Permission to Scale

The critical role of trust in successful AI adoption



How leading Australian businesses are building the trust needed to scale AI in 2026

© 2026 Cisco and/or its affiliates. All rights reserved.

Chapters

01 Chapter one
Introduction:
The AI trust dynamic 3

02 Chapter two
Key findings 8

03 Chapter three
Research findings
and recommendations 12

04 Chapter four
Conclusion 40

© 2026 Cisco and/or its affiliates. All rights reserved.



01

Introduction:

The AI trust dynamic

"Trust is the rate limiter on scale, not compute or capability. The technical ceiling on what AI can do is far above the governance ceiling on what we are prepared to let it do – and that ceiling rises only as evidence accumulates."



Sanjeev Gupta,
Chief Information and Transformation Officer,
HBF Health Insurance

© 2026 Cisco and/or its affiliates. All rights reserved.

Chapter one

Introduction:
The AI trust dynamic

Artificial intelligence can generate, analyse, automate, and act. But the pace at which those abilities are adopted within organisations depends on something other than technology.

Early adopters are finding that one of the most important barriers to broader AI adoption is quintessentially human.

That factor is trust.

Organisations can buy AI tools, build AI models, and provide AI training. But if the people using them are not confident that AI is reliable, secure, and accountable – or see it as a tool for their own replacement – adoption quickly stalls.

Trust has become an essential condition for scaling AI. It shapes how willing people are to use AI, influences whether organisations can move beyond experimentation and apply AI to more consequential work and helps determine whether AI changes how the organisation operates.Trust acts as an organisation-wide permission mechanism for AI adoption and is a critical, concurrent element of successful AI strategy.

So how can leaders build appropriate confidence in a technology that many people do not fully trust?

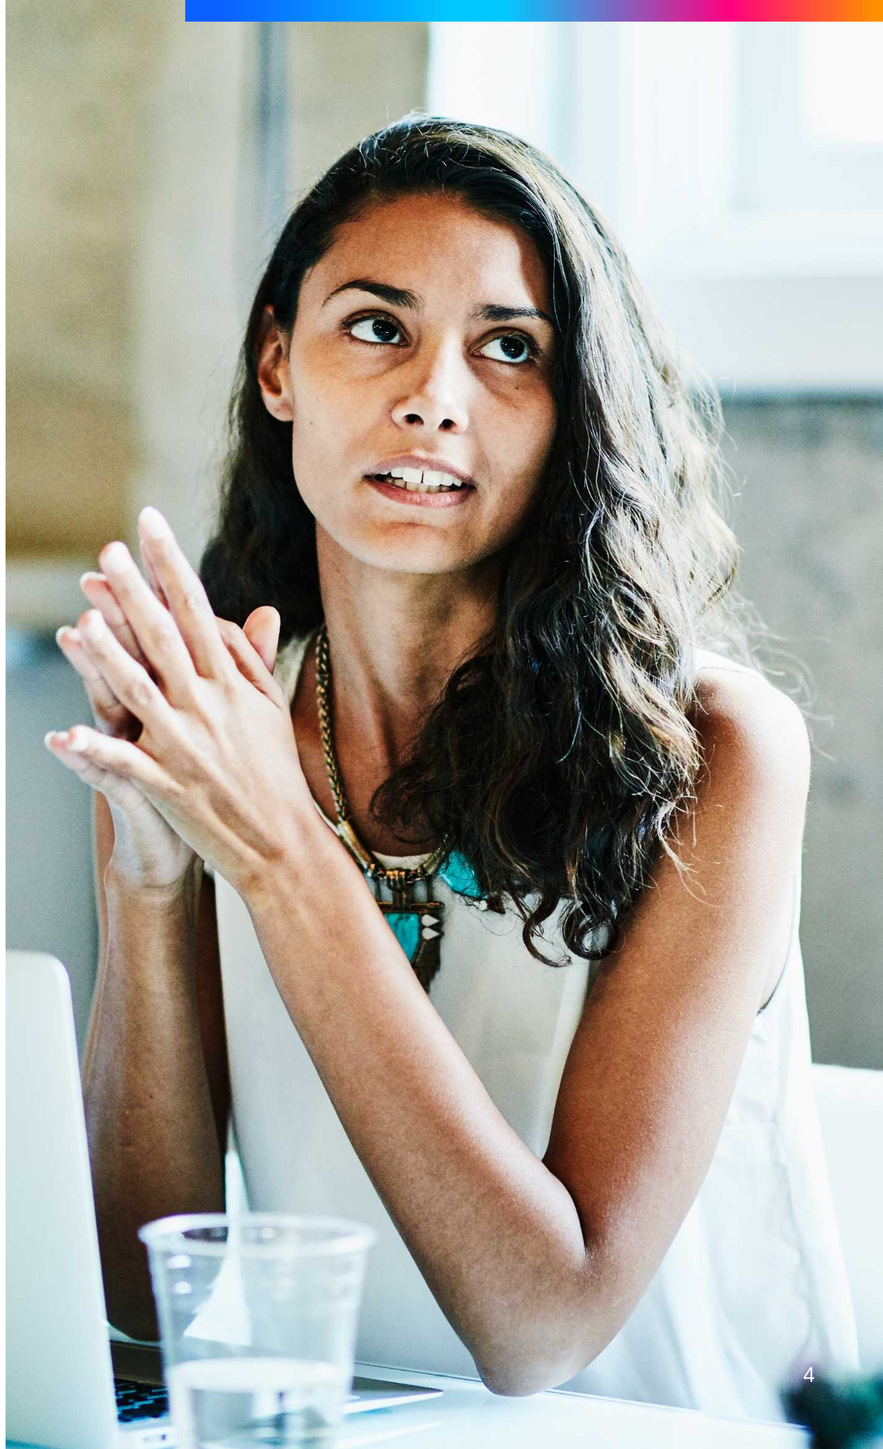
This report examines the relationship between trust and AI adoption within organisations, based on insights gathered from business leaders at the forefront of AI adoption.

It investigates what builds trust among employees, customers, and other stakeholders, what can weaken it, and how organisations can maintain confidence as AI use expands. It also provides practical guidance for leaders seeking to build trust in AI within their businesses.

As this report shows, the goal is not unconditional trust. Organisations need calibrated trust, where confidence grows as evidence, safeguards, and accountability give people reason to rely on AI. At a time when organisations are seeking to maximise the advantages AI can offer, the ability to adopt it quickly, responsibly, and at scale is critical.

Success with AI adoption depends on more than just technology. It depends on whether employees, leaders, and customers have enough trust in AI to use it, rely on it, and permit it to play a greater role in how the organisation operates. While trust is critical to successful AI adoption, deployment can still begin before it is fully established. Trust can develop in parallel with an organisation's AI strategy, as people build skills and fluency, gain confidence through experience, and develop an understanding of the technology's capability, role, and value.

As evidence of reliability, safeguards, and accountability builds, that confidence can strengthen into trust. The development of trust in AI should therefore be integrated into the AI strategy from the outset, evolving alongside its deployment rather than being treated as a separate step that follows it.



© 2026 Cisco and/or its affiliates. All rights reserved.

“If you haven’t built trust, you’re not going to get the outcomes or benefits you predicted because people won’t use the system.”



Sarah Carney,
Head of Artificial Intelligence,
AustralianSuper

Chapter one

Premise:
Australia’s AI trust paradox

AI use is already widespread in Australia, but trust, optimism, and AI literacy remain comparatively low.

This creates a practical challenge for organisations. The technology may be available, and employees may already be using it personally or professionally, yet the trust needed before you can embed AI into important workflows can still be missing. Exposure alone doesn't build it. It needs experience, evidence and effective safeguards.

This is not simply due to a lack of exposure or experience. According to a 2025 University of Melbourne and KPMG study, half of Australians are using AI regularly, but only 36 per cent are willing to trust it. Just 30 per cent believe the benefits of AI outweigh the risks, the lowest result among the 47 countries surveyed.

That gap between capability and adoption for enterprises is already visible in enterprise deployments. **Deloitte’s 2026 State of AI in the Enterprise** research found that only 28 per cent of Australian respondents had moved at least 40 per cent of their AI pilots into production, and just 12 per cent said generative AI was already transforming their business or industry.

© 2026 Cisco and/or its affiliates. All rights reserved.

Background

This report follows on from research conducted by Cisco in 2025 titled **"Turning Hesitation Into Action: How Risk Leaders can Unlock AI's Potential"**, developed in conjunction with the Governance Institute of Australia. That work examined the risks inherent in AI adoption within organisations and the role risk professionals play in their identification, assessment, and mitigation. Its conclusions highlighted the need for stronger governance and practical controls as AI becomes more deeply embedded in business.

While risk mitigation is critical to the safe adoption of any new technology, it does not automatically lead to widespread adoption.

This report investigates the role that trust plays in supporting the adoption of AI. It draws on a survey and interview process involving 25 Australian organisations, designed to understand their perspectives on the relationship between AI trust and adoption and the techniques they use to build, maintain and measure trust. We connected with business leaders guiding AI adoption across varying roles, including business, technology and HR leadership. The patterns identified in these discussions have uncovered a series of compelling insights on the state of AI trust in Australia today.



© 2026 Cisco and/or its affiliates. All rights reserved.

The authors would like to express their deep gratitude to the many executives who offered their time and insights for this report.

This includes those who gave permission for their thoughts to be included:

Prof Cori Stewart, Chief Executive Officer	**ARM Hub (Professional Services)**
Sarah Carney, Head of Artificial Intelligence	**AustralianSuper (Financial Services)**
Paul Keen, Vice President of Engineering	**Culture Amp (Technology)**
Anne Tosky, Chief People Officer	**ELMO Software (Technology)**
Joe Waller, Chief Technology Officer	**Finder (Technology)**
Chris Locke, Chief Technology Officer	**Flight Centre Travel Group (Travel and Entertainment)**
Naz Wilson, Group Culture and Enablement	**Funlab (Entertainment)**
Sanjeev Gupta, Chief Information and Transformation Officer	**HBF Health Insurance (Healthcare and Financial Services)**
David Fischl, Legal Digital Transformation Lead Partner	**Hicksons (Professional Services)**
Claudine Ogilvie, Co-founder and Chief Executive Officer	**HivePix (Technology)**
Devesh Maheshwari, Chief Technology Officer	**Lendi Group (Financial Services)**
Matthew Nation, Director	**Nation Partners (Professional Services)**
Kathryn Giudes, Founder and Managing Director	**ORCA Opti (Technology)**
Simon Goodrich, Co-founder	**Portable (Media and Technology)**
Lucio Ribeiro, Chief AI and Innovation Officer	**TBWA Australia (Media and Advertising)**

© 2026 Cisco and/or its affiliates. All rights reserved.

02

Key findings



© 2026 Cisco and/or its affiliates. All rights reserved.

Chapter two

Key findings

01

Trust enables adoption and successful adoption reinforces trust

Trust and AI adoption form a reinforcing cycle. Trust gives organisations permission to adopt AI, while successful use creates evidence that can deepen that trust. Organisations that begin building trust early can create the conditions for broader deployment and higher-value applications.

"Training, usage, and successful outcomes start to build a flywheel of trust. By training your people, getting them to use AI, and seeing good outcomes, you create the confidence for them to keep going."

– Chris Locke,
Chief Technology Officer, Flight Centre Travel Group

02

Trust matters most when AI asks people to change how they work

Trust becomes more important as AI begins to alter workflows, responsibilities, job design, and decision-making. Basic AI use cases, such as summarising meeting transcripts or identifying action items from emails, can require relatively little confidence. The stakes rise as AI takes on more complex tasks, changes established ways of working or is granted greater autonomy through agentic systems. As the level of change, risk, or autonomy increases, so too does the level of trust and assurance people require.

"The employee is the change agent that can make or break this. We've done the foundational work around policy, role modelling, training and making it safe. Now we need to move to automating and making sustainable, repeatable changes in our processes."

– Anne Tosky,
Chief People Officer, ELMO Software

© 2026 Cisco and/or its affiliates. All rights reserved.

03

Trust is built on visible foundations of reliability, responsible data use, and accountability

Trust in AI arises from practical evidence. People want to know that AI works, that information is protected, and that responsibility remains clear when something goes wrong. Confidence grows when these safeguards are visible in everyday use rather than existing only in policies or principles.

"When AI does not respond consistently, people start to question its ability. Consistency and predictability matter because every variation becomes part of the trust judgement the user is making."

– Kathryn Giudes,
Founder and Managing Director, ORCA Opti

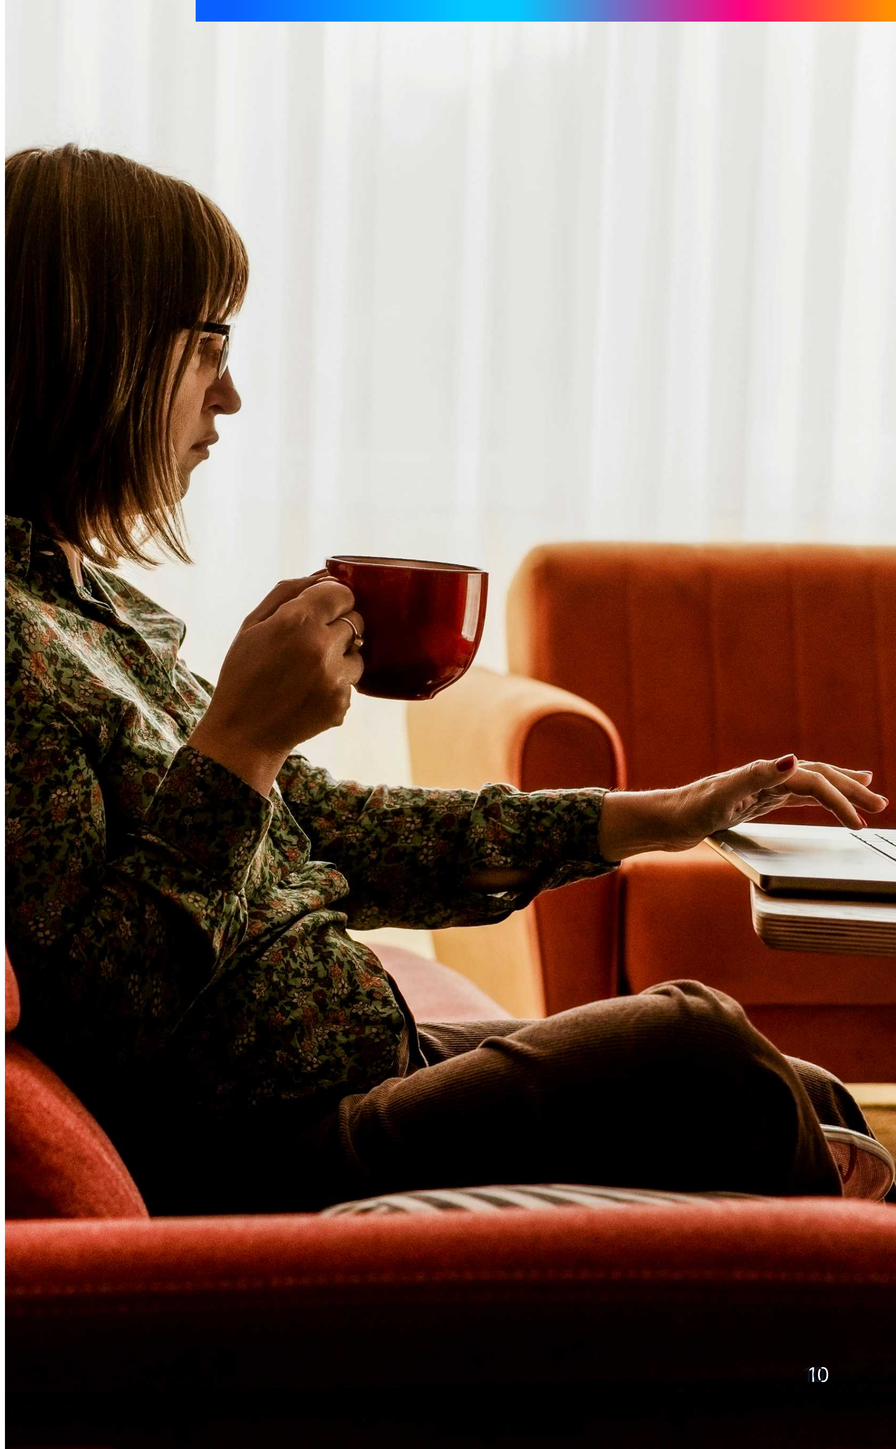
04

The goal should be calibrated trust, not maximum trust

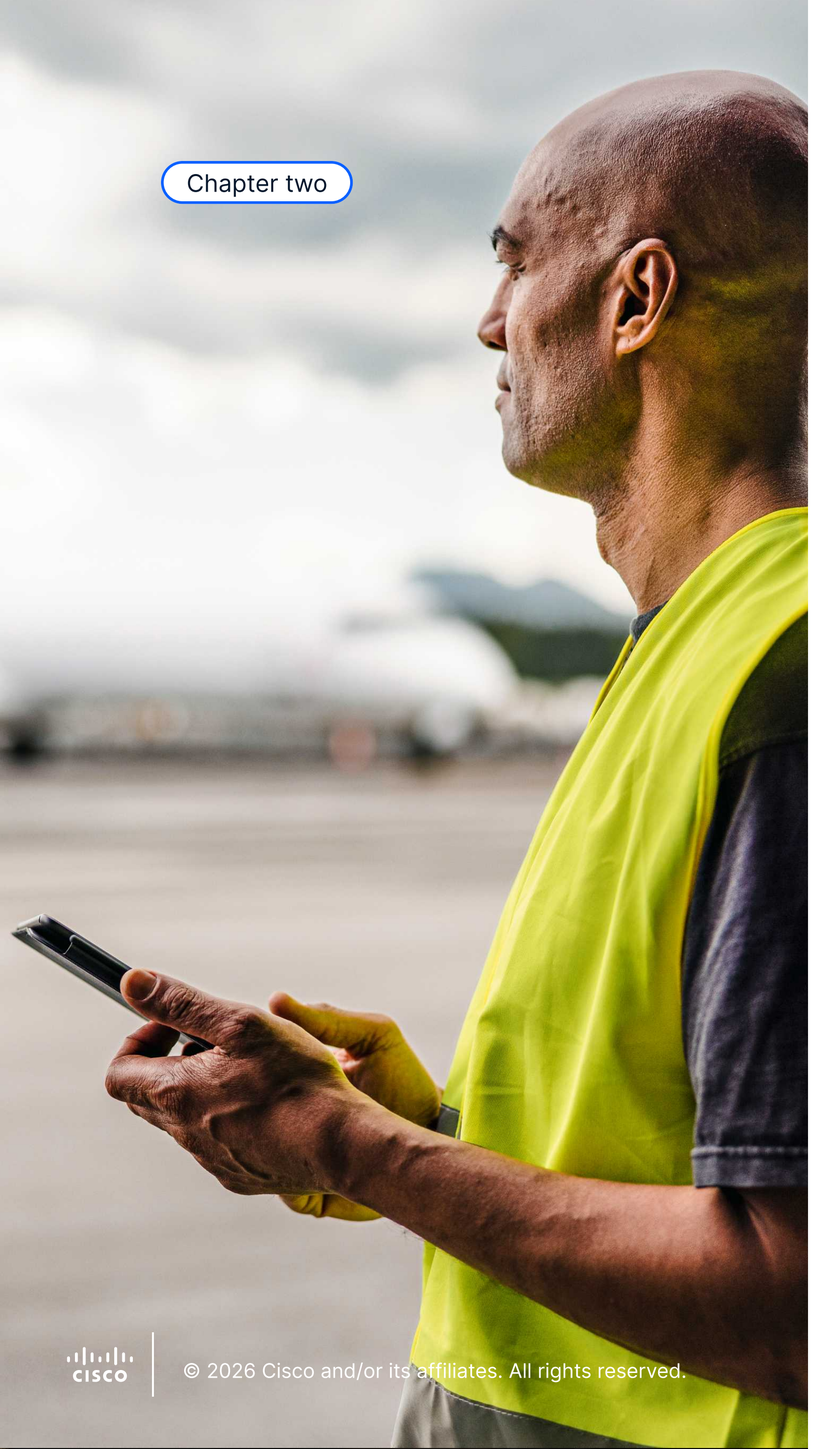
More trust is not always better. Excessive trust can lead people to accept AI outputs without enough scrutiny, while too little trust can stop useful adoption altogether. The goal is calibrated trust, where confidence is matched to the capability and risk of the system, and people understand when AI can be relied upon, when outputs should be checked, and when human judgement should take over.

"Blindly trusting these systems and their outputs is a huge risk. Trust has to include knowing when to challenge, when to check and when judgement still needs to sit with a person."

– Sarah Carney,
Head of Artificial Intelligence, AustralianSuper



© 2026 Cisco and/or its affiliates. All rights reserved.



Chapter two

05

Privacy and security increasingly define what organisations will permit AI to do

As AI connects to more organisational data and systems, the trust question shifts from whether the technology is secure to what information and systems it should be allowed to access, and what actions it should be permitted to take. Privacy and security failures, unreliable outputs, and unclear human oversight can quickly undermine confidence. Stronger identity, access, and monitoring controls can expand what organisations are prepared to permit, while weak controls can restrict adoption.

"We are actively managing the intersection of AI and our cybersecurity and have been cautious around what we give AI access to. We are using AI to safely access information from various systems, although access to our shared network and project files is not permitted. We are vigilant around how AI platforms fit within our broader cybersecurity strategy."

– Matthew Nation,
Director, Nation Partners

06

Future AI capabilities may advance faster than organisational readiness

Organisations are gaining access to increasingly autonomous AI while many are still building the foundations needed to govern it with confidence. The next phase of AI adoption will depend not just on what the technology can do, but whether governance, security, and accountability can keep pace.

"Security protects the system, governance provides the constraints, and reliability earns the trust. As we move further into AI, those three things have to work together rather than be treated as separate activities."

– Devesh Maheshwari,
Chief Technology Officer, Lendi Group

© 2026 Cisco and/or its affiliates. All rights reserved.



03

Research findings and recommendations

© 2026 Cisco and/or its affiliates. All rights reserved.

Trust turns AI use into a reinforcing cycle

Chapter three

Experience converts uncertainty into evidence

Respondents consistently described trust as an important factor in successful adoption and scale.

Our investigations revealed:

01

Trust is either very important or essential to generating business value from AI.

02

Stronger trust is important for accelerating AI adoption.

03

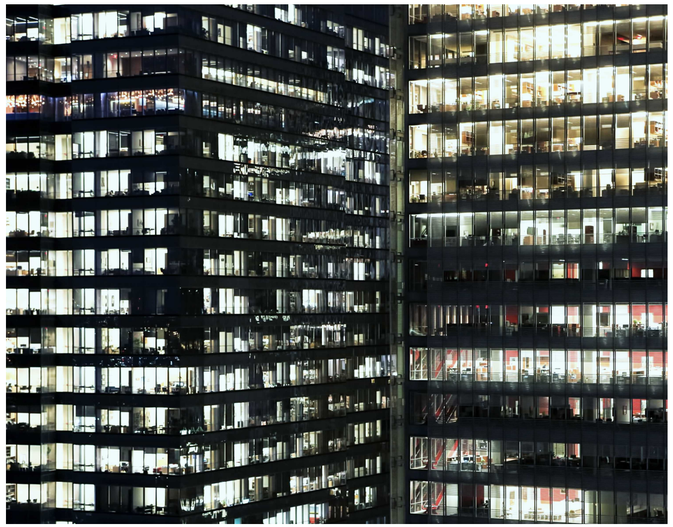
Trust is a prerequisite for successful AI adoption and something that continuously reinforces successful adoption.

© 2026 Cisco and/or its affiliates. All rights reserved.

Chapter three

Trusts turns AI use into a reinforcing cycle

Trust and AI adoption do not have a simple one-way relationship. They reinforce each other. Successful AI use creates evidence that AI can be reliable, useful, and appropriately controlled, supporting its application to higher-value or higher-risk work. That experience strengthens confidence and encourages further adoption.



Broader industry research shows a similar association, with **Deloitte's Asia-Pacific "AI at a Crossroads"** research finding that organisations with more mature AI governance reported 28 per cent higher staff use of AI solutions and close to 5 per cent higher revenue growth. The association is consistent with a central finding of this report – strong governance and trust-building mechanisms can support adoption and value rather than constraining deployment.

Organisations seeking greater AI adoption should not focus solely on technical capability or employee training. They also need to create the conditions that give people confidence to use AI more broadly. Confidence and trust are closely related. Confidence reflects a belief in capability, experience, authority or training, and among the more AI-fluent organisations we spoke to, it emerged as an important foundation for acceptance and trust.

Across the research, reliability, responsible data use, clear accountability and appropriate oversight repeatedly emerged as important foundations of trust, further strengthened by a clear understanding of the technology, its role and the value it creates for employees and the organisation.

The earlier organisations start building these confidence and trust foundations, the sooner the reinforcing cycle can begin.

Respondents described confidence growing when people could see AI working in their specific environment, compare outcomes with previous ways of working, and understand where the technology was likely to fail. In several organisations, safe experimentation was itself part of the trust-building process.

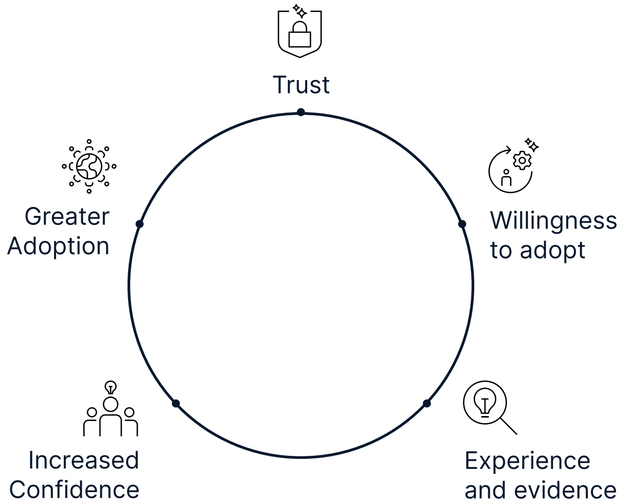
© 2026 Cisco and/or its affiliates. All rights reserved.

“Start small and develop trust. Measure the impact, expand the use of AI, and gradually increase the responsibility it is given.”



Devesh Maheshwari,
Chief Technology Officer,
Lendi Group

The AI trust flywheel



AI introduction with clear use cases and training — **Evidence** — **Calibrated trust (and growing user confidence)**

Permission to expand — **Higher-value use** — **Provable ROI**

Confidence and capability are the foundations for trust, which in turn creates the permission for businesses to use AI. Successful, confident use generates the evidence and ROI needed to strengthen trust and support broader adoption.

© 2026 Cisco and/or its affiliates. All rights reserved.

“We could see that the people using the tools were three times faster, while error rates were exactly the same.
That gave us statistical evidence to back up the benefits, and that built trust around the adoption of AI.”



Paul Keen,
Vice President of Engineering,
Culture Amp

How trust and adoption interact

When asked how they would describe the relationship between trust and AI adoption in their organisation, respondents ranked the following three responses most highly:

01	**Trust and successful AI adoption continuously reinforce each other.**
02	**Trust is a prerequisite and that AI cannot succeed without it.**
03	**Trust is a byproduct of successful AI outcomes.**

Organisations seeking to accelerate AI adoption should deliberately build the conditions that make its use safe, credible, and worthwhile. Trust is not separate from the adoption journey; it is built through it.

© 2026 Cisco and/or its affiliates. All rights reserved.

Findings:

Implications for adoption

01

Building confidence and capability should be part of the adoption strategy

Investment in AI capability alone may not be enough to drive widespread adoption. Employees, leaders, and customers also need the skills and confidence to use AI appropriately. Training, role-modelling, reliability, responsible data use, accountability, and oversight need to develop alongside technical capability as part of the adoption strategy.

02

Good governance creates confidence

Governance committees, human oversight, workforce training, risk assessments, and executive reporting can provide the assurance needed to move AI from experimentation into production.

03

Low adoption may not be a capability problem

Reluctance to use AI should not automatically be attributed to inadequate training, poor functionality, or resistance to change. Employee willingness to adopt and use AI was the area where respondents most frequently said trust influences adoption and scale.

04

The importance of trust rises with the stakes

Trust becomes more important as the value, risk, and autonomy of AI use increase. Using AI to summarise a document requires less confidence than allowing it to automate a complex process, analyse and provide recommendations from private business data, interact with customers, or operate in a high-risk environment. That dynamic becomes more complex with agentic AI, where trust extends beyond the quality of an output to what a system can access, what actions it can take, and whether it will remain within intended boundaries.

© 2026 Cisco and/or its affiliates. All rights reserved.

Recommendations:

Techniques for building trust in AI

Chapter three



Start where trust can be visibly earned

Choose applications where accuracy, value, and risk can be demonstrated quickly. Successful early use gives people evidence that AI can work in their environment. Move from lower-risk experimentation to routine use, and then towards higher-value or more automated applications as evidence and confidence develop.

Make safeguards visible

People should understand what data AI can access, where human review occurs, how errors are handled, and who remains accountable. Safeguards should be experienced as enablers of safe use rather than simply compliance requirements.

Show evidence rather than relying on enthusiasm

Use performance data, successful deployment examples, known limitations, and lessons from incidents to demonstrate where AI can and cannot be relied upon.

Give employees permission to question AI

Trust is stronger when employees know they are expected to question outputs, escalate concerns, and retain professional judgement rather than simply comply with an AI-generated recommendation. Create trusted spaces to share and capture this feedback.

Train for judgement as well as tool use

AI training should cover limitations, verification, privacy, appropriate reliance and accountability rather than focusing solely on how to operate the technology.

© 2026 Cisco and/or its affiliates. All rights reserved.

“Clients usually require us to prove
that the process works with their own data
and delivers the **expected return**.
By the time we reach implementation,
they expect almost no chance of failure.
There is very little room for **a leap of faith** anymore.”



Professor Cori Stewart
Chief Executive Officer, ARM Hub

© 2026 Cisco and/or its affiliates. All rights reserved.

Trust spreads peer to peer

Entertainment business Funlab has taken a deliberately peer-led approach to building AI confidence within its workforce. Its “AI Super Squad” brings together representatives from different departments every four weeks to share what is working, discuss concerns and take practical knowledge back into their teams. The group creates a two-way flow of information between employees, the people and culture function and Funlab’s data and AI leadership.

Participants are given dedicated time to experiment, learn, and share what they discover. The group also helps ensure AI use remains aligned with security requirements and organisational priorities. Funlab sees this peer-to-peer structure as a way to build AI fluency and confidence from within the workforce, with employees learning from colleagues who are already using AI effectively rather than simply being told to adopt it.

“The AI Super Squad is really a bottom-up approach. It is designed for peer-to-peer learning, with people already using AI sharing information back within their own departments. It is not top-down communication - it is **two-way**, with information coming back in as well.”



Naz Wilson,
Group Culture and Enablement,
Funlab

© 2026 Cisco and/or its affiliates. All rights reserved.

Trust becomes critical when work has to change

Chapter three

Trust becomes more consequential as AI moves from an optional productivity tool and is integrated into workflows, decisions, and processes that require people to change established behaviour, give up control, or rely on a system in new ways.

When asked where trust has influence, respondents most commonly selected the following factors as critical to successful AI-led change:

Capability defines what AI can do, but trust defines what organisations will let it do

© 2026 Cisco and/or its affiliates. All rights reserved.

Chapter three

While an organisation can achieve widespread use of low-risk AI tools without fundamentally changing how work gets done, deeper transformation requires people to trust AI enough to redesign workflows, rely on its outputs, or delegate parts of the work to it.

Because of this, the business return from AI may depend more on whether employees trust it enough to change established ways of working than on how much they are using it. Employees are judging not only the technology, but also the organisation's intentions for how it will be used.

"It's not just adoption and use of the tools. It's value generation through that. You have to understand whether people are using it in the right way and whether it is creating the outcomes you intended, rather than simply looking at the number of people who have opened the tool."

– Sarah Carney,
Head of Artificial Intelligence, AustralianSuper



© 2026 Cisco and/or its affiliates. All rights reserved.



Chapter three

People and culture leaders repeatedly connected people's feelings regarding AI with their confidence in job security, the future of work, and whether efficiency gains would increase human impact or simply remove roles.

"The more important group are the workers. We need them to trust AI, not fear they are going to be replaced by it. A large focus for us is how we keep them relevant and help them use the tools to enhance what they do."

– Chris Locke,
Chief Technology Officer, Flight Centre Travel Group

Cisco's 2026 CEO Survey identifies upskilling teams for AI readiness as the leading priority among Australian respondents. Yet 52 per cent say their own understanding of AI is a barrier to informed decision-making. This highlights the need to build AI understanding across leadership as well as the workforce.

If employees hear ambitious AI messages but do not receive the skills or support to work differently, their confidence will not advance at the same pace as the technology.

"People need confidence that AI is helping them do better work - not simply different work. If we can remove the things that drain time and create capacity for higher-impact work, that builds trust in what AI is there to do."

– Naz Wilson,
Group Culture and Enablement, Funlab

© 2026 Cisco and/or its affiliates. All rights reserved.

Chapter three

The same dynamic applies externally. Confidence is greater when customers understand when AI is being used, how their information is handled, where human oversight remains available, and who is accountable for the outcome.

This makes trust a critical element of AI-powered customer experience. A system may be technically capable and internally approved, but needs stronger evidence, transparency, or human oversight before customers are willing to accept it.

"Human involvement is a big part of reinforcing trust from a customer perspective. Meaningful oversight of the outcome helps build confidence."

– Anne Tosky,
Chief People Officer, ELMO Software



© 2026 Cisco and/or its affiliates. All rights reserved.

Findings:

Implications for change

Chapter three

AI adoption should be treated as behavioural change, not simply technology deployment

As AI begins to alter roles, workflows, and decision-making, organisations need to manage the human transition alongside the technical deployment. Changing the technology is only part of the job. Organisations also have to manage how people experience and respond to the change.

"Even when people see colleagues using AI and getting the benefits, adoption does not always happen naturally. At some point you have to raise the floor, give people the tools and make the expected behaviour clear."

– Paul Keen,
Vice President of Engineering, Culture Amp

Resistance may be rational rather than cultural

Reluctance to use AI should not automatically be labelled as resistance to change. People may have legitimate concerns about accuracy, privacy, job impact, loss of control, or responsibility when something goes wrong.

"People don't resist AI because they hate it. They hold back because they're unsure what's allowed and who carries the risk if it goes wrong. Remove that uncertainty and adoption follows. Human skills like Intention, taste, and critical judgment are incredibly valuable."

– Lucio Ribeiro,
Chief AI and Innovation Officer, TBWA Australia

Trust can determine whether AI remains incremental or becomes transformational

Use of AI for basic tasks may require relatively little confidence. Deeper adoption requires people to rely on AI in more complex workflows, higher-risk decisions, and more automated processes. The greater the operational change, the greater the assurance people are likely to require.

"If employees do not have the willingness to adopt, those changes will either fail or drag on. You first need to earn employee trust. Then it flows into leadership and board confidence, because they need observability, repeatability, and auditability before they will support scaling."

– Sanjeev Gupta,
Chief Information and Transformation Officer,
HBF Health Insurance

© 2026 Cisco and/or its affiliates. All rights reserved.

Chapter three

Recommendations:

Techniques for facilitating AI-driven change

Explain what is changing and what is not

People need clarity about which tasks AI will support, which decisions remain human, and how responsibilities may shift.

Involve employees in redesigning work

The people doing the work are often best placed to identify where AI adds value, where judgement remains essential, and where it should not be used.

"A way to accelerate adoption is to show how humans and AI work best together, and how AI can augment human talent. But to do that, organisations need to step back and understand where the value is shifting for them specifically, rather than looking at AI only as task automation."

– Claudine Ogilvie,
Co-founder and Chief Executive Officer, HivePix

Use peer networks and visible role modelling

Trust is social. Cross-functional champions, communities of practice, and leaders who use AI visibly can make experimentation feel safer and turn adoption into a shared learning process rather than a top-down directive.



Make human authority explicit

Employees should know when they are expected to question an AI output, override it, or escalate a concern. Trust is stronger when people understand that professional judgement remains valued.

Measure confidence throughout the change

Usage alone does not reveal whether people are comfortable with the new way of working. Organisations should track confidence, perceived reliability and willingness to rely on AI as adoption deepens. Measures should also test whether people are relying on AI appropriately, including whether they know when to challenge, override, or escalate an output.

"Honestly, we mostly rely on proxies like usage. The signal I'd trust more is calibrated trust – people who use AI on real client work and still push back when the output is wrong. Blind acceptance and avoidance are both trust failures."

– Lucio Ribeiro,
Chief AI and Innovation Officer, TBWA Australia

© 2026 Cisco and/or its affiliates. All rights reserved.

Chapter three

Building trust is a team sport

There is no single organisational definition of what makes AI trustworthy. Different stakeholders require different forms of evidence before they will support, use, or rely on the same AI system.

Different parts of the organisation tend to emphasise different dimensions of trust:

01

People and culture leaders tend to focus on workforce confidence, job impact, psychological safety, communication, and whether AI enhances rather than diminishes human contributions.

02

Technology and AI leaders place greater emphasis on reliability, data quality, security, identity, permissions, observability, and whether systems remain observable and controllable as autonomy increases.

03

Risk and governance leaders focus on boundaries, accountability, proportionality, auditability, regulatory alignment, and clarity about what AI should not be allowed to do.

04

Business and customer leaders are more likely to focus on demonstrated value, return on investment, customer confidence, and whether there is enough evidence to support broader deployment.

This makes it difficult for any single function to build trust on its own. The strongest approaches bring together technical assurance, governance, workforce confidence, and business value through governance committees, cross-functional risk assessment, peer networks, and shared accountability. Two thirds of respondents had an active AI governance committee, and the same proportion provided mandatory AI workforce training.

Trust-building should be designed for the audience. The evidence that reassures an employee may not be the same evidence required by a board, a technology leader, or a customer. No single function can supply all the evidence people need to trust AI. Confidence is stronger when these perspectives are brought together.

© 2026 Cisco and/or its affiliates. All rights reserved.

Chapter three

Trust is built on visible foundations

"Trust is a byproduct of putting frameworks and guardrails in place. It emerges from the way the system is governed and experienced, rather than from simply asking people to believe in it."



Kathryn Giudes,
Founder and Managing Director,
ORCA Opti

Trust demands proof, not promises

Top three trust factors influencing adoption, acceptance and ongoing use:

01
Accuracy and reliability.

02
Strict privacy and responsible data use.

03
Clear human accountability and oversight.

© 2026 Cisco and/or its affiliates. All rights reserved.

Trust is built from evidence that AI works reliably, that information is protected, and that responsibility remains clear when something goes wrong. Trust can fall quickly when outputs contain plausible errors, lose important context, or cannot be verified.

Organisations place particular emphasis on dependable performance, responsible data handling, and human accountability when deciding whether AI can be relied upon. Confidence grows when users can check outputs against authoritative information and retain professional judgement over the final result.

"AI doesn't work for everything, but there are ways to get genuinely useful outcomes. People find them by breaking their own work down into tasks and testing AI on each one, to see for themselves what works and what doesn't."

– David Fischl,
Legal Digital Transformation Lead Partner, Hicksons

Confidence in AI also depends on confidence in the organisation's underlying data. Some organisations described AI projects forcing them to confront basic questions about where data sits, who owns it, who can authorise its use, and whether the data itself is reliable enough to support automation.

"People might see data as the boring plumbing work, but it's also the enabler. If you get that right, you get the trust of the company, and then you've got permission to do more adventurous work, because people have an awareness of how secure their environment is and why it works."

– Prof Cori Stewart,
Chief Executive Officer, ARM Hub

© 2026 Cisco and/or its affiliates. All rights reserved.

Chapter three

Findings:

What builds trust in practice

Trust needs to be demonstrated rather than simply declared

Policies, principles, and communications can establish expectations, but confidence is more likely to grow when people can see AI operating reliably within clear safeguards. Leaders who ‘walk the walk’ and advocate for AI from practical experience can have a powerful influence on acceptance.

“Our founder, Graham ‘Skroo’ Turner, was one of the first to say, ‘This AI thing is different, and I need to adopt it.’ When a 77-year-old founder is showing you the way, you kind of have to look at it or you look pretty silly.”

– Chris Locke,
Flight Centre Travel Group

Governance becomes credible when it is visible

Testing, monitoring, access controls, human review, and clear escalation paths turn governance into something users can experience.

“Trust in AI comes back to the tools, the guardrails, and the controls, as well as senior leadership and how AI is being deployed.”

– Anne Tosky,
Chief People Officer, ELMO Software

Human accountability needs to remain clear

People need to know where responsibility sits when AI contributes to an outcome. Confidence can weaken when accountability becomes unclear or appears to have shifted from a person to the technology.

The required assurance depends on the use case

A low-risk application may require relatively modest controls, but customer-facing, regulated, or autonomous uses will require stronger evidence of reliability, privacy protection, and accountability.

For agentic systems, assurance also needs to cover permissions, actions, escalation, and whether the system remains within the boundaries it has been given.

© 2026 Cisco and/or its affiliates. All rights reserved.

Chapter three

Recommendations:

Creating visible foundations for trust

Assign clear ownership for AI outcomes

Responsibility cannot disappear into the technology. People should know who owns an AI-supported outcome and where responsibility sits when something goes wrong.

"Trust does not attach to a model – it attaches to a named person who is answerable for the outcome."

– Sanjeev Gupta,
Chief Information and Transformation Officer,
HBF Health Insurance

Set clear boundaries for AI

Define where AI can assist, where it can act, and where human approval remains mandatory.

Prove reliability before expanding use

Trust grows when people can see consistent evidence that the AI system performs as expected. Test AI against real business scenarios, then monitor performance and be transparent about known limitations.

"My trust in AI and my ability to use it comes from the mistakes I made while learning and testing it. You learn where the limits are, what needs checking and which tasks need more supervision. That experience is what makes the trust justified rather than blind."

– David Fischl,
Legal Digital Transformation Lead Partner, Hicksons

Build confidence in the data beneath the AI

Organisations should understand where relevant data resides, who owns it, whether it is fit for purpose, and what permissions apply. Weak data foundations can undermine trust even when the AI system itself performs as designed.

Design privacy into the workflow

Make it clear what data AI can access, what information should not be entered, where data goes, and what is retained, along with whether AI services introduce data-sovereignty, regulatory, or third-party supply-chain considerations. For sensitive or strategically important data, confidence may be greater when AI operates close to where that data resides, reducing the need to move information across external platforms or third-party environments. Data classification and labelling can help ensure sensitive information is recognised and handled according to the appropriate permissions and controls. Data protection should be visible to users rather than just existing in policy documents.

Make oversight operational rather than theoretical

Users should have simple ways to question outputs, report failures, and trigger human review.

"A lot of it is about bringing in their subject matter expertise and being able to communicate where we have human intervention and where we don't. Because people understood it, that's when they felt more comfortable with it."

– Simon Goodrich
Co-founder, Portable

Use failures to strengthen trust

Failures can damage trust, but effective detection, response, and improvement can help rebuild it.

© 2026 Cisco and/or its affiliates. All rights reserved.

Privacy and security set the boundaries of trust

Chapter three

Permission depends on protection

The five most cited risks to trust in AI

01
Privacy or security breaches.

02
Incorrect or unreliable outputs.

03
Insufficient human oversight.

04
Legacy infrastructure as a risk to maintaining trust.

05
Black-box decision-making or lack of transparency.

Respondents identified concerns affecting trust that extended beyond the technology itself to the leadership, governance and controls needed to use AI safely and effectively.

© 2026 Cisco and/or its affiliates. All rights reserved.



Chapter three

Cisco's 2026 Relentless Defence Report found that 36 per cent of Australian respondents applied consistent controls to AI agents and non-human identities. This provides further context for the importance of clear permissions, oversight and accountability as organisations expand their use of AI. It found that 36 per cent of Australian respondents applied consistent controls to AI agents and non-human identities. This provides further context for the importance of clear permissions, oversight and accountability as organisations expand their use of AI.

"You can automate everything and implement a hundred agents, but who is accountable for them? You can't just say the agents did it. Someone has to be accountable, and many organisations have not formalised that yet."

– Claudine Ogilvie,
Co-founder and Chief Executive Officer, HivePix

The regulatory context within which trust is being built is also constantly shifting. Australia's privacy reforms are placing greater emphasis on how organisations collect, use, and protect personal information, in addition to requiring disclosure of AI use that impacts or influences individuals in a meaningful way. Similarly, government security guidance increasingly highlights privacy, jurisdictional, and cyber risks across third-party supply chains. For organisations adopting AI through external providers, confidence will depend on understanding both their own controls and how data and permissions are managed across the wider technology ecosystem.

Respondents were particularly concerned about AI exposing information, using data in unexpected ways, or operating without sufficient oversight. Those concerns were reinforced by examples involving excessive permissions, unexpected access to information, unreliable outputs, and AI behaviour that users did not anticipate.

Security is not simply a technical or compliance issue. It helps determine whether employees, customers and leaders are willing to trust AI at all.

© 2026 Cisco and/or its affiliates. All rights reserved.

Findings:

Where security meets trust

Chapter three



Security directly affects whether people are willing to rely on

Privacy and security influence adoption. Confidence that information will not be exposed, misused, or accessed inappropriately directly affects whether people are willing to rely on AI.

AI introduces new forms of security failure

Trust can be damaged even without a conventional cyberattack. An AI system surfacing information that users did not expect it to access, or acting outside an assumed boundary, may be enough to lead organisations to restrict use, impose additional controls, or halt deployment altogether. With agentic AI, the trust question expands from “Can I trust this output?” to “Can I trust this system to act appropriately on my behalf?”

“We think of AI not as a tool but as a coworker, so it needs access rights into systems as well. That creates a privileged-access-management challenge.”

– Paul Keen
Vice President of Engineering, Culture Amp

Permissions increasingly define the extent of acceptable AI use

Stronger assurance around identity, access, and monitoring can expand what organisations are prepared to permit AI to do, such as working with public information but not customer data, or recommending actions but not executing them. Agentic AI further raises the importance of identity and authority, as organisations need confidence not only in what an AI system can do, but in which agent is acting, whose authority it is acting under, what it is permitted to access, and whether its actions can be traced.

“Trust determines the autonomy you are prepared to give AI. The aim is bounded autonomy, where the agent can act on your behalf but cannot go beyond the boundaries you have established.”

– Devesh Maheshwari
Chief Technology Officer, Lendi Group

© 2026 Cisco and/or its affiliates. All rights reserved.

Chapter three

Recommendations:

Building trust through security

"It is an ongoing accountability check. You have to keep making visible where human intervention occurs, which use cases are low risk, which are higher risk, and what is happening to the data."



Simon Goodrich,
Co-founder,
Portable

Set clear data and access boundaries

Users should be deliberate in deciding what information AI can access and where it is stored based on purpose, sensitivity, compliance requirements and business risk. Agentic systems can inherit or accumulate permissions in ways that are not obvious from the original design. Testing needs to examine not only what an agent is intended to access, but the permissions available through the tools and sub-agents it can call. AI systems should receive only the permissions needed to perform their intended role.

"Create a clear approved-use document so people know what they can put into which tool."

– Joe Waller,
Chief Technology Officer, Finder

Monitor access and actions continuously

Logging, monitoring, and audit trails should allow organisations to determine what an AI system accessed, the actions it took and the rationale behind them, allowing the sequence of events to be reconstructed whenever necessary. As AI systems become more autonomous, organisations can strengthen trust by translating governance requirements into technical controls across the underlying network and security infrastructure that enforce permissions, detect anomalous behaviour and contain systems that operate outside approved boundaries. As the speed of AI activity increases, these controls will increasingly need to operate at machine speed.

Make protections visible

Employees and customers will have greater confidence in AI when they understand the safeguards surrounding data, access, and accountability.

Test unexpected behaviour, not just model accuracy

Testing should look for breaches of permission boundaries, unintended information exposure, and actions outside the expected workflow.

© 2026 Cisco and/or its affiliates. All rights reserved.

Today’s trust may not survive tomorrow’s ambition

Chapter three

Many organisations that are comfortable with AI today feel less certain that their current foundations are ready for more autonomous systems.

Responses pointed to a mixed picture of future readiness:

- More than half said their infrastructure requires either moderate upgrades or a complete overhaul to support the next generation of AI.
- Fewer than half were very or extremely confident that their security, governance and infrastructure foundations were ready for more autonomous AI.

Readiness for today’s AI is not readiness for tomorrow’s autonomy

This suggests a growing awareness that the infrastructure and security settings that are appropriate today may not remain so in the near future. The result is a widening gap between what organisations can deploy and what they are prepared to permit. As AI ambitions grow, governance, identity, security and accountability have to keep pace.

“AI will keep evolving, and I think that creates a trust loop. Every new capability changes what people are prepared to rely on, and the organisation then has to revisit the controls and the confidence around it.”

– Matthew Nation,
Director, Nation Partners



© 2026 Cisco and/or its affiliates. All rights reserved.

Findings:

Where the readiness gap appears

Chapter three

Current success may create false reassurance

An organisation can use generative AI successfully without having the identity, access, monitoring, and escalation capabilities required for more autonomous systems. Trusting AI to generate content for human review is different from trusting it to access systems, initiate actions, or make decisions with limited oversight.

Readiness is not purely a technology issue

Infrastructure, governance, accountability, and security need to mature together. Otherwise organisations may become technically capable of deploying systems they are not yet ready to trust.

© 2026 Cisco and/or its affiliates. All rights reserved.

Recommendations:

Retaining trust in an agentic world

Chapter three

Separate technical capability from organisational permission

The fact that an AI system can perform a task does not mean the organisation is ready to let it do so.

Assess readiness against future use, not current use

Evaluate identity, permissions, monitoring, escalation, and governance against the demands of more autonomous AI.

Strengthen controls before increasing autonomy

Security, governance, and accountability should be upgraded before organisations move into higher-risk agentic use cases. The aim is not to remove human judgement from the system, but to use controls in ways that preserve clear responsibility for consequential choices.

Define levels of autonomy and human oversight

Define where AI can advise, recommend, execute with approval or act independently, and where human review remains necessary. Controls and oversight should increase as autonomy and consequence rise.

Prepare for faster-moving risk

As AI systems act faster and across more systems, monitoring and controls will need to operate at comparable speed. As AI becomes more autonomous, organisations need to be able to reconstruct what happens when something goes wrong, retaining sufficient information regarding inputs, decisions, permissions, and actions to investigate and, where possible, replay the sequence that produced the outcome.

© 2026 Cisco and/or its affiliates. All rights reserved.

“Not everything needs the **human in the loop.**
The question is where the consequential decision sits, where a human needs to make or approve that decision, and where an agent can execute within the boundaries that have already been established.”



Devesh Maheshwari
Chief Technology Officer, Lendi Group

© 2026 Cisco and/or its affiliates. All rights reserved.

04

Conclusion



© 2026 Cisco and/or its affiliates. All rights reserved.



From trust to permission

Chapter four

The goal in building trust in AI is to create enough justified confidence for people and organisations to allow AI to take on more valuable work.

This is achieved progressively, as organisations demonstrate that AI works, that data is protected, that accountability remains clear, and that people understand when human judgement is still needed.

The organisations that scale AI successfully will be those that keep building, testing and maintaining trust alongside the technology.

© 2026 Cisco and/or its affiliates. All rights reserved.

Recommendations:

Chapter four



Start where trust can be earned

Begin with use cases where value, reliability, and risk can be demonstrated clearly. Give people opportunities to use AI in bounded environments, compare its performance with existing approaches, and understand where it succeeds and fails.

Successful experience creates evidence, which in turn creates confidence, and confidence creates permission to move into more complex or consequential uses.

Make the safeguards visible

Policies and governance frameworks matter, but people also need to experience the protections surrounding AI. Create clear policies about what information AI can and cannot access, where human review occurs, how errors are handled, and who remains accountable. The more consequential the use case, the stronger and more visible those safeguards should become. Governance should give people confidence to proceed, not simply tell them what they cannot do. Organisations should also apply the same scrutiny to external AI providers, models, and connected services as they do to internally managed systems, including how data, permissions, and downstream dependencies are managed.

Match oversight to the level of autonomy

The level of human oversight should change as AI takes on more responsibilities. Organisations can increase autonomous operations progressively as evidence and safeguards develop.

Assisted	AI generates drafts, summaries, or recommendations, with a human reviewing and approving the output before action is taken.
Augmented	AI executes lower-risk, pre-approved tasks within defined boundaries, with humans monitoring outcomes and reviewing exceptions or anomalies.
Autonomous	AI acts within tightly defined policy and permission boundaries, supported by stronger identity, access, monitoring, and containment controls.

© 2026 Cisco and/or its affiliates. All rights reserved.

Recommendations:

Chapter four

Make trust a shared responsibility

Trust cannot be delegated to just the technology team, risk function, or people and culture team, although each has an important role.

Bringing these perspectives together enables organisations to manage the different reasons employees, customers, and leaders may hesitate to trust AI.

01

Technology leaders must demonstrate reliability and security.

02

Risk and governance teams need to establish boundaries and accountability.

03

People and culture leaders need to build workers' confidence and skills and understand the impact on work.

04

Business leaders need to show value.

© 2026 Cisco and/or its affiliates. All rights reserved.

Measure whether trust is appropriate

High usage should not automatically be interpreted as high trust. Tool log-ins, seat utilisation, and similar activity metrics can show adoption, but they do not reveal whether people are confident in the system or rely on it appropriately.

Organisations can supplement usage measures with indicators such as:

Challenge and override rates	Whether users question, correct, or escalate AI outputs rather than accepting them uncritically.
Governed versus unsanctioned use	Whether employees are choosing approved AI environments or bypassing organisational controls.
Agent control coverage	The proportion of autonomous agents operating with defined identity, least-privilege access, and monitoring.
Detection and response time	How quickly material AI failures, inappropriate access, or data exposure are identified and contained.

Raise the safeguards as AI gains autonomy

The fact that an AI system can perform a task does not mean an organisation should immediately allow it to do so.

As AI gains access to more data, systems, and capabilities, organisations should strengthen identity, permissions, monitoring, accountability, and response in proportion to the autonomy granted.

Human involvement must be deliberate. Some decisions will demand human approval, while others may be safely delegated once sufficient evidence and controls are in place.

Earn trust and sustain support as AI's role grows

Trust should not be treated as a one-off hurdle that disappears after AI is deployed. Systems change, models improve, new capabilities emerge and incidents occur. Each can change what people are prepared to rely on.

Organisations need to keep demonstrating that trust remains warranted. They need to monitor performance, respond visibly when failures occur, and reassess whether safeguards remain appropriate as AI evolves.

© 2026 Cisco and/or its affiliates. All rights reserved.

Chapter four

The final word on creating trust in AI



The organisations that succeed with AI will not always be those with access to the most capable AI.

AI capability determines what an organisation can automate, but it is **trust** that determines what people and organisations are prepared to automate.

As capability increases, closing that gap becomes an organisational challenge requiring evidence, governance, security, accountability, and workforce confidence.

That is what trust ultimately provides for AI – ***the permission to scale.***

© 2026 Cisco and/or its affiliates. All rights reserved.

Thank you



© 2026 Cisco and/or its affiliates. All rights reserved.