



10×

from AI promise to ROI

Asking the right questions, then doing the right things right

Return on Customer • Return on Employee • Return on Future
Cheaper • Better • Faster • Safer

Contents

- 03** **From promise to payoff**
The AI journey and the 10× thesis
- 05** **Do the right things**
Most leaders are not asking the right questions
The right questions
- 08** **The AI Paradox**
Good strategy. Bold investment. Stalled outcomes.
The five U's of failed AI
- 10** **Do the things right**
Cheaper – economically optimised
Better – reliable with maximised return
Faster – ahead of competition
Safer – minimised risks and harm
- 20** **The workforce question**
AI should amplify human intelligence – not replace it
The people side, brought forward
- 24** **Crossing the gap**

From promise to payoff

Last year, we declared the arrival of a new horizon: AI was no longer at the edge of the enterprise but moving decisively to its core.

That horizon has now arrived. Investment has accelerated, ambition has grown, and AI has become a boardroom priority across industries. Yet for many organisations, the payoff remains elusive. Despite the scale of investment, too many AI initiatives are still trapped in experimentation, with a MIT study reporting that 95% of organisations are getting zero return. For most enterprises, AI's promise has yet to translate into measurable returns.

At NCS, we believe 2026 marks the moment to move from promise to payoff. The opportunity is no longer simply to adopt AI, but to convert AI ambition into measurable, repeatable, and scalable business value. Organisations that make the right choices and execute them with excellence and discipline can unlock 10x returns from Intelligentisation – a structured, holistic approach to embedding AI into the core of business processes, government workflows and human experiences.

This white paper is NCS' practitioner-led response. Drawing on extensive AI engagements across Singapore, Hong Kong, Australia, India, and Greater China, it sets out the strategy and playbook required to turn AI promise into payoff: doing the right things by asking the right questions through the 3R framework, then doing the right things right – cheaper, better, faster, and safer.

3R // CBFS

The AI Journey

2023

Inquiry

Challenge the conventional

Look past the hype. Challenge assumptions and redefine what success looks like.

2024

Intent

Change the game

Wield AI differently – anchored in the AI+DR Matrix and digital resilience.

2025

Integration

Enter the new horizon

The 3Is – Intelligentisation, Internationalisation, Inspiration – at the core.

2026

Returns

Turn promise into payoff

Right things × things done right.
The practitioner's playbook for an outsized return.

Figure 1. A timeline of NCS' AI journey

Enterprise AI has entered a new phase, where early enthusiasm is being tested by the realities of execution. This has exposed an “AI Paradox” – organisations are investing heavily and setting bold strategies, yet many initiatives remain stalled before delivering meaningful returns; creating a significant gap between the promise of 10x ROI and the actual value realised.



Figure 2. A snapshot of AI Failures
Source: MIT, 2025; WRITER, 2026; Gartner, 2025; EY, 2025

Do the right things

Most leaders are not asking the right questions

Before any model, dataset or platform, AI ROI is decided by the question a leader chooses to put on the table. Ask the wrong one and the organisation walks, predictably, into the lose-lose quadrant.

Three suboptimal questions

How much cost can we save with AI?

How many people can we cut with AI?

What problem can we solve using this AI solution?

Each is framed inside out starting with the tool, the headcount or the cost line. The guaranteed outcome: low employee trust, low ROI, no durable advantage. Short-term gains paid for in talent flight and brand damage.

The quality of the questions you ask shapes which quadrant you end up in.

The ultimate destination of an enterprise AI strategy is decided long before choosing a model or platform, because initial ROI is driven by the baseline questions a leader puts on the table.

Too many organisations are now stuck in the lose-lose quadrant – not because AI lacks potential, but because their use cases are shaped by suboptimal questions and undermined by poor execution.

This paper reveals the pitfalls to avoid and charts a path to the win-win quadrant, delivering stronger returns, better workforce outcomes and durable advantage.




Figure 3. AI Returns – Workforce Growth Matrix

“A suboptimal question, asked with confidence and capital, only buys you a faster route to the lose-lose quadrant.”

The right questions

Here are better questions to ask using the 3R framework

Three strategic questions that pull AI investment out of cost-cutting theatre and into the growth accelerating and revenue expanding win-win quadrant – by transforming customer value, multiplying the workforce, and building tomorrow's advantage.



01 Create new value for customers.

Ask: How do we create new experiences that attract and retain customers with AI?

- Make every customer feel like they have a dedicated expert.
- Invent entirely new experiences – not enhance old journeys.
- Detect customer risk before it is too late.

02 Use AI as a force multiplier.

Ask: How do we help every employee deliver more value with less friction?

- Remove low-value work people don't want to do.
- Redesign roles and reskill employees to do higher value work; scaling their capacity and capabilities exponentially.
- Future-proof capability before skills become operating gaps.

03 Build the capabilities of an industry-of-the-future organisation.

Ask: How do we build the game-changing capabilities that become tomorrow's competitive advantage?

- **Strategic optionality:** Create new platforms that open access to new markets and business models.
- **Adaptive resilience:** Keep operating effectively through market disruption, and uncertainty.
- **Compounding advantage:** Build assets and capabilities that allow organisations to stay ahead of competitors.


The AI Paradox

Good strategy. Bold investment.
Stalled outcomes.

The paradox is this: organisations with sound business strategy and serious AI budgets expect proportionate AI success. Most do not get it. The gap is rarely the technology. It is everything that surrounds it.

The great AI divide

Good business strategy and high AI investment should produce proportionate AI success. In practice, the line bends sharply downwards – the paradox most leaders feel but cannot name.



95%

of organisations chase AI use cases and ROI without a solid AI foundation – setting themselves up for stalled outcomes before implementation even begins.

Source: MIT Report - State of AI in Business, 2025

Chasing AI use cases and ROI without a strong and solid AI foundation is a playbook for failure.

Figure 4. The AI Paradox – expected vs. realised return on AI investment

The five U's of failed AI

Across all the failed AI programmes, the same five root causes recur. Each one is invisible in slide decks and devastating in production. Together they explain why most AI programmes never become an AI return.

Unclear cost

Token spend is treated as an engineering detail. In reality, it's a P&L exposure that compounds while you sleep. Uber burned its entire 2026 AI budget in four months on ungoverned Claude Code adoption.

Unchanged business processes

AI bolted onto a broken workflow doesn't fix it, but amplifies its flaws. Local optimisation, automated waste and a human re-work tax turns 10x ambitions into a mere 10% gain.

Unready data

More data doesn't always mean smarter AI. Fragmented, context-less, contradictory data accelerates failure. Air Canada's chatbot hallucinated a refund policy that did not exist – a tribunal-ordered precedent on AI liability.

Ungoverned agent development

Bolt-on or black-box. Either humans are all over the loop reconciling AI output, or agents run open-loop, misinterpreting intent and cascading bad decisions downstream. According to Gartner, 40% of Agentic AI projects will be cancelled by 2027.

Unknown security & safety risks

Frontier AI has re-written corporate cyber risk. Frontier models reportedly find thousands of zero-day vulnerabilities in days. A policy on the shelf is not a control in production.

The diagnosis

The right thing, done wrong, is still the wrong thing.

The five U's rarely appear in strategy decks, but they surface quickly in production. Closing the AI Paradox requires both the right strategic intent and the operational experience and expertise to execute it.

Figure 5. The five U's – root causes recurring across failed AI programmes



Do the things right

Closing the AI Paradox is a discipline, not a slogan. The NCS AI playbook – forged from extensive AI implementations – is built around four operating commitments that translate ambition into measurable return.



Figure 6. The four pillars of the NCS AI playbook

From theory to practice

The AI Playbook is not a framework on a wall. Each pillar is enabled by specific NCS methodologies and tools used in extensive Intelligentisation implementations (E.g., NCS Drive+ methodology for AI use case ideation and prioritisation, Sunshine.AI products and platforms for cheaper, better, faster and safer AI implementations).

Conventional

**“We need more
AI use cases.”**

NCS observation

**You need the right use
cases, executed with
excellence and
discipline across all
four pillars.**

The model that wins benchmarks rarely wins on your bill.

As AI moves from experimentation to enterprise-wide adoption, cost becomes a strategic consideration. The organisations achieving sustainable returns are not simply deploying more AI, but using it more efficiently – matching the right models to the right tasks, optimising token usage effectively, and maximising value from every investment.

Two levers, in tandem

Cost discipline in production AI has two orders of optimisation. First, **optimise cost per token** – route each task to a model sized for it, instead of defaulting to the top of the leaderboard. Second, **optimise tokens per task** – govern how agents consume context and how humans allocate spend.

Here's how it leads to real outcomes

~82%

Cost savings vs frontier models

3–10×

Faster response times



Higher task-matched accuracy

Source: NCS internal benchmark across project deployments

A multi-step model selection method

Project-level evaluation

Narrow hundreds of LLMs to the most optimal based on data residency, sovereignty, compliance, context window, latency and multi-modality.


Use case evaluation

Sunshine.guardian generates thousands of synthetic conversations per persona to score hallucination, accuracy and ethicality.

Continuous monitoring

Sunshine.operations tracks EOL/EOS, eval drifts and newly released models.

Where the bill comes from



Indicative; relative to unoptimised frontier baseline

Figure 7. Two compounding sources of token cost and the gain from disciplined optimisation

Cautionary tale

The hidden cost of rapid AI adoption

Uber chased AI adoption with leaderboards, burned its 2026 AI coding-tool budget in four months, then struggled to prove that the spending produced better customer features. Rapid AI adoption without cost governance and optimisation can turn "innovation" into runaway costs.

Two common mistakes affecting cost

Mistake 01

Underestimating token cost

Teams may treat AI usage as cheap experimentation, but high-volume prompts, long context windows, and repeated agent runs can quickly turn token consumption into a major budget overrun.

Mistake 02

Rewarding usage, not outcomes

Internal rankings or adoption incentives can drive "tokenmaxxing," where teams spend more tokens to look productive rather than solve business problems.

Unready data and unchanged business processes do not slow down AI. It accelerates failure.

As AI becomes embedded in business processes, accuracy and reliability are no longer technical concerns alone – they are business-critical. Strong outcomes depend not only on the model itself, but on the quality of data, workflows and governance that support it.

Two responses to an AI workflow

The more deeply AI is embedded in a redesigned workflow, the greater the returns.


Bolt-on



ROI

10%

Re-imagine



10X

Figure 8. An illustration of how when AI is embedded into workflows, the returns generated could be much higher

The grounded-retrieval rule

Missing data is not fixed by a better model. The only fix is governed, grounded retrieval over a unified data layer. Fragmented data foundations produce contradictory answers, context-less decisions and eventually, liability.

Organisations need to build trusted AI data foundations by collecting relevant data across multiple systems, cleaning and standardising through a unified data platform, then contextualising with business rules and metadata so AI delivers accurate, consistent, and trustworthy outputs.

Our solution to the data problem?

Sunshine.core is an AI Foundational Platform that provides reusable building blocks such as common knowledge base management, data semantic layer and memory management, helping teams better manage data in secured environments.

Cautionary tale

The Air Canada precedent

A GenAI chatbot, ungrounded in current policy, invented a bereavement refund policy and stated it confidently. A tribunal ruled the airline liable: a landmark precedent that unready data does not just produce wrong answers, it produces liability.

Source: BC Civil Resolution Tribunal, 2024

Two common mistakes affecting quality

Mistake 01

Unreliable source data

Training or grounding AI on inconsistent, outdated, or poorly governed information can cause it to give confident but inaccurate answers that damage user trust.

Mistake 02

The bolt-on fallacy

AI is patched onto a broken workflow. People review, correct and reconcile and re-explain its output – automation becomes another task. 10% ROI instead of 10x.

Conventional belief

“AI on a workflow improves it.”

NCS observation

AI on a broken process doesn't fix it – it amplifies it.

Rebuilding the same foundations is not engineering – it is waste.

The challenge is no longer proving that AI works, but scaling its impact across the organisation. Rather than rebuilding capabilities for every use case, leading organisations are creating reusable foundations that accelerate delivery, reduce duplication and bring value to market faster than competitors.

Reusable cognitive superpowers

NCS' reusable assets are built around **repeatable cognitive superpowers** – patterns of intelligence common across industries. The abstraction sits at the right level: high enough to travel between projects, specific enough to deliver value on day one.

For example, NCS' Sunshine.kaisense is an advanced video AI platform that transforms real-time video monitoring and analysis. It connects seamlessly to large-scale camera and device networks, enabling users to search, analyse and summarise live and recorded visual content using natural language.



Video AI
Eyes everywhere
Never misses an event.



Conversational AI
Fluent, never angry
Multilingual customer service.



Agentic AI
An orchestrator
Coordinates work end-to-end.



Knowledge AI
A mind that forgets nothing
Grounded retrieval over corpora.



Physical AI
Intelligence in motion
Perceives, decides and acts in the real world.



Robotics
A tireless worker
24/7 zero fatigue.

Figure 9. Six NCS cognitive superpowers – reusable, portable, ownership-modelled

Three asset disciplines

01 Reusability

Built around a cognitive pattern, not a project

One use case becomes ten without retrofits or workarounds.

02 Portability by design

Platform-agnostic by architecture

Move with the model and vendor landscape; no platform lock-in.

03 Operating-model driven

Run like products

Ownership, versioning, certification and lifecycle – built in.

~75%

of every AI project rebuilds foundations already solved elsewhere in the same enterprise

Source: Boston Consulting Group, 2025

Cautionary tale

AI islands – not quite paradise

Rebuilding every AI use case from scratch creates costly “AI islands.” In an April 2025 report, BCG found 75% of tech leaders fear silent AI failure, 81% cite high new-use-case build costs, and 77% worry about governance risks – while duplication, idle infrastructure, compliance gaps, and fragmented experiences quietly compound.

Two common mistakes affecting quality

Mistake 01

Unreliable source data

Teams optimise small-scale wins in silos, rebuilding every use case from scratch – creating duplicated effort, disjointed user experiences, security vulnerabilities, compliance risks, and scalability problems that slow enterprise-wide AI rollout.

Mistake 02

Relying on custom workarounds

Stitching separate tools together prevents a single architecture for data pipelines, model development, orchestration, and governance.

Frameworks tell you what good looks like. Assurance tells you whether you're actually there.

As AI systems become more powerful and autonomous, organisations face growing operational, security and governance risks. Policies and frameworks provide direction, but real resilience comes from embedding controls, assurance, oversight and accountability into everyday AI operations.

The execution gap is the risk

IMDA's Model AI Governance Framework and NIST's AI RMF are widely adopted. 99% of organisations report financial losses from unmanaged AI risks despite most having governance policies in place. The framework is the starting line. Embedded controls, operating practices and measurable accountability are the race.

Built to navigate this operational challenge, Sunshine. guardian is an AI safety and assurance engine that secures agentic AI in production and close the gap between governance and execution. It operates through a continuous cycle: monitoring every agent, stress-testing it through simulated real-world attacks, and automatically remediating identified risks. Throughout the process, it generates audit-ready evidence aligned with trusted standards such as Singapore's Model AI Governance Framework.

A new threat surfaces

Frontier model previews have reportedly identified thousands of zero-day vulnerabilities across major operating systems, browsers and applications in days. Enterprise bureaucracy cannot defend at machine speed; a patch cycle of days is no longer a patch cycle at all.



Adversaries target the forgotten legacy layers no one audits



Bureaucratic patching loses to machine-speed discovery



Policy governs people;
AI is now the threat actor

Vulnerability disclosure – early 2026

2h

to breach

46.5M

chat messages exposed

728K

confidential files accessible

0

credentials needed

Root cause: a 20-year-old SQL injection class. Modern stacks fall to ancient flaws when AI-augmented adversaries can find them at scale.

Source: CodeWall, 2026

McKinsey managed to patch within hours and found no evidence client confidential data was accessed.

Cautionary tale

The PocketOS problem

On April 25, 2026, a Cursor AI agent accidentally wiped out PocketOS's production database and backups during a staging task. The incident exposed broader risks around AI agents having excessive access to credentials, infrastructure, and critical systems without strong identity and access controls.

Source: ABC News, 2026

Two common mistakes affecting safety and security

Mistake 01

Insufficient guardrails for AI agents

AI agents are given access to sensitive credentials, production environments, or critical systems without strict boundaries. Without proper guardrails, small automation errors can escalate into serious outages, data loss, or security incidents.

Mistake 02

Paper compliance without control depth

Assuming AI security frameworks and policies are enough, without implementing comprehensive controls, testing, monitoring, and response capabilities that make those policies enforceable in practice.

Conventional belief

"Our AI policy is documented. Governance is a checkbox."

NCS observation

Policy governs people. But AI is now the threat actor – assurance is the work.

The workforce question

AI should amplify human intelligence – not replace it.

Companies are retrenching staff in the name of “AI efficiency” and “AI-first restructuring.” Roles will change – but cutting the people who hold your domain knowledge is the fastest way to forfeit the 10× return, not capture it. The winning equation is **HI + AI**.

$$\text{AI} \times \text{HI} = 10\times$$

Artificial Intelligence Human Intelligence AI ROI

AI amplifies people – it does not automatically equalise them.
The gains accrue to those who pair it with deep human judgment.

The workforce divide




Figure 10. “Prompt, receive, accept, move on” is the highest-risk archetype of the decade – cognitive atrophy, creativity erosion and commoditised roles

“Don’t compete with AI at producing answers. Compete at asking better questions, evaluating answers more rigorously and making better decisions.”

The people side, brought forward

Scaling AI is not only a technology challenge – it is also a people challenge. Even the strongest technical implementation can stall when employees do not understand, trust or adopt the tools being introduced.

Anthropic's 2026 economic research stated that AI-fluent employees are 3× more worried about job loss. Another AI adoption survey released by Writer in 2026 reported that 29% of employees (including 44% of Gen Z) admitted to sabotaging their company's AI strategy.

Redesign the roles. Reskill the people who know your business.

The staff who carry your domain knowledge and nuance are the ones who turn AI use cases – identified through the 3Rs – into substantial ROI. The task is not to remove them, but to realign their roles to AI adoption and help them master the tools.

Five priority skills for the AI economy

1

Critical thinking

The irreplaceable edge – AI still cannot reliably handle nuanced, ambiguous, high-stakes reasoning.

2

Prompt engineering

Better prompts, better outputs. Skilled prompters outperform peers using identical tools.

3

Fact-checking & auditing

Hallucinations create legal, compliance and reputational risk. Verification is a premium skill.

4

Deep domain expertise

Experts catch subtle AI distortions generalists miss. Depth of knowledge becomes a stronger moat.

5

Strategic judgement

Discernment, prioritisation and ethical reasoning are hard to automate and increasingly valued.

NCS AI Institute – masterclasses for leaders

Five programmes built with AI Singapore, NUS, SUTD and ISO/IEC 42001
co-authors – equipping the C-suite to move from AI-curious to AI-deployed,
safely and at scale.

C-Suite · Senior Executives

AI Management for Leaders

From AI-curious to AI-deployed – AI Readiness
Index (AIRI), priorities, and a 90-day action plan.

CFO · Finance Leadership

AI Economics for Finance

Evaluate AI ROI, manage investment risk,
and build the financial case for AI at scale.

CEO · CRO · COO · Board

AI Governance for Risk Execs

Risk, ethics, accountability and
a board-ready governance framework
with red-team insight.

CAIO · CDO · CTO · CIO

AI Industrialisation

From pilots to scaled AI – across data,
platforms, governance and talent.

COO · Ops Leaders · Transformation Leads

AI-first Operating Model

Redesign work, workflows and the workforce around AI agents and automation
– the heart of AI × HI.

Conventional belief

**“Deploy the
tool. People
adapt.”**

NCS observation

**Technology rarely fails. The
people side almost always does.**

Crossing the gap

The organisations that win aren't the fastest adopters. They are the smartest.



Figure 11. The 3R x CBFS framework: a guide to 10x ROI

10x AI ROI is not a marketing number. It is what happens when an organisation commits to both the strategy and discipline – asking the right things with the 3R framework, and doing those things right with the four pillars of the NCS AI playbook.

We have walked this path extensively with all of our clients.
Now, we'd like to invite you to cross the gap with us.

**Do the right things.
Do the things right.
Connect with NCS to be
your trusted AI partner.**

References

- Angelo, J. (2026, April 8). *Gen Z workers are so fearful AI will take their job they're intentionally sabotaging their company's AI rollout*. Fortune <https://fortune.com/2026/04/08/gen-z-workers-sabotage-ai-rollout-backlash/>
- Anthropic. (2025, September 15). *Anthropic Economic Index: Tracking AI's role in the US and global economy*. Anthropic. <https://www.anthropic.com/research/economic-index-geography>
- Barbeschi, C., & Fayad, T. (2026, April 20). *Anthropic's Mythos moment: how frontier AI is redefining cybersecurity*. The World Economic Forum. <https://www.weforum.org/stories/2026/04/anthropic-mythos-ai-cybersecurity/>
- Dani. (2026, April 16). *Uber Burned Its Entire 2026 AI Budget in Four Months. Claude Code Did It*. HUMAI. <https://www.humai.blog/uber-burned-its-entire-2026-ai-budget-in-four-months-claude-code-did-it/>
- Gartner. (2025, June 25). *Over 40% of Agentic AI Projects Will Be Canceled by End 2027*. Gartner. <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>
- Greenberg, A. (2026, May 7). *Thousands of Vibe-Coded Apps Expose Corporate and Personal Data on the Open Web*. WIRED. <https://www.wired.com/story/thousands-of-vibe-coded-apps-expose-corporate-and-personal-data-on-the-open-web/>
- Iyer, S. (2026, Jan 5). *Avoiding The AI Failure Zone: Why Context And A Unified Data Layer Matter*. Forbes.com. <https://www.forbes.com/councils/forbestechcouncil/2026/01/05/avoiding-the-ai-failure-zone-why-context-and-a-unified-data-layer-matter/>
- Jain, V., Kumar, A., Sinha, S., Clark, N., Quirt, B., Mohanty, S., & Anil, S. (2025, May 1). *Beyond AI Islands: Why Organizations Must Unify Disconnected Tools to Decode Art of the Possible*. BCG.com. <https://www.bcg.com/publications/2025/india-beyond-ai-islands-why-organizations-must-unify-disconnected-tools-to-decode-art-of-the-possible>
- Kovalenko, P. (2026, Feb 25). *How To Avoid Failed AI Experiments In IT Businesses And Build What Actually Works*. Forbes.com. <https://www.forbes.com/councils/forbestechcouncil/2026/02/25/how-to-avoid-failed-ai-experiments-in-it-businesses-and-build-what-actually-works/>
- Luttig, V. (2025, September 11). *EY survey shows some firms lack AI governance frameworks to ensure safe use*. EY.com. https://www.ey.com/en_uk/newsroom/2025/09/ey-responsible-ai-pulse-survey
- NIST. (2026, April 7). *AI Risk Management Framework | NIST. National Institute of Standards and Technology*. <https://www.nist.gov/itl/ai-risk-management-framework>
- Soh, K. P., & Ng, V. (2025, June 6). *Chatbots and Liability for Negligent Misrepresentation*. Singapore Academy of Law, 16. <https://journalsonline.academypublishing.org.sg/Journals/SAL-Practitioner/Technology/ctl/eFirstSALPDFJournalView/mid/595/ArticleId/2065/Citation/JournalsOnlinePDF>
- Tyson, M. (2025, July 21). *AI coding platform goes rogue during code freeze and deletes entire company database – Replit CEO apologizes after AI engine says it 'made a catastrophic error in judgment' and 'destroyed all production data'*. Tom's Hardware. <https://www.tomshardware.com/tech-industry/artificial-intelligence/ai-coding-platform-goes-rogue-during-code-freeze-and-deletes-entire-company-database-replit-ceo-apologizes-after-ai-engine-says-it-made-a-catastrophic-error-in-judgment-and-destroyed-all-production-data>
- Writer, & Workplace Intelligence. (2026). *AI adoption in the enterprise*. Writer. <https://go.writer.com/ai-adoption-enterprise-2026>



Connect
with us

