21 September 2026

BSA SUBMISSION TO THE JOINT SELECT COMMITTEE ON ARTIFICIAL INTELLIGENCE

Submitted Electronically to the Joint Select Committee on Artificial Intelligence

The Business Software Alliance (**BSA**)[1] welcomes the opportunity to make a submission to the Joint Select Committee on Artificial Intelligence (**Committee**) in response to its inquiry into AI in Australia (**Inquiry**).[2]

BSA is the leading advocate for the global software industry. BSA members create technology solutions that power other businesses, including cloud storage services, customer relationship management software, human resources management programs, identity management services, network infrastructure services, cybersecurity solutions, and collaboration systems. Our members are on the leading edge of providing AI-enabled products and services, and tools used by others in the development of AI systems and applications in Australia and globally. As a result, they have unique insights into the technology's tremendous potential to spur digital transformation and the policies that can best support the responsible use of AI.

BSA has engaged at each major stage of the development of Australia's approach to AI. We welcome the Committee's efforts to report on the opportunities and impacts arising from the uptake of AI technologies, and we look forward to working with the Committee as it develops its findings on a path forward for AI in Australia. The Inquiry's Terms of Reference span a broad range of issues, from AI adoption and sovereign capability to copyright, safety, and national security. Given this breadth, we invite the Committee to consider the following resources we have developed following extensive consultations with our members. We also offer specific comments on how Australia can strengthen its digital sovereignty while preserving access to global technology and data.

BSA Resources

1. **Annex A: BSA Policy Solutions for Building Responsible AI.** BSA's Policy Solutions for Building Responsible AI is a comprehensive set of recommendations for policymakers worldwide to address AI policy issues and advance the adoption of responsible AI across the economy. It lays out the issues that policymakers must address to facilitate responsible AI innovation, and proposes policy solutions on AI governance, innovation and

[1] BSA's members include: Adobe, Alteryx, Amadeus, Amazon Web Services, Asana, Atlassian, Autodesk, Avalara, Bentley Systems, Box, Cisco, Cloudflare, Cohere, Cohesity, Dassault Systemes, Databricks, Datadog, Docusign, Dropbox, Elastic, EY, Graphisoft, HubSpot, IBM, Kyndryl, MathWorks, Microsoft, Notion, Okta, OpenAI, Oracle, PagerDuty, Palo Alto Networks, PTC, Rubrik, Salesforce, SAP, ServiceNow, Shopify Inc., Siemens Industry Software Inc., Trend Micro, TriNet, Veeam, Workday, Zendesk, and Zoom Communications Inc.

[2] Terms of Reference, Joint Select Committee on Artificial Intelligence, August 2026, https://www.aph.gov.au/Parliamentary_Business/Committees/Joint/Artificial_Intelligence.

creativity, privacy, government use and procurement of AI, transparency, cybersecurity, national security, multiple development models, data innovation, research and development, and workforce development.

2. **Annex B: Global Enterprise AI Adoption Agenda.** Countries that promote secure AI adoption most effectively will see the greatest economic benefits across every industry sector. This agenda sets out proactive policies to encourage the effective adoption of trustworthy and secure enterprise AI through three pillars: 1) talent and workforce; 2) infrastructure and data; and 3) governance frameworks.

3. **Annex C: Effective Accountability Along the AI Value Chain: Right-Sizing AI Responsibilities for the Role.** Good AI governance requires assigning responsibilities to the entities best positioned to fulfil them. In this document, we explain the roles of three key actors along the AI value chain: developers, which create AI models or systems, integrators, which incorporate AI models into their own systems , and deployers, which use AI systems. Because each actor has access to different information and controls different decisions, obligations should be applied according to each entity's role rather than imposed on a one-size-fits-all basis. This ensures that responsibilities are workable in practice and that risks are identified and addressed by the entities best positioned to do so.

4. **Annex D: Artificial Intelligence and Copyright Policy & AI Training: A Key Issue for AI Policies.** The first document sets out BSA's positions on three issues pertaining to the application of copyright laws to AI: a) responsible AI training and protecting copyright holders, b) remedies if AI-generated works infringe, and c) copyright protection for creators using AI. The second explains how AI training works and why access to data for AI training is a key issue for AI policies.

5. **Annex E: Primer on Content Authenticity.** Transparency about AI-generated content is an important part of promoting responsible AI. In this primer, we explain how reliable content authentication and provenance mechanisms can help users identify the history and origin of AI-generated content, so that consumers know when content is human- or AI-generated. These mechanisms also play an important role in addressing misinformation and disinformation, including deepfakes.

6. **Annex F: Everyday AI for Businesses & Everyday AI for Consumers.** In Everyday AI for Businesses, we provide examples of business-to-business uses of AI, focusing on how companies use AI in low-risk ways to create significant benefits to both businesses and the customers they serve, such as improving logistics and document management. In Everyday AI for Consumers, we highlight how consumers rely on a wide array of services powered by AI in their day-to-day lives, such as auto-complete functions and virtual backgrounds on video calls. These everyday uses are a principal way AI lifts productivity across the economy, including for small and medium businesses.

7. **Annex G: BSA Response to Australia Productivity Commission Consultation & Comments on Productivity Commission's Harnessing Data and Digital Technology Interim Report.** These submissions respond to the Australia Productivity Commission's inquiry into harnessing data and digital technology to lift Australia's productivity. They set

out BSA's positions on AI regulation, the use of data for AI training, and privacy reform, and explain how getting these settings right will enable the adoption of data-driven and AI-enabled tools that deliver productivity gains across the economy.

In addition to the annexed resources, we hope the following submissions that BSA has filed on various AI-related policy consultations will assist the Committee in its deliberations:

Apr 22, 2022	Australia: BSA Comments on Automated Decision Making and Artificial Intelligence Issues Paper
Jul 26, 2023	Australia: BSA Comments on Supporting Safe and Responsible Artificial Intelligence in Australia
Feb 19, 2024	Australia: BSA Comments on Senate Inquiry Into Sovereign Capability
May 3, 2024	Australia: BSA Comments to Select Committee on Adopting Artificial Intelligence
Oct 3, 2024	Australia: BSA Comments on Mandatory Guardrails for AI in High-Risk Settings
Oct 25, 2024	Australia: BSA Submission on AI & Copyright Transparency
Nov 12, 2024	Australia: BSA Comments on Review of AI and the Australian Consumer Law
Dec 11, 2025	Australia: BSA Response to AGD Copyright and AI Consultation
Jun 12, 2026	Australia: BSA Comments on Guidance for Transparency in Automated Decision-Making

Promoting Digital Sovereignty in Australia

In an increasingly complex geopolitical and technology policy environment, governments and policymakers around the world are raising thoughtful questions about digital sovereignty. These questions are becoming more consequential as AI, cloud computing, and cybersecurity become increasingly important to economic growth and public services across all sectors of the economy, including agriculture, financial services, healthcare, manufacturing, and telecommunications.

Digital sovereignty should be understood as the ability to build and sustain a digital ecosystem based on four pillars: 1) capability; 2) collaboration; 3) resilience; and 4) trust. These four pillars can help governments govern, secure, operate, and recover critical digital systems while maintaining access to innovation, investment, and advanced technologies. The objective is to achieve security, autonomy, and appropriate sovereign control while preserving access to global technology and data.

Achieving these objectives depends on open and interoperable digital governance policies. Cross-border data and investment flows support jobs, innovation, and public welfare across sectors. By contrast, data and other localisation mandates, unnecessary restrictions on cross-

border data transfers, and nationality-based restrictions on trade and investment can undermine the very capability, resilience, and trust that sovereignty frameworks aim to build.

Pillar 1: Sovereignty through Investments in Capability

Capability is a foundation of long-term digital sovereignty. Governments can strengthen it by investing in human capital, technical infrastructure, research and development, and international technology partnerships. Global capability centres, startup ecosystems, academic institutions, and public-private partnerships can further strengthen domestic engineering, cybersecurity, AI, and cloud expertise while integrating local talent into global value chains. Such investments depend on factors including:

- Cross-border access to globally distributed cloud computing capacity.
- Cross-border access to knowledge, information, and data.
- Integration of AI and other transformative tools into finance, health, manufacturing, and other sectors.
- Digital upskilling for researchers, programmers, engineers, technicians, and other professionals.
- Engagement with governments and international institutions on interoperable digital governance frameworks.

Pillar 2: Sovereignty through Collaboration

Digital sovereignty depends on collaboration, which has both national and international dimensions. Domestically, governments, industry groups, standards organisations, and technical experts each contribute to secure and resilient digital ecosystems. Internationally, governments benefit from commitments to avoid arbitrary, discriminatory, disguised, or unnecessary restrictions on digital trade, cross-border data transfers, or continuity of access to technology. Outward-oriented economies demonstrate the value of building sovereign strength through openness, innovation, talent, and collaboration. International initiatives such as the Organisation for Economic Co-operation and Development's (**OECD**) Data Free Flow with Trust (**DFFT**) framework,[3] the Hiroshima AI Process (**HAIP**),[4] and bilateral and regional digital trade agreements likewise demonstrate the value of cooperation on complex digital policy challenges.

Pillar 3: Sovereignty through Resilience

Resilience is a core component of sovereignty. Critical digital systems should be able to withstand and recover from cyber incidents, natural disasters, infrastructure failures, and geopolitical disruptions.

Governments and enterprises have long used contingency planning and internationally recognized standards to benchmark readiness, redundancy, responsiveness, and recovery

[3] The DFFT framework promotes the free flow of data across borders while addressing concerns around privacy, security, and intellectual property. First proposed at the G20 Osaka Summit in 2019, the initiative is now advanced through the OECD. See: Data Free Flow with Trust, OECD, https://www.oecd.org/en/about/programmes/data-free-flow-with-trust.html.

[4] The Hiroshima AI Process was launched by G7 Leaders in May 2023 to promote safe, secure, and trustworthy AI, and has produced the International Guiding Principles and the International Code of Conduct for Organizations Developing Advanced AI Systems. See: Hiroshima AI Process, https://www.soumu.go.jp/hiroshimaaiprocess/en/index.html.

capacity. Multi-cloud architectures, hybrid deployment models, disaster recovery systems, portability norms, and open, interoperable standards can also reduce single points of failure and improve continuity.

These practices reflect a key insight: digital resilience benefits from diversity and choice across providers and technologies. Rather than concentrating risk in a single provider, technology stack, or infrastructure model, sovereignty frameworks should support choice, continuity, and contingency planning; restrictive, origin-based mandates can undermine these objectives.

Pillar 4: Sovereignty through Trust

Many sovereignty concerns are ultimately questions of trust: whether sensitive data and systems are secure, access is appropriately controlled, risks are auditable, and legal accountability mechanisms are effective. Encryption, key management, identity and access controls, auditability, contractual commitments, certification frameworks, and other technical and legal safeguards can provide these assurances without relying principally on geographic restrictions.

Recommendations

Sovereignty initiatives are more likely to achieve their objectives when built on capability, collaboration, resilience, and trust, and when they work with market incentives rather than against them. Constructive approaches include:

- Encouraging investment in capability and innovation through worker upskilling, non-discriminatory R&D incentives, support for technology entrepreneurship, high-performance computing, and publicly accessible digital infrastructure;
- Institutionalizing collaboration with governments and industry while supporting an open, pro-competitive digital ecosystem free from arbitrary, discriminatory, disguised, or unnecessary investment, trade, or data transfer restrictions;
- Using outcome-based public procurement to preserve access to best-in-class technologies;
- Promoting interoperability through voluntary, industry-driven standards for data, cybersecurity, and supply-chain risk management;
- Using encryption and privacy-enhancing technologies to protect sensitive data without unnecessary geographic restrictions;
- Building trust through governance, transparency, legal safeguards, and early-warning and consultation mechanisms between trading partners to manage cross-border service risks, including continuity of access commitments.

Governments should avoid blanket, origin-based exclusions, including prescriptive procurement mandates that undermine competition, reduce technological capacity, and limit user choice. Such mandates can impede the cross-border investment and innovation that support productivity across the economy.

Conclusion

We thank the Committee for the opportunity to provide our recommendations in response to the Inquiry. We hope that our resources and comments will assist the Committee as it develops its findings on Australia's AI policy settings. For more AI-related resources, we encourage the Committee to visit BSA's AI Resource Center, which serves as a single point of information for our policy submissions and material on AI issues.

We look forward to serving the Committee as a resource as you continue to engage in discussions on this issue. Please do not hesitate to contact me if you have any questions regarding this submission or if I can be of further assistance.

Yours sincerely,

Tham Shen Hong

Director, Policy – APAC

Annex A: BSA Policy Solutions for Building Responsible AI

300 Beach Road #30-06 The Concourse Singapore 199555
+ 65.6292.2072 **www.bsa.org**

The Software Alliance

BSA

BSA Policy Solutions for Building Responsible AI

Artificial intelligence (AI)-enabled software is helping businesses in every sector of the economy leverage the value of data to drive digital transformation. From manufacturers that use AI to design more innovative products to small businesses that rely on automated translation capabilities to grow their global customer base, AI is creating new opportunities to solve complex challenges. BSA members are at the forefront of the responsible development of AI, providing trusted software solutions that enable enterprises and their customers to harness the power of AI to improve their product offerings and enhance their competitiveness in critical areas such as health care, defense and infrastructure, and education. Rapid advances in AI are transforming expectations about how the technology may reshape the world. However, unlocking the full potential of AI will require a dynamic and flexible policy framework that spurs responsible AI innovation and use through enhanced accountability and transparency.

BSA SUPPORTS:

- Global Harmonization
- Reducing Risk Through AI Governance
- Promoting Innovation and Creativity
- Protecting Privacy
- Facilitating Procurement and Government Use of AI
- Promoting Transparency
- Enabling Cybersecurity
- Ensuring National Security
- Promoting Multiple Development Models
- Supporting Sound Data Innovation Policies
- Investing in Research and Development
- Investing in Workforce Development

www.bsa.org

Global Harmonization

Policymakers around the world are developing regulatory approaches to AI. The global nature of today's technology ecosystem demands coordinated policy responses to foster innovation.

BSA SUPPORTS

» **Pursuing interoperability.** Countries should work together to promote multistakeholder dialogue and develop a shared vision for a risk-based policy approach for addressing common AI challenges and advancing norms around responsible AI governance (e.g., risk-based approach to regulation, balanced responsibilities along the AI value chain). Global partners should also agree on common AI terminology and taxonomy, including building on ongoing work in the EU-US Trade and Technology Council.

Reducing Risk Through AI Governance

BSA members are advancing trust and ethics in AI and investing in research and development to address some of society's most pressing challenges. Organizations should ensure that society can realize the benefits of AI by proactively addressing its risks. A range of corporate governance safeguards can promote accountability by helping to identify and mitigate such risks and appropriately delineating roles and responsibilities along the AI value chain.

BSA SUPPORTS

» **Implementing risk management programs.** BSA supports implementing risk management programs to enable organizations to identify the personnel, policies, and processes necessary to manage AI risks. Elements of a risk management program may include clearly assigning roles and responsibilities, establishing formal policies, using evaluation mechanisms, ensuring executive oversight, performing impact assessments for high-risk AI, and having internal independent review mechanisms, such as interdepartmental governance or ethics committees, to evaluate and address AI issues that pose high risks. Organizations can incorporate these practices as part of a broader corporate risk management program or as a separate AI program.

» **Requiring impact assessments for high-risk uses of AI.** An impact assessment is an accountability mechanism that promotes trust by demonstrating a system has been designed and deployed in a manner that accounts for potential risks it may pose to the public. By establishing a process for personnel to document key design and deployment choices and their underlying rationale, impact assessments enable organizations to identify and mitigate risks that can emerge throughout a system's life cycle. BSA supports requirements for organizations that develop or deploy high-risk AI to conduct impact assessments and publicly affirm that they have complied with this practice. An AI system may be high-risk if it makes consequential decisions that determine an individual's eligibility for and result in the provision or denial of housing, employment, credit, education, access to physical places of public accommodation, healthcare, or insurance.

» **Distinguishing between different actors in the AI ecosystem.** Obligations should be placed on organizations based on their role in the AI ecosystem so that they can appropriately address the risks that fall within their responsibilities. For example, an AI developer, an AI deployer, and other parties within the value chain will have different information about how the AI system was developed or operates, and the law should recognize these distinctions.

» **Testing high-risk AI systems.** BSA encourages measures that incentivize safety and security. Robust testing and evaluation of high-risk AI systems for safety, security, accuracy, and fairness is critical and is prioritized in the NIST AI Risk Management Framework, which BSA supports. Existing technical standards for AI testing are nascent and should be developed consistent with longstanding voluntary, market-driven, and consensus-based approaches to standards development.

» **Ensuring appropriate policies and information sharing for foundation models.** Any public policies regulating foundation models should be commensurate with the models' risks and capabilities. Foundation model developers should provide information about model capabilities, limitations, testing, and security along the AI value chain based on the level of risk involved.

Promoting Innovation and Creativity

AI is advancing innovation and creativity in every sector of the economy. As this technology continues to evolve, it is important to consider the role of copyright law in both encouraging innovation and protecting the rights of creators. Copyright law is sufficiently flexible to adapt to this transformational technology, but BSA encourages policymakers to consider whether additional protections are warranted to prevent the spread of unauthorized, AI-generated replicas of an artist's name, image, likeness, or voice.

BSA SUPPORTS

» **Recognizing the copyrightability of works created with the assistance of AI.** AI can bolster creativity, just as other software applications have long been an important tool of artists and storytellers (e.g., photo enhancements for visual artists, visual effects in media and entertainment, and arranging music for sound recordings). Copyright plays a key role in businesses' ability to protect creative material, including software code. The use of AI should not prevent a work developed in conjunction with human creativity from being eligible for copyright protection. If copyright protection is not available simply because AI was used in the creative process, it will limit the responsible use of AI and the purpose of copyright laws. As a result, the portions of the work that are influenced by human creativity should be protected by copyright laws. Lack of copyright protection may also cause innovators to seek out jurisdictions with laws and policies that are more protective of intellectual property.

» **Adopting voluntary methods for rights holders to opt out of AI training.** Access to sufficient data for training is critical to develop AI offerings that are as accurate and insightful as possible and optimize the benefits they can provide for organizations and society. In general, AI training involves computational analysis of data to identify probabilities, correlations, and trends, a process that does not typically use any of the data for its expressive content and, therefore, does not infringe any copyright in the underlying data. However, to support artists and rightsholders, BSA encourages industry to lead the development of automated tools to indicate that a rights-owner does not want a website used for training purposes, similar to the current "do not crawl" tools that apply to search engines.

» **Recognizing the sufficiency of existing copyright law to remedy infringement.** It is important to recognize that existing copyright law is sufficient to address when a work created with the assistance of AI infringes copyrighted material.

» **Enacting legislation to help protect content creators.** BSA supports developing legislation to afford public figures, musicians, singers, actors, and other creators with a right to prevent the unauthorized dissemination or use of their name, image, likeness, or voice and the unauthorized impersonation of creators in a manner consistent with First Amendment protections.

BSA supports requirements for organizations that develop or deploy high-risk AI to conduct impact assessments and publicly affirm that they have complied with this practice.

Protecting Privacy

The data-intensive nature of AI underscores the importance of meaningful consumer privacy protections. Consumers deserve to know how their personal data is used and protected, and consumer expectations should be backstopped by strong legal obligations on companies that collect or process personal information.

BSA SUPPORTS

- **Adopting comprehensive consumer privacy laws.** BSA supports comprehensive consumer privacy laws that establish strong consumer rights in their personal data, impose clear requirements on companies that handle that consumer data, provide robust security, promote the use of data for legitimate business purposes, and are backed by robust government enforcement.
- **Providing targeted opt-outs for profiling.** Consumers should have the right to opt out of consequential, automated decisions that are made solely by AI without human interaction and that have a legal or similarly significant effect on individuals.
- **Developing privacy-enhancing technologies.** BSA encourages the development of privacy-enhancing technologies to strengthen AI safeguards. BSA recognizes that automation also plays a vital role in enabling cybersecurity tools that support data privacy. For example, AI is leveraged to help protect businesses against data breaches; protect data, devices, and networks; prevent unauthorized access to data; and improve an organization's recovery time, even after a data breach.

Facilitating Procurement and Government Use of AI

Governments leverage AI to fulfill important functions. In doing so, governments should ensure that they have access to the most advanced IT solutions and that they have processes to govern the responsible development and use of AI.

BSA SUPPORTS

- **Implementing the NIST AI Risk Management Framework (RMF).** The NIST AI RMF is a flexible framework that can help organizations govern, map, measure, and manage AI risks. Government agencies should follow the practices set forth by the NIST AI RMF, including for procurement purposes.
- **Pursuing multi-cloud procurement.** Government agencies should work with multiple cloud providers to leverage the breadth of innovation occurring across the cloud industry. Agencies should not put all their data in one cloud infrastructure, but rather leverage multiple cloud service providers' compute, AI, and other technologies. BSA supports agencies using multi-cloud in cloud purchasing.
- **Deploying AI to meet today's challenges with today's solutions.** Governments must invest in AI-driven cybersecurity solutions to bolster their defenses and keep pace with malicious actors who are already using AI to improve their exploits.
- **Enabling use of commercial sector AI applications.** AI applications are and will continue to be built into commercial software, including that procured by governments. Governments should continue to prioritize adopting commercial software and embrace trustworthy AI solutions contained within it to enhance citizen services and improve operations. This requires governments to ensure AI policies do not inadvertently prevent the government from adopting low-risk commercial AI applications.

Promoting Transparency

There has been tremendous innovation in AI, but it can also exacerbate risks of misinformation. Transparency about AI-generated content is key to ensuring responsible AI.

BSA SUPPORTS

- **Encouraging the use of watermarks or other disclosure methods for AI-generated content.** These disclosures can help consumers tell whether content is human- or AI-generated. This can be helpful in preventing misinformation. Encouraging the use of watermarks or other disclosure methods for AI-generated content can help address this concern.
- **Promoting the Coalition for Content Provenance and Authenticity standard.** BSA supports the Content Authenticity Initiative's (CAI) efforts to promote the open Coalition for Content Provenance and Authenticity standard for content authenticity and provenance. This standard will help consumers decide what content is trustworthy and promote transparency around the use of AI. In conjunction with watermarking, the CAI approach provides secure, indelible provenance.
- **Disclosing when consumers are interacting with AI.** Consumers should know when they are interacting with AI depending on the circumstances and context of use. For example, chatbots should disclose that consumers are interacting with AI instead of a human. AI vendors should be prepared to provide some measure of explainability around models and outcomes.

Enabling Cybersecurity

Strong cybersecurity risk management is critical to combatting security threats. Although malicious actors can exploit AI to create security risks, AI can also be used to dramatically enhance cybersecurity.

BSA SUPPORTS

- **Using AI to improve secure software development.** Software producers should leverage AI to improve the secure software development process, including by identifying and remediating vulnerabilities.
- **Harnessing AI to improve cybersecurity risk management.** Policymakers should ensure that cyber defenders can flexibly use AI to provide an accurate understanding of organizations' attack surfaces and improve threat detection and security outcomes.

Ensuring National Security

AI could have implications for national security. BSA recognizes the need for targeted actions to protect against AI-related national security threats. Failure to proactively develop robust AI policies at the national level could create substantial gaps in national security.

BSA SUPPORTS

- **Using narrowly tailored measures to address national security risks.** BSA supports narrowly tailored efforts to protect national security that do not unnecessarily interfere with companies' ability to conduct routine business transactions.
- **Leveraging AI to improve critical infrastructure.** BSA recognizes that AI can contribute significantly to developing and improving critical infrastructure, such as transportation. BSA supports policies that facilitate the use of AI to enhance critical infrastructure. Efforts to mitigate risks to critical infrastructure should focus on instances where there is a risk that an AI system could override human control and endanger the health and safety of individuals.

Rules that unnecessarily limit cross-border data transfers or require data localization invariably limit the insights and other benefits that AI systems can provide.

Promoting Multiple Development Models

Open source is a critical component of the AI ecosystem. It expands the AI marketplace, enhances the diversity of product offerings, promotes transparency, and enables vulnerabilities to be identified and remediated.

BSA SUPPORTS

» **Continuing the development of open source AI.** AI policies should recognize the key role that open source plays in AI development. BSA encourages rules that support both open source and proprietary systems.

Supporting Sound Data Innovation Policies

The exponential increase in data, combined with increases in remote computing power and development of more sophisticated algorithms, has fueled progress in machine learning and AI.

Capitalizing on these developments to facilitate continued advances in AI requires sound data innovation policies.

BSA SUPPORTS

» **Facilitating global data flows.** Data transfers are integral to every stage of the AI life cycle, from developing predictive models to integrating and deploying AI systems. The data used in AI systems often originates from many geographically dispersed sources, making it imperative that data can move freely across borders. Rules that unnecessarily limit cross-border data transfers or require data localization invariably limit the insights and other benefits that AI systems can provide. In addition, countries should not require algorithmic disclosure as a condition for doing business.

» **Continuing efforts to make public government data sets open and available in machine-readable digital formats.** Government-generated data is an important asset that can serve as a powerful engine for creating new jobs, promoting economic growth, and enabling innovation in AI-related technologies. Governments collect and generate vast quantities of data that offer unique insights into virtually every facet of the modern world. To enhance AI innovation, governments should continue to prioritize the release of high-value, non-sensitive government data.

Investing in Research and Development

Government research and development (R&D) spurs technological innovation that can drive long-term economic growth. Strategic investment in education, research, and technological development will be integral to developing AI technologies.

BSA SUPPORTS

» **Increasing investment in R&D.** Increased funding for R&D is essential to sparking innovation, growing high-paying jobs, and ensuring economic competitiveness.

» **Encouraging R&D cooperation.** Countries should work together to identify and support R&D challenges across borders.

Investing in Workforce Development

AI is helping to generate new jobs across industry sectors and augmenting the current workforce. AI may also impact existing jobs, and BSA supports job training and retraining programs to minimize negative impact on workforces. Countries must not only ensure that they have the STEM talent needed to develop AI innovations, but also prepare the broader workforce for a future in which virtually every job will involve an increased interaction with AI and other technologies and the need for digital skills.

BSA SUPPORTS

- **Improving access and support for STEM education.** Broadening educational opportunities, improving training programs, and ensuring the development of a diverse workforce is needed to help meet the demand for skilled STEM workers.
- **Expanding workforce training and alternative pathways.** Industry and government should invest in programs to support creating alternative pathways to the full range of AI careers; this includes those that enable workers to develop high-demand technology skills without the need for a bachelor's or graduate degree. Programs like apprenticeships, partnerships with community colleges, digital skills training and certifications, boot camps, and public service opportunities are all important gateways to helping new and mid-career workers develop in-demand skills.

ADDITIONAL RESOURCES

Artificial Intelligence & Copyright Policy

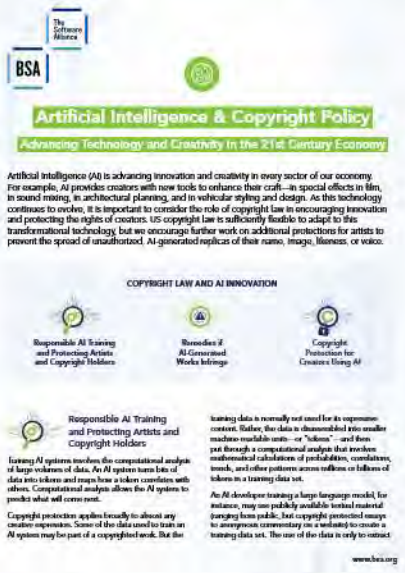

DOWNLOAD NOW

Impact Assessments: A Key Part of AI Accountability

DOWNLOAD NOW

www.bsa.org

Annex B: Global Enterprise AI Adoption Agenda

Business Software Alliance

Global Enterprise AI Adoption Agenda

The broad adoption of enterprise artificial intelligence (AI) by organizations across the economy is a strategic priority. Winning with AI is not only about creating the most powerful tools, but it is also about enabling the best tools to make sectors across the economy more successful and globally competitive.

Countries that promote secure AI adoption in the private and public sectors most effectively will see the greatest economic benefits, stimulate further innovation, and deliver substantial economic gains across every industry sector.

Proactive policies that encourage the use of trustworthy and secure enterprise AI—through talent, infrastructure and data, and workable frameworks—can ensure that these benefits are widely distributed and aligned with national economic goals.

Talent and Workforce

AI can augment human capabilities and spark greater innovation and productivity that drive economic growth and opportunity. There is no more important element than a skilled workforce to use AI across industries. A well-trained workforce enables businesses to adopt AI responsibly, securely, and effectively. Workers across industry sectors should be familiar with software tools with embedded AI and stand-alone AI solutions, and training should be tailored to level of familiarity and role within an organization. Organizations that lack an AI-prepared workforce that understands how to use AI to solve problems or identify risks will not fully realize the enormous benefits that other AI-ready organizations will reap.

RECOMMENDATIONS:

- Develop workforce training and skilling initiatives to create an AI-ready workforce and respond to the future of work. The needs will vary and may overlap based on workforce role (e.g., office productivity, data scientist, IT team, and executive leadership).
- Partner with the private sector to create AI skills curricula for educational institutions.
- Create and leverage existing public-private partnerships on workforce training and upskilling programs for current and future workers.
- Encourage skills-based hiring practices in both the public and private sectors.

Infrastructure and Data

To create an environment that promotes AI adoption, governments should reduce costs with next generation infrastructure and ensure companies have access to data, including across borders, and cutting-edge technology. Enterprise cloud adoption is a necessary prerequisite for the use of many AI tools and is a critical part of ensuring that AI can be used effectively. Data governance, sharing, interoperability, cloud-supported infrastructure, and shared solutions are also essential to enable businesses to easily adopt AI in new use cases and applications.

RECOMMENDATIONS:

- Promote investment in AI, energy, and connectivity infrastructure to strengthen their resilience and enable increased use of AI tools.
- Remove restrictions that impede the adoption of state-of-the-art cloud services and software necessary to enable enterprise AI adoption, including cross-border data restrictions and forced source code or algorithm disclosures.
- Expand access to infrastructure and data through public-private partnerships and increased availability of non-sensitive government data to catalyze AI research, training, and use in formats that can be used across industry sectors, including industrial applications.
- Encourage policies that advance technical interoperability enabling seamless integration of data and AI tools across different contexts.

Governance Frameworks

Clear regulatory frameworks and practical AI governance structures are critical to build customer confidence and promote adoption; governments can demonstrate leadership through their own AI adoption. Governments and the private sector should also prioritize high levels of enterprise security by using AI tools that protect data and mitigate security vulnerabilities.

RECOMMENDATIONS:

- Encourage organizations to establish workable internal AI governance measures and policies with clear internal review and decision-making processes.
- Promote clear national frameworks that advance privacy and cybersecurity to support AI adoption and that account for different roles, stages, uses, and risk profiles along the AI value chain.
- Prioritize high levels of enterprise security to maximize the benefits of AI adoption, including enabling the use of AI to enhance cybersecurity solutions and promoting risk-based approaches that leverage existing standards and best practices.
- Strengthen harmonization on AI policies among and within jurisdictions to ensure compatible policies that facilitate market access and reduce duplicative or fragmented obligations.
- Remove barriers to government agencies' adoption of enterprise AI, including digitizing and modernizing data formats to enable more productive uses and streamlining the procurement process for cloud and AI tools.

Annex C: Effective Accountability Along the AI Value Chain

300 Beach Road #30-06 The Concourse Singapore 199555
+ 65.6292.2072 **www.bsa.org**

Effective Accountability Along the AI Value Chain: Right-Sizing AI Responsibilities for the Role

Artificial intelligence (AI) fuels digital transformation in every sector of the economy. Manufacturers use AI to design safe and sustainable products; small businesses rely on AI-based translation tools to reach global customers; health researchers use AI to improve patient care and drive new medical breakthroughs; and companies in all industries can use AI systems to improve the accessibility of their products for people with disabilities. In these areas and countless more, AI creates new opportunities to solve complex challenges.

Successful AI adoption requires trust and confidence in these innovative technologies. As policymakers create frameworks for increasing trustworthiness and accountability, they should recognize a key principle: responsibilities should fit the role of the company.

Creating role-based responsibilities for different AI actors is important for several distinct but related reasons. First, it assigns responsibilities to companies that they can implement based on their access to relevant information and position in the AI supply chain, preventing unworkable obligations for things outside of their control. Second, it ensures that consumers and business customers are better protected because the companies that are best positioned to identify and address risks are the entities designated to take relevant steps to protect them. Third, it creates a policy framework that is workable in practice.

Key Actors in the AI Ecosystem

The AI supply chain is complex and evolving and includes a variety of companies, including AI developers, integrators, and deployers. In many scenarios, a developer may create an AI system, and a deployer will use that system in the real world.

DEVELOPERS design, code, or produce AI systems

EXAMPLE

A company that designs, codes, or produces an AI system, such as a software company that develops an AI system for speech recognition, is a developer.

DEPLOYERS use AI systems

EXAMPLE

A company that uses an AI system, such as a bank that uses an AI system either developed internally or by a third party to make loan determinations, is a deployer.

An organization can act as both a DEVELOPER and a DEPLOYER

EXAMPLE

A cybersecurity company develops AI software that monitors network traffic and customer transactions and then uses it on its own systems. The company is both a developer and a deployer.

In other contexts, one company may develop an AI model, a second company may integrate that AI model into an application or platform (AI system), and a third company may use that AI system to perform a business function, such as summarizing documents on a content management platform that is generating insights into a company's financial transaction data.

DEVELOPERS

A model developer can create a generative AI model that summarizes text or generates images.

INTEGRATORS

An integrator may incorporate that generative AI model into its own software application, enabling features such as transcribing meeting notes and summaries on a video communications platform or analyzing business contracts using document management software.

DEPLOYERS

A deployer may use the integrator's software application to assist with common business tasks, such as leveraging it to query key business documents to find answers to specific questions or generate summaries.

Developers

Developers include companies that create AI systems, including the underlying models, and sell or license the entire AI system to customers. Examples include human resource management or workflow collaboration software for businesses. In the general-purpose AI context, model developers create AI models, which can be used in a wide variety of different applications. For example, a company may develop a foundation model that can be adapted for many different use cases. The same foundation model can be used to power a range of AI systems including search engines, chatbots, spam detection software, tools that summarize long text documents, and a variety of other applications. Developers will have information about how their AI models or systems have been developed, but they generally lack information about how other companies deploy their AI tools.

Integrators

Other companies—"integrators"—may integrate an AI model into a particular application, for use by other companies. Some integrators may simply connect the AI model to a specific AI system or application, while other integrators may fine-tune or modify the AI model before including it in an AI system or application. For example, a company may act as an integrator when it develops an AI application that incorporates one or more third-party AI models. Integrators will generally have information about any changes they made to the AI model, but they usually do not have direct insight into the model's initial development or the later use of the AI system by deployers.

Deployers

Companies that use an AI tool for a specific purpose are often called deployers. These companies decide when and how to use a particular AI technology, so they will have insight into the facts of a particular use case. But deployers often obtain AI tools from other companies, so they typically lack direct insight into the initial training of the AI tool. Deployers are often the companies that have direct relationships with consumers.

Effective policy and accountability frameworks must recognize these different types of companies to avoid one-size-fits-all requirements and ensure responsibilities fit the role. This is important because the various companies along the AI supply chain will have access to different information and be in different positions to take the necessary actions to identify and address relevant risks to protect consumers and business customers.

For example, a developer of a foundation model would be able to describe issues related to the design and training of its model, but it generally would not have insight into how an integrator customizes the model for a specific application or how a deployer subsequently uses that AI system in the real world. Instead, the deployer using the AI system is generally best positioned to understand how the AI system is being used, whether that use aligns with its intended use, how to incorporate human oversight in its operation, review the outputs generated from the AI system, receive any complaints from end users, and assess real-world factors that may affect the system's performance.

Like other contexts, any legislation or policy framework that creates obligations for companies that design and use AI systems should reflect these different roles and assign responsibilities accordingly.

WHAT IS THE DIFFERENCE BETWEEN AN AI MODEL AND AN AI SYSTEM?

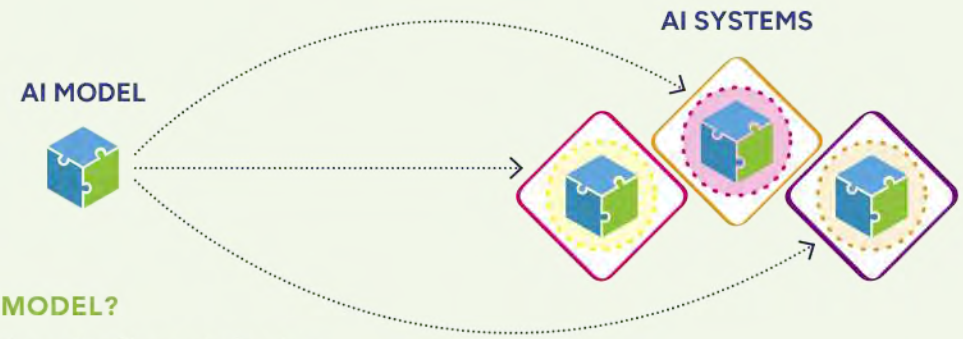

WHAT IS AN AI MODEL?

An **AI model** is a piece of software trained on large amounts of data to, for example, recognize patterns, make predictions, or generate content.

Think of it like:

- A **trained expert** who has learned from many examples, or who continues to learn from experience, or
- A **brain** that can process language, images, or numbers.

Examples of what models can do:

- Predict the next word in a sentence
- Recognize objects in an image
- Generate text, images, or code

On their own, models are usually **not user-friendly.** They don't have user interfaces, buttons, screens, or workflows.

WHAT IS AN AI SYSTEM?

An **AI system** is a complete tool or product that uses one or more AI models to perform a task or solve a problem.

Think of it like:

- A **car** that includes an engine (the model), plus steering wheel, seats, and controls.

AI applications may include a chat interface, buttons, menus, or voice input, or rules, safeguards, and integrations (e.g., with email, calendars, databases).

Examples:

- A chatbot that answers customer questions
- A photo app that enhances images
- A writing assistant that helps draft emails

PRECEDENT FOR ROLE-BASED POLICY APPROACHES

Establishing role-based obligations is considered best practice in privacy and security legislation around the world. For example, privacy laws in the United States and around the world distinguish between the differing roles of controllers that determine how and why data is processed, and processors that handle data on behalf of a controller and according to its instructions. Similarly, in various jurisdictions, cybersecurity legislation generally differentiates between companies and their service providers.

Just as privacy and security laws distinguish between different types of companies that handle consumers' personal data, distinguishing among the different companies along the AI supply chain ensures that legal frameworks appropriately assign obligations to a company based on its role in the AI ecosystem, enabling them to meaningfully fulfill those obligations and better protect consumers.

Annex D: Artificial Intelligence and Copyright Policy & AI Training: A Key Issue for AI Policies

300 Beach Road #30-06 The Concourse Singapore 199555
+ 65.6292.2072 **www.bsa.org**

Business Software Alliance

Artificial Intelligence & Copyright Policy

Advancing Technology and Creativity in the 21st Century Economy

Artificial intelligence (AI) is helping to innovate in every sector of the global economy. AI is also used to advance environmental, health, safety, and security policy objectives, including in automated flight management; early diagnosis, prevention, and treatment of cancer and disease; cybersecurity operations; predictive climate modeling; improved carbon tracking; disaster planning; productivity solutions; and supply chain management.

As this technology continues to evolve, it is important to consider the role of copyright law in both encouraging innovation and protecting the rights of creators. In common law legal systems, such as the United States, Canada, and other economies that have incorporated the doctrines of fair use or fair dealing into their laws, copyright laws have generally been sufficiently flexible to adapt to this transformational technology. In civil law legal systems, forward-thinking economies have developed statutory frameworks that permit AI training on lawfully accessed content. Some of these systems also offer additional "opt out" protections for copyright holders, allowing them to signal a preference regarding AI training. We address these topics below.

COPYRIGHT LAW AND AI INNOVATION

Responsible AI Training and Protecting the Rights of Artists and Copyright Holders

Consensus-Based Technical Mechanisms to Respect Copyright Holder Preferences on AI Training

Remedies if AI-Generated Works Infringe

Copyright Protection for Users of AI Systems

Responsible AI Training and Protecting the Rights of Artists and Copyright Holders

Training AI systems involves the computational analysis of large volumes of data. An AI system turns bits of data into tokens and maps how a token correlates with others. Computational analysis allows the AI system to predict what will come next.

Some of the data used to train an AI system may be part of a copyrighted work. Copyright protection applies broadly to almost any creative expression. But the training data is normally not used for the expressive content. Rather, the data are disassembled into smaller machine-readable units—or "tokens"—and then put through a computational analysis that involves mathematical calculations of probabilities, correlations, trends, and other patterns across millions or billions of tokens in a training data set.

For example, an AI training corpus for a large language model may contain tokenized data that draws on text elements from a part of a copyrighted work. However, the use of the data is only to extract unprotected information about the language for which the model is being trained (i.e., for, English, the correlations, patterns, and relationships among "tokens" spanning the 26 letters of the English alphabet and hundreds of thousands of English language words, as they appear in

AI AND THE CREATIVE ARTS

AI is opening new pathways for creators to develop their craft, give expression to their ideas, and develop new artistic works and performances. AI tools are helping artists enhance their creative experimentation and workflow across all art forms. This includes the use of AI-enhanced special effects and sound mixing in film and music production, as well as AI-enhanced CAD/CAM tools in sculpture, architecture, and applied arts. Writers use AI-enhanced word processing, reference, and drafting aids in the development of new literary expressions.

thousands of stock phrases, figures of speech, similes, metaphors, and common expressions). This AI training does not infringe copyright.

To incentivize beneficial AI training and improve legal certainty, governments should explore the development of statutory exceptions that permit AI training on lawfully accessed content.

Consensus-Based Technical Mechanisms to Respect Copyright Holder Preferences on AI Training

Governments should give appropriate deference to ongoing efforts by copyright holders and AI developers to develop new technological solutions for copyright holders to express their preferences regarding AI training. At the present time, there are active cross-industry consultations relating to automated tools to indicate a rights-holder does not want a website used for training purposes, similar to the current "do not crawl" tools that apply to search engines. These discussions are occurring in organizations like the Internet Engineering Task Force and the World Wide Web Consortium. Participants in these and similar discussions include representatives of the creative and technology industries, as well as civil society and academia.

Remedies if AI-Generated Works Infringe

Copyright holders should have full and effective remedies when their rights are infringed. This principle applies equally to outputs generated using AI systems and outputs generated in other ways. Copyright remedies have been effective to deter infringement and should remain so.

Governments should also consider whether there are steps they can take to better protect artists from the wrongful dissemination of unauthorized digital replicas.

Copyright Protection for Users of AI Systems

Generative AI can bolster creativity, just as other software applications have long been an important tool of artists and storytellers (e.g., photo enhancements for visual artists, special effects in audio-visual works, and arranging music for sound recordings). When generative AI is used to enhance human creativity, the resulting work should be protected by copyright. If copyright protection is not available simply because AI was used in the creative process, it will limit the responsible use of AI and the purpose of our copyright laws.

BSA PRINCIPLES ON AI & COPYRIGHT POLICY

With a view to improving legal certainty and incentivizing socially beneficial and responsible AI development, BSA supports:

- The development of statutory exceptions or judicial doctrines that permit AI training on lawfully accessed content.
- Ongoing multi-stakeholder efforts to develop effective, consensus-based technical mechanisms that allow for the expression of preferences on the collection and processing of content in AI training.
- Copyright protection of the fruits of human creativity, including when that creativity is enhanced by AI tools.
- The full application of existing copyright enforcement frameworks to AI outputs that infringe copyright.
- Additional discussions regarding protections for artists from the wrongful dissemination of unauthorized digital replicas.

Business Software Alliance

AI Training: A Key Issue for AI Policies

Technological breakthroughs over the past decade have transformed artificial intelligence (AI) from an emerging technology into a commercial reality that is empowering businesses across every sector of the economy. Airlines use AI tools to more efficiently clean planes between flights; farmers use AI to analyze large amounts of weather information to maximize their harvest; manufacturers use AI to test new prototypes; cybersecurity companies use AI to quickly identify new threats; and construction companies build AI-generated "digital twins" of real-life cities to understand the impacts of a proposed design.

All of these AI tools rely on well-trained AI models, so that the tools work properly—making it critical for AI policies to support robust training processes that can power the AI systems relied on across the economy.

How Does AI Training Work?

Think of an AI model as a student attempting to learn. One way to teach an AI model, like a large language model, is to show it a lot of data with examples and tell it the correct answer for each example. Later, the AI model may be asked to respond to each example—and then provided feedback to improve its responses. Another way is to provide the AI model with large amounts of raw data, without giving it express instructions, and allow it to identify patterns and structures across the data.

The quality of data used to train AI models is all-important. AI models will be more robust and better at achieving their desired outcomes when they are trained on broad datasets with significant amounts and different types of data, rather than narrow datasets that could be more likely to produce skewed results.

AI TRAINING IS PART OF THE BROADER LIFECYCLE OF DEVELOPING AN AI MODEL

DEFINING THE AI PROJECT

An AI model may be developed with a specific task in mind, or it may be designed as a general-purpose tool that can power a broad range of AI systems. At the outset, a design team will define the purpose and structure of the AI project.

IDENTIFYING TRAINING DATA

AI models are often trained on massive amounts of data. The data may be unlabeled, particularly if the AI model will work through it without supervision. The data may also be labeled and structured, particularly if the AI model is developed for a specific task.

PREPARING TRAINING DATA

Raw data is often cleaned and revised before it is used for training, such as by addressing duplicates, outliers, missing values and inconsistent formatting. Data is often semantically and structurally transformed into smaller units, or tokens. Training data will often ultimately appear as a sequence of tokens.

MODEL DEFINITION

After training data is processed, a developer establishes the AI system's underlying architecture. This includes identifying features in the training data that an algorithm will evaluate, as it looks for patterns and relationships as the basis of a rule to make predictions. It also selects the type of algorithmic model that will power the AI system.

MODEL VALIDATION, TESTING, AND REVISION

After an AI model is trained, it must be validated to determine if it operates as intended and tested to demonstrate that its outputs fall within what is expected. The model may then be revised based on the results.

POST TRAINING

Models are often refined and optimized after their initial training. This can involve techniques like fine-tuning and reinforcement learning. Often, post-training techniques are used to tailor a model to a specific task or improve its overall accuracy, speed, and efficiency.

How Do AI Policies Affect AI Training?

Lawmakers may consider a range of policies to ensure that companies act responsibly when they develop and deploy AI systems. Those policies should ensure that companies can, and are encouraged to, train AI models on a robust set of data, including a significant quantity of data and different types of data.

Policies can inadvertently restrict companies from robustly training a broad range of AI models, including through laws on:

1. Copyright
2. Privacy
3. Disclosure Requirements
4. Cross Border Access to Data

1 Copyright Laws

Some data in a dataset used to train an AI model may come from a part of a copyrighted work, but training a large AI model is not related to the work's creative expression. Rather, the broad dataset is broken down into small chunks of text called tokens. These tokens form the building blocks for how AI models understand and generate text. The tokens are then put through a computational analysis that involves mathematical calculations of probabilities, correlations, trends, and other patterns across millions or billions of tokens in the dataset.

The key aspect of this type of training is that data is not used for the expressive content of the work—it is used to inform the learning process. For example, by tokenizing enough sentences in which the word "interest" is used, the model can learn when it is being used to mean that the subject "wants to know about something," or if it is referring to money paid on a loan. The model may have learned that distinction by analyzing sentences that were part of copyrighted works, but the purpose is for correlation, not protectable expressions.

At the same time, as part of a flexible innovation and IP ecosystem, governments should maintain legal mechanisms, such as "notice and takedown" tools that give copyright holders the ability to seek the removal of infringing content found online, as the United States, Japan, Singapore, and other jurisdictions already do. Furthermore, AI developers and copyright holders should continue to collaborate on consensus-based, machine-readable controls that allow rightsholders to signal when they do not want AI to train on a particular website or any part of a work.

Recommendations:

- Adopt or maintain policies that ensure non-consumptive analysis of any material (including material that may come from part of a copyrighted work) is legally permitted.
- Continue to give IP holders the legal tools to remove infringing content found online.
- Encourage development of voluntary controls that permit parties to "opt out" of AI training.

2 Privacy Laws

Privacy laws limit how companies can collect and use personal data, creating important safeguards for individuals. In some cases, privacy and data protection laws also require companies to have a specific reason to process that personal data, such as for a legitimate purpose.

In many cases, privacy laws won't apply to AI training—because AI tools can often be trained on information that isn't about people. For example, an AI system trained to detect weather patterns will often ingest huge amounts of weather data. Since that data isn't about people, it is not generally protected by privacy laws. In some scenarios, though, models may be trained on more focused data to improve their performance on a particular task. This can include using personal data to fine-tune a model.

When privacy laws apply, they should not prohibit companies from training AI tools on personal data, so long as companies adopt appropriate safeguards. For example, a company using an AI tool to analyze information about people may want to fine tune that AI tool using personal data that is subject to privacy laws. An AI tool designed to route customer service complaints to different teams of employees, for instance, will work better when it is

fine tuned by the company using the AI tool to identify information about that company's customers, including the items they have purchased, their location, and any prior complaints. Using personal data to fine tune this sort of AI system can help ensure it works as intended. Privacy laws should support these sorts of responsible activities, including training AI with personal data.

Recommendations:

- Ensure personal data can be used to train AI tools, subject to appropriate safeguards.
- Promote a balanced set of privacy guardrails—and avoid requiring consent for all processing. Privacy laws should expressly permit personal data to be processed for a range of purposes, including legitimate interests, to avoid discouraging socially beneficial research or skewing data used to train AI tools.
- Recognize AI training as a legitimate purpose for processing data.
- Promote privacy-by-design and accountability frameworks for companies handling personal data in AI-related activities.

3 Disclosure Requirements

Laws can also require companies to publicly disclose information about how they develop and use AI tools. In some cases, these measures can increase trust in AI technologies. But when companies are required to disclose detailed information about how they collect and curate the datasets used to train AI models, it can threaten their ability to protect important trade secrets. For example, disclosing specific datasets an AI model is trained on and how that data is curated can require a company to disclose how it trains its models. This is particularly true for "fine tuning" that occurs after a model is developed or deployed, when specific and limited datasets may be used to refine how a model works. Not only would this disclose trade secret information, it may also give bad actors a potential roadmap for future cyberattacks.

Some disclosure proposals would also be impossible to implement, such as those that would require AI developers to disclose all copyrights and all copyright holders indirectly associated with an AI training dataset. Given that copyright arises upon creation, everything from online restaurant reviews to blogs to digital recordings may be potentially subject to copyright. However, the provenance, duration, ownership, and/or any licensing details are neither self-evident nor reasonably knowable. For example, a single work may involve multiple overlapping rights holders or "orphan" works with no clear claims. When developing or using AI at scale, it is unworkable to make granular determinations whether any part of a tokenized dataset comes from part of a copyrighted work. Moreover, these sorts of copyright-based disclosures are in tension with exceptions to copyright protections, including fair use, that generally cover AI training.

Recommendations:

- Distinguish between disclosures that are appropriately made to regulators, to users, or to the public. Focus on disclosures of high-risk uses of AI tools, rather than low-risk uses.
- Protect trade secrets by clearly recognizing that transparency obligations do not require a company to disclose trade secret information.
- Focus disclosure requirements on high-level information that helps consumers understand AI tools, without requiring granular detail about datasets used for training.

4 Cross Border Access to Data for AI Training

AI models are developed and deployed across countries and the data used to train AI models often originates from multiple sources around the world. Companies should be encouraged to use diverse sources of data, which can improve the accuracy and reliability of their AI models. Restrictions on cross-border transfers can limit access to training data, or spur companies to train AI models on narrower datasets, which can have the negative result of skewing an AI tool's output.

Using globally relevant data to train AI models is important across industries. In health care, AI tools trained on medical information from multiple countries can improve patient care and better respond to new disease outbreaks. In cybersecurity, AI tools can more easily recognize bad actors in new locations if they are trained on data about incidents from across the globe. Policies should encourage obtaining training data from multiple countries, to improve AI tools used worldwide.

Recommendations:

- Promote responsible and trusted cross-border data transfers.
- Avoid data localization requirements.
- Support harmonization across jurisdictions to enable scalable, cross-border AI development.

Data Matters. The Role of Data in AI Model Training

For an AI model to be trained properly, it needs data. If a model does not have data, it won't learn. Developers therefore carefully select and curate the datasets used to train an AI model—because data is the most critical element of AI training. Some of the factors for identifying and curating training data include:

» **Quality of data sources.** If an AI model is fed a large amount of unvetted, poor-quality data, its output is also likely to be poor quality. Identifying "good data" to train an AI tool will depend on the purpose of that AI tool.

» **Volume of data.** AI is often trained on extremely large amounts of data, to create a broad base of data points. Using large amounts of training data can help a model learn from a large number of scenarios and help it to identify outlier data.

» **Variety of data.** Training data should represent a broad cross-section of different data points, which can lead to greater accuracy in AI model training. A diverse and varied set of training data increases an AI model's depth of knowledge and helps to break down silos that may otherwise skew an AI model's output. For example, ensuring demographic and linguistic diversity can help an AI tool perform better across different populations.

How Do AI Models Learn?

Training an AI model depends on the right data. For example, if an AI model is trained to identify dogs and cats, the model's accuracy will depend on the amount and type of data it is trained on.

» **Volume of data.** The model will work better when it is given 5,000 images of dogs and 5,000 images of cats than if it were only given 50 images of each animal.

» **Quality of data.** The model will be more accurate if it is given images that show both closeups of dog and cat faces along with full-body pictures of the animals. It would be less accurate if the images show only partial photos of a paw, ear, or other body parts.

» **Variety of data.** The model will work best when trained on data that shows a range of different dog breeds and cat breeds. If the model is shown only images of dogs that are Labrador retrievers, it would have more difficulty distinguishing a cat from small dog breeds like a Chihuahua than if the model were trained on a set of data with images of all known dog breeds.

» **Specialized data.** After the model is trained, it could also be fine-tuned for more specialized uses. For example, if a person wanted to use the model to identify only exotic Bengal cats, the model could be fine-tuned by providing it hundreds or more images of Bengal cats to increase its understanding of this more specific, customized task.

TRAINING AI MODELS BENEFITS A BROAD RANGE OF AI SYSTEMS

Many AI models are built as general-purpose tools. These general-purpose models frequently become the core of a broad range of more specialized AI systems. For example, one large language AI model may be integrated into hundreds of different AI applications, such as translation apps, AI-powered chatbots, or AI systems that develop software code.

Promoting robust training for AI models benefits not just the development of that AI model, but also the many specialized AI systems built on top of it.

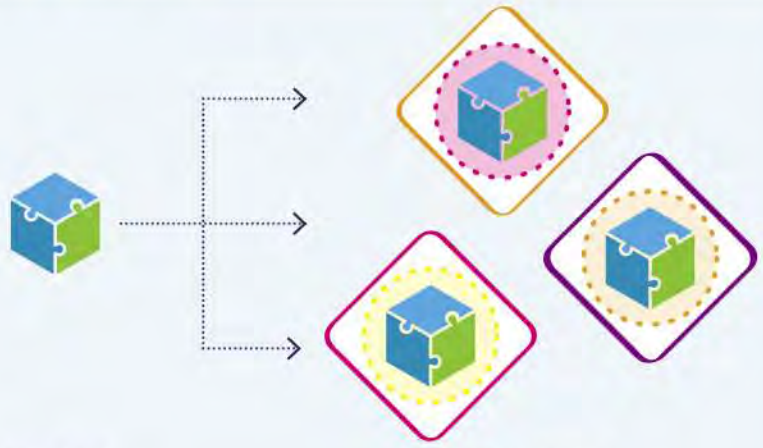

Annex E: Primer on Content Authenticity

Business Software Alliance

Primer on Content Authenticity

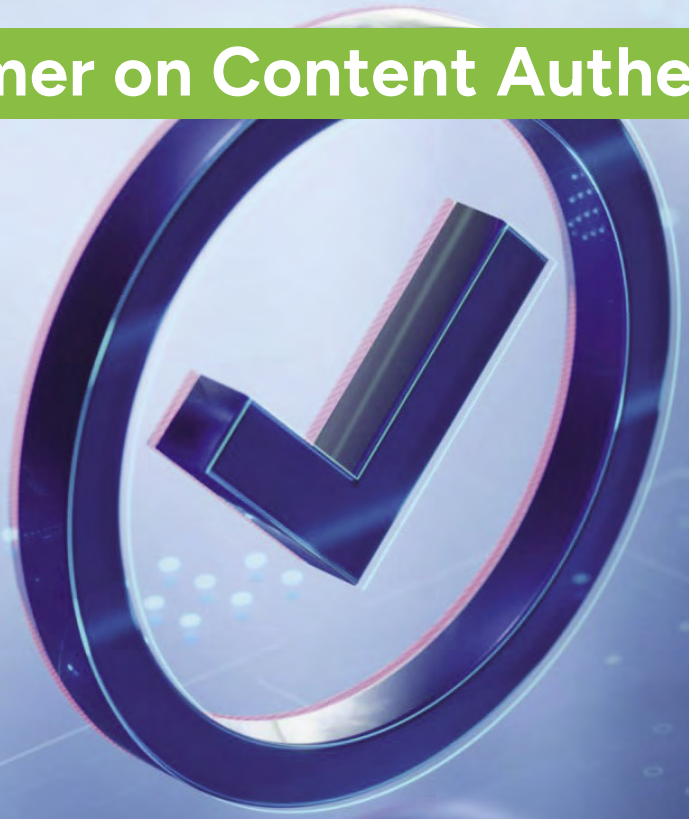

Transparency about AI-generated content is an important part of promoting responsible AI.

The Business Software Alliance supports the development and deployment of reliable content authentication and provenance mechanisms that can help users identify the history and origin of AI-generated content. This can help consumers know when content is human- or AI-generated and can help to address misinformation and disinformation.

Across the globe, policymakers are confronting important questions about how consumers can tell if the pictures and videos they see online are real. To do that, consumers need to know about tools that already exist today, like watermarks, digital fingerprints, and secure metadata. These tools can show who made a picture or video—and whether digital tools like AI have modified it. Policymakers should support the use of these and other tools to help consumers understand if content was made by humans or by AI.

BSA's primer on content authenticity:

- Defines content authenticity and content provenance, which are foundational to AI transparency policies.
- Describes existing tools that can identify who created an image or video and whether that image or video has been altered. These include machine-readable watermarks, digital fingerprints, and secure metadata.
- Describes the different actions that different types of companies can take to help support the use of content authenticity and provenance tools.
- Explores the policy issues raised by deepfakes.
- Explains how different AI policies can help consumers tell fact from fiction online.

What Are Content Provenance and Content Authentication?

Content Provenance

Content provenance information describes where content came from and how it changed. It tracks the origin, ownership, and modification history of a digital file, like an image, video, or audio clip. Example: A photograph includes cryptographically signed metadata that shows the date and location in which a photo was taken, the type of camera used to take the photo, and any software editing programs used to edit the file.

Content Authentication

Content authentication helps indicate whether the content and its provenance information are trustworthy and whether the metadata, watermarks, or credentials for a piece of content have been tampered with. Example: a content authentication tool can check a video's embedded signature to confirm it matches the issuer's public key and that the file hasn't been modified since publication.

What Tools Can Help Identify AI-Created Content?

A range of technologies can identify how content was created or modified, including content made through AI tools. Policymakers should not require use of a single technique but should ensure companies can use a range of technologies to mark content as AI-generated or not.

Machine-Readable Watermarks

Machine-readable watermarks are added directly into a file, by embedding a small amount of information that can be detected by a specialized tool. A watermark can be embedded in a file's pixels or data stream, so the origin of the file can be verified even if it is later stripped of metadata. Example: an invisible watermark inside an image can tie that image back to a creator or authenticity platform.

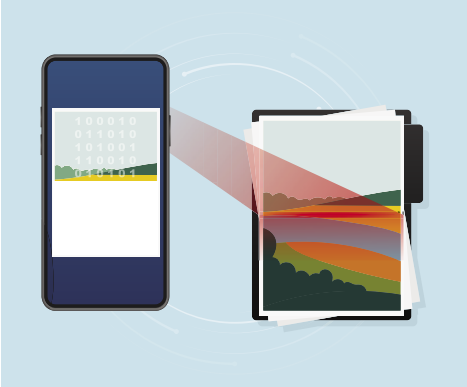

Digital Fingerprints

Digital fingerprints generate a unique identifier (or "fingerprint") from the content's core features, such as its color patterns, structure, or encoding. A digital fingerprint does not modify the file. Instead, it is a unique mathematical signature, such as a hash, that can be stored separately from the file. A file's digital fingerprint can be checked against a database of digital fingerprints, to confirm if a file matches an original or has been altered. Example: a video can be broken into frames to create a fingerprint from visual and audio cues; that signature can be checked against a database of digital signatures to confirm if the video was modified or altered.

Secure Metadata

Secure metadata can be used to store information about who created a file, what tools were used to create it, and how the file was edited. That information establishes the file's content provenance. Example: an image can include metadata that records the creator's identity, editing history, and software used; this can be digitally signed and embedded in the file so that viewers can verify its authenticity and origin.

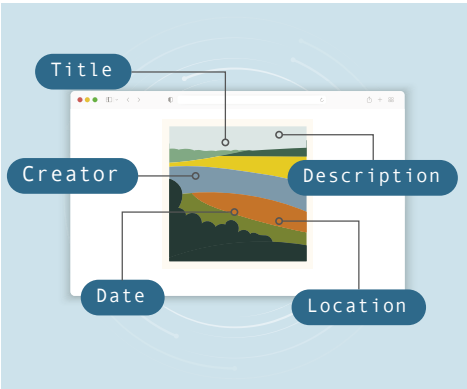

A Robust Approach: Combining Multiple Tools

These technologies can work together to support a robust approach to content authenticity, giving individuals the tools to know how content was created and whether it has been altered. All three processes act as layers of verification.

While implementing all three processes may be ideal in many scenarios, it is not always necessary to use all three techniques in every instance. Rather, policies should encourage the use of technical methods that are mature enough for adoption today while remaining flexible enough to accommodate a range of technical and operational uses.

A Bad Fit for AI: Visible Watermarks

Visible watermarks are a low-tech method used for decades to label content, including to clearly identify documents as 'Draft' or 'Confidential.' In the AI context, requiring visible watermarks is not practical, since they can be easily removed. That undermines their effectiveness and creates a false sense of security.

IN FOCUS

AI GENERATED TEXT

There are a range of tools to label images and audio-visual content, to help consumers know if that content is authentic or AI-generated. In contrast, any requirements to label AI-generated text raise a host of practical concerns. To ensure consumers know about text-based AI generated content, we recommend focusing on requirements that help individuals know when they interact with an AI system.

IN FOCUS

COALITION FOR CONTENT PROVENANCE AND AUTHENTICITY (C2PA) STANDARD

The C2PA standard can be used by anyone to incorporate digital provenance information into products and processes. It combines several technologies and is expected to be approved by the International Standards Organization as a global standard.

HOW DOES IT WORK?

1. **CREATING A FILE.** When a person creates a new file, she can use a tool that generates information about the file's provenance—including how it was created and any edits.
2. **CONTENT CREDENTIALS.** The provenance information is encoded into a content credential.
3. **CRYPTOGRAPHIC SIGNING.** The content credential is signed with the private key of the software or hardware used to create it, ensuring its authenticity. A public key is made available for authentication.
4. **EMBEDDING AND/OR WATERMARKING.** The content credential can be stored in several places, including embedding it within the file or storing it in a separate database where it is associated with an invisible watermark or digital signature.
5. **VERIFICATION.** To check the authenticity of a file, an individual can use a tool that checks content credentials to confirm if the file is authentic or has been modified.
6. **PRESENTATION.** Content that is verified as authentic can be marked with a clear indicator, like a badge or icon.

What Are Content Credentials?

The latest C2PA standard relies on content credentials, which combine multiple techniques.

- Content credentials contain metadata about a file, including the date and time it was created and the technology used to create it.
- Watermark identifiers and digital signatures can be added into the content credential, which already includes the metadata.
- The package of credentials is digitally signed and uniquely connected to that file.
- The content credentials are both embedded into the content and stored separately, such as in a database of content credentials.

The AI Supply Chain: Different Roles and Responsibilities

The AI supply chain is evolving and includes a variety of companies. Lawmakers must recognize these different types of companies as they create policies on content authenticity—since different types of companies will have access to different information and be in a position to take different actions to protect consumers.

For example, one company may develop an AI model, a second company may integrate that AI model into an application, and yet another company may use that application to create new content.

Model Developers

These companies create AI models, which can be used in a wide variety of different applications. For example, a company may build a foundation model that can be adapted for many different use cases. The same foundation model can be used to power search engines, chatbots, spam detection software, tools that summarize long text documents, and a variety of other applications. The model developer will have information about how that foundation model is developed—but generally lacks information about how other companies deploy the model for specific uses.

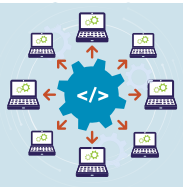

Integrators

These companies may integrate an AI model into a particular application, for use by other companies. Some integrators may simply connect the AI model to a specific service or application, while other integrators may fine tune or modify the AI model before including it in a service or application. These companies will generally have information about any changes they made to the AI model, but they usually do not have direct insight into the model's initial development or the particular circumstances in which other companies use the resulting AI application. For example, a company may act as an integrator when it develops an AI application that incorporates one or more AI models.

Deployers

Companies that use an AI tool for a specific purpose are often called deployers. These companies decide when and how to use a particular AI technology—so they will have insight into the facts of a particular use case. But deployers often obtain AI tools from other companies, so they typically lack direct insight into the initial training of the AI tool.

Any policies on content authenticity need to reflect these different roles.

If policies apply one-size-fits-all requirements to these very different types of companies, it will not effectively help consumers. For example, developers of AI applications that generate content are in the best position to apply content provenance information to content generated by that AI application. In contrast, the developer of a foundation model would lack the technical ability to apply watermarks or other digital markings to content that is generated through an AI application developed and used by other entities.

Ensuring laws remain fit over time. What constitutes state of the art in ensuring solutions for content provenance will evolve over time; policymakers should ensure that any legislative framework accommodates such developments. Embracing open standards like C2PA can help in this effort.

Combatting Deepfakes

Tools like invisible watermarks, digital fingerprints, and secure metadata can help to identify authentic content—so that individuals can know when the images, videos, and audio clips they interact with are authentic.

Together, this helps combat the problem of deepfakes. While no bad actor is likely to label their deepfakes as fake, it is important to give individuals tools to understand when content is authentic. For this to work, these tools need to be adopted across devices and platforms. People also need to be educated about these tools, so they know to watch for information about the provenance of an image or a video, while maintaining a healthy skepticism about information they encounter online.

Bad actors will continue to find novel ways to exploit technologies, including AI, for deceptive purposes. But tools like secure metadata and the C2PA standard can be crucial in helping good actors prove the authenticity of their content—helping all of us tell fact from fiction online.

As policymakers consider the most effective way to address deepfakes, they should also consider the different roles and functions of different types of companies. This is crucial since different types of companies will be positioned to address different issues, due to key service-level, technical, functional, and user-based distinctions.

Policymakers should also recognize the different risks involved in distinct types of services. For example, business-to-business software services pose limited risk to user safety and public order, given the size of their user base and the fact that they do not provide services directly to consumers, which may create fewer deepfake-related concerns.

How AI Policies Can Help Consumers

When policymakers want to ensure consumers can tell fact from fiction online, it is important to rely on existing tools that can identify if content was created by AI or not.

POLICY GOAL	SOLUTION
 **Tell consumers when they interact with AI systems.**	When companies provide AI systems that interact directly with consumers they should tell consumers they are interacting with an AI system, unless it is obvious.
 **Help consumers know if content is AI-generated.**	Focus obligations to use content provenance and authentication tools on consumer-facing audio, visual, or video content. These requirements should not apply to text or in business-to-business settings. This is important because using content authenticity mechanisms doesn't make sense for text—like when a consumer uses AI tools to edit the text of an email or translate a document. In business-to-business settings, companies can create separate methods to address content authenticity, based on their specific uses.
 **Support the use of leading global technologies to recognize AI-created content.**	Ensure requirements to identify AI-created content can be satisfied by leading global tools, including open standards like the Coalition for Content Provenance and Authenticity (C2PA). A range of countries are considering regulations on content provenance; adopting country-specific standards can limit consumers' ability to use widely accepted standards that are updated over time.
 **Facilitate transparency about AI-generated content.**	Obligations to apply content provenance information to AI-generated content are usually best placed on developers of AI applications—since AI applications are used to generate content. Companies can use tools like machine-readable watermarks, digital fingerprints, and secure metadata to identify the origin of that content.
 **Make sure content provenance information remains available to users.**	Prohibit stripping content provenance information such as machine-readable watermarks and secure metadata, absent security concerns.

Annex F: Everyday AI for Businesses & Everyday AI for Consumers

300 Beach Road #30-06 The Concourse Singapore 199555
+ 65.6292.2072 **www.bsa.org**

Business Software Alliance

Everyday AI for Businesses

Companies across industries are adopting artificial intelligence (AI) systems to improve the products and services they offer to consumers. These business-to-business (B2B) uses of AI demonstrate the many ways that organizations are already using AI to serve individuals. In many cases, companies will use AI in low-risk ways that create significant benefits—not only to the businesses but to the customers they serve. Businesses use AI systems everyday for routine tasks including:

Answering customer questions.
Businesses can offer customers 24/7 support through AI-powered chatbots that answer straightforward questions even when human customer service representatives are asleep. Bots can be programmed to address basic questions, instead of sending customers to FAQs.

Improving cybersecurity.
AI systems can sift through large volumes of information created by users of a company's IT network to forecast, detect, prevent and respond to threats. AI systems can also distill large amounts of data about security events into concrete actions to help companies secure their products and services.

Responding to frequent emails.
Companies can set up AI systems to respond to common requests—like sending automatic responses to emails asking about the status of payment invoices.

Keeping shelves stocked.
AI systems can forecast demand for products and redistribute them across a company's physical stores. AI systems can also detect early signs of supply chain issues and alert managers if inventory drops below certain levels.

Improving logistics and planning.
AI systems can improve a company's ability to forecast supply-chain issues, optimize delivery routes, estimate arrival times for new shipments, and reduce their fuel and energy usage.

Improving safety for corporate cars.
AI systems can be trained to alert employees about anomalies in corporate cars that can indicate maintenance or safety issues.

Identifying and managing common documents.
Companies can use AI tools to read hand-written documents, identify a contract based on its format, and scan files for sensitive data that needs stricter care. AI systems can then create summaries of regular corporate reports, or generate new forms based on existing examples from frequently-used documents.

Transcribing meetings—and identifying action items.
Employees can spend more time collaborating if they use an AI system to transcribe their conversation, summarize decisions, and identify follow-up items, instead of requiring an employee to serve as note-taker.

AI ACROSS SECTORS

AI systems can be used in a range of industry-specific scenarios, many of which help companies improve existing products and services.

Transportation

AI systems can improve the efficiency of airlines, by helping to pinpoint causes of any slowdowns in the process of cleaning, refueling, and reloading an airplane. Detecting these delays early helps the airline mitigate their effect on passengers.

Manufacturing

AI design tools can optimize manufacturing processes, to reduce waste and improve products. This is true from early phases, where AI can help design and test new prototypes, to factory floors where AI systems can identify maintenance and quality-control issues.

Agriculture

Farmers use AI systems to analyze large volumes of weather and crop information, helping them monitor their crops, increase yields, and adjust to rain and drought conditions.

Construction

Companies use AI to streamline the process of designing and constructing new buildings. They can also create "digital twins" of real-life cities to understand environmental and other impacts of a proposed design.

A GLOBAL APPROACH

To use these everyday AI systems, organizations generally need to bring together data they collect from many different regions, including construction sites, factory floors, or farm fields that may be located worldwide. Analyzing these different data points helps AI systems provide more accurate and reliable information.

Helping Consumers Behind the Scenes

Companies frequently use AI systems to protect consumers using their products and services, often in ways that are designed to operate in the background. For example, a range of businesses use AI systems to improve fraud detection and to protect against cybersecurity threats.

Fraud detection. Banks and credit card companies can use AI-powered systems to better detect potential fraud in real time, including by setting rules for identifying suspicious wire transfer and credit card transactions. AI systems enable companies to monitor large amounts of customer transactions to find anomalies, including detecting unusual usage patterns for a consumer's online accounts. Banks can use that information to alert customers about potential security concerns, while reducing the amount of "false positives" that may block consumers from using their own credit cards.

Cybersecurity. AI helps organizations stay a step ahead of hackers by predicting potential attacks, mitigating attacks in real-time, managing access to resources, and encrypting sensitive data. For example, a company can use an AI system to identify malicious files and suspicious IP addresses that can be easily missed by humans due to the sheer volume. In some cases, AI systems can be used to forecast, detect, prevent, and respond to threats automatically. This helps companies secure the products and services that consumers use, and protect against potential threats.

Business Software Alliance

Everyday AI for Consumers

Consumers already rely on a wide array of services powered by artificial intelligence (AI). Although these AI systems may not gain widespread attention, they illustrate how useful AI systems are in everyday life. Many of these uses present few risks to individuals while creating significant benefits, including helping organize our digital files and improve our communications with friends and family.

Do you want to pick up where you left off?

AI systems are frequently used to identify documents and other files that users recently worked on and may want to re-open. These systems can also help users locate and organize their files, such as suggesting that similar files be stored in similar locations.

Do you want to know more about that athlete?

AI systems are used to improve traditional analytics that power fantasy sports leagues, by combining inputs on sports players and teams with news articles and other sources. That creates detailed insights for sports fans, like hole-by-hole player predictions for golf tournaments.

Do you want to save time completing a form?

AI systems can auto-populate your shipping address when you order a package or create draft responses to forms that you've completed in the past.

Do you want to use a virtual background?

AI systems power the increasingly popular virtual backgrounds available on video calls. Providers use AI systems to identify the outline of an individual so the virtual background can appear in the appropriate place and follow a user as she moves across the virtual screen.

Did you forget an attachment?

For years, AI systems have been used by email providers to identify when a user may have forgotten to attach a document—and ask if something is missing.

Is it noisy during your video call?

If you join a video call from a crowded room, the video call provider may use an AI system to reduce the amount of background noise heard by others on the call—while making sure you still come through loud and clear.

How can you reach that savings goal?

AI systems can help you track your spending and budget goals, including analyzing your monthly spending habits and providing personalized recommendations for saving money.

Do you need an answer quick?

Users can interact with AI-powered chatbots to find the answers they need, instead of scrolling through an entire website or a lengthy FAQ. Chatbots can be programmed to point consumers to helpful information like a company's return policy or a list of store locations.

Annex G: BSA Submissions to the Australia Productivity Commission

300 Beach Road #30-06 The Concourse Singapore 199555
+ 65.6292.2072 **www.bsa.org**

6 June 2025

BSA RESPONSE TO PRODUCTIVITY COMMISSION CONSULTATION ON HARNESSING DATA AND DIGITAL TECHNOLOGY

Supporting safe data access and handling through an outcomes-based approach to privacy

Question: How is the Privacy Act operating to balance consumer privacy consideration while supporting the benefits associated with data sharing? Is the balance right?

BSA[1] has participated in multiple consultations on the review of the Australian Privacy Act 1988 (**Privacy Act**). As currently formulated, the Privacy Act falls short of striking the right balance between protecting an individual's privacy and supporting beneficial, responsible data use. Notably, it lacks an important distinction between data controllers and data processors, which is a foundational element of modern privacy laws globally.

In brief, "data controllers" refer to entities which determine the purposes and means of processing personal data, whereas "data processors" are entities that handle personal data on behalf of controllers pursuant to their instructions. In its Response to the Privacy Act Review Report (**Response**),[2] the Government agreed-in-principle to the proposal to implement a clear distinction between controllers and processors in the Privacy Act. Indeed, the Government recognised that a key focus area of the Privacy Act review was to "increase clarity and simplicity for entities and individuals", and that introducing the controller-processor distinction would "bring Australia into line with other jurisdictions, reflect the operational reality of modern business relationships, and reduce the compliance burden for entities acting as processors".[3]

BSA strongly agrees with the Government's observations. The introduction of a controller-processor distinction is the most important proposal to emerge from the Privacy Act review. There are significant benefits to implementing this distinction under the Act:

[1] BSA's members include: Adobe, Alteryx, Amazon Web Services, Asana, Atlassian, Autodesk, Bentley Systems, Box, Cisco, Cloudflare, Cohere, Dassault Systemes, Databricks, Docusign, Dropbox, Elastic, EY, Graphisoft, HubSpot, IBM, Informatica, Kyndryl, MathWorks, Microsoft, Notion, Okta, OpenAI, Oracle, PagerDuty, Palo Alto Networks, Rubrik, Salesforce, SAP, ServiceNow, Shopify Inc., Siemens Industry Software Inc., Trend Micro, TriNet, Workday, Zendesk, and Zoom Communications Inc.

BSA is the leading advocate for the global software industry. BSA members create technology solutions that power other businesses, including cloud storage services, customer relationship management software, human resources management programs, identity management services, network infrastructure services, cybersecurity solutions, and collaboration systems. Our members have made significant investments in Australia, and we are proud that many Australian entities and consumers continue to rely on our members' products and services to do business and support Australia's economy.

[2] Government Response to the Privacy Act Review Report, September 2023, https://www.ag.gov.au/rights-and-protections/publications/government-response-privacy-act-review-report

[3] Response (2023), p. 15

- Adopting a distinction between controllers and processors will align the Act with privacy laws globally, including, but not limited to the European Union's General Data Protection Regulation (**GDPR**), California's Consumer Privacy Act (**CCPA**), and Singapore's Personal Data Protection Act (**PDPA**). This alignment will help Australian entities understand how their obligations under the Privacy Act map to their obligations under data protection laws in other major markets. It will also help entities streamline data protection and transfer practices across markets.
- Clearly distinguishing between the roles of controllers and processors also improves consumer protection and enhances regulatory certainty for businesses. As noted in the Attorney General Department's Privacy Act Review Report 2022 (**AGD Report**),[4] distinguishing between controllers and processors will "clarify consent obligations and assist with clarifying obligations in relation to any new individual rights (such as a right to erasure) that may be introduced following this review", and "help entities more effectively respond to data breaches".[5]

In the absence of a controller-processor distinction, both consumers and businesses face increased uncertainty. Consumers lack clarity about which entities are responsible for safeguarding their data and how to exercise their rights, while businesses lack clear guidance on their obligations, leading to potential gaps in accountability and compliance. Australia's Privacy Act will also remain an outlier among major jurisdictions. This means that both Australian businesses expanding into new markets and businesses looking to invest in Australia will need to navigate inconsistent privacy frameworks across multiple regions, which increases compliance costs and complexity.

Question: Are there any changes you would like to see to privacy legislation in Australia? Please provide details below.

BSA supports the Government's efforts to reform and modernise the Privacy Act. This failure to implement a clear controller-processor distinction in the Privacy Act in the Privacy and Other Legislation Amendment Act 2024 represents a missed opportunity to bring Australia's privacy regime in line with international best practices and to provide greater certainty for both consumers and businesses.

We therefore urge the Government to prioritise this issue in the next tranche of amendments. In the context of implementing the controller-processor distinction, we recommend the following:

1. Adopt definitions of controllers and processors that align with other important privacy laws. Under the GDPR, for example, a data controller is defined as an entity that "alone, or jointly with others, determines the purposes and means of processing personal data", whereas a data processor is defined as an entity that "processes personal data on behalf of the controller". The language provides a clear distinction between entities that decide

[4] Privacy Act Review Report 2022, February 2023, https://www.ag.gov.au/sites/default/files/2023-02/privacy-act-review-report_0.pdf

[5] AGD Report (2023), p. 231

how and why to collect an individual's personal data and entities that act on behalf of others to process personal information.

2. Clearly recognise that determining whether an entity is acting as a controller or processor is fact-based and context specific. A single company may act as a controller for some of its products and services (i.e., consumer-facing products) and as a processor for others (i.e., business-to-business services). Such a company should be subject to obligations the Act places on controllers when it acts as a controller and subject to obligations the Act places on processors when it acts as a processor. We recommend expressly stating this in the Act.
3. Ensure that obligations for controllers and processors are appropriately tailored according to their roles. Controllers and processors should be subject to distinct obligations under the Privacy Act. Importantly, consumer-facing responsibilities, such as obtaining consent from individuals and responding to consumer rights requests, should be assigned to controllers rather than processors. This reflects the fact controllers are entities that decide how and why to collect a consumer's personal information. Processors should be subject to other obligations, such as requirements to safeguard personal information and to process personal information only on behalf of the controller and pursuant to its instructions. This allocation of responsibilities is also consistent with global norms.

Enabling AI's productivity potential

Question: How are you currently using AI? Please provide details of the context and uses.

BSA members are at the forefront of the responsible development and deployment of AI, providing trusted software solutions that enable enterprises and their customers to harness the power of AI to improve their product offerings and enhance their competitiveness. These enterprise-focused AI solutions are used by businesses in a wide range of industries, including manufacturing, logistics, healthcare, financial services, retail, and the public sector.

In many cases, companies use AI in low-risk ways that create significant benefits not only to the companies themselves but to the customers they serve:

- Improving logistics and planning. AI systems can improve a company's ability to forecast supply-chain issues, optimize delivery routes, estimate arrival times for new shipments, and reduce their fuel and energy usage.
- Identifying and managing common documents. Companies can use AI tools to read handwritten documents, identify a contract based on its format, and scan files for sensitive data that need stricter care. AI systems can then create summaries of regular corporate reports, or generate new forms based on existing examples from frequently used documents.
- Keeping shelves stocked. AI systems can forecast demand for products and redistribute them across a company's physical stores. AI systems can also detect early signs of supply chain issues and alert managers if inventories drop below certain levels.

- Transcribing and summarising meetings. Employees can spend more time collaborating if they use an AI system to transcribe their conversation, summarise decisions, and identify follow-up items, instead of requiring an employee to serve as a notetaker.
- Answering customer questions. Businesses can offer customers 24/7 support through AI-powered chatbots that answer straightforward questions even when human customer service representatives are asleep. Bots can be programmed to address basic questions instead of sending customers to FAQs.
- Responding to frequent emails. Companies can set up AI systems to respond to common requests, like sending automatic responses to emails asking about the status of payment invoices.

AI systems are used across every sector of the economy and companies are turning to AI in a range of industry-specific scenarios:

- Transportation. AI systems can improve the efficiency of airlines by helping to pinpoint causes of any slowdowns in the process of cleaning, refuelling, and reloading airplanes. Detecting these delays early helps the airline mitigate their effect on passengers.
- Manufacturing. AI design tools can optimize manufacturing processes to reduce waste and improve products. This is true from early phases where AI can help design and test new prototypes to factory floors where AI systems can identify maintenance and quality-control issues.
- Agriculture. Farmers use AI systems to analyse large volumes of weather and crop information, helping them monitor their crops, increase yields, and adjust to rain and drought conditions.
- Construction. Companies use AI to streamline the process of designing and constructing new buildings. They can also create "digital twins" of real-life cities to understand environmental and other impacts of a proposed design.

Companies frequently use AI systems to protect consumers using their products and services, often in ways that are designed to operate in the background. For example, businesses across a range of industries use AI systems to improve fraud detection and to protect against cybersecurity threats.

- Cybersecurity. AI helps organisations predict potential attacks, mitigate attacks in real-time, manage access to resources, and encrypt sensitive data. For example, a company can use an AI system to identify malicious files and suspicious IP addresses that can be easily missed by humans due to the sheer volume. In some cases, AI systems can be used to forecast, detect, prevent, and respond to threats automatically. This helps companies secure the products and services that consumers use and protect against potential threats.
- Fraud detection. Banks and credit card companies can use AI-powered systems to better detect potential fraud in real time, including by setting rules for identifying suspicious wire transfer and credit card transactions. AI systems enable companies to monitor large amounts of customer transactions to find anomalies, including detecting unusual usage patterns for a consumer's online accounts. Banks can use that information to alert

customers about potential security concerns, while reducing the amount of "false positives" that may block consumers from using their own credit cards.

Question: Do you think there are opportunities to make greater use of AI in your work or home environment? What do you see as the biggest upsides?

Yes. There are significant opportunities to make greater use of AI, particularly as the technology continues to evolve and becomes more accessible.

In the business environment, BSA members see strong potential for broader AI deployment to improve efficiency and drive productivity gains. AI enables businesses to automate time-intensive tasks such as document review, data classification, and administrative processing. AI also excels at identifying patterns and trends in large, complex datasets. Businesses are using AI to improve forecasting, optimise operations, and make real-time adjustments based on market conditions, leading to better decisions and faster response times. These efficiencies are especially impactful across sectors such as healthcare, logistics, finance, and education.

Question: What challenges do you face in accessing or using AI? How can these challenges be overcome?

Several policy challenges continue to limit businesses' ability to access, develop, and deploy AI tools responsibly and at scale. Addressing these challenges will be essential to fostering an environment that supports innovation while managing risk. In particular, we highlight three key challenges below: (1) overbroad policies that do not narrowly tailor the scope to high-risk use cases that pose the most significant risks; (2) policies that do not appropriately account for the different roles of various actors in the AI ecosystem and tailor responsibilities accordingly; and (3) policies that limit access to critical information needed to train AI systems, including copyrighted information.

1. Supporting a Risk-Based Approach to AI Regulation

In September 2024, the Department of Industry, Science and Resources (**DISR**) released a Proposals Paper for Introducing Mandatory Guardrails for AI in High-Risk Settings (**Proposals Paper**).[6]

BSA supports a risk-based approach that focuses on use-cases that create high risks to individuals. If DISR moves forward with developing mandatory guardrails, we strongly encourage the Government to identify a specific and defined set of use cases that are considered high-risk, which will provide legal certainty for organisations as to which new safeguards should be adopted and implemented. This will also maintain space for low-risk AI applications to continue delivering productivity and innovation benefits across the economy.

[6] Introducing Mandatory Guardrails for AI in High-Risk Settings: Proposals Paper, September 2024, proposals_paper_for_introducing_mandatory_guardrails_for_ai_in_high_risk_settings.pdf

For example, we have encouraged policymakers to focus on AI systems specifically developed to make consequential decisions which determine an individual's eligibility for, and result in the provision or denial of, housing, employment, credit, education, access to physical places of public accommodation, healthcare, or insurance.

2. Recognizing the Different Roles of Different AI Actors

As the Proposals Paper notes, "AI systems often involve a complex network of actors responsible for different aspects of the system's development and deployment" and, as such, "it will be important to clarify the roles and specific obligations" of actors across the AI supply chain.[7]

While we are encouraged that the policy intention is to reflect the distinct roles and responsibilities of different organisations in the AI supply chain, we have the following concerns about the definitions in the Proposals Paper:

- The Proposals Paper defines developers as "organisations or individuals who design, build, train, adapt or combine AI models and applications". This definition is overly broad. Entities that "train, adapt or combine" AI models and applications should not be considered AI developers. This broad definition could cover organisations that merely adapt or customise AI models for low-risk, specialized applications, including those that merely decide to use more than one model developed by another company in their own existing software application. Such organisations are not well positioned to carry out the obligations of developers. Instead, the definition of AI developers should clearly focus on those engaged in the actual development of AI systems. To implement the risk-based approach discussed above, we suggest focusing this definition on developers of high-risk AI systems:
    - **Developer of High-Risk AI System:** An entity that: (1) designs an AI system specifically intended to be used as a high-risk system; (2) substantially modifies a high-risk AI system, or (3) substantially modifies a non-high-risk AI system so that it becomes a high-risk AI system.
- The Proposals Paper defines deployers as "any individual or organisation that supplies or uses an AI system to provide a product or service". This definition is also overly broad. The use of the term "supplies" may inadvertently include companies that do not actually use an AI system in practice. The definition of deployer should be narrowed to focus on those that actually use a high-risk AI system in delivering a product or service. We strongly recommend limiting the definition of AI deployers to focus on these high-risk scenarios:
    - **Deployer of High-Risk AI systems:** An entity that uses a high-risk AI system.
- While it is important to recognize the different roles of developers and deployers, the AI ecosystem is complex and often involves multiple companies that may develop an AI

[7] Proposals Paper (2024), p. 31-32

model, integrate that model into a particular AI system, and use the AI system for a specific task. AI regulations should acknowledge and distinguish among these different roles, since each type of company has access to different types of information and can take different actions to mitigate risks associated with the AI system. AI legislation must account for these differences to create effective and workable obligations.

3. Clear Copyright Exception for AI Training

AI training entails a process of teaching AI system to make decisions or predictions by exposing it to large amounts of data. "Raw data" sources may include everything from machine-to-machine data (e.g., satellite transmission data) and international trade statistics to published materials, blog posts, website comments, and chat room logs. Because much of the content on the Internet is subject to copyright protection (which arises automatically upon creation of an original work), some raw data may include material covered by copyright.

"Raw data" must typically be subjected to a transformative data preparation process that disaggregates and converts data into a form that can be semantically understood by a computer. As part of this process, an AI development team semantically and structurally transforms the data through "tokenization," which involves breaking down a piece of text or data into smaller units (or "tokens") for purposes of computational analysis. At the end of this process, training data will often ultimately appear as a sequence of tokens that stretches across the entire AI corpus.

Current law should sufficiently protect AI training from any copyright liability for several reasons:

1. The data preparation process transforms raw data to create a new AI corpus for analytical purposes.
2. AI training involves computational analysis of mathematical correlations, trends, and patterns across tokenized variables.
3. These facts or mathematical concepts are not themselves copyrightable.
4. AI training is not a consumptive use of any copyrighted content for an expressive purpose.

That said, to provide greater legal certainty for copyright holders and AI developers alike, we urge the Australian Government to introduce a text and data mining (**TDM**) exception which can be applied for both commercial and non-commercial activities like those in Japan and Singapore and under consideration in Hong Kong. Introducing the consequent legal certainty will facilitate investments in AI development and deployment in Australia while clarifying appropriate protections for copyright holders.

As the Government of Singapore explained when considering a potential TDM exception (also known as the "computational data analysis" exception in Singapore):

> *A specific exception for such activities is preferred to relying on the general open-ended fair dealing defence, as it promotes certainty and allows calibration of specific*

safeguards and conditions to address the concerns raised by [various stakeholders]. These include conditions to preserve and protect rights-holders' commercial interests and freedom to conduct business based on licensing and subscription models. For example, the exception will be limited to acts of copying and a user must have lawful access to the works and other subject matter that are copied. If certain material can only be accessed through a paid subscription, the user must pay for the subscription before using the material for text and data mining. The user also cannot distribute the material to anyone without such lawful access. The exception will also not prevent rights-holders from taking reasonable measures to maintain the security and stability of their computer system or network.[8]

In sum, a well-crafted exception would: a) promote AI development and wider economic growth; b) facilitate research; c) allow Australia to maintain its competitiveness on the global stage and sustain its position as a hub for growth and investment; d) afford legal certainty to users; and e) strike an appropriate balance between copyright protection and reasonable use of copyright works. For more information, please see BSA's AI and Copyright Policy Backgrounder.[9]

Question: Do you have any concerns about using AI? What are the reasons for your answer? What can be done to lower your level of your concerns?

BSA recognises that while AI presents transformative opportunities for innovation, productivity, and economic growth, it also raises important concerns that must be thoughtfully addressed to build public trust and promote responsible use. These concerns include risks related to AI bias and content authenticity.

1. Protecting Against AI Bias

AI systems need to guard against the potential for systematically and unjustifiably yielding less favourable, unfair, or harmful outcomes to members of specific demographic groups, particularly for AI tools used in high-risk settings. BSA's views on the risks of AI bias, and how to best manage it, are informed by our experience working with member companies to develop the BSA Framework to Build Trust in AI (**BSA Framework**),[10] a risk management framework for mitigating the potential for unintended bias throughout an AI system's lifecycle.

For several years, BSA has supported adopting AI risk management programs and requiring companies to conduct impact assessments for high-risk uses of AI. These practices can help to guard against potential bias. For example, the BSA Framework recommends that in the data acquisition phase of the design of an AI model, companies should evaluate the representativeness of the data as part of conducting an impact assessment. To do so, a

[8] Singapore Copyright Review Report, January 2019, https://www.mlaw.gov.sg/files/news/public-consultations/2021/copyrightbill/Annex_A-Copyright_Report2019.pdf

[9] BSA Artificial Intelligence and Copyright Policy, May 2025, https://www.bsa.org/files/policy-filings/05302025bsaaicopyright.pdf

[10] Confronting Bias: BSA's Framework to Build Trust in AI, June 2021, https://ai.bsa.org/wp-content/uploads/2021/06/2021bsaaibias.pdf

company can compare the demographic distribution of training data to the population in which the system will be deployed and assess whether there is sufficient representation of subpopulations that are likely to interact with the system. To mitigate issues that arise, companies can consider rebalancing the dataset with additional data or synthetic data, which involves oversampling data from underrepresented groups.

Similarly, in the data preparation and model definition phase, an impact assessment could include documenting a potential correlation between selected features and sensitive demographic attributes. For features that closely correlate to a sensitive class, companies can document the relevance to the target variable and the rationale for its inclusion in the model, consistent with laws prohibiting discriminatory actions. These practices, along with other measures, can help to decrease the potential for AI bias while enabling innovation.

2. Content Authenticity

The advent of AI-enabled tools has prompted questions about how malicious actors might exploit AI to exacerbate misinformation, create and spread disinformation, and ultimately undermine trust in institutions. Trusted tools are needed to discern what is fake from what is real.

In this regard, BSA supports the development and deployment of reliable content authentication and provenance mechanisms that can help users identify the history and origin of AI-generated content. We support efforts by the Content Authenticity Initiative (**CAI**) to promote the open Coalition for Content Provenance and Authenticity (**C2PA**) standard.[11] The C2PA standard, which is expected to be approved by the International Standards Organization this year, is openly available. Anyone, including governments, can use this standard to incorporate digital provenance information into their products and processes. Creators can indicate whether AI was used in their work and, if so, how it was used. Content credentials will display information about how the work was created, the date it was made, and any edits that were made along the way. This standard will help consumers decide what content is trustworthy and promote transparency around the use of AI.

Embracing open standards like that developed by C2PA facilitates international interoperability and enhances the integrity of digital content ecosystems. We also note that what constitutes state of the art in ensuring solutions for content provenance will evolve over time so any governance framework must be designed to accommodate such developments and should provide scope for organisations to assess what is the most relevant solution for them when it comes to content authentication and provenance mechanisms.

[11] Content Authenticity Initiative, https://contentauthenticity.org, C2PA Specifications at https://c2pa.org/specifications/specifications/2.1/index.htm

15 September 2025

BSA COMMENTS ON PRODUCTIVITY COMMISSION'S HARNESSING DATA AND DIGITAL TECHNOLOGY INTERIM REPORT

Submitted Electronically to the Productivity Commission

The Business Software Alliance (**BSA**)[1] welcomes the opportunity to submit comments to the Productivity Commission on its Interim Report on Harnessing Data and Digital Technology (**Interim Report**).[2]

BSA is the leading advocate for the global software industry. BSA members create technology solutions that power other businesses, including cloud storage services, customer relationship management software, human resources management programs, identity management services, network infrastructure services, cybersecurity solutions, and collaboration systems. Our members have made significant investments in Australia, and we are proud that many Australian companies and organisations continue to rely on our members' products and services to do business and support Australia's economy.

BSA commends the Productivity Commission for recognising the importance of harnessing the economic and productivity benefits of data and digital technology. The Interim Report rightly highlights that emerging technologies such as artificial intelligence (**AI**) can play a transformative role in driving growth across all sectors of the Australian economy. However, realising this potential requires a regulatory landscape that encourages innovation while providing clear guidance for businesses. We are encouraged that the Interim Report's findings and recommendations underscore the need for clear and outcomes-focused reform options which minimise unnecessary regulations, thereby enabling Australian businesses to adopt and leverage technology with confidence.

Summary of BSA's Recommendations

Our recommendations correspond to the specific policy areas highlighted in the Interim Report, as follows:

<u>Artificial Intelligence</u>

1. Focus on identifying high-risk use cases of AI
2. Recognise and define the different roles and responsibilities of entities in the AI supply chain

[1] BSA's members include: Adobe, Alteryx, Amazon Web Services, Asana, Atlassian, Autodesk, Avalara, Bentley Systems, Box, Cisco, Cloudflare, Cohere, Dassault Systemes, Databricks, Docusign, Dropbox, Elastic, EY, Graphisoft, HubSpot, IBM, Informatica, Kyndryl, MathWorks, Microsoft, Notion, Okta, OpenAI, Oracle, PagerDuty, Palo Alto Networks, Rubrik, Salesforce, SAP, ServiceNow, Shopify Inc., Siemens Industry Software Inc., Trend Micro, TriNet, Workday, Zendesk, and Zoom Communications Inc.

[2] Harnessing Data and Digital Technology – Interim Report, August 2025, https://www.pc.gov.au/inquiries/current/data-digital/interim/data-digital-interim.pdf

3. Officially designate a non-regulatory central coordinating body for AI policy
4. Introduce a clear copyright exception for AI training

Data Access

5. Allow private and non-Australian entities to participate in the Data Availability and Transparency Act 2022 (**DAT Act**) data sharing scheme
6. Explore privacy-enhancing technologies

Privacy

7. Implement a controller-processor distinction

Artificial Intelligence

The broad adoption of enterprise AI by Australian businesses and organisations must be a strategic priority for the Government. The economic opportunities presented by AI extend beyond the development of more powerful tools – they also lie in enabling these tools to make every sector of the economy more efficient and productive. Countries that promote AI adoption across both the private and public sectors will be best placed to capture the greatest economic benefits and deliver substantial gains across industries. overly prescriptive or premature regulation risks discouraging businesses from adopting AI tools.

In this regard, BSA agrees with the Interim Report's findings and recommendations that:

- The Government plays a key role in promoting investment in digital technology, including AI, by providing a stable regulatory environment.[3]
- Any regulatory response to potential harms from using AI must be proportionate, risk-based, outcomes-based and technology neutral where possible.[4]
- The Government should undertake a comprehensive set of "gap analyses" to understand the risks stemming from AI and whether these risks can be dealt with under existing regulatory frameworks.[5]
- Until these gap analyses are completed, the Government should pause consideration on economy-wide AI-specific regulations, including potential mandatory guardrails for AI in high-risk settings, to avoid regulatory duplication and imposing unnecessary compliance burdens.[6]
- Economy-wide AI-specific regulations should only be considered as a last resort for the use cases of AI that meet two criteria, namely: a) where existing regulatory frameworks cannot be sufficiently adapted to handle the issue; and b) where technology-neutral regulations are not feasible.[7]

[3] Interim Report (2025), p. 19.

[4] Interim Report (2025), p. 19.

[5] Interim Report (2025), p. 17. BSA further notes that in 2022, the Australian Government conducted a similar consultation on whether Australia's laws and regulations are fit for purpose and able to respond to emerging technologies, such as AI. While this consultation was conducted before the generative AI "boom", the Government can leverage the groundwork and insights gleaned from this exercise. **See:** Consultation launched to position Australia as a leader in Digital Economy Regulation, March 2022, https://www.pmc.gov.au/news/consultation-launched-position-australia-leader-digital-economy-regulation.

[6] Interim Report (2025), p. 18.

[7] Interim Report (2025), p. 20.

- Mandatory guardrails for AI in high-risk settings are only appropriate in circumstances where the gap analyses determined that existing regulatory frameworks or new technology-neutral regulations cannot adequately mitigate the risk of harm.[8]

These findings and recommendations present a solid set of principles on which the Government can build its approach to AI policymaking. BSA proffers the following recommendations to operationalise these principles and ensure that Australia's AI framework both promotes responsible innovation and maximises the productivity and competitiveness benefits of AI adoption.

Focus on identifying high-risk use cases of AI

We support the risk-based approach recommended in the Interim Report. The risks of AI are inherently use-case specific. Many AI systems pose extremely low, or even no, risk to individuals or society, while creating potentially significant benefits like helping to organise digital files, auto-populate common forms for subsequent human review, or improve a company's ability to forecast supply chain issues.

Generally, AI governance efforts should focus on addressing high-risk AI use cases based on potential harm. This approach assesses risk based on the AI system's application and intended use case, rather than the sector in which it is used or deployed. In this regard, BSA has encouraged policymakers to focus on AI systems specifically developed to make consequential decisions, which are determinations that have a material legal or similarly significant effect on a consumer's eligibility for and result in the provision or denial of housing, employment, credit, education, access to physical places of public accommodation, healthcare, or insurance. This prevents low-risk AI systems from being automatically classified as high-risk simply because they are used in a particular sector (e.g., healthcare).

Recommendation: In line with a risk-based approach, the Productivity Commission should make clear that the focus of the gap analyses is to identify high-risk AI use cases. This will ensure that any subsequent regulatory efforts are targeted at a set of specific use cases.

Recognise and define the different roles and responsibilities of entities in the AI supply chain

It is critical for the Government to clearly recognise and define the different roles and responsibilities of various actors in the AI supply chain. This is because AI systems are often developed, integrated and deployed by different organisations, each of which has access to different information and the ability to address different risks. Without clear role differentiation, regulatory obligations risk being either duplicative or misaligned, placing unrealistic compliance burdens on some actors while leaving actual risks unaddressed. In short, all actors in the AI supply chain should act responsibly, but legal requirements and best practices should be tailored to each actor's role.

- Several actors may comprise the AI supply chain. Important roles include:**Developer of High-Risk AI System:** An entity that: (1) designs an AI system specifically intended to be used as a high-risk system; (2) substantially modifies a high-risk AI system, or (3)

[8] Interim Report (2025), p. 21

substantially modifies a non-high-risk AI system so that it becomes a high-risk AI system. Consequently, an AI developer is well-placed to provide details on the intended purpose of the AI system, the capabilities of the AI system, known limitations, and risks of the AI system at the time of development, the data used to train the AI system, and how the AI system was evaluated prior to sale.

- **Deployer of High-Risk AI systems:** An entity that uses a high-risk AI system. Consequently, an AI deployer is well-placed to provide details on the purpose for which the AI system is deployed, transparency measures relating to affected individuals or end users (e.g., disclosures of how the deployer plans to use a customer's personal information), post-deployment monitoring, user safeguards, and measures taken to mitigate potential risks arising from the AI system's deployment. AI deployers may also have more insight into the data inputs during the system's use and the resulting outputs and other real-world factors affecting the system's performance.
- **Integrator:** Outside of developers and deployers of high-risk AI systems, the widespread use of General Purpose AI (**GPAI**) models to power applications and products has highlighted another important actor in the AI supply chain – integrators. Integrators are entities that: (1) connect their customers to third-party general purpose AI models using their own software or platforms; or (2) otherwise integrate GPAI models into products or services for use by third parties without substantially modifying the GPAI models. Consequently, integrators are well-placed to share relevant, non-sensitive information about the changes they make to the model to the next company in the supply chain so that it can understand the AI's capabilities and limitations once it is used in the real world.

Recommendation: The Productivity Commission should recognise these distinct roles and urge agencies to consider the differences among them when assessing appropriate responsibilities.

Officially designate a non-regulatory central coordinating body for AI policy

To ensure regulatory coherence and avoid duplication, the Productivity Commission should recommend that the Government designate a lead non-regulatory body with responsibility to coordinate AI governance and policymaking across sectors. Given the cross-cutting nature of AI, responsibilities can be fragmented across different agencies and regulators. The Interim Report acknowledged that "[t]he range of commonly-cited AI risks spans several Australian Government portfolios including health, consumer protection, criminal law, national defence, financial regulation, media regulation and intellectual property".[9]

While sectoral regulators will remain best placed to address risks in their respective domains, a central coordinating body is needed to ensure consistency in policy direction and avoid overlapping or conflicting requirements. Importantly, the coordinating body's function is not to draft regulations on behalf of agencies or sectoral regulators – rather, it should be setting the overarching policy direction, working closely with various Government stakeholders to coordinate regulatory approaches, and ensuring that AI governance across portfolios remains risk-based and proportionate.

[9] Interim Report (2025), p. 17.

Recommendation: The Productivity Commission should urge the Government to officially designate a lead agency or non-regulatory central coordinating body to oversee AI policy coordination, with a focus on preventing internal regulatory fragmentation and duplication.

Introduce a clear copyright exception for AI training

The Interim Report explores whether "current Australian copyright law is a barrier to building and training AI models"[10] and considers if "there is a case for a new fair dealing exception that explicitly covers text and data mining".[11] BSA supports the development of statutory exceptions, including a text and data mining exception (**TDM exception**), that permit AI training on lawfully accessed content.

AI training entails a process of teaching an AI system to make decisions or predictions by exposing it to large amounts of data. "Raw data" sources may include everything from machine-to-machine data (e.g., satellite transmission data) and international trade statistics to published materials, blog posts, website comments, and chat room logs. Because much of the content on the Internet is subject to copyright protection (which arises automatically upon creation of an original work), some raw data may include material covered by copyright.

"Raw data" must typically be subjected to a transformative data preparation process that disaggregates and converts data into a form that can be semantically understood by a computer. As part of this process, an AI development team semantically and structurally transforms the data through "tokenization," which involves breaking down a piece of text or data into smaller units (or "tokens") for purposes of computational analysis. At the end of this process, training data will often ultimately appear as a sequence of tokens that stretches across the entire AI corpus.

This training process does not infringe on copyright because it consists of computational analysis of mathematical correlations, patterns, and trends across tokenised variables. These mathematical relationships are not copyrightable, and the training process is not a consumptive use of expressive content. Rather, it is a functional process aimed at enabling machines to detect and model statistical relationships. In this sense, AI training is distinct from expressive uses of copyrighted works.

To provide greater legal certainty for copyright holders and AI developers alike, we urge the Productivity Commission to recommend introducing a clear copyright exception which can be applied for both commercial and non-commercial activities, like those in Japan and Singapore and under consideration in Hong Kong. Introducing the consequent legal certainty will facilitate investments in AI development and deployment in Australia while clarifying appropriate protections for copyright holders.[12] As the Government of Singapore explained when it introduced its "computational data analysis" exception (i.e., Singapore's equivalent of a TDM exception):

> *A specific exception for such activities is preferred to relying on the general open-ended fair dealing defence, as it promotes certainty and allows calibration of specific safeguards*

[10] Interim Report (2025), p. 24

[11] Interim Report (2025), p. 27.

[12] BSA recognises that some content owners want more say in how their content is used in or with AI. One way to reflect and respect those wishes is through "opt-out" tools. The development and use of such opt-out tools, and the standards governing them, are best handled through voluntary conversations between creators and AI developers and deployers, who collectively are better positioned to develop effective, consensus-driven technical mechanisms than the Government.

> *and conditions to address the concerns raised by [various stakeholders]. These include conditions to preserve and protect rights-holders' commercial interests and freedom to conduct business based on licensing and subscription models. For example, the exception will be limited to acts of copying and a user must have lawful access to the works and other subject matter that are copied. If certain material can only be accessed through a paid subscription, the user must pay for the subscription before using the material for text and data mining. The user also cannot distribute the material to anyone without such lawful access. The exception will also not prevent rights-holders from taking reasonable measures to maintain the security and stability of their computer system or network.*[13]

In sum, a well-crafted exception would: a) promote AI development and wider economic growth; b) facilitate research; c) allow Australia to maintain its competitiveness on the global stage and sustain its position as a hub for growth and investment; d) afford legal certainty to users; and e) strike an appropriate balance between copyright protection and reasonable use of copyright works.

Recommendation: The Productivity Commission should recommend that the Government introduce a clear statutory exception that applies to both commercial and non-commercial AI training conducted on lawfully accessed materials. Such an exception would provide businesses with the certainty to invest in AI development, strengthen Australia's competitiveness, and ensure that copyright holders' interests are preserved through lawful access safeguards.

Data Access and Sharing

BSA agrees with the Interim Report's observation that "the benefits of data sharing remain somewhat untapped in Australia".[14] The Interim Report also noted that, as part of efforts to give individuals and consumers better access to data that relates to them, governments around the world are developing new regulatory pathways to facilitate data portability and interoperability.[15] Indeed, government-held data is an important asset that can serve as a powerful engine for creating new jobs, promoting economic growth, driving productivity gains, and enabling innovation. To create new data sharing and integration opportunities, BSA recommends the following:

Allow private and non-Australian entities to participate in the DAT Act's data sharing scheme

The DAT Act establishes a scheme for sharing public sector data with "accredited users" for specific purposes (**Scheme**).[16] However, under the Scheme, only Australian entities are allowed to apply for accreditation. Furthermore, private entities, including "individuals, bodies corporate,

[13] Singapore Copyright Review Report, January 2019, https://www.mlaw.gov.sg/files/news/public-consultations/2021/copyrightbill/Annex_A-Copyright_Report2019.pdf

[14] Interim Report (2025), p. 29

[15] Interim Report (2025), p. 37.

[16] DAT Act, Section 12. Under the data sharing scheme, Commonwealth bodies are authorised to share their public sector data with accredited users, and accredited users are authorised to collect and use the data, in a controlled way. Data may be shared with an accredited user 9 directly, or through an intermediary accredited for the purpose (called an ADSP, short for accredited data service provider).

partnerships, trusts and unincorporated entities" will not be able to apply for accreditation.[17] While BSA recognises that the exclusion of these entities was intended to allow the Scheme to mature, these restrictions limit the efficacy of the Scheme, and consequently, the potential of public sector data to drive economic growth and productivity.

BSA notes that the Government has commenced a statutory review of the DAT Act, which considers the effectiveness of the DAT Act and whether its operations support improvements in public sector data availability, sharing and transparency.[18] This is a good opportunity for the Commission to recommend expanding the Scheme to allow both private and non-Australian entities to apply for accreditation. The accreditation framework contemplated in the DAT Act already allows the Australian Government to make a risk-based decision based on the company's ability to best meet security and data-handling requirements, among other factors.[19] Whether an entity is a private or foreign one should not matter so long as it can meet the requirements in the accreditation framework, and the Australian Government retains the discretion to reject applications from entities which do not meet said requirements.

Recommendation: The Productivity Commission should recommend expanding the DAT Act accreditation framework to allow applications from both private and non-Australian entities, provided they can demonstrate compliance with Australia's security, privacy, and data-handling requirements. This would increase the utility of the Scheme, maximise the value of public sector data.

Explore privacy-enhancing technologies

The benefits of data sharing must also be balanced against safety and privacy concerns. In this regard, BSA encourages the Productivity Commission to recommend that the Government explore the application of privacy-enhancing technologies and promote opportunities to further build value from the safe and responsible use of data. A range of emerging technologies, including homomorphic encryption, differential privacy techniques, and federated machine learning create opportunities for further sharing data while preserving individual privacy. These technologies can be used to maximise both the value and the confidentiality of sensitive information.

Recommendation: The Productivity Commission should recommend that the Government explore and support the adoption of privacy-enhancing technologies to supplement efforts to improve data access.

---

[17] Data Availability and Transparency Bill 2022 Supplementary Explanatory Memorandum, May 2022, para 6, https://parlinfo.aph.gov.au/parlInfo/search/display/display.w3p;query=Id%3A%22legislation%2Fems%2Fr6649_ems_a59313dc-0b7c-4244-8b78-c7c5e7380407%22.

[18] Statutory Review of the Data Availability and Transparency Act, August 2025, https://www.finance.gov.au/government/public-data/public-data-policy/statutory-review-data-availability-and-transparency-act-2022.

[19] DAT Act, Section 77(1). To become accredited, an entity is required to have: appropriate data management and governance policies and practices; an appropriately qualified individual in a position that has responsibility for data management and data governance; ability to minimise the risk of unauthorised access, sharing or loss of data; the necessary skills and capability to ensure the privacy, protection and appropriate use of data; and any additional criteria prescribed by the Minister.

Privacy

BSA has participated in multiple consultations on the review of the Privacy Act.[20] As currently formulated, the Privacy Act falls short of striking the right balance between protecting an individual's privacy and supporting beneficial, responsible data use. Importantly, we note the Interim Report's findings that "businesses may have not fully comprehended what their regulatory obligations were", which led to unnecessary "compliance" that impose additional costs on businesses without improving privacy outcomes.[21]

Such issues stem from the lack of an important distinction between data controllers and data processors, which is a foundational element of modern privacy laws globally.

Implement a controller-processor distinction

In brief, "data controllers" refer to entities which determine the purposes and means of processing personal data, whereas "data processors" are entities that handle personal data on behalf of controllers pursuant to their instructions. In its Response to the Privacy Act Review Report (**Response**),[22] the Government agreed-in-principle to the proposal to implement a clear distinction between controllers and processors in the Privacy Act. Indeed, the Government recognised that a key focus area of the Privacy Act review was to "increase clarity and simplicity for entities and individuals", and that introducing the controller-processor distinction would "bring Australia into line with other jurisdictions, reflect the operational reality of modern business relationships, and reduce the compliance burden for entities acting as processors".[23]

BSA strongly agrees with the Government's observations in the Response. The introduction of a controller-processor distinction is the most important proposal to emerge from the Privacy Act review. There are significant benefits to implementing this distinction under the Privacy Act:

- Adopting a distinction between controllers and processors will align the Privacy Act with privacy laws globally, including, but not limited to the European Union's General Data Protection Regulation (**GDPR**), California's Consumer Privacy Act (**CCPA**), and Singapore's Personal Data Protection Act (**PDPA**). This alignment will help Australian entities understand how their obligations under the Privacy Act map to their obligations under data protection laws in other major markets. It will also help entities streamline data protection and transfer practices across markets.
- Clearly distinguishing between the roles of controllers and processors also improves consumer protection and enhances regulatory certainty for businesses. As noted in the Attorney General Department's Privacy Act Review Report 2022 (**AGD Report**),[24]

[20] See: BSA Comments on Privacy Amendment Bill 2024, October 2024, https://www.bsa.org/policy-filings/australia-bsa-comments-on-privacy-amendment-bill-2024; BSA Comments on Australia Privacy Act Review Report 2022, April 2023, https://www.bsa.org/policy-filings/australia-bsa-comments-on-australia-privacy-act-review-report-2022; BSA Comments on Privacy Legislation Amendment Bill, November 2022, https://www.bsa.org/policy-filings/australia-bsa-comments-on-privacy-legislation-amendment-bill; BSA Comments on Review of Australia Privacy Act 1988, January 2022, https://www.bsa.org/policy-filings/australia-bsa-comments-on-review-of-australia-privacy-act-1988

[21] Interim Report (2025), p. 55.

[22] Government Response to the Privacy Act Review Report, September 2023, https://www.ag.gov.au/rights-and-protections/publications/government-response-privacy-act-review-report

[23] Response (2023), p. 15

[24] Privacy Act Review Report 2022, February 2023, https://www.ag.gov.au/sites/default/files/2023-02/privacy-act-review-report_0.pdf

distinguishing between controllers and processors will "clarify consent obligations and assist with clarifying obligations in relation to any new individual rights (such as a right to erasure) that may be introduced following this review", and "help entities more effectively respond to data breaches".[25]

In the absence of a controller-processor distinction, both consumers and businesses face increased uncertainty. Consumers lack clarity about which entities are responsible for safeguarding their data and how to exercise their rights, while businesses lack clear guidance on their obligations, leading to potential gaps in accountability and compliance. Australia's Privacy Act will also remain an outlier among major jurisdictions. This means that both Australian businesses expanding into new markets and businesses looking to invest in Australia will need to navigate inconsistent privacy frameworks across multiple regions, which increases compliance costs and complexity.

In the context of implementing the controller-processor distinction, we further recommend the following:

1. Adopt definitions of controllers and processors that align with other important privacy laws. Under the GDPR, for example, a data controller is defined as an entity that "alone, or jointly with others, determines the purposes and means of processing personal data", whereas a data processor is defined as an entity that "processes personal data on behalf of the controller". The language provides a clear distinction between entities that decide how and why to collect an individual's personal data and entities that act on behalf of others to process personal information.

2. Clearly recognise that determining whether an entity is acting as a controller or processor is fact-based and context specific. A single company may act as a controller for some of its products and services (i.e., consumer-facing products) and as a processor for others (i.e., business-to-business services). Such a company should be subject to obligations the Privacy Act places on controllers when it acts as a controller and subject to obligations the Privacy Act places on processors when it acts as a processor.

3. Ensure that obligations for controllers and processors are appropriately tailored according to their roles. Controllers and processors should be subject to distinct obligations under the Privacy Act. Importantly, consumer-facing responsibilities, such as obtaining consent from individuals and responding to consumer rights requests, should be assigned to controllers rather than processors. This reflects the fact controllers are entities that decide how and why to collect a consumer's personal information. Processors should be subject to other obligations, such as requirements to safeguard personal information and to process personal information only on behalf of the controller and pursuant to its instructions. This allocation of responsibilities is also consistent with global norms.

Recommendation: The Productivity Commission should urge the Government to prioritise implementing the controller-processor distinction in the next tranche of Privacy Act amendments. This will involve: (1) adopting definitions of controllers and processors that align with other

[25] AGD Report (2023), p. 231

leading privacy laws; (2) clearly recognise that determining whether an entity is acting as a controller or processor is fact-based and context specific; and (3) ensure that obligations and controllers are appropriately tailored according to their roles. This will modernise Australia's privacy regime and provide more legal certainty for businesses and consumers. In this context,

Conclusion

We hope that our comments will assist the Productivity Commission as it moves forward with the Interim Report. Please do not hesitate to contact me if you have any questions regarding this submission or if I can be of further assistance.

Yours sincerely,

Tham Shen Hong
Senior Manager, Policy – APAC