loop oversight where required.
4. Exceptions & Incidents: Create simple workflows for requesting policy exceptions and reporting AI-related incidents.
5. Policy & Education: Update existing InfoSec and Privacy policies with AI references and deliver initial staff training.

Outputs: Operational inventory and incident management, active exception process, 80% policy integration, and ready-to-use educational materials.

7.3 Phase 3: Maintenance & Continuous Improvement

The focus shifts to ensuring governance remains adaptive as AI risks and organizational needs evolve.

Key Activities:

1. Iterative Expansion: Update the AI inventory with new tools and refined risk levels.
2. Reporting Cycles: Establish bi-annual reporting on AI incidents, usage trends, and inventory changes.
3. Policy Review: Conduct annual reviews of the AI policy to align with existing governance cycles. Refer to 1.2 Strategic Context: The AI Governance Imperative
4. Feedback Loops: Use exception data and incident logs to improve governance rules.
5. Culture Building: Reinforce responsible AI use through ongoing communication and capability training.

Outputs: Mature monitoring processes, regular reporting cycles, updated policies, and an established culture of safe AI use.

8. CONCLUSION & OUTLOOK

"Building the institutional 'muscle memory' required to navigate a future where AI is an inseparable partner in professional success."

The Minimum Viable AI Governance (MV-AIG) model demonstrates that responsible AI adoption does not require heavy bureaucracy, expansive teams, or overly complex frameworks. By narrowing the scope to high-value, high-frequency use cases, specifically document analysis and meeting transcription, even the smallest organizations can meaningfully mitigate risk. This streamlined approach ensures that governance acts as a catalyst for innovation rather than a barrier, providing the clarity and confidence necessary to leverage AI-powered tools safely.

For larger organizations, MV-AIG serves as a high-velocity starting point. It provides a clean, low-friction foundation that cuts through the complexity often associated with enterprise-scale AI initiatives. Because the model is modular, it can be expanded and matured into more comprehensive frameworks as regulatory environments evolve from voluntary guidelines to mandatory compliance.

Ultimately, the success of MV-AIG rests on its cultural integration. By embedding AI oversight into existing policies and daily workflows, governance becomes a natural extension of operational excellence. The model's emphasis on proportionality and a "just culture" ensures that staff feel supported and empowered rather than policed. As organizations progress through the phased implementation of establishment, operationalization, and continuous improvement, they build a resilient foundation capable of evolving alongside the rapidly changing AI landscape.

Outlook - Navigating the Next Frontier of AI

The release of this Minimum Viable AI Governance model marks a foundational milestone, but it is not the final destination. As AI capabilities evolve from static document analysis to autonomous agents and generative multimodal systems, the risks and opportunities will continue to shift.

Looking ahead, we anticipate three key trends that will shape the next phase of this framework:

1. **Adaptive Regulatory Alignment:** With the anticipated formalization of Malaysia's AI legal frameworks, MV-AIG is designed to be the "on-ramp" for compliance. Future iterations will focus on mapping these lightweight mechanisms to mandatory national standards as they emerge.
2. **Automated Governance Tools:** We foresee a transition where "AI monitors AI." Future updates will explore how organizations can use basic automated tools as virtual watchdogs, to provide real-time oversight of AI usage, further reducing the manual burden on staff.

3. **Cross-Sector Applicability:** While this model currently focuses on high-value document and transcription tasks for NGOs and small organization, the outlook involves expanding these "Safe Path" guidelines to more significant sectors like finance, healthcare, and manufacturing, ensuring that the principles of proportionality and human-in-the-loop oversight remain universal.

By adopting MV-AIG today, organizations are not just solving for current risks; they are building the institutional "muscle memory" required to navigate a future where AI is an inseparable partner in professional success.

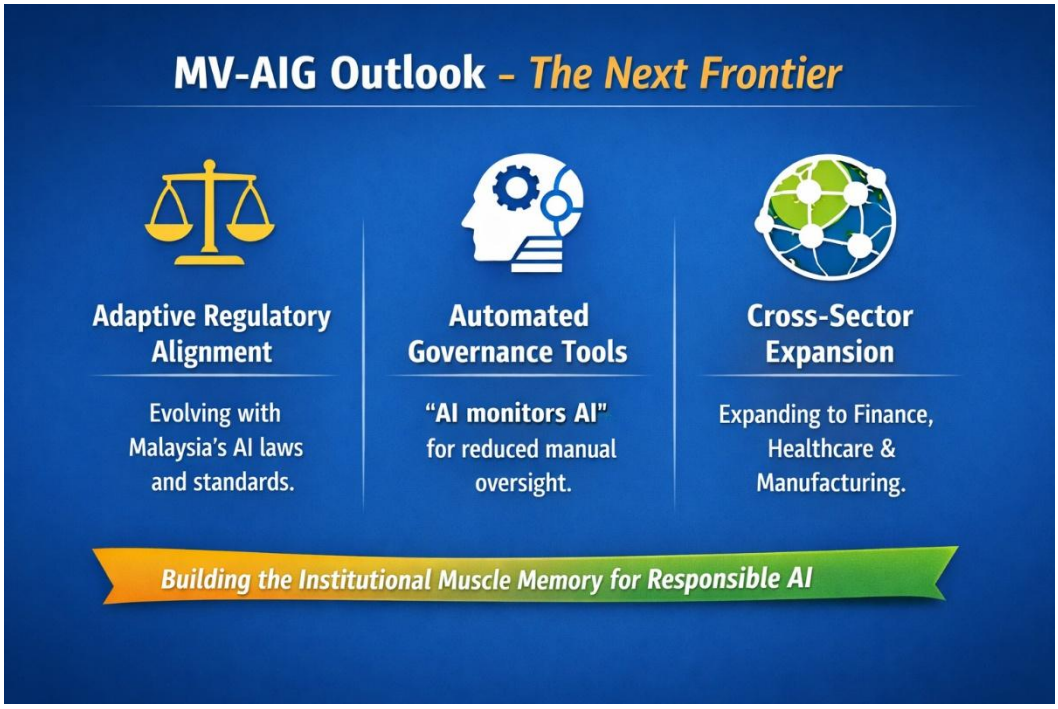

ACKNOWLEDGEMENT & REFERENCES

This report gratefully acknowledges AI Career Pro© for the methodologies, templates, and practical governance approaches that form the backbone of the Minimum Viable AI Governance model presented here. Their dynamic, real-world framework and streamlined policy —designed specifically to help organizations adopt responsible AI without unnecessary complexity—has been instrumental in shaping the structure, content, and implementation strategies described in this document. The clarity, practicality, and proportionality embedded throughout this work reflect the influence of AI Career Pro's guidance and resources, which have significantly strengthened the organization's early AI governance journey. For more information about AI Career Pro, go to https://governance.aicareer.pro/start or contact them at grow@aicareer.pro .

The authors would also like to thank the PIKOM leadership, council and staff for their unwavering support, interest and commitment to the implementation of MV-AIG within PIKOM.

This would not have been made real without everyone's support.

REFERENCES

i. https://mastic.mosti.gov.my/publication/the-national-guidelines-on-ai-governance-ethics/

ii. https://www.nacsa.gov.my/act854.php

iii. https://www.pdp.gov.my/ppdpv1/en/akta/pdp-act-2010-en/

iv. For a detailed analysis of the real-world incidents and case studies mentioned above (including the Samsung and Otter.ai breaches), please refer to Chapter 6. "CASE STUDIES: REAL-WORLD APPLICATIONS & INCIDENTS."

v. https://governance.aicareer.pro/

vi. https://aws.amazon.com/blogs/enterprise-strategy/tenets-supercharging-decision-making/

vii. https://www.pikom.org.my/2025/PIKOM_AI_ethic_and_governance_2025.pdf

viii. https://governance.aicareer.pro/blog/creating-your-ai-risk-management-policy

ix. https://governance.aicareer.pro/

x. Risk in risk management is typically considered as a combination of impact and likelihood of the impact. However, for AI risk classification, this is based on impact. For this document, the use of the word "risk" means "impact" in the standard risk management terminology.

xi. https://incidentdatabase.ai/cite/768/

xii. https://www.fieldlaw.com/insights/publication/how-an-ai-transcription-tool-triggered-a-privacy-breach

PIKOM

E1, Empire Damansara, No. 2, Jalan PJU 8/8 A,
Damansara Perdana, 47820 Petaling Jaya, Selangor
T : +(603) 7622 0079
E : info@pikom.org.my
W : www.pikom.org.my