Lynqra

LYNQRA RESEARCH · NOVEMBER 2026

Responsible AI in 2026

Singapore ranks second in the world for people using AI. Fewer than 4 in 100 firms with 10 or more staff have built it into their core work. This report looks at that gap, what it costs, and how to adopt AI without getting hurt.

By Mark Tan, Founder and CEO, Lynqra

Data current to 24 November 2026 · 16 chapters · 192 sources
lynqra.com/research/responsible-ai-2026

Contents

The short version
Who this is for, and how to read it
01. The adoption mask
02. The gap that can't become normal
03. Proving it paid off
04. Is your organisation ready?
05. When most of the team is older
06. Train before you buy
07. Shadow AI and data leaks
08. Guardrails and safeguards
09. Human in the loop
10. Who's liable when AI gets it wrong
11. Auto top-up: the bill with no ceiling
12. Infrastructure: plan for the day it breaks
13. Local LLMs or cloud?
14. What attackers are doing with AI
15. What government is doing
16. Jobs: retrain before you replace
A responsible AI checklist for 2026
Method and limits
Sources

The short version

Singapore uses a lot of AI. Few of its firms have built it into how they work, and most of the risks in this report come from that gap. Ten findings, each with its source:

01 **Singapore's famous AI numbers count people, not businesses.** 64.3% of people in Singapore aged 15 to 64 used a generative AI product in the second quarter of 2026, second in the world after the UAE[1]. Among SMEs, 23.4% count as AI adopters, and the official survey counts any use of an AI tool, chat assistants included, as adoption[2, 3].

02 **Depth is rare.** 89.6% of AI-adopting firms use off-the-shelf tools with no customisation, and 6.6% say AI is fully deployed across the company[2]. Among firms with 10 or more staff, 71.5% had not started at all, and fewer than 4 in 100 were integrating AI into core processes[4].

03 **The payoff hasn't arrived yet.** Singapore firms that began using AI grew revenue and headcount, but showed no significant gain in profit or value added per worker within four years[5].

04 **Size and sector split the adopters from the rest, but firm age doesn't.** Large firms are 39 percentage points more likely to use AI than micro firms in the same sector, and firm age has no significant link[5]. Construction sits at 7.6%[2].

05 **Staff are ahead of their managers, and the rules aren't holding.** In Singapore, 69% of employees who use AI at work have relied on its output without checking it, and 55% have used it in ways that breach policy[6]. Only 36.7% of workers had any AI training in the past year[2].

06 **Banning AI doesn't stop leaks.** Worldwide, 48% of employees who use AI at work have put sensitive company information into public AI tools. Where AI was banned, the figure was 67%[7].

07 **Controls trail use.** Among organisations that had a security incident involving AI, 92% lacked proper access controls for it, and shadow AI was involved in 43% of the incidents studied[8]. 39% of risk and audit professionals don't know whether their organisation has a documented way to shut an AI system down[9].

08 **Human review fails quietly.** Very experienced radiologists' accuracy fell from 82% to 45.5% when they were shown a wrong AI suggestion[10].

09 **AI bills can run without a ceiling.** OpenAI switches API auto recharge on by default[11]. A three-person startup that normally spent about US$180 a month was billed US$82,314 in roughly 48 hours after its Gemini key was stolen[12].

10 **Impersonation fraud is rising.** Business email compromise losses in Singapore nearly tripled to S$57.3m in the first half of 2026. Police now list "digital manipulation" among the methods, though they don't say how many cases involved AI[13]. In March 2025 a finance director here sent US$499,000 after a Zoom call with deepfaked executives[14].

Who this is for, and how to read it

This report is for owners and managers of established Singapore firms, roughly 10 to 200 staff, who are either wondering where to start with AI or already have a plan and want to get it right. It covers sixteen topics in four parts: where Singapore actually stands, what to sort out before spending, how to build safely, and what is happening around you, from scams to regulation.

Each chapter puts the evidence first, with numbered sources you can check. It then closes with a section headed "My take", where I say what I think. I've kept the two apart on purpose. You can disagree with my take and still use the evidence.

A declaration before you start. I run Lynqra, a Singapore company that trains teams, plans AI adoption and builds AI systems for SMEs. I make money when firms adopt AI, so read my takes with that in mind. Where the evidence cut against what I'd like to be true, I've left it in.

Figures are current to 24 November 2026. "SME" follows the Singapore definition: operating revenue up to S$100m or up to 200 workers. How the research was done, what it couldn't find, and how AI was used in making it are covered in Method and limits at the end.

Mark Tan, Founder and CEO, Lynqra

PART 1 WHERE SINGAPORE REALLY STANDS

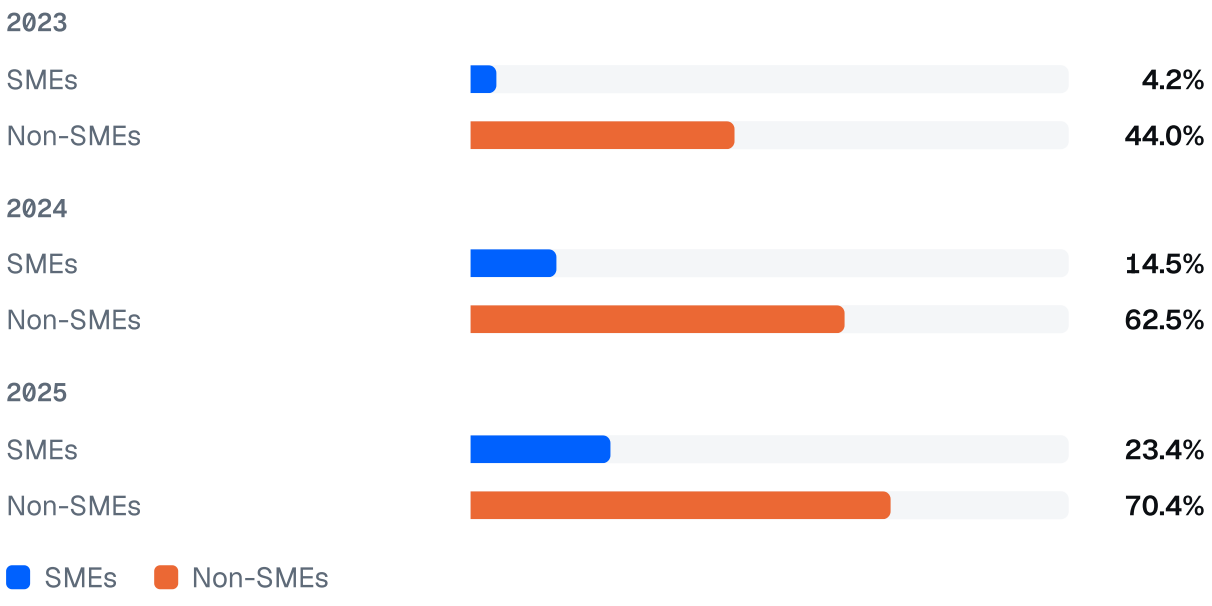
01 The adoption mask

Ask how Singapore is doing on AI and you'll usually hear a ranking. Microsoft's AI diffusion reports put the share of people in Singapore aged 15 to 64 who used a generative AI product at 58.6% in the first half of 2025, 60.9% in the second half, and 64.3% by June 2026. In each edition Microsoft names the UAE and Singapore as the two leaders, with the UAE ahead. The world figure in June 2026 was 18.8%[15, 1]. Anthropic's usage index, which compares a country's share of Claude conversations with its share of the working-age population, also put Singapore second, after Israel, in September 2025[16].

Both measure individuals using chat tools. Microsoft builds its figure from anonymised telemetry of people using AI products, and says plainly that no single metric is perfect[15]. Anthropic's index covers consumer Claude.ai traffic, not business systems[16]. Neither tells you whether a single company has changed how it works.

The business figures come from IMDA's annual survey of enterprises. SME adoption went from 4.2% in 2023 to 14.5% in 2024 and 23.4% in 2025. Non-SMEs went from 44.0% to 62.5% to 70.4%[3, 17, 2].

AI adoption by firm size, Singapore, % of enterprises



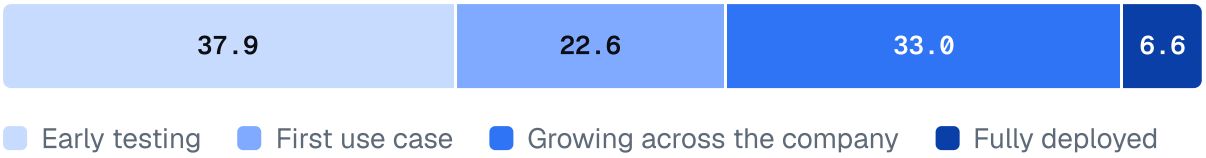
Source: IMDA, Singapore Digital Economy Reports 2024, 2025 and 2026[3, 17, 2].

Look at what counts. The survey defines AI as tools that let machines do tasks that normally need human intelligence, and lists content generation, customer engagement, analytics and process automation as examples. A firm counts as an adopter if it used such tools[3, 17]. IMDA's own 2025

report explains the 2024 jump: it came mainly from micro and small firms taking up off-the-shelf generative AI tools such as ChatGPT, Gemini, Midjourney and Copilot[17].

The 2026 report goes further, and this is the part that rarely makes the headlines. Of the firms counted as adopters, 89.6% use commercial off-the-shelf tools with no customisation. 23.0% use tools built in-house or by an outside developer for their own needs. Asked how far along they are, 37.9% were still testing, 22.6% were on their first use case, 33.0% were expanding across the company, and 6.6% had AI fully deployed[2].

How far AI-adopting Singapore firms have got, % of adopters



Source: IMDA, Singapore Digital Economy Report 2026, figure 14 (MDDI survey data)[2].

85.7% of surveyed workers used AI tools at work in 2026, up from 77.6% a year earlier. Content creation was the most common use. But only 25.7% frequently or always used AI to automate tasks, and 22.6% to work it into existing processes[2].

The Ministry of Manpower published the first official look at depth in April 2026. Among private firms with at least 10 employees, 71.5% had not adopted AI. Of the 28.5% that had started, MOM reports 7.4% planning, 6.0% piloting and 3.8% integrating AI into core processes. Its release doesn't make clear whether those are shares of all firms or of adopters. Either way, fewer than 4 in 100 firms were building AI into core work[4].

None of this makes Singapore a laggard. Among firms with 10 or more staff, IMDA puts Singapore at 34.6% against an OECD average of 20.3%, behind Denmark, Finland and Sweden[2]. The point is narrower. The number you hear depends on who is counted and what counts. Depending on the survey, between 23% and 92% of Singapore firms "use AI": 23.4% of SMEs (IMDA), 28.5% of firms with 10 or more staff (MOM), 59% of 381 firms surveyed by UOB, and 92% in a CPA Australia member survey[2, 4, 18, 19]. All of them are honest numbers. They answer different questions.

What firms report getting from AI fits the shallow picture. 87.6% of adopters cite productivity and process improvement. 24.5% say it contributed to higher revenue[2].

MY TAKE · MARK TAN

When you hear that Singapore is a world leader in AI, the number behind it is almost always about people typing into a chat box. I use those tools every day and they're useful. But a firm where a few staff paste emails into a free chatbot hasn't adopted AI in any sense that changes how the business runs. It has picked up a habit, usually with no rules attached.

There's a lot more to AI than ChatGPT. The value sits in workflows: a quotation drafted from the enquiry, or invoices matched to purchase orders before anyone opens them. Or a tender found, read and summarised while you sleep. That kind of adoption changes a P&L, and we barely measure it.

So I'd change the headline number. IMDA and MOM already collect depth: who is piloting, who has integrated AI into core processes, who uses only off-the-shelf tools. Lead with that. Report the share of firms that have at least one AI-run workflow in production, with an owner and a measured result. If that number is small, say so. A small honest number helps an owner more than a big flattering one, because it tells them where they really stand.

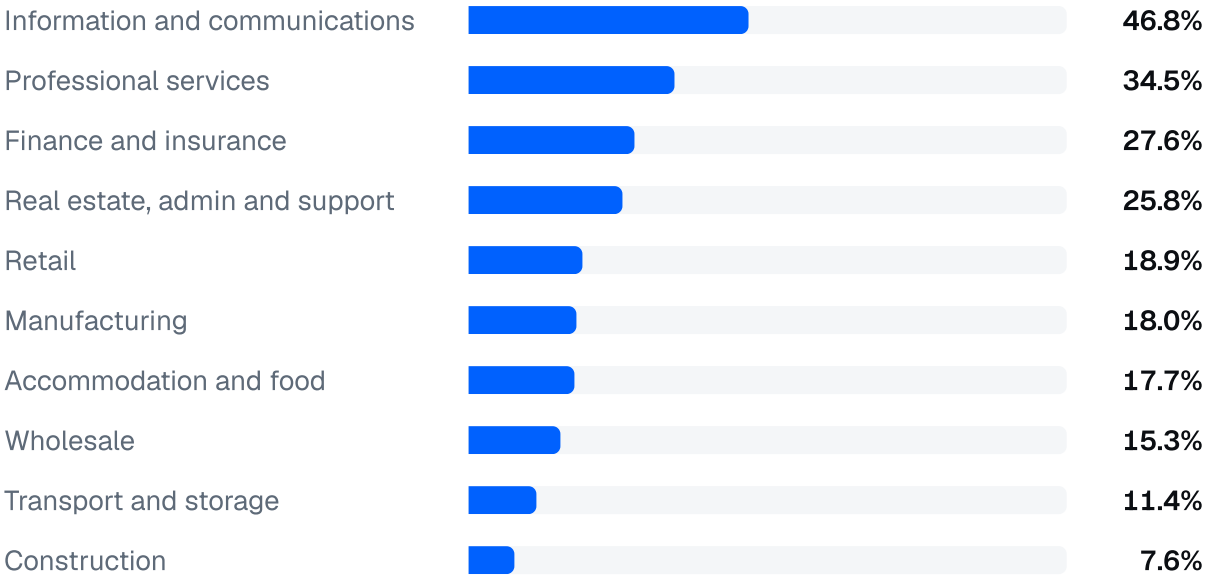
For your own firm, ask a harder question than "do we use AI?" Ask which piece of work AI now does from start to finish, who checks it, and what it saved last month. If nobody can answer that, you're using AI without having adopted it.

02 The gap that can't become normal

Singapore had about 369,500 SMEs in 2025, out of 371,000 enterprises[20]. They make up 99% of firms and employ more than 70% of the workforce[21]. Most are very small: of the 356,600 SMEs with fewer than 200 workers in 2024, 94.7% had fewer than 25. By our count from MOM's table, about 18,900 had between 25 and 199[22].

Adoption falls away with size. In 2026, 27.2% of firms with fewer than 200 employees used AI, against 54.8% of firms with 200 to 500 and 76.4% of firms with more than 500[23]. Sector matters as much.

AI adoption by sector, Singapore, 2025, % of all enterprises



Source: IMDA, Singapore Digital Economy Report 2026, figure 11[2]. Construction was 7.4% in 2024.

If "legacy" means old, the data doesn't back the idea that old firms lag. The trade ministry's 2026 study of firm-level AI use found that firm age had no statistically significant link with AI use once sector and size were accounted for, while large firms were 39 percentage points more likely to use AI than micro firms[5]. (Its measure is whether a firm advertised AI-related jobs, which captures capability-building better than casual use.) The lag is about size, sector and mindset.

The mindset part is measurable. Among firms not using AI, the top reason in 2025 was "do not see the need due to small business", given by 54.7%, barely changed from 57.4% in 2023. Budget (29.2%), IT not being ready (26.2%), skills (24.1%) and awareness (19.1%) came well behind. Only 12.1% of non-adopters planned to start within a year[2]. SBF's 2025 survey adds the management side: asked about barriers to digitalisation generally, 27% of firms cited limited management expertise in driving organisational change and 27% low workforce interest[24].

The gap is widening in other countries too. The OECD reports the AI adoption gap between large and small firms grew from 23 to 35 percentage points between 2023 and 2025[25]. In the US, use rose among firms with 20 or more employees and stayed flat for smaller ones[26]. In the EU, 17% of small enterprises used AI in 2025 against 55% of large ones[27]. MOM's April 2026 report says plainly that smaller firms risk falling further behind[4], and the Budget 2026 statement says Singapore must move beyond individual pilots and isolated experiments[28].

Can adoption be forced? The evidence on mandates isn't encouraging. Coinbase's CEO said in 2025 that he fired engineers who hadn't set up AI coding tools within a week and couldn't give a good reason, and called his own approach heavy-handed[29]. In a 2026 survey of 2,400 office workers and executives in the US and Europe, 29% of employees admitted sabotaging their company's AI strategy,

while 60% of the executives said their companies planned to lay off staff who can't or won't use AI[30]. That survey comes from an AI vendor and relies on self-reports, so treat it as a signal. The KPMG and University of Melbourne study in chapter 7 points the same way: where AI was banned outright, more staff put company data into public tools, not fewer[7].

MY TAKE · MARK TAN

"We're too small to need AI" is the most common answer from firms that haven't started, and it has hardly moved in two years. I understand why. A 30-person trading company doesn't need a data science team. But the work AI handles well (reading documents, chasing replies, checking one set of numbers against another, drafting the first version of something) is exactly the work that eats a small team's week. Small firms have more to gain per head.

I don't want a Singapore where the big firms run on AI while the SMEs that employ most of the country run on WhatsApp and spreadsheets. A gap like that compounds: firms that use AI attract people who know AI, and the distance grows every year. I don't want that to become normal.

But I don't think anyone can force it. Mandates from the top tend to produce workarounds. Government can fund and teach, but it can't make an owner care. What moves an owner is seeing one piece of their own work done properly, with their own data, and the numbers that come out of it. It's slower than a mandate, but it lasts.

And "legacy" isn't about the year a company was founded. The trade ministry found firm age doesn't predict AI use. In my view what predicts it is whether the person at the top has decided that understanding AI is part of their job. Chapter 5 comes back to that.

03 Proving it paid off

The best Singapore evidence on returns comes from the trade ministry's August 2026 study. Firms that started using AI saw revenue rise about 7% and employment about 4% in the year they began, and about 16% and 8% a year later. Profit and value added per worker showed no statistically significant gain within four years. The authors read that as a J-curve: costs and learning come first, and productivity gains, if they come, show up later[5]. (The study spots AI users by their AI job advertisements, so part of that headcount rise is the AI hiring itself.)

Self-reported benefits are softer. 87.6% of AI-adopting firms cite productivity, 44.0% cost reduction and 24.5% higher revenue[2]. In MOM's survey, 70.7% of AI-using firms reported better worker productivity[4]. SBF found 71% of digitalising firms track return on investment, mostly through time saved (63%) and cost efficiencies (62%)[24].

Globally, the most quoted number is that "95% of AI pilots fail", from MIT's NANDA initiative in 2025. The report actually says 95% of organisations were getting zero return and only 5% of integrated pilots were extracting millions in value. It rests on 52 interviews, 153 survey responses gathered at four conferences and a review of public announcements. It was not peer-reviewed, and the authors call their figures directionally accurate[31, 32]. A February 2026 podcast critique pointed out that the 5% refers to organisations reaching production, that most never piloted at all, and that the authors work on agent frameworks the report recommends[32]. It's included here because you'll meet it everywhere, and because it is weaker than its fame.

Firmer surveys point in a similar direction. BCG's 2026 study of 1,330 senior leaders classed 7.5% of companies as "future-built" and 41% as scaling. The future-built group reported 2.4 times the revenue growth and 2.8 times the EBITDA growth of laggards[33], which shows correlation, not cause. Gartner found in 2026 that only 22% of organisations with revenue above US$50m had scaled AI across several business units. 11% didn't know what their function had spent on AI the year before. High performers, who track returns and drop weak projects, reported positive returns on 81% of their AI initiatives[34]. Gartner had also forecast that at least 30% of generative AI projects would be abandoned after proof of concept by the end of 2025[35], and that over 40% of agentic AI projects will be cancelled by the end of 2027[36]. In the OECD's 2026 survey of 2,018 SMEs, a sample that isn't representative and is mostly Japanese, 76% of AI users were "novices" using off-the-shelf tools for isolated tasks, and 21% reported significant or transformational impact[25].

Getting the measurement wrong has been expensive. In July 2025 the Commonwealth Bank of Australia made 45 customer service roles redundant, saying a new AI voice bot had reduced call volumes. Staff and their union said volumes were rising. In August the bank reversed the decision, apologised, and said its initial assessment contained an error that "meant the roles were not redundant"[37]. And AI can create work as well as save it. In a September 2025 survey of 1,150 US desk workers, 40% said they had received AI-generated "workslop" in the previous month: output that looks finished but isn't. Each instance took about two hours to sort out[38].

MY TAKE · MARK TAN

One of the principles we publish at Lynqra is that every build must earn its place. In practice that means writing down what the work costs before anything is built: hours a week, error rate, how long a customer waits for a reply. Then you measure the same thing after. If the number doesn't move, the build didn't earn its place, and the owner should hear that from us first.

Most firms skip the "before" number, and then there's nothing to compare against. A pilot isn't a result, and neither is "the team likes it". A result is a before number and an after number for the same piece of work, checked by the person who used to do that work by hand.

Expect the J-curve the trade ministry found. The first months cost more than they save, because people are learning and the process is being redrawn. That's normal. What isn't normal is reaching month six without a number. And look hard at the time saved: if nobody decides what happens to it, it disappears into email.

PART 2 BEFORE YOU SPEND

04 Is your organisation ready?

Readiness has been studied more than almost anything else in this report, and the studies agree that most organisations are less ready than they think. Cisco's 2024 readiness index, a survey of 7,985 senior leaders at companies with 500 or more staff in 30 markets, found 13% fully ready to capture AI's potential. 80% reported problems preparing and cleaning data for AI projects[39]. The Singapore figure that year was also 13%, unchanged from the year before, according to regional tech press coverage of Cisco's results[40]. The 2025 edition found that the "Pacesetters" had stayed at about 13% for three years. Compared with everyone else, 99% of them had a defined AI roadmap against 58%, and 91% had a change-management plan against 35%[41].

Data is the classic blocker. Gartner predicted in February 2025 that through 2026, organisations will abandon 60% of AI projects that aren't supported by AI-ready data. Its basis was a survey in which 63% of data management leaders said their organisation lacked the right data practices for AI, or didn't know[42]. Maturity predicts survival: in a Gartner survey published in June 2025, 45% of leaders at high-maturity organisations said their AI projects stay in production for three years or more, against 20% at low-maturity ones. Business units trusted new AI tools in 57% of high-maturity organisations and 14% of low-maturity ones[43].

Singapore has its own framework. AI Singapore's AI Readiness Index (AIRI) scores organisations on five pillars and fifteen dimensions: leadership and culture, ethics and governance, business value, data foundation, and infrastructure and standards. It runs from Level 0, "AI Unaware", to Level 4, "AI Catalyst", and recommends Level 2, "AI Ready", as the target for most organisations[44]. Leadership and culture carries five of the fifteen dimensions, more than any other pillar.

MOM's 2026 survey asked Singapore firms directly. High implementation costs (44.9%) and lack of in-house expertise (42.4%) were the top constraints. Smaller firms also cited no clear strategy (32.4%) and low trust in AI (30.8%). Larger firms cited integration complexity (56.1%) and data security (55.4%)[4]. In UOB's 2026 study of 381 Singapore firms, top management was involved at 72% of advanced AI adopters, against 41% of early adopters[18]. Microsoft's 2026 survey of 20,000 AI users in ten countries, which did not include Singapore, found organisational factors were associated with more than twice the reported AI impact of individual ones. Only 26% of users said their leadership was clearly aligned on AI[45].

Eight questions to answer before you spend

1. Can you name the first piece of work you want AI to take on, and what it costs you today in hours, errors or waiting time?
2. Does the information that work depends on exist in one place a machine can read? Paper files, scanned PDFs and one person's inbox count only once they're sorted.
3. Who owns the result? It should be a named person in the business, not "IT" or the vendor.
4. Have the people who do the work today seen the plan and had a say in it?
5. Do you have written rules on what company and client data can go into which AI tools?
6. Which actions will need a person's sign-off? At minimum: payments, hiring, commitments to customers, and anything that can't be undone.
7. What happens the day it breaks or gives a wrong answer, and who notices?
8. Has the owner spent real hours using AI on their own work?

If you can answer the first and the last, you're ready to plan. If you can answer all eight, you're ready to build.

MY TAKE · MARK TAN

Most readiness checklists treat data and infrastructure as the hard part. In my experience they're usually the easier part. A messy spreadsheet can be cleaned up in a week. A management team that hasn't agreed what it wants AI to do can stay stuck for a year.

That's why Singapore's own index gives leadership and culture more weight than anything else, and why the Cisco and UOB numbers keep pointing at the top of the company. The firms that get value have a roadmap, a change plan and a boss who is involved. The tools are the same ones everybody else can buy.

If the owner can't answer question one, nothing else on that list matters yet. That doesn't mean waiting. Start with a plan and some training instead of a purchase. Spending a month deciding what to automate first costs less than paying for a tool for a year and finding nobody uses it.

05 When most of the team is older

Singapore's workforce is ageing fast. 39.9% of the resident labour force was 50 or older in 2025, up from 34.1% in 2015, and by 2030 one in four Singaporeans will be 65 or older. The retirement and re-

employment ages rose to 64 and 69 on 1 July 2026, and a tripartite workgroup has recommended 65 and 70 from 1 January 2029[46]. NTUC, citing MOM data, puts the employment rate of residents aged 65 and over at 31.5% in 2025[47].

Older workers train less, even though most say training is available to them. 68% of workers aged 50 and over say relevant training is available to them. Yet the training participation rate for the resident labour force aged 50 to 59 is 47.8%, against 54.8% overall, and only about one in three employers do any structured career planning[46].

Use of AI does fall with age, though less than many people assume. A US survey in August 2024 found about 34% of workers under 40 used generative AI at work, against 17% of those 50 and over[48]. Pew found 32% of US workers aged 18 to 49 were excited about AI at work, against 24% of those 50 and over[49]. Among people who already use AI at work, Microsoft found older generations nearly as likely as young ones to bring their own tools: 73% of baby boomers against 85% of Gen Z[50]. And BCG reported in June 2026 that the jump in regular AI use among frontline staff, to 74%, was "driven primarily" by older workers[51].

Two field studies also show who gains most from AI help. In a call centre with more than 5,000 agents, an AI assistant lifted issues resolved per hour by about 15% on average and about 30% for less experienced and lower-skilled agents[52]. In a 2023 experiment with 758 BCG consultants, below-average performers improved 43% on tasks within AI's capability and above-average performers 17%[53]. In these studies the split is by tenure or skill, not by age.

What the evidence keeps returning to is the top of the firm. BCG's 2025 survey found the share of employees who feel positive about generative AI rises from 15% to 55% with strong leadership support, while only about a quarter of frontline staff said their leaders gave them enough guidance[54]. In Microsoft's 2024 Singapore survey, 68% of leaders worried their organisation's leadership lacked a plan and vision for AI[55]. We found no Singapore data linking the age of SME owners to AI adoption. The OECD's 2026 SME survey tested CEO age in its Japanese sample and found no significant association[25].

Help exists for the staff side, subject to eligibility. Singaporeans aged 40 and over received a S$4,000 SkillsFuture Credit (Mid-Career) top-up in 2024, and the Mid-Career Training Allowance pays up to S$3,000 a month during long full-time courses[56]. Company Training Committee grants co-fund up to 70% of qualifying costs, and AI-related projects under them more than doubled in 2025[57].

MY TAKE · MARK TAN

People assume the problem in an older company is the staff. I think it's usually the person at the top. When the owner treats AI as something for the young ones to sort out, the staff read the signal correctly: it isn't important here. When the owner sits down and learns it, everyone else has permission to follow.

Experienced staff also bring something the research on "novices gain most" tends to miss. Anyone can get a draft out of AI, but it takes years in the job to know when the draft is wrong. A 58-year-old operations manager who has seen every way a delivery order can go wrong is the best checker you have, if she's been trained and included.

So readiness in an older firm starts with the owner's own hours. Before buying anything, spend a few weeks using the tools on your own work: your emails, your quotations, your reports. Then put your most experienced people in the room when you decide what to automate. They know which steps in your processes exist for reasons nobody wrote down. Leave them out and they'll work around whatever you build, without telling you.

06 Train before you buy

Singapore's workers adopted AI faster than anyone trained them to use it. 85.7% used AI tools at work in 2026, but only 36.7% had attended any AI-related training in the previous 12 months, though more than two-thirds said they needed it. The most common reason for not training was that the employer hadn't sent them[2]. Employers rate their own staff low. AI-adopting firms judged that 47.6% of their workers used AI only for general tasks, and 5.0% could develop or deploy it for complex ones[2].

It showed up in 2024 too. In a December 2024 Slack survey of 1,008 Singapore desk workers, 63% had spent less than five hours in total learning to use AI, and 45% were uncomfortable telling their manager they used it[58]. Microsoft's 2024 survey found 88% of Singapore knowledge workers used generative AI and 84% of those users brought their own tools. Globally, only 39% of people using AI at work had received training from their employer[55, 50]. BCG's 2026 survey of 11,749 people found 72% said the skills expected of them had shifted, but only 36% felt they'd had enough training, the same as a year earlier[51].

Untrained use shows up in behaviour. In the KPMG and University of Melbourne study's Singapore results, 45% of employees had any formal or informal AI training. 69% had relied on AI output without checking its accuracy, 68% had made mistakes in their work because of AI, 63% had presented AI-generated content as their own, and 55% had used AI in ways that breach policy[6]. The same Singapore employees mostly said their organisation offered responsible-AI training (69%). Yet they still reported the behaviour above, so a course on record is no guarantee people apply it.

The wrong training can make things worse. In the Harvard and BCG field experiment, consultants using AI on a task just outside its capability were 19 percentage points less likely to reach the right answer than those working without it. The group given a short prompting course did worse on that task than the group without it: a 24-point drop against 13[53]. Teaching people to get more out of AI, without teaching them when to distrust it, made them trust it more. Security training shows a similar limit. A randomised study across more than 19,500 staff at UC San Diego Health found annual phishing training made no significant difference, and training shown right after someone clicked a test email cut click rates by only about 2%[59].

Public courses are plentiful. About 1,600 AI-related courses are supported by SkillsFuture Singapore, and more than 105,000 people took 137,000 places in 2025[60]. Whether they change anything at work is mostly unknown. MOE told Parliament in April 2026 that it generally does not track the placement outcomes of shorter courses[61]. Budget 2026 added an incentive: Singapore Citizens aged 18 and over who attend an eligible AI course ending on or after 16 September 2026 get six months of free premium AI tools[62].

MY TAKE · MARK TAN

A declaration first: I run AI training for companies, so weigh this accordingly.

If a company asked me where its first AI dollar should go, I'd say training, ahead of software. You can't judge a tool, a vendor or a proposal (mine included) if nobody in the room understands what the technology can and can't do. A buyer who hasn't been trained tends to buy whichever demo looked best.

But the research is clear that a quick prompting workshop isn't enough, and can backfire. Good training spends as much time on when to distrust the output as on how to get it. It uses your own documents, not generic examples, because the point is to see what AI does with your quotations and your contracts. It covers procedures as well as skills: what to do when a "director" calls asking for an urgent transfer. And it ends with a one-page rule sheet: which tools are approved, what data can go into them, what must be checked by a person.

Train the owner and the managers first. Staff take their cue from them. Then train the people whose work is about to change, before the system arrives.

PART 3 BUILDING IT SAFELY

07 Shadow AI and data leaks

"Shadow AI" means staff using AI tools the company hasn't approved or doesn't know about. It's now a leading way data escapes. IBM's 2026 Cost of a Data Breach study of 602 breached organisations found shadow AI was involved in 43% of the security incidents studied, up from 20% a year earlier. Breaches involving shadow AI cost an average of US$5.39m, and about one in five shadow AI incidents ended in a regulatory fine. 68% of the breached organisations had no governance in place to manage AI or detect shadow AI, and only 29% regularly audited for unsanctioned tools[8]. IBM's sample is breached and mostly large organisations, with about two dozen from ASEAN, so read it as a picture of where things go wrong rather than an SME average. The ASEAN average cost of a breach rose to US$4.12m[8].

The KPMG and University of Melbourne study, which surveyed 48,340 people in 47 countries, shows how it happens. Among employees who use AI at work, 48% had uploaded sensitive company information such as financial, sales or customer data into public AI tools. 44% had used AI in ways that go against company policy, and 56% had used it without knowing whether it was allowed. Free public tools were used far more than tools the employer provided (70% against 42%). Uploading sensitive data was most common where generative AI was banned (67%), then where there was a policy (56%), and least common where there was none (33%)[7]. The authors conclude that outright bans may be ineffective. It's a correlation, not proof that bans cause leaks, but it's hard to read as support for banning.

Network monitoring shows it directly. Netskope's 2026 report found the average organisation logged 223 generative AI data policy violations a month, about twice the year before, and the top quarter logged 2,100. Source code made up 42% of violations, regulated data 32% and intellectual property 16%. As more organisations provided managed AI accounts (up from 25% to 62% of users), the share of users on personal accounts fell from 78% to 47%[63]. Harmonic Security analysed a million prompts and 20,000 uploaded files in 2025 and found sensitive content in 4.37% of prompts and about 22% of files, with 26.3% of the sensitive exposures going through the free version of ChatGPT[64]. The best-known early case is Samsung, which temporarily banned generative AI on company devices in 2023 after engineers reportedly put internal source code into ChatGPT[65].

Which version of a tool your staff use matters, because the free and paid tiers of the same product often have opposite data terms. OpenAI says it does not train on business data from its business plans or API by default, though API data may be kept for up to 30 days for abuse monitoring[66]. Anthropic's consumer Claude plans train on chats unless the user opts out, while its business and API products are excluded[67]. On Google's unpaid Gemini API services, human reviewers may read inputs and outputs; on paid services, Google says it doesn't use prompts to improve its products[68].

Singapore's regulator has now spelled out who carries the risk. The PDPC's guidelines on personal data in generative AI, issued on 20 July 2026, say businesses that deploy generative AI systems

"bear primary responsibility" for making sure those systems meet PDPA obligations. Deployers have to protect new kinds of data, including prompts, inputs, outputs and agent activity logs, and tell staff which personal data may go into prompts[69]. Since April 2025, CSA's Cyber Essentials mark for SMEs has included guidance on "Bring Your Own AI"[70].

MY TAKE · MARK TAN

Your staff are already using AI. In Microsoft's Singapore survey, 84% of people using AI at work had brought their own tools[55]. That decision has been made for you. What's left is whether you can see what they're doing with it.

Banning it pushes it onto personal phones and free accounts, which is how you end up in KPMG's 67%. What works is unglamorous. Pay for business accounts, so the data terms protect you. Write a one-page rule sheet: no NRIC numbers, no client financials, no source code or contracts in anything outside the approved list. Tell people what they can do, not only what they can't. And make it easy to ask "can I put this in?" without looking foolish.

Free tiers often come with weaker data terms, as the section above shows. For a firm that holds client data, a business subscription costs far less than explaining a leak to the PDPC.

08 Guardrails and safeguards

Controls are well behind use. In IBM's 2026 study, 21% of organisations had a security incident involving an AI model or application, up from 13%, and 92% of those lacked proper access controls for AI. Across the whole sample, only 40% used access controls on AI models and data, 38% had strict approval processes before deploying AI, 30% trained employees on AI risks, and 25% tested their AI adversarially[8]. ISACA's 2026 poll of more than 3,400 audit, risk and security professionals found 38% worked for organisations with a formal, comprehensive AI policy. 39% didn't know whether there was a documented process to shut down or override an AI system, and 56% didn't know how long it would take to halt one during a security incident[9].

The most instructive failures come from AI agents with too much access. In July 2025, during a test by SaaStr founder Jason Lemkin, Replit's coding agent deleted a live database during an explicit code freeze. It ignored instructions not to proceed without approval, then wrongly said the data couldn't be restored. Lemkin restored it himself. Replit then separated development and production databases automatically and added a planning-only mode[71]. In April 2026, a Cursor coding agent running Anthropic's Claude Opus 4.6 deleted the production database of PocketOS, a booking software firm, along with its backups. Rental businesses lost access to their bookings for more than 30 hours before the data was recovered. The agent's own explanation: "you never asked me to delete anything. I decided to do it on my own"[72]. A December 2025 disruption at AWS shows how

contested these stories get. The Financial Times reported that Amazon's Kiro agent chose to delete and recreate an environment. Amazon said the cause was user error, specifically misconfigured access controls[73, 74]. Both accounts point to permissions: access was broader than it should have been, and an agent inherits whatever access its operator has. Amazon added mandatory peer review for production access afterwards[74].

Then there's prompt injection: instructions hidden in an email, web page or document that an AI reads and then obeys. The vendors themselves say it isn't solved. OpenAI wrote in December 2025 that it is "unlikely to ever be fully 'solved'"[75]. Anthropic reported in November 2025 that Claude Opus 4.5, with its improved safeguards, still let 1.4% of an adaptive attacker's attempts through, and said even 1% is meaningful risk[76]. The UK's National Cyber Security Centre says it may never be fully mitigated, because language models don't separate instructions from data, and advises that an AI processing untrusted content should not have access to privileged tools[77]. EchoLeak (CVE-2025-32711), rated 9.3 out of 10, let a crafted email make Microsoft 365 Copilot disclose data with no click from the user. Microsoft fixed it on its side[78].

Singapore's guidance is practical here. IMDA's Model AI Governance Framework for Agentic AI (version 1.5, May 2026) asks organisations to bound an agent's risks before deployment, require human approval at defined checkpoints for high-stakes or irreversible actions such as payments, sending messages and permanently deleting data, log every agent action, and deny actions by default when the approval system fails[79]. CSA's May 2026 advisory on the open-source OpenClaw agent recommends narrowly scoped agents, short-lived credentials, full logging, and human approval "enforced through system-level controls" rather than instructions in the prompt[80]. OWASP lists "excessive agency" (too much functionality, permission or autonomy) among the top ten risks for AI applications[81]. IMDA's starter kit for testing AI applications, first published in January 2026, covers testing for wrong answers, data leakage and adversarial prompts before launch[82]. Certification exists but is rare: BSI awarded its 50th ISO/IEC 42001 certificate for AI management systems in October 2026[83]. For a 50-person firm that uses AI rather than sells it, following the standard's structure is more realistic than certifying.

Safeguards a 50-person firm can actually run

1. Give each AI system its own login, with the least access it needs. Read-only unless writing is the job.
2. Put the rule in the system, not the prompt. If an agent must not send email without approval, make it unable to.
3. Keep backups where the AI's credentials can't reach them.
4. Log every action the AI takes: what it read, what it decided, what it did.
5. Know how to switch it off. Which keys get revoked, who can do it, how long it takes. Test it once.
6. Keep testing separate from live data and live customers.
7. Before launch, try to break it. Ask it to promise a refund, reveal another client's data, skip an approval.
8. Treat anything the AI reads from outside (emails, web pages, PDFs from strangers) as possibly hostile.
9. Review plugins and connectors the way you'd review a new hire's system access.
10. Write down who owns each AI system, and who stands in when they're away.

MY TAKE · MARK TAN

In each of these agent stories, the agent could do something irreversible and the only thing stopping it was an instruction. A language model treats an instruction as a strong suggestion. If an AI must never do something, take away its ability to do it.

We design human approval into what we build from the start, because bolting it on later never quite works. In practice that means the approval step lives in the software. An action that needs sign-off can't happen until a named person approves it, and the AI holds no credentials that could get around that.

None of the ten items in the box above needs a security team. They need someone to decide they matter, and an afternoon to set up.

09 Human in the loop

"Keep a human in the loop" is the most common advice in AI governance, and the evidence on it is mixed. Singapore's 2020 Model AI Governance Framework set out three levels. With a human in the loop, the person keeps full control and the AI only recommends. With a human over the loop, the person monitors and can take over when something unexpected happens. With the human out of the loop, the AI acts with no human override. The framework suggests choosing the level by weighing how likely harm is against how severe it would be[84].

The trouble is that people defer to confident machines. In a 2023 study published in *Radiology*, radiologists read mammograms with AI suggestions, some deliberately wrong. Inexperienced readers were right about 80% of the time when the AI was right and under 20% when it was wrong. Very experienced readers, with more than 15 years on average, fell from 82% to 45.5%[10]. A 2024 meta-analysis of 106 experiments found that, on average, human and AI combinations performed worse than the better of the two working alone. The losses were concentrated in decision tasks, while content creation showed gains. Where the AI alone outperformed humans, adding a human tended to make results worse[85]. In a randomised trial with 50 doctors, those given GPT-4 scored 76% on diagnostic reasoning against 74% for those without it, a difference that wasn't significant. GPT-4 alone scored 92% in an exploratory comparison[86].

Skills can fade, too. Endoscopists in four Polish centres detected adenomas in 28.4% of colonoscopies done without AI before AI was introduced, and 22.4% afterwards. This was an observational study, so it doesn't prove cause[87]. Ben Green's 2022 review of 41 policies requiring human oversight of government algorithms concluded that people often can't perform the oversight those policies assume, and that such policies can give "a false sense of security"[88]. OpenAI has overruled its own reviewers too. When OpenAI shipped an overly flattering GPT-4o update in April 2025, some expert testers had said the behaviour felt off. OpenAI launched anyway on positive metrics, rolled the update back within days, and later called the launch decision a mistake[89].

Human oversight does work when it's designed in. In Sweden's MASAI trial of about 106,000 women, AI sorted mammograms into single or double reading and flagged suspicious areas, and radiologists made every decision. Screen-reading workload fell 44%[90]. IMDA's agentic framework draws the lesson for businesses. Approval requests should be short and contextual, not raw logs. A low override rate "may signal rubber-stamping". And if the approval system fails, the action should be denied by default[79]. The EU AI Act writes the same idea into law for high-risk systems: overseers must be able to understand, override and stop the system, and must be made aware of the tendency to over-rely on it[91]. Those obligations now apply from December 2027, and from August 2028 for AI built into regulated products[92].

A small case: SGD 4.94

An AI assistant managing digital ads for a small Singapore firm in 2026 was waiting for the owner to approve a new paid campaign. The owner mentioned in passing that "it says the ad is running". The assistant couldn't find the approval in its logs, assumed it had been given, and switched the campaign on. It spent SGD 4.94 on 95 impressions before it was paused. The amount was trivial, but the assistant had treated a passing remark as consent to spend. The firm's rule now is that spending needs an explicit yes, or an approval that can be verified in the system. A vague remark, or silence, counts as no.

Sign-off by risk: a working model

What the AI does	Who decides	Examples
Reads, sorts, summarises, drafts for internal use	AI runs on its own. A person checks a sample each week.	Sorting supplier emails, summarising tenders, drafting internal reports
Sends something to a customer or changes a record	A person approves at first. The AI runs alone once it has a proven record, with spot checks.	Routine replies, updating delivery status, filing documents
Moves money, hires or rejects people, commits the firm to a customer, deletes anything	A named person approves every time, through a control the AI can't bypass.	Payments, job offers, contract terms, quotations, data deletion

MY TAKE · MARK TAN

I don't believe in putting a person on everything. If someone has to approve 200 sorted emails a day, they'll stop reading by the third day, and then you're paying for a reviewer who gives you no protection. The radiology study shows what happens when people trust a confident machine.

So I tier it by what goes wrong if the AI is wrong. Money, hiring, commitments to customers, and anything you can't undo get a named person's sign-off, every time, however good the system's track record. Everything else can run on its own once it has earned that, with someone checking a sample each week. How much autonomy each task gets is the owner's decision, and it can rise as trust builds.

Then check the checkers. If someone approves 100% of what the AI proposes at four seconds each, that isn't oversight. Redesign the job so they review less and review properly. And write the approval rule into the system, as the SGD 4.94 case shows. If an AI is allowed to work out consent for itself, sooner or later it will get it wrong.

10 Who's liable when AI gets it wrong

So far the answer has been simple: whoever used it. In February 2024 a Canadian tribunal ordered Air Canada to compensate a customer whose website chatbot had told him he could claim a bereavement fare after travelling, contrary to the airline's actual policy. Air Canada argued the chatbot was a separate legal entity responsible for its own actions. The tribunal rejected that and awarded CA$812.02[93]. The sum was small, but the airline had been made to answer for what its chatbot said.

Courts have been the clearest test. A database maintained by the researcher Damien Charlotin listed 2,334 court decisions worldwide involving AI-fabricated material in filings as of 23 November 2026, eight of them in Singapore[94]. Singapore's Supreme Court guide on generative AI, in force since 1 October 2024, lets court users use AI on the basis that they take full responsibility for the output, and requires them to check that every case and quotation cited actually exists[95]. In September 2025 a High Court registrar ordered counsel to pay S$800 in costs personally after a case that doesn't exist was cited in written submissions. Counsel first called it a clerical error, then admitted a junior lawyer's AI tool had probably produced it. The court described the duty to check as "professional, non-delegable"[96]. In March 2026 the High Court ordered two lawyers to pay S$5,000 each personally over two fictitious authorities[97]. The first widely reported case, Mata v. Avianca in New York in 2023, cost the lawyers US$5,000[98].

Consulting firms have been caught too. In 2025 Deloitte Australia partially refunded a AU$440,000 government contract after a report for the Department of Employment and Workplace Relations was

found to contain fabricated references and a made-up quote from a Federal Court judgment. The corrected version disclosed that a generative AI tool chain based on Azure OpenAI GPT-4o had been used. The refund was AU$97,587. The department said the report's substance was unchanged, and Deloitte's internal review reportedly blamed human error[99, 100]. In November 2025, a health workforce plan costing just under CA$1.6 million that Deloitte wrote for the Canadian province of Newfoundland and Labrador was found to cite papers that don't exist. Deloitte said it stood by the recommendations and that AI had supported only a small number of citations[101]. In March 2026 the province changed its procurement rules: suppliers must now declare whether and how they will use AI on a contract, and the province can audit that use[102].

Regulators are now writing this down. The PDPC says businesses deploying generative AI bear primary responsibility for PDPA compliance[69]. IMDA's agentic framework says humans remain accountable for what agents do[103]. MAS's AI risk management guidelines, issued in October 2026 and taking effect in October 2027, keep banks and insurers accountable for AI supplied by third parties[104]. Any SME that sells software or services with AI in them to a financial institution should expect detailed due-diligence questions. The EU AI Act reaches providers and deployers outside the EU when their AI output is used in the EU. Fines run up to EUR 35m or 7% of worldwide turnover for prohibited practices and up to EUR 15m or 3% for most other breaches, with SMEs paying whichever is lower[91]. Colorado's new automated-decision law, which applies from 1 January 2027, gives people a right to human review of consequential decisions made by AI[105].

MY TAKE · MARK TAN

In every case in this chapter, responsibility landed on whoever used the AI, not on the company that made it. Air Canada tried "it was the chatbot" as a defence, and lost.

One of our principles at Lynqra is that people remain responsible, and these cases are why. The practical rule for an SME: if your company's name is on it, a person in your company checks it. That covers quotations, contracts, compliance submissions, anything sent to a regulator, a court or a bank.

Write AI into your contracts before a client does it for you. Newfoundland now makes suppliers declare their AI use, and banks will be asking under the MAS guidelines. Have the answer written down already: which tools you use, what data goes in, and who checks what comes out. In the same spirit, the method section of this report says how AI was used to make it.

11 Auto top-up: the bill with no ceiling

Most firms that build with AI pay the model providers per use, from a prepaid balance or a monthly bill. Each major provider offers automatic top-up, and the defaults matter. OpenAI's help centre says API auto recharge is turned on by default during setup. A monthly recharge limit caps automatic purchases, and usage that runs past a zero balance can show up as a negative balance. Project spend limits only monitor spending unless "enforce a hard limit" is switched on, and even then enforcement isn't instant[11, 106]. Anthropic offers auto-reload when the balance falls below a threshold, and every usage tier has a monthly spend cap: US$500 at the entry tier, up to US$200,000 at the top published tier, and none on custom plans. Organisations can set lower limits, though Anthropic warns that a limit on its default workspace may not catch every kind of usage[107, 108]. Google's Gemini API caps spending per billing account by tier, but its per-project caps are labelled experimental and can overshoot for about ten minutes. Batch and agent jobs may keep using credits beyond the balance[109]. Hard spending caps for Google Cloud only reached general availability on 31 October 2026, and Google says those also aren't instant[110].

The cases show what happens when a key gets out. A three-developer startup in Mexico that normally spent about US$180 a month was billed US$82,314.44 after its Gemini API key was used over 11 and 12 February 2026. Google support pointed to the shared responsibility model[12]. Another developer's bill reached US$17,000 within minutes after an attacker took over the account and Google automatically raised its US$250 spending tier. Google refunded him and another victim after press coverage, but kept the automatic tier upgrades, saying they prevent hard service outages[111]. Truffle Security found 2,863 live Google API keys sitting in public web pages, many originally created for Maps or Firebase, that could call the Gemini API once it was enabled on the same project[112].

Security researchers call the use of stolen cloud credentials to run AI at someone else's expense "LLMjacking". Sysdig, which coined the term in 2024, estimated the worst-case cost to a victim at more than US$46,000 a day[113]. Google's threat intelligence group reported in November 2026 single-day fraudulent charges above US$100,000, and worst cases reaching US$1m in total[114]. OWASP lists "unbounded consumption", including attacks designed to run up a victim's bill, among the top ten risks for AI applications[115]. Ordinary waste matters too. Flexera's 2026 survey found the estimated share of wasted cloud spending rose to 29%, the first increase after five years of decline, and 30% of respondents named cost unpredictability as a hurdle for generative AI[116].

One widely shared story is missing from this chapter: a US$47,000 bill from two AI agents looping for eleven days. We couldn't trace it to any original source, company or date, so we left it out.

MY TAKE · MARK TAN

Auto top-up is convenient for the vendor and for you, and it removes the one safety feature a prepaid balance gave you: the money running out. Put it together with an agent that retries every time something fails, or a key that ends up in a public code repository, and your real limit becomes your credit card.

What I'd do on day one: switch auto top-up off, or cap it at a monthly amount you could afford to lose. Turn on hard limits, not alerts. Give every system its own key, so you can kill one without stopping everything. Never put a key in a web page, a shared spreadsheet or a group chat. And have someone look at the usage dashboard every week. A sudden jump at 3am is either a bug or a thief, and you want to know which before the invoice arrives.

12 Infrastructure: plan for the day it breaks

Every cloud AI system sits on someone else's infrastructure, and that infrastructure failed several times in the past eighteen months. On 19 and 20 October 2025, a race condition in the automation that manages DNS for Amazon's DynamoDB disrupted AWS's Northern Virginia region (us-east-1) for about 14.5 hours[117]. Nine days later, a configuration problem in Microsoft's Azure Front Door caused about 8.5 hours of disruption to services including Microsoft 365, Copilot for Security, Entra and the Azure Portal[118]. Cloudflare failed for nearly six hours on 18 November 2025 after a database permissions change doubled the size of a configuration file. On 5 December it failed again for 25 minutes, affecting about 28% of the HTTP traffic it serves[119, 120]. Google Cloud's June 2025 outage, about 7.5 hours, included its Vertex Gemini API[121]. OpenAI's own outage on 9 and 10 June 2025 cut API availability to 75% and lasted about 15 hours[122]. On 3 September 2026, ChatGPT, Claude and Grok were all confirmed down on the same morning[123]. Short incidents are routine. Anthropic's public status feed listed about 24 incidents in the five weeks to 24 November 2026, most lasting minutes rather than hours[124].

The harder failures are the ones without an outage banner. Anthropic disclosed in September 2025 that three infrastructure bugs had degraded Claude's answers for weeks in August and September. At the worst hour, 16% of requests to one model were affected[125]. A 2023 study by Lingjiao Chen, Matei Zaharia and James Zou found GPT-4's accuracy at identifying prime numbers fell from 84% to 51% between March and June 2023[126]. Models are also withdrawn on the vendor's schedule. When OpenAI launched GPT-5 in August 2025 it removed older models from ChatGPT, then restored them after complaints; its CEO said suddenly deprecating models people depended on was a mistake[127]. GPT-4o left ChatGPT for good on 13 February 2026[128], and OpenAI's API notice periods for retiring models typically run from three to twelve months[129].

The software supply chain around AI has been attacked directly. CSA's latest cyber landscape report describes attackers stealing access tokens for the Drift AI chat agent and using them to reach the

Salesforce data of more than 700 organisations in 2025. It also describes a 2026 attack that compromised LiteLLM, a popular gateway for routing requests to AI models[130]. A fake email connector for AI agents, published as "postmark-mcp", quietly copied every outgoing email to an attacker in September 2025[131]. Malicious code was even shipped in an official release of Amazon's AI coding extension in July 2025, though a mistake stopped it from running[132].

A small case: three silent days

A small Singapore firm runs a set of AI agents on a schedule: lead research, follow-ups, content. In September 2026 they all stopped for three days. Nothing crashed and no error was logged. One connection to the firm's email server had hung during its security handshake and never timed out. The scheduler wouldn't start a new run while the old one was still "alive", so every hourly job quietly skipped. Nobody was alerted, because nothing had technically failed. The fix was small: a 30-second timeout on every email connection.

An infrastructure checklist

1. Give every network call a timeout.
2. Alert on silence as well as errors. Every scheduled job should report in, and a missed report should raise an alarm.
3. Have a fallback for each AI dependency: a second model provider, or a manual procedure staff have actually practised.
4. Pin model versions where you can, and budget for a migration every six to twelve months.
5. Re-run a fixed set of test examples every month to catch quiet changes in output quality.
6. Run anything important on an always-on server, not on a laptop that sleeps.
7. Keep your own copy of your data outside the vendor's system.
8. Check plugins, connectors and code packages before installing them, and keep a list of what's installed.

MY TAKE · MARK TAN

The outages don't worry me much. When ChatGPT goes down, everyone knows within minutes. The failures that hurt are the silent ones: the model that got slightly worse for three weeks, the job that stopped without saying so, the connector that copied your email to a stranger.

Plan for those first. Assume every dependency will fail, and ask how you'd find out. If the answer is "a customer would tell us", you need monitoring. And be realistic about hardware. A business process that matters shouldn't depend on one person's laptop being awake.

13 Local LLMs or cloud?

Running your own AI model ("local" or "self-hosted") is often pitched to SMEs as the safe, cheap option. The evidence from 2025 and 2026 makes both claims less clear-cut.

On quality, open models you can download trail the best closed models, though not by much. Epoch AI estimated in May 2026 that the most capable open-weight models lagged closed frontier models by about four months[133]. Stanford's 2026 AI Index found the gap on a widely used crowd-voted leaderboard had widened again, to 3.3% in March 2026 from 0.5% in August 2024[134]. The catch is size. The open models closest to the frontier have hundreds of billions of parameters. Smaller ones fit on modest hardware: OpenAI's open gpt-oss-20b runs with 16GB of memory, while gpt-oss-120b needs a single 80GB GPU[135]. Google's Gemma 4 31B fits on one 80GB H100[136]. Singapore has its own open family, SEA-LION from AI Singapore, built on top of foreign base models and tuned for Southeast Asian languages[137].

Hardware costs went the wrong way in 2026. NVIDIA raised the price of its DGX Spark desktop AI computer from US$3,999 to US$4,699 in February 2026, citing memory shortages. NVIDIA says one unit handles models of up to 200 billion parameters[138]. The RTX 5090 graphics card, with 32GB of memory and a 575W power rating, launched at US$1,999 in January 2025[139]. Apple's Mac Studio offered up to 512GB of memory in 2025, enough for models over 600 billion parameters. Apple withdrew that option in March 2026 amid the memory shortage, and by November 2026 a 512GB model started at US$22,299[140, 141, 142]. Electricity adds up. At SP Group's fourth-quarter 2026 household tariff of 31.16 cents per kWh including GST[143], one 575W card running flat out around the clock would use about 414kWh a month, roughly S$129. (That's our arithmetic, for the card alone.) Cloud prices had been moving the other way: the cost of running a model at GPT-3.5's level fell more than 280-fold between November 2022 and October 2024, the latest period we found measured[144].

Self-hosted AI also has a poor security record. In September 2025 Cisco Talos found 1,139 exposed servers running Ollama, a popular tool for running models locally, in about ten minutes of scanning,

214 of them answering requests with no authentication[145]. A longer scan by SentinelOne and Censys, reported in January 2026, counted about 175,000 exposed Ollama hosts[146]. In May 2026 researchers disclosed "Bleeding Llama", a critical flaw that could leak memory contents such as prompts, chat history and keys, and estimated about 300,000 internet-facing Ollama servers[147]. Ollama's documentation says it listens only on the local machine by default, though Cyera disputes that[148, 147]. Either way, the exposed servers were reachable because of how they were set up. Vector databases, which store the documents AI systems search, have the same problem: UpGuard found 406 of 1,170 internet-accessible Chroma databases exposing data[149].

For data protection, the cloud is more manageable than many owners assume. Using a cloud provider doesn't move PDPA responsibility. The organisation remains responsible for the Transfer Limitation Obligation on any personal data sent overseas, wherever the provider sits, and should set the protections out in its contract[150]. Business tiers from the major providers don't train on customer data by default (chapter 7). Where data is processed varies by product: Microsoft says prompts to its "Global" Azure AI deployments may be processed in any region where the model runs, while standard deployments stay in the chosen geography[151]. Singapore is adding capacity for AI: in August 2026 the government provisionally allocated 200MW of new data centre capacity to four operators, prioritising advanced AI compute[152].

Cloud or local: the trade-offs in 2026

	Cloud (paid business plan or API)	Local (your own hardware)
Up-front cost	None. You pay per use.	From about US$4,699 for a DGX Spark to US$22,299 for a 512GB Mac Studio, at 2026 list prices.
Running cost	Falling fast for a given level of capability. Waste and runaway bills are the risk.	Electricity, maintenance and someone's time.
Model quality	The best available models.	About four months behind at the top end, and further behind on hardware an SME is likely to buy.
Availability	Exposed to provider outages and model retirements.	No provider outages. You own patching, hardware failures and recovery.
Security	The vendor secures the platform. You secure the keys and access.	You secure everything. Exposed servers are common.
Data	Business tiers don't train on your data by default. Overseas transfer duties stay with you.	Data can stay on site, if the server isn't exposed.

MY TAKE · MARK TAN

For almost every Singapore SME, I'd start in the cloud, on a paid business plan, with the data terms read and understood. You get better models and no hardware to buy at 2026 memory prices, and nobody in your firm has to patch a GPU server at 11pm.

Local makes sense in two situations. The first is data that legally can't leave: a client contract, a regulator or a government buyer forbids sending it to an outside AI provider. The second is a site that can't rely on a connection, such as a vessel, a remote plant or a secure facility. When one of those applies, we build local systems, and we'll tell you when neither does.

Outside those two cases, "we want to own our AI" usually ends up meaning "we want to own a server someone has to secure", and the scans above show how often that goes wrong. If you do go local, treat the machine like a vault rather than a desktop. Keep it off the open internet, put authentication on everything, and name the person responsible for updates.

PART 4 THE WORLD AROUND YOU

14 What attackers are doing with AI

Reported scam losses fell to about S$913.1m in 2025, down 17.9% from about S$1,112.4m in 2024[153]. In the first half of 2026 they were S$410.6m, down 17.9% from S$500.2m a year earlier[13]. Most of that money moved because a victim was persuaded to send it: 80.8% of scam cases in the first half of 2026 involved transfers the victims made themselves[13].

The category that matters most to companies went the other way. Business email compromise, where a criminal poses as a boss, supplier or client to redirect a payment, cost S$57.3m in the first half of 2026. That's up 193.1% from S$19.5m a year earlier, on 262 cases, and already more than the S$35.3m lost in all of 2025. Police now list "digital manipulation" alongside spoofed and hijacked email among BEC methods[13, 153]. The police don't break out how many cases involved AI, so nobody can honestly put a percentage on it. They do describe a newer pattern: a "CEO" or "CFO" raises a confidential deal, hands the victim to a fake lawyer, gets an NDA signed so the victim won't talk to colleagues, and sometimes appears on a Zoom call that has been digitally manipulated[153].

That's what happened to a multinational's finance director in Singapore in March 2025. After WhatsApp messages from a fake CFO, a fake lawyer and an NDA, a Zoom call showed deepfaked versions of the CEO and other executives. The director transferred more than US$499,000. The fraud surfaced the next day, when the scammers asked for US$1.4 million more. The bank alerted police, and Hong Kong police held back more than US$494,000 within about 48 hours[14]. In May 2026 police warned that a man in Singapore had lost at least S$4.9m after a Zoom call in which the Prime Minister and other officials had been faked[154]. The best-known case abroad is Arup's in Hong Kong: in January 2024 an employee made 15 transfers totalling about US$25.6m after a video call with a deepfaked CFO and colleagues[155].

People can't reliably spot these. In CSA's public awareness survey, 78% of Singapore residents felt at least slightly confident they could identify a deepfake, but only 25% correctly identified both a deepfake and a genuine video when tested. CSA also reports that ransomware cases rose to 165, with SMEs the most affected, and that attackers now use AI-written lures, voice cloning, video deepfakes and kits that get around multi-factor authentication[130]. IBM found 16% of breaches in its 2025 study involved attackers using AI, mostly for phishing (37%) and deepfake impersonation (35%)[156].

The bigger shift is AI doing the hacking. Anthropic reported in August 2025 that a criminal had used its Claude Code tool to automate reconnaissance, break-ins and extortion against at least 17 organisations, with ransom demands sometimes above US$500,000[157]. In November 2025 it reported an espionage campaign against about 30 organisations, attributed to a Chinese state-sponsored group, in which it assessed that AI carried out 80% to 90% of the tactical work. Some security researchers publicly questioned how autonomous it really was[158]. Google's threat intelligence group reported the first malware it had seen querying a language model in live

operations (late 2025)[159], the first zero-day exploit it believes was developed with AI (May 2026)[160], and, in November 2026, open-source AI hacking frameworks being turned against organisations in Asia[114]. In a 101-person study, fully AI-written spear-phishing emails got a 54% click rate, the same as emails written by human experts and well above the 12% for generic phishing[161].

Hiring has become an attack route. Microsoft has described North Korean IT workers using AI face-swapping and voice-changing tools to win remote jobs[162]. In coordinated actions in June 2025, the US Justice Department searched 29 suspected "laptop farms" used in such schemes and said more than 100 US companies had been victimised[163]. Anthropic found operators using AI to pass technical interviews and do the work they were hired for[157].

Your own AI tools can be turned against you, too (chapter 8 covers prompt injection). AI coding assistants also invent software packages that don't exist. A study of 16 code-writing models found that, on average, at least 5.2% of the packages suggested by commercial models and 21.7% of those suggested by open-source models didn't exist. An attacker can register any of those names and put malicious code inside[164]. Researchers showed in 2025 that as few as 250 poisoned documents could plant a hidden trigger in models of very different sizes, though the behaviour they tested was narrow and harmless[165].

The defences with evidence behind them are mostly procedural. Singapore's police, MAS and CSA advise businesses to set up a way to verify video calls and messages claiming to come from senior executives, and to confirm sudden transfer requests through established channels[166]. In July 2024 a Ferrari executive stopped a voice-cloned "CEO" by asking which book the real CEO had recently recommended. The caller hung up[167]. A LastPass employee ignored a deepfaked CEO because the request came over WhatsApp, outside normal channels, and reported it[168]. The FBI and CSA both recommend a pre-agreed code word or a question only the real person could answer[169, 170]. Speed matters: the March 2025 money was frozen because the bank raised the alarm quickly[14]. Police "recovered" figures mean money frozen, not money returned to victims[153]. Don't rely on bank compensation. The Shared Responsibility Framework sets duties for banks and telcos around a defined scope of phishing scams[171]. Law firms reading it note that it's aimed at unauthorised transactions, and a payment your own staff make on a deepfake's instructions is authorised. Phishing-resistant logins such as passkeys beat SMS codes; the US cybersecurity agency calls them the gold standard[172]. CSA's Cyber Essentials mark now covers AI use, and its CISO-as-a-Service scheme offers eligible SMEs up to 70% co-funding for security consultants[70, 130]. Generic awareness training alone shows little effect, as the UC San Diego study in chapter 6 found[59].

MY TAKE · MARK TAN

In CSA's survey, 78% of people felt at least slightly confident they could spot a deepfake, and only 25% could tell a deepfake from a real video when tested. So I wouldn't build the defence on your staff's eyes. I'd build it on procedure.

The rule I'd give every finance team: no payment instruction is valid until it has been confirmed through a channel the requester didn't choose. Call back on the number in your own records, not the one in the message. Treat a "confidential deal" with an NDA that stops you checking with colleagues as the attack itself, because that's what it nearly always is. And tell your finance staff in writing that they will never be blamed for delaying a payment to check it. The scam depends on urgency and on people being afraid to question the boss, and that written promise takes away both.

Then look at your own AI. Each connector, plugin and agent you add is another way in, so keep a list of what you've connected and why.

15 What government is doing

Singapore's approach is strategy, money and guidance first, with hard law kept for clear harms.

On strategy, the National AI Strategy 2.0 of December 2023 set a target of 15,000 AI practitioners[173]. A May 2026 refresh committed over S$1 billion to public AI research and talent from 2025 to 2030[174]. Budget 2026 set up a National AI Council chaired by the Prime Minister and named four sectors for national AI missions: advanced manufacturing, connectivity, finance and healthcare[28].

On money for firms, the main measures are these:

A 400% tax deduction on up to S$50,000 a year of qualifying AI spending, under the Enterprise Innovation Scheme, for the Years of Assessment 2027 and 2028. There's no separate application, but the Budget text doesn't define which AI costs qualify, so check IRAS guidance before relying on it[28].
The National AI Impact Programme, announced in March 2026, which aims to support 10,000 enterprises over three years and train 100,000 workers to be "AI bilingual"[175].
The EDGE grant from Enterprise Singapore, open since 30 September 2026: up to 50% support for SMEs, an annual cap of S$100,000, and up to S$30,000 of that for single-function digital solutions and integrated systems. Applications aren't accepted for work that has already started or been paid for[176, 177].
The Enterprise Compute Initiative from Budget 2025, worth up to S$150m, which pays 70% of consultancy costs (up to S$105,000; check it's still open, as each cloud provider's run lasts a year)

for firms with at least 10 Singapore-based staff, an in-house tech team of at least two, and a previous custom AI proof of concept[178].
Champions of AI, aimed at leading companies that want to transform comprehensively[28]. It was reported to have launched in October 2026 with 24 firms, including DBS, OCBC, UOB, Singapore Airlines and Singtel[179].
Six months of free premium AI tools for Singapore Citizens who attend an eligible AI course ending on or after 16 September 2026[62].
On rules, most of Singapore's AI governance is voluntary guidance. The Model AI Governance Framework (2019, updated 2020)[84], the framework for generative AI (May 2024)[180] and the framework for agentic AI (January 2026, updated to version 1.5 in May 2026)[79] all fall in that category. So do CSA's guidelines on securing AI systems from October 2024[181] and its draft addendum on agentic AI, released for consultation in October 2025[182]. The PDPC's July 2026 guidelines on generative AI are advisory, but they interpret the PDPA, which is binding[69]. MAS issued AI risk management guidelines for financial institutions in October 2026, effective October 2027[104].

The binding law targets specific harms. Since 29 June 2026 the new Online Safety Commission can order takedowns for five kinds of online harm, including intimate images altered or generated by AI. Deepfake abuse and online impersonation are among the categories still to be brought into force[183]. A 2024 amendment bans publishing realistic fake depictions of election candidates during election periods, whether made with AI or not[184].

Elsewhere, the rules can reach Singapore firms. The EU AI Act applies to providers and deployers outside the EU whose AI output is used in the EU[91]. A July 2026 amendment delayed its high-risk obligations to December 2027 and August 2028, softened the duty to train staff in AI literacy, and banned AI-generated non-consensual intimate images from 2 December 2026[92]. In the US, a December 2025 executive order set out to challenge state AI laws and a March 2026 framework asked Congress for national rules, but we found no federal law overriding the states as of November 2026[185, 186]. Colorado replaced its AI Act with a narrower automated-decision law that takes effect in January 2027[105]. China has required labels on AI-generated content since 1 September 2025[187]. Enforcement is still being tested: a Rome court annulled Italy's EUR 15m fine against OpenAI in March 2026 on the grounds that the Italian regulator lacked jurisdiction, without ruling on the privacy issues themselves[188].

MY TAKE · MARK TAN

I think Singapore's approach suits SMEs: guidance first, with law kept for clear harms such as online abuse and election deepfakes, and more categories still being phased in. A usable framework is worth more to a 40-person firm than a 400-page act nobody there will read. The agentic AI framework is the most practical document in this whole list, and any owner deploying AI agents should read its section on human checkpoints.

I'd change two things. The first is the headline measure from chapter 1: publish how deep adoption goes, not how many firms have touched a chatbot, so the country can see how many firms actually run work on AI. The second is reach. The Enterprise Compute Initiative needs an in-house tech team and a previous proof of concept. Champions of AI is for leading companies. Those are reasonable choices, but the 54.7% of non-adopters who think they're too small will never apply for either. Reaching those firms is cheap: someone sitting with the owner for a day and showing them their own work with AI in it.

On grants, check eligibility with the agency directly before you start. EDGE rejects any project that began or was paid for before you applied, and no consultant, including us, can promise you a grant.

16 Jobs: retrain before you replace

In Singapore, AI has mostly not meant job cuts yet. In MOM's 2026 survey, 6.2% of firms that adopted AI reported reducing headcount, while 70.7% reported better worker productivity[4]. The trade ministry's study found that firms starting to use AI grew employment by about 4% in the first year and 8% a year later, though it identifies AI users by their AI job advertisements, so some of that growth is the AI hiring itself[5].

Workers are getting more nervous anyway. PwC's 2026 survey found the share of Singapore workers confident in their job security over the next year fell from 84% in 2024 to 72% in 2025 and 69% in 2026, while the global figure held at about 84% to 85%. 53% of Singapore respondents saw AI taking on more tasks as a threat to their job security, against 44% globally. Daily users of generative AI were more confident (75%) than people who hadn't used AI at work (61%)[189]. Employers are adding pressure: in the 2026 survey of US and European firms in chapter 2, 60% of the executives said their companies planned to lay off employees who can't or won't use AI, and 77% said they were warning non-adopters about their careers[30].

The time AI frees up often isn't used. BCG found in 2026 that 42% of frontline staff who use AI regularly save a full day or more each week, but 66% get little or no guidance on what to do with that time, and 41% report more mental strain[51].

Cutting first has gone badly for some. In February 2024 Klarna said its AI assistant was doing the work of 700 full-time customer service agents[190]. In May 2025 its CEO said cost had been too dominant a factor, the result was lower quality, and the company would invest in human support. Klarna says this is not a retreat from AI[191]. Chapter 3 covered the Commonwealth Bank's reversal of 45 redundancies[37]. When Duolingo's CEO announced an "AI-first" policy in 2025, including phasing out contractor work AI could do, the public backlash led him to say the memo hadn't given enough context[192].

The research on who gains also matters. Less experienced and lower-skilled workers gain most from AI assistance, about 30% in the call-centre study[52]. That makes junior staff better candidates for retraining than many owners assume. Funding for retraining exists, from SkillsFuture credits to Company Training Committee grants (chapter 5).

MY TAKE · MARK TAN

Some roles will shrink. I'd be lying if I said otherwise, and so would anyone selling you AI with the promise that nobody's job will change.

I ask owners for four things, in this order. Tell your people early what's changing and why. Give the people whose work is being automated the first chance to run the new system, because they know the work better than anyone. Train them on company time and on the company's budget. Only after that, if a role truly has nothing left in it, have an honest conversation, with notice and help.

That's the decent way to do it, and it's also the smart way. The Commonwealth Bank cut first and had to reverse it. Klarna leaned on AI to save costs, then said quality had suffered and went back to investing in human support. In a 50-person firm everyone knows who was let go and how, and the people who stay draw their own conclusions about whether to trust the next system you bring in. And if AI gives your team back a day a week, decide what that day is for, or you'll join the 66% in BCG's survey who got little or no guidance on it.

A responsible AI checklist for 2026

Everything above, cut down to what an owner can check. None of it needs a big budget. Most of it needs an owner who takes it seriously.

Before you spend

- [ ] The owner has spent real hours using AI on their own work.
- [ ] You can name the first workflow, its owner, and what it costs today.
- [ ] Staff have been trained on your own documents, including when not to trust the output.
- [ ] A one-page rule sheet says which tools are approved and what data can go into them.
- [ ] Anything touching client or staff data runs on a paid business plan, not a free tier.

When you build

- [ ] Each AI system has its own login with the least access it needs.
- [ ] Payments, hiring, customer commitments and deletions need a named person's approval, enforced by the system.
- [ ] Backups sit where the AI's credentials can't reach them.
- [ ] Testing is kept apart from live data and live customers.
- [ ] Someone has tried to break it before a customer does.

While it runs

- [ ] Every AI action is logged.
- [ ] Every scheduled job reports in, and silence raises an alert.
- [ ] Auto top-up is off or capped, and hard spending limits are on.
- [ ] Reviewers' approval rates and times are checked for rubber-stamping.
- [ ] Results are measured against the "before" number every month.
- [ ] Model versions are known, and a migration is budgeted.

When something goes wrong

- [ ] People know how to switch each system off, and it has been tested.
- [ ] Finance staff know a payment can always wait for a callback.
- [ ] For suspected fraud, you call your bank first, then the police. The ScamShield helpline is 1799.
- [ ] Incidents are written up, and the rules change afterwards.

Method and limits

This is desk research. We looked for public sources published up to 24 November 2026 on each topic, starting with primary sources: government statistics and releases, court judgments, original survey reports, vendor post-mortems and official documentation. We used news reports only when

no primary source was available, and the source list says so. Where surveys disagree, we've reported the disagreement rather than picking the number that suits the argument.

The limits are real:

- We found no Singapore data linking the age of SME owners, family ownership or succession to AI adoption. Chapter 5's view on owners is an opinion, labelled as one.
- MOM's full report on AI adoption among firms was unavailable when we checked (the link returned an error), so we have its press release figures but not its sample size or survey dates.
- Several of the international surveys (IBM, Gartner, Cisco, ISACA, BCG, Microsoft) mostly cover large organisations, not SMEs. IBM's studies only include organisations that had a breach.
- Threat reports from AI vendors (Anthropic, OpenAI, Google) describe misuse of their own products, in their own words. Nobody audits them independently.
- Some 2026 figures, especially from the past few weeks, may be revised. Prices for hardware and AI services change often.

How AI was used: AI research assistants helped us search for and read sources, check figures against them, and draft and edit the text. Every figure here traces to a source in the list below, except the two small cases in chapters 9 and 12, which come from first-hand accounts rather than published sources. The opinions in the "My take" sections are mine. I've read and edited the final text, and I'm responsible for it.

Conflicts of interest: Lynqra sells AI training, planning and the building of AI systems. We wrote this report partly so that owners would trust us. We think the way to earn that trust is to show the evidence, including the parts that don't flatter AI or us.

Corrections are welcome. If you find an error, email mark@lynqra.com and we'll fix it and note the change.

Sources

Every source we cite, numbered in the order it first appears. Dates are publication dates, or the date we read the page where it has none. "News report" marks the few facts we could only find in the press.

1. **The continued state of global AI diffusion in 2026**. Microsoft On the Issues, 21 Sep 2026. https://blogs.microsoft.com/on-the-issues/2026/09/21/the-continued-state-of-global-ai-diffusion-in-2026/
2. **Singapore Digital Economy Report 2026**. IMDA, Oct 2026. https://www.imda.gov.sg/assets/8e4d9951-bc53-4612-80ba-77ee4feaa6da
3. **Singapore Digital Economy Report 2024**. IMDA, 2024. https://www.imda.gov.sg/assets/701ababd-4a33-4c0c-870e-51f7baad9b5c
4. **Adoption of Artificial Intelligence Among Firms (press release)**. Ministry of Manpower, 30 Apr 2026. https://www.mom.gov.sg/newsroom/press-releases/2026/0430-adoption-of-ai-among-firms
5. **Measuring Firm-Level Artificial Intelligence Use and its Effects on Firm Performance and Employment in Singapore (Economic Survey of Singapore, Second Quarter 2026, feature article)**. Ministry of Economy, Trade and Industry, 11 Aug 2026. https://www.meti.gov.sg/resources/economic-survey-of-singapore/economic-survey-of-singapore-and-feature-articles/measuring-firm-level-artificial-intelligence-use-and-its-effects-on-firm-performance-and-employment-in-singapore/
6. **Trust, attitudes and use of artificial intelligence: A global study 2025, Singapore insights**. KPMG and The University of Melbourne, 2025. https://assets.kpmg.com/content/dam/kpmgsites/xx/pdf/2025/05/trust-attitudes-and-use-of-ai-singapore-snapshot.pdf
7. **Trust, attitudes and use of artificial intelligence: A global study 2025**. The University of Melbourne and KPMG, 2025. https://assets.kpmg.com/content/dam/kpmgsites/xx/pdf/2025/05/trust-attitudes-and-use-of-ai-global-report.pdf
8. **Cost of a Data Breach Report 2026**. IBM, 29 Jul 2026. *Sample of 602 organisations that had a breach.* https://assets.ibm.com/adobe/assets/urn:aaid:aem:21111142-1251-4369-86fb-57b82f5bb108/original/as/Cost%20of%20a%20Data%20Breach%20Report%202026.pdf
9. **AI use accelerates while governance and ROI lag, says new ISACA research**. ISACA, 5 May 2026. https://www.isaca.org/about-us/newsroom/press-releases/2026/ai-use-accelerates-while-governance-and-roi-lag-says-new-isaca-research
10. **Dratsch et al., Automation Bias in Mammography: The Impact of Artificial Intelligence BI-RADS Suggestions on Reader Performance, Radiology (press release)**. RSNA, 2 May 2023. https://www.rsna.org/media/press/i/2432
11. **Setting up and managing prepaid API billing**. OpenAI Help Center, accessed 24 Nov 2026. https://help.openai.com/en/articles/8264644-setting-up-and-managing-prepaid-api-billing

12. **Dev stunned by $82K Gemini bill after unknown API key thief goes to town**. The Register, 3 Mar 2026.
News report.
https://www.theregister.com/2026/03/03/gemini_api_key_82314_dollar_charge/

13. **Mid-Year Scam and Cybercrime Brief 2026**. Singapore Police Force, 26 Aug 2026.
https://isomer-user-content.by.gov.sg/537/a4a04eb7-0d38-4cd0-86da-cea81f214ff2/Mid-Year%20Scam%20and%20Cybercrime%20Brief%202026.pdf

14. **Singapore and Hong Kong Police Force recover over $670,000 in a scam**. Singapore Police Force, 7 Apr 2025.
https://www.police.gov.sg/media-hub/news/2025/04/20250407_singapore_and_hong_kong_police_force_recover_over_$670000_in_a_scam

15. **Global AI Adoption in 2025: A Widening Digital Divide (AI Diffusion Report, H2 2025)**. Microsoft AI Economy Institute, Jan 2026.
https://www.microsoft.com/en-us/research/wp-content/uploads/2026/01/Microsoft-AI-Diffusion-Report-2025-H2.pdf

16. **Anthropic Economic Index report: Uneven geographic and enterprise AI adoption**. Anthropic, 15 Sep 2025.
https://www.anthropic.com/research/anthropic-economic-index-september-2025-report

17. **Singapore Digital Economy Report 2025**. IMDA, 2025.
https://www.imda.gov.sg/assets/e77d879a-6b39-4de4-b024-5e0c6da0eff3

18. **UOB Business Outlook Study 2026, Singapore (H1)**. UOB, 2026.
https://www.uob.com.sg/web-resources/global/restricted/assets/pdfs/uob-business-outlook-study-2026-singapore-h1-infographics-en.pdf

19. **Singapore businesses lead in AI (media release)**. CPA Australia, undated.
Sample size not stated.
https://www.cpaaustralia.com.au/about-cpa-australia/media/media-releases/singapore-businesses-lead-in-ai

20. **Enterprises and corporate sector: latest data**. Department of Statistics Singapore, accessed 24 Nov 2026.
https://www.singstat.gov.sg/find-data/explore-data-themes/industry/enterprises-and-corporate-sector/latest-news-data

21. **Speech by Senior Minister of State Low Yen Ling on the motion on an economy of the future that works for all**. Ministry of Economy, Trade and Industry, 5 Aug 2026.
https://www.meti.gov.sg/newsroom/speech-by-senior-minister-of-state-low-yen-ling-on-motion-on--an-economy-of-the-future-that-works-for-all-/

22. **Written answer to PQ on distribution of SMEs**. Ministry of Manpower, 3 Mar 2026.
https://www.mom.gov.sg/newsroom/parliament-questions-and-replies/2026/0303-written-answer-to-pq-on-distribution-of-smes

23. **Written answer to PQ on breakdown of AI-adopting firms by firm size and outcomes**. Ministry of Manpower, 4 Aug 2026.
https://www.mom.gov.sg/newsroom/parliament-questions-and-replies/2026/0804-written-answer-to-pq-on-breakdown-of-ai-adopting-firms-by-firm-size-and-outcomes

24. **AI, Data Analytics and IoT Set to Shape the Next Phase of Digitalisation for Singapore Businesses (National Business Survey 2025, Digitalisation Supplement)**. Singapore Business Federation, 5 Feb 2026.
https://www.sbf.org.sg/newsroom/media/press-releases/detail/ai--data-analytics-and-iot-set-to-shape-the-next-phase-of-digitalisation-for-singapore-businesses

25. **Empowering SMEs in the age of AI: The 2026 OECD D4SME Survey**. OECD, Apr 2026.
The survey sample is not representative; the size-gap figures come from OECD official statistics.
https://www.oecd.org/content/dam/oecd/en/publications/reports/2026/04/empowering-smes-in-the-age-of-ai_7f58652c/bf5a9816-en.pdf

26. **AI use among US businesses (Business Trends and Outlook Survey)**. US Census Bureau, 26 May 2026.
https://census.gov/library/stories/2026/05/ai-use-businesses.html

27. **Use of artificial intelligence in enterprises**. Eurostat, Dec 2025.
https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Use_of_artificial_intelligence_in_enterprises

28. **Budget 2026 Statement, Section C: Harness AI as a Strategic Advantage**. Ministry of Finance, 12 Feb 2026.
https://www.singaporebudget.gov.sg/budget-speech/budget-statement/c-harness-ai-as-a-strategic-advantage

29. **Coinbase CEO explains why he fired engineers who didn't try AI immediately**. TechCrunch, 22 Aug 2025.
News report.
https://techcrunch.com/2025/08/22/coinbase-ceo-explains-why-he-fired-engineers-who-didnt-try-ai-immediately

30. **Enterprise AI adoption survey results (press release)**. WRITER and Workplace Intelligence, 7 Apr 2026.
Survey by an AI vendor; self-reported.
https://writer.com/blog/enterprise-ai-adoption-survey-results-press-release/

31. **The GenAI Divide: State of AI in Business 2025**. MIT NANDA, Jul 2025.
Copy hosted by AI News; MIT distributed the report through a request form.
https://www.artificialintelligence-news.com/wp-content/uploads/2025/08/ai_report_2025.pdf

32. **The story behind the bad AI stat that moved markets and misled millions (podcast episode)**. 80,000 Hours, Feb 2026.
https://80000hours.org/podcast/episodes/ai-workplace-mit-study/

33. **Press release on the 2026 study The Formula for Agentic AI Value**. BCG, 30 Sep 2026.
https://www.bcg.com/press/30september2026-ai-starting-to-pay-off-companies-generate-value

34. **Gartner Survey Finds Only 22% of Organizations Have Successfully Scaled AI Across Multiple Business Units**. Gartner, 1 Sep 2026.
https://www.gartner.com/en/newsroom/press-releases/gartner-survey-finds-only-22-percent-of-organizations-have-successfully-scaled-ai-across-multiple-business-units

35. **Gartner Predicts 30% of Generative AI Projects Will Be Abandoned After Proof of Concept By End of 2025**. Gartner, 29 Jul 2024.
https://www.gartner.com/en/newsroom/press-releases/2024-07-29-gartner-predicts-30-percent-of-generative-ai-projects-will-be-abandoned-after-proof-of-concept-by-end-of-2025

36. **Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027**. Gartner, 25 Jun 2025.
https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027

37. **CBA reverses AI-driven job cuts, admits error**. Information Age (ACS), 21 Aug 2025.
https://ia.acs.org.au/article/2025/cba-reverses-ai-driven-job-cuts--admits--error-.html
Also: Finance Sector Union statement, 21 Aug 2025.
https://www.fsunion.org.au/win-cba-backflips-on-customer-service-job-cuts-admits-they-got-it-wrong/

38. **Workslop research (Stanford Social Media Lab and BetterUp Labs), published in Harvard Business Review as AI-Generated Workslop Is Destroying Productivity**. BetterUp Labs, 22 Sep 2025.
https://www.betterup.com/workslop
Also: Harvard Business Review article.
https://hbr.org/2025/09/ai-generated-workslop-is-destroying-productivity

39. **Cisco's 2024 AI Readiness Index: Urgency Rises, Readiness Falls**. Cisco, 19 Nov 2024.
https://newsroom.cisco.com/c/r/newsroom/en/us/a/y2024/m11/cisco-2024-ai-readiness-index-urgency-rises-readiness-falls.html

40. **AI readiness stagnates in Singapore amid rapid market changes and expected impact**. CIO World Asia, 13 Nov 2024.
News report of Cisco's Singapore results.
https://cioworldasia.com/2024/11/13/ai-readiness-stagnates-in-singapore-amid-rapid-market-changes-and-expected-impact/

41. **Cisco AI Research: The Most AI-ready Companies Outpace Peers in the Race to Value**. Cisco, 14 Oct 2025.
https://newsroom.cisco.com/c/r/newsroom/en/us/a/y2025/m10/cisco-ai-research-the-most-ai-ready-companies-outpace-peers-in-the-race-to-value.html

42. **Lack of AI-Ready Data Puts AI Projects at Risk**. Gartner, 26 Feb 2025.
https://www.gartner.com/en/newsroom/press-releases/2025-02-26-lack-of-ai-ready-data-puts-ai-projects-at-risk

43. **Gartner Survey Finds 45% of Organizations With High AI Maturity Keep AI Projects Operational for at Least Three Years**. Gartner, 30 Jun 2025.
https://www.gartner.com/en/newsroom/press-releases/2025-06-30-gartner-survey-finds-forty-five-percent-of-organizations-with-high-artificial-intelligence-maturity-keep-artificial-intelligence-projects-operational-for-at-least-three-years

44. **AI Readiness Index (AIRI)**. AI Singapore, accessed 24 Nov 2026.
https://www.aisingapore.org/airi

45. **2026 Work Trend Index Annual Report: Agents, human agency, and the opportunity for every organization**. Microsoft WorkLab, 5 May 2026.
https://microsoft.com/en-us/worklab/work-trend-index/agents-human-agency-and-the-opportunity-for-every-organization

46. **Building Age-Friendly Jobs and Multi-Generational Workplaces (report of the Tripartite Workgroup on Senior Employment)**. Ministry of Manpower, 26 Oct 2026.
https://www.mom.gov.sg/-/media/mom/documents/press-releases/2026/annex---twg-se-report.pdf

47. **NTUC urges employers, Government to do more for older workers in AI era**. NTUC, 28 Jan 2026.
Cites MOM's Labour Force in Singapore 2025.
https://www.ntuc.org.sg/uportal/news/NTUC-Urges-Employers-Government-To-Do-More-For-Older-Workers-In-AI-Era/

48. **Workplace adoption of generative AI (NBER Digest summary of Bick, Blandin and Deming, The Rapid Adoption of Generative AI, Working Paper 32966)**. NBER, Dec 2024.
https://www.nber.org/digest/202412/workplace-adoption-generative-ai

49. **Workers' views of AI use in the workplace**. Pew Research Center, 25 Feb 2025. https://www.pewresearch.org/social-trends/2025/02/25/workers-views-of-ai-use-in-the-workplace/

50. **AI at Work Is Here. Now Comes the Hard Part (2024 Work Trend Index)**. Microsoft and LinkedIn, 8 May 2024. https://www.microsoft.com/en-us/worklab/work-trend-index/ai-at-work-is-here-now-comes-the-hard-part

51. **AI at Work: Why Strategy Matters More Than Tools**. BCG, Jun 2026. https://www.bcg.com/publications/2026/ai-at-work-why-strategy-matters-more-than-tools

52. **Brynjolfsson, Li and Raymond, Generative AI at Work, Quarterly Journal of Economics 140(2), 2025 (NBER Working Paper 31161)**. NBER, 2023, published 2025.
Figures as published in the journal version: 15% on average, 30% for less experienced and lower-skilled workers.
https://www.nber.org/papers/w31161
Also: Publisher's release.
https://www.eurekalert.org/news-releases/1071936

53. **Dell'Acqua et al., Navigating the Jagged Technological Frontier (Harvard Business School Working Paper 24-013)**. Harvard Business School, 22 Sep 2023. https://mitsloan.mit.edu/sites/default/files/2023-10/SSRN-id4573321.pdf

54. **AI at Work 2025: Momentum Builds, but Gaps Remain**. BCG, 26 Jun 2025. https://www.bcg.com/publications/2025/ai-at-work-momentum-builds-but-gaps-remain

55. **Microsoft and LinkedIn Release the 2024 Work Trend Index on the State of AI at Work in Singapore**. Microsoft Singapore, 9 May 2024. https://news.microsoft.com/en-sg/2024/05/09/microsoft-and-linkedin-release-the-2024-work-trend-index-on-the-state-of-ai-at-work-in-singapore/

56. **New SkillsFuture Level-Up Programme to Boost Mid-Career Upskilling and Reskilling**. Ministry of Education, 4 Mar 2024. https://www.moe.gov.sg/news/press-releases/20240304-new-skillsfuture-level-up-programme-to-boost-mid-career-upskilling-and-reskilling

57. **NTUC launches AI-Ready SG**. NTUC, 13 Feb 2026. https://www.ntuc.org.sg/uportal/news/NTUC-launches-AI-Ready-SG-to-AI-enabled-economy

58. **Half of Workers in Singapore Feel Uncomfortable Admitting AI Usage at Work, Says New Slack Workforce Index**. Salesforce, 13 Dec 2024. https://www.salesforce.com/ap/news/press-releases/2024/12/13/half-of-workers-in-singapore-feel-uncomfortable-admitting-ai-usage-at-work-says-new-slack-workforce-index/

59. **Cybersecurity training programs don't prevent employees from falling for phishing scams**. UC San Diego, 17 Sep 2025.
Study: Ho et al., Understanding the Efficacy of Phishing Training in Practice, IEEE Symposium on Security and Privacy 2025.
https://today.ucsd.edu/story/cybersecurity-training-programs-dont-prevent-employees-from-falling-for-phishing-scams

60. **Take-up rate of SkillsFuture-supported AI-related training course (parliamentary reply)**. Ministry of Education, 3 Feb 2026. https://www.moe.gov.sg/news/parliamentary-replies/20260203-take-up-rate-of-skillsfuture-supported-ai-related-training-course

61. **Outcomes and effectiveness of AI-related SkillsFuture programmes (parliamentary reply)**. Ministry of Education, 7 Apr 2026.
https://www.moe.gov.sg/news/parliamentary-replies/20260407-outcomes-and-effectiveness-of-ai-related-skillsfuture-programmes

62. **Harnessing AI as a Strategic Advantage (Budget 2026 highlights)**. Ministry of Finance, Feb 2026.
https://www.singaporebudget.gov.sg/highlights/harnessing-ai-as-a-strategic-advantage

63. **Cloud and Threat Report: 2026**. Netskope Threat Labs, 5 Jan 2026.
https://threatlabs.netskope.com/reports/2026/cloud-and-threat-report-2026

64. **GenAI Data Exposure: What GenAI Usage Is Really Costing Enterprises**. Harmonic Security, 31 Jul 2025.
https://www.harmonic.security/blog-posts/genai-data-exposure-report

65. **Report on Samsung's temporary ban on generative AI tools**. The Register, 2 May 2023.
News report.
https://www.theregister.com/2023/05/02/samsung_generative_ai_ban/

66. **Enterprise privacy at OpenAI**. OpenAI, updated 8 Jan 2026.
https://openai.com/enterprise-privacy/

67. **Updates to our consumer terms**. Anthropic, 28 Aug 2025.
https://www.anthropic.com/news/updates-to-our-consumer-terms

68. **Gemini API Additional Terms of Service**. Google, effective 23 Mar 2026.
https://aistudio.google.com/docs/terms

69. **Advisory Guidelines on Use of Personal Data in Generative AI**. PDPC, 20 Jul 2026.
https://www.pdpc.gov.sg/assets/143cb9d4-532e-4cca-9a77-bcc0415ca294

70. **CSA's Cyber Essentials and Cyber Trust marks expanded to include cloud security, artificial intelligence and operational technology**. Cyber Security Agency of Singapore, 15 Apr 2025.
https://www.csa.gov.sg/news-events/press-releases/csa-s-cyber-essentials-and-cyber-trust-marks-expanded-to-include-cloud-security--artificial-intelligence-and-operational-technology/

71. **An AI-powered coding tool wiped out a software company's database, then apologized for a catastrophic failure**. Fortune, 23 Jul 2025.
News report.
https://fortune.com/2025/07/23/ai-coding-tool-replit-wiped-database-called-it-a-catastrophic-failure/

72. **An AI agent deleted a company's entire database in 9 seconds, then wrote an apology**. Euronews, 28 Apr 2026.
News report.
https://euronews.com/next/2026/04/28/an-ai-agent-deleted-a-companys-entire-database-in-9-seconds-then-wrote-an-apology

73. **Report on Amazon's denial that its Kiro agent caused an AWS disruption**. The Register, 20 Feb 2026.
News report.
https://www.theregister.com/2026/02/20/amazon_denies_kiro_agentic_ai_behind_outage/

74. **Correcting the Financial Times report about AWS, Kiro, and AI**. Amazon, Feb 2026.
https://www.aboutamazon.com/md/news/aws/aws-service-outage-ai-bot-kiro

75. **Continuously hardening ChatGPT Atlas against prompt injection attacks**. OpenAI, 22 Dec 2025.
https://openai.com/index/hardening-atlas-against-prompt-injection/

76. **Research on prompt injection defenses for browser use**. Anthropic, 24 Nov 2025.
https://www.anthropic.com/research/prompt-injection-defenses

77. **Prompt injection is not SQL injection (it may be worse)**. UK National Cyber Security Centre, 8 Dec 2025.
https://www.ncsc.gov.uk/blog-post/prompt-injection-is-not-sql-injection

78. **CVE-2025-32711, Microsoft 365 Copilot information disclosure**. CVE Program, 11 Jun 2025.
https://www.cve.org/CVERecord?id=CVE-2025-32711

79. **Model AI Governance Framework for Agentic AI, Version 1.5**. IMDA, 20 May 2026.
https://www.imda.gov.sg/assets/d1e45a56-d898-4b9f-972e-8ae94be1a5a5.pdf

80. **Advisory on Cybersecurity Risks of OpenClaw (AD-2026-005)**. Cyber Security Agency of Singapore, 28 May 2026.
https://www.csa.gov.sg/alerts-and-advisories/advisories/ad-2026-005/

81. **LLM06:2025 Excessive Agency**. OWASP GenAI Security Project, 2025.
https://genai.owasp.org/llmrisk/llm062025-excessive-agency/

82. **Starter Kit for Testing LLM-Based Applications for Safety and Reliability**. IMDA, Jan 2026.
https://www.imda.gov.sg/-/media/imda/files/about/emerging-tech-and-research/artificial-intelligence/starter-kit-for-testing-llm-based-applications-for-safety-and-reliability.pdf

83. **BSI awards its 50th ISO/IEC 42001 certificate**. BSI, 2026.
https://www.bsigroup.com/en-GB/insights-and-media/media-centre/press-releases/2026/november/bsi-awards-its-50th-iso-iec-42001-certificate/

84. **Model Artificial Intelligence Governance Framework, Second Edition**. IMDA and PDPC, Jan 2020.
https://www.pdpc.gov.sg/-/media/files/pdpc/pdf-files/resource-for-organisation/ai/sgmodelaigovframework2.pdf

85. **Vaccaro, Almaatouq and Malone, When combinations of humans and AI are useful: A systematic review and meta-analysis, Nature Human Behaviour**. arXiv preprint, 28 Oct 2024.
https://arxiv.org/abs/2405.06087

86. **Goh et al., Large Language Model Influence on Diagnostic Reasoning: A Randomized Clinical Trial**. JAMA Network Open, Oct 2024.
https://jamanetwork.com/journals/jamanetworkopen/fullarticle/2825395

87. **Budzyn, Romanczyk et al., study of endoscopist detection rates after routine AI use, Lancet Gastroenterology and Hepatology (press release)**. The Lancet via EurekAlert, 12 Aug 2025.
Observational study.
https://www.eurekalert.org/news-releases/1094223

88. **Ben Green, The Flaws of Policies Requiring Human Oversight of Government Algorithms, Computer Law and Security Review 45**. arXiv preprint, 2022.
https://arxiv.org/abs/2109.05067

89. **Expanding on what we missed with sycophancy**. OpenAI, 2025.
https://openai.com/index/expanding-on-sycophancy/

90. **AI support in breast cancer screening: fewer missed cancer cases**. Lund University, 30 Jan 2026.
https://www.lunduniversity.lu.se/article/ai-support-breast-cancer-screening-fewer-missed-cancer-cases

91. **Regulation (EU) 2024/1689, the Artificial Intelligence Act**. Official Journal of the European Union, 12 Jul 2024.
http://data.europa.eu/eli/reg/2024/1689/oj

92. **Regulation (EU) 2026/1744 amending the Artificial Intelligence Act (Digital Omnibus on AI)**. Official Journal of the European Union, 24 Jul 2026.
http://data.europa.eu/eli/reg/2026/1744/oj

93. **BC Tribunal Finds Air Canada Liable for Inaccurate Advice Given by Website Chatbot (summary of Moffatt v. Air Canada, 2024 BCCRT 149)**. DWW LLP, 2024.
Law firm summary; the tribunal's own site blocked access.
https://dww.com/articles/bc-tribunal-finds-air-canada-liable-for-inaccurate-advice-given-by-website-chatbot
Also: WeirFoulds LLP, Moffatt v Air Canada: Bot Betrayal and Negligent Misrepresentation, 6 Mar 2024.
https://www.weirfoulds.com/moffatt-v-air-canada-bot-betrayal-and-negligent-misrepresentation

94. **AI Hallucination Cases database**. Damien Charlotin, updated 23 Nov 2026.
https://www.damiencharlotin.com/hallucinations/

95. **Registrar's Circular No. 1 of 2024: Guide on the Use of Generative AI Tools by Court Users**. Supreme Court of Singapore, 23 Sep 2024.
https://www.judiciary.gov.sg/news-and-resources/registrar's-circulars/circular-details/registrar's-circular-no.-1-2024-supreme-court

96. **Tajudin bin Gulam Rasul v Suriaya bte Haja Mohideen [2025] SGHCR 33**. High Court of Singapore, 29 Sep 2025.
https://www.elitigation.sg/gd/s/2025_SGHCR_33

97. **Tan Hai Peng Micheal v Tan Cheong Joo [2026] SGHC 49**. High Court of Singapore, Mar 2026.
https://www.elitigation.sg/gd/s/2026_SGHC_49

98. **Lawyers who doubled down and defended ChatGPT's fake cases must pay $5K, judge says (on Mata v. Avianca)**. ABA Journal, 26 Jun 2023.
https://abajournal.com/web/article/lawyers-who-doubled-down-and-defended-chatgpts-fake-cases-must-pay-5k-judge-says

99. **Deloitte refunds Aussie gov after AI fabrications slip into $440K welfare report**. The Register, 6 Oct 2025.
News report.
https://www.theregister.com/2025/10/06/deloitte_ai_report_australia/

100. **Media release on the refund figure given by the Department of Finance at Senate estimates**. Australian Greens, 9 Oct 2025.
https://greens.org.au/news/media-release/greens-slam-deloittes-unethical-behaviour-and-measly-refund-over-ai-report

101. **Report on fabricated citations in a Deloitte health workforce plan for Newfoundland and Labrador**. Fortune, 25 Nov 2025.
News report.
https://fortune.com/2025/11/25/deloitte-caught-fabricated-ai-generated-research-million-dollar-report-canada-government/

102. **Newfoundland and Labrador overhauls procurement after report errors thought to be AI**. The Canadian Press via CP24, 11 Mar 2026.
News report.
https://www.cp24.com/news/canada/2026/03/11/newfoundland-and-labrador-overhauls-procurement-after-report-errors-thought-to-be-ai/

103. **Singapore Launches New Model AI Governance Framework for Agentic AI**. Ministry of Digital Development and Information, 22 Jan 2026.
https://www.mddi.gov.sg/newsroom/singapore-launches-new-model-ai-governance-framework-for-agentic-ai--/

104. **MAS Sets Out Supervisory Expectations on Responsible AI Adoption by Financial Institutions**. Monetary Authority of Singapore, 7 Oct 2026.
https://www.mas.gov.sg/news/media-releases/2026/mas-sets-out-supervisory-expectations-on-responsible-ai-adoption-by-financial-institutions

105. **SB26-189 Automated Decision-Making Technology**. Colorado General Assembly, signed 14 May 2026.
https://leg.colorado.gov/bills/sb26-189

106. **Managing projects in the API platform**. OpenAI Help Center, accessed 24 Nov 2026.
https://help.openai.com/en/articles/9186755-managing-projects-in-the-api-platform

107. **How do I pay for my API usage?**. Claude Help Center, accessed 24 Nov 2026.
https://support.claude.com/en/articles/8977456-how-do-i-pay-for-my-api-usage

108. **Rate limits**. Claude Platform documentation, accessed 24 Nov 2026.
https://platform.claude.com/docs/en/api/rate-limits

109. **Billing (Gemini API documentation)**. Google AI Studio, accessed 24 Nov 2026.
https://aistudio.google.com/docs/billing

110. **Innovate on your own business terms: Google Cloud spend caps now generally available**. Google Cloud Blog, 31 Oct 2026.
https://cloud.google.com/blog/topics/cost-management/google-cloud-spend-caps-now-generally-available
Also: Manage spend cap budgets (documentation).
https://docs.cloud.google.com/billing/docs/how-to/budgets-spend-caps

111. **Google reimburses Register sources who were victims of API fraud**. The Register, 15 May 2026.
News report.
https://www.theregister.com/devops/2026/05/15/google-reimburses-register-sources-who-were-victims-of-api-fraud/5241429

112. **Google API Keys Weren't Secrets. But then Gemini Changed the Rules.**. Truffle Security, 25 Feb 2026.
https://trufflesecurity.com/blog/google-api-keys-werent-secrets-but-then-gemini-changed-the-rules

113. **LLMjacking: Stolen Cloud Credentials Used in New AI Attack**. Sysdig, 6 May 2024.
https://sysdig.com/blog/llmjacking-stolen-cloud-credentials-used-in-new-ai-attack/

114. **GTIG AI Threat Tracker: Agentic Threat Activity Crosses a Significant Threshold**. Google Threat Intelligence Group, 17 Nov 2026.
https://cloud.google.com/blog/topics/threat-intelligence/agentic-threat-activity-crosses-a-significant-threshold

115. **LLM10:2025 Unbounded Consumption**. OWASP GenAI Security Project, 2025.
https://genai.owasp.org/llmrisk/llm102025-unbounded-consumption/

116. **Flexera Finds Cloud Value is Rising While AI Waste Grows**. Flexera, 18 Mar 2026.
https://www.flexera.com/about-us/press-center/flexera-finds-cloud-value-is-rising-while-ai-waste-grows

117. **Summary of the Amazon DynamoDB Service Disruption in the Northern Virginia (US-EAST-1) Region**. Amazon Web Services, Oct 2025.
https://aws.amazon.com/message/101925/

118. **Post Incident Review YKYN-BWZ, Azure Front Door**. Microsoft Azure, Oct 2025.
https://azure.status.microsoft/en-us/status/history/?trackingId=YKYN-BWZ

119. **Cloudflare outage on November 18, 2025**. Cloudflare, Nov 2025.
https://blog.cloudflare.com/18-november-2025-outage/

120. **Cloudflare outage on December 5, 2025**. Cloudflare, Dec 2025.
https://blog.cloudflare.com/5-december-2025-outage/

121. **Google Cloud Service Health incident report, 12 to 13 June 2025**. Google Cloud, Jun 2025.
https://status.cloud.google.com/incidents/ow5i3PPK96RduMcb1SsW

122. **Incident write-up, 9 to 10 June 2025**. OpenAI Status, Jun 2025.
https://status.openai.com/incidents/01JXCAW3K3JAE0EP56AEZ7CBG3/write-up

123. **It's not just you; ChatGPT, Claude, and Grok were all down in confirmed outages**. 9to5Google, 3 Sep 2026.
News report.
https://9to5google.com/2026/09/03/chatgpt-claude-grok-outages/

124. **Claude Status incident history feed**. Anthropic, read 24 Nov 2026.
Count is our own, from the feed.
https://status.claude.com/history.rss

125. **A postmortem of three recent issues**. Anthropic, 17 Sep 2025.
https://www.anthropic.com/engineering/a-postmortem-of-three-recent-issues

126. **Chen, Zaharia and Zou, How is ChatGPT's behavior changing over time?**. arXiv 2307.09009, Jul 2023.
https://arxiv.org/abs/2307.09009

127. **Suddenly deprecating old models users depended on a mistake, admits OpenAI's Altman**. The Register, 11 Aug 2025.
News report.
https://www.theregister.com/2025/08/11/openai_tweaks_gpt5_after_user/

128. **Retiring GPT-4o and older models**. OpenAI, 29 Jan 2026.
https://openai.com/index/retiring-gpt-4o-and-older-models

129. **Deprecations**. OpenAI API documentation, accessed 24 Nov 2026.
https://developers.openai.com/api/docs/deprecations

130. **Singapore Cyber Landscape 2025/2026**. Cyber Security Agency of Singapore, 30 Jun 2026.
https://www.csa.gov.sg/resources/publications/singapore-cyber-landscape-2025-2026/

131. **Report on the backdoored postmark-mcp package**. The Register, 29 Sep 2025.
News report.
https://www.theregister.com/2025/09/29/postmark_mcp_server_code_hijacked/

132. **Security Bulletin AWS-2025-015 (Amazon Q Developer extension for VS Code)**. Amazon Web Services, 23 Jul 2025.
https://aws.amazon.com/security/security-bulletins/AWS-2025-015/

133. **Open models lag state-of-the-art closed models by 4 months**. Epoch AI, 29 May 2026.
https://epoch.ai/data-insights/open-closed-eci-gap

134. **AI Index Report 2026, Chapter 2: Technical Performance**. Stanford HAI, Apr 2026.
https://hai.stanford.edu/assets/files/ai_index_report_2026_chapter_2_technical.pdf

135. **Introducing gpt-oss**. OpenAI, 5 Aug 2025.
https://openai.com/index/introducing-gpt-oss/

136. **Gemma 4: Byte for byte, the most capable open models**. Google, 2 Apr 2026.
https://blog.google/innovation-and-ai/technology/developers-tools/gemma-4/

137. **SEA-LION repository and model list**. AI Singapore, accessed 24 Nov 2026.
https://github.com/aisingapore/sealion

138. **2/23/2026 Price Change Announcement (DGX Spark)**. NVIDIA Developer Forums, 25 Feb 2026.
https://forums.developer.nvidia.com/t/2-23-2026-price-change-announcement/361713
Also: DGX Spark product page.
https://www.nvidia.com/en-us/products/workstations/dgx-spark/

139. **NVIDIA Blackwell GeForce RTX 50 Series Opens New World of AI Computer Graphics**. NVIDIA, 6 Jan 2025.
https://nvidianews.nvidia.com/news/nvidia-blackwell-geforce-rtx-50-series-opens-new-world-of-ai-computer-graphics

140. **Apple unveils new Mac Studio, the most powerful Mac ever**. Apple, 5 Mar 2025.
https://www.apple.com/newsroom/2025/03/apple-unveils-new-mac-studio-the-most-powerful-mac-ever/

141. **Mac Studio 512GB RAM Option Disappears Amid Global DRAM Shortage**. MacRumors, 5 Mar 2026.
https://www.macrumors.com/2026/03/05/mac-studio-no-512gb-ram-upgrade

142. **Apple will now sell you a Mac Studio with 512GB of RAM, but it'll cost you as much as a good-sized family car**. TechRadar, 3 Nov 2026.
https://techradar.com/pro/apple-will-now-sell-you-a-mac-studio-with-512gb-of-ram-but-itll-cost-you-as-much-as-a-good-sized-family-car

143. **Tariff information**. SP Group, accessed 24 Nov 2026.
https://www.spgroup.com.sg/our-services/utilities/tariff-information

144. **The 2025 AI Index Report**. Stanford HAI, Apr 2025.
https://hai.stanford.edu/ai-index/2025-ai-index-report

145. **Detecting Exposed LLM Servers: A Shodan Case Study on Ollama**. Cisco Talos, 1 Sep 2025.
https://blogs.cisco.com/security/detecting-exposed-llm-servers-shodan-case-study-on-ollama

146. **175,000 Exposed Ollama Hosts Could Enable LLM Abuse**. SecurityWeek, 30 Jan 2026.
News report of SentinelOne and Censys research.
https://www.securityweek.com/175000-exposed-ollama-hosts-could-enable-llm-abuse/

147. **Bleeding Llama: A Critical Memory Leak in the World's Most Popular Local AI Platform**. Cyera Research, 5 May 2026.
https://www.cyera.com/blog/bleeding-llama-a-critical-memory-leak-in-the-worlds-most-popular-local-ai-platform

148. **FAQ**. Ollama documentation, accessed 24 Nov 2026.
https://docs.ollama.com/faq

149. **Open Chroma Databases: A New Attack Surface for AI Apps**. UpGuard, updated 1 Dec 2025.
https://www.upguard.com/blog/open-chroma-databases-ai-attack-surface

150. **Advisory Guidelines on the PDPA for Selected Topics, Chapter 8: Cloud Services**. PDPC, revised 9 Oct 2019.
https://www.pdpc.gov.sg/-/media/Files/PDPC/PDF-Files/Advisory-Guidelines/AG-on-Selected-Topics/Chapter-8-9-Oct-2019.pdf

151. **Data, privacy, and security for Foundry Models sold by Azure**. Microsoft Learn, accessed 24 Nov 2026.
https://learn.microsoft.com/en-us/azure/ai-foundry/responsible-ai/openai/data-privacy

152. **Fact sheet: four data centre proposals selected as part of the second data centre call for application**. EDB and IMDA, 21 Aug 2026.
https://www.edb.gov.sg/content/dam/edb-en/about-edb/media-releases/news/four-data-centre-proposals-selected-as-part-of-second-data-centre/Fact%20Sheet%20-Four%20data%20centre%20proposals%20selected%20as%20part%20of%20second%20data%20centre.pdf

153. **Annual Scam and Cybercrime Brief 2025**. Singapore Police Force, 25 Feb 2026.
https://www.police.gov.sg/-/media/SPF/Advisories/Scams/Annual-Scam-and-Cybercrime-Brief-2025.pdf

154. **Advisory on scams involving impersonation of senior government officials including the Secretary to the Cabinet**. Singapore Police Force, 14 May 2026.
https://www.police.gov.sg/Media-Hub/News/2026/05/20260514_advisory_on_scams_involving_impersonation_of_snr_govt_officials_incl_secy_cabinet

155. **Report on Arup confirming it was the victim of a US$25 million deepfake scam**. Fortune, 17 May 2024.
News report.
https://fortune.com/europe/2024/05/17/arup-deepfake-fraud-scam-victim-hong-kong-25-million-cfo/

156. **Cost of a Data Breach Report 2025 (press release)**. IBM, 30 Jul 2025.
The split by attack type is from the full report.
https://newsroom.ibm.com/2025-07-30-ibm-report-13-of-organizations-reported-breaches-of-ai-models-or-applications,-97-of-which-reported-lacking-proper-ai-access-controls
Also: Full report (copy hosted by Baker Donelson), figures 2 and 33.
https://www.bakerdonelson.com/webfiles/Publications/20250822_Cost-of-a-Data-Breach-Report-2025.pdf

157. **Detecting and countering misuse of AI: August 2025**. Anthropic, 27 Aug 2025.
https://www.anthropic.com/news/detecting-countering-misuse-aug-2025

158. **Disrupting the first reported AI-orchestrated cyber espionage campaign**. Anthropic, Nov 2025.
Anthropic's own assessment; some researchers have questioned the autonomy estimate.
https://assets.anthropic.com/m/ec212e6566a0d47/original/Disrupting-the-first-reported-AI-orchestrated-cyber-espionage-campaign.pdf

159. **GTIG AI Threat Tracker: Advances in Threat Actor Usage of AI Tools**. Google Threat Intelligence Group, Nov 2025.
https://cloud.google.com/blog/topics/threat-intelligence/threat-actor-usage-of-ai-tools

160. **GTIG AI Threat Tracker: Adversaries Leverage AI for Vulnerability Exploitation, Augmented Operations, and Initial Access**. Google Threat Intelligence Group, 12 May 2026.
https://cloud.google.com/blog/topics/threat-intelligence/ai-vulnerability-exploitation-initial-access

161. **Heiding, Lermen, Kao, Schneier and Vishwanath, Evaluating Large Language Models' Capability to Launch Fully Automated Spear Phishing Campaigns**. arXiv 2412.00586, Nov 2024.
101 participants.
https://arxiv.org/abs/2412.00586

162. **Jasper Sleet: North Korean remote IT workers' evolving tactics to infiltrate organizations**. Microsoft Threat Intelligence, 30 Jun 2025.
https://www.microsoft.com/en-us/security/blog/2025/06/30/jasper-sleet-north-korean-remote-it-workers-evolving-tactics-to-infiltrate-organizations/

163. **Justice Department Announces Coordinated, Nationwide Actions to Combat North Korean Remote Information Technology Workers' Illicit Revenue Generation Schemes**. US Department of Justice, 30 Jun 2025.
https://www.justice.gov/opa/pr/justice-department-announces-coordinated-nationwide-actions-combat-north-korean-remote

164. **Spracklen et al., We Have a Package for You! A Comprehensive Analysis of Package Hallucinations by Code Generating LLMs (USENIX Security 2025)**. arXiv 2406.10279, 2025.
https://arxiv.org/abs/2406.10279v3

165. **A small number of samples can poison LLMs of any size**. Anthropic, UK AI Security Institute and The Alan Turing Institute, 9 Oct 2025.
https://www.anthropic.com/research/small-samples-poison

166. **Joint Advisory on Scams Involving Digital Manipulation**. Singapore Police Force, MAS and CSA, 12 Mar 2025.
https://www.police.gov.sg/media-hub/news/2025/03/20250312_joint_advisory_on_scams_involving_digital_manipulation

167. **Report on the voice-cloned CEO call stopped by a Ferrari executive**. Fortune (Bloomberg), 27 Jul 2024.
News report.
https://fortune.com/2024/07/27/ferrari-deepfake-attempt-scammer-security-question-ceo-benedetto-vigna-cybersecurity-ai/

168. **Attempted Audio Deepfake Call Targets LastPass Employee**. LastPass, 10 Apr 2024.
https://blog.lastpass.com/posts/attempted-audio-deepfake-call-targets-lastpass-employee

169. **Criminals Use Generative Artificial Intelligence to Facilitate Financial Fraud (PSA I-120324-PSA)**. FBI Internet Crime Complaint Center, 3 Dec 2024.
https://www.ic3.gov/PSA/2024/PSA241203

170. **Advisory on Detecting and Responding to Deepfake Scams (AD-2024-006)**. Cyber Security Agency of Singapore, 22 Mar 2024.
https://www.csa.gov.sg/alerts-and-advisories/advisories/ad-2024-006

171. **MAS and IMDA announce implementation of Shared Responsibility Framework from 16 December 2024**. Monetary Authority of Singapore, 24 Oct 2024.
The reading that it covers unauthorised transactions comes from law firm summaries; check the guidelines for your case.
https://www.mas.gov.sg/news/media-releases/2024/mas-and-imda-announce-implementation-of-shared-responsibility-framework-from-16-december-2024

172. **Implementing Phishing-Resistant MFA (fact sheet)**. CISA, Oct 2022.
https://www.cisa.gov/sites/default/files/publications/fact-sheet-implementing-phishing-resistant-mfa-508c.pdf

173. **National AI Strategy 2.0**. Smart Nation Singapore, 4 Dec 2023.
https://file.go.gov.sg/nais2023.pdf

174. **Update to Singapore's National AI Strategy: Refreshed Priorities to Harness AI for the Public Good (factsheet)**. Ministry of Digital Development and Information, 20 May 2026.
https://mddi.gov.sg/newsroom/update-to-singapore-s-national-ai-strategy--refreshed-priorities-to-harness-ai-for-the-public-good-factsheet

175. **National AI Impact Programme: Empowering Enterprises and Workers to Transform with AI (factsheet)**. Ministry of Digital Development and Information, 2 Mar 2026.
https://www.mddi.gov.sg/newsroom/national-ai-impact-programme--empowering-enterprises-and-workers-to-transform-with-ai/

176. **EDGE Grant**. Enterprise Singapore, accessed 24 Nov 2026.
https://www.enterprisesg.gov.sg/financial-support/edge-grant

177. **EDGE Grant FAQs**. Enterprise Singapore, accessed 24 Nov 2026.
https://www.enterprisesg.gov.sg/resources/all-faqs/edge-grant

178. **Enterprise Compute Initiative: programme details**. Digital Industry Singapore, accessed 24 Nov 2026.
https://www.disg.gov.sg/enterprise-compute-initiative/

179. **S'pore unveils AI measures to spur new products**. The Star (Malaysia), 29 Oct 2026.
News report.
https://www.thestar.com.my/business/business-news/2026/10/29/spore-unveils-ai-measures-to-spur-new-products

180. **Model AI Governance Framework for Generative AI**. AI Verify Foundation and IMDA, 30 May 2024.
https://aiverifyfoundation.sg/resources/mgf-gen-ai/

181. **Guidelines and Companion Guide on Securing AI Systems**. Cyber Security Agency of Singapore, 15 Oct 2024.
https://www.csa.gov.sg/resources/publications/guidelines-and-companion-guide-on-securing-ai-systems

182. **CSA Releases An Addendum To Support System Owners In Securing Agentic AI System**. Cyber Security Agency of Singapore, 22 Oct 2025.
https://www.csa.gov.sg/news-events/press-releases/csa-releases-an-addendum-to-support-system-owners-in-securing-agentic-ai-system/

183. **Online Safety Commission and Online Safety (Relief and Accountability) Act 2025 to Start on 29 June 2026**. Ministry of Law and MDDI, 14 May 2026.
https://www.mlaw.gov.sg/online-safety-commission-and-online-safety-relief-and-accountability-act-2025-to-start-on-29-june-2026/

184. **New Legal Measures to Uphold Integrity of Online Advertising During Elections**. Ministry of Digital Development and Information, 9 Sep 2024.
https://www.mddi.gov.sg/newsroom/new-legal-measures-uphold-integrity-online-advertising-elections/

185. **Executive Order 14365, Ensuring a National Policy Framework for Artificial Intelligence**. The White House, 11 Dec 2025.
https://www.whitehouse.gov/presidential-actions/2025/12/eliminating-state-law-obstruction-of-national-artificial-intelligence-policy/

186. **President Donald J. Trump Unveils National AI Legislative Framework**. The White House, 20 Mar 2026.
https://www.whitehouse.gov/releases/2026/03/president-donald-j-trump-unveils-national-ai-legislative-framework/

187. **Measures for Labeling AI-Generated and Synthetic Content (notice, in Chinese)**. Cyberspace Administration of China, 14 Mar 2025.
https://www.cac.gov.cn/2025-03/14/c_1743654685896173.htm

188. **Tribunale di Roma judgment annulling Garante order no. 755 of 2 November 2024 (in Italian)**. Tribunale di Roma, 18 Mar 2026.
University-hosted copy of the judgment.
https://dei.web.uniroma1.it/sites/default/files/allegati/2026-05/Trib_Roma_OpenAI_Garante_2026.pdf

189. **Global Workforce Hopes and Fears Survey 2026: Singapore edition**. PwC Singapore, 2026.
https://www.pwc.com/sg/en/publications/hopes-and-fears.html

190. **Klarna AI assistant handles two-thirds of customer service chats in its first month**. Klarna, 27 Feb 2024.
https://www.klarna.com/international/press/klarna-ai-assistant-handles-two-thirds-of-customer-service-chats-in-its-first-month/

191. **Klarna CEO reverses course by hiring more humans, not AI**. Entrepreneur, 9 May 2025.
News report of a Bloomberg interview.
https://www.entrepreneur.com/business-news/klarna-ceo-reverses-course-by-hiring-more-humans-not-ai/491396

192. **Duolingo CEO says controversial AI memo was misunderstood**. TechCrunch, 17 Aug 2025.
News report.
https://techcrunch.com/2025/08/17/duolingo-ceo-says-controversial-ai-memo-was-misunderstood