



## RESPONSE PAPER

Policy Makers

# From Adoption to Sovereignty: A Capability Framework for AI-Ready Economies

A response paper proposing a national capability framework for AI-ready economies. Co-authored with Dominic Vincent Ligot of CirroLytx, bridging AI leadership theory with data science and AI policy.



**Ethan Seow & Dominic Ligot** — Centre for AI Leadership

25 March 2026

c4ail.org • Centre for AI Leadership

# From Adoption to Sovereignty: A Capability Framework for AI-Ready Economies

Ethan Seow Yi Zhe - Centre for AI Leadership (C4AIL) / Verixiom Dominic Vincent Ligot - CirroLytix Working Paper - March 2026

---

## Abstract

---

The Philippines exemplifies a paradox now visible across the global economy: high AI adoption rates coexisting with near-zero institutional readiness. Ligot (2026) documents this through IBPAP survey data and labor market scenarios - 86% of Filipino knowledge workers use generative AI, yet the country lacks the compute infrastructure, talent pipelines, and governance frameworks to move from consumer to builder. This paper provides the missing diagnostic and intervention framework. Drawing on C4AIL's maturity model (a 0-6 capability scale developed through engagement with organisations across multiple sectors), the ARGS discipline framework (Agency, Architecture, Governance, Scaling), and cybersecurity governance research on AI misclassification, we argue that the adoption-capability gap is not a resource problem but a maturity problem - measurable, diagnosable, and buildable. For services-dependent economies like the Philippines, the stakes are existential: staying at Level 0-2 does not merely limit productivity gains, it erodes the trust premium that justifies the entire offshoring value proposition. We propose a capability-building intervention that maps Ligot's four workforce roles (builders, users, planners, trainers) to maturity bands, and outline how the ARGS framework translates into actionable policy for industry bodies, regulators, and educational institutions.

**Keywords:** AI maturity, AI governance, labor markets, Philippines, IT-BPM, capability building, ARGS framework, cybersecurity governance

---

## 1. The Paradox: Fastest Adopters, Least Ready

---

In February 2026, Ligot presented data that should have alarmed every stakeholder in the Philippine IT-BPM sector. The numbers told two stories simultaneously - and neither was the one being

celebrated.

The first story was adoption. The Philippines had become one of the fastest AI-adopting economies in ASEAN. Knowledge workers were using generative AI tools at rates exceeding 86%. Enterprise adoption was accelerating. The dashboards looked green.

The second story was readiness. The IBPAP member survey revealed that 26% of organisations required significant reskilling to integrate AI into operations. Only 13% reported headcount increases attributable to AI, against 8% reporting reductions. Institutional capacity to build, govern, or verify AI systems remained near zero. The country had no sovereign compute infrastructure. Its data protection framework - the Data Privacy Act of 2012 - predated large language models by a decade.

Ligot named what he was seeing: the Philippines was a "fast adopter" but not a "sovereign builder." His four-role workforce model - builders, users, planners, and trainers - revealed the structural imbalance: the vast majority of the workforce occupied the "user" role, with minimal capacity in the other three.

This paradox is not uniquely Philippine. It is the default condition of 95% of organisations globally. US firms spent \$40 billion on AI in 2024; 95% saw zero measurable bottom-line impact (MIT, 2025). McKinsey's 2025 State of AI survey found that only approximately 1% of organisations describe themselves as "mature" in AI deployment. Eighty per cent of large enterprises claim AI governance initiatives, but fewer than half demonstrate measurable maturity (Gartner, 2025).

The Philippine case is distinctive not because its adoption-capability gap is wider than the global average - it is roughly the same. It is distinctive because the consequences of that gap are structurally different for a services-dependent economy.

---

## 2. Why the Gap Matters More for Services Economies

---

The IT-BPM sector contributed \$35.5 billion to the Philippine economy in 2023, employing 1.82 million workers directly and supporting millions more in adjacent industries. This is an economy built on exported knowledge work - and knowledge work is precisely the category that generative AI disrupts most directly.

For a manufacturing economy, AI adoption without capability is an efficiency miss. For a services economy, it is an existential threat operating through three compounding mechanisms.

## The Eloquence Trap at National Scale

The Eloquence Trap (C4AIL, 2026) describes what happens when AI output is fluent, confident, and professionally structured - regardless of whether it is accurate. The format signals competence even when the substance is wrong. A 2025 clinical study demonstrated this with devastating clarity: when physicians received AI-generated diagnostic advice that was eloquently worded and factually wrong, their accuracy dropped by 14 percentage points - and the most experienced physicians felt the hardest (Jia et al., 2025).

For the Philippine IT-BPM sector, the Eloquence Trap operates at industrial scale. When a Manila-based service delivery team uses AI to draft client reports, proposals, or analyses, the output is polished, professional, and fast. But if the content contains unverified claims, wrong regulatory references, or logical leaps that sound right but are not, the damage is not internal rework. It is the erosion of the trust premium that justifies the client's decision to offshore that work in the first place.

The competitive logic of BPO rests on a specific value proposition: skilled human workers delivering quality knowledge work at a cost advantage. If AI can generate the same polished output without the cost of human labour, the only remaining justification for the human workforce is quality - the ability to catch what AI misses, to apply domain judgment, to verify. If Filipino knowledge workers are using AI without verification, they are competing on speed against the machine itself. That is a race the machine wins.

## The Reliability Trap in Multi-Tier Workflows

BPO operations are inherently multi-step. A client engagement might flow through intake, processing, quality check, review, and delivery - each step potentially AI-assisted. If each step achieves 95% accuracy independently, the compound accuracy across five steps is not 95%. It is 77%. One in four end-to-end outputs contains an error.

The client does not see per-step accuracy. They see end-product quality. At the scale of Philippine BPO operations - millions of transactions per day - a 23% error rate is not a quality issue. It is a contract termination event.

Without structured verification at each step (what C4AIL calls a "verification chain"), multi-step AI-assisted workflows operate below the quality threshold that most service-level agreements specify. The speed gains from AI adoption are real. The quality degradation is invisible until it is catastrophic.

## The Confidence Plateau as Structural Dependency

Ligot's "consumer not builder" framing maps precisely to what C4AIL terms the Confidence Plateau. When a workforce becomes fluent in using AI tools without developing the capability to build, govern,

or verify them, three things happen simultaneously:

First, **skill erosion**. Anthropic's own 2025 study found that developers using AI coding assistants scored 17% lower on skill assessments than those who learned without AI. The largest gap was in debugging - the skill you need most when AI code breaks. A lead developer wrote publicly about removing all AI integrations from his editors after noticing he had stopped reading documentation, stopped thinking through problems, and felt "worse at his own craft than a year before." Extrapolate this to a workforce of 1.82 million.

Second, **verification atrophy**. The METR study found experienced developers were 19% slower with AI on familiar codebases - while believing they were 20% faster. The perception-reality gap means the workforce is not merely failing to verify; it has lost the impulse to check.

Third, **permanent dependency**. A workforce that uses AI tools but cannot build or govern them is dependent on the providers of those tools. If the provider changes pricing, terms, or capability, the economy has no fallback. This is not sovereignty. It is tenancy.

---

### 3. The Maturity Framework: Diagnosing Where You Stand

---

The adoption-capability gap becomes actionable when you can measure it. C4AIL's maturity model provides a 0-6 scale that maps to three capability bands, each with distinct economic characteristics.

#### Band 1: Explorers (Levels 0-2)

Usage is high. Verification is low. Leaders celebrate adoption metrics while the quality indicators tell a different story. AI is a faster typewriter - it accelerates production without improving judgment.

This is where approximately 90% of organisations sit globally, and where the vast majority of the Philippine IT-BPM sector operates today. The IBPAP survey data maps cleanly: high tool adoption, minimal process redesign, no systematic verification of AI output.

The economic characteristic of Band 1 is **linear returns**. Output increases proportionally to the number of users. There is no leverage effect. Each additional AI user adds roughly the same marginal value - and roughly the same marginal risk.

Addy Osmani, who leads developer experience at Google Chrome, named the experience of Band 1 the "70% Problem" in late 2024: AI gets you 70% of the way to a finished product rapidly, but the remaining 30% takes just as long as it ever did. Worse, the AI-generated 70% often introduces problems that would not have existed if a human had built it from scratch. The 70% Problem is not

limited to software. It describes any knowledge work where "looks done" and "is done" are two different things.

## Band 2: Architects (Levels 3-4)

The organisation has shifted from buying tools to building systems. Workers are trained not just to use AI but to question it. The critical shift: leaders begin to own their AI-informed decisions rather than pointing at the machine.

Level 3-4 workers - what C4AIL calls "Amplifiers" - share specific behaviours:

- They provide their own logic before asking AI to generate. They say "here is my analysis, implement this and flag where it might break," not "analyse this for me."
- They verify against their own mental model. They already know what the output should look like because they designed it.
- They catch what they have seen before. Five years of domain experience means knowing the edge cases that AI will miss.
- They improve the system. When they catch a recurring AI failure, they update the template so the AI handles it correctly next time.

The economic characteristic of Band 2 is **compound returns**. Each improvement to the system makes the next improvement faster. Templates become reusable. Verification patterns standardise. New Amplifiers emerge from the workforce as people outgrow the structured systems.

## Band 3: Orchestrators (Levels 5-6)

Output is decoupled from headcount. One domain expert - someone who deeply understands the business and has learned to design AI systems - manages the verified output of what previously required a team. Not because AI replaced the team, but because the expert's knowledge now scales through the architecture they built.

The economic characteristic of Band 3 is **power-law returns**. A single orchestrator can produce verified output at the scale of a department. This is where Ligot's "builders" sit - the people who design the systems that everyone else operates within.

## Mapping Ligot's Four Roles

Ligot's workforce model (builders, users, planners, trainers) maps to the maturity bands:

Ligot Role	Maturity Band	C4AIL Term	Economic Function
Users	L0-2 (Explorer)	Explorers	Production (linear returns)
Planners	L3-4 (Architect)	Amplifiers	System improvement (compound returns)
Trainers	L3-4 (Architect)	Translators	Capability development
Builders	L5-6 (Orchestrator)	Orchestrators	Architecture design (power-law returns)

The Philippine workforce is overwhelmingly concentrated in the "users" category. The national capability challenge is not to produce more users. It is to develop the planners, trainers, and builders who create the systems that make users productive.

---

## 4. The Four Disciplines: ARGS for a Services Economy

---

C4AIL's framework identifies four disciplines that separate the 5% of organisations capturing AI value from the 95% that are not. Each has specific implications for the Philippine IT-BPM sector.

### Agency: The Decision to Verify

Agency is the shift from accepting what AI provides to questioning it. It means recognising that no matter how polished the output, it is a first draft - never a final answer. It means providing the context, the intent, and the domain knowledge that AI does not have and cannot generate.

**Philippine application:** BPO workers trained to question AI output, not just use it. This is a cultural shift, not a technology one. It requires reskilling programmes that teach verification, not just prompting. The Harvard-BCG study (2023) demonstrated the stakes: on tasks that fell outside AI's capability boundary, passive users were 19 percentage points more likely to produce incorrect answers than people not using AI at all. The tool made the passive users worse. Active users - those who maintained what the researchers called an "active posture" - knew when to override it. **Policy lever:** IBPAP-endorsed training standards that include verification competency, not just tool proficiency.

## Architecture: The Floor That Catches Everyone

Architecture means building structured systems with built-in verification - what C4AIL calls the "Floor." The Floor is not "give everyone Copilot and hope for the best." It is engineering infrastructure that embeds institutional knowledge into the AI interaction so that output meets quality standards by default.

A Floor has four components:

- **Context injection** - the AI receives the specific situation, client requirements, and domain constraints before generating
- **Standards alignment** - quality criteria, formatting rules, and compliance requirements are constraints the AI operates within
- **Objective scoping** - each AI interaction has a defined scope and a definition of done
- **Verification chain** - automated checks catch obvious errors before a human reviews; human review focuses on judgment, not formatting

**Philippine application:** Industry-standard templates for common BPO workflows. When a service delivery team generates a client report, the Floor ensures the AI already knows the client's industry, the engagement's regulatory context, and the firm's quality standards. A secondary check flags inconsistencies before the output reaches a reviewer.

**Policy lever:** IBPAP developing shared Floor templates for high-volume, high-risk service categories. Open-sourced to member organisations. Each firm customises for their domain; the architecture is shared.

## Governance: The Weekly Learning Loop

Governance is not compliance. It is a living practice that gets better every week. When the system produces an error that reaches a client, the question is not "why did the AI get it wrong?" (it is probabilistic; it will sometimes get things wrong). The question is: "Why did our system not catch it?"

Werner Vogels, AWS's CTO, named the accumulating cost of skipping this: "verification debt." His framing is precise: "When you write code yourself, comprehension comes with the act of creation."

When the machine writes it, you have to rebuild that comprehension during review." Verification debt is the AI era's equivalent of technical debt - invisible, compounding, and eventually catastrophic.

**Philippine application:** Weekly governance reviews within delivery teams - not strategy discussions, but specific failure analysis. Which errors reached the client? Where did the verification chain break? What template update would prevent this next time? **Policy lever:** Industry reporting standards for AI-assisted service delivery quality. First-time-right rates. Error catch rates. Verification time trends. These metrics replace adoption rates as the measure of AI maturity.

## Scaling: Decoupling Output from Headcount

Scaling is where the investment pays off. The economics of AI have inverted the cost of production - generating a draft costs almost nothing. But verifying that output remains human. The bottleneck is not production. It is judgment.

Scaling means designing systems where one domain expert can manage the verified output of what previously required a team. The machine handles the language. The human owns the meaning.

**Philippine application:** This is the sovereign builder path. Instead of scaling by adding more L0-2 users (the current model), scale by developing L3-4 Amplifiers who each multiply the output of the users operating within their Floor. **Policy lever:** Protected development time for identified Amplifiers. Not "when things are quiet" - actual budgeted time for system improvement. The Harvard-BCG evidence is clear: organisations that combine workflow redesign with human development see 25-30% productivity gains. Those that deploy tools alone see 10-15%.

---

## 5. The Security Twin: Governance Without Cybersecurity Is Governance Without Walls

---

The adoption-capability gap has a twin that most AI governance discussions ignore: the cybersecurity governance gap.

Seow's misclassification thesis (2025, 2026) argues that AI security is systematically treated as an IT controls problem - extend existing access management, data loss prevention, and endpoint protection to cover AI tools - when it is actually a contextualisation and integration problem. The diagnostic question:

*"What data has been fed to which AI, who authorised it, and what is the blast radius if it is wrong?"*

If the answer is "we do not know," the organisation has an AI governance problem that no amount of IT security will solve.

For the Philippine IT-BPM sector, this manifests in three specific ways:

**Shadow AI at industrial scale.** The 2025 State of Shadow AI Report found that the average enterprise hosts 1,200 unauthorised applications, and 86% of organisations are blind to AI data flows. In a BPO operation handling client data across multiple jurisdictions, uncontrolled AI tool usage is not just a security risk - it is a contractual liability. Client data fed into an unauthorised AI tool may be absorbed into model weights, reconstructed from outputs, or leaked through side channels. The traditional CIA triad (Confidentiality, Integrity, Availability) was designed for data at rest and in transit. AI transforms, generates, and acts on data - a fundamentally different security model. **The paradigm mismatch.** When AI security is classified as IT, the CISO leads it with an IT background. Risk is calculated as probability times data exposure. The security team evaluates AI tools. Shadow AI is treated as a policy violation. But when AI security is classified as contextualisation, a cross-functional governance board leads. Risk is calculated as probability times decision impact times reversibility. Business owners evaluate AI tools with security as a constraint. Shadow AI is treated as a governance gap - if people need AI and the organisation has not provided it, the failure is the organisation's. **Regulatory gap.** The Philippine Data Privacy Act of 2012 was written before large language models existed. It addresses personal data processing but not AI-generated content, model training on proprietary data, or the liability chain when AI-assisted decisions cause harm. Singapore's Model AI Governance Framework (three generations: traditional AI 2019, generative AI 2024, agentic AI 2026) demonstrates what a contemporaneous regulatory approach looks like. The gap is not merely technical - it leaves Filipino organisations and their international clients without a governance floor for AI-assisted service delivery.

---

## 6. The Intervention: A National Capability-Building Programme

---

The maturity framework makes the intervention specific. The Philippines does not need more AI adoption. It needs a structured programme that moves workers from Band 1 (Explorers) to Band 2 (Architects) - from users to Amplifiers.

## Phase 1: Diagnostic (Month 1-2)

Establish baseline measurements across the IT-BPM sector:

- What percentage of output is AI-generated?
- What percentage is verified before delivery?
- What is the first-time-right rate for AI-assisted deliverables?
- How many organisations can answer the cybersecurity diagnostic question?

C4AIL's diagnostic tool ([assess.c4ail.org](https://assess.c4ail.org)) provides an individual and organisational assessment in three minutes. A sector-wide deployment through IBPAP would produce the first rigorous baseline of AI maturity for any national services sector globally.

## Phase 2: Build Floors (Month 3-6)

Develop structured AI systems for the five highest-volume BPO workflow categories. Each Floor embeds:

- Industry-specific context injection
- Client-appropriate quality standards
- Verification chains with automated pre-checks
- Domain-expert review gates

These Floors are developed by experienced practitioners (L3-4 minimum) and deployed to L0-2 users. The goal is not to train users to be better prompters. It is to build systems that produce reliable output regardless of the user's individual AI skill.

### Phase 3: Develop Amplifiers (Month 4-12)

Identify 5% of the workforce who demonstrate architectural curiosity - the people who, when shown a Floor, start asking how to improve it. These are the natural L3-4 candidates. Invest in them:

- 20% protected time for system improvement
- Cross-company Amplifier communities (via IBPAP)
- Mentorship from L5-6 practitioners (domestic and international)

### Phase 4: Compound Cycle (Month 12+)

Each Floor the Amplifiers build makes the next one faster. Templates become reusable across domains. Verification patterns standardise. New Amplifiers emerge from the workforce. The gap between Philippine IT-BPM and competitor markets widens every quarter - not because of AI adoption rates, but because of AI capability depth.

This is the sovereign builder path. Not sovereignty through building AI models (which requires compute infrastructure the Philippines currently lacks). Sovereignty through building the human systems that make AI productive - which requires only investment in people, processes, and governance.

---

## 7. Implications

---

### For Industry (IBPAP and Member Organisations)

Stop measuring adoption. The Philippines has already adopted. Start measuring capability: first-time-right rates, verification time trends, error catch rates, Amplifier density per thousand workers.

Develop industry Floors for high-volume workflows. Share the architecture; compete on domain expertise. The firm that builds the best verification chain for financial services BPO wins the contracts that require trust - which, in an AI-saturated market, will be all of them.

## For Policy

Reskilling programmes must target L3-4 capability, not L1-2 tool use. Teaching workers to prompt better is teaching them to be better passengers. Teaching them to verify, design systems, and improve templates is teaching them to be Amplifiers - which is what the economy needs.

National AI strategy should include cybersecurity governance explicitly. Data protection is necessary but insufficient. The question is not "is the data protected?" but "what decisions are being made on AI-generated output, who authorised those decisions, and what is the blast radius if the output is wrong?"

The regulatory gap between the Data Privacy Act of 2012 and contemporary AI governance frameworks (EU AI Act, Singapore Model AI Governance) should be addressed - not through piecemeal legislation but through a principles-based framework that treats AI governance as a business integration problem, not a technology regulation problem.

## For Education

Curriculum should teach AI verification as a core professional competency. The Eloquence Trap - the cognitive vulnerability that causes experienced professionals to trust polished AI output without checking - should be taught alongside critical thinking in every professional programme.

The four-role model (builders, users, planners, trainers) should inform curriculum design. The Philippines produces users at scale. It needs to produce planners and trainers at scale. Builders will emerge from the most capable planners - but only if the educational pipeline develops the prerequisite judgment.

---

## 8. Conclusion

---

The Philippines stands at a decision point that will define its economic trajectory for the next decade. The country has already adopted AI - faster than most. What it has not done is build the human systems that make AI productive rather than merely fast.

The maturity framework shows that this gap is not about resources or technology. It is about capability - and capability is buildable. The ARGS disciplines provide the architecture. The Floor/Ceiling model provides the implementation pattern. The compound expansion cycle provides the economic logic.

The path from AI consumer to sovereign builder does not require sovereign compute. It requires sovereign judgment - the institutional capability to verify, govern, and improve AI-assisted work. That capability lives in people, not in data centres.

The organisations and economies that build this capability in 2026 will own the value chain by 2028. Not because they adopted first - but because compound returns accelerate over time, and every quarter of delay makes the gap harder to close.

Nobody has all the answers yet - but someone has to go first.

---

## References

---

Anthropic. (2025).

*The impact of AI coding assistants on developer skill development.*

Internal research publication.

BCG & Harvard Business School. (2023).

*Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality.*

Working paper.

C4AIL. (2026).

*Sovereign Command: Leadership in the Age of Intellectual Automation.*

Centre for AI Leadership.

Gartner. (2025).

*Strategic predictions for 2026.*

Gartner Research.

Jia, Z. et al. (2025).

*The impact of AI-generated clinical advice on physician diagnostic accuracy.*

Clinical study.

Ligot, D. V. (2026a).

*From adoption to capability: Enterprise AI, labor structure, and talent constraints in the Philippines in a global context.*

SSRN 6163747.

Ligot, D. V. (2026b).

*Labor futures under artificial intelligence: Scenarios for the Philippine economy.*

SSRN 6171288.

McKinsey. (2025).

*The state of AI: How organisations are rewiring to capture value.*

McKinsey Global Survey.

METR. (2025).

*Measuring the impact of AI coding assistants on developer productivity.*

Autonomy Evaluation Resources.

MIT. (2025).

*Enterprise AI adoption and revenue impact.*

MIT Sloan Management Review.

Osmani, A. (2024).

*The 70% problem: Hard truths about AI-assisted coding.*

Blog post.

Seow, E. (2025).

*Practical cybersecurity decisions.*

Verixiom. v3.6.

Seow, E. (2026).

*Practical cybersecurity decisions for the AI age.*

Verixiom. First edition.

Vogels, W. (2025).

*Verification debt.*

AWS re:Invent keynote.