Foundations of Effective AI Governance

Core standards and frameworks for safer AI development

AI Governance at a glance

What AI governance means

AI governance is the way you make your AI safe, clear, and under control—from idea to live use. It rests on four pillars: data governance, model lifecycle management, risk & compliance, and monitoring & auditing. Two forces sit across all of this: transparency and accountability. Together, they help you show what your model does, why it does it, and who is responsible.

Why should this matter to you?

Good governance reduces real-world harm, legal risk, and lost trust. Buyers and regulators now expect proof, not promises. Teams that can explain and evidence their models win more deals and can scale to new markets faster.

What "good" looks like (practical steps)

- **Build governance in from day one.** Prioritise as an integral part of your development planning and system lifecycle, rather than an afterthought.
- **Create lean, reusable documentation:** model cards, system cards, data lineage, impact and risk logs, change logs, audit trails. Aim to "document once, comply to many."
- **Shift-left testing and checks:** bias, safety, privacy, and security at every stage, not just pre-launch.
- **Set up continuous monitoring:** watch for drift, data poisoning, regression, and privacy issues; keep records live and up to date.
- **Make ownership explicit:** who fixes issues, who answers users, who maintains evidence, and who handles redress.

Standards and independent assurance

Use common frameworks to align teams and prove trust: ISO/IEC 42001 (AI Management Systems), NIST AI Risk Management Framework, third-party AI assurance service, plus broader ISO/NIST security and privacy baselines (e.g., ISO/IEC 27001, 27701). Independent validation boosts buyer and regulator confidence.

Global policy signals you should track

Regulators differ (EU = stricter, US = patchwork, APAC = mixed), but the asks are similar: transparency, accountability, privacy, bias testing, and documentation. OECD and UNESCO principles sit underneath many rules, so building to these makes you more "portable" across markets.

What startups and enterprises both need

Startups: ship fast, but add proof—templated docs, simple governance kits, and audit-ready evidence to enter regulated sectors.

Enterprises: break silos, move checks earlier, and add AI observability across business units and legacy stacks.

Legal and market reality

Risk isn't only "AI laws." General laws on privacy, IP, safety, and discrimination also apply. Failing on transparency and controls can lead to lawsuits and bans. Buyers have raised the bar and often switch vendors for stronger AI assurance.

Bottom line for developers

Design for governability. Keep clear records. Monitor always. Map owners. Align to known standards. Get independent checks. Do this, and you gain both compliance and performance assurance—and build a product others can trust.

Contents

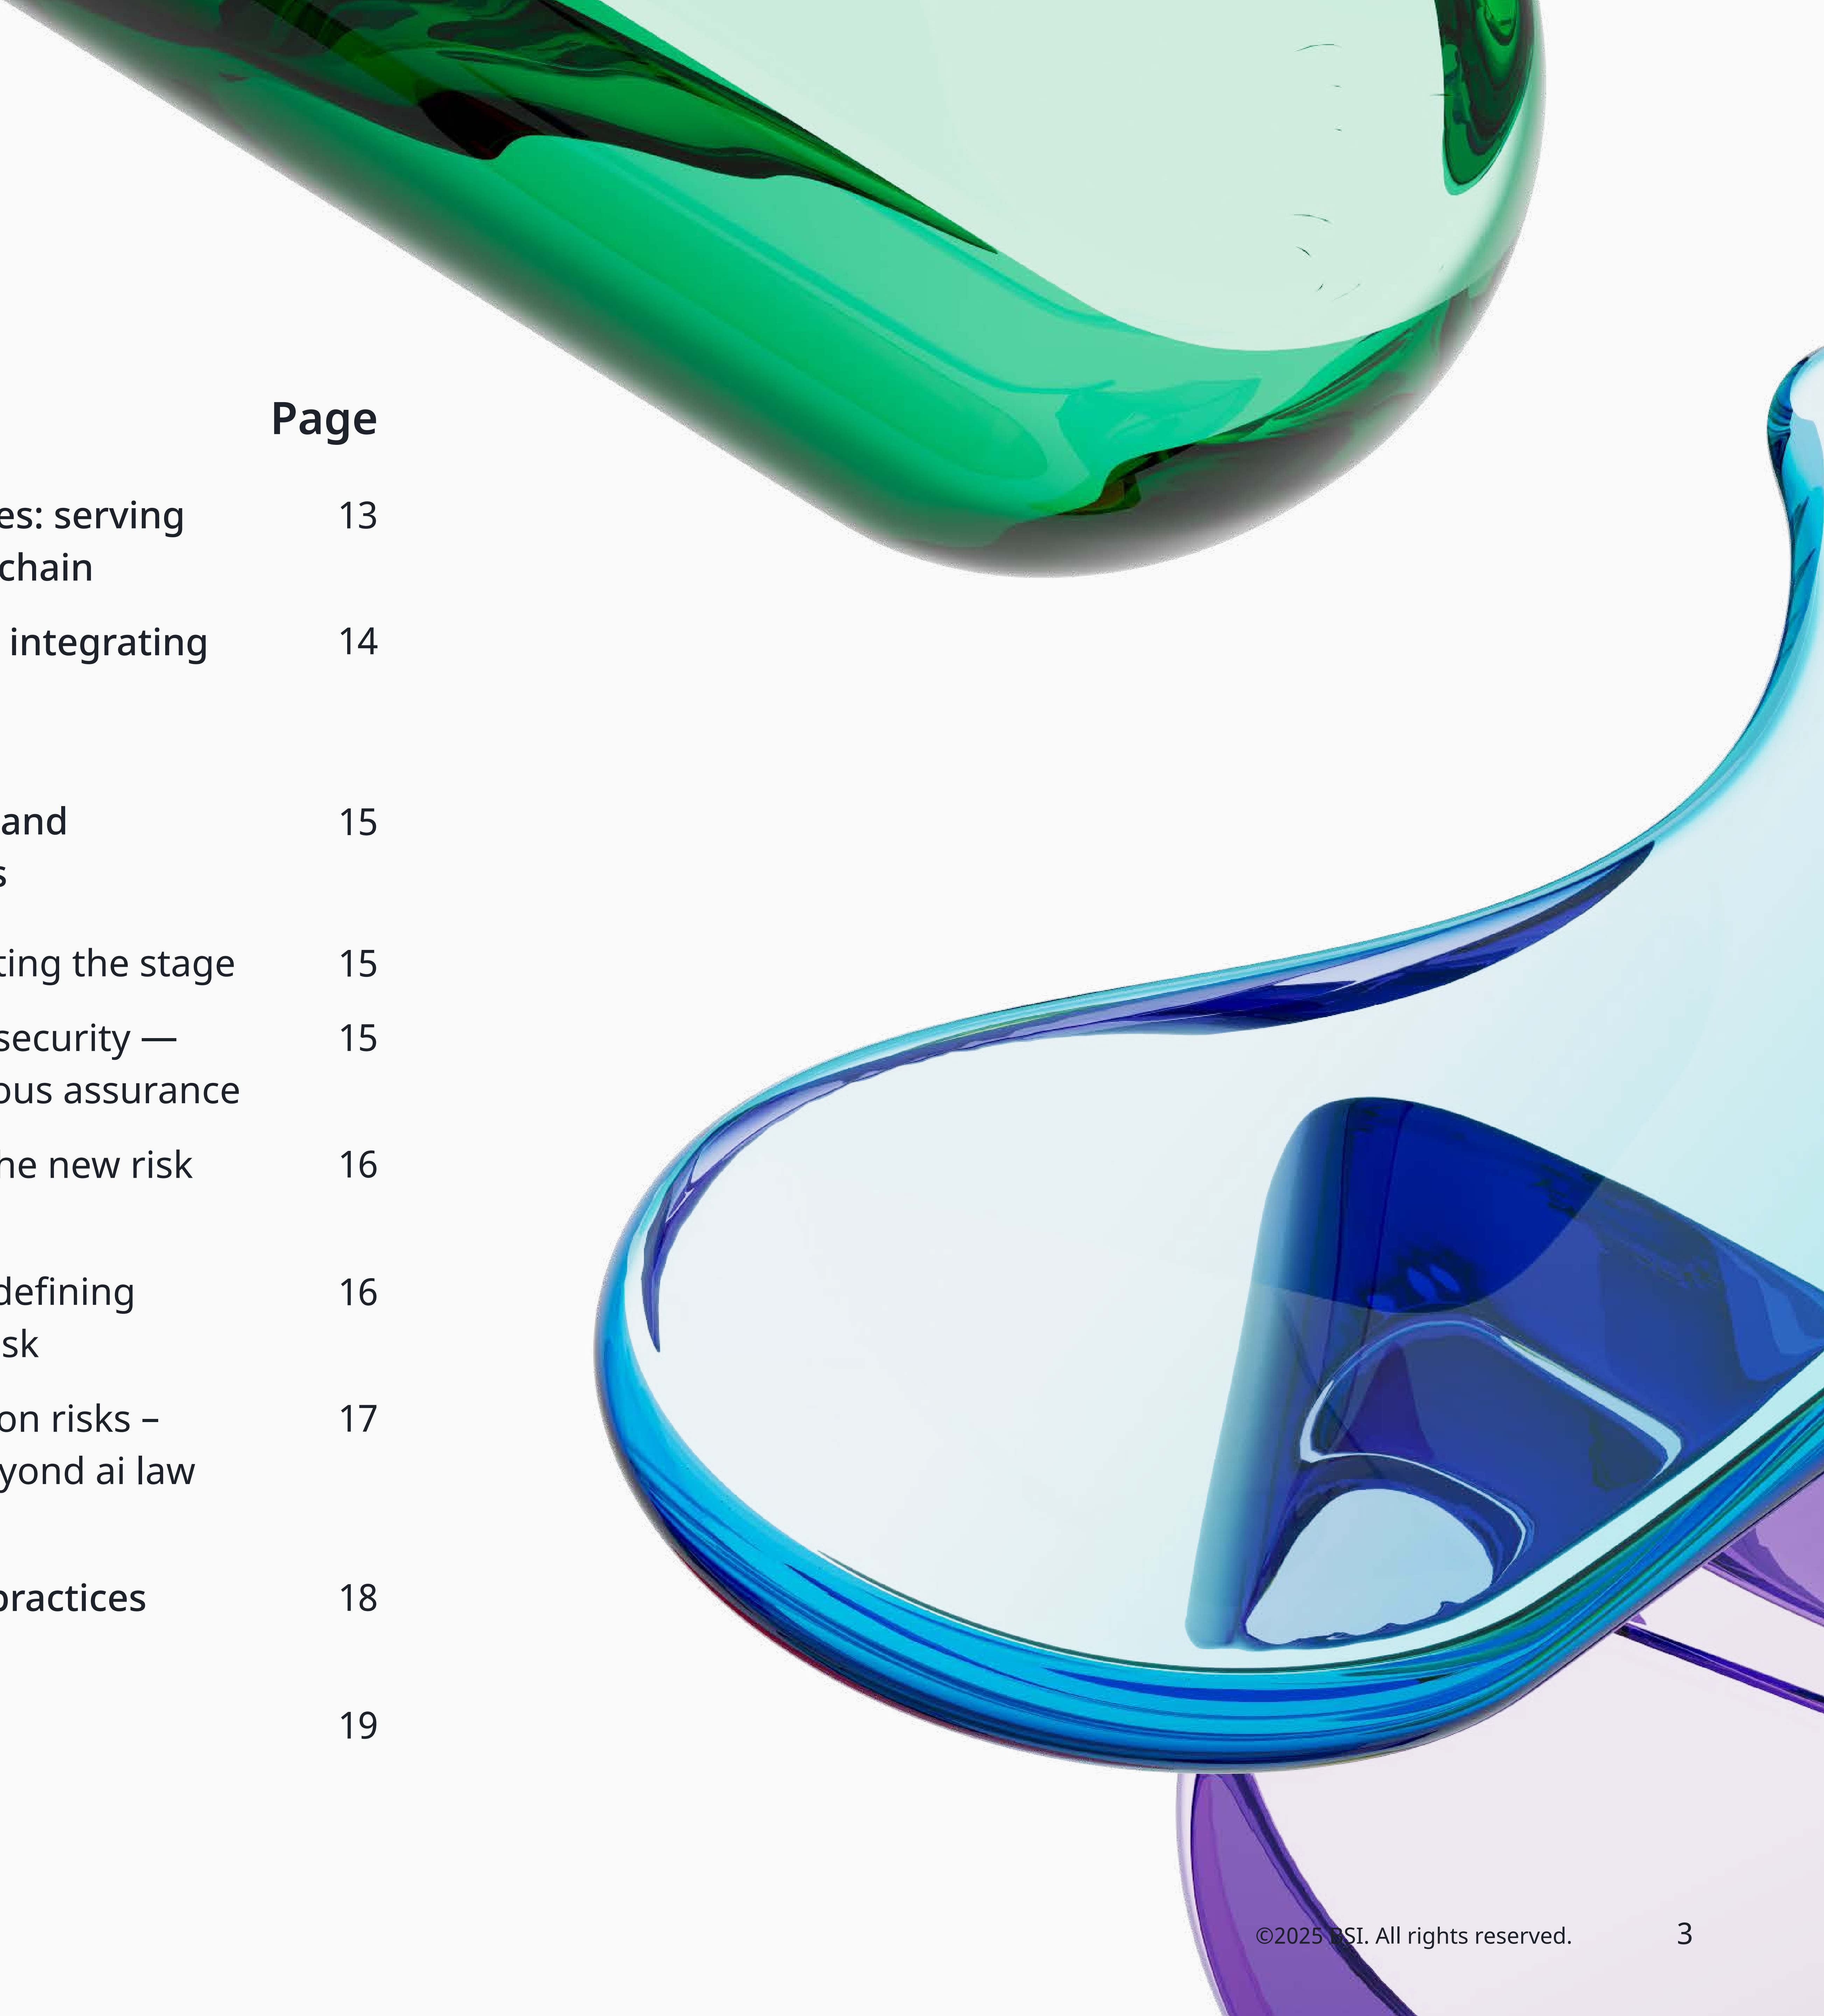

1.0 What is AI governance?

AI governance is about bringing structure, transparency, and accountability to the design, build and use of artificial intelligence (AI) models and tools. Whether it's machine learning, neural networks or other techniques, all forms of AI can benefit from robust governance. Rather than being an afterthought, this is an opportunity to provide a continuous thread that runs from the earliest ideas right through to end-user interactions. AI governance is designed to push organizations to ask tough questions: do we know what our AI is doing, can we explain it, and do we have guardrails, particularly if things go wrong?

These four pillars are supported by two overarching drivers: transparency and accountability. They permeate through every aspect of AI Governance ensuring explainability and responsibility of models and tools.

Everyone from model developers, data scientists and business leaders to system integrators, solution architects and ultimately end users need clear insights into how AI decisions are made. This includes confidence that those building technologies take appropriate responsibility. This clarity not only manages risk but also helps builds trust with clients, regulators and wider society.

Importantly, AI governance is about life cycle. It should not start and stop at launch or sale but instead be an everyday part of how teams improve models, monitor performance and respond to new challenges. This can bring huge upsides to small start-ups and multinationals alike. While the scale might differ, the expectation of good governance stretches across all functions, including risk, compliance, operations and client-facing roles.

Four core pillars

1. Data governance:

Ensuring data quality, privacy, and security.

2. Model lifecycle management:

Managing AI models from development to deployment.

3. Risk management and compliance:

Identifying and mitigating risks while ensuring compliance with legal requirements and ethical standards.

4. Monitoring and auditing:

Continuous oversight and evaluation of AI systems for model performance and compliance.

Whilst privacy and security (by design) have always been regulatory and business drivers (in high-risk sectors), there is now renewed impetus following record fines and litigations, that they are built in from day one. Organizations can't bolt on fixes after the fact. Instead, they can take the opportunity to make sure that every decision, from data collection to final outputs, protects individual rights and keeps systems secure from day one.

AI governance is a shared responsibility, ensuring that AI solutions are developed using clearly defined principles and frameworks. This approach effectively manages risks throughout the product lifecycle while providing clarity on stakeholder roles. With transparency and legal requirements in flux, both locally and globally, robust governance structures can help to demonstrate commitment to ethical and safe AI while maintaining flexibility. This can enable organizations to adapt and prove their compliance within evolving ethical and legal boundaries.

2.0 The need for AI governance

AI is shaping our world, from medical imaging and robotic process automation to smart retail and facial recognition access. Good governance is what balances smart innovation with real risks. When the rules are unclear, things can go wrong: Air Canada's chatbot giving misleading legal advice (2024) or GM's driverless cars causing harm (2024). These are not just technical errors, but real-world problems that have an immediate impact on people and lead to direct legal action.

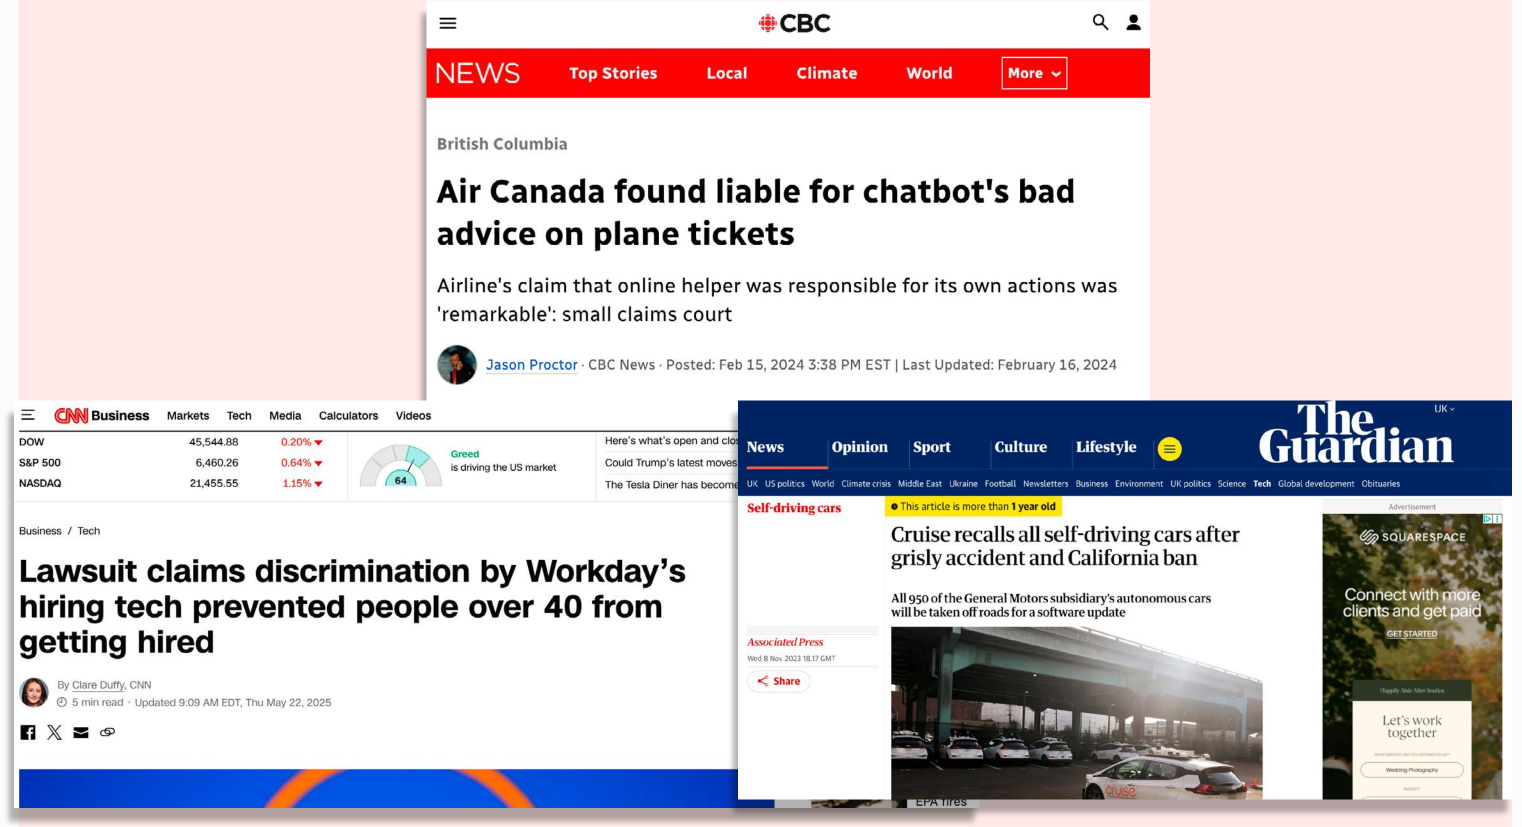

British Columbia

Air Canada found liable for chatbot's bad advice on plane tickets

Airline's claim that online helper was responsible for its own actions was 'remarkable': small claims court

Jason Proctor · CBC News · Posted: Feb 15, 2024 3:38 PM EST | Last Updated: February 16, 2024

DOW	45,544.88	0.20% ▼
S&P 500	6,460.26	0.64% ▼
NASDAQ	21,455.55	1.15% ▼

64 Greed is driving the US market

Here's what's open and clos

Could Trump's latest moves

The Tesla Diner has become

Business / Tech

Lawsuit claims discrimination by Workday's hiring tech prevented people over 40 from getting hired

By Clare Duffy, CNN

5 min read · Updated 9:09 AM EDT, Thu May 22, 2025

Self-driving cars

This article is more than 1 year old

Cruise recalls all self-driving cars after grisly accident and California ban

All 950 of the General Motors subsidiary's autonomous cars will be taken off roads for a software update

Associated Press

Wed 8 Nov 2023 18.17 GMT

Share

References:

https://edition.cnn.com/2025/05/22/tech/workday-ai-hiring-discrimination-lawsuit

https://www.theguardian.com/world/2024/feb/16/air-canada-chatbot-lawsuit

https://www.theguardian.com/technology/2023/nov/08/cruise-recall-self-driving-cars-gm

In light of these risks, the value of building trust in AI systems becomes a powerful driver for governance. Recent BSI research, found that over one in five senior leaders interviewed reported that their organizations have established an AI governance programme. Beyond risk mitigation, robust AI governance frameworks increasingly serve as proof points for private equity investors, customers and partners seeking assurance about market readiness and compliance.

Globally, groups like the OECD (Organization for Economic Co-operation and Development) and the UNESCO (United Nations Educational, Scientific and Cultural Organization) have set out widely adopted principles to address these challenges. The OECD's AI Principles focus on transparency, fairness and human rights — pushing for innovation that's explainable, safe and genuinely trustworthy. UNESCO's guidance brings in ethics, diversity and environmental care, making sure AI is not just smart, but fair and sustainable.

These foundational agreements now appear in most AI policy, regulation and best practices often embedded directly into operational standards and frameworks. For example, ISO/IEC 42001, which launched at the end of 2023, provides a certifiable management system for AI and draws on these global principles to help organisations design, develop and deploy new models responsibly. Keeping up with these evolving global rules is key. They reveal what drives regulatory thinking, helping organizations set strategy and avoid surprises in a fast-moving market.

OECD thread

The OECD's AI Principles offer a common starting point, shaping how governments and businesses think about AI risk and opportunity. Signed by 47 countries and initially adopted in 2019 with a 2024 update, these principles put human rights and societal well-being front and centre, aiming for AI that powers inclusive growth rather than division. The principles focus on transparency, fairness and safety. Importantly, they call for a fine balance, encouraging AI-driven innovation while making sure systems are explainable and risks are managed. The OECD's approach isn't just about compliance; it's about making AI genuinely trustworthy across all sectors.

UNESCO's layer

Building on this, UNESCO's global recommendations on the ethics of AI, adopted in 2021 and applicable to all 194 member states, brings in a strong human rights and environmental perspective. They underline that AI should respect dignity, support fairness, reflect diversity and safeguard the environment. Key points like proportionality, privacy throughout the AI lifecycle and inclusiveness are part of the core message. This means even as regulatory frameworks differ and new rules are written, nearly all of them share this common thread: AI should be accountable and help wider society whilst reducing environmental impact.

3.0 Legislation: EU vs US vs APAC

3.1 Introduction: the accelerating patchwork — why legislation matters

In a rapidly evolving global landscape, AI has become a transformative force with the potential to unlock immense growth and innovation. Yet, with over 70 pieces of legislation worldwide, navigating this regulatory maze presents significant challenges. While some organizations embed AI into enterprise systems, many are now harnessing AI to revolutionize product functions and redefine entire industries — much like Uber transforming taxi services.

Amidst this complexity, it’s crucial to recognize that AI doesn’t exist in a vacuum. It intersects with existing laws on cybersecurity, privacy, and consumer protection, broadening the regulatory scope. This dynamic requires organizations to adopt robust, flexible governance frameworks that can adapt as legal requirements evolve.
When addressing these challenges, regions like the EU, US, and APAC offer diverse approaches to AI regulation. The EU’s stringent requirements contrast with the US’s innovation-driven patchwork and APAC’s spectrum of regulations.

We will explore how the OECD and UNESCO provide guiding principles that serve as a global foundation for policy, standards and best practice. By grounding strategies in these principles, businesses can stay ahead of compliance, meet new benchmarks such as ISO/IEC 42001 and unlock innovation with greater confidence.

3.2 Global lessons, overlaps and practical takeaways

Across the globe, AI governance reveals a tapestry of lessons and shared principles. The EU stands as a beacon of stringent regulation, whereas the US operates as a principle-led “living lab,” and APAC showcases a diverse regulatory spectrum. Despite these differences, common compliance demands, transparency, accountability, privacy, bias testing and documentation emerge consistently.

These align with the foundational principles of the OECD and UNESCO, serving as a global glue that binds diverse regulatory landscapes and permeates through global standards. For businesses, the ability to adopt agile, flexible governance, designing for governability and employing a “document once, comply many” approach can be critical to thriving in rapidly changing environments.

48% of professionals feel confident they understand AI regulation across jurisdictions (60% C-suite).”

BSI Research

3.3 Considerations for developers and suppliers

The interconnectedness of global AI governance is driving organisations towards comprehensive, proactive practices right across the value chain, whether you’re developing, supplying or integrating AI solutions.

Practical actions for developers and suppliers:
Align with global expectations by embedding the core pillars of AI governance; Data Governance, Model Lifecycle Management, Risk Management and Monitoring & Auditing; into your processes from the outset.

Prioritise clear, minimum documentation, effective labelling and strong privacy controls early, to build transparency and accountability step-by-step. Prepare for regulatory handover and independent audit, ensuring you can demonstrate compliance at each key stage and manage contagion risks when operating in regulated sectors.

Taking these actions helps developers and suppliers stay ahead of shifting requirements, builds trust with partners and customers, and makes scaling or entering new markets much smoother.

USA

- **No single AI law:** Relies on Executive Orders, voluntary NIST guidance, sector-specific rules in key regions.
- **State-by-state approach:** Standalone state bills and agency enforcement eg California (certification shield proposal, bias audits), Colorado AI Act, Texas AI Governance Act.
- **NIST AI Risk Management Framework** widely recommended but not required.
- **Patchwork compliance:** Tech firms navigate overlapping state, sector and federal expectations.
- **Regulatory "whiplash":** Sudden changes in state law can create cost and confusion.
- **Lack of harmonisation:** No federal pre-emption, so local rules can diverge.

US AI developers face a mix of guidance and rules, using NIST and best practise as their toolkit, but must constantly monitor local changes and prepare for fast-moving updates.

Common threads in global AI governance

Transparency and accountability
Regular disclosures, fairness checks and decision traceability.

Consumer protection
Privacy by design and safety measures, such as GDPR or HIPAA.

Risk management
Continuous risk reviews and incident response, not just at launch.

Principles-based approach
Guided by OECD and UNESCO values around human rights and social benefit.

Ethical AI development
Adoption of industry codes and ISO 42001 for ethical practise.

Adaptive governance
Responsive updates and readiness for new tech and regulations.

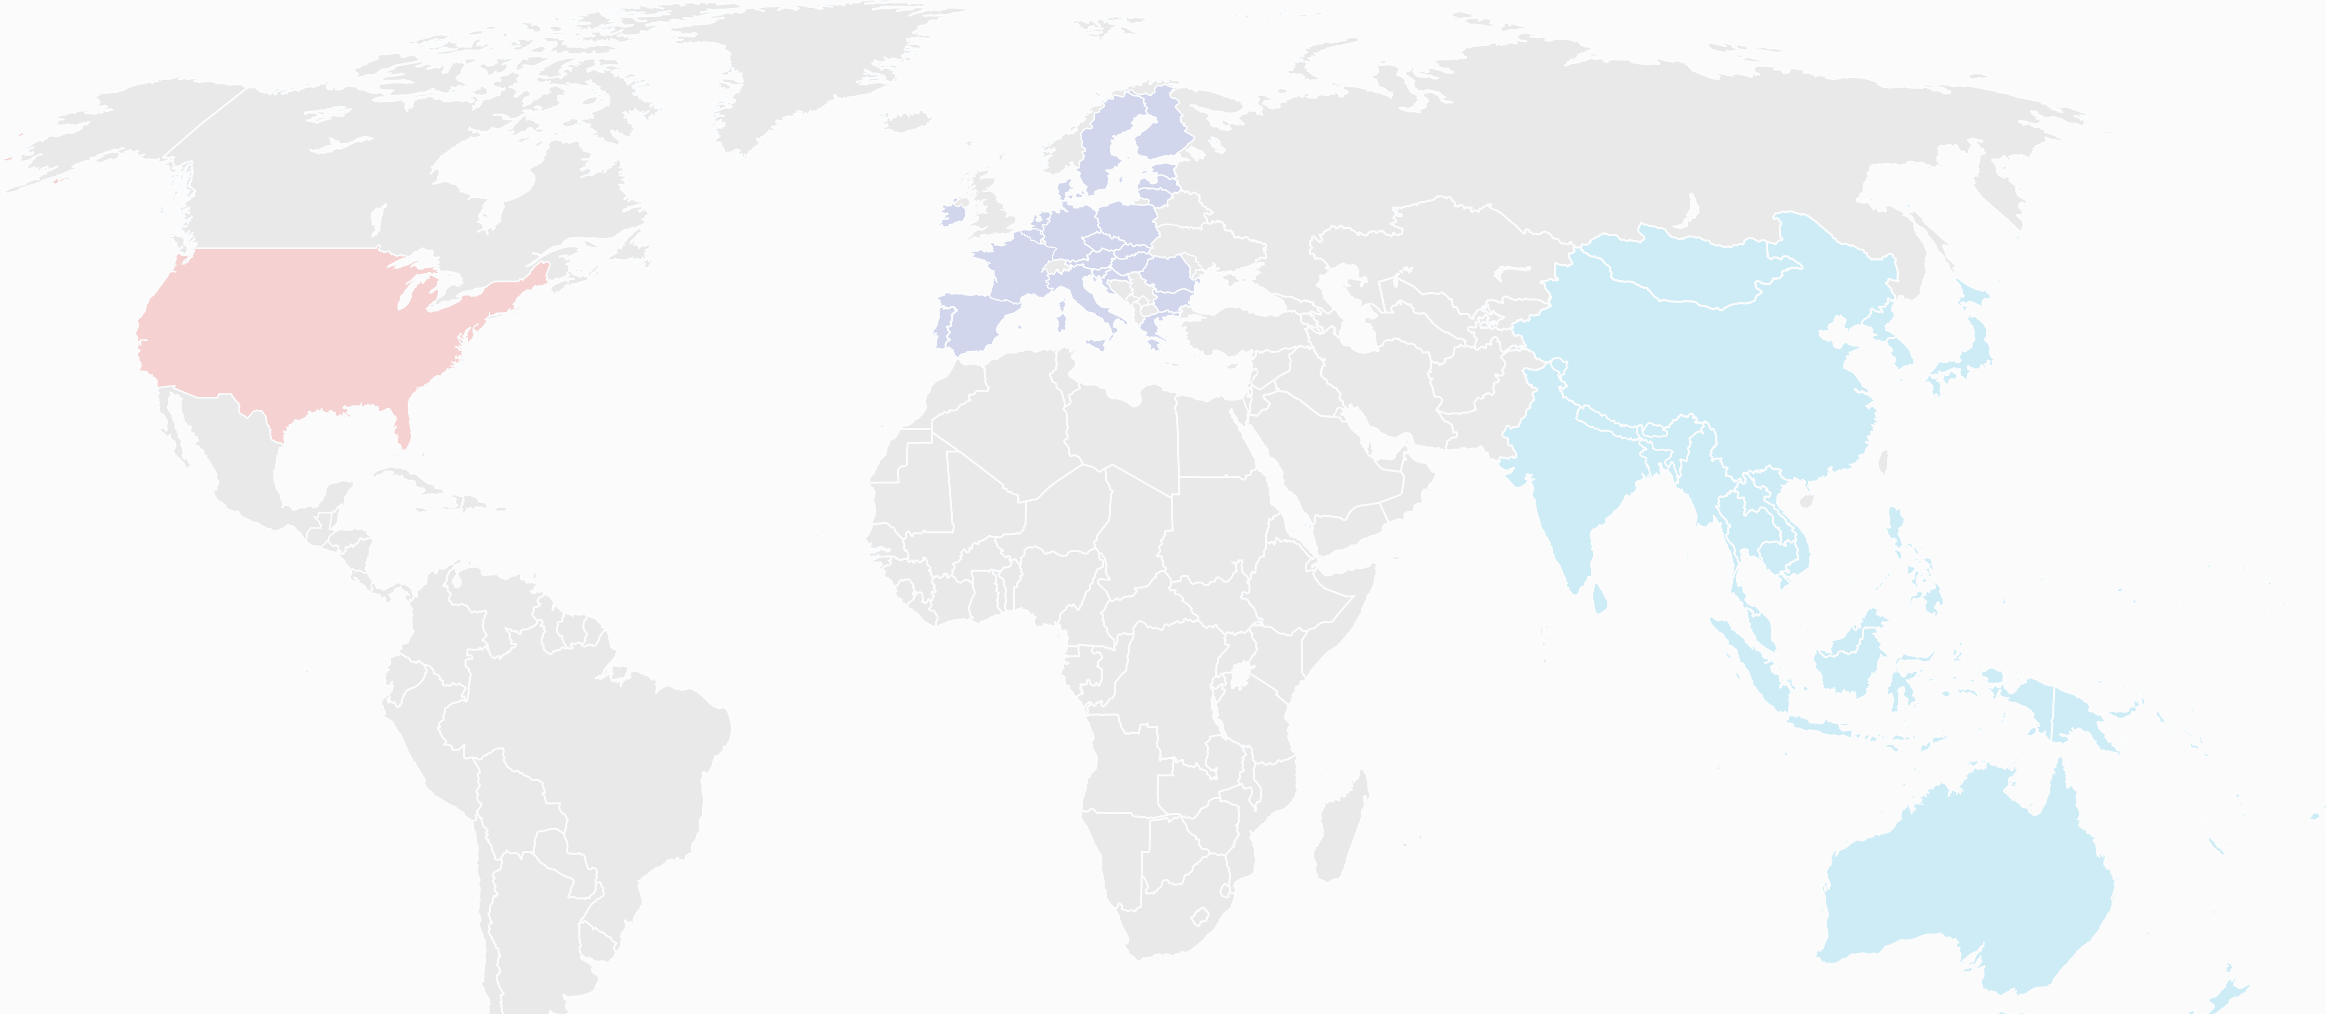

EU

- **EU AI Act** in force from August 2024.
- **Risk-based:** bans on "unacceptable" AI, high-risk system rules, general transparency for others.
- **Phased compliance:** major deadlines in 2025, most obligations live by 2026–27.
- **Push for harmonised standards** aligned to the AI Act; ISO/IEC standards development underway to define the 'how' of compliance.
- **Stringent documentation and logs:** Providers must keep detailed records, label data and outputs.
- **Rigorous conformity and oversight:** Systems face explicit assessment and ongoing reporting.
- **Adapting to national differences:** Each EU country sets up its own AI enforcement, causing uneven rollout.

Businesses suppling AI in the EU must follow strict checklists, submit to regular audits and support "privacy by design" to meet region-wide safety and rights requirements.

APAC

- **Mixed landscape:** China uses strict, targeted rules (content labelling, banned uses); Singapore, Japan and others lean on voluntary codes, guidance and sandboxes.
- **Regional standards and "soft law"** framework, eg. Singapore's' FEAT (finance), sandboxes and AI Verify Foundation toolkit.
- **Ongoing response** to OECD principles and international peer influence.
- **Compliance signals vary:** China mandates audits and labelling, SE Asia prefers trust-based frameworks.
- **Sector-driven requirements:** Banking, health and content face higher scrutiny.
- **Law in motion:** Regulation and guidance evolve quickly, especially in fast-developing economies.

APAC AI governance ranges from China's top-down enforcement to neighbour states' flexible models, meaning that tech providers often adapt their approach by market and sector while watching for sudden regulatory shifts.

4.0 Current industry challenges

4.1 Introduction: the changing risk and governance landscape

AI and digital technology are evolving at speed, bringing a mix of new models and modular tools from APIs to SDKs, many from new and early-stage vendors, into daily business. This rapid change is reshaping the risk landscape. It's no longer enough to focus just on the technical side; effective governance now has to cover practical, organizational and increasingly regulatory areas at every stage.

However, the challenge is not just a technical one. AI governance increasingly means putting people, end-user needs and real-world outcomes at the centre, which is critical for sustainable adoption. This is about doing the right thing today, that also makes strong business sense, as regulators and investors are now prioritizing these concerns and introducing new requirements that directly impact market access and reputation.

Importantly, AI governance isn't just for big companies or the tech team, it is relevant everyone, from startups to established firms. These risks and responsibilities don't sit in silos, so solutions will work best when they operate across the full value chain and at every scale.

Understanding the real challenges faced by customers, suppliers and partners is now a core part of getting governance right. Good practice is about working both up and down the supply chain, sharing responsibility and building partnerships that address each other's risks and market needs. This joined-up approach doesn't just support individual companies, or business units, it helps everyone work together to adapt and thrive as the rules and technology evolve.

4.2.1 Early-stage and emerging business: building for trust and scale

Startups and scaleups thrive on speed and ingenuity, but stepping into regulated sectors can mean trust and evidence become as important as technology. Whether it's imaging diagnostics in healthcare, biometric checks in finance or public sector monitoring, proving reliability and transparency appears essential for growth.

Moving from agility to assurance:

- Breaking into regulated industries means showing clear governance and risk management early.
- Limited staff and skills can make audits, documentation and compliance feel tough to juggle with product delivery.
- Clients increasingly expect validation through audits or certification, not just promises.

Formal governance can feel like an extra burden, especially for small teams focused on survival and innovation. Practical tools and sensible processes are needed.

Practical actions for early teams:

- Scalable templates and guides for documenting models and risks.
- Assurance and learning built into busy roadmaps.
- Quick ways to share evidence and open market opportunities.

4.2.2 The deployment dilemma for large enterprises

For big organizations, bringing in new AI isn't just a technical upgrade, it's about fitting new tools into complex systems and cultures. Legacy technology and established processes can slow innovation, especially when rolling out models, APIs or SDKs across multiple business units.

Tackling scale and complexity:

- Deploying AI across teams can require serious coordination, integrating with legacy systems often limits flexibility.
- Moving governance and risk checks earlier ("shift left") is the goal but can be hard to achieve with habitual late-stage sign-offs.
- Siloed departments lead to patchy risk management and slower progress.

There is the potential risk of governance becoming a tick-box exercise seen as someone else's job. Senior leaders also report an expectation to be spending more time handling AI compliance, distracting from broader strategy. Recognising this, standards like ISO/IEC 38507 now offer boards and executive teams practical guidance for aligning AI deployment and governance with their organisation's values, ethics and long-term goals.

Practical actions for large organizations:

- Practical, business-wide frameworks for ongoing monitoring and "AI observability".
- Tools and clear documentation that follow models through their journey.
- Ways to connect teams and embed continuous learning, avoiding both silos and unnecessary bureaucracy.

Getting these basics right helps large businesses innovate safely while adapting to shifting rules and client needs.

There are useful lessons from sectors such as Fintech for getting ahead on governance, where rapid innovation outpaced governance in the early years — leading to regulatory action and fines for poor controls, as recently seen in the UK with Monzo. Learning from this, AI teams of all sizes are encouraged to prioritise assurance as early as possible to avoid similar pitfalls down the line.

49.5% of enterprise buyers switched providers last year specifically for stronger AI features and compliance confidence.

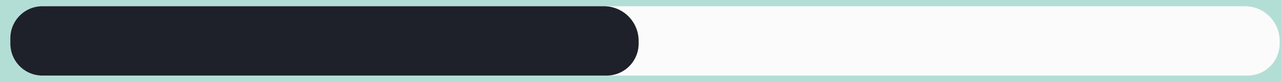

4.2.3 Market challenges: serving the whole value chain

Market access increasingly hinges on more than just effective technology: The latest Buyer Behavior Report (G2, 2025) highlights a significant shift — nearly three out of four companies now impose stricter requirements when evaluating AI-powered solutions compared to non-AI software. This includes demands for robust certification, model documentation and stronger risk management, especially as AI becomes a standard part of operational decision.

Meeting these rising expectations appears essential for both established organizations and startups, shaping who gets to do business and where.

Partnerships and the need for transparency

Strong market collaboration means suppliers and clients work together to manage data quality, model reliability and ongoing compliance. This is often achieved through joint controls, open documentation and clear pathways for assurance across three of the four main pillars:

1. Data Governance:

It's not just about an organization's data: it's about understanding and documenting how data flows down to end-users and beyond, reflecting real-world use and risk.

2. Model Lifecycle Management:

Embed reliable documentation and sound processes throughout an AI model's journey, especially for modular or evolving solutions.

3. Risk Management and Compliance:

Share responsibility for risk and compliance — both supplier and client commit to staying on top of changing obligations. According to G2, 40% of IT security and 37% of legal teams describe their requirements for AI procurement as "much stricter" than for previous software deals. (G2, 2025)

Transparency expectations are climbing fast, with buyers seeking continuous oversight and evidence trails. As regulations and market expectations keep shifting, having the ability to adapt processes and collate evidence to suit different clients and regions has the potential to turn strong governance into a shared business asset — demonstrated by the fact that 49.5% of enterprise buyers switched providers last year specifically for stronger AI features and compliance confidence. (G2, 2025)

The goal: make responsible AI manageable for your organization. Turn strong governance into something that creates value and trust across the whole value chain and keeps you ready for whatever's next.

4.3 Moving forward: integrating solutions

Effective AI governance shouldn't slow innovation, teams or companies down: it has the potential to be a significant business advantage when done right.

Practical integration actions

1. **Use modular toolkits, clear documentation and bite-sized learning to support ongoing work.** Focus on reusable artefacts that support the pillars of AI Governance, OECD/UNESCO principles and that are common to compliance standards and frameworks: things like model cards, audit trails, impact assessments and change logs.
2. **"Document once, comply many" so the same records work up and down the supply chain.** Think about shared responsibilities — make governance cross supplier and client boundaries, with clear handover of documentation. This can make adapting for new markets or audiences much easier and more robust.
3. **Have clear handover processes** to make it easy to modify governance artefacts for different markets or user groups. Continuous checks, straightforward labelling and a collaborative mindset means you can keep up with changing regulations without constant rework.

5.0 Emerging risks, safety and governance challenges

5.1 Introduction/setting the stage

Emerging risks, safety and governance challenges with AI aren't always brand new, but the speed and scale of AI development means existing risks are now presenting in new ways and with renewed importance. It's no longer enough to "just comply". To keep customers and regulators on side, responsible organizations increasingly need to evidence ongoing good governance.

5.2.1 Quality, privacy, security — towards continuous assurance

With rapid updates and evolving risks like data poisoning or model drift, it's not enough to certify a system once and move on. Continuous assurance offers an opportunity to monitor, validate and update quality, privacy and security controls throughout the AI's lifecycle. Clear documentation, strong process labelling and alignment across the whole value chain can help. Buyers and regulators increasingly expect independent validation too with 63% of senior leaders saying they would trust AI far more when it's backed by independent review (BSI Research).

Key takeaways

- **Embed continuous monitoring and validation:** Set up systems for live checks on quality, privacy and security, with alerts and periodic reviews through the full AI lifecycle.
- **Prioritize live documentation:** Keep technical records current so updates, patches and issues are always visible to the right people, not lost in version history.
- **Involve independent validation:** Bring in external reviewers regularly, checking performance, risk and compliance against not just AI-specific frameworks (like ISO/IEC 42001, NIST AI Risk Managment Framework (RMF)) but also broader frameworks for quality, security and privacy (such as ISO/IEC 9001, 27001, 27701, NIST CSF, Singapore IMDA Data Protection Trustmark SS 714:2025 or SOC21), whatever your local legal requirements.

63% of senior leaders saying they would trust AI far more when it's backed by independent review

5.2.2 Transparency – the new risk and safeguard

Transparency in AI means being able to clearly see and explain how data, decisions and processes have shaped your models. Without clear insight, teams can face growing challenges in debugging, tracing and improving models. Hidden issues can become trickier to find and fix, especially when those models are built into other products or managed by third parties.

Internal transparency means everyone, from developers to auditors, can understand what data or logic influenced the outcomes. Externally, the right level of openness builds trust with customers, partners and regulators. With high-level principles from the OECD, regulations such as the EU AI Act and ISO, IEC, CEN & CENELEC / NIST/AI Verify frameworks all demanding proof, a failure to manage, document and evidence transparency in AI systems becomes a business risk, increasing exposure to legislative challenges and impacting trust.

Key takeaways

- **Use direct labelling and traceability tools (model cards, system cards):** Make sure anyone can track what data, assumptions and logic underpin your models.
- **Ensure open processes for issue resolution:** Make transparent who is notified, who fixes issues and how feedback is captured at each stage. This can help both internally and with external confidence.

5.2.3 Accountability – defining ownership and risk

Accountability in AI goes beyond just signing off paperwork, it's about knowing exactly who is responsible when different modules, SDKs, APIs and newly deployed tools are in play. With technical components now passing between so many hands, the old ways of shared accountability or loose agreements aren't enough.

This also extends to new considerations specific to AI, such as ongoing checks for bias and discrimination in algorithms and data, and a clear need to provide pathways for contestability and redress if consumers feel they've been treated unfairly.

Good practice now means organisations should map out explicit ownership for every area of risk including responsibility for regular bias assessments, maintaining audit trails and ensuring clear escalation points for handling consumer complaints or redress. Using tools like TAIBOM to declare known biases, boundaries and intended usage upfront can help signal transparency and set expectations.

When accountability is visible, contract management, insurance and downstream compliance can become easier, and organizations can be better protected if legal or regulatory questions come up later.

Key takeaways

- **Map explicit ownership for risk and responsibility:** Define who is in charge for each part; development, maintenance, deployment and supplier relationships.
- **Clarify accountability contracts and policies:** Make sure supplier agreements, roles and insurance include who handles what risk, to be best placed if issues arise downstream.

$243 m
damages awarded by Autopilot system

$50 m
paid by Clearview AI for scraping and selling facial images without consent

5.2.4 Legal and litigation risks – why this goes beyond AI law

Legal risks in AI are shaped just as much by existing laws around negligence, liability, privacy, IP and discrimination laws as they are by emerging AI-specific regulation. Recent cases say it all: Tesla was found partially liable in a fatal crash involving its Autopilot system with $243 million in damages awarded, highlighting the complexity of deploying autonomous technology.

State Farm still faces a large lawsuit claiming its AI system discriminated against Black homeowners. DeepSeek was pulled from German app stores over weak data protection evidence and Clearview AI paid $50 million to settle claims of scraping and selling facial images without consent. These aren't just regulatory technicalities around AI, they reinforce the need for good governance, transparency or clear accountability to mitigate major business, legal and reputational fallout.

Key takeaways

- **Review against broader legal exposure and not just AI-specific law:** Include privacy, IP and discrimination in your regular risk mapping and crisis run throughs.
- **Bring legal and compliance in early:** Don't wait for AI-specific rules to catch up: review your governance for exposures under general law, including discrimination, negligence and IP. Engage legal and compliance as part of your ongoing process, not just at crisis points.

6.0 Global standards and practices

Common ground in leading AI governance frameworks

Alignment with global principles
Consistent focus on fairness, human oversight and social benefit, reflecting OECD and UNESCO values across all leading standards.

Risk management at every stage
All major frameworks require continuous risk identification, mitigation and review: helping catch problems early and adapt as systems evolve.

Transparency and accountability
Clear responsibilities, traceable decision-making and open reporting are core mandates, supporting trust for clients, regulators and society.

Comprehensive documentation and audit trails
Ongoing technical records, model cards and audit logs underpin robust governance and smooth compliance handovers to support “document once, comply many” approach.

Interoperability for global trade
Increasing crosswalks (e.g. AI Verify with ISO/IEC 42001 and NIST AI RMF) allow teams to map controls and evidence across frameworks to support exports, supply chain assurance and easier market access worldwide.

Action for developers and leaders

- Choose framework to match client and regulatory needs
- Use portable artefacts (model cards, logs, audit trails) for “document once, comply many”
- Track updates: standards and regulations evolve to ensure continuous alignment
- Consider third-party validation for stronger market position, especially in regulated sectors.

7.0 Looking to the future

AI governance is changing quickly, shaped by new tech and regulations that don't always move at the same pace. The smartest organizations now treat governance as an ongoing habit, not just a compliance task. With only a small share of companies fully embedding these practises (Trustmarque, 2025), there's real value in moving early and making governance a part of everyday decision-making.

Building governance in from the start means regular reviews, well-documented processes and early problem finding — not patching up gaps after launch. Teams that put feedback and adaptation at the centre of their work are ready for fast-changing risks, better able to reassure buyers, clients and regulators. Standards and frameworks such as ISO/IEC 42001 AI Management System, NIST AI Risk Management Framework and CSA Star for AI converge around these pillars and offer a global "language" for responsible AI. Teams use them to watch for local differences, regulatory application and opportunities for market expansion. Tools for monitoring, traceability and documentation are helping bring these frameworks to life and keep organizations ready for real-world demands.

Learning and team confidence

Governance isn't just compliance — it's a shared skill. Regular team training, short workshops and hands-on resources can help everyone get to grips with core principles (data, privacy, risk, transparency). Spreading knowledge means fewer surprises and fewer bottlenecks, so innovation and safety go hand in hand.

What does Agile governance look like?

Data stays at the centre:

Routine data checks are everyone's business, not just for the data team. By working privacy and data quality into every iteration, the whole team can understand where risk sits.

Model lifecycle steps are visible:

Progress and changes in AI models are tracked in the open so nothing slips through the cracks, from initial build to updates in production.

Risks and compliance are always in view:

Instead of waiting for problems to emerge, teams review risks and compliance issues as part of their everyday workflow. Checks for bias, discrimination or misuse are flagged as work progresses.

Continuous monitoring and audit-friendly habits:

Monitoring isn't a "sometime" activity. Automated checks and regular review points mean issues are caught early and evidence for audits builds up naturally.

Consider

Review your governance habits — are your reviews, documentation and audits regular and joined-up?

Involve independent experts early — don't wait for a problem or a contract deadline.

Bring governance learning into team meetings or sprint "retro" sessions, making it practical and ongoing.

The principles, frameworks, standards and tools discussed here offer a launchpad for ongoing, practical AI governance — flexible enough to meet today's demands and ready to adapt for whatever's next. The aim is to make governance a driver of opportunity, enabling your organization to innovate, scale and build enduring trust in a rapidly changing world.

NIST AI RMF

Scope

Voluntary risk management framework providing practical guidelines for identifying, measuring and managing AI risks. Principles-based and widely used for procurement, compliance and internal controls.

Supporting standards and resources

Includes NIST AI RMF Playbook, Roadmap, and crosswalks to ISO/IEC 42001 and CSA Star for interoperability.

Launch/status

Officially released January 2023 by National Institute of Standards and Technology (USA).

Market traction

- Widely adopted by large US tech, healthcare and defence firms. (Microsoft aligned)
- Forms the backbone for US federal procurement requirements and trusted in external audits.
- Referenced in US State and Executive Orders as the recommended baseline.

ISO/IEC 42001

Scope

The first certifiable AI management system standard, covering governance across AI lifecycles—risk, accountability, documentation and improvement. Applies across all sectors and organization sizes.

Supporting standards and resources

Includes ISO/IEC 23894 (AI Risk Management) ISO/IEC 42005 (Impact Assessment), ISO/IEC 42006 (Audit/ Certification requirements) and crosswalk alignment with ISO/IEC 27001 (security).

Launch/status

Published in December 2023 by ISO/IEC.

Market traction

- Adopted by early-mover multinationals (AWS, 2023 & Devoteam, 2025)
- Reference point for EU AI Act compliance.
- Required or expected for major contracts in finance, health and regulated supply chains.
- Third-party certification active globally.

CSA STAR for AI

Scope

STAR for AI is designed to extend CSA's established STAR (Security, Trust, Assurance, Risk) program, historically focused on cloud security, to cover AI systems. It is intended for a broad range of AI actors: model providers, platform/ orchestrator providers, application providers, and AI-using customers.

Supporting standards and resources

STAR for AI builds on several existing standards and resources, as well as new ones including an AI Controls Matrix (AICM) built on CSA's Cloud Controls Matrix. but specifically adapted for AI. The AICM is mapped to international and widely used standards, including ISO/IEC 42001, ISO/IEC 27001, NIST AI RMF and other global standards.

Launch/status

Official launch announcement for STAR for AI was made on 23 October 2025.

Market traction

Although only recently launched, Zendesk were one of the first to fully meet the prerequisites for Level 2 and a set of organizations that have signed the AI Trustworthy Pledge, though CSA's publicly listed "Committed Organizations" on their website shows a growing list.

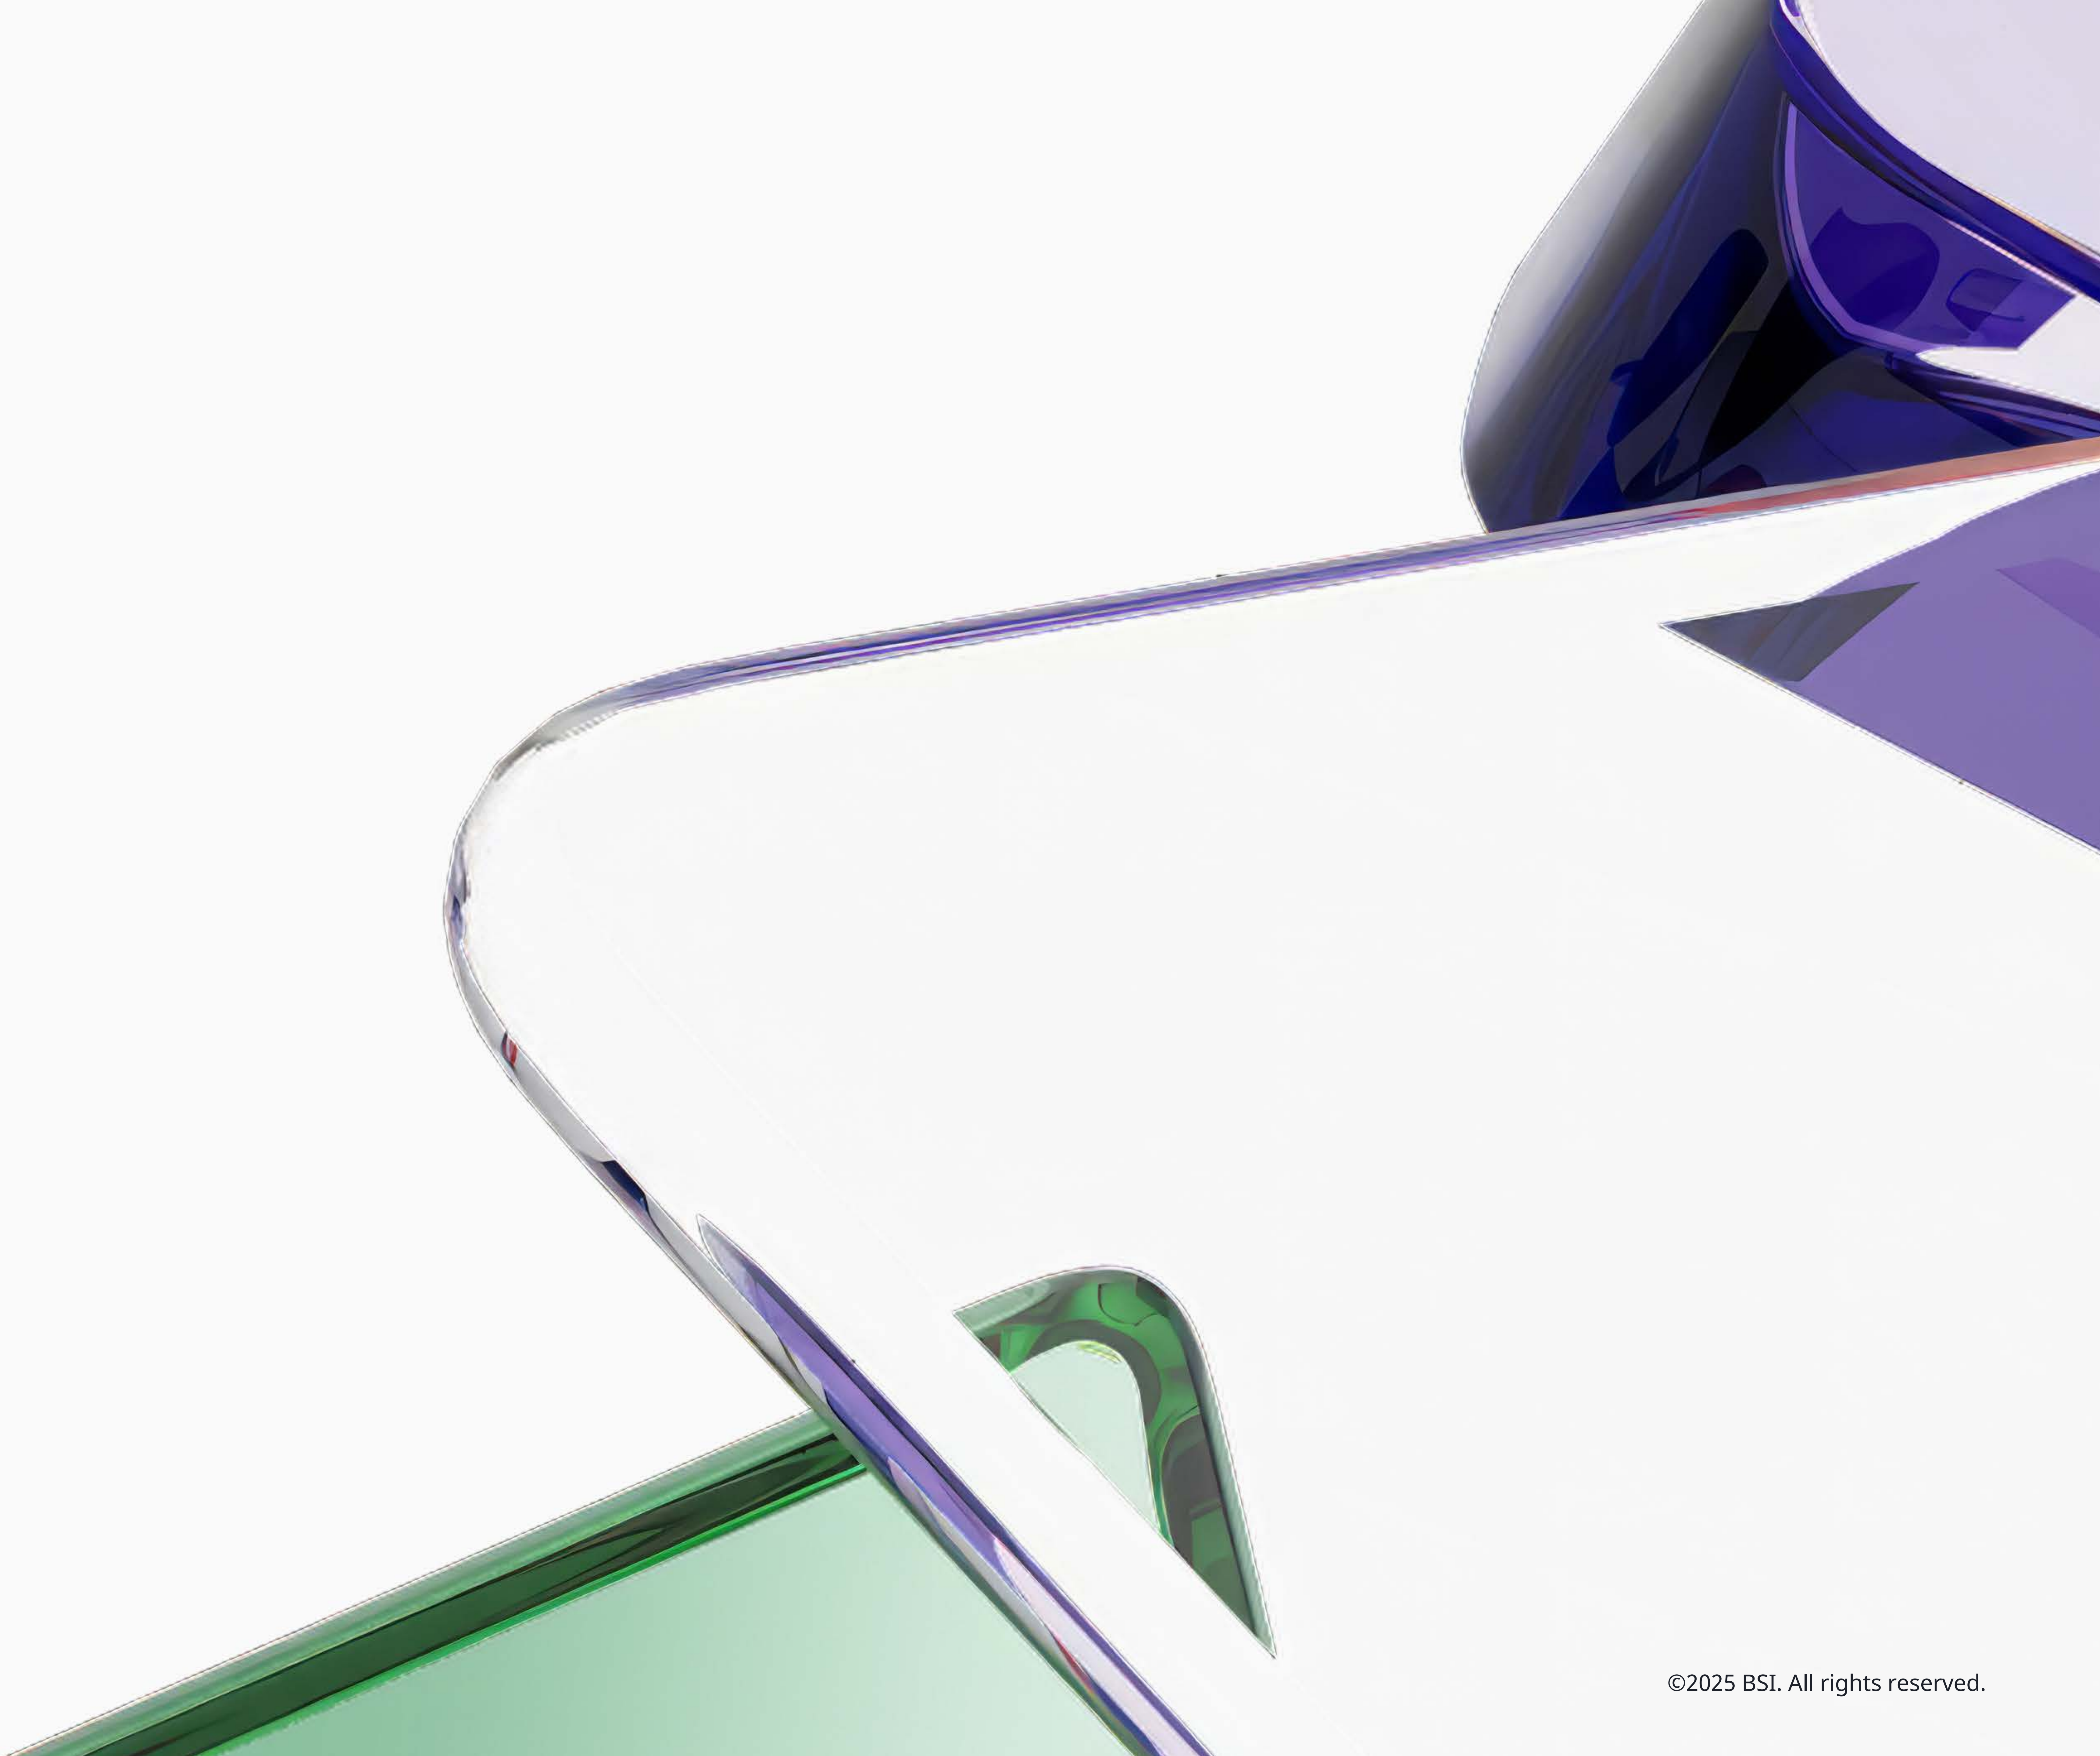

Contact us for more information at
aibsicontactpage@bsigroup.com

bsigroup.com