Past its use-by-date:

Turning end-of-life technology risk into national advantage

REPORT

Jason Van der Schyff

ASPI
AUSTRALIAN STRATEGIC POLICY INSTITUTE

25 YEARS
2001–2026

JULY 2026

About the author

Jason Van der Schyff is a Fellow with ASPI's Cyber, Technology and Security Program.

AI contributed no ideas to this report.

Acknowledgements

ASPI wishes to thank those external stakeholders—across governments and industry—who reviewed this report and provided comments that strengthened its analytical rigour.

Sincere thanks also to all ASPI staff who were involved with this project. The author especially thanks James Corera, whose guidance shaped both the framing and the substance of the argument, and to Justin Bassi for careful editorial work. Separate to ASPI, sincere thanks also to Angela Suriyasenee for fact-checking and research support.

ASPI would like to also acknowledge the sponsorship and support for this project from Cisco, including specifically Cisco's Head of Cybersecurity Policy for APAC, Sarah Sloan, without which this report would not have been possible. While Cisco has funded the report, the views and conclusions expressed herein are those of the author and do not necessarily reflect the official policy or position of Cisco.

About ASPI

The Australian Strategic Policy Institute was formed in 2001 as an independent, non-partisan think tank. Its core aim is to provide the Australian Government with fresh ideas on Australia's defence, security and strategic policy choices. ASPI is responsible for informing the public on a range of strategic issues, generating new thinking for government and harnessing strategic thinking internationally.

ASPI's sources of funding are identified in our Annual Report, online at www.aspi.org.au. ASPI remains independent in the content of the research and in all editorial judgements. It is incorporated as a company, and is governed by a Council with broad membership. ASPI's core values are collegiality, originality & innovation, quality & excellence and independence.

ASPI's publications—including this paper—are not intended in any way to express or reflect the views of the Australian Government. The opinions and recommendations in this paper are published by ASPI to promote public debate and understanding of strategic and defence issues. They reflect the personal views of the author(s) and should not be seen as representing the formal position of ASPI on any particular issue.

About Reports

Reports deliver original and deeply researched analysis and judgements on issues of strategic significance and ongoing importance. Reports aim to inform and influence national strategic choices with original, deeply researched analysis that surfaces implications, tests assumptions, and may consider future courses of action.

First published July 2026
Published in Australia by the Australian Strategic Policy Institute
ISSN: 3083-2853 (Online), 3083-2861 (Print)

ASPI
Level 2
40 Macquarie Street
Barton ACT 2600
Australia

Tel Canberra + 61 2 6270 5100
Email enquiries@aspi.org.au
www.aspi.org.au
www.aspistrategist.org.au

Facebook.com/ASPI.org

@ASPI_org

Cover image: Created by author using ChatGPT.

Contents

Foreword

Every generation of technology arrives more secure than the one before it. Richer telemetry. Stronger identity. Better encryption. Architectures secure-by-design rather than secured after the fact. This is not an accident — it is the product of deliberate investment and hard-won operational experience, and it is one of the genuine success stories of the digital age.

And yet the systems running much of the world were not built in this generation, but in the last one or the one before that. They still function and carry critical operational demands across government, infrastructure, and industry — but functionality is not the same as defensibility. The gap between the two is where this report begins.

At Cisco, we see networks as strategic infrastructure: the foundation on which organisations deliver essential services, protect data, enable innovation and participate in the digital economy. That is why the lifecycle of that infrastructure — how it is built, maintained, retired, and replaced — is no longer a back-office concern. It is a board-level conversation about effective risk management with national consequences.

This report poses a defining question: will governments, regulators, boards and operators build the governance architecture a more secure digital future requires — or will they wait for incidents, outages or adversaries to force the terms of transition? History suggests that without deliberate action, the second path wins by default.

Across the Indo-Pacific and beyond, systems often get patched until they can't be. Budgets defer the hard decision. Ownership of the problem quietly evaporates. And the technology that no longer receives security patches, cannot support modern identity standards or zero-trust architectures, or cannot prepare for the post-quantum transition, keeps running anyway — because turning it off was never anyone's job.

AI accelerates this risk. Advanced models compress the time between vulnerability and exploitation, lowering the bar for attackers and scaling the search for legacy weaknesses well beyond what human adversaries could manage alone. For end-of-life systems that cannot be patched, monitored or upgraded, risk no longer accumulates — it compounds at machine speed.

Cisco Talos and other industry researchers see this play out in practice: threat actors targeting unpatched, ageing and end-of-life infrastructure, particularly at the network edge, because that is where defenders often have the least visibility and room to manoeuvre.

But the same forces reshaping the threat are reshaping the defence. Modern technology gives defenders the visibility, identity and segmentation that zero-trust requires. It detects anomalies earlier and restores operations faster. And it opens the door to AI-enabled defence of its own —automation and machine-speed analysis that can identify exposure before an adversary gets there first. Modernisation, in this light, is a deliberate shift from those seeking to disrupt toward those protecting the systems on which societies depend.

Which brings the question back to cost — but not the cost most boards are used to weighing. The question is not whether an organisation can afford to modernise. It is whether it can afford the compounding cost of delay: downtime, incident response, operational fragility and constrained innovation. As AI sharpens both the threat environment and defensive capability, delay will increasingly be measured not only in financial terms, but in lost strategic optionality — the options a country or company no longer has because it waited too long to build them.

Seen this way, the risk and opportunity take a similar shape. Well-governed technology transition strengthens sovereign capability, improves resilience, unlocks data for better decision-making, supports AI-ready operations and creates new markets for assurance, migration and managed modernisation services. Countries that embed lifecycle governance early — through procurement, standards and clear accountability—will be better positioned to shape trusted digital infrastructure, reduce dependency and build regional confidence.

The most effective path forward will be built through partnership: shared standards, transparent accountability, trusted suppliers and practical transition pathways that reflect operational realities.

What ASPI has produced is rigorous, timely and highly relevant to decision-makers across government and industry. It names the governance failure constructively and for what it is — not a gap, but an absence: of assigned ownership, of funded exits and of enforceable thresholds. And it reaches a conclusion that is genuinely practical: lifecycle governance, done well, can turn inherited exposure into national advantage.

I remain optimistic because I have seen what technology can do when trust is built into it from the start. This report makes the case for choosing that future deliberately—and for acting before a crisis makes the choice for us.

Jeff Campbell,
SVP & Chief Government Strategy Officer, Cisco

1 Executive summary

We're quietly failing a test most don't know we're sitting. Across governments and industry, including critical infrastructure sectors such as energy, health and telecommunications, the everyday reality is that technology systems designed for the threat environments of the 1990s or 2000s continue to carry live operational demands into the 2030s. Some systems no longer receive security patches. Others remain operational but have no road map for post-quantum cryptography, modern identity standards, secure management protocols, secure boot, zero-trust integration or advanced telemetry. The result isn't managed risk. It's inherited exposure, accumulating without a named owner, a funded exit, or a threshold for action.

What's changed isn't the problem's existence but the speed and scale at which vulnerable systems can now be identified, targeted and exploited. Leading cybersecurity experts suggest we've crossed the Rubicon on cyber risk—that the threshold between human-paced attack and machine-speed, network-scaled attack has passed a point of no return with Anthropic's Mythos, OpenAI's GPT-5.5 and whatever comes next. Advances in offensive artificial intelligence (AI) capability are compressing the time between vulnerability discovery and operational exploitation—in some cases to a window measured in hours, not days. Cisco Talos finds that nearly 40% of the most actively targeted vulnerabilities affect end-of-life devices.[1] For systems that will never receive another patch, this isn't a narrowing window; it's no window at all. State-sponsored actors have already demonstrated persistent access to critical infrastructure through precisely this vector, and AI-enabled tooling is accelerating this asymmetry: what once required deep technical expertise is now more accessible, adaptable and deployable by a wider range of threat actors at significantly greater speed.

Legacy ICT debt operates across three interconnected layers: underlying infrastructure, the applications built on top of it, and the workforce responsible for securing and sustaining the whole stack. The most intractable problems often sit in the middle layer. Replacing ageing infrastructure risks disrupting dependent applications that are costly, operationally sensitive and/or difficult to migrate. The result is a compounding dynamic: systems that should be retired are repeatedly deferred; and the cost and operational risk of replacement grows faster than the capability or willingness to act. Unsupported systems persist not simply because organisations are negligent, but because the surrounding incentives reward continuity over transition. The data bears this out. By 2020, nearly half of all global enterprise network infrastructure was classified as ageing or obsolete.[2] In the UK, more than one in four of 228 legacy government systems was rated a critical risk in 2024.[3] The US Government directs four-fifths of its IT budget to sustaining existing systems rather than replacing them.[4] The more that's spent keeping unsupported systems operational, the less capacity remains to replace them.

This paper reframes that problem. End-of-life technology is not, at its core, a technical deficit. It's a governance failure: a policy, regulatory and assurance failure that allows unsupported systems to persist because decision rights are blurred, ownership is unassigned and the incentive to act falls on no one in particular. The determinative questions aren't whether better defensive tools exist. They're about who controls our legacy infrastructure, on what terms and under what governance structures.[5]

A framing familiar to defence planners treats early technology transition not as cost, but as strategic investment. Japan's approach to naval force structure, highlighted in analysis by ASPI's Richard Gray and documented by the Center for International Maritime Security, illustrates the principle: early decommissioning and replacement short of full service life are favoured to avoid the compounding costs and constrained options that accompany emergency replacement under operational pressure.[6] The same logic applies to digital infrastructure. The real choice isn't between spending and not spending. It's between spending reactively, after an incident, under pressure and at higher cost, or investing proactively on strategically advantageous terms. A forced technology refresh, structured deliberately, can strengthen capability, improve resilience and create new sources of commercial return. And it's most consequential at the moment of acquisition, not the moment of crisis.

Across the Indo-Pacific, the governance challenge takes difference forms depending on the stage of technological development and the capability, capacity and resources of public institutions.

In Southeast Asia generally, where digital infrastructure build-out is rapid, governance frameworks are still maturing and technology dependency creates strategic exposure, life-cycle governance is an immediate operational and sovereignty question. Jurisdictions that establish credible frameworks now are better positioned to shape their own procurement environments, reduce external dependency and build the assurance capacity that regional partners and investors increasingly expect.

In the Pacific and some parts of Southeast Asia, the constraint is different. Smaller state capacity, higher aid dependence and legacy exposure often reflect infrastructure scarcity rather than governance failure—and in many jurisdictions, that exposure is compounded by operational realities that sit beneath the headline numbers: chronic skills shortages, deeply interdependent legacy architectures that resist clean separation, and heavy reliance on external vendors whose commercial incentives rarely align with managed transition. For those governments, the most actionable contribution from more capable partners isn't market access but practical capacity—shared standards, technical assistance and procurement frameworks designed to prevent the next generation of infrastructure from inheriting the same life-cycle vulnerabilities.

For Australia, engagement contexts carries genuine strategic return. That's because countries with more mature domestic technology ecosystems and tighter public–private integration hold structural advantages in managing life-cycle transitions—and Australia is comparatively well positioned. A strong domestic technology sector, established enterprise partnerships with global providers and deepening public–private integration create a foundation that many regional partners can't yet draw on. Where that gap is widest, the most valuable contribution isn't technology transfer alone but the governance architecture that makes sustained transition possible: shared standards, assurance frameworks and procurement disciplines that build indigenous capacity rather than replicate dependency. Supporting regional partners along that path shapes the standards environment, strengthens assurance markets and builds the trust relationships on which long-term regional digital security depends. This has flow-on advantages in critical-infrastructure assurance, secure migration services, operational technology modernisation and post-quantum transition support. The common thread across every jurisdiction is the strategic prudence of early investment, before a crisis forces the terms.

Key findings

The governance gap is structural, not incidental

- End-of-life technology persists not because the risk is unknown, but because no governance mechanism assigns ownership, sets a threshold for action, or compels a decision. Responsibility is often distributed across operators, boards, regulators, procurement authorities and, in nationally significant cases, ministers, with no mechanism to force resolution where those roles are blurred.
- Policy and regulatory frameworks largely assume continued vendor support. When that assumption fails, structured alternatives are absent and operators shift, without explicit decision, from managed risk to improvised risk acceptance.
- Across this paper's primary case studies (South Korea, the Philippines, India and Australia), cybersecurity obligations are well established, but life-cycle governance isn't. Patching and auditing are routine requirements, but the absence of vendor support triggers no equivalent response. That gap is sharpest where operational constraints compound the governance deficit—for example, in the Philippines where skills shortages, deeply interdependent architectures and vendor dependency allow systems to drift into unsupported states long before any formal retirement decision is made.
- The capacity to operationalise life-cycle governance varies materially across jurisdictions, and ecosystem maturity is a significant driver. Countries with stronger domestic technology sectors and tighter public–private integration benefit from greater continuity in product-support commitments, clearer upgrade pathways and more predictable access to the technical expertise that migration, patching and system replacement require.
- Where vendor relationships are more transactional or fragmented, life-cycle management becomes harder to sustain consistently, even where policy frameworks exist. That structural asymmetry shapes not only who manages the transition well, but who's best positioned to support others in doing so.

The threat environment has materially shifted

- AI-accelerated exploitation has eliminated the response window for unpatched systems. Vulnerabilities are now discovered continuously, not episodically. Legacy systems are exposed not just to known weaknesses but to a machine-assisted discovery process that expands the attack surface over time.
- State-sponsored actors—including, but not limited to, the China-linked Volt Typhoon and Salt Typhoon groups—have demonstrated persistent access to critical infrastructure through legacy and end-of-life systems. The co-sealed advice from the Cybersecurity and Infrastructure Security Agency (CISA) is to patch where possible, but plan for the end of life. Patching buys time but it doesn't resolve the underlying life-cycle problem.[7]
- Unsupported systems accumulate not only known vulnerabilities but strategic incompatibility with the future security environment. Harvest-now, decrypt-later operations mean that adversaries are collecting encrypted data today against future quantum decryption capability, and end-of-life systems have no post-quantum migration pathway. The same logic extends to legacy platforms that typically can't adopt secure boot, modern identity integration, encrypted management protocols or contemporary monitoring standards. The capability gap compounds over time.
- Legacy systems are increasingly also creating a data sovereignty exposure. Organisations holding sensitive personal or strategic data will find it difficult to demonstrate reasonable security safeguards where critical systems can't adopt future cryptographic standards—a compliance gap that will widen as those standards are mandated.
- The convergence of IT and operational technology (OT) has widened attack surfaces beyond what legacy systems were designed to withstand. OT failure carries physical consequences—not merely data compromise. Legacy OT life-cycle decisions need to be governed by consequence,

not age or maintenance cost. Where wholesale replacement is unrealistic, a deliberate shift is required towards isolating, segmenting and surveilling what can't be patched or replaced, even where containment is second-best to remediation.

Inaction isn't cost-neutral

- Every dollar spent keeping unsupported systems operational is unavailable for new security innovations such as zero-trust architectures, AI-enabled defensive capability or post-quantum readiness. Deferred modernisation degrades the conditions under which later investment will have to occur.
- Splunk estimates that system downtime costs the world's largest companies US$400 billion annually, and 56% of incidents are attributable to cybersecurity causes. Fifty-four per cent of executives surveyed report deliberately leaving root causes unaddressed to avoid legacy remediation costs—behaviour that entrenches the conditions that make major failure more likely and eventual transition more expensive.[8]
- The distinction that matters is between forced recovery and governed renewal. The 2020 Hackney ransomware attack, costing over £12 million in recovery expenditure in its second year alone, didn't produce a stronger system by design. That's the structural difference between reactive and proactive transition.[9]

Modernisation as a strategic investment, not just a liability

- Well-governed transition lowers whole-of-life cost, improves resilience, unlocks decades of operational data trapped in closed architectures and creates new sources of commercial return. In the AI era, analytics and decision-support systems are only as effective as the data they can reach. Closed legacy architectures systematically block that access.
- Large-scale replacement programs create sovereign demand. Structured deliberately—through local capability requirements, anchor contracts and standards that reward trustworthy suppliers—legacy replacement can stimulate domestic industry and build enduring capability pipelines.
- For jurisdictions where life-cycle governance is still developing, the path to a fully integrated system is necessarily incremental. The prerequisite foundations—a complete and regularly maintained asset inventory, classification of systems by support status and sufficient institutional capacity to apply common standards across agencies—are themselves governance achievements that take time to consolidate. Recognising this, the strategic value of early investment lies not only in the transitions it enables but in the institutional infrastructure it builds: each step towards full life-cycle visibility reduces the conditions under which systems drift unmanaged into unsupported states.

Recommendations

1. Require life-cycle registers for high-consequence systems.

   Governments and sector regulators should require operators of critical and high-consequence systems to maintain continuously updated life-cycle registers. These should record support status, end-of-support dates and consequence tier, with end-of-life exposure in critical functions treated as a notifiable condition, not a quietly managed one. Australia's *Security of Critical Infrastructure Act 2018* (the SOCI Act) reforms and Protective Security Policy Framework provide examples of the legislative vehicles that can support this approach.

2. Set a consequence-based life-cycle standard.

   Governments and sector regulators should restrict or prohibit end-of-life ICT in the highest consequence operational environments, while permitting structured, time-limited exceptions in lower risk environments, subject to compensating controls and formal review.

3. Embed life-cycle obligations into procurement.

   Governments should incorporate vendor life-cycle disclosure, minimum support periods, software bills of materials (SBOMs) and transition planning requirements into government contracts and supplier assurance at the point of acquisition, before unsupported exposure accumulates as an unfunded liability.

4. Create structured transition pathways.

   Governments and sector regulators should publish phased transition road maps with clear compliance timelines, baseline technical expectations and sector-specific implementation guidance. Financial incentives, co-funding mechanisms and coordinated procurement programs may be necessary to reduce transition fragmentation and offset replacement costs for high-consequence operators. Workforce and training initiatives should support organisations managing large-scale technology transitions.

5. Require formal replace-or-mitigate decisions before end-of-support milestones.

   Industry and system operations should conduct formal transition reviews assessing replacement, isolation, compensating controls or retirement pathways at defined intervals before systems reach end-of-support status. Consistent with the OT constraint above, tolerating some end-of-life exposure may be unavoidable for now—but only as an affirmative, documented decision, with the cost of compensating controls explicitly weighed against the cost of replacement. Those determinations should identify the accountable executive, accepted residual risk, compensating controls, transition timeline and review date. Continuation by default shouldn't be permitted.

2 Introduction

End-of-life technology is no longer a backroom technical problem. It's a strategic governance problem with consequences for service continuity, economic resilience and national security. Across governments and critical infrastructure—such as energy, telecommunications, health care and public administration—systems are unsupported, unpatchable and, in some cases, no longer well understood. Those systems are embedding structural vulnerabilities into the critical infrastructure upon which our societies run due to a lack of governance mechanisms assigning ownership, forcing a decision or setting a threshold for action.

But the core failure isn't technical. Many policy, regulatory and assurance frameworks still assume a world in which systems can be patched, updated and secured within normal life cycles. That assumption no longer holds consistently. As vendor support ends, patching ceases and replacement becomes operationally difficult or financially prohibitive. Institutional risk shifts from governed risk management to improvised risk acceptance, usually without anyone explicitly deciding the risk is acceptable.

That ownership gap compounds the problem. Responsibility for end-of-life risk is typically distributed across the operator running the system, the board accountable for enterprise risk, the regulator setting minimum obligations, the procurement authority shaping replacement pathways and, in nationally significant cases, the ministerial system owner. Where those roles are blurred, legacy exposure persists not because the risk is unknown, but because decision rights, escalation points and funding responsibilities remain unclear. Operational constraints exacerbate this governance deficit—skills shortages, deeply interdependent legacy architectures and heavy reliance on external vendors create conditions in which systems drift into unsupported states well before any formal decision to retire or replace them has been made.

Adding to the implications of this is that the environment in which these systems now operate has shifted. Increasing connectivity, deeper convergence of IT–OT and the democratising effect of emerging technologies—which are lowering the barriers to entry for malicious cyber activity—have raised both the number of potential adversaries and the sophistication of the threats they can generate. Legacy exposure now carries greater strategic consequence. What was once a manageable constraint is now a material point of failure. But not all legacy risk is equal. Some exposures produce inconvenience from slower service delivery to higher maintenance overheads to localised outages. Others disrupt essential services at scale or carry safety consequences where OT and digital control systems intersect. At the highest end, legacy exposure can create strategic dependence by locking operators and governments into brittle systems and constrained replacement pathways. It increases the likelihood that major decisions are made reactively, after incidents, outages or acute operational pressure, rather than on planned and strategically advantageous terms.

Most policy responses tend to frame legacy technology as a cost, a compliance burden or a remediation problem. That framing is too narrow. It directs attention to liability and expenditure while obscuring the more important question of whether a forced technology refresh can also strengthen capability, resilience and market position. The real choice isn't between spending and not spending. Cost is just a matter of timing. The guarantee of rapidly developing offensive capabilities and their impact on security is simply not free of cost. Rather, the choice is between delayed spending to defensively preserve a deteriorating baseline, or proactive spending to reduce exposure early while creating durable value. In some instances, not spending also means that certain future pathways can't be taken; for example, the deployment of new leading-edge capabilities such as AI technologies, or defensive measures such as protecting against emerging quantum threats.

Commercial advantage here should be understood precisely. It may mean lower whole-of-life cost through planned transition rather than emergency replacement; improved resilience that reduces downtime, recovery costs and insurance exposure; stronger procurement leverage by avoiding last-minute vendor dependence; the development of exportable assurance; or new managed services that help other operators address the same problem. On this view, legacy modernisation can strengthen institutional performance and open adjacent sources of commercial return.

This paper treats legacy technology as both a governance gap and a potential catalyst for capability renewal. Its central argument is that policy and investment frameworks that treat end-of-life technology solely as a cyber hygiene problem will underperform, because they may reduce exposure but do so without improving decision utility, investment discipline or transition outcomes.

The paper uses South Korea, the Philippines and India as primary case studies—alongside a spotlight on Australia—due to the scale, pace and structural complexity of their digitisation programs, which make governance tensions more visible. The paper identifies two consistent patterns across the case studies. First, cybersecurity obligations are well established, but the governance of technology life-cycle risk isn't. Second, requirements for patching, auditing and risk management are common, but generally assume continued vendor support. When that assumption fails, structured alternatives are fragmented, indirect or absent.

These jurisdictions aren't identical to Australia, but they're analytically useful as they illuminate different versions of a commonly shared policy problem. They're all rapidly digitising states trying to manage technology dependence, transition pressure and uneven governance maturity. For the wider Indo-Pacific region, they offer practical comparative insight into

how governance gaps emerge when digitisation outpaces life-cycle planning. With similar conditions facing the entire region, those gaps are likely to repeat across jurisdictions, and lessons should be identified, shared and heeded.

The known consequences don't just involve annoying one-off interruptions but represent a national security issue. Unsupported and unpatched systems create persistent entry points for adversaries whose capabilities are rising, particularly when complemented with new AI capabilities. As digital systems become more central to national capability and economic resilience, unmanaged life-cycle risk scales with them. But so too does the opportunity. That's because the more central these systems are, the greater the value of replacing them well.

Importantly, this paper identifies the legacy problems and offers viable long-term solutions.

Defining the problem

For the purposes of this paper, end-of-life technology refers to hardware or software that's no longer supported by the vendor, or that can't practicably be updated or defended to a standard commensurate with its operational role. This includes systems that no longer receive security patches, systems that no longer receive substantive security capability development, and systems whose architecture or dependencies prevent practical modernisation. Once a product reaches that point, security patches—the primary means of remediating known vulnerabilities—are no longer assured. The system may continue to function. What changes is its defensibility: vulnerabilities discovered after support ends may remain exposed for the rest of the asset's operational life.

This paper uses end of life (EoL) as its primary analytical category, distinguishing where relevant between end of support (EoS), when regular updates cease, and end of security support, when even critical security patches are no longer provided. That distinction matters. Many systems remain operational after support arrangements degrade, making the key governance question not whether or not they still function, but whether they can still evolve and be defended against contemporary threats. This is particularly relevant for capabilities such as post-quantum cryptography, secure boot, modern identity integration, encrypted management protocols and advanced telemetry, which many EoS systems will never support.

This paper also uses the term 'legacy technology' more broadly than simple product age. A system may become 'legacy' because it's unsupported, unpatchable, operationally isolated, excessively difficult to maintain or unable to meet contemporary security and assurance expectations. In this sense, legacy risk is closely linked to technical debt, although the concept must be expanded beyond software engineering usage alone. The problem isn't merely ageing hardware or outdated code. It's the accumulation of maintenance and compliance debt: systems that continue operating while drifting further from the security baseline required by the contemporary threat environment. A modern server running an unpatched operating system, or a platform unable to adopt required security controls, can create the same governance and security liabilities as formally end-of-life infrastructure.

In this way, end of life shouldn't be confused with obsolescence in the colloquial sense. Power grid control systems, hospital records platforms and telecommunications infrastructure can remain functional for years after vendor support ends, and often do. The core risk here isn't immediate functional collapse. It's that those systems become progressively harder to assure against threats that are evolving more rapidly, particularly in the AI-era, while continuing to absorb capital, management attention and technical effort that could otherwise be directed towards transition. The OpenEoX framework, an OASIS technical report, proposes a four-stage life-cycle taxonomy: general availability, end of sale, end of security support, and end of life. While consistent vocabulary alone won't solve the problem, it can improve procurement design, transition planning and regulatory clarity. It also creates a shared basis for deciding when heightened oversight, compensating controls, replacement planning or risk acceptance should apply.[10]

3 Framing opportunity alongside risk

Legacy technology is ageing faster that it's being replaced, while geo-economic pressure, fiscal constraint, growing threats and operational dependencies are narrowing the room for governments, regulators and operators to manoeuvre. The policy debate has largely treated this as a risk to be managed. That framing is necessary but insufficient.

First, the prevailing policy framing captures exposure but not the full decision context in which operators act. Choices about legacy transition are shaped not only by cyber risk but also constrained capital budgets, board and shareholder pressure to preserve continuity, regulatory obligations, contractual commitments, insurance costs and the operational risk of replacing critical systems while they're still carrying live demand—a tension that plays out inside organisations as security teams press for investment while finance functions demand demonstrable return on investment under cost-control pressure.

Second, the same framing treats legacy systems primarily as liabilities and pays insufficient attention to the accumulated value they embody: embedded operational knowledge, long-run data holdings, installed demand and future procurement leverage. That value remains locked inside closed architectures and governance arrangements built to preserve continuity rather than enable transition. At the same time, organisations continue to incur a double cost: sustaining end-of-life systems that absorb significant expenditure, while the same resources are unavailable for innovation, modernisation and capability development. The result is not only higher operating cost, but reduced future competitiveness.

The more useful policy question for boards, regulators and system owners is therefore not only how much risk legacy systems create, but what follows if transition is treated as a structured investment program rather than a recurring cost centre. A well-governed transition can lower whole-of-life cost, improve resilience, strengthen procurement leverage and create new sources of value. An operator that develops a managed migration capability to modernise its own legacy estate may, for example, convert that experience into a repeatable service for other asset owners facing the same problem—turning a defensive expense into a source of operational resilience and commercial return.

Scale of debt

The scale of legacy exposure is substantial and still growing.[11] In this context, 'legacy' is used as a broader operational category that includes, but isn't limited to, EoL systems—capturing technologies that remain in service beyond their intended life cycle, including those approaching or operating without vendor support. Between 2017 and 2019, the proportion of old or obsolete assets in enterprise networks tripled.[12] By 2020, almost half of all global business network infrastructure was classified as ageing or obsolete.[13] These aren't marginal back-office systems—and this baseline of accumulated exposure predates the additional risk that quantum advances now pose to classical cryptography.

In the UK, the National Audit Office identified 228 legacy IT systems across government departments in 2024, one in four of which were rated 'red'—indicating a critical risk of failure or security compromise.[14] The UK Police National Computer, which underpins criminal records and border checks, has been in continuous service since 1974. In the US, the Government Accountability Office found that the 10 most critical federal legacy systems were between 8 years and 51 years old, and several are still dependent on COBOL—a programming language that predates the internet.[15] These examples show that legacy exposure is entrenched even in states with mature digital governance frameworks and significant IT budgets.

Increasing threat

The Indo-Pacific is the fastest-growing region for application modernisation services. That growth doesn't necessarily signal the problem is being solved—instead it reflects delayed recognition of accumulated legacy exposure and the beginning of a long investment cycle.[16] For many Indo-Pacific countries, where digitisation accelerated rapidly in the 2000s and 2010s under compressed delivery timelines, legacy risk is likely to be more acute and less visible.

That exposure is reflected in how adversaries operate. According to Cisco Talos, threat actors continue to exploit unpatched, years-old vulnerabilities—particularly in networking and edge devices—to establish persistent access, highlighting how legacy and end-of-life systems remain a primary attack vector rather than a residual risk.[17]

The region also recorded more than 57,000 ransomware attacks in the first half of 2024, with networking and edge devices among the most frequently exploited targets.[18] The urgency of acting on the legacy technology is sharpened by a threat environment that's more permissive, more crowded and more capable. In the European Union, major incidents in financial services almost trebled between 2020 and 2023.[19] In the UK, the National Cyber Security Centre reported a 50% increase in nationally significant incidents in 2024, alongside a threefold rise in severity (compared to the previous year).[20] Google's threat analysis has found that the median time to exploit vulnerabilities fell from 63 days in 2018–19 to as few as five days by 2023–24.[21] *For systems that will never receive another patch, this doesn't mean a narrowing window—it means no response window at all.*

This is because the barriers to disruptive cyber activity are falling. With the rapid development and adoption of AI-enabled tools, techniques and effects that once required deep technical expertise are now more accessible, adaptable and deployable (see box). That democratising effect expands the pool of actors able to generate

meaningful harm, regardless of their capacity and capabilities. That said, state-backed actors remain central to this threat picture. In 2024, CISA, along with the Federal Bureau of Investigation and the National Security Agency, disclosed that the China-linked group Volt Typhoon had compromised multiple critical-infrastructure sectors by exploiting legacy systems and maintaining persistent access to networks that defenders struggled to monitor or remediate.[22] The sequencing of CISA's advice was telling—patch where possible, but plan for EoL—reinforcing that patching may buy time, but it doesn't solve the underlying life-cycle problem.

Industry views on the threat environment

Industry reporting reinforces that legacy exposure isn't merely a maintenance problem but instead a standing security liability that sophisticated threat actors continue to convert into operational access across government, critical infrastructure and telecommunications networks.

- Cisco Talos analysis shows that legacy systems remain disproportionately exposed, and that nearly 40% of the most actively targeted vulnerabilities affect EoL devices. It also finds that a significant share of exploited vulnerabilities is more than a decade old, highlighting both persistent patch delays and the enduring relevance of flaws embedded deep within core enterprise and government systems.[23]
- Palo Alto Networks' research indicates that a state-aligned group operating out of Asia conducted a global espionage campaign that has compromised government and critical-infrastructure organisations in nearly one in five countries—including several here in the Indo-Pacific region including Afghanistan, Saudi Arabia and Uzbekistan are typically not considered part of the Asia-Pacific region.[24]
- Google Threat Intelligence Group (GTIG) assesses that this pattern isn't isolated to one campaign or one sector. In its GRIDTIDE disruption write-up, Google said industry partners moved against a global espionage campaign targeting telecommunications and government organisations across four continents. Google assesses UNC2814 to be a suspected China-linked threat actor tracked since 2017, with a long record of targeting governments and major telecommunications providers across Africa, Asia and the Americas. As of 18 February 2026, GTIG reported 53 victims in 42 countries and suspected infections in at least 20 more.[25]

The cost of inaction and the case for structured investment

Legacy technology imposes well-documented direct costs. Less well captured is the opportunity cost of leaving those systems in place. Both matter, because policy choices that appear fiscally prudent in the short term prove strategically and financially inefficient over time. In 2023, the US Government spent approximately US$100 billion on IT and cybersecurity, of which around US$80 billion—four-fifths—went to operating and maintaining existing systems rather than developing new ones.[26] The UK directed nearly half of planned IT spending in 2019 towards sustaining legacy infrastructure rather than replacing it.[27]

This 'build versus run' dynamic is self-reinforcing: the more an organisation spends keeping ageing systems alive, the less capacity it retains to replace them, and the more fragile the estate becomes. Every dollar spent keeping unsupported systems operational is a dollar unavailable for new security innovations such as zero-trust architectures, AI-enabled defensive capability, post-quantum readiness or the next generation of digital public services. Deferred modernisation doesn't merely delay capability uplift—it degrades the conditions under which later investment will have to occur.[28] The operational consequences are already visible. In Southeast Asia, the average cost of a data breach reached US$3.3 million in 2024, and financial services breaches averaged US$5.7 million. More than a third of Indo-Pacific organisations reported breaches costing between US$1 million and US$20 million in 2023.[29] Splunk estimates that system downtime costs the world's largest companies US$400 billion annually, at an average rate of US$9,000 per minute, with 56% of incidents attributable to cybersecurity causes.[30] More telling: 54% of executives surveyed reported deliberately leaving root causes unaddressed to avoid the cost and disruption of remediating legacy systems.[31] Such behaviour is understandable in the short term, but it entrenches the conditions that make major failure more likely and eventual transition more costly and forgoes the possibility of value-generating transitions.

In October 2020, a ransomware attack on the London borough of Hackney disrupted council services for almost a year at a reported cost of over £12 million in recovery expenditure in its second year alone.[32] In July 2023, a LockBit ransomware attack on the Port of Nagoya—Japan's largest port, handling roughly 10% of the country's trade—disrupted container terminal operations for two days after compromising legacy OT management systems.[33] Applied to nationally significant systems—power grids, financial clearing infrastructure, telecommunications backbones—the same dynamics would produce consequences of a different order. Hackney didn't emerge from crisis with a stronger system by design. Instead, it incurred major costs to restore compromised functions under pressure. That's the difference between forced recovery and governed renewal.

Three forms of opportunity deserve explicit recognition here.

- *First, procurement as a market-shaping instrument.* Large-scale replacement programs create substantial sovereign demand. When structured deliberately—through local capability requirements, anchor contracts, phased transition pathways or standards that reward trustworthy suppliers—legacy replacement can stimulate industry and build domestic capability pipelines.
- *Second, transition as a mechanism to unlock data value.* Many legacy systems function as value traps: they contain decades of operational data that can't be integrated, analysed or reused effectively because the surrounding architecture is

closed or brittle. Modernisation converts that data into public, operational and commercial assets. This becomes more acute in the AI era: analytics, automation and decision-support systems are only as effective as the data they can reach, and closed or brittle legacy architectures systematically block that access. Unlocking it is a precondition for AI-ready operations, not an incidental benefit.[34] In health, energy, and transport, the forgone value of continued lock-in is likely to be substantial.

- *Third, life-cycle governance as strategic and commercial advantage.* Jurisdictions that develop credible EoL standards, transition pathways, assurance models and certification regimes create governance architecture that others may adopt or adapt. In the Indo-Pacific, where many states face similar life-cycle pressure but uneven institutional capacity, first movers can carry policy value, diplomatic weight and, potentially, exportable service opportunities in assurance, transition support, post-quantum transition capability and managed modernisation. This is particularly consequential given the post-quantum cryptographic transition now underway: jurisdictions with mature life-cycle governance are structurally better placed to identify, sequence and replace non-upgradeable systems before they become the weak point in national cryptographic resilience. The same governance infrastructure that supports EoL transition is, in effect, the infrastructure that supports post-quantum cryptography readiness.[35]

This section shows that the weight of legacy debt is real, but so is the opportunity embedded in addressing it well. The difference between those two outcomes lies less in the scale of spending than in the design of the framework governing it. Frameworks built only to close vulnerabilities produce a narrower return. Frameworks built to improve resilience, sharpen accountability and build capability do more than reduce risk. The box below illustrates that distinction through a sample board-level decision paper for a critical-infrastructure operator.

Investment case in practice: a sample board paper translating legacy risk and opportunity into an executive decision framework

Classification: Commercial-in-Confidence

Paper type: For Decision

Sponsor: CEO | **Co-sponsors:** CIO, CFO, CISO

DECISION SOUGHT

The board is asked to approve:

1. a budget envelope of [X]–[X] million over three to five years, drawn from existing capital allocation and new investment
2. establishment of a Legacy Modernisation Governance Committee, chaired by the CEO and reporting to the board six-monthly
3. commencement of Phase 1 (asset visibility and triage) within 60 days of approval, at a cost of $[X] million.

EXECUTIVE SUMMARY

A significant and growing proportion of the organisation's technology estate is operating beyond vendor support life cycles. These systems can't be fully secured, absorb a disproportionate share of operating expenditure and constrain the organisation's ability to develop new capability.

This investment will be made—either proactively, through a structured program that generates capability and return alongside risk reduction, or reactively, following an incident that forces unplanned expenditure at higher cost and lower strategic return. Modelling in this paper indicates that a structured program will reduce the total cost of ownership by approximately $[X] million over five years, generate $[X] million in value through data and operational capability currently locked in legacy architecture, and reduce the probability of a material incident from [X]% to [X]% over the program period.

STRATEGIC CONTEXT

An asset life-cycle audit conducted in [quarter/year] identified [X] systems operating beyond vendor end-of-life, [X] beyond end-of-security-support, and a further [X] within 24 months of EoS. Together, these account for [X]% of critical operational systems and [X]% of current IT operating expenditure. Legacy systems also contain [X] years of operational data that can't currently be integrated, analysed, or used for commercial or regulatory purposes—value that remains locked until transition occurs.

THE INVESTMENT CASE

Current cost position

The organisation currently spends approximately $[X] million annually maintaining systems beyond vendor support or within 24 months of EoS. This includes direct maintenance, compensating security controls, manual workarounds, and incident response costs attributable to legacy exposure. Of this, approximately $[X] million maintains systems providing no incremental capability—immediate targets for cost reduction. Probability-weighted incident cost represents a contingent liability of $[X] million not currently fully reflected in the risk register.

Structured investment return

A proactive modernisation program would generate return across three vectors:

- *Direct cost reduction.* Transitioning [X] legacy systems over three years eliminates an estimated $[X] million in annual maintenance expenditure by year four. Capital payback period on this basis alone is [X] years.
- *Data and operational value unlocked.* [X] legacy systems contain operational datasets currently inaccessible for analytics, regulatory reporting or commercial use. Transition makes

this data available, supporting an estimated $[X] million in operational efficiency gains and $[X] million in new revenue capability over five years, through for example improved asset utilisation modelling / customer service quality / regulatory reporting efficiency.

- *Strategic and governance positioning.* Early development of a mature life-cycle management framework positions the organisation ahead of mandatory requirements, reduces compliance cost, and creates potential for leadership in industry standard-setting—with value in both cost avoidance and regulatory relationship capital.

RISK ANALYSIS

Execution risk—operational disruption during system transition. Mitigation: phased implementation, parallel running where feasible, compensating controls for systems awaiting transition and defined rollback procedures. Residual risk: moderate.

Funding risk—sustained capital commitment across multiple budget cycles. Mitigation: milestone-gated reapproval after Phase 1. Residual risk: low to moderate.

Risk of not proceeding—structural and increasing. Probability-weighted incident cost of $[X] million over five years is the floor of this exposure, excluding regulatory consequence and reputational impact.

GOVERNANCE AND REGULATORY CONSIDERATIONS

The program establishes formal life-cycle governance as an organisational capability, not a project. This includes a maintained asset life-cycle register, periodic board reporting on exposure and remediation progress, and a defined 'replace or mitigate' decision framework for systems approaching EoS.

RECOMMENDATION

The board is asked to approve the three decisions set out above. The case for proceeding is strong on risk grounds alone; on value and return grounds, it's compelling. Deferral isn't cost-neutral—it's a choice to incur higher costs, accept higher risk and forgo the strategic value of modernisation undertaken on the organisation's own terms.

4 Emerging technology acceleration

Three technological shifts are changing the strategic context in which EoL technology risk now sits. Each accelerates exposure, shortens the window for reactive response and strengthens the case for structured investment. Together, they demonstrate that legacy technology can no longer be treated as a static maintenance problem—and that, instead, the shifts are now tied to capability competition, industrial positioning and the future resilience of critical systems.

AI-accelerated vulnerability discovery and modernisation

AI is changing the speed, scale and economics of vulnerability discovery. When vendors release patches for supported products, they also reveal information about the underlying vulnerabilities. Because those vulnerabilities often exist across shared codebases, the effect can extend to unsupported versions that will never receive a fix. For legacy systems, each new advisory can become a new avenue for exploitation with no corresponding path to remediation.

In October 2024, Google Project Zero and DeepMind's 'Big Sleep' AI agent identified the first publicly documented case of an AI-discovered zero-day vulnerability in production software—a stack buffer underflow in the SQLite database engine—through automated variant analysis.[36] By September 2025, researchers had also demonstrated that AI systems could generate working proof-of-concept exploits within minutes of a security advisory being published.[37] The *Foresight 2030* report of the European Union Agency for Cybersecurity (ENISA) ranks the exploitation of legacy systems among the most significant cybersecurity threats of the decade, ahead of both the abuse of AI and smart device attacks.[38]

This trajectory is unlikely to materially slow in the near term. Emerging frontier AI models are demonstrating the ability not only to analyse code, but to reason across systems, adapt techniques and iterate towards exploit development with minimal human intervention. As those capabilities mature, vulnerability discovery is likely to become more automated, scalable and persistent. For legacy and EoL systems, this has a specific implication: vulnerabilities are no longer discovered episodically, but continuously. Systems that can't be patched are therefore exposed not just to known weaknesses, but to a growing and machine-assisted discovery process that systematically expands the attack surface over time.

The strategic problem isn't simply faster exploitation. It's the widening asymmetry between how quickly malicious actors can identify opportunity and how slowly operators can reduce inherited exposure.

But the same AI capabilities can also improve transition decisions. AI-enabled scanning, asset discovery and triage can map legacy exposure across large and complex estates, identify highest-consequence systems, and support disciplined remediation sequencing—converting periodic, labour-intensive audit into continuous assurance. Industry is already beginning to organise around this inversion. In 2026, Anthropic PBC launched Project Glasswing, bringing together a consortium of founding partners—Amazon Web Services, Apple, Broadcom, Cisco, CrowdStrike,

Google, JPMorgan Chase, the Linux Foundation, Microsoft, NVIDIA and Palo Alto Networks—to apply frontier AI models to vulnerability discovery and remediation at scale, explicitly framed as defensive deployment ahead of equivalent capability being weaponised.[39] Just one month following its release, Anthropic reported 1,451 vulnerability findings, acknowledged by the source code maintainers and resulting in 97 upstream patches and 88 published security advisories.[40] What requires governance is not only the deployment of those tools but the terms of their use: validation standards, procurement rules, human oversight and the conditions under which model outputs can inform assurance and replacement decisions.

Governments, regulators and operators that deploy these tools early—and set credible standards for their use in critical-infrastructure assurance—won't just manage risk more effectively. They also have the opportunity to build sovereign capability and domestic expertise in an emerging market for assurance, transition support and managed modernisation services.

Harvest-now, decrypt-later: the post-quantum transition

Quantum risk introduces a different problem: not fast exploitation, but longer liability. State actors are already assessed to be intercepting and storing encrypted data in anticipation of future decryption by cryptographically relevant quantum computers. Where legacy systems rely on outdated or non-upgradeable encryption, data captured today may become readable later—even if the operational consequences aren't felt until years after the original compromise.

Harvest-now, decrypt-later is only one dimension of the challenge posed by quantum computing. Equally significant is the emerging 'trusted now, forged later' threat, in which cryptographic credentials, signatures and authentication mechanisms trusted today may be vulnerable to forgery once sufficiently capable quantum systems emerge. In complex, interconnected environments, cryptographic trust underpins far more than data confidentiality. It supports authentication, system integrity and the reliable operation of critical infrastructure. If that trust is degraded, the effects aren't isolated—they can propagate across identity systems, network control and signalling layers (including position, navigation and timing services such as global navigation satellite systems[41]), secure orchestration platforms and service assurance functions. Once realised, quantum capabilities could therefore undermine both the confidentiality and integrity of data-in-transit and authentication mechanisms at scale, with the potential for cascading operational disruption and loss of trust across systems.[42]

The common limitation across the three case-study countries isn't recognition of the post-quantum challenge but the governance of systems that can't transition. South Korea has set a 2035 migration target; India, in line with a number of hyperscalers such as Google, has adopted a 2029 deadline for critical infrastructure;[43] while the Philippines acknowledges the quantum threat without publishing a comparable migration timetable.[44] In all three cases, the treatment of non-upgradeable systems remains unclear, uneven or underdeveloped. India recommends isolation; South Korea's road map appears largely silent; the Philippines doesn't set an equivalent approach.[45] The consequence is the same across each case: transition planning exists to varying degrees, but residual risk from systems that can't be upgraded remains insufficiently owned.

This creates a risk profile that conventional cost models handle poorly. The liability is incurred now, through ongoing exposure and data capture, but consequences might surface much later and might not be readily attributed to the original governance failure. That temporal disconnect makes underinvestment more likely. It also complicates compliance with emerging data protection and critical-infrastructure obligations, including requirements to maintain 'reasonable security safeguards' for sensitive personal and operational data—for example the Digital Personal Data Protection Act in India.

Yet, the strategic opportunity is larger than current policy discussion allows. The post-quantum transition is one of the largest infrastructure transitions affecting any sector that depends on encrypted communications, secure identities or protected data storage. Those that build capability around that transition—in cryptographic software, secure hardware, assurance services, standards implementation or managed migration—should expect lasting industrial benefits, not merely short-term compliance advantages. For Indo-Pacific states, early movement may, again, also shape regional expectations about what credible transition governance looks like.

OT convergence: expanded attack surface and operational value

The growing convergence of IT and OT is widening the gap between legacy design assumptions and current operating conditions. Remote management, digital monitoring and automation have increased the connectivity of OT environments well beyond what many of those systems were designed to withstand. Much OT infrastructure continues to rely on software and firmware that predate modern cybersecurity practices and will never be updated. CISA has repeatedly warned that legacy OT infrastructure—particularly in essential services with limited defensive capacity, such as water systems—remains acutely vulnerable.[46]

This risk is qualitatively different in ways that matter for governance. That's because OT system failure isn't limited to data compromise or temporary digital disruption. It can involve physical effects, operational safety issues and visible service failure. A compromised water system, electricity network or industrial control environment can create harm that's immediate, public and difficult to reverse. OT life-cycle decisions need to therefore be governed by consequence, not merely by age or maintenance cost.

Again, the opportunity side of this equation is routinely underplayed. Legacy OT is not only harder to defend—it's harder to integrate, analyse and optimise. Closed architectures can't readily support

predictive maintenance, real-time operational analytics, flexible energy management or the wider efficiency gains that modern infrastructure operators are already capturing. Operators constrained by legacy OT face higher risk and are locked out of performance improvements that better governed competitors are gaining.

Across parts of the Indo–Pacific power sector, for example, grid operators are investing in distributed-energy integration, demand–response capability and adaptive network management. Operators running legacy OT are less able to participate in those shifts on favourable terms, even where the commercial and system benefits are clear.

IT–OT convergence is usually framed as a source of new cyber risk, which it is, but it's also the enabling condition for the next generation of operational capability. Smart-grid functions, the industrial 'internet of things', advanced maintenance analytics and more adaptive infrastructure all depend on integrated and modernised OT environments. The hardest layer to replace and harden is often the physical edge—the sensors and actuators—which makes it all the more important that the IT systems around them are secured, current and able to protect the components that can't protect themselves. Governments, regulators and operators that approach OT modernisation only as a security-remediation task risk under-scoping it, underfunding it and deferring it for too long. Those that approach it as a resilience and capability investment are more likely to make sounder decisions about procurement, sequencing and system design.

A compressing window, a stronger case for action

These three trends don't alter the basic character of EoL risk. Instead, they alter its urgency, its economics and its governance implications. AI-accelerated exploitation, harvest-now collection and OT convergence each narrow the space for reactive response. Each also increases the value of proactive investment—whether in automated assurance and modernisation tooling, post-quantum transition capability or OT infrastructure designed for both security and performance.

This means the real test is whether governance frameworks help governments, regulators and operators decide which systems most need to be replaced, which can be contained, who carries the risk in the interim, and how transition can be structured to improve resilience, capability and return. The section that follows examines how the governance frameworks in South Korea, the Philippines and India—alongside Australia—perform against that test.

5 Comparative case studies: three pathways to the same governance challenge

The three primary case studies here are best understood not as separate country narratives, but as three distinct paths to the same governance challenge: namely, that unsupported technology remains in service when visibility, ownership, transition mechanisms and accountability don't align.

South Korea illustrates policy-created lock-in compounded by deferred replacement and fragmented oversight. The Philippines illustrates procurement-driven unsupportability, where the administrative system can't sustain vendor support relationships. India illustrates fragmentation at scale: a large and increasingly capable regulatory system with important sectoral progress, but no unified life-cycle governance architecture. A 'spotlight' on Australia is also provided as a distinct section for comparative purposes.

Read comparatively, the cases show that EoL risk has multiple origins—vendor discontinuation, earlier policy choices, procurement settings, budget distortion, incomplete replacement programs, or the absence of an authority able to force a decision. The central policy question in each case is the same: whether the risk can be seen clearly, ownership assigned, and recognition converted into funded and enforceable transition.

Trigger incidents and strategic overlays

In each jurisdiction, trigger incidents exposed not simply a technical failure, but strategic constraints in the broader operating environment.

In South Korea, the SK Telecom breach in April 2025 (23 million users affected; servers running an operating system with a 2016 patch never applied; three years of undetected access)[47] and the Daejeon data-centre fire five months later exposed the same structural problem from two directions: unsupported software and deferred infrastructure replacement.[48] The deeper strategic issue, however, is actually concentration risk—a highly digitised state in which too many critical functions depend on common localised infrastructure, the continuity assumptions of which haven't been fully tested.[49]

In the Philippines, the 2023 PhilHealth breach[50] and the 2021 BDO Unibank[51] hack illustrated procurement-driven limitations and EoL infrastructure exploitation, respectively.[52] Both sit within a wider strategic overlap: life-cycle decisions in energy and telecommunications can't be separated from vendor-trust and geopolitical questions.[53]

In India, the pattern is breadth rather than concentration. Incidents across banking,[54] health care,[55] nuclear and energy[56]

show EoL-related vulnerability appearing across multiple sectors with different regulators and life-cycle expectations. This shows that the governance challenge scales with the country: life-cycle management needs to operate across diverse ownership models and federal relationships, from advanced national infrastructure to lower capacity state and local systems.[57]

A 'spotlight' on Australia to complement this analysis is on page 17.

Dominant challenge

South Korea: policy-created lock-in

South Korea leads the Organisation for Economic Co-operation and Development (OECD) Digital Government Index, has near-universal broadband penetration, and built comprehensive e-government infrastructure early. That success produced what researchers at UC Berkeley's Center for Long-Term Cybersecurity have called the 'South Korean Digital Paradox': the policies that accelerated digitisation from the late 1990s also created systemic legacy debt that's now costly to unwind.[58]

The paradox has its roots in policy design. Former government mandates required the use of ActiveX controls and associated proprietary security plug-ins to accelerate online banking and e-government services, locking the ecosystem into specific vendor platforms and browser dependencies.[59] When those technologies became obsolete, the transition cost became substantial. Physical network separation requirements in the financial sector, while providing a security layer, also perpetuated isolated systems that received fewer updates and became progressively less suited to a fast-evolving threat environment.[60] Legacy, here, was produced by earlier policy success, not neglect.

The Philippines: procurement-driven unsupportability

The defining challenge for the Philippines isn't a vendor discontinuing support—it's an administrative system constrained in its capacity and capability to maintain a vendor relationship in time to preserve support. Capacity constraints also shape the problem. Agencies must often manage ageing and highly interdependent systems with limited specialist personnel, uneven access to technical expertise and significant dependence on external vendors and service providers. Those factors make replacement and modernisation difficult even when risks are recognised.

The PhilHealth Medusa ransomware breach—one of the most consequential cyber incidents in recent Philippine government history—began not with an unpatchable system but with a security control the antivirus subscription of which expired in April 2023 because of a procurement cycle that didn't treat security continuity as an explicit requirement.[61] Five months later, Medusa ransomware exfiltrated 734 gigabytes of data from an effectively unprotected system.[62] The instruments that might be expected to govern this risk—the national cybersecurity strategy, the privacy commission's data security circular, and the Bangko Sentral ng Pilipinas IT risk regime, which applies only to financial institutions supervised by the central bank and doesn't extend to government corporations—don't address it holistically.[63] The most consequential reform attempt of the past decade is a procurement statute, the Republic Act No. 12009, rather than a cybersecurity law.[64] It showed that procurement reform can remove one transactional barrier to maintaining supported technology. But it's still not the same as life-cycle governance, which needs to also define obligations for visibility, supportability, retirement, exception and transition.

India: fragmentation at scale

India's EoL technology challenge is inseparable from its size: a population of 1.4 billion, a federal structure distributing authority across central and state governments, and a critical-infrastructure landscape spanning from nuclear power plants to rural water systems. There's no unified national policy addressing technology obsolescence. Relevant provisions are distributed across multiple regulatory bodies, each with different mandates, enforcement mechanisms and levels of specificity. Some sectors—notably power, finance and elements of critical-infrastructure regulation—have developed comparatively advanced cybersecurity and resilience frameworks. The Indian case isn't one of institutional absence. It's uneven progress without integrated coherence.

Governance maturity and institutional ownership

None of the jurisdictions examined has implemented a fully integrated governance framework for EoL technology risk. The comparative distinction lies instead in how far life-cycle governance has developed within existing cybersecurity, procurement and critical-infrastructure systems, and whether any institution possesses the authority to convert visibility into enforceable replacement, remediation or risk-acceptance decisions. Across all three case studies, life-cycle risk is still managed indirectly through adjacent mechanisms such as vulnerability management, procurement controls, operational resilience obligations and sector-specific assurance regimes, rather than through dedicated EoL governance frameworks.

South Korea

South Korea demonstrates a structurally mature but life-cycle-blind model. The country possesses one of the Indo-Pacific's most sophisticated cybersecurity architectures, supported by extensive digital infrastructure, high levels of connectivity and a dense network of sector-specific regulators. Yet responsibility for cybersecurity governance remains fragmented across the National Intelligence Service, the Ministry of Science and ICT, the Korea Internet and Security Agency, the Personal Information Protection Commission and multiple sector-specific regulators.[65] No institution holds end-to-end authority for unsupported technology risk across government and critical infrastructure.

Existing obligations address life-cycle risk only indirectly. The Information and Communications Infrastructure Protection Act requires designated critical information infrastructure operators to conduct annual vulnerability assessments, while the Network Act and related sectoral frameworks impose general security safeguard obligations.[66] However, none of these instruments defines EoL or EoS status as a distinct governance category, establishes mandatory replacement timelines or requires operators to maintain inventories of unsupported technology. The 2024 National Cybersecurity Strategy signals growing recognition of the problem through its proposed 'minimum-security requirements for national infrastructure operation systems' and plans to phase out certain legacy security software from 2026, but those initiatives remain partial and implementation detail remains limited.[67]

The resulting governance gap is increasingly visible in practice. A February 2026 Board of Audit and Inspection penetration test reportedly compromised all seven tested public-sector systems, and investigators identified legacy infrastructure, fragmented identity controls and uneven cyber maturity as recurring weaknesses.[68] Likewise, the 2025 SK Telecom breach exposed the operational consequences of weak life-cycle governance: systems running obsolete or unpatched operating systems, poor network segmentation, and a security patch released in 2016 that had never been applied reportedly enabled attackers to maintain access for nearly three years.

The nuclear sector illustrates the same pattern in a more mature regulatory environment. KINAC Regulatory Standard RS-015 establishes detailed cybersecurity controls for nuclear facilities and critical digital assets, but a 2022 Korea Atomic Energy Research Institute study (22A-096) reviewing the framework confirmed that none of the 182 underlying technical control requirements explicitly addresses technology obsolescence, EoL, EoS or digital asset life-cycle management.[69] A critical digital asset can therefore satisfy the sector's cybersecurity control framework while remaining operationally unsupported. This reflects a broader characteristic of the South Korean model: extensive operational and compliance-focused cybersecurity governance without a corresponding mechanism for governing the life-cycle state of the underlying technology base.

South Korea's challenge is compounded by what researchers have described as the country's 'digital paradox': rapid state-led digitisation and strong domestic technology ecosystems created early-mover advantages that now generate concentrated legacy exposure. Mandatory or quasi-mandatory domestic software ecosystems, including the long-running Korea Security Applications regime, created forms of technology lock-in that persisted long after the underlying technologies became obsolete. The result isn't merely fragmented responsibility, but the absence of any authority capable of forcing retirement, replacement or formal risk-acceptance decisions across the system.[70]

The Philippines

The Philippines represents a life-cycle-governance environment in which unsupported technology risk is addressed only indirectly, and often through procurement or administrative processes rather than dedicated cybersecurity obligations. The National Cybersecurity Plan 2023–2028 contains no reference to EoL, EoS, obsolescence, technology refresh, vendor support, unsupported software or legacy systems. Its asset-visibility provisions focus primarily on scanning internet-exposed assets for vulnerabilities rather than maintaining life-cycle inventories of deployed technology.[71] Likewise, the National Privacy Commission's Circular 2023-06 refers to 'end-of-life' only in the context of physical storage-device destruction, while the Bangko Sentral ng Pilipinas IT risk framework requires inventories, patch management and media-disposal procedures without imposing obligations to identify, report or remediate systems operating beyond vendor support.[72]

The result is a governance model in which unsupported technology can remain operational without triggering a formal life-cycle-management response. Responsibility is distributed rather than unified: the Department of Information and Communications Technology plays a central role in ICT standards and digital government strategy, but individual agencies manage their own procurement life cycles and capital planning; technology upgrades are funded through the national budgeting process rather than any dedicated replacement pipeline. Most actionable obligations remain sector-specific, and institutional alignment across procurement, cybersecurity and operational ownership is limited. That fragmentation is compounded by practical constraints (skills shortages, reliance on external vendors and the complexity of deeply interdependent legacy architectures) that limit the ability of agencies to translate awareness of life-cycle risk into timely remediation or replacement decisions.

In practice, procurement authorities function as *de facto* life-cycle governors without being explicitly designed for that role. As PhilHealth's account of the 2023 incident illustrates, the ability to maintain supported and operational cybersecurity controls can depend as much on procurement timelines, renewal authorities and administrative execution as on technical security policy itself. Republic Act No. 12009, the New Government Procurement Act, begins to partially address this gap. Section 13 introduces life-cycle cost analysis and life-cycle assessment requirements in government procurement, while section 35(m) permits direct retail purchase of online subscriptions, including 'system protection software', without competitive bidding under negotiated procurement procedures. Those provisions reduce some procurement friction but don't establish affirmative vendor-support obligations, mandatory replacement timelines or an integrated governance regime for unsupported technology.[73] The PhilHealth case establishes a broader comparative principle: where procurement, cybersecurity and operational ownership aren't institutionally aligned, systems can become functionally unsupported without any actor explicitly deciding that outcome.

India

India presents the most operationally mature but institutionally uneven life-cycle governance model of the three case studies. Rather than a uniform national framework, the Indian system is characterised by sector-specific advancement, particularly in electricity generation and elements of the broader cybersecurity audit and software-assurance regime. CERT-In's 2025 cybersecurity audit guidelines require legacy systems to remain within audit scope, be documented through formal risk exceptions, and remain visible to boards and auditors through SBOMs obligations.[74] The power sector has progressed further still. Article 7 of the Central Electricity Authority's 2021 cybersecurity guidelines requires responsible entities to maintain inventories of equipment nearing end of life or unsupported by original equipment manufacturers (OEMs), secure board approval for replacement plans, implement compensatory controls in consultation with OEMs until replacement occurs and maintain procedures for the secure disposal of legacy systems.[75] A September 2022 amendment further requires CERT-In-empanelled OT auditors to assess IT systems at least every six months and OT systems at least annually, with critical and high-risk vulnerabilities remediated within one month.[76]

The draft Central Electricity Authority (CEA) (Cyber Security in Power Sector) Regulations 2025, which are currently under consultation and not yet in force, would extend this approach further by requiring vendors to provide security patches and updates for the contract period or useful life of the system, whichever is longer; notify operators when products reach EoS or EoL status; and provide CERT-In-compliant SBOMs covering firmware and software deployed within critical systems.[77] Telecommunications licensing frameworks similarly require annual cybersecurity audits and infrastructure modernisation activities by service providers, reinforcing the broader direction of travel towards life-cycle-based security assurance.

Those developments demonstrate that high-consequence sectors can achieve comparatively sophisticated life-cycle governance where operational risk, regulatory authority and sector-specific oversight align. However, that maturity remains uneven across the wider governance architecture. India's procurement framework continues to prioritise acquisition efficiency and value for money, but impose no general obligation for vendors to disclose life-cycle timelines, minimum support durations or EoS pathways at the point of procurement.[78] Likewise, the National Critical Information Infrastructure Protection Centre's publicly available *Guidelines for Critical Information Infrastructure Protection* (version 2) don't define EoL or EoS, require life-cycle-status tracking within asset inventories, or establish mandatory retirement timelines or EoL-specific compensating controls.[79] Telecom policy has similarly focused on trusted-source requirements for new equipment while explicitly avoiding mandatory replacement obligations for existing deployments.[80] The Reserve Bank of India mandated phased ATM migration requirements, but no public compliance data has been identified.[81]

India's central governance challenge is therefore not an absence of activity, but the difficulty of converting sector-specific progress into horizontally coherent accountability. Responsibility is dispersed across multiple regulators with different authorities, enforcement tools and degrees of life-cycle specificity. Some sectors have developed comparatively mature governance mechanisms for unsupported technology, while others continue to rely primarily on conventional cybersecurity or procurement controls. India consequently illustrates a broader comparative point: life-cycle governance can advance materially within strategically important sectors even while the wider national governance system remains fragmented and uneven.

Transition mechanisms: budgets, procurement and parallel running

Recognising the problem isn't the same as having a credible transition pathway. In each case, the binding constraint lies in a different part of the system. For many organisations, that constraint is structural: capital planning cycles that weren't designed to absorb the pace of technology obsolescence, multi-year appropriations processes that lock in platform decisions years before EoL risk materialises, and approval chains that treat modernisation as a discrete project rather than a continuous life-cycle obligation.

South Korea

The transition challenge is shaped by budget choices and continuity gaps. Approximately 69% of South Korea's public-sector software project budget goes to operations and maintenance, while only 23% is directed to new development.[82] Yet, the government has simultaneously cut maintenance budgets to fund flagship digital transformation and AI initiatives: for example, the Administrative Information Sharing System saw a 57% maintenance budget cut between 2022 and 2024.[83] That underinvestment is starting to register on the cost side. Following the September 2025 fire at the National Information Resources Service (NIRS) data centre, the National Assembly Budget Office's review of the FY2026 budget records a ₩252.3 billion increase to the Ministry of the Interior and Safety's 'Integrated Construction of Aging Central Government Facilities (Informatization)' line — attributed by budget office to accelerating dual-system redundancy for major government information systems as a direct follow-up to the fire.[84]

The consequences materialised in September 2025 when a lithium-ion battery fire at the NIRS data centre in Daejeon destroyed 96 systems and forced the precautionary shutdown of a further 551. The G-Drive platform – the only one of the 96 destroyed systems without a backup, owing to its size – alone lost 858 terabytes of government data.[85] At the same time, 34.6% of the centre's equipment was outdated due to delayed replacement; ₩25.1 billion allocated for IT infrastructure in 2024 went entirely unspent; and a disaster recovery centre completed in May 2023 never became operational because of budget constraints.[86] The Digital Platform Government initiative, which targeted migrating 10,000 government systems to cloud-native by 2030, has since been

subsumed under an 'AI Government' agenda—part of the recurring pattern the OECD noted of successive Korean administrations recasting digital priorities, requiring careful balancing of novelty against long-term strategic continuity.[87] South Korea demonstrates how a digitally advanced state can still struggle to modernise safely when replacement budgets, operational continuity and strategic digital priorities pull in different directions.

The Philippines

In the Philippines, the transition challenge is administrative and operational, as well as a capability challenge: agencies may recognise the need to modernise yet still face shortages of specialist personnel, dependence on external vendors, and the technical complexity of replacing systems deeply embedded in business processes over many years. No fully integrated national mechanism consistently ensures that technology replacement is planned, funded and enforced before systems become high-risk or unsupported—and the Government Procurement Policy Board wasn't designed to fill that role. The result is that life-cycle decisions default to sector-specific processes, individual agency budget cycles and *ad hoc* procurement authorities operating without a common standard or enforcement backstop.

National Telecommunications Commission Memorandum Circular No. 003-09-2025 mandates the phased decommissioning of 2G and 3G mobile networks and a nationwide 3G switch-off by 31 December 2026, making it the most explicit mandatory-retirement directive identified in the three-country study.[88] But the broader challenge is better illustrated by the Land Transportation Office's Land Transportation Management System: a commission of audit (COA) found that the office incurred ₱438.7 million in avoidable fees in 2024 because it ran its legacy Stradcom IT system in parallel with an incomplete Dermalog replacement whose data centre connectivity fell short of contracted specifications, delivering 80 megabits/second per line against a requirement for three redundant 200 megabits/second connections, leaving the system vulnerable to the nationwide outage that occurred in October 2024.[89]

The starkest example is the weather warning network of the Philippine Atmospheric Geophysical and Astronomical Services Administration. Only 10 of its 19 Doppler weather radars were operational as of 31 December 2024;[90] a station destroyed by Typhoon Rolly in 2020 was still non-functional four years later because no adequate budget provision was made across successive appropriations cycles and spare parts carry an eight-month lead time.[91] COA recommended direct contracting to bypass standard procurement delays.[92] In a country facing 15–20 typhoons per year entering territorial waters, a weather warning system operating at half capacity isn't just a technical shortfall.[93]

India

In India, the transition challenge is uneven sectoral maturity. Electricity is the most advanced sector, in which draft CEA regulations point towards binding vendor life-cycle obligations and stronger software-assurance requirements.[94] Parts of the financial and securities environment are also moving: the Securities and Exchange Board of India's 2024 Cybersecurity and Cyber Resilience Framework introduced SBOM requirements for market entities,[95] and CERT-In's 2025 SBOM technical guidelines specify 'EOL Date' as a named minimum data field in the SBOM schema (Table 5; Section 4.2(11); worked example in Table 6), placing EoL status at the level of first-class metadata rather than ancillary detail.[96]

Whether that specification translates into systematic life-cycle tracking across these regimes is a separate question: the guidelines are advisory rather than mandatory, so take-up depends on voluntary adoption rather than enforcement. Even so, they provide a visibility foundation that could enable life-cycle tracking.[97]

Progress outside priority sectors remains more fragmented. India's federal and ministerial budgeting structure means that departments and state governments largely manage IT modernisation independently, producing uneven replacement cycles, inconsistent uplift priorities and limited cross-government visibility of ageing systems. By contrast, cross-government procurement settings, telecom legacy replacement and nationally consistent EoL visibility remain less developed.[98] No centralised mechanism tracks EoL across government. India shows that transition capability can improve meaningfully in priority sectors while the wider system remains uneven—making national coherence difficult to achieve.

Comparative implications

While the three case studies are analytically distinct, they converge on the same takeaway: no single reform lever is the solution. Strategy without inventory has limited effect. Visibility without obligations is insufficient. Retirement mandates without workable transition mechanisms produce paralysis or unsafe parallel running. Equally, transition mechanisms require sufficient workforce capability, technical expertise and vendor engagement if recognised risks are to be converted into successful replacement programs. Procurement reform without vendor life-cycle obligations leaves structural exposure in place. Sectoral progress without national coherence creates islands of maturity within a broader field of unmanaged risk.

What the comparison establishes isn't yet a full model of good governance, but a clear capture of the various instruments needed to address transition (Table 1). An effective governance framework needs to manage vendor-driven obsolescence, procurement-driven unsupportability and fragmentation-driven inaction—concurrently—while aligning budget, procurement and continuity planning to make transition credible in practice. The conclusion builds from these findings towards a practical model for decision and action—answering the question 'What does good look like?'

Table 1: Three pathways to the same governance problem

Jurisdiction	Dominant pathway	Why it persists	What blocks transition
South Korea	Policy-created lock-in	Life-cycle risk is governed indirectly; no single body can force replacement or exception across sectors.	Deferred replacement, budget discontinuity, and weak continuity execution.
Philippines	Procurement-driven unsupportability	Procurement, cybersecurity and operational ownership aren't aligned; systems can become unsupported without any actor intending that outcome.	Renewal delays, administrative friction, and prolonged parallel running of incomplete replacements.
India	Fragmentation at scale	Sectoral progress exists, but responsibility is dispersed across regulators and sectors without a unified life-cycle framework.	Uneven sectoral maturity, weak procurement settings, and no centralised visibility of EoL exposure.

Spotlight on Australia: governance maturity, persisting exposure

Australia is the analytical inverse of this report's three case studies. Where South Korea, the Philippines and India illustrate what happens when governance frameworks are absent, fragmented or structurally misaligned with operational reality, Australia illustrates what happens when maturity is concentrated largely at the federal government core but doesn't yet extend uniformly to the critical-infrastructure operators whose systems also carry national-security consequences , though that has begun to change for high-risk asset classes. That distinction reframes the challenge. It shows that EoL technology risk isn't simply a problem of drafting policies, standards or regulatory frameworks intended to achieve governance maturity. It's an investment, procurement and institutional execution problem. It remains so even in jurisdictions that have largely solved the drafting.

Government obligations

At the apex of the national framework sits the Protective Security Policy Framework (PSPF), which requires federal government entities to maintain secure ICT systems across all stages of the technology life cycle. PSPF Direction 002-2024 further required entities to complete technology asset stocktakes of internet-facing systems and develop technology security risk management plans by June 2025, explicitly elevating life-cycle visibility and legacy IT management into mandatory governance activity.[99]

The PSPF now carries a formal definition of legacy IT and a dedicated control (0093) requiring agencies to identify, manage and report on legacy technology risks through the annual PSPF reporting cycle. Technical implementation guidance is provided through the Australian Signals Directorate's (ASD's) *Information security manual* (ISM), which operationalises those obligations through granular controls addressing unsupported software, operating systems, network devices and compensating controls. They include requirements to remove unsupported online services (ISM-1905), unsupported office productivity, browser and email software (ISM-1704) and residual unsupported applications (ISM-0304); the replacement of unsupported operating systems (ISM-1501), internet-facing network devices (ISM-1753), non-internet-facing network devices (ISM-1981) and networked IT equipment (ISM-1982); and compensating controls where immediate replacement isn't feasible (ISM-1809).[100]

These controls are reinforced by ASD's 'Essential Eight' maturity model (Figure 1). Critically, the core obligations relating to unsupported online services, office productivity software and operating systems sit at Maturity Level 1, meaning the operation of unsupported systems in these categories constitutes a baseline cyber resilience failure, not an aspirational uplift target.[101] For non-corporate federal entities, this architecture establishes a reasonably robust and enforceable EoL governance floor.

Critical infrastructure

The governance picture for critical-infrastructure operators is substantially different. The operative legislative framework is the SOCI Act and related reforms under which prescribed operators in the 11 defined asset classes are required to maintain and implement a critical infrastructure risk management program (CIRMP). The CIRMP obligations require operators to identify and manage hazards—including cyber and information security hazards—that could materially affect their assets. The exception is telecommunications, which operates under a separate regulatory instrument entirely.[102]

Until recently, the SOCI/CIRMP framework contained no EoL-specific controls. That is changing for the highest-risk asset classes. New enhanced cyber and information security hazard requirements oblige responsible entities for the specified assets to maintain a process that, so far as reasonably practicable, minimises or eliminates the material risks of failing to patch systems in a timely manner and of failing to replace legacy systems or mitigate risks from redundant, unsupported, obsolete or discontinued technology. This is the first time the framework names legacy and unsupported technology as a material risk rather than leaving it to general risk management. The obligation is outcome-based, not prescriptive, it sets a risk to minimise rather than a control equivalent to ISM-1501 or ISM-1905, and it reaches only the defined high-risk subset, not all CI operators. For the broader CIRMP population, implementation remains principles-based.[103]

The framework remains largely principles-based for most asset classes, leaving operators to identify life-cycle risk through

Figure 2: Australian Government's ICT life-cycle governance architecture

Protective Security Policy Framework (PSPF) Apex obligation—secure ICT systems across all life-cycle stages Direction 002-2024: technology asset stocktakes and risk management plans required by June 2025	
Information security manual (ISM) Operationalises PSPF life-cycle obligations	**'Essential Eight maturity model'** Staged uplift. Extends to critical infrastructure and the broader economy
ISM controls—unsupported systems	
Action	**Control**
Remove	ISM-1905—unsupported online services
	ISM-1704—unsupported office productivity, browser and email software
	ISM-0304—residual unsupported applications
Replace	ISM-1501—unsupported operating systems
	ISM-1753—internet-facing network devices
	ISM-1981—non-internet-facing network devices
	ISM-1982—networked IT equipment
Compensate	ISM-1809—compensating controls where immediate replacement isn't feasible

whatever risk-management and cybersecurity practices they choose. That design carries a predictable cost. Legacy systems are rarely decommissioned the moment a replacement program begins—operational dependencies, integration requirements and continuity concerns all push the other way—so organisations often run old and new systems in parallel for extended periods. The result is a wider attack surface and higher operating costs, diverting resources that could otherwise fund security uplift and modernisation. The assumption underpinning this approach is that sound cyber risk-management will flag unsupported or obsolete systems as a matter of course. In practice, that judgement sits with operators, producing wide variation in how life-cycle risk is interpreted, prioritised and governed across sectors.

The Essential Eight model is available to critical-infrastructure operators as a voluntary uplift framework, and ASD actively encourages its adoption. ASD has, for example, sought to elevate the issue publicly, naming 'manage legacy IT risks' as one of four headline recommendations in its 2024–25 *Annual cyber threat report*— reinforcing the legacy-IT guidance it had already issued through a dedicated cyber.gov.au management hub, with separate executive and practitioner publications aimed at boards and IT operators.[104] Behind this sits the 2023–2030 Australian Cyber Security Strategy and its A$586.9 million committed across the strategy's six 'cyber shields'—including critical-infrastructure protection and federal government system uplift.[105] While welcome, voluntary adoption isn't equivalent to mandatory obligation. The constraint here's structural, not rhetorical. As a technical authority, rather than a regulator, ASD's *Fortify* guidance can sharpen how operators understand the risks, but it can't compel replacement. Extending the framework into the critical-infrastructure sector doesn't presently create an enforceable EoL governance floor—it creates a benchmark that operators may reference, partially adopt, or disregard.

The governance gap

The result is a structural inversion. Federal government entities managing administrative systems are subject to specific, enforceable, technically granular EoL obligations. Critical-infrastructure operators managing systems the compromise of which could affect essential services, public safety or national security have until recently been subject to broad risk-management obligations with no equivalent technical specificity on life-cycle management. The enhanced requirements narrow that gap for high-risk asset classes but do not close it for the broader CI population.

State and territory government entities and corporate federal entities, including government business enterprises, compound this further, operating under separate arrangements with varying degrees of maturity and enforcement that sit outside the PSPF framework entirely (see box).

The reality in numbers

The federal Digital Transformation Agency's *APS digital workforce insights report* states that the digital landscape within government is 'legacy-heavy'. Over 40% of systems surveyed by the agency self-identified as needing significant investment, at either 'approaching legacy' or 'at legacy' stage.[106]

The 2025 *Commonwealth Cyber Security Posture report*[107]—the government's own annual measurement to parliament—found that 59% of federal entities reported that their use of legacy technology had impeded their ability to implement Maturity Level 2 of the Essential Eight—an improvement on 71% the year before but still a majority of the public sector facing challenges. Only 22% of entities achieved Maturity Level 2 across all eight mitigation strategies—a figure lower than the 25% reported in an earlier 2023 report, in part due to the tightening of requirements to achieve Maturity Level 2. Asked to identify the single most significant

reason for their continued use of legacy IT, entities pointed to insufficient dedicated funding (34%), lack of a viable replacement (18%) and lack of prioritisation (16%), with the remainder spilt across skilled-personnel constraints and unclassified or other reasons.[108] Industry modelling estimates that more than 70% of federal entities reported legacy IT as a barrier to Essential Eight compliance, and that accelerated public cloud adoption alone could yield A$13.5 billion in cumulative savings over 2026–2035.[109]

An Australian National Audit Office (ANAO) audit found that, by mid-2024, only around 5% of Defence organisation IT systems requiring authorisation had been recorded in Defence's authorisation management system.[110] Even within that limited pool, nearly half held an 'Expired' or 'No accreditation' status, and Defence was frequently unable to determine whether noncompliant systems were still operational. The ANAO further concluded that the dataset itself was incomplete and unreliable, meaning the accreditation status of the overwhelming majority of systems was effectively unknown.[111] The Defence Digital Strategy and Roadmap targets continued retirement of legacy across all major capability streams through to the 2027–28 financial year; the audit record suggests that program is still very much in progress.

OT compounds the exposure. ASD's *CI Fortify* guidance, issued in October 2025, observes that OT environments 'often have legacy devices, outdated protocols and slow replacement cycles' and that many critical-infrastructure operators remain 'tied to long-standing engineering contracts that weren't designed with cyber security in mind'. The challenge is rarely technical alone: operational processes, workforce practices and contractual arrangements often evolve around long-lived industrial systems, making replacement significantly more complex than deploying a newer technology. The result, ASD concludes, is 'a widespread, persistent cyber risk to Australia's most vital services'.[112] That's the lead national cyber agency acknowledging that architecture and security controls can't, by themselves, overcome the organisational and operational barriers that delay modernisation.

The gap between policy maturity and operational delivery isn't a new observation in Australian cyber governance. Rachael Falk, a former CEO of the Cyber Security Cooperative Research Centre and a co-author of the Australian Institute of Company Directors' Cyber Security Governance Principles, has observed publicly that getting cybersecurity consistently onto board agendas 'in a significant way' remained difficult until recently—and that it 'should be on every agenda of every board meeting' for organisations of all sizes.[113] The qualifier matters because it recognises that 'should be' isn't the same as 'is'. It's the same gap that the compliance data above records—between an architecture that creates the obligation and an institutional culture that hasn't yet fully absorbed it.

What this means

Australia's experience shows that, while architecture is necessary, alone it's insufficient. Execution is key. Four implications follow.

1. First, funding structures need to align with policy obligations requiring replacement and life-cycle management. Annual budget cycles that treat legacy modernisation as discretionary capital expenditure will continue to defer the decisions that mandatory frameworks nominally require. Multi-year funding envelopes tied to consequence tiers, with independent progress reporting, are the minimum condition for systematic remediation.
2. Second, accountability needs to be personalised and enforced. Where the barriers are prioritisation and leadership commitment, anonymous aggregate reporting achieves little. Accountable officers' obligations for high-consequence legacy exposure—named, documented and subject to regulatory consequence—convert governance architecture from aspiration into obligation.
3. Third, operational technology must be governed by consequence, not procurement convention. OT replacement cycles embedded in engineering contracts that predate modern cyber requirements need override mechanisms keyed to risk tiers, not contract tenure.
4. Fourth, Australia should treat the demonstrable reduction of legacy exposure as a strategic governance objective, not merely a compliance exercise. Australia is increasingly viewed internationally as a reference point for critical-infrastructure governance, meaning other jurisdictions are likely to look to its life-cycle governance settings as a regional benchmark. The export and strategic value of that model will depend less on the sophistication of architecture than on its ability to demonstrate measurable execution, sustained replacement activity and declining residual exposure over time.

6 Conclusion: What 'good' looks like

Legacy technology governance currently falls short at two distinct levels, and addressing only one risks producing limited results. At the policy and regulatory level, governments lack the instruments to make life-cycle governance obligatory, visible and consistently enforced across sectors and ownership models. At the enterprise level, operators lack the frameworks, decision structures and investment logic to manage EoL risk as a governed liability rather than an inherited condition.

Neither level is sufficient on its own. Enterprise frameworks without regulatory reinforcement will be adopted selectively; that is, by operators already motivated to act, in sectors already under scrutiny and at a pace determined by internal appetite rather than systemic risk. Regulatory frameworks without organisational architecture to implement them produce compliance theatre where obligations can't be met because the underlying governance capability doesn't exist.

The 'Essential Eight' offers a relevant analogy. Developed by ASD from its experience producing cyber threat intelligence, responding to cyber incidents, conducting penetration testing and assisting organisations to improve their security posture, the Essential Eight is a prioritised set of baseline mitigations that's well known to Australian operators. It works not because any single control is transformative but because the combination is coherent. Each mitigation reinforces the others, the set is prioritised by impact, and organisations can progress against a defined maturity model. Critically, it shifted cybersecurity from a best-practice aspiration to a governable, auditable baseline, giving operators, boards, auditors and regulators a shared reference point for measuring progress.

EoL governance needs the same thing. This report proposes the 'Legacy Five' (Figure 2) as a framework built around five levers for action, but with a twist: two applications of the same core framework; one for governments and regulators and one for enterprise operators. Both work from the same underlying logic, reducing the translation gap between policy intent and operational practice.

The first section sets out the government-facing levers—a consolidated set of actions to make life-cycle governance mandatory, funded and enforceable.

The second section sets out enterprise levers that operators and boards can implement in response. Regulation creates the obligation while enterprise governance provides the mechanism for meeting it. In each instance, the five levers aren't all-or-nothing. They can be applied progressively, calibrated by sector and risk tier, and sequenced to build on existing regulatory architecture. While partial adoption produces partial results, the model is explicitly designed for a staged approach, with near-term actions achievable within existing authority and medium-term actions that may require new instruments or legislative change.

Figure 1: The 'Legacy Five'—two sides, one coin

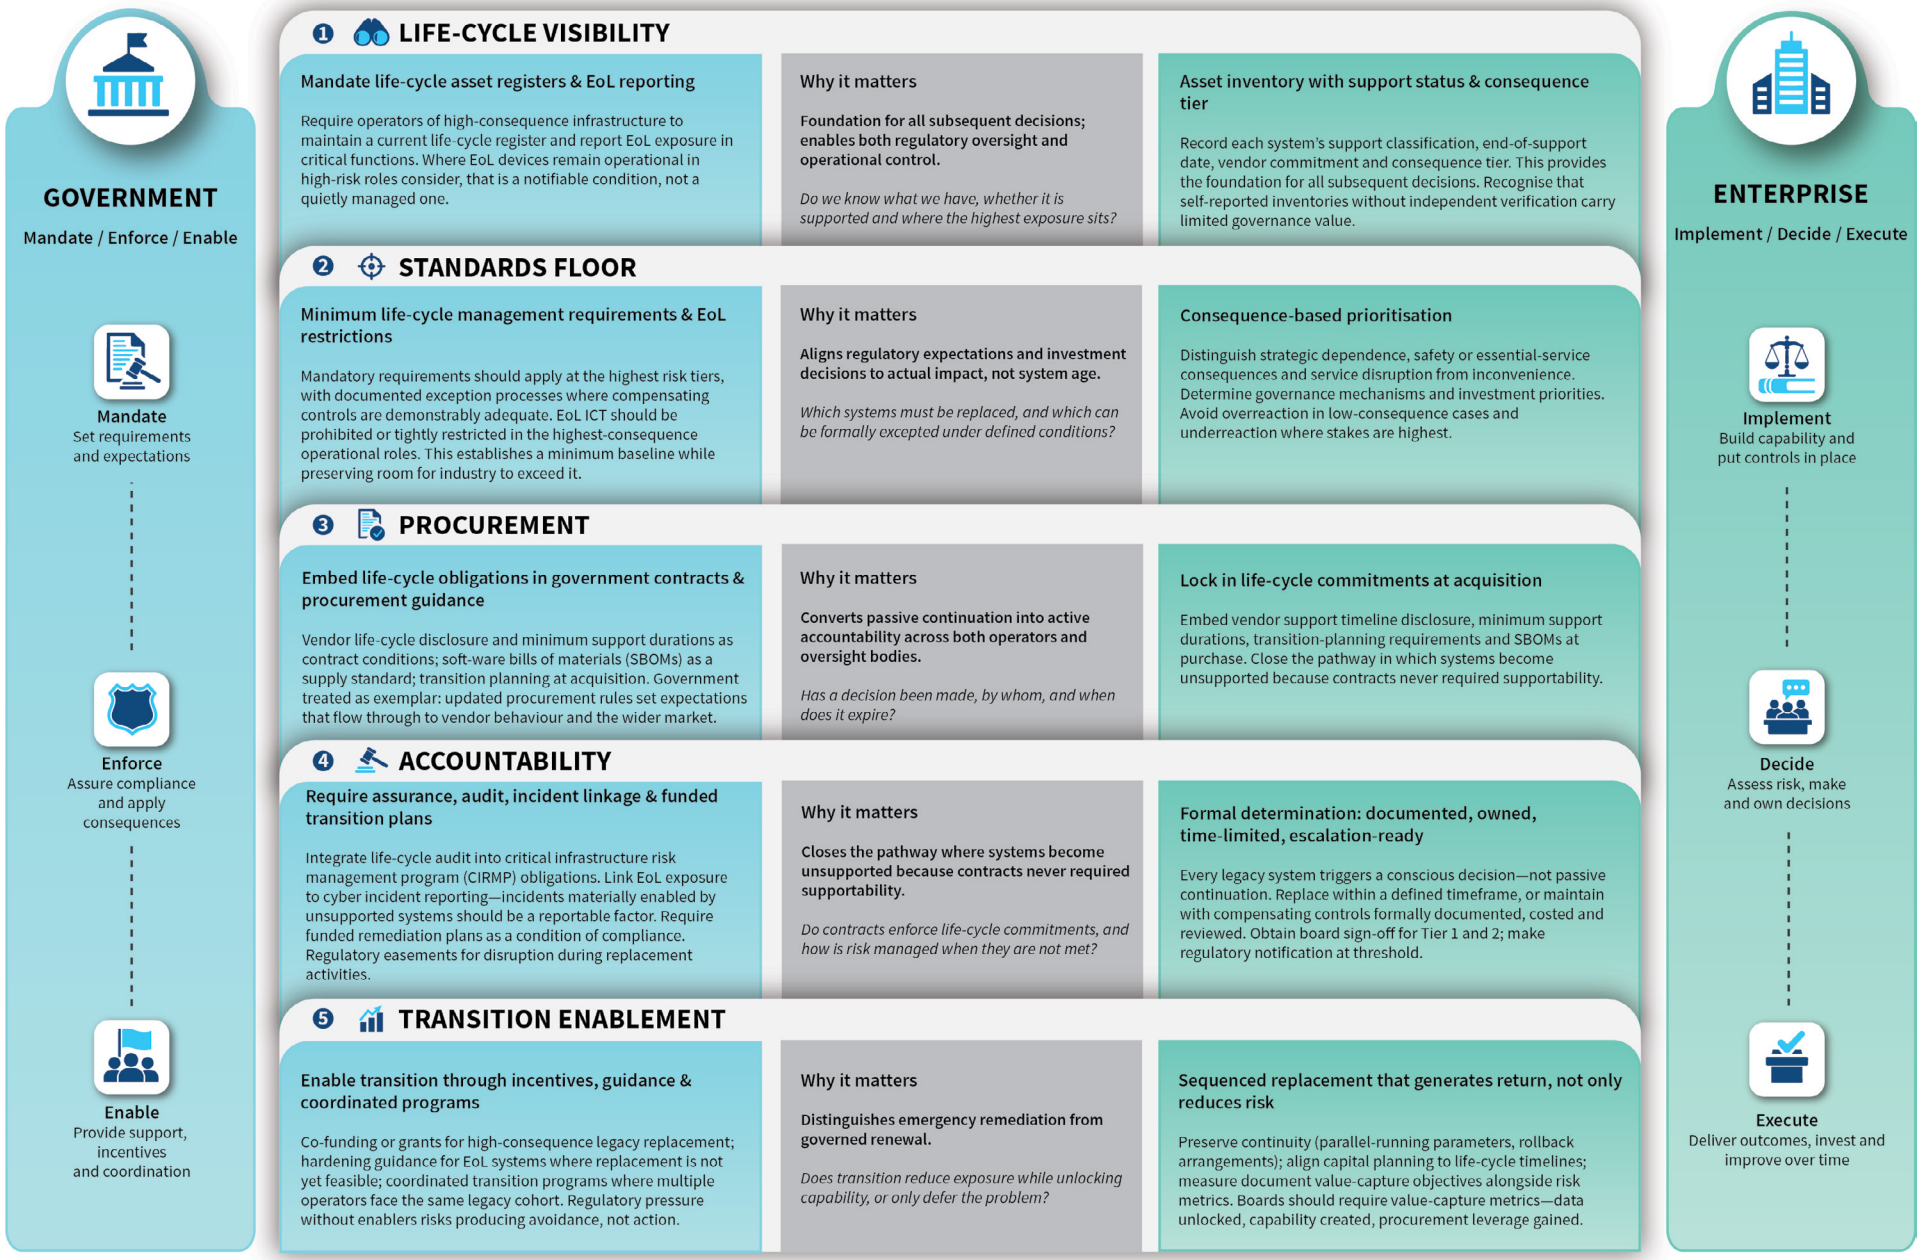

Policy levers: What governments need to do

In many cases, operators don't manage EoL risk adequately because they lack visibility of their technology estate or because existing governance frameworks don't force action. Governments have the regulatory, procurement and assurance levers to change that. The challenge is to apply them in a coordinated and scalable way that reflects operational consequence while creating a realistic pathway for uplift rather than binary compliance.

'Legacy Five' for governments

Lever 1: Make it visible—mandate life-cycle asset registers and EoL reporting

Governments can't manage unsupported technology risk they can't see. Many public- and private-sector operators still don't maintain accurate technology inventories, let alone records of support status, EoS dates or operational consequence. Governments should treat life-cycle visibility as a baseline governance obligation.

Operators of high-consequence systems should be required to maintain life-cycle registers that record support status, EoS dates, vendor commitment status and consequence tier, and to report material EoL exposure affecting critical functions. Where unsupported systems remain in high-risk operational roles, that exposure should become a notifiable condition subject to escalation and review.

Existing frameworks already provide practical implementation pathways. Australia's PSPF and the SOCI Act's critical-infrastructure risk-management obligations can be extended to incorporate life-cycle requirements, while CISA's binding operational directives provide a reference model for implementing visibility and remediation requirements at scale.

Specifically: require designated high-consequence operators to submit a standardised 'life-cycle exposure return' to the relevant sector regulator on an annual basis, disclosing all systems operating beyond vendor EoS, their consequence tier, and the controls or transition plans in place. Require the responsible Chief Information Security Officer (CISO) or equivalent to attest to the accuracy of that return personally, creating a direct accountability link between the register and named leadership.

Lever 2: Set the floor—minimum life-cycle management requirements and EoL restrictions

Visibility alone doesn't reduce exposure. Governments need to establish minimum life-cycle management standards for high-consequence systems, including obligations to maintain vendor-supported technology in critical functions, document and escalate unsupported-system exceptions, and restrict or prohibit EoL ICT in the highest risk operational environments.

These requirements need to be calibrated to operational consequence. Standards that can't realistically be met at scale will produce administrative noncompliance rather than measurable risk reduction. Governments should therefore apply stricter obligations to the highest consequence systems while allowing structured, time-limited exceptions for lower tier environments where compensating controls are demonstrably effective.

This approach is already emerging internationally. Elements of the logic are visible in the Monetary Authority of Singapore's framework and CISA guidance, where unsupported technology exposure increasingly triggers enhanced governance, remediation and reporting expectations. The task now is to make those expectations explicit, enforceable and consistently applied across critical sectors.

Specifically: establish a two-tier floor. For Tier 1 and Tier 2 systems, prohibit unsupported technology in operational roles except under a formal exception approved at ministerial or equivalent authority level, time-limited to 12 months and subject to mandatory compensating controls. For Tier 2 systems, require documented risk acceptance renewed annually. Any exception at any tier that isn't renewed within the required period defaults to a mandatory replacement directive.

Lever 3: Fix procurement—embed life-cycle obligations in government contracts and guidance

Governments should intervene at acquisition, before unsupported technology becomes embedded operational debt. Systems that become unsupported because procurement frameworks never required life-cycle commitments represent a governance failure at purchase, not an operational failure years later. Governments hold direct leverage here, both as regulators defining procurement expectations and as major buyers shaping market behaviour through contracting requirements.

At a minimum, procurement frameworks should require vendors to disclose support timelines and EoS dates as a condition of contract, alongside minimum support-duration requirements calibrated to operational consequence. For higher consequence systems, transition planning should be built into acquisition from the outset rather than deferred until systems approach obsolescence. SBOMs should similarly become a standard condition of supply for critical systems and services.

These measures shift life-cycle risk from an unmanaged downstream problem to an explicit procurement and investment consideration. They ensure that unsupported technology exposure is identified, understood and priced during acquisition rather than emerging later as an unfunded operational liability.

The direction of travel is already visible. India's draft CEA regulations and the Philippines' Republic Act No. 12009 both point in this direction. In Australia, the Commonwealth Procurement Rules and standard contract templates provide a practical vehicle for embedding the same expectations at scale, with government purchasing power shaping supplier behaviour across the market.

Specifically: amend standard government ICT contract templates to include three mandatory fields at the point of procurement: (1) a 'life-cycle support declaration' requiring vendors to disclose EoS dates for all supplied components at the time of contracting; (2) a

minimum support-duration commitment scaled to a consequence tier, with contractual remedies for early withdrawal of support; and (3) a transition-planning clause for higher consequence systems requiring a vendor-supported migration pathway as a condition of the contract. Require a bills of materials covering software, hardware and cryptographic components as a standard condition of supply for all critical systems and services. These requirements should apply across government procurement frameworks and flow through to subcontractors in high-consequence supply chains.

Lever 4: Require accountability—assurance, audit, incident linkage and funded transition plans

Requirements without accountability will produce uneven implementation and growing exception backlogs. Governments should integrate life-cycle exposure directly into existing assurance, audit and incident-reporting frameworks rather than treating unsupported technology as a separate compliance regime. Incident reporting is a critical part of this process. Where a cyber incident is materially enabled by unsupported, unpatched or EoL systems, that exposure should be explicitly identified in reporting and post-incident review processes. This creates a direct feedback loop between operational incidents, governance oversight and investment decisions.

High-consequence EoL exposure should also trigger mandatory transition planning. Operators maintaining unsupported systems in critical operational roles should be required to demonstrate a funded and time-bound pathway for remediation, replacement or managed mitigation. An operator carrying the highest tiers of legacy exposure without a credible transition plan hasn't met its governance obligation. Governments also need to recognise that transition itself can introduce operational risk. Regulatory frameworks should therefore allow managed disruption and staged migration approaches where necessary to support safe modernisation. Without that flexibility, operators will continue to defer replacement decisions in favour of maintaining ageing systems that appear operationally safer in the short term.

Lever 5: Enable transition—incentives, guidance and coordinated programs

Regulatory pressure without credible transition pathways produces avoidance rather than action. Operators often defer replacement not because they dismiss the risk, but because practical conditions—funding, capability, continuity, coordination—aren't sufficiently aligned. Governments therefore need to move beyond compliance and create structured pathways that make transition achievable.

Transition support shouldn't rely solely on subsidies or one-off procurement interventions. Governments should instead treat life-cycle uplift as a participation model, in which funding access, accreditation pathways and market opportunity are increasingly tied to demonstrable progress against life-cycle-governance expectations. Australia's sovereign capability and security-uplift programs provide a useful reference point: they create pathways towards defined security baselines while linking modernisation effort to accreditation, eligibility and commercial opportunity. The same logic can be applied to life-cycle governance.

The objective is to make transition credible at scale. Where operators see a viable pathway— supported by funding, capability and coordination—compliance becomes achievable rather than aspirational. Where they can't, requirements will continue to be deferred, exceptions will accumulate, and legacy exposure will persist regardless of regulatory intent.

Specifically: establish two distinct enablement mechanisms, each addressing a different constraint. First, a co-funding facility for high-consequence legacy replacement—targeted at Tier 1 and Tier 2 operators where replacement barriers are genuine and transition risk is nationally significant, with access conditional on a board-approved transition plan and regulatory notification of exposure. Second, a coordinated cohort program for operators managing similar legacy technology across a sector—allowing shared procurement, common tooling, aligned assurance activity and coordinated vendor negotiation, reducing duplication and improving market leverage.

Implementation road map

The five levers above aren't sequential in a strict sense—visibility work can begin immediately, while procurement reform and accountability mechanisms are progressed in parallel. The sequencing should also build in an explicit offramp: where assets can't realistically be replaced in the near term, that should be acknowledged formally rather than treated as failure, and be paired with compensating controls (isolation, segmentation, enhanced protection), a costing of those controls, a periodic requirement to weigh them against the cost of replacement, and a documented transition and reporting plan. Table 2 provides proposed sequencing across three time-horizons.

The critical governance test occurs at the end of Phase 2. At that point, every Tier 1 and Tier 2 operator either has a formally approved exception on file or has commenced a funded transition. Passive continuation—systems persisting simply because no decision has been forced—is no longer possible.

Table 2: Phased implementation

Phase	Priority actions	Decision gate
1. Mandate	Amend ICT contract templates to include life-cycle support declaration, minimum support-duration commitments and SBOMs requirements. Designate high-consequence operators required to submit annual life-cycle exposure returns. Issue CISO attestation requirement.	Procurement amendments in force. Operator notification issued. First life-cycle exposure returns scheduled.
2. Enforce	Activate Tier 1 and Tier 2 exception regime: all unsupported systems require formal approval, compensating controls and a time-limited exception. Amend incident reporting obligations to require legacy-exposure field. Publish hardening guidance library for highest prevalence legacy platforms.	No Tier 1 or Tier 2 system continues without a formal exception on file. First incident reports with legacy-exposure field received and reviewed.
3. Enable	Stand up co-funding options for high-consequence legacy replacement. Launch first sector cohort programs. Conduct first independent audit of life-cycle exposure return submissions. Review exception register for patterns requiring regulatory escalation or legislative change.	Cohort programs operational. Audit findings reported to relevant minister. Legislative amendments identified and tabled where required.

Commercial levers: what enterprises need to do

What follows is the implementation architecture that operators (as distinct from governments) need to meet the regulatory requirements set out in the preceding section—and to capture the transition value that a purely compliance-driven response will leave unrealised. This achieves two objectives concurrently: it (1) establishes a clear view of system support status across the enterprise, and (2) creates persistent mechanisms for enterprises to act on that information in a timely and prioritised way.

'Legacy Five' for enterprises

Lever 1: Make it visible—asset inventory with support status and consequence tier

Establish a life-cycle register. This isn't a replacement for a general asset inventory, which remains essential, but a structured extension of it. It records each system's support status, EoS date, vendor commitment, consequence tier and dependency relationships. Without this layer, risk can't be prioritised, ownership can't be assigned and transition can't be planned.

This is where most operators underestimate the work. Self-reported inventories without verification carry limited governance value. A credible life-cycle register will surface more unsupported exposure than most boards have been told exists. That's its purpose.

The case studies in this report highlight three cautionary experiences in this regard: South Korea maintains sectoral asset tracking but lacks a cross-sector life-cycle register. India incorporates life-cycle concepts in some sectoral guidance but doesn't require system-wide reporting of support status. The Philippines has asset-protection frameworks but no explicit life-cycle-visibility mandate. In each case, visibility exists in fragments, not as a decision-making foundation.

This points to a consistent requirement: visibility should be explicit, standardised and independently assured. Emerging approaches such as OpenEoX, developed under the OASIS Open framework, provide a basis standardised exchange of EoL and EoS information between vendors and operators.[114] While still maturing, they offer a practical path to improving data quality and reducing reliance on manual reporting.

Lever 2: Set the floor—consequence-based prioritisation

Ensure governance is determined by consequence, not by system age or maintenance cost. The practical danger to avoid is uniformity. Frameworks that treat all legacy systems as equivalent tend to overreact in lower consequence cases and underreact where the stakes are highest. Make distinctions based on four consequence tiers:

- Tier 1—Strategic dependence: failure creates strategic dependence, compromises national or sector-wide functions, or locks the operator into constrained vendor or replacement pathways. Requires mandatory retirement or formal exception with ministerial or board accountability.
- Tier 2—Safety or essential-service consequence: failure interrupts essential services or carries physical safety effects. Requires formal replace-or-mitigate determination with board sign-off and regulatory notification above a defined threshold.
- Tier 3—Service disruption: Failure interrupts business functions at scale but without safety consequences. Requires documented risk acceptance and compensating controls on a fixed review cycle.
- Tier 4—Inconvenience: Failure produces localised disruption. Managed through standard risk-management processes.

Consequence-based prioritisation sharpens investment discipline. Systems with the highest operational or strategic consequences are typically those where modernisation generates the greatest return through improved data, stronger procurement leverage and reduced dependency. A board approving a transition program structured around consequence tiers is therefore approving both a risk-reduction sequence and an investment sequence.

Lever 3: Fix procurement—lock in life-cycle commitments at acquisition

Embed life-cycle obligations in procurement from the outset. Specifically:

- vendor disclosure of support timelines and EoS dates
- minimum support-duration requirements scaled to consequence tier
- transition-planning requirements for higher consequence systems
- SBOMs as a condition of supply.

Systems rarely become unsupported unexpectedly. More often, procurement frameworks fail to require life-cycle commitments at acquisition. Vendors aren't required to declare support timelines, contracts don't specify minimum support periods, and transition pathways are neither costed nor planned. While this component doesn't resolve existing exposure, it prevents new exposure from accumulating. At scale, it also reshapes incentives, favouring vendors that invest in long-term supportability.

Procurement decisions should account for both life-cycle commitments and the consequences of failing to meet them. Examples from this report point in this direction. India's draft CEA regulations require vendors to provide security updates throughout the product life cycle.[115] Section 13 of the Philippines' Republic Act No. 12009 introduces life-cycle cost analysis into government procurement.[116] The challenge is extending this logic into a consistent, cross-sector standard.

Lever 4: Require accountability—formal determination: documented, owned, time-limited, escalation-ready

Require a decision, not passive continuation. Effective governance should require a formal determination that's documented, owned, time-limited and subject to escalation. Any replace-or-mitigate determination should require:

- a documented decision: replace within a defined time frame or maintain with compensating controls
- named ownership: a specific individual or body accountable for the decision and its execution
- a time limit: the decision expires and must be renewed; it doesn't continue indefinitely
- an escalation path: board sign-off for Tier 1 and Tier 2 systems; regulatory notification beyond a defined threshold; ministerial accountability for nationally significant systems.

Where replacement is feasible, a transition plan should be mandatory. Where it isn't, compensating controls such as segmentation, enhanced monitoring and restricted access should be formally documented, costed and reviewed on a fixed cycle.

The critical design feature isn't the content of the decision but the requirement that a decision be made at all. A board that has approved and is actively reviewing a replace-or-mitigate register has materially reduced its exposure. A board that hasn't is accepting risk without structured oversight.

Lever 5: Enable transition—sequenced replacement generating returns, not only reducing risk

Ensure that any transition pathway performs three functions concurrently:

1. Preserve operational continuity: parallel-running parameters, rollback mechanisms and defined load-transfer thresholds are governance requirements. A transition that introduces more risk than the system it replaces creates a second exposure.
2. Align capital planning with life-cycle timelines: legacy exposure should be recognised as a liability within long-term financial and asset-management frameworks, with a funded resolution pathway.
3. Capture value, not only close exposure: modernisation should unlock data, improve operational capability and reduce dependency. Boards should require value-capture metrics alongside risk-reduction metrics.

Implementation road map

The proposed governance framework doesn't require all five components to be implemented simultaneously. The practical entry point is the first phase: establish what exists. Everything else follows from that. The three-phase sequence shown in Table 3 provides a board-approvable structure, with explicit decision gates between phases.

Phase 1 can be approved and commenced immediately at a defined cost, with a specific deliverable: the life-cycle register. That register becomes the basis for Phase 2 triage, which in turn defines the scope and sequencing of the Phase 3 program. No phase requires the next to be fully designed in advance. Each produces a concrete output and a decision point for the board.

The first meaningful governance test occurs at the end of Phase 2. At that point, every Tier 1 and Tier 2 system either has a funded transition plan or a formally approved exception with compensating controls, named ownership and a defined review cycle. Passive continuation is no longer possible; systems persist only through explicit decision.

Enabling conditions

Four conditions determine whether a life-cycle-governance framework is workable in practice, rather than remaining a conceptual model.

- *Common vocabulary is foundational*

  Without shared definitions of EoL, EoS and end-of-security-support, visibility is inconsistent, comparison is unreliable, and procurement requirements can't be enforced. The adoption of a standardised taxonomy, for example, drawing on OpenEoX or a comparable regional approach, is a low-cost, high-leverage intervention that improves data quality across operators and suppliers.

- *Transparent reporting creates accountability*

  Frameworks without visibility produce compliance theatre rather than measurable risk reduction. Reporting life-cycle exposure, remediation progress and formal exception approvals to boards and, where relevant, regulators converts governance obligation into tracked outcomes. The act of reporting is itself a forcing function for decision-making.

- *Capability should match ambition*

  The gap between policy intent and operational outcome is often a capacity problem. Organisations that can't identify life-cycle exposure can't manage it. This applies equally to operators and regulators. Investment in life-cycle assessment capability, including automated asset discovery and triage, is a practical precondition for implementation.

- *Coordination amplifies impact*

  Life-cycle risk doesn't respect organisational or sectoral boundaries. Operators that align approaches across supply chains, sectors or jurisdictions reduce duplication and increase leverage over vendors. In the Indo-Pacific context, the early adoption of credible life-cycle-governance practices also carries commercial and strategic advantage, shaping supplier behaviour and influencing emerging standards.

Differentiated application

This report recognises that any single framework applied mechanically won't produce consistent outcomes across all organisations. Effective implementation needs to vary according to the dominant constraint, even where the core governance model remains constant. The case studies in this report point to three recurring operator patterns.

- *Policy-created lock-in*

  The central requirement of this environment is decision authority. Systems become unsupported not because support is unavailable, but because it isn't maintained. The priority is to treat supportability as an operational obligation rather than an administrative task. That requires dedicated transition budgets insulated from routine operational IT spending, explicit retirement mandates for the highest-consequence legacy cohorts and a coordinating authority able to force decisions across sectoral boundaries.

- *Procurement-driven unsupportability*

  The priority is alignment between procurement, cybersecurity and operational ownership. That means vendor life-cycle disclosure at procurement, support renewal treated as a governed operational obligation and clear accountability for supportability as a distinct risk category. One option is to have the actor responsible for a system's security also be responsible for maintaining its supportability.

- *Fragmentation at scale*

  The priority is coherence without false uniformity. Full standardisation is often neither feasible nor necessary. What's required is a minimum common architecture: shared definitions of EoL and EoS, a common consequence-tier model, cross-sector visibility reporting to a central authority, and a mechanism by which high-consequence sectors can pull forward national standards rather than wait for whole-of-system convergence. The core governance requirement is coordination design: sectoral progress needs to produce transferable standards, not isolated islands of maturity.

Table 3: Phased implementation

Phase	Time frame	Board deliverable	Decision gate
1. Implement	60–120 days	Life-cycle register: all systems, support status, consequence tier, dependency map.	Approve Phase 2 scope and budget. Refer highest consequence unsupported systems for immediate replace-or-mitigate determination.
2. Decide	90–180 days following Phase 1	Replace-or-mitigate determination for all Tier 1 and Tier 2 systems. Funded transition plans or formal exception approvals with compensating controls documented.	No Tier 1 system continues without a formal board decision. Escalation log established for regulatory notification threshold.
3. Execute	Rolling program; annual board review	Transition program underway for highest consequence cohort. Capital plan aligned to life-cycle timelines. Value-capture metrics reported alongside risk-reduction metrics.	Annual board sign-off on residual exception register. Procurement standards updated to embed life-cycle obligations for all new acquisitions.

Governance standard

Done well, the approach set out here does more than reduce exposure. It creates clearer accountability, stronger transition discipline and a credible basis for resilience, capability and value capture. Done badly, or not at all, it preserves the pattern that the case studies document: unsupported systems remain in service because no framework is strong enough to turn recognition into action.

The minimum standard is this: no unsupported system continues by default. Every system that operates beyond EoS does so because a named decision-maker, accountable to the board, has formally accepted that position—with documented controls, a funded exit pathway and a defined review date. Ultimately, that's the difference between inherited exposure and managed risk.

Notes

1 Cisco Talos, *2025 year in review*, 2026, online.
2 NTT Ltd, *2020 global network insights report*, 2020, online.
3 National Audit Office (NAO), *Government cyber resilience*, UK Government, January 2025, online.
4 US Government Accountability Office (USGAO), *Information technology: agencies need to develop modernization plans for critical legacy systems*, GAO-19-471, US Government, June 2019, online.
5 Alastair MacGibbon, 'Mythos has business worried. The rest are lulled into false cybersecurity', *Australian Financial Review*, 4 May 2026, online.
6 Richard Gray, *Hedging our bets: a potential Japanese option for managing risk in the AUKUS Optimal Pathway*, ASPI, Canberra, 4 May 2026, online; Jeong Soo Kim, 'Japan's submarine industrial base and infrastructure—unique and stable', Center for International Maritime Security (CIMSEC), 15 July 2024, online.
7 Cybersecurity and Infrastructure Security Agency (CISA), 'PRC state-sponsored actors compromise and maintain persistent access to US critical infrastructure', Advisory AA24-038A, US Government, 7 February 2024, online.
8 Splunk Inc., Oxford Economics, *The hidden costs of downtime: the $400B problem facing the Global 2000*, 23 July 2024, online.
9 NAO, *Government cyber resilience*; London borough of Hackney, cyberattack recovery reporting (2021–2022) and ICO reprimand, July 2024.
10 OpenEoX Technical Committee, *A standardized framework for managing end of life and other product lifecycle information*, OASIS technical report, 24 April 2025.
11 Forrester, 'Asia Pacific tech market forecast, 2024 to 2027', 2024; Forrester, 'Predictions 2025', *Technology and Security*, October 2024.
12 NTT Ltd, *2020 global network insights report*.
13 NTT Ltd, *2020 global network insights report*.
14 NAO, *Government cyber resilience*.
15 USGAO, *Information technology: agencies need to develop modernization plans for critical legacy systems*.
16 'Application modernization services market', *Markets and Markets*, 2025.
17 Cisco Talos, *2025 year in review*.
18 *The changing cyber threat landscape: Asia–Pacific region*, vol. 3, CYFIRMA, 26 September 2024, online.
19 European Union Agency for Cybersecurity (ENISA), *ENISA threat landscape 2024*, European Union, 19 September 2024, online.
20 National Cyber Security Centre, *Annual review 2024*, UK Government, November 2024.
21 'M-Trends 2025', Google Threat Intelligence Group (GTIG), April 2025; '2024 Zero-day exploitation analysis', GTIG, April 2025.
22 CISA, 'PRC state-sponsored actors compromise and maintain persistent access to US critical infrastructure'.
23 Cisco Talos, *2025 year in review*.
24 'The shadow campaigns: uncovering global espionage', *Unit 42*, Palo Alto Networks, 5 February 2026, online.
25 GTIG, 'Exposing the undercurrent: disrupting the GRIDTIDE global cyber espionage campaign', 26 February 2026, online.
26 US Office of Management and Budget, federal IT spending data, IT Dashboard (itdashboard.gov), FY2023 figures; see also USGAO, 'Agencies need to continue addressing critical legacy systems', GAO-23-106821, June 2023.
27 WPI Strategy, 'Update critical: why "end of life" technology is a cybersecurity risk for critical national infrastructure', November 2025, 20.
28 WPI Strategy, 'Update critical: why "end of life" technology is a cybersecurity risk for critical national infrastructure'.
29 'Cost of a data breach report 2024', IBM, July 2024; 'Global Digital Trust Insights 2024: Asia Pacific', PwC, October 2023.
30 Splunk Inc. & Oxford Economics, *The hidden costs of downtime*.
31 Splunk Inc. & Oxford Economics, *The hidden costs of downtime*.
32 NAO, *Government cyber resilience*; London borough of Hackney, cyberattack recovery reporting (2021–2022) and ICO reprimand (July 2024).
33 'OT cybersecurity breach disrupts operations at the Port of Nagoya, Japan', Dragos, July 2023.
34 'AI Readiness Index 2025: realizing the value of AI', Cisco, October 2025.
35 Australian Signals Directorate (ASD), Australian Cyber Security Centre, *Planning for post-quantum cryptography*, Australian Government, July 2022, updated August 2023.
36 'From naptime to big sleep: using large language models to catch vulnerabilities in real-world code', Google Project Zero and DeepMind, November 2024.
37 'PoC code in 15 minutes? AI turbocharges exploitation', *Dark Reading*, 2 September 2025.
38 ENISA, *Foresight cybersecurity threats for 2030*, March 2024 update, online.
39 'Project Glasswing: securing critical software for the AI era', Anthropic, 7 April 2026, online.
40 'Project Glasswing: an initial update', Anthropic, 22 May 2026, online.
41 Department of Homeland Security, 'Resilient Positioning, Navigation, and Timing (PNT) Conformance Framework', US Government, 2020, updated 2023.
42 'Quantum-readiness: migration to post-quantum cryptography', joint cybersecurity information sheet, CISA, National Security Agency, National Institute of Standards and Technology, August 2023; 'Transitioning to a quantum-secure economy', World Economic Forum with Deloitte, September 2022.
43 South Korea: National Intelligence Service, 'Post-Quantum Cryptography Master Plan', 2023; India: Department of Science and Technology, 'Implementation of quantum safe ecosystem in India', taskforce report, 4 February 2026.
44 Department of Information and Communications Technology, *National Cybersecurity Plan 2023–2028*, Philippines Government, 2024, online; Department of Space and Technology, Philippine Council for Emerging Technology Research and Development, *Quantum Technology Roadmap*, Philippines Government, February 2025.
45 Department of Science and Technology, *Implementation of Quantum Safe Ecosystem in India*, Indian Government, February 2026, 13.
46 CISA, FBI, EPA, *Incident response guide: water and wastewater systems sector*, US Government, January 2024; 'Top cyber actions for securing water systems', fact sheet, CISA, March 2024.
47 Jo He-rim, 'SK Telecom hit with record privacy fine after massive data leak', *Korea Herald*, 28 August 2025, online.
48 Lee Jung-joo, 'Recovery from data center fire to take longer than expected', *Korea Herald*, 29 September 2025, online; Bruno Ferreira, 'South Korean Government learns the importance of backups the hard way after catastrophic fire—858 terabytes of data goes up in magic smoke', *Tom's Hardware*, 10 October 2025, online.
49 Nick Merrill, *The South Korean digital paradox: how South Korea's internet development model creates unique cybersecurity vulnerabilities*, UC Berkeley Center for Long-Term Cybersecurity, April 2025, online.
50 Gelo Gonzales, 'PhilHealth hit by ransomware—report', *Rappler*, 22 September 2023, online.
51 Gelo Gonzales, 'Privacy commission to probe role of BDO's 10-year-old system in hack', *Rappler*, 22 December 2021, online.
52 'PhilHealth confirms antivirus subscription had expired amid ransomware attack', *GMA News Online*, 9 October 2023, online; National Privacy Commission, *In the matter of the alleged personal data breach of BDO Unibank, Inc.*, NPC SS 21-023, Philippines Government, 5 June 2024, online.
53 'Philippines mulls retaking control of national power grid amid concerns over Chinese firm's ownership', *South China Morning Post*, 17 May 2023, online; Harrison Prétat, Yasir Atalan, Gregory B Poling, Benjamin Jensen, 'Energy security and the US-Philippine Alliance', Center for Strategic and International Studies, 21 October 2024, online; Department of Information and Communications Technology (DICT), *National Digital Connectivity Plan*, Philippines Government, 2026, online; Elyssa Lopez, 'Philippines gives China-backed telecom firm green light to build cell towers on military bases', *South China Morning Post*, 14 September 2020, online.

54 'WannaCry ransomware cyber attack: Indian ATMs could be at high risk as most run on Windows XP', *Business Today*, 15 May 2017, online.

55 'AIIMS ransomware attack', Cyber Management Alliance, 5 July 2023, online; Rajeswari Pillai Rajagopalan, 'The AIIMS cyberattack reflects India's critical vulnerabilities', Observer Research Foundation, 3 December 2022, online.

56 Alexander Campbell, Vickram Singh, 'Lessons from the cyberattack on India's largest nuclear power plant', *Bulletin of the Atomic Scientists*, 14 November 2019, online; Melissa Robbins, 'Cyberattack hits Indian nuclear plant', *Arms Control Today*, December 2019, online.

57 Tejas Bharadwaj, 'Mapping India's cybersecurity administration in 2025', Carnegie Endowment for International Peace, 1 September 2025, online; Rajagopalan, 'The AIIMS cyberattack reflects India's critical vulnerabilities'; Priyanka Sharma, 'AIIMS computers, IT system not upgraded for at least three decades', *Tech Circle*, 5 December 2022, online.

58 Merrill, *The South Korean digital paradox: how South Korea's internet development model creates unique cybersecurity vulnerabilities*; Hun Myoung Park, 'The pitfalls of the certificate-based user authentication scheme on Korean public websites', *Journal of Cases on Information Technology*, 2024, 26(1), online.

59 'Galapagos regulation', editorial, *Korea Herald*, 25 March 2014, online.

60 Hyun-il Hwang, 'Data privacy laws in Korea's fintech landscape', *LawAsia*, 16 May 2025, online.

61 Government Procurement Policy Board, 'Resolution no. 06-2022: Approving the Guidelines on the Renewal of Regular and Recurring Services', Philippines Government, 2022, online; 'PhilHealth insists new procurement rules affected renewal of antivirus software', *Bilyonaryo*, 12 October 2023, online.

62 Aaron Raj, 'Could the Philippine Health Insurance Corporation cyberattack have been avoided?', *Tech Wire Asia*, 11 October 2023, online.

63 DICT, *National Cybersecurity Plan 2023–2028*, Philippines Government, February 2024, online; National Privacy Commission, 'NPC circular no. 2023-06: Security of personal data in the government and the private sector', Philippines Government, 1 December 2023, online; 'Circular no. 1213: Amendments to regulations on information technology risk management to implement section 6 of the Anti-Financial Account Scamming Act (AFASA)', Bangko Sentral ng Pilipinas, 10 June 2025, online.

64 *Republic Act No. 12009 (New Government Procurement Act)*, Philippines Congress, 24 July 2023, online; Government Procurement Policy Board, 'New Government Procurement Act or Republic Act No. 12009', Philippines Government, 20 July 2024, online.

65 *Act on the Protection of Information and Communications Infrastructure*, Act no. 6383, Republic of Korea, 26 January 2001 (as amended), online.

66 Jinyong Park, Tae-Sung Kim, 'A framework to improve the compliance guideline for critical ICT infrastructure security', *Journal of Open Innovation: Technology, Market, and Complexity*, 2025, 11(2):100547, online.

67 Natasha Wood, 'South Korea's 2024 Cyber Strategy: a primer', Center for Strategic and International Studies, 2 August 2024, online.

68 Martin Dale Bolima, 'Korea's BAI audit exposes digital weak points as simulations breach all tested public systems', *Cybersecurity Asia*, 9 February 2026, online.

69 Jae-Gu Song, Cheol-Kwon Lee, Jung-Woon Lee, 'A study on the RS-015 based checklist and acceptance criteria for support consistent review in regulation', paper presented at Transactions of the Korean Nuclear Society Autumn meeting, Changwon, Korea, 20–21 October 2022, online; Yun-Hyuk Choi, Sang-Jin Lee, 'A study on the implementation of technical security control for critical digital asset of nuclear facilities', *Journal of the Korea Institute of Information Security and Cryptology*, 2019, 29(4):877.

70 National Intelligence Service, *Regulations on Cybersecurity Services*, presidential decree no. 34287, South Korean Government, 5 March 2024, art. 9(5), 13, online; Financial Services Commission, *Electronic Financial Transactions Act*, South Korean Government, art. 21-3, online; Ministry of Science and ICT, *Act on the Protection of Information and Communications Infrastructure*, Act no. 20068, South Korean Government, 23 January 2024, art. 9, online; Korea Communications Commission, *Act on Promotion of Information and Communications Network Utilization and Information Protection*, Act no. 21066, South Korean Government, 1 October 2025, online; Cecille Suerte Felipe, Elijah Felice Rosales, 'Medusa hackers release stolen PhilHealth data', *The Philippine Star*, 6 October 2023, online.

71 DICT, *National Cybersecurity Plan 2023–2028*, Philippines Government, February 2024, s. 1.7, online; Office of the President of the Philippines, 'Executive order no. 58: Adopting the National Cybersecurity Plan 2023–2028, and directing the implementation thereof', 4 April 2024, online.

72 National Privacy Commission, 'NPC circular no. 2023-06: Security of Personal data in the government and the private sector', Philippines Government, 1 December 2023 (effective 30 March 2024), s. 30(3), online; Bangko Sentral ng Pilipinas, 'Circular no. 982: Enhanced guidelines on information security management', 2 November 2017, appendix 75b/Q-59b, ss. 3.2.2, online.

73 Government Procurement Policy Board, 'Implementing rules and regulations of Republic Act no. 12009 (New Government Procurement Act)', Philippines Government, 2025, ss. 13, 35(m), online; Government Procurement Policy Board, 'Lifecycle assessment (LCA) and lifecycle cost analysis (LCCA) under RA 12009', Philippines Government, 5 August 2024, online.

74 Indian Computer Emergency Response Team (CERT-In), *Comprehensive cyber security audit policy guidelines* (CISG-2025-02), version 1.0, 25 July 2025, ss. 6(xxvi), 9.1–9.2, 9.6.

75 Central Electricity Authority, 'CEA (Cyber Security in Power Sector) Guidelines 2021', Ministry of Power, Indian Government, 2021, online; CERT-In, *Technical guidelines on software bill of materials (SBOM), quantum BOM (QBOM), cryptography BOM (CBOM), AI BOM (AIBOM) and hardware BOM (HBOM)*, version 2.0, 9 July 2025, online.

76 Ministry of Power, *Amendment in CEA (Cyber Security in Power Sector) Guidelines 2021* (no. CEA-CH-13-12/13/2021-IT Division), Indian Government, September 2022, online.

77 Draft Central Electricity Authority (Cyber Security in Power Sector) Regulations 2025 (released for public consultation; comments closed 7 November 2025; not yet notified in the official gazette and therefore not in force), regs 16(2), 16(4), 16(5), online.

78 *General Financial Rules 2017*, Indian Government, updated to 31 July 2024, rule 169, online; Department of Expenditure, Ministry of Finance, *Manual for procurement of goods*, 2nd ed., Indian Government, 2024, online.

79 National Critical Information Infrastructure Protection Centre, *Guidelines for critical information infrastructure protection*, version 2.0, Indian Government, 16 January 2015, IC1 (§7.1), IC6 (§7.6), OC6 (§8.6), OC7 (§8.7), PC9 (§6.9), online.

80 Ayyappan Rajesh, 'A strategic framework for mitigating electronic hardware-related national security risks in India', ORF occasional paper no. 530, Observer Research Foundation, 16 March 2026, online.

81 Reserve Bank of India, *Control measures for ATMs—timeline for compliance* (RBI/2017-18/206, DBS(CO).CSITE/BC.5/31.01.015/2017-18), 21 June 2018, online.

82 Ministry of Science and ICT, National IT Industry Promotion Agency, *2026*년 공공부문 SW사업 수요예보(확정) 발표자료 [Announcement materials on the 2026 Public Sector SW Project demand forecast (finalized)], South Korean Government, 31 March 2026, slide 2-2.

83 Lee Sang-eun, Kang Kyung-min, Choi Hae-ryeon, '정부, 레거시 유지보수 예산 삭감해 신규 사업 재원 마련 [Government cutting legacy maintenance budgets to fund new initiatives]', *Hankyung*, 20 November 2023, online; Kang Ju-ri, Yoo Seung-hyeok, '[Exclusive] Next year's E-government support budget cut by 74% … overseas cooperation projects increased', *Seoul Shinmun*, 23 November 2023, online.

84 National Assembly Budget Office, *2026*년도 예산안 심의결과 *[FY2026 Budget Bill Deliberation Results]*, South Korean Government, 17 December 2025, online.

85 Lee Jung-joo, 'Recovery from data center fire to take longer than expected', *Korea Herald*, 29 September 2025, online.

86 Yoon Sung-min, 'State audit board finds beleaguered NIRS failed to manage old equipment properly', *Korea JoongAng Daily*, 29 September 2025, online.

87 Eunju Hong, 'Digital government innovation, preparing for hyperscale AI and cloud', Samsung SDS, 16 December 2024, online; OECD Digital Government Studies, *Digital Government Review of Korea: Harnessing Digital and Data to Transform Government*, Organisation for Economic Co-operation and Development (OECD), OECD Publishing, Paris, 2025, 51, online.

88 National Telecommunications Commission, 'Memorandum circular no. 003-09-2025: Decommissioning of third generation (3G) and second generation (2G) mobile network services', Philippines Government, 2025, online.

89 Maria Bernadette Romero, 'COA flags LTO for LTMS delays, P438M fees', *The Manila Tribune*, 28 April 2025, online; 'COA orders LTO, foreign contractor to address audit findings by February 28', *Philstar.com*, 22 February 2025, online.

90 Peter Tabingo, 'Only 10 of 19 PAGASA radars operational—2024 audit', *Malaya Business Insight*, 29 December 2025, online.

91 Pauie Benavidez, 'Why PAGASA never gets the budget it really needs', *Catanduanes Tribune*, 7 November 2024, online; 'COA: PAGASA should speed up "crucial" doppler radar repairs', *Rappler*, 29 December 2025, online.

92 Benavidez, 'Why PAGASA never gets the budget it really needs'.

93 Philippine Atmospheric, Geophysical and Astronomical Services Administration (PAGASA), 'Tropical Cyclone Information', Philippines Government, online.

94 Draft Central Electricity Authority (Cyber Security in Power Sector) Regulations 2025 (released for public consultation; comments closed 7 November 2025; not yet notified in the official gazette and therefore not in force), regs 16(2), 16(4), 16(5), online.

95 Securities and Exchange Board of India, 'Circular no. SEBI/HO/ITD-1/ITD_CSC_EXT/P/CIR/2024/113: Cybersecurity and Cyber Resilience Framework (CSCRF) for SEBI regulated entities', 20 August 2024, online; Reserve Bank of India, 'Master direction on information technology governance, risk, controls and assurance practices', RBI/2023-24/107, DoS.CO.CSITEG/SEC.7/31.01.015/2023-24), 7 November 2023, online.

96 CERT-In, *Technical guidelines on software bill of materials (SBOM), quantum BOM (QBOM), cryptography BOM (CBOM), AI BOM (AIBOM) and hardware BOM (HBOM)*.

97 CERT-In, *Technical guidelines on SBOM, QBOM, CBOM, AIBOM and HBOM*, Table 5, s. 4.2(11), Table 6.

98 *General Financial Rules 2017* (updated to 31 July 2024), Indian Government, online; Department of Expenditure, Ministry of Finance, *Manual for procurement of goods*, 2nd ed., Indian Government, 2024, online; Ministry of Electronics and Information Technology, 'Public Procurement (Preference to Make in India) Order 2019 for Cyber Security Products', file no. 1(10)/2017-CLES, Indian Government, 2019, online; Ministry of Electronics and Information Technology, *India Enterprise Architecture (IndEA) Framework*, version 1.0, Indian Government, October 2018, ss. 6.9, 8.6, online.

99 Department of Home Affairs, *Protective Security Policy Framework*, July 2025, ch. 13, 'Technology lifecycle management', Australian Government, online; Department of Home Affairs, 'PSPF direction 002-2024: Technology asset stocktake', July 2024, online.

100 ASD, *Information security manual*, Australian Government, March 2026, online; ASD, *Essential Eight maturity model and ISM mapping*, Australian Government, October 2024, online.

101 ASD, *Essential Eight maturity model*.

102 Department of Home Affairs, 'Guidance for the Critical Infrastructure Risk Management Program', Australian Government, March 2025, online; Department of Home Affairs, 'Critical Infrastructure Risk Management Program: factsheet', Australian Government, April 2025, online; *Security of Critical Infrastructure Act 2018* (Cth).

103 Security of Critical Infrastructure Legislation Amendment (Enhanced Critical Infrastructure Risk Management Program) Rules 2026 (Cth) (F2026L00701), registered 9 June 2026, amending the Security of Critical Infrastructure (Critical Infrastructure Risk Management Program) Rules (LIN 23/006) 2023 (Cth), ss 4A, 8A, online.

104 ASD, 'Managing the risks of legacy ICT: executive guidance', Australian Government, April 2024, online; ASD, 'Managing the risks of legacy IT: practitioner guidance', Australian Government, June 2024, online; ASD, *Annual cyber threat report 2024–25*, Australian Government, October 2025, online.

105 Australian Trade and Investment Commission, 'Australia's strategy to become a global cyber leader by 2030', Australian Government, 15 January 2024, online.

106 Digital Transformation Agency, *Australian Public Service (APS) digital workforce insights report*, Australian Government, November 2025, online.

107 ASD, 'The Commonwealth cyber security posture in 2025', Australian Government, November 2025, online.

108 ASD, 'The Commonwealth cyber security posture in 2025'.

109 Mandala Partners, *Unlocking the productivity dividend of digital government*, September 2025, online; ASD, 'The Commonwealth cyber security posture in 2024'.

110 Australian National Audit Office (ANAO), *Defence's management of ICT systems security authorisations*, Auditor-General report no. 2 of 2024–25, Australian Government, September 2024, online.

111 Defence Department, *Defence Digital Strategy and Roadmap*, Australian Government, August 2024, online; ANAO, *Defence's management of ICT systems security authorisations*.

112 ASD, *CI Fortify: guidance for Australian critical infrastructure service continuity and resilience*, Australian Government, October 2025, online.

113 'Building a world leading cyber security nation', *auDA*, 29 February 2024, online.

114 OpenEoX Technical Committee, *A standardized framework for managing end of life and other product lifecycle information*.

115 Central Electricity Authority, 'Draft Cyber Security in Power Sector Regulations', Indian Government, October 2025.

116 Republic Act No. 12009, 'An Act revising Republic Act No. 9184, otherwise known as the "Government Procurement Reform Act", and for other purposes', Philippines Government, 2024.

Acronyms and abbreviations

AI	artificial intelligence
ANAO	Australian National Audit Office
ASD	Australian Signals Directorate
CEA	Central Electricity Authority (India)
CERT-In	Computer Emergency Response Team – India
CIRMP	critical infrastructure risk management program
CISA	Cybersecurity and Infrastructure Security Agency (United States)
CISO	Chief Information Security Officer
COA	commission of audit
ENISA	European Union Agency for Cybersecurity
EoL	end of life
EoS	end of support
GTIG	Google Threat Intelligence Group
ICT	information and communications technology
ISM	Information security manual
IT	information technology
NIRS	National Information Resources Service (South Korea)
OECD	Organisation for Economic Co-operation and Development
OEM	original equipment manufacturers
OT	operational technology
PSPF	Protective Security Policy Framework
SBOMs	software bills of materials
SOCI Act	Security of Critical Infrastructure Act 2018
UK	United Kingdom


ASPI
AUSTRALIAN
STRATEGIC
POLICY
INSTITUTE
25
YEARS
2001–2026