Allens

In alliance with Linklaters

Guide to AI procurement

JUNE 2025

As generative artificial intelligence moves from pilot to production, organisations are grappling with the push to rapidly procure AI systems while appropriately managing the risks.

'Market' positions and practices are still emerging, as both customers and suppliers continue to test their approach and their risk appetite, and as AI regulation evolves at pace. AI offerings – in terms of functionality, vendor commitments and pricing – are, as a result, still wildly inconsistent.

In this guide, we share our insights from recent AI procurements and answer frequently asked questions, including:

- Is AI procurement different from any other tech procurement? *(Yes!)*
- What questions should we be asking our business (and should they be asking AI vendors)? *(page 4)*
- How should we manage AI risk throughout the lifecycle of our arrangement? *(page 3)*
- What should we include in our contract? Can we just use our existing tech contract and bolt on an AI clause? *(page 7)*

Note: For the purposes of this guide, references to AI include automated decision-making systems, even though these systems may not always involve AI.

Eight tips for procuring AI systems

1 Tailor your supplier engagement policy for AI procurements

Your supplier engagement policy should (among other things) contemplate AI-specific risks and processes throughout the procurement lifecycle.
See page 3 for AI lifecycle considerations.

2 Prioritise governance and relationships

AI development is necessarily iterative, and performance needs to be monitored, tested and recalibrated on an ongoing basis. This means that the relationship between customers and their AI vendor (including to manage cost containment) throughout the arrangement is critical. Customers should select suppliers they trust to work constructively and collaboratively. The contractual structure (including governance arrangements) underpinning the arrangement should support this.

3 Get the right people in the room from the start

AI procurement should involve a cross-functional and diverse team, comprising representatives from Procurement, Legal, Technology, Compliance, Risk, Product Development, Data, Privacy, Cyber and (where relevant) Customer Service. This team should work together to identify the risks involved in the procurement and the controls needed to ensure those risks remain within the organisation's risk tolerance, from early in the procurement process.

4 Define outcomes, rather than specifications

Request for proposal (RFP) scopes should articulate the problem to be solved, rather than prescribing particular specifications. Focusing on outcomes, rather than the tool, can encourage innovation and allow vendors to provide the best solution/s for the problem. Once a solution is selected, statements of work and related specifications should be detailed to protect against arguments from vendors that the scope is expanding and additional fees are payable.

5 Understand upstream software and data dependencies

Diligence should include enquiries as to dependencies on fourth-party software (including open source and off-the-shelf solutions) and datasets.

6 Traditional tech contract terms are inadequate to address AI risks

AI contracts should address AI-specific risks, the changing nature of AI solutions, the evolving regulatory environment, and the positions AI vendors are adopting as the AI landscape and their own risks change. Bolting on an 'AI clause' to existing tech arrangements will rarely be sufficient to address these issues, particularly where an AI solution is being developed or customised.

7 Expect the unexpected

Have contingency plans in place, and mechanisms for unexpected events or outcomes during implementation and the operation of the AI system. Focus on robust testing and hypercare periods, as well as operational redundancies, to ensure your AI solution cannot present a single point of failure. Build in notification triggers for AI incidents or identified hazards.

8 Stay flexible

Since AI use cases, outcomes and features change quickly, arrangements should be structured to maintain flexibility, adapt to different needs and legal requirements, and ensure that performance doesn't degrade as more data is processed over time.

The AI procurement lifecycle

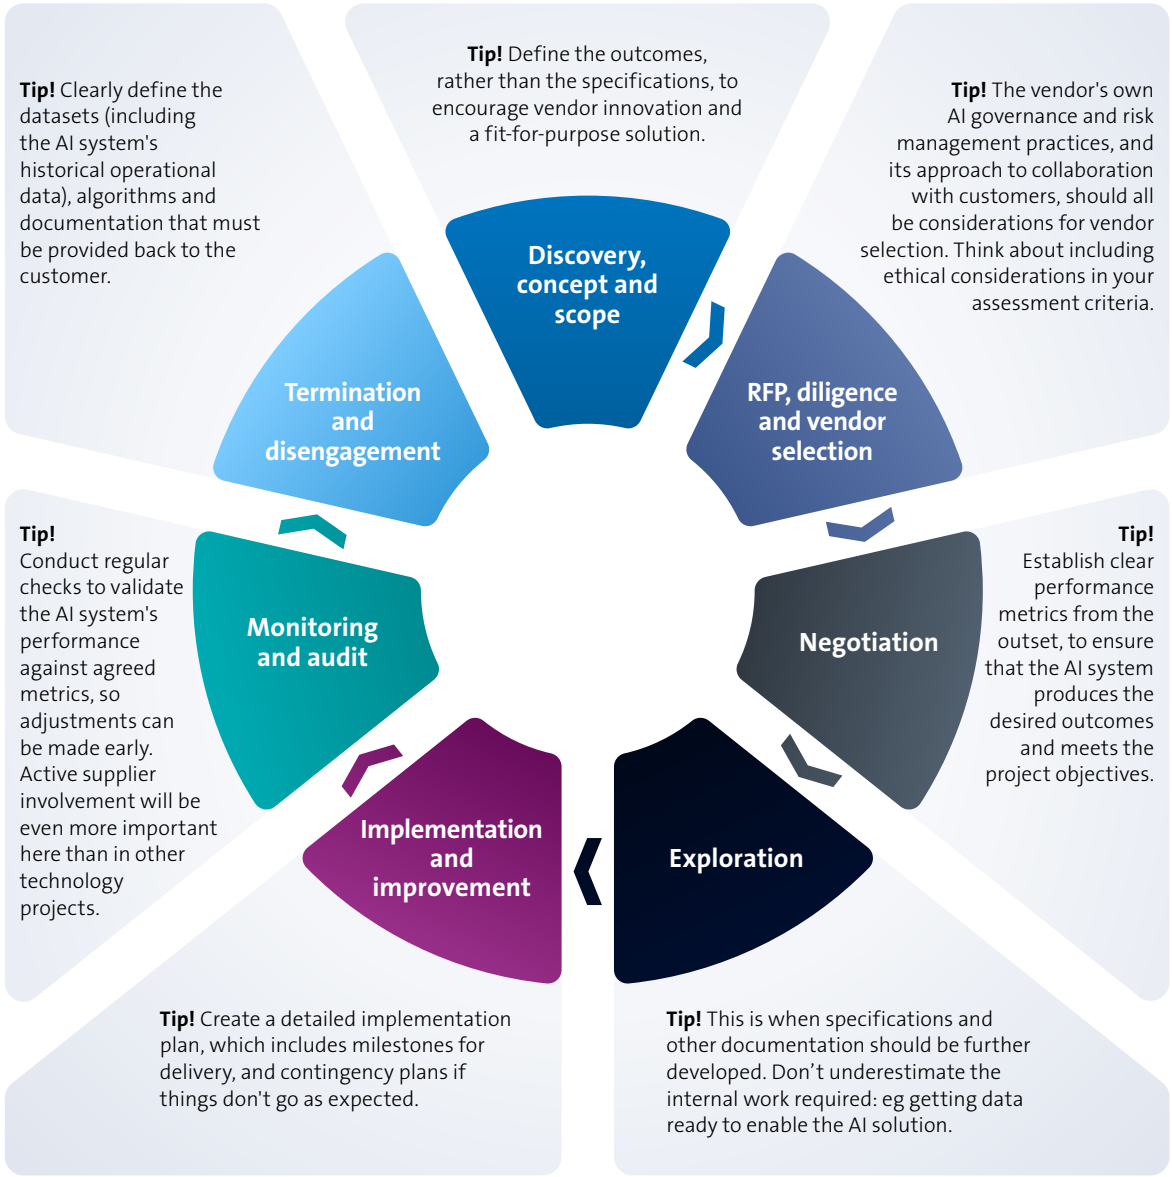

Make these three changes to your procurement policy

Your AI procurement policy can be standalone or incorporated into your existing procurement policy – either way, it should address the specific risks inherent in AI procurements.

1 Tailor your approach to supplier due diligence

Supplier due diligence processes and the artefacts (eg questionnaires) required to support supplier due diligence should facilitate your understanding of testing and training data provenance and suitability; intellectual property (***IP***) ownership and use considerations; and the supplier's own approach to governance, ethical issues, and the safe and responsible development and use of AI.

2 Require integrated AI impact assessments

In traditional technology procurements, it is common for subject matter experts to undertake their own risk or impact assessments (eg to address legal, privacy, indigenous data sovereignty, cyber, modern slavery or sustainability risks), sequentially and in isolation. This siloed approach is particularly problematic when it comes to the assessment of AI risks. Consistent with the Government's Voluntary AI Safety Standard and ISO/IEC DIS 42005 (AI system impact assessment), we recommend that organisations take a more integrated approach to AI impact assessments, and that these are undertaken by a multi-disciplinary team of experts working together.

3 Define triggers for procurement health checks

Your policy should identify both the project-specific and non-project specific triggers for the multi-disciplinary team to review the project and revisit the AI impact assessment. Triggers could include:

- changes in the organisational risk appetite, strategy or regulatory landscape;
- changes in the use or complexity of the AI system, or the sensitivity of the data;
- performance or compliance issues; and
- certain testing outcomes (eg where the tool needs to be retrained or recalibrated due to data drift).

Questions to ask the business about high-risk AI procurements

If you ask your business nothing else, ask...

1. What problem are we trying to solve and why does AI offer the best solution?
2. What are the key risks and potential impacts associated with this procurement?
3. What safeguards (contractual, governance, technical and operational) do we and/or the vendor have in place to address these risks?
4. Are we comfortable with the vendor's approach to AI, IP and data governance, and risk management?

Business case

- [ ] Is there an existing arrangement in place with the relevant vendor, or is this a new procurement?
  Tip! Although amending existing agreements to cover AI systems may be workable in some low-risk cases, higher-risk AI procurements are a different proposition and will typically warrant a customised contract.
- [ ] What type of AI is being utilised (eg automation, machine learning, automated decision-making, generative AI?) What are the various components of the AI system?
- [ ] What were our criteria for selecting this AI system and what are our alternatives?
- [ ] Will our AI system leverage an existing pre-trained model, or will we need to finetune?
- [ ] How much will this procurement cost?
  Tip! Consider whole-of-life costs, and don't forget internal costs including of: (i) obtaining and 'cleaning' the right data for training and testing on an ongoing basis; (ii) procuring the expertise required to develop, train, test and maintain the AI system; and (iii) integrating vendor technology with other systems.
- [ ] Does the vendor adhere to any recognised AI governance standards?[1]
  Tip! This can provide additional assurance but does not obviate the need for customers to undertake their own impact assessments.

1 For example, ISO/IEC 42001 – Information technology – Artificial intelligence – Management system.

Questions to ask the business about high-risk AI procurements

Data management

- [ ] What data will be / has been used to train, validate and test the AI system? Who is procuring or providing that data?
  Tip! Poor-quality, inappropriate or unrepresentative data may give rise to poor-quality decisions by the AI technology, and unpredictability, bias and discrimination.

- [ ] Has that data been legally obtained? Has it been cleared for copyright and other IP infringements? Do these datasets use 'protected characteristics' or characteristics that act as proxies for protected characteristics?
  Tip! Consider whether data has been scraped from public sources, as this can create regulatory and copyright risk. Consider also whether the collection, use or generation of training data and outputs could contravene privacy, IP or competition law, indigenous data sovereignty, or discrimination and contractual requirements and prohibitions.

- [ ] What are the limits or deficiencies of the datasets? How will we / the vendor address these?
  Tip! Consider data quality, bias, currency and gaps in the data used to build and train the system, as well as the data to be input into the system after training (including for ongoing testing). Consider also whether the training environment closely matches the environment in which the AI system will be used.

- [ ] Who is responsible for cleaning / converting the data into a usable form?

- [ ] Are we comfortable with our data being used to improve the supplier's AI or is our own quarantined instance required? Are we making available confidential information to train the AI system?
  Tip! Where data relates to a third party, consider whether it is confidential, and whether any contractual and equitable duties of confidence may prevent the information from being used for developmental purposes or shared with third-party collaborators.

- [ ] Who is responsible for monitoring and testing for data drift on an ongoing basis? What controls are in place to do this effectively?
  Tip! Data drift refers to the change in input data patterns over time. This is important because AI models are usually trained on a specific dataset at a particular point in time. If the real-world data that the model is processing starts to deviate, or 'drift', from the original training data, it can lead to decreased accuracy and reliability of predictions.

Performance and explainability

- [ ] Are there limitations on commercial use under the relevant licence terms?

- [ ] What is the inference speed? What are the parameters, and how might these be adjusted for our particular business purpose? What impact do these parameters have on our ability to deploy the AI system on commodity hardware?
  Tip! Parameters of an AI system refer to the variables that determine how input data is translated into the desired output, essentially defining an AI model's behaviour. The nature of an AI model's learning operations may influence whether it can be deployed on commodity hardware, rather than requiring AI-specialised hardware.

- [ ] How will we measure performance?
  Tip! Consider uptime availability, accuracy, reliability, scalability and speed. More qualitative metrics may also be appropriate, depending on the AI system's purpose.

- [ ] Do we have the right skills or domain expertise to assess the AI solution?

- [ ] What documentation do we need from the supplier to enable us to measure performance and understand – and potentially explain to auditors, regulators, affected individuals and others – how the AI system works, its capacities and limitations, and how to use it and interpret outputs?

- [ ] What reports should we be receiving from suppliers and how frequently?
  Tip! Consider how these reports may need to feed into internal reporting obligations (eg to the board), as well as regulatory reporting obligations (eg to regulators or the market).

- [ ] Will knowledge transfer, training, code walkthroughs and support be part of the services the vendor provides?

- [ ] What is the process for updating the solution (eg to scale it, and allow for feedback, improvement and change in requirements)?

- [ ] Who will be liable if the system fails to perform or causes harm?

Questions to ask the business about high-risk AI procurements

AI safety, security and resilience

- [ ] How is the algorithm protected?

- [ ] What safeguards are in place to prevent manipulation of training datasets (ie data poisoning) or the use of inputs designed to cause the model to make a mistake? Where is this data located?
  Tip! Data is often moved to a different location for training purposes – consider whether the protections in relation to the new environment are sufficiently robust.

- [ ] What are the supplier's processes for managing harmful outcomes?
  Tip! Implement internal processes for following instructions set by developers to manage risks specific to your organisation's use case.

- [ ] What circuit-breakers, human oversight or human-in-the-loop and other controls are in place to identify issues, and to override the output or operation of AI systems, especially for high-risk use cases?

Rights over system and datasets

- [ ] Was the system fully developed in house by the supplier, or does it incorporate third-party components? Can the supplier provide records that identify third-party inputs and components, including open source models/code/datasets?

- [ ] If the AI system/model/data is (wholly or partly) licensed in from a third party, does the supplier have adequate rights to pass this on to us?

- [ ] Who should own the AI inputs and AI outputs?
  Tip! Consider the rights you might want in any AI inputs (including training data, AI models and their improvements, prompt templates and the prompts themselves), AI outputs and related documentation. This might be ownership or a licence to use the IP rights, such as copyright and patents, as well as contractual tools, like confidentiality and exclusivity.

- [ ] What assurances / remedies are available from the supplier (and/or its upstream suppliers) to address IP infringement risk? Do we need additional assurances / remedies?

Termination

- [ ] What termination rights are required, particularly if the AI system does not function as desired?
  Tip! Termination rights should be tailored to the situation, particularly as termination for breach may be difficult if the agreement does not contain clear, measurable requirements for the AI system.

- [ ] What disengagement assistance is required, and for how long?
  Tip! A replacement AI system may take longer to implement than a typical IT system, particularly if extensive model customisation/training is required. A longer period of disengagement support may be required.

- [ ] What data, algorithms and other artefacts will be delivered up on termination?

ESG

- [ ] Who is responsible for training the AI? What safeguards are in place for protecting these workers?
  Tip! Modern slavery and other human rights factors should be a key consideration in the responsible procurement of AI, particularly in light of reports of workers being exploited to undertake the labour-intensive work of training models and developing safeguards.

- [ ] How might this procurement impact our carbon footprint? Does it simply shift emissions to a different part of the supply chain?
  Tip! The training and use of AI models is highly power intensive, and can have a material impact on an organisation's overall carbon footprint. It will be important to consider this impact, and balance it against any decarbonisation gains (eg if the AI solution is being deployed to assist with emissions reduction / more efficient operations).

- [ ] Have any other key ethical and sector-specific considerations been considered in (eg social impacts of relevant decision-making)?

Insurance

- [ ] What impact might this procurement have on our insurance coverage?

10 considerations when drafting and negotiating AI contracts

How to get prepared for AI procurements

1 Ensure procurement and IT teams are forewarned to run **clickwrap,** or equivalent standard terms, past Legal – they are often not fit for purpose. Beware hyperlinked terms of use that can be updated by AI vendors at will.

2 AI system development and higher-risk AI deployments may warrant a **bespoke contract** rather than amending an existing agreement to cover AI.

3 Your contract is just one way to manage your AI procurement's risks. Consider what **other organisational steps** can be taken to mitigate risk.

4 Be mindful that AI systems may be rolled out across your supply chain without your knowledge. **Review existing contractual arrangements** to ensure they require appropriate notification and approval where AI systems are being deployed.

5 Consider, as part of your standard data-handling practices, what **data you may wish to retain** (and in what form) for potential AI training, validation and testing purposes, and update your data retention and deletion programs accordingly.

This is not an exhaustive list – contractual positions should, as always, depend on:

- whether the solution is a commercial-off-the-shelf (***COTS***) vs a bespoke or customised system
- whether it is provided on-premises or is software-as-a-service (***SaaS***)
- the leverage of each party
- the risk profile of the AI system, having regard to its intended use

1 Governance and reporting

Significance: By its nature, AI is uncertain and iterative. This means that the relationship between customers and their AI vendor throughout the arrangement is particularly important.

Finding a partner that you trust to work constructively with you to figure things out – especially as the system, the risks and the regulatory environment evolve – should be a critical consideration during diligence and vendor selection. Governance and reporting requirements do the heavy lifting while the contract is on foot.

Actions:

- Governance provisions should establish a committee comprising representatives from both parties, responsible for managing strategic direction, risk management, issues resolution and overall service performance.
- This committee should meet regularly, and provide a forum for transparent and constructive discussion regarding performance (including any bias, defects, safety or ethical issues), risk management, controls effectiveness, any corrective action and the evolving regulatory environment.
- Regular reporting on these matters should address the same issues.

10 considerations when drafting and negotiating AI contracts

2 Customer requirements, performance standards and warranties

Significance: It is common for AI developers and vendors to insist that their AI system is provided 'as is', and that any improvements the customer makes and any output the system generates are at the customer's risk. Suppliers also often resist warranting compliance with documentation because, by their nature, AI models are continually evolving, and require an iterative approach to development and training.

This is where context matters – eg the level of assurance and warranties for a COTS tool will be quite different from what should be provided for an AI managed service. Consider whether you are procuring an AI system, or a particular outcome, and tailor your contract accordingly.

Actions:

- Consider the following performance metrics:
  - **Uptime availability** – This is a measure of how long a system is working and available (usually expressed as a percentage of time).
  - **Accuracy and reliability** – This refers to the tool's predictive power, including how consistently it performs over time. The terms should clearly define what constitutes a 'correct' result in the context of the intended use.
  - **Speed** – The response time of an AI system can be critical, particularly for real-time applications. The contract should specify acceptable response times under normal operating conditions.
  - **Resilience** – This refers to the capacity to withstand, and quickly identify and recover from, unexpected adverse events and disruptions. Contracts may need to include predefined tolerance levels for periods of disruption.
  - **Scalability** – This metric assesses how well the system can handle increased loads or larger datasets. It is crucial to ensure that performance doesn't degrade as more data is processed.

    Remember that these metrics will need to be tailored to each specific situation – what works for one application or industry might not work for another.
- Consider requirements to ensure the AI system includes functionality to log / record performance and other information.
- The AI system's interoperability with other platforms should also be addressed.

3 Training, testing, support, maintenance and remediation

Significance: AI systems should be continually monitored and tested throughout their lifecycle to identify whether: (1) changes in the underlying model, data or the production environment are affecting the anticipated results; and (2) data governance and other controls remain effective. Where issues are identified, AI systems may need to be recalibrated and/or retrained.

Actions:

- Testing, validation, support, maintenance, recalibration and remediation requirements should also be more prescriptive than under traditional technology procurements, and will need to continue throughout the implementation phase. Consider whether adversarial testing to identify dangerous capabilities is appropriate, particularly in high-risk use cases (as proposed by the Government's draft mandatory AI guardrails).
- The contract should also address responsibility for training, acceptance testing and recalibration / retraining and, where appropriate, contemplate close collaboration between the customer and supplier.
- Contracts may need to contemplate what happens (and the cost allocation) if new data sources or the introduction of a new environment mean that the algorithm needs to be retrained.

4 Data

Significance: The scope, provenance, accuracy and completeness of data used to train and test AI systems will have a direct correlation to an AI system's potential outputs and performance. This will also have direct implications for explainability and IP, as discussed below. Discussions around what data is being used, parties' rights and responsibilities regarding that data, and the data governance processes to be implemented, should be a high priority.

Actions:

- Consider (and, where appropriate, specify) requirements regarding the provenance, rights / permissions, format, accuracy, suitability, completeness, sensitivity and representativeness of training and testing data. This will need to be assessed on an ongoing basis.
- Stress-test how your vendor will address limits or deficiencies in datasets.

10 considerations when drafting and negotiating AI contracts

5 Oversight and explainability

Significance: AI systems must be able to be effectively overseen and monitored by humans, so that signs of anomalies, dysfunctions and unexplained performance can be detected and addressed as soon as possible. Customers need to understand – and potentially explain to auditors, regulators, affected individuals and others – how the AI system works (ie how outcomes are derived), and how to use it and interpret outputs. This is particularly critical for high-risk applications and those subject to regulation.

Actions:

- For bespoke systems, oversight tools and explainability may be able to be built into the design. For SaaS, the parties will need to negotiate/discuss what level of oversight and explainability is feasible. Importantly, the level of explainability may be particularly difficult to achieve for more 'highly intelligent' AI systems.
- Parties should consider what audit rights may be required, who should carry out the audit, and access requirements in connection with any audit.

6 Transparency and documentation

Significance: Supplier documentation for AI systems should not only operate as the specification against which the AI system will be assessed (presenting challenges for AI, which, by its nature, continues to evolve), but should also facilitate transparency and correct use.

Actions:

Consider the types of documentation that may be required, including:

- **technical specifications** – outlining the system's functionalities, how it works, its system requirements and integration capabilities;
- **training methodology** – including information on the initial dataset, feature selection process and model selection process;
- **testing reports** – including performance metrics used for validation/testing of the model and their results; and
- **a user manual** – including clear instructions on how to use the AI system effectively, along with troubleshooting tips for common issues and prompt templates.

10 considerations when drafting and negotiating AI contracts

7 Rights to AI inputs and AI outputs

Significance: IP clauses in traditional technology contracts tend to focus on the distinction between background IP, foreground IP, and rights to access and control particular datasets. However, in AI procurements, parties need to consider the value and limitations associated with their ownership of, and rights across, more granular categories, including AI inputs (eg training data, testing data, validation data, AI models / algorithms and prompts), the AI system's components (including source code), and their improvements, AI outputs and related documentation.

Actions:

- IP rights in relation to each of these categories need to be clearly articulated and considered, to ensure that the position taken does not inadvertently impact other protections (eg warranties).
- Vesting and assignment mechanisms and/or licences need to be put in place to give legal effect to the parties' commercial intent regarding IP treatment. Keep in mind that open source models / datasets / code are owned by third parties.
- Contractual mechanisms may need to be put in place to protect the intended owner's position where there is a risk that commercially valuable materials, such as some generative AI outputs or data, are not protected by IP rights. These mechanisms may include delivery/escrow and confidentiality obligations, and restrictions on the supplier's ability to use inputs/outputs.

8 IP infringement and other third party claims

Significance: Parties should consider the risk of IP infringement claims by third parties not only in relation to the model itself, but also the use of any third-party content or other data to train or improve the AI model. Third parties may make infringement claims against customers, alleging that they have created IP-infringing AI outputs, even if this was unintentional.

Non-infringement warranties for AI systems tend to be hotly contested, given these systems often source data and algorithms from a variety of sources. The nature of machine learning also means that the AI system may evolve over time, which can lead to unintentional infringement. The way generative AI systems operate means that instructions or prompts the user enters can increase the risk of AI outputs infringing in ways that are difficult for the system developer to foresee.

Aside from IP infringement claims, consider what other third party claims could arise in the context of how the AI system will be used - for instance, defamation, misleading and deceptive conduct or privacy complaints or claims - and how these can be addressed in your contract.

Actions:

In addition to requiring a warranty from the supplier that it has not infringed any third-party IP and that the customer will not infringe any third-party IP by deploying the AI system, customers can seek additional protection by:

- requiring the supplier to pass on the benefit of any warranties or indemnities provided by their upstream suppliers (including any model or training data vendors);
- limiting (contractually) the data that can be used for training and improvements;
- requiring that suppliers warrant that they have the necessary licences and permissions in place; and
- undertaking due diligence (including by requesting information about the supplier's own IP clearance processes) and conducting IP searches.

Consider what other contractual rights you may require to manage risks of defamation, misleading and deceptive conduct or privacy claims (as relevant). For example, monitoring of outputs, training of personnel using the AI, and inclusion of disclaimers or watermarking when outputs are generated.

10 considerations when drafting and negotiating AI contracts

9 AI safety and security

Significance: Geopolitical tensions, the evolving cyber threat environment and the accelerated mainstream uptake of generative AI have triggered an intense focus on AI safety and security. Regulators, policymakers, shareholders and the public will hold companies to account for failures to address AI-related security risks, including where they originate in the supply chain.

This requires consideration of AI-specific threats, including:

- **Data poisoning** – ie attempts to manipulate data or introduce misleading data into the training set, so as to corrupt the learning process, leading to inaccurate or biased outputs.
- **AI model flaws** eg (i) *model inversion* – which involves a threat actor using outputs from a model to infer the model's architecture or details about the original training data; (ii) *model tampering* – which involves manipulating the parameters, to generate inaccurate or biased results; and (iii) *backdoors embedded in models* – which cause a model to produce a threat actor's desired output when a trigger is introduced into the model's input data: eg via a malicious prompt.

Actions:

- Consider what contractual, technical and operational safeguards are required (and the allocation of responsibility and risk) to protect against threats to the AI system's confidentiality, integrity and availability.
- Consider requirements for the implementation of tripwires and controls to override, reverse or halt the output or operation of AI systems, as well as clearly defined requirements to notify, respond quickly and cooperate in connection with AI incidents.
- Consider whether certain AI models or systems should be prohibited, such that they cannot be used at all.
- Regularly review your processes for monitoring and evaluating the AI system and ensure such processes remain fit for purpose.
- Document how decisions are to made in the event of an actual or potential AI incident.

For more on AI safety and security, see our Insight: Why everyone is talking about AI safety and cybersecurity.

10 Change in laws

Significance: With the pace of AI regulatory change outstripping traditional contract cycles, and the consequences of non-compliance posing an enormous financial and operational risk, a general compliance with law clause may not suffice.

Actions:

Consider:

- addressing compliance with AI laws separately from general compliance with laws (similar to the approach now taken to compliance with privacy and work health and safety laws);
- including clauses that contemplate regular regulatory reviews and adjustments in line with evolving regulatory requirements; and
- addressing the potential consequences of regulatory enforcement action. Contracts should include clearly defined responsibilities and risk allocation for instances where non-compliance leads to financial penalties or the need to unwind transactions (including disgorgement), as well as obligations to cooperate in response to regulator scrutiny.

In March 2025 the Digital Transformation Agency released the latest version of its AI model clauses, intended for use by the Australian public service.

The clauses are designed to address three use cases:

1. **procuring services that may use AI;**
2. **procuring the development of AI tools; and**
3. **procuring software with embedded AI capabilities (these clauses are yet to be provided).**

While the clauses for use case 1 are relatively high-level (focused on approval rights, accuracy and record-keeping) the clauses for use case 2 are much more detailed, and cover the development, deployment, security, intellectual property rights, data management, record keeping, and training/testing/monitoring requirements for bespoke AI systems.

Key contacts

Valeska Bloch
Partner, Head of Cyber
T +61 2 9230 4030
Valeska.Bloch@allens.com.au

Jessica Mottau
Partner
T +61 2 9230 4587
Jessica.Mottau@allens.com.au

Gavin Smith
Partner, Co-head of Corporate, Head of Technology, Media and Telecommunications
T +61 2 9230 4891
Gavin.Smith@allens.com.au

Phil O'Sullivan
Partner
T +61 2 9230 4393
Phil.O'Sullivan@allens.com.au

David Rountree
Partner
T +61 7 3334 3368
David.Rountree@allens.com.au

Miriam Stiel
Partner, Practice Group Leader, Intellectual Property, Patent & Trade Mark Attorneys
T +61 2 9230 4614
Miriam.Stiel@allens.com.au

Elyse Adams
Partner
T +61 3 9613 8534
Elyse.Adams@allens.com.au

Dominic Anderson
Partner
T +61 2 9230 4099
Dominic.Anderson@allens.com.au

Tommy Chen
Managing Associate
T +61 2 9230 5303
Tommy.Chen@allens.com.au

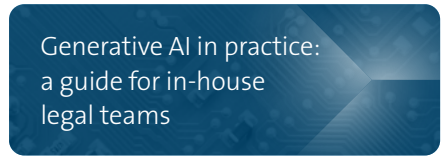

19904D