BIMA

BIMA AI Standards

1. Human first, and last
2. Uphold the rights of creators
3. Design for transparency
4. Innovate with purpose
5. Safeguard fairness and accurate representation
6. Preserve trust through privacy & security

Drafted by the AI Council:

Lisa Talia Moretti,
Digital Sociologist

Olly Veysey
Strategic Content Director

Katie Debrah
Gen AI & Marketing Consultant

Tom Good
Head of UX

Prateek Jain
Director Digital Experience Services

Rachel Baines
Fractional CMO and Strategist

Jedrzej Osinski
Snr Tech Analyst and AI Team Lead

Dora Moldovan
Co-founder and MD

Introduction

AI is being developed and implemented at an exponential rate. But there has been a lag in the introduction of meaningful standards, especially for the creative industries, to guide how we design, develop and deploy this technology. In response to the BIMA membership asking for more guidance, the AI Council has researched and led the development of six BIMA AI Standards.

Why standards?

Standards are written to encourage action. The AI Council decided that principles alone, in this context, would be too high level, difficult to implement, and not offer enough support on how to achieve a desired outcome. These standards will include guidance on how to achieve the stated goals, making them easier to 'action'. This is especially important when it comes to issues like transparency, privacy, fairness and sustainability. These are specialist topics and many of us require clear guidance to implement best practice.

With these standards, the AI Council aims to provide a blueprint of priorities, to help our members reflect on whether we're meeting the expectations set out by our community. And we seek to create a framework within which our members can work successfully, because setting out what good looks like, helps us all make better, more-informed decisions.

What research informed the standards?

The development of these standards were informed by research and activities taking place in government, academia and the technology and creative industries. This included reviewing new legislation, strategies and policies created by the US, EU, UK, Australian, Canadian and Chinese governments. We also read AI ethics and governance policies by technology companies including Adobe, Anthropic, Google, Meta, Microsoft, OpenAI, Stability AI and as well as principles and standards created by media organisations like CNN and the BBC. Early input from the BIMA membership was also captured via an online survey. This activity concluded the first phase, Discovery, and the creation of draft AI standards.

On Thursday 21st November 2024, the BIMA AI Council presented the final draft of the AI Standards on a BIMA AI CoLab call. We collected a second round of feedback from the BIMA membership during this 60 minute online session. This feedback was used to iterate the standards and progress them to the second phase, Alpha.

The Alpha AI Standards document was then made public so that feedback could be collected from a wider community of people. The AI Council led a public consultation from 17 February - 10 March 2025. Comments and additions were reviewed by the Council members and where relevant, incorporated into the final Beta v1.

An always-Beta approach

Innovation is ever-changing and for that reason, the AI Standards will always be in final Beta. We want to encourage current and future members of the AI Council, BIMA team and membership to ensure they always reflect best and next practice.

As a result of this approach, you'll be able to share feedback with us at any time. Feedback will be reviewed regularly and considered by the AI Council which is currently Co-Chaired by Oliver Veysey and Lisa Talia Moretti. To make it easier for the BIMA AI Council to incorporate your feedback and iterate the standards please reflect on the following questions and use them to shape your comments:

1. Does this standard resonate with me?
2. Does this standard relate to my work?
3. Does the information, research and guidance in this standard feel complete?
4. After reading the supporting sections, i.e. *Why it's important*, *What you can do*, *What to ask yourself*, do I feel confident enough to implement this standard in my work?
5. What else should I tell the BIMA AI Council that would make it easier to implement this standard?

FAQs

Q: I am not a BIMA member. Can I provide feedback?
A: Yes!

Q: I would like to stay up to date with what happens to the BIMA AI Standards. How do I do that?
A: Updates on the standards will be shared on the BIMA Slack channel #member_opportunities and BIMA's official LinkedIn page.

Q: Who will maintain the BIMA AI Standards?
A: This is a collaborative effort between the AI Council and BIMA, but we can't do it alone. As ever, we look to the BIMA membership and wider community for feedback, constructive discourse, interesting provocations and big questions to ensure that the Standards represent a horizon we are all sailing toward.

Q: As a BIMA member, do I have to use the BIMA AI Standards?

A: The BIMA AI Standards are not mandatory and there are no penalties for not using them. BIMA and the AI Council have created them to support best practice and hope that they are useful, helpful and meaningful to many.

BIMA AI Standard #1

Human first, and last

Prioritise human creativity, thinking, and oversight. Not to preserve outdated processes, but for the connections your businesses hope to make in future.

Why it's important

- The great AI disruption can feel personal. Many of us are designers or writers—crafts that generative AI is uniquely suited to disrupt. But this standard is not about protecting processes just because *we* like our jobs.
- Human centricity helps to preserve quality, and protect against homogenisation. We do not promote racing to AI-powered efficiencies at the cost of that meaningful connection between customer and your business that is - fundamentally - human.
- One thing AI cannot do is experience life. AI has no sense of how that experience feels, or the way it shapes us. Only another human can know that; only another human can connect with you through that shared experience. And creativity is about forming connections.
- AI interactions are designed to mimic human emotions, but are fundamentally synthetic. Removing true emotion and connection from communication risks creativity being flattened and false attachments being created with AI 'characters'.
- Businesses that jump to efficiency-first approaches will save money in the short term. But at what cost? Loyalty and trust requires connection.
- While AI-powered tools can generate countless ideas, finding a truly effective solution requires more than a large volume of output. It demands thoughtfully crafted human prompts, as well as the ability to evaluate, contextualise, and align those ideas with the audience's need.

What you can do

- Experiment with how AI can enhance your creative processes. But remember that people remain fundamental to creativity and connection with customers in the short, medium and longer term.
- This means making conscious choices about how and when to integrate AI now, soon, and in future. AI should not dictate our future because of hype and short-sighted decisions made today.
- Put the right KPIs and success metrics in place, across different time horizons, to guard against efficiency ruling them all. Traffic, engagement, and conversions are key but so too are brand awareness, perception and longer term customer loyalty.
- Stay nimble, embrace the disruption, encourage your teams to do the same.
- Support the evolution of creatives' skills to maintain their knowledge and ability to learn as a core value of your business
- Put smart, passionate people in a room together, charge them with identifying and solving problems, and encourage them to be candid. Then consider how AI helps them do those things, and not how AI replaces them doing those things.

Supporting research & data

- Conversational Persuasiveness of Large Language Models shows Chat GPT-4 to be 82% more likely to change your mind in direct conversation with you, than a human. Based on knowing very little about you. Wait til it knows everything about you - your preferences, history, biology.
- AI-powered creativity tools are leaping forward. Project Concept lets you source, organize, ideate, edit, and showcase concepts.
- A study conducted at Cornell in early 2024 used advanced AI models to generate prompts for text-to-image generators and found their prompts to be up to 20% more effective than those crafted by human creatives.
- Consider this 2024 study published by PhD students from the University of Arkansas: "The current state of artificial intelligence generative language models is more creative than humans on divergent thinking tasks" Consider also its limitations. For example, divergent thinking is essential to creativity, but not the totality. Creativity is more than just generating many ideas.

What to ask yourself

- Does your organisation have a human-centred approach to projects, including AI-enhanced ways of working? At what points are humans meaningfully inputting to the process?
- Do you fully understand the impact of having AI mediate your creative processes?
- Is the way performance and success are measured relevant for human and AI application? Can they be compared directly on the same metrics?
- Have you done your due diligence on what you will give up to gain any perceived advantage?
- What does testing, iterating, improving look like for your business? At what points are humans involved in your workflow?
- What human oversight is required to ensure you don't find yourself on the wrong side of regulatory lines thanks to using AI tools?

- "Artistic creativity is about making choices and develops in the context in which an artist works... You cannot – for now – transfer life experience into data." (Anne Ploin, Oxford Internet Institute, University of Oxford)
- Mango creates the first campaign generated by artificial intelligence for its Teen line. The results are... for you to decide.

BIMA AI Standard #2

Uphold the rights of creators

Protect and enforce the economic and moral rights of creatives, and support current and future efforts that allow livelihoods to flourish and businesses to grow.

Why it's important

- The UK creative industries are recognised internationally for creative talent and collectively, they generate billions of pounds. Government statistics show that the creative industries grew by 6.8% in 2022 and contributed £124.6bn to the national economy; that's more than the UK car industry, or aerospace, oil and gas. Protecting copyright contributes to a thriving creative and innovative economy.
- It's been reported that many free and enterprise large language models (LLMs) have used copyrighted materials without permission. We can't hit rewind but going forward we need to hold ourselves and each other to higher standards. In order to protect the livelihoods of creative communities, new normals urgently need to be created. This includes the contractual licensing of creative works in the technology industry, creating new revenue streams and ensuring creators are correctly attributed and fairly compensated.
- The principles of rewarding creators, preventing use of work without permission and encouraging innovation are enshrined in copyright law. In current and future legislation the rights of creative stakeholders must continue to be fairly considered, protected and enforced.

What you can do

- When developing an AI system, whether they are LLMs or other types of AI, train it on content that you have the rights and permissions to. Do not use copyrighted or private content without permission.
- As you embrace these tools, check which companies and models have taken a consent-based approach to training. Fairly Trained provides certifications and makes it easier to find and work with these organisations and models. For example KL3M is a family of models that doesn't violate copyright and has clearly documented the provenance of their data.
- To minimise risk of copyright infringement, consider using gen AI platforms as a way to develop ideas and iterate generated outputs to avoid exact duplication or substantial replication of a previously original idea.
- Fairly compensate creatives for their time and effort and do not use generative AI platforms as a bargaining chip to drive down the cost and value of human creative capital.
- Strengthen the collective efforts of the creative industries and their partners in protecting and advancing creative rights by contributing to critical discussions, research and the gathering and publishing of evidence. Support, spread awareness and encourage the use of tools like Content Credentials, Nightshade and Glaze.
- Ensure transparency in content provenance, consider supporting initiatives like the Coalition for Content Provenance and Authenticity which work to ensure transparency in content provenance like who created and owns the content as well as how it was created.

Supporting research & data

- In an essay on his blog, former OpenAI researcher Suchir Balaji notes that industry practices often do not qualify as 'fair use'.
- Thousands of creatives have signed a statement saying that creative works must be protected.
- The rights granted by copyright are online guidance from the UK's Intellectual Property Office that outlines the economic and moral rights of copyright holders.
- In February 2025, a US judge ruled in favour of Thomson Reuters against Ross Intelligence. The judge found that Ross Intelligence was not permitted to use Thomson Reuters' content to train a competing AI platform. This is a landmark case in which the "fair use" argument was not upheld.
- In January 2023, Getty Images commenced legal proceedings against Stability AI in the UK, alleging infringement of its IP rights. The case is still ongoing and could be years till a judgement is reached.

What to ask yourself

- Has due diligence been carried out to understand the data supply chain of the technology partners involved? Are there any controversies surrounding how they sourced their data? What assurances or statements are they making on the issue?
- Have efforts been made to examine potential licensing requirements? Is my organisation using someone else's IP and if yes, do I hold the right licence to do so?
- Have the authors and creators involved in the work been correctly attributed?
- Have all relevant authors and creators been fairly compensated for their work?

BIMA AI Standard #3

Design for transparency

Demonstrating integrity through transparency is essential to meet changing customer demands. More than ever, loyalty is earned by allowing customers to give informed, dynamic consent about when and how they engage with what you offer.

Why it's important

- Chatbots and automated systems are often used to answer questions or to deliver parts or all of an experience or service. As AI gets better at having conversations, it's not always clear to people if they are engaging with a human or machine. Be transparent up front to remove confusion. Feeling duped is horrible and bad for business, and worse, people won't stand for it.
- When people are engaging with an AI experience, a lot of data can be collected and substantial insight can be inferred. Be transparent about what data is being collected and why, and if user data will be used to train an AI system. Don't forget to note what data people should avoid sharing with the system. Honesty can help to build and maintain trusted relationships between organisations and their customers.
- Explainable artificial intelligence (XAI) is a set of processes and methods that supports understanding and trust in the outputs created by AI. Unless prompted through chain-of-thought prompts, black box AI models like ChatGPT or Stability AI provide an output without providing an explanation of how they got there. When using models like this, it's important to share with people that there could be potential risks to accuracy and that clearly identifying owners of content is not possible. New reasoning models, launched end of 2024, have enhanced reasoning capabilities in technical topics like maths, coding and science however there is limited evidence that this has improved the explainability of the outputs.

What you can do

- When conversational AI or a chatbot is being used to deliver a service or experience, make it clear to people who or what they are engaging with. Is this a human or machine? Spell it out.
- When using AI to deliver a service, ensure that documentation that explains how AI is used to deliver the service is available and always accurately reflects reality.
- Use straight-forward and accessible language when communicating with people about what data is being collected, for what purposes and if third parties will have access to that data. Ensure that consent is informed, dynamic, accurately recorded and simple to withdraw. People must be able to understand the specific implications of giving consent and how that may change over time.
- Make people aware of any risks associated with the outputs the solution provides. Where possible provide helpful guidance and advice on how outputs can be interpreted, understood and used.
- Where possible, opt to use AI models, tools and services that have transparent and explainable decision-making processes to foster stakeholder accountability and trust. Almost all future innovations will depend on AI technologies that are developed by other organisations. As a result, make every effort to choose (and support) tools and services that offer transparency.

Supporting research & data

- This article published in the Harvard Business Review, talks about the importance of building transparency into AI projects as well as how to do it. This includes decreasing the risk of error and misuse, distributing responsibility and enabling internal and external oversight.
- Mozilla Foundation and Thoughtworks have created a guide, AI Transparency in Practice, that shares research, practices and recommendations on how to implement meaningful transparency. The report highlights the importance of providing useful and actionable information tailored to the literacy and needs of specific stakeholders.
- Supported by their annual research, Hubspot created 'The Complete Guide to AI Transparency' that presents the benefits of being transparent, suggestions on how to create a transparent AI policy, and AI transparency best practices.

What to ask yourself

- Has straight-forward and accessible language been used to communicate clearly that AI is delivering some or all of the experience / service?
- Are users made aware of any risks associated with the outputs and has appropriate guidance been shared on how to interpret, understand and use the outputs?
- Is it clear to people what data is being collected, why and who else will have access to their data? Is it clear to people that data collected from them is or is not being used to train AI systems?
- How is the AI system and the collected data audited, measured, and proactively enhanced for privacy, security and transparency on a continual basis?
- If the use of AI is not disclosed, has the potential reputational damage to the business been considered?

BIMA AI Standard #4

Innovate with Purpose

Focus on the business of solving tangible challenges. Practitioners and businesses alike should evaluate how AI supports more effective outcomes, instead of being distracted by novelty and hype.

Why it's important

- Creative industries naturally adopt emerging technologies, usually by necessity, though often driven by hype. Creators' excitement about technologies like AI, AR, or NFT, under the influence of 'Hype Culture', can lead to commitment before choices are fully considered. While AI is a powerful tool, it is also resource-intensive, unpredictable, currently unproven over the long term, as well as being costly to develop, run, and maintain.
- Training LLMs is carbon intensive. Users are responsible for the demand as much as owners are for the supply.
- While not all AI is LLM-based, data centres supporting AI operations consume vast amounts of energy, and training a single large AI model can emit over 284 tons of CO_2, equivalent to the lifetime emissions of five cars.
- Creating purposefully creates less waste. AI is still an emerging technology, and whilst some projects can create impressive results, most AI projects fail. The time, resource, and environmental impact of failure can be huge.
- A problem-centric approach ensures that AI is used responsibly, where its benefits clearly justify its costs. By focusing on well-defined problems first, creative teams can unlock AI's full potential while avoiding unnecessary compromises.

What you can do

- Start by clearly defining the problem or opportunity and understand the specific needs of users or clients. Be conscious in avoiding choices driven by hype culture.
- Determine if AI is the most effective way to address the problem at hand, consider how the overall success of the project is measured, and if the application of AI is a meaningful and direct driver of its success.
- Make sure the disadvantages of AI are also considered, such as how it can be unpredictable, how machine learning models can get worse over time (data drift) or how outputs can't always be explained.
- Understand if there are other methods to achieve a similar outcome. Create a comparison of options and review which are most efficient, resource-intensive, sustainable, and produce the highest value. Prioritise ease of application and holistic long-term efficiencies.
- Weigh the sustainability, financial, ethical and accountability implications of using AI.
- Be transparent about the carbon impact of using your AI services, so users can make up their own minds about whether or not to use them.

Supporting research & data

- According to Gartner, "30% of Generative AI Projects Will Be Abandoned After Proof of Concept By End of 2025".
- Harvard Business Review reports that most AI projects fail. Some estimates place the failure rate as high as 80% – almost double the rate of corporate IT project failures from a decade ago. Analysis by Goldman Sachs calculates that data centres account for about 1-2% of global electricity demand, and data centre power demands will increase by 160% by 2030.
- Researchers from Rand interviewed data scientists and engineers and identified five leading root causes for the failure of AI projects. The first root cause for failure? Industry stakeholders often misunderstand — or miscommunicate — what problem needs to be solved using AI.
- Example:The CEO of Indian food delivery app Zomato asked restaurant partners to stop using AI-generated food images for marketing purposes. "AI-generated food/dish images are misleading, and we have received numerous customer complaints on this issue. Customers say that this leads to breach of trust, and it leads to higher complaints and refunds, and also, lower ratings",

What to ask yourself

- What is the problem that needs to be solved? Has it been defined in a way that doesn't bias one solution over another?
- Has an AI solution been selected as a selling point, or is it really the right solution to solve the problem at hand?
- Can AI be directly linked to the measure of a successful project outcome?
- Will choosing an AI solution impact the project approach and timings, or distract from delivering against other project outcomes?
- Does your AI solution lock you into a dependency on 3rd party providers that may not meet your future needs or standards?
- Are there simpler solutions that could achieve similar outcomes?
- Have the costs and sustainability impacts of AI been considered against its benefits?
- Is AI being treated as a magic wand, with the assumption it can solve all problems without rigorous interrogation or a proof of concept for success?

BIMA AI Standard #5

Safeguard fairness and accurate representation

AI tools, often developed by homogenous teams and trained on biased data, risk perpetuating harmful stereotypes and excluding diverse perspectives. Be proactive in your work to be fair and inclusive.

Why it's important

- AI is rapidly becoming ubiquitous and integrated into countless products and services, which become central to our lives. Unrepresentative data sets can result in unfair outcomes, poor experiences or error-prone results that can make it more difficult for some people to go about their lives. As AI advances and data sets get bigger and more complex, the biases contained within data sets can become harder to detect and address. This has significant potential to cause unintended harm.
- In order for AI to be beneficial to all, it is vital to grasp both the capabilities and the shortcomings of the technology.. AI operates without a moral framework and doesn't inherently know what is helpful, harmful and inaccurate. It doesn't understand fairness or what it means to champion diverse experiences.
- Language models often reinforce gender stereotypes, while image generation tools produce racially skewed results. Disability representation is also underrepresented by AI, often defaulting to narrow portrayals of human life.

What you can do

- Conduct a fairness and inclusion review, analyse the outputs for diversity and cultural sensitivity - to ensure they align with your chosen definition of fairness. Understand the risk to your brand of not conducting this review. Most important is to take action and make changes where necessary, based on the findings.
- Data diversity is essential, but so are testing training methods. A model's ability to generalise,i.e. learn to recognise objects it has not seen before, is influenced by both the diversity of the data and the way the model is trained. Ensure your team is up to date with the latest methods and check that the training data covers the representative time period for your interrogation to avoid "misplaced" results.
- Learn from case studies by looking at real-world examples like FN Meka to see how AI bias can have real-world consequences. By learning from these mistakes, we can create a more inclusive and ethical AI future.
- Be inspired by real people's stories and experiences to improve the quality of prompts used to create AI-generated content.

Supporting research & data

- In Kantar's Brand Inclusion Index 2024, 75% of consumers say that a brand's diversity and inclusion reputation influences their purchase decisions.
- The FN Meka case, where a virtual rapper rose and then quickly fell from fame due to a cultural appropriation backlash, highlights the importance of cultural sensitivity and real representation in AI.
- Dr. Joy Buolamwini is the founder of the Algorithmic Justice League (AJL), an AI bias expert, an artist and author. At SXSW 2024, she shares her research into studying algorithmic biases. Tests conducted by AJL found that facial recognition systems are more likely to be accurate on male-labelled faces with lighter skin tones. AJL found that facial recognition systems perform worst on female-labelled faces with darker skin tones.
- Dr. Michael Li,the founder and CEO of The Data Incubator, writes about how to build less-biased AI. The key he says? Hire a more diverse team.
- "Data diversity is key to overcoming bias, but it is not like more data diversity is always better." In this article on MIT News, two researchers discuss how thinking like a neuroscientist changed their approach to training a neural network. The article discusses how data diversity and adopting new training methods can help to overcome bias.

What to ask yourself

- Has the use of AI resulted in authentically representing our audience's diverse experiences? Is the AI generated output fair to all communities the initiative is intended for?
- Has a human reviewed the results of the AI output and sense-checked against the cultural and inclusivity requirements for the project?
- Have different participatory research methods (e.g. peers-observing-peers, self-documentation, co-creation) been used so that customers can have direct input into the development of the product or service?
- Have marginal groups or communities who are currently underserved been identified? What are their specific needs and what are the specific barriers and pain points that face them?
- What processes and tools are in place to measure diversity and understand the biases in tools?
- Is there shared awareness of the training dataset used for the AI system we're using? Is it well documented? *[Note: this question is for those teams who are building their own dataset for use]*
- When any proprietary data is included within AI solutions using techniques like RAG, was it validated for accuracy and reliability before implementation?-

BIMA AI Standard #6

Preserve trust through privacy & security

AI is changing the landscape for bad actors online, and therefore privacy and security cannot be afterthoughts. Your reputation, and trust in your AI-powered products and services, depends not only from what you deliver, but also from how your business protects customers.

Why it's important

- Governments around the world are increasingly aligning AI regulation with data privacy and protection legislation. Ensuring that privacy and security strategies are at the heart of every business is no longer a nice to have, it's an imperative. Outputs often involve sensitive data, such as insights into audiences and customers, as well as proprietary information. Upholding privacy and security as core principles are essential to protecting trust and brand reputation, but equally, to being a law-abiding corporate citizen.
- Society is experiencing a significant shift as people gain greater awareness of how data and technology can be misused—such as hacks, data theft, misinformation, and deep fakes. As understanding of these risks grows, people are becoming more cautious, privacy-conscious, and aware of the value of their personal data. In response, governments worldwide are increasingly mandating robust data protection measures, issuing warnings and imposing fines on organisations that fail to comply.
- Owning large amounts of data is a burden, over-collecting data not only increases storage and security risks but also complicates governance. Ensuring data privacy in AI is more challenging than in regular software as multiple independently secure sources can be merged, potentially exposing private information. Adopting a data-proportionate approach can mitigate this risk: collect only what is necessary and delete data once it is no longer needed..

What you can do

- Learn about and adopt privacy-conscious practices like avoiding inputting sensitive information into AI tools, and anonymise prompts and outputs wherever possible. This ensures privacy even when working with data-intensive creative projects.
- Review your current practices that are required under GDPR and ensure you're following up-to-date guidance on how to comply with any legislation.
- Put an AI Policy in place so that employees know what they can and shouldn't do when engaging with publicly available AI platforms. In regulated industries, where the risks to non-compliance could be higher, organisations should consider having a bank of accepted and tested prompts included in such a policy as a way to encourage safe and productive use of AI platforms.
- Ensure the required compliance measures are integrated into your organisation's data practices. Organisations that regularly process sensitive or enterprise data, whether for training AI solutions or employing enhanced prompting techniques such as RAG, should consider obtaining certifications like ISO27001.
- Include comprehensive data processing clauses in contracts to clearly define how 3rd party data is collected, processed, and stored.More data doesn't mean better. Implement a data proportionate approach by only asking for and using the data needed, for the duration it is required. This protects organisations, clients and the people engaging with project outcomes.
- When storing data used for machine learning models, use state-of-the-art encryption techniques to secure the data at rest and in transit.
- Design for trust by going beyond compliance and advocate for users. Include proactive data protection, transparent permission protocols, and clear disclosures about data usage.

Supporting research & data

- A report from the Human-Centered Artificial Intelligence Centre at Stanford University looks at privacy in an era increasingly dominated by AI and asks the question, *how do we protect our personal information?* The report puts forward recommendations like shifting from opt-in to opt-out consent models, taking a supply chain approach to data privacy and considering collective solutions as opposed to only focusing on the privacy rights of individuals.
- The UK's National Cyber Security Center has shared their helpful advice in a number of blogs and guidance documents that cover topics like the security of AI systems, and the risk of LLMs like ChatGPT. For example, their updated machine

What to ask yourself

- Is your organisation abiding by the most up-to-date rules and regulations regarding data protection and privacy acts that are enforceable by law?
- Are you going beyond compliance by ensuring that you're doing the right thing and not just ticking off the things you have to do?
- Do you have a framework in place to make sure customers can trust every interaction they have across data, content and transactions?
- Do employees/colleagues have the appropriate training and guidance so that they can perform any data activities required in their role in an ethical and legal way?

learning security principles reinforce the importance of supply chain security and place more focus on 'security by design'.

- Axelerant, an experience agency, ran a webinar exploring RAG AI. In the webinar they discuss how organisations can create their own LLM, host it on their network and ensure customer data is kept private.

- What data is being collected across the organisation? Is it certain that this type of data collection meaningfully improves the quality of the product/service?
- Is all collected and stored data being protected with techniques (e.g. anonymisation) and security features that are appropriate for the kind of data you're safeguarding?