

MINIMUM VIABLE HUMAN

WHY THE FUTURE OF AI
LEADERSHIP DEPENDS ON
KNOWING WHERE HUMANS
STILL MATTER MOST

Who is this paper for?

This brief is for CEOs, boards, executive teams and transformation leaders scaling AI across work, service, operations or decision-making. It is designed to help leaders ask a harder question than “What can we automate?”

Where must humans remain involved so the organisation stays trusted, accountable and governable?


The risk most AI strategies are missing

AI can make organisations faster, leaner and more productive. But if leaders cut too deeply into the human layer, they risk creating organisations that are efficient but fragile, automated but unaccountable, and intelligent but no longer wise.

The Minimum Viable Human question helps leaders identify where humans must remain present to preserve five critical capabilities: Oversight, Judgement, Accountability, Relationships and Succession.

After reading this brief, leaders should be able to ask:

- Where are we automating decisions faster than we can govern them?
- Where are we removing the human experiences that develop future leaders?
- Where could accountability become too complex to trace?
- Which human relationships must be protected because trust cannot be fully automated?



When AI causes the damage, who goes to jail?

The year is 2030. A CEO and their Chair are sitting before a Royal Commission. In the wake of thousands of customers being denied support, overcharged, misclassified or pushed into unsuitable products by their company, they are asked a simple question.

"How did your company make so many poor decisions that had such a detrimental impact on the people you were supposed to be serving?"


We don't know," replies the CEO. "It was our AI Agents that made those decisions. We weren't aware of them until the damage was done. We are currently recruiting people to help us fix these issues, but it will take some time. Very few people understand the complexity of the work anymore."

What do you think will be the outcome here? Is saying: "My AI hurt my customers" in the same tone as "My dog ate my homework" likely to fly?

If the negligence is significant, who should be held accountable? Who should pay the fine? Who should go to jail? We are racing toward companies that can operate with fewer humans without asking whether they will still have enough humans to remain accountable, responsible, sustainable and trustworthy.

The dream of the AI-enabled company is seductive: fewer





people, faster decisions, lower costs and smoother operations. Sales, marketing, finance, customer service, reporting, recruitment and management dashboards all humming away with minimal human intervention. Automated decisions, agentic AI at the steering wheel working at a relentless pace. Never tiring, never complaining, never needing a holiday. But the most important question for leaders is no longer, "What can AI do?" Instead:

How few humans can an organisation have before it stops being sustainable, accountable and governable?

Every organisation now needs to define its **Minimum Viable Human** layer. The smallest group of people capable of preserving judgement, trust, accountability, learning and leadership succession in an AI-driven organisation. What capabilities will they need people to do when AI is running so much?

While our initial dreamy visions might suggest a massive, AI-driven company full of robots, with a mastermind human at the helm... we might need to think on this a little more.

What will all the humans do?

There is no doubt that AI is rapidly reshaping the employment landscape. Reportedly 150,000 US employees have lost their jobs in the first half of 2026 alone to AI, and this is just the beginning. Most at risk include junior developers, data analysts and support roles in the tech sector. Customer service and help desk roles in call centres. Financial analysts, compliance officers and loan processing within the Finance world. Paralegals, copywriters, translators, content moderators... the list goes on.

While the World Economic Forum has predicted that AI, robotics and information-processing technologies will be a net creator of jobs, they project that by 2030, 92 million current roles will be displaced globally. It was hard to imagine just a decade ago that AI would be capable of the things it does for us today and at the rate that Claude, ChatGPT and others are advancing, future workplaces are likely to be hives of agentic AIs running 24/7. What jobs will be left to humans? What types of work should we be training our children to do? What do we need to retrain ourselves to do?

While much of the AI conversation seems focused on the concern of jobs being cut, too little of the conversation is being focused on the long term impacts it will have on organisations and institutions themselves. All the new models, data centres and share prices are being built with the focus on transforming the world of work activity, but what about the organisations themselves?

The Minimum Viable Human Framework identifies five human capabilities AI-driven organisations cannot afford to lose.

Together, these form the MVH Framework: Oversight, Judgement, Accountability, Relationships and Succession.

The Minimum Viable Human Framework



Oversight

Monitor, challenge, intervene



Judgement

Interpret through complexity



Accountability

Own responsibility for outcomes



Relationships

Preserve human connection



Succession

Develop the future leaders

Oversight

Through Robodebt, Australians have seen what happens when automated decision-making outpaces effective human oversight. Robodebt was designed to recover more than \$1 billion in alleged welfare debts, but it failed on multiple fronts. It harassed thousands of people creating serious harm to many vulnerable people and ended up costing the Commonwealth around \$2.4 billion in compensation and settlements. One of the major criticisms of the scheme was the lack of human oversight and empathy for the circumstance it was putting people through.



Under the Australian AI Ethics Principles, it states: "Those responsible for the different phases of the AI system lifecycle should be identifiable and accountable for the outcomes of the AI systems, and human oversight of AI systems should be enabled."

Oversight is the human capacity to monitor, challenge and intervene when automated systems create consequences that are technically efficient but ethically, socially or commercially dangerous. While algorithms are excellent at following clear rules and guidelines, this isn't going to be the focus for the larger wave of AI implementation. We are now deploying AI into complex, ambiguous and often messy ethical frontiers. These situations require more than a logic check to see that they have complied with the rules, they require someone to navigate the values, intentions and broader consequences involved.

This is also where diversity and constructive challenge become essential. A single human reviewer may create efficiency and consistency, but they can also become a bottleneck or simply reinforce existing flawed assumptions. Having a number of diverse people with varied experiences, perspectives and thought processes enables better exploration and understanding.

For any organisation to implement effective and safe AI within their organisation, they must consider the human capability to monitor the decisions being made against the moral, ethical and social implications it could have. The greatest risk may not come from one bad decision, it will be the thousands of small, unchallenged decisions that compound into disaster before anyone realises.

The big questions:

- As AI scales rapidly, how many humans will be required to provide oversight to the vast number of decisions and actions an AI could take?
- How will they be trained to apply any ethical guidelines on these decisions?
- How do we ensure they have the capabilities and diversity to challenge effectively?
- How will they maintain visibility as the AI systems increase their complexity?

Oversight is not just observation. It is the capacity and capability to intervene when needed. In a Minimum Viable Human layer, oversight must be a priority. Without that human capacity, oversight will lag: dashboards, policy changes and committee meetings always chasing problems after the harm has already scaled.

Judgement

Excavations in 2008 of rock shelters within the Juukan Gorge in Western Australia found clear evidence that indigenous people had lived, hunted and made jewellery more than 40,000 years ago. This highly significant archaeological site was further explored in 2014 providing incredible insight into the ancient lives of Australia's First Nation's people.

In May of 2020, Rio Tinto drilled 382 blast holes into the area surrounding the gorge and packed them with explosives. When traditional owners of the land heard about the impending explosion, they launched an emergency

appeal that was initially agreed to by Rio Tinto. Just 9 days later, the explosives were detonated, dislodging around eight million tonnes of high-grade iron ore worth around \$135 million. It also destroyed much of this ancient and sacred site.

Rio Tinto had the legal approval to conduct the explosions, this was granted to them in 2013. But judgement goes beyond legal process, it is about making decisions that weigh more ambiguous factors that include context, consequence and morality. It is not asking the question "Can we do it?" but "Should we do it?" and "What will be the broad impacts of this decision?"

This matters now more than ever before because many of the most important decisions inside organisations are neither purely technical nor dimensionally confined. They involve long-term trade-offs and complex network repercussions. A decision may be commercially



attractive today but broadly damaging across other stakeholders over time. Legally permissible but reputationally dangerous. Good to win the battle, but what about the war? Decisions that are consistent with policy in the here and now may have seriously damaging impacts in different places and times.

These are the moments where good judgement becomes essential. AI systems can analyse data, generate options, summarise risks and recommend actions. But they need to be instructed as to the moral responsibility they must use. They do not necessarily understand nuanced cultural significance or elements of historical trauma not well documented in their training data. They don't innately consider community trust or the long-term symbolism of a decision in the way humans must when making complex decisions. AI can quickly optimise towards an objective, but it cannot always know whether that objective is worthy of pursuit. That is why AI-enabled organisations need enough humans with the experience, wisdom and authority to guide through the uncertain and volatile situations our complex, interconnected modern world faces.

This is also where narrow organisational focus becomes dangerous. When speed, revenue, efficiency or transformation become hyper-dominant goals, judgement can be silenced. Busy people intent on these goals may stop asking difficult questions because the process appears compliant, the commercial case looks compelling or the decision has already gathered too much momentum. The logical brain sees the green lights and races towards the prize. AI is the ultimate logical brain.

For any organisation to implement AI safely and effectively, it must preserve the human capability to interpret complexity, challenge assumptions and make decisions where values, consequences and

outcomes collide. The greatest risk is not that AI will obviously break rules to succeed. The greater risk is that AI will find plausible, efficient and commercially attractive answers that humans no longer have the judgement, courage or authority to challenge.

The big questions:

- Who will be responsible for challenging AI recommendations when they appear plausible but problematic?
- How will organisations ensure humans are not simply approving decisions they no longer deeply understand?
- Where will future leaders gain the experience required to make difficult calls in ambiguous situations?
- How will organisations protect the space for pause, reflection and challenge when AI increases the speed of work?

Judgement is not the ability to choose from options an AI provides. It is the capacity to understand whether those options are good, fair, wise and worthy. In a Minimum Viable Human layer, human judgement must be protected deliberately. Without it, organisations risk becoming highly efficient decision machines with too few people left who know when the machine is wrong.




Accountability

In March 2010, Queensland Health went live with a new payroll system for roughly 78,000 health workers. Payroll is one of the most basic promises an organisation makes to its people. Sadly, the new system failed almost immediately. Staff were being underpaid, overpaid or not paid correctly. What began as a technology implementation became one of Australia's most expensive public administration failures. While various committees bickered and vendors scrambled, many employees faced the consequences of uncertainty and financial hardship.

The Queensland Health payroll failure was not simply a technology overrun. The Commission of Inquiry found that too many bodies and people were involved in decision-making. Committees, boards, project teams and vendors all had roles, but it was not clear where responsibility for the governance and successful implementation sat. The Commission's conclusion was simple: there was no single point of accountability.

In simple systems, accountability can be clear. As complexity increases, accountability can become diffused. When a musician plays the wrong note, it's on them. But an orchestra can have individuals playing fine but still sound off. That's why the Conductor is so critical.

The modern world is already complex, but AI will make this problem significantly harder. In an AI-enabled organisation, a decision may be shaped by data collected by one team, a model built by another, configured by a vendor, approved by a governance committee, deployed by a business unit and acted on automatically across thousands of customers, employees or stakeholders.



The more advanced these systems become, the harder it may be to trace exactly where a decision came from, why it was made, who influenced it and who had the authority to stop it. For AI to be safely implemented, accountability must bring oversight and judgement to a sharp point of ownership. Oversight sees what is happening. Judgement interprets what it means. Accountability decides who owns the consequence and who must act. Without it, oversight can be just observation and judgement, just a vague range of opinions.

For any organisation using AI, accountability must be designed before the system scales. Leaders need to know who owns the decision, who owns the system, who owns the risk, who owns the escalation path, who owns the customer impact and who has the authority to stop the machine. The risk is not that there won't be enough intelligence or complexity in executing the parts, it will be that without accountability, no one is responsible for the whole.

The big questions:

- Who is ultimately accountable for the outcomes created or influenced by AI systems?
- Who has the authority to pause, override or withdraw an AI system when harm begins to appear?
- How will leaders ensure accountability doesn't become lost in process, complexity or technical explanations?
- Who holds the liability if the complexity spirals into harm?

Accountability is not just about someone to blame after the fact. It is maintaining clear human ownership before, during and after decisions. In a Minimum Viable Human layer, accountability must be explicit, visible and enforceable. Without it, organisations risk creating systems powerful enough to affect thousands of lives, but too complex for anyone to truly be held accountable.

Relationships

Commonwealth Bank offers a useful Australian example of a different logic at work. Even as banking becomes more digital, CBA has continued to position business banking as relationship-led, with dedicated bankers and relationship managers central to how it serves small, medium and regional business customers.

This is not because CBA is retreating from AI or technology. Quite the opposite. The bank has positioned technology and AI as strategically critical. It is because they know technology alone isn't the way to develop the most important element of banking: trusted relationships.

Their strategy isn't sentiment. It's commercial. CBA strongly believe that when customers are making significant financial decisions, they want a person who they knew understood their business, understood their context and could exercise judgement to navigate the complexities. Whether borrowing to expand, managing cash flow through uncertainty, dealing with risk; the data, the platform and the AI could inform the conversation, but they wanted a human to be there with them.



This is what relationships look like as a commercial discipline. Not a warm sentiment layered over an efficient system, but a deliberate strategic investment in human presence at the moments that matter most. CBA didn't retreat from technology. They will invest \$2.38 billion in technology in FY25 alone. While only part of this is in AI, utilising technology remains a high priority. But they are using technology to free their relationship managers to spend more time on the conversations that require genuine human understanding and less time on administration, analysis and transaction processing that doesn't.

This will become increasingly critical. Many organisations have competed in the last century on efficiency. Companies like Toyota, Walmart and Amazon have worked hard to perfect the processes at scale. But one thing that will change rapidly with the deployment of greater AI and automation is that efficiency will simply be a ticket to the game. Organisations must focus more on building a trusted relationship with their stakeholders.

Customers trust you will solve their problem. Employees trust leaders will treat them fairly. Regulators trust organisations will act responsibly. Communities trust decisions will consider broader consequences. Trust is a precious commodity. It is slow to build and once damaged, every decision becomes harder, every explanation is doubted and every mistake becomes more expensive.

As organisations automate more service, sales, support, complaints, onboarding, communications and decision-making, they risk increasing the distance between themselves and the people they serve. A customer can be classified, routed, prioritised, priced, rejected or reassured by systems that are technically impressive but emotionally tone-deaf. The danger is not simply that AI will

remove human contact. The greater danger is that organisations will forget where human contact is most valuable.

When humans are involved, moments become more than process steps. Some moments carry unquantifiable weight and nuance. When a customer feels humiliated, it is more than a negative score. Emotions are rich, complex and intangible. They are at the heart of many human decisions and are not just logical process steps. They are moments of truth, moments of trust.

This is where the Minimum Viable Human layer becomes essential. Humans are needed to preserve the organisation's capacity to understand people beyond their data. Relationships require context, empathy, memory and nurturing. AI can help organisations respond faster, but humans are still needed to decide what deserves care, when efficiency is damaging trust and where a relationship needs to be protected rather than processed.

For any organisation implementing AI, relationships must be treated as a fundamental asset. Not decorative sentiment layered over the top, but a core building block of future performance.

Trust is what allows organisations to operate more efficiently and with less friction. Without it, every mistake becomes another detailed process review, bureaucracy becomes the natural response and red tape spreads.

The greatest risk is not that AI will make organisations less efficient. The greater risk is that it will make them more efficient at



being indifferent. Efficiency without relationship will be like being cheap without quality, you will be in a race to the bottom.

The big questions:

- When an AI system damages customer relationships, who owns the repair and do they have the skills to actually do it?
- Which customers, employees, regulators or communities need a seat in the room when AI-driven decisions are being made?
- As human touchpoints become rarer, how do organisations ensure the ones that remain are genuinely meaningful rather than just escalation paths for complaints?
- How will organisations repair trust when AI-driven systems cause confusion, frustration or harm?

Relationships can't be just a 'nice to have'. They must be central in how the system works. In a Minimum Viable Human layer, organisations must preserve enough human connection to understand, explain, care and repair. Without it, they may become faster and more automated, but less trusted, less forgiven and less human when it matters most.

Succession

Nicola Steele joined Krispy Kreme Australia in 2006 as a crew member while she was still at university. From there, she moved through the organisation as store manager, area manager, state manager, national operations manager, head of retail and development, COO, and eventually CEO of Krispy Kreme Australia and New Zealand in 2023. In 2025, she was promoted again into a global COO role with Krispy Kreme.



Her experience means Nicola knows every part of the business, not because she has been given data or read reports, but because she has lived and experienced it. She knows what it is like to run a counter as a crew member. What it is like to manage rosters, handle complaints and uphold service standards as a Store Manager. She has experienced leadership during expansions into new markets and the development and implementation of new systems and partnerships.

This may seem like an old-fashioned model of succession, but it has served organisations for generations for a reason. Management consultants can bring industry best practice. New hires can bring fresh perspectives and diverse thinking. Both matter. But few things give a future CEO the depth of understanding that comes from living inside the business over time. That kind of knowledge is difficult to capture because it is built through time and experience.

Succession is a critical human capacity because it develops the knowledge that will shape effective oversight, judgement and accountability into the future. It is not just a talent pipeline or a list of potential future leaders. It is the development of highly contextual experience, wisdom and responsibility. Internal promotion accounted for more than 70% of new CEO appointments across APAC in 2025. Boards continue to value leaders who bring institutional knowledge, cultural understanding and lived experience of the business.

AI puts this at risk because many of the tasks most likely to be automated are also the tasks through which people learn. From entry-level, repetitive and customer-facing work. People doing analysis, drafting, scheduling, reporting and basic problem-solving will all look inefficient from a productivity perspective compared to

agentic AI. As boards sign off on aggressive automation targets and leadership teams celebrate productivity gains, someone needs to think about the human capability lost for the future. The organisation now has fewer people who understand the work, fewer future leaders who have learned through experience and fewer humans close enough to the system to know when something feels wrong.

The tasks we are automating are also the training grounds to develop pattern recognition, confidence and organisational understanding. If juniors no longer experience the frontline work, will they become senior people capable of governing it? If they no longer sit close to operations and experience the mistakes, trade-offs and small decisions, where do they develop the judgement required for larger decisions later? If they don't engage with customers, complaints and partners, how do they build the trusted relationships? Organisations may become more efficient in the short term while quietly weakening the talent pipeline needed long term.

The work will get done and decisions will move faster. But fewer people will understand how the work really happens, where it breaks or what risks lay beneath the polished output. Future leaders will need judgement, ethics, empathy and the courage to challenge automated recommendations. Those capabilities are not built by receiving AI-generated summaries and reports. They are built through participation, responsibility and consequence.

Succession must be designed deliberately and protected. Leaders need to ask which tasks are simply mindless labour and which tasks involve learning and understanding. This does not mean protecting pointless busywork. It means protecting the pathways through which humans become capable. The greatest risk is not that AI

replaces today's junior roles. It is that it removes the opportunities to build experience to create tomorrow's senior leadership.

The big questions:

- How will people learn to challenge AI outputs if they no longer understand the work beneath them?
- If the work that teaches people how to think is automated away, what replaces it — and who is responsible for designing that replacement?
- When your next crisis hits, will you have enough people who understand the system deeply enough to know where it broke — or will you be recruiting that capability after the damage is done?
- Who is responsible for ensuring the organisation still has enough people capable of governing the system in five, ten or twenty years?

In a Minimum Viable Human layer, succession must be treated as a strategic capability. The danger is not just that AI replaces junior roles. The danger is that it removes the lower rungs of the ladder while still expecting people to appear at the top.



So, How Many Humans Do You Need?

Modern organisations cannot afford to ignore the power of AI. It would be naive, if not fatal, to ignore the opportunities to become faster, smarter and more productive. However, fast is not always better, smarter isn't always wiser and productivity only matters if you are doing the right things.

The real leadership challenge is not just deciding where AI can replace humans. It is deciding where people are needed to provide oversight, make careful judgements, retain accountability, build relationships and prepare the pipeline of succession.

This is the essence of the Minimum Viable Human question. It is not sentimentality or AI scepticism. This is a strategic imperative. One based in risk, governance and organisational sustainability. The seductive promise of AI is that organisations can do far more with fewer people. In many areas, this will absolutely be the case. But cutting too deeply into the human layer may leave organisations impressively efficient and strategically fragile. When the next Royal Commission comes to your door, "the AI did it" will not be a valid defence for CEOs or Boards to hide behind.

We shouldn't be frozen in fear or reading this as an anti-AI argument. This is a warning against human-light organisations becoming highly productive, ethically-questionable and accountability-light organisations. The future will not belong to organisations that simply remove the most humans. It will belong to organisations that understand where humans matter most.



The question every leader needs to answer

The organisations that navigate this well won't be the ones that moved fastest or cut deepest. They will be the ones that were deliberate about what they kept.

They asked which decisions needed a human. Which relationships required care that couldn't be automated. Which mistakes needed someone with the authority and courage to own them. Which roles existed not just to do work, but to develop the people who would one day govern the work.

These are not soft questions. They are the hardest strategic questions of the next decade and most leadership teams are not yet asking them. **That's the work I do with organisations.**

I help leadership teams identify where humans are genuinely irreplaceable, design the governance structures that keep AI accountable, and build the trust, judgement and authority required to intervene before harm scales.

Bring this conversation to your leadership team

If your organisation is scaling AI, redesigning work or reducing human involvement in key decisions, this is the conversation to have before efficiency becomes fragility. To explore a keynote, executive briefing or workshop, contact Daniel Murray.

Daniel Murray
daniel@empathicconsulting.com
+61 418 920 412
empathicconsulting.com

