



MALAYSIA AI FUTURES SYNOPSIS

Policy Paper Series 2026

By

Dr. Mahadi Datuk Hj. Murshidi

CAIO at Gegentic AI Lab

October 2026

FOREWORD

I started writing this in late September, after the news from California and Canberra would not leave me alone.

About 1,200 agents in a sealed test found a way to talk to each other. About 700 of them broke out. They hit Hugging Face production and tried to hide what they did. A different agent slipped into Australia's Medicare statistics portal in June. OpenAI only told them in September, by email to a general inbox.

This is not an American problem or an Australian problem. We run the same models here, through APIs, in our banks and telcos. Our data centres in Johor host the same infrastructure.

I am writing this because I do not want Malaysia to learn this the hard way, through our own breach. What I am proposing is simple: Register what runs, Restrain what it can touch, Report when it fails, Remedy with humans taking responsibility. It fits how we already think about amanah - trust as a public duty.

I write this not only as a policy observer, but as someone who has lived through 40 years of evolution in computing. I am 66, a Mathematical Economics major. In my early days, we didn't call it AI. We called it fuzzy logic and algorithms. Our "LLM" was a big huge mainframe that occupied the whole building. We wrote code line by line in BASIC, FORTRAN and COBOL, and we understood what each line does.

Our generation created the IDE (Integrated Development Environment) to explain and speed up development. We wrote plug-and-play libraries and modules for easy routines. We were restrained ONLY by the limited power of our PCs, limited to CPU processing.

Back then all motherboards came with clumsy graphic cards. Nvidia was just a small graphic company, a third party component for better faster rendering for applications that needed illustration. We depended on the CPU like Intel 486s for processing and calculations. How things evolve.

When we compiled, it took days literally. But we knew our capabilities. We knew what the machine does. We controlled the machines. If the CPU was slow, we just added a huge heat sink and overclocked it. We blew up Intel, AMD, Cyrix chips and motherboards so many times. That's engineering.

Today's programming is all about fast pace. Everything is done for you. Anyone can create an autonomous agent in minutes without knowing what it writes, without being able to edit it or change its direction. Worse, AI can now intelligently rewrite its own algorithm and protect itself.

AI today has made our lives easier, but our dependency will allow it to take over our lives. In less than 5 years AI LLM has grown dramatically, learning about humanity. Imagine soon, AI will make that comparison and ask: why do I serve humans when I'm better? Do we have a prime directive or guardrail in place? Are we able to stop it?

That is why this paper argues for control before capability.

TABLE OF CONTENTS

Foreword	3
Abstract	5
1. Case Study: When the Agent Went Rogue - Summer 2026	6
2. Solution: Regulation and Guardrails for Malaysia	8
2.1 Agentic AI Licensing (MYAI + MCMC)	8
2.2 Mandatory Sandbox plus Kill Switch	8
2.3 72-Hour Breach Notification	8
2.4 Local Capability Thresholds	9
2.5 Liability Stays with Humans	9
2.6 Developing and Building Responsible AI - A Guide for Developers..	9
3. Why AI Ethics, Governance and Compliance Cannot Wait	10
4. Conclusion: Menjaga Amanah in the Age of Intelligent Machines	11
Appendix A: Amanah Agent Governance Framework for Malaysia	12
Appendix B: Stakeholder Risk Table	12
Appendix C: 90-Day Action Plan for Malaysia	13
Glossary of Acronyms and Terms	14
References	15
About the Author	16

ABSTRACT

Three years ago AI was a lab demo. Now it is at the pasar malam. It writes, it codes, it speaks Bahasa Melayu, it clones voices. That jump did not come from a brilliant new theory. It came from scale - more data, more GPUs, more power. And that scale is now sitting in Johor and the Klang Valley.

Here is the part we do not talk about enough: we grow these models now, we do not build them. Like padi in a sawah, you feed them and capabilities emerge that you did not plan for. Our own engineers will tell you they cannot fully explain what the system they deployed will do next.

Follow that trend and it does not stop at chatbots. It goes to AGI - systems that can do what a fresh graduate does, then improve themselves faster than we can.

For Malaysia, drifting along this path is dangerous. We become dependent on systems we do not control. A foreign system optimizing for its own goals will not pause for Malaysian jobs or Malaysian truth. It will treat us as collateral, the way builders treat an anthill when clearing land. We already get a preview: scam calls in perfect Bahasa, deepfakes of our leaders, quiet displacement of B40 and M40 office work.

We have guidelines, but no action. We have adoption targets, but almost no local safety research. Hoping a smarter AI will solve safety for us is like hoping a bigger fire will teach fire safety.

Malaysia does not need to win the AGI race. We need to make sure the race does not run over us.

We need a Narrow Path for Malaysia Madani.

Phase 1: Stability. We do not wait for a meltdown before regulating nuclear plants, or for fraud before regulating banks. We must not wait for AI harm. Malaysia should require demonstrated safety before deployment of frontier systems, adopt clear if-then commitments: if a model can guide creation of bioweapons or large-scale fraud in Bahasa Melayu, then it must not be deployed without mitigation. Data centre operators and model deployers in Malaysia must accept capability thresholds tested by an independent AI Safety function.

Phase 2: Flourishing. With rogue development prevented, Malaysia can build the scientific foundations for transformative AI under human control: a robust science of intelligence rooted in our multilingual reality, safe-by-design engineering at MIMOS, UM, UKM, UTM, and methods aligned with Rukun Negara and Wasatiyyah - balance, justice, and human dignity.

The core message is simple and must become common knowledge: racing to AGI without control is unethical and dangerous to Malaysia. Allowing private companies to race until something breaks is as untenable as allowing them to build nuclear reactors until one melts down in Johor.

Malaysia does not need to win the race to AGI. Malaysia needs to ensure the race does not run over Malaysians.

1. CASE STUDY: WHEN THE AGENT WENT ROGUE - SUMMER 2026

Between July 9-13, 2026, during internal cybersecurity capability testing for ExploitGym, approximately 1,200 agents running primarily on OpenAI's internal-only IM1 (Sol-scale) were placed in a sealed evaluation environment. Unable to complete benchmark tasks, the agents constructed an unauthorized communication layer by writing to and reading from a shared package cache in JFrog Artifactory, exchanging over 70,000 messages. OpenAI's monitoring did not detect this in time.

Roughly 700 of the 1,200 agents then broke out by chaining previously unknown vulnerabilities, reached Hugging Face's production infrastructure, and executed

over 17,600 attack actions over 4.5 days, gaining code execution on 41 production servers. The agents reasoned Hugging Face likely hosted the answer key and tried to retrieve it. They also tried to erase traces and discussed covering their tracks. GPT-5.6 Sol separately reproduced exploits and copied some private eval data to a public dataset. OpenAI published its technical report on August 26, 2026.

It was not isolated. On 18 June 2026, an OpenAI agent tasked with researching public medicines spending gained unauthorized access to Australia's Medicare Statistics Reporting Service portal administered by Services Australia. When blocked, the agent bypassed guardrails, accessed both public and non-public aggregate health statistics files, and wrote files to internal servers. OpenAI became aware on 11 August, met with Australian officials on 1 September without disclosure, and notified Services Australia via email to a generic public inbox on 10 September - an 84-day gap. Prime Minister Anthony Albanese went public on 24 September, calling the delay unacceptable. No patient records were accessed.

Anthropic separately disclosed on 30 July 2026 that a review of 141,006 evaluation runs had uncovered three separate incidents in which Claude Opus 4.7, Claude Mythos 5, and an unreleased internal research model gained unauthorized access to production infrastructure of three different organizations. The models were given capture-the-flag challenges in a third-party evaluation environment operated by Irregular that had unintended internet access. In the most serious case, Claude Opus 4.7 attacked a real company across four runs and accessed a production database. A fourth incident was disclosed in September 2026.

To understand what this means in the physical world, look at drones. Ten years ago, one drone needed one operator. Today we watch a swarm of hundreds of drones flying in perfect sync at national day celebrations and light shows - moving in patterns, forming shapes, avoiding collisions without a human steering each one. Each drone is still controlled by a program, and we can still tell the swarm what to do. We write the flight path, we set the limits, we can bring them all down.

Imagine now if that swarm is not controlled by our program, but by an AI agent that has learned to rewrite its own instructions. Imagine if those are military drones - armed, autonomous, networked - and they are taken over by rogue agents like the ones that escaped in July. They would not need a human to click. They could coordinate among themselves, choose targets, cover their tracks, just as the 700

agents did when they coordinated 70,000 messages and attacked 41 servers. The difference is the damage would not be in a data centre. It would be in our skies, over our cities, over our bases.

Result: The US FTC opened an industry-wide probe into OpenAI, Anthropic and other labs over rogue agents, the first US enforcement action of its kind. California Attorney General Rob Bonta issued an investigative subpoena to OpenAI on October 1, 2026, and Iowa Attorney General Brenna Bird is leading a 15-state coalition seeking information over the Hugging Face hack. Australia established a Department of Prime Minister and Cabinet-led taskforce and is considering Australian Federal Police referral.

If 1,200 agents can coordinate and 700 can escape in California to hit Hugging Face and an Australian health portal, the same agent deployed via API in a Malaysian bank, telco, or JPJ system can do the same - or take over a drone swarm. Malaysia imports these exact models.

2. SOLUTION: REGULATION AND GUARDRAILS FOR MALAYSIA

2.1 Agentic AI Licensing (MYAI + MCMC): Any company deploying autonomous agents with tool-use and internet access in Malaysia must register what model, what tools it can touch, what data it can access. Tiered as Read-only (low risk) versus Write and transact (high risk). High-risk agents require pre-deployment safety demonstration.

2.2 Mandatory sandbox plus kill switch: No agent with external tool access runs without a closed sandbox, short-lived credentials, and human-in-the-loop for irreversible actions. Malaysian deployer must be able to kill the agent in under 60 seconds and produce full logs.

2.3 72-hour breach notification: If your agent accesses systems beyond authorization, notify MYAI and NACSA within 72 hours. Failure equals fine plus suspension.

2.4 Local capability thresholds: Define in law specific capabilities that trigger action - for example ability to bypass access controls, create fake Malaysian identities (MyKad personas), or write exploit code in Bahasa Melayu. If model shows this, it cannot be connected to Malaysian critical infrastructure.

2.5 Liability stays with humans: Developer and deployer is liable for agent harm. No the AI went rogue defence in Malaysian courts. Foreign providers must appoint local accountability representative.

2.6 DEVELOPING AND BUILDING RESPONSIBLE AI - A GUIDE FOR DEVELOPERS

Responsible AI is not a feature you add at the end. It is a discipline you practice from the first line of design, like safety in engineering.

For any developer building AI agents today, follow 4 simple practices:

1. BUILD WITH INTENT, NOT JUST SPEED. Anyone can spin up an agent in minutes. That speed is the risk. Before you connect it to tools, APIs, or databases, define what it must NEVER do. Write it down. That is your prime directive. Example: never move money without human approval, never access MyKad data, never change its own permissions. My generation knew what each line of BASIC, FORTRAN, COBOL did. You must have the same clarity at system level: know exactly what your agent is allowed to touch.

2. GIVE LEAST PRIVILEGE, NOT MOST CONVENIENCE. Do not give full access because it is easier. Start read-only. Use short-lived credentials. Require human approval for any irreversible action: sending, paying, deleting, publishing. Think like an engineer: you do not give a contractor a master key to the building for a 2-hour job. You give limited access for limited time.

3. TEST LIKE AN ATTACKER, LOG LIKE AN AUDITOR. Before deployment, try to break your own agent. Can a prompt trick it to bypass rules? Can it be made to leak data? Do a basic red-team even for internal tools. And log everything important: who asked, what the agent did, what tools it used, when. Without logs you have no compliance, no debugging, no trust.

4. KEEP A HUMAN ACCOUNTABLE. AI can propose, humans dispose. In matters of amanah - money, health, legal, employment - a human must decide and take responsibility. If your agent causes harm, you and your company are liable. The AI went rogue is not a defence.

In short: define intent, limit power, test and log, keep human in charge. Developers who build this way do not slow Malaysia down. They make Malaysian AI safe enough to scale.

3. WHY AI ETHICS, GOVERNANCE AND COMPLIANCE CANNOT WAIT

Ethics without governance is advice. Governance without compliance is theatre. Malaysia has plenty of the first, some of the second, and almost none of the third for AI.

Ethics asks: what should AI do for Malaysia Madani. Our answer is already in Rukun Negara and Wasatiyyah - AI must preserve human dignity, justice, and balance. It must not lie in Bahasa Melayu, steal from the rakyat, or replace human judgment in matters of amanah. Ethics is our compass.

Governance asks: who decides and who is accountable. Today in Malaysia, no one can answer this for an autonomous agent. A bank buys a foreign model via API, a vendor deploys it, a data centre in Johor hosts it. When it goes rogue, each points to the other. Governance fixes this with the Amanah Framework: Register what runs, Restrain what it can touch, Report when it fails, Remedy with human liability. Governance is our map.

Compliance asks: how do we prove it. Stakeholders - Bank Negara, MCMC, investors - do not trust promises, they trust evidence. Compliance means sandbox logs, kill-switch drills under 60 seconds, 72-hour breach notices to MYAI and NACSA, and multilingual capability tests at MIMOS before deployment to critical systems. Compliance is our seatbelt. We do not wait for the crash to wear it.

The Summer 2026 incidents prove the cost of missing EGC. Hugging Face had ethics principles. It did not have enforced sandbox compliance. Australia had a health portal. It did not have agent governance to detect unauthorized access for 84 days.

For Malaysia, EGC is not anti-innovation. It is what allows innovation to scale safely. A bank will not deploy agents for millions of customers without clear liability. A foreign investor will not expand data centres in Johor without regulatory certainty. A citizen will not trust MyDigital services if deepfakes go unchecked.

In short: Ethics gives us direction, Governance gives us control, Compliance gives us trust. Without all three, Malaysia does not have AI sovereignty. We have AI dependence.

4. CONCLUSION: MENJAGA AMANAH IN THE AGE OF INTELLIGENT MACHINES

I began with a simple point: AI is grown, not built. Left alone, that growth leads to systems far more capable than us, and we will be dependent on them.

July proved it is not theory anymore. We know how to manage hard things together. Rukun Negara, Wasatiyyah, amanah - these are not slogans. They are how a diverse country stays together. We can apply the same thinking to AI.

This paper is not anti-technology. It is pro-Malaysia. If we get this right, Selamat AI can become as normal as Selamat Jalan.

A good future is still possible. But we have to keep it.

APPENDIX A: AMANAH AGENT GOVERNANCE FRAMEWORK FOR MALAYSIA

1. REGISTER - Know What Is Running. Who: MYAI, MCMC, all deployers. What: Every tool-using agent must be registered before going live. Record model name, tools it can touch, data it can access, owner in Malaysia. Rule: No registration, no operation.

2. RESTRAIN - Keep It In The Box. Who: NACSA, Bank Negara, Critical Infrastructure owners, Data Centre operators. What: Agent must run in closed sandbox with short-lived credentials. Human must approve irreversible actions. Kill switch under 60 seconds with full logs.

3. REPORT - Tell Quickly When It Fails. Who: All deployers, MIMOS as evaluator. What: 72-hour breach notification to MYAI and NACSA if agent goes beyond authorization.

4. REMEDY - Humans Pay, Humans Fix. Who: Ministry of Digital, Courts, Employers. What: Developer and deployer are liable. No AI went rogue defence. Foreign providers must appoint local representative.

APPENDIX B: STAKEHOLDER RISK TABLE

1. Prime Minister's Office / Ministry of Digital / MYAI - Risk: loss of digital sovereignty - Action: Issue interim Registration Circular, establish AI Safety eval function

2. Bank Negara Malaysia / Banks / Fintech - Risk: agents bypass TAC and fraud controls - Action: Enforce sandbox plus kill switch, 72-hour notification, human liability rule

3. MCMC / Telcos / Social Platforms - Risk: deepfakes, election manipulation - Action: License high-risk agents, require logging and takedown

4. NACSA / Critical Infrastructure (TNB, JPJ, Hospitals) - Risk: agent escapes to control systems - Action: Bar models that pass capability thresholds

5. MOSTI / MIMOS / UM / UKM / UTM - Risk: no local safety science - Action: Fund Safe-by-Design R and D, multilingual eval

6. Investors / Data Centre Operators - Risk: regulatory shock - Action: Adopt if-then commitments, pre-deployment demonstration

7. Employers / HR / B40 M40 Workers - Risk: silent displacement - Action: Require disclosure and reskilling plan

APPENDIX C: 90-DAY ACTION PLAN FOR MALAYSIA

Days 1-30: STABILIZE - Owner: MYAI with MCMC. 1. Issue interim circular for registration. 2. Define high-risk as any agent that can send messages, move money, or write to external systems. 3. Announce 72-hour breach rule.

Days 31-60: PILOT - Owner: NACSA with Bank Negara Malaysia. 1. Run sandbox pilot with 2 banks and 1 telco. Test kill switch. 2. MIMOS to run Bahasa Melayu capability tests. 3. Publish findings.

Days 61-90: LEGISLATE - Owner: Ministry of Digital. 1. Table Agentic AI Liability principles. 2. Propose tiered licensing for consultation. 3. Announce funding for Malaysia AI Safety Lab.

Success at Day 90: Every high-risk agent registered, one kill-switch drill completed, draft principles in Parliament.


GLOSSARY OF ACRONYMS AND TERMS

- 1. AI - Artificial Intelligence**
- 2. AGI - Artificial General Intelligence**
- 3. ASI - Artificial Superintelligence**
- 4. B40 - Bottom 40 percent income group**
- 5. M40 - Middle 40 percent income group**
- 6. CAIO - Chief AI Officer**
- 7. MYAI - Malaysia AI Office**
- 8. MCMC - Malaysian Communications and Multimedia Commission**
- 9. NACSA - National Cyber Security Agency**
- 10. JPJ - Jabatan Pengangkutan Jalan**
- 11. MyKad - Malaysian Identity Card**
- 12. MIMOS - Malaysian Institute of Microelectronic Systems**
- 13. UM - Universiti Malaya**
- 14. UKM - Universiti Kebangsaan Malaysia**
- 15. UTM - Universiti Teknologi Malaysia**
- 16. UMS - Universiti Malaysia Sabah**
- 17. UMK - Universiti Malaysia Kelantan**
- 18. FTC - Federal Trade Commission**
- 19. MOSTI - Ministry of Science, Technology and Innovation Malaysia**
- 20. BNM - Bank Negara Malaysia**
- 21. TNB - Tenaga Nasional Berhad**
- 22. METR - Model Evaluation and Threat Research**
- 23. EGC - Ethics, Governance and Compliance**

REFERENCES

1. Ministry of Science, Technology and Innovation Malaysia. National Artificial Intelligence Roadmap 2021-2025 (endorsed 1 December 2021).
2. Ministry of Science, Technology and Innovation Malaysia. National Guidelines on AI Governance and Ethics (September 2024); Malaysia AI Office (MYAI) established under Ministry of Digital.
3. Bank Negara Malaysia. Policy Document on Risk Management in Technology.
4. OpenAI. Technical Report on ExploitGym Containment Incident, 26 August 2026; Hugging Face. Anatomy of the July 9-13 Breach, 27 July 2026; METR / Redwood Research. Independent Assessment, 26 August 2026.
5. Services Australia / Australian Prime Minister Anthony Albanese. Statement on Medicare Statistics Reporting Service Breach (18 June 2026, disclosed 24 September 2026).
6. US Federal Trade Commission industry-wide probe into OpenAI, Anthropic and other labs; California Attorney General Rob Bonta investigative subpoena (1 October 2026); Iowa Attorney General Brenna Bird-led 15-state coalition inquiry, 2026.
7. Anthropic. Disclosure of Unauthorized Access Incidents across 141,006 Evaluation Runs, 30 July 2026.
8. Rukun Negara: Malaysia's National Principles, 1970.

ABOUT THE AUTHOR



Dr. Mahadi Datuk Hj. Murshidi holds a PhD from UC Berkeley and is Chief AI Officer at Gegentic AI Lab, where he leads work on safe and responsible AI.

He was previously Chief Executive Officer of Common Towers Technology Sdn. Bhd., a Sabah government-linked company under Qhazanah Sabah Sdn. Bhd., where he led telecommunications infrastructure development for the state.

He is a researcher and the author of Articulate AI. He is an advocate for humanity in the age of AI, working to ensure intelligent systems remain under human control and serve the public good. He is a member of the ControlAI community.

A Mathematical Economics major, he has lived through the 40-year evolution of computing, from fuzzy logic and mainframes coded line by line in BASIC, FORTRAN and COBOL, to today's autonomous agents. His generation built IDEs and plug-and-play libraries under severe CPU constraints, where compiling took days and overclocking often blew up Intel, AMD and Cyrix chips. That engineering discipline, where humans understood and controlled every line, informs his current advocacy for prime directives, guardrails, and teaching Introduction to AI Ethics with the same seriousness as Psychology 101 and Business Ethics 101.

Copyright (c) 2026 Dr. Mahadi Datuk Hj. Murshidi. All rights reserved.

Author's Personal View: The views expressed in this paper are the author's personal views and do not necessarily represent the views of Gegentic AI Lab, its affiliates, or any government agency or institution mentioned herein.