BCS, The Chartered Institute for IT's submission to the Department for Science, Innovation and Technology Cyber Security of AI Call for views

August 2024

Compiled by BCS Senior Policy and Public Affairs Manager, Claire Penketh

BCS
The Chartered Institute for Information Technology
3 Newbridge House,
Newbridge Square,
Swindon SN1 1BY

Introduction

The Department for Science, Innovation and Technology on Cyber Security of AI - is a two part intervention, including a voluntary Code of Practice on AI security that will form a new global standard.

As part of the consultation, BCS contributed to a round-table organised by DSIT's Issy Hall, the Cyber Security, AI and EmTech Standards Lead. The panel included a BCS Fellow, Adam Leon Smith, who is an AI standards expert, and experts from the BCS Information Security Specialist Group, (ISSG), and the Software Testing Group (SIGIST)

The panel's responses will form part of this submission. BCS has over 50 specialist interest member groups. Several of our recommendations are also drawn from our experts' reactions to the CrowdStrike outages, and our responses on software resilience and security for businesses and organisations, the cyber security code of practice, and the cyber resilience of the UK's critical national infrastructure.

Executive Summary

Overall, the group agreed with the main principles outlined in this call for views, in particular aligning the wording of the Code's content with the future standard developed in the European Telecommunications Standards Institute (ETSI).

There were also additional recommendations:

- Introduce a Code of Practice with mandatory breach reporting and quarterly reporting on risk (including third-party risk) rather than a voluntary code, as currently. This would ensure effective take-up and a level playing field for all firms.
- There must be an accountable person on a company's board (or at senior management level) who is a competent, ethical, registered technical professional, preferably with chartered status.
- Standards should be adopted and developed that set out clear basic common terminology for best practice.
- Companies need to be encouraged to focus on their business continuity and disaster recovery plans to deal with the consequences of an outage.
- Develop an ongoing government-led awareness-raising campaign on cyber-security and cyber resilience, backed by industry partners.
- The Code of Practice needs to be accessible, understandable and relevant for SMEs who may need extra support to implement it. We recommend the government set up an easily accessible one-stop shop on gov.uk for SMEs need to meet all their cybersecurity information and advice needs

Who we are/demographics

BCS, the Chartered Institute for IT is a charity and is the professional body for the technology sector. Our purpose, as defined by a Royal Charter, is to promote and advance the education and practice of

computing for the benefit of the public. BCS has over 70,000 members and brings together academics, practitioners, industry and government to share knowledge, promote new thinking, inform the design of new curricula, and shape government policy.

BCS is an organisation of medium size, based and headquartered in England.

Call for Views Questions

Q7. In the Call for Views document, the Government has set out our rationale for why we advocate for a two-part intervention involving the development of a voluntary Code of Practice as part of our efforts to create a global standard focused on baseline cyber security requirements for AI models and systems. The Government intends to align the wording of the voluntary Code's content with the future standard developed in the European Telecommunications Standards Institute (ETSI).

Do you agree with this proposed approach?

Yes

Q8. In the proposed Code of Practice, we refer to and define four stakeholders that are primarily responsible for implementing the Code. These are Developers, System Operators, Data Controllers (and End-users).

Do you agree with this approach?

Yes

Please outline the reasons for your answer.

Whilst the panel agreed with the four stakeholders proposed, the panel thought two more stakeholders should also be represented:

1. Society at Large - as the impact of cyber security failures impacts everyone.
2. Internal auditors - i.e. those accountable at a board level/ senior management level for cyber/AI.

Q9. Do the actions for Developers, System Operators and Data Controllers within the Code of Practice provide stakeholders with enough detail to support an increase in the cyber security of AI models and systems?

No

Please outline the reasons for your answer:

Steve Sands, Chair of the BCS Information Security Specialist Group said there's no one size fits all when it comes to governing AI, which is developing rapidly: "The companies creating AI solutions and systems are of very different sizes, ranging from the tech giants, down to the one-man band companies. Different AI systems will have differing risks and impacts on the users and the public.

"Many of the standard practises for securing any system will apply equally to AI systems as they do to everything else, and the development of AI obviously follows what is becoming a reasonably problem path now in terms of the phases of it, and that introduces new dimensions and new intricacies, new complexities. But there are established cybersecurity principles and frameworks in place that will handle much, certainly not all, of the practices."

Nicola Martin, Chair of the BCS Software Testing Group said: “It’s a very hard question to give either a yes or no answer to.” She said clarification was needed around whether the guidance, when it came to AI, would be iterative or normative, and she emphasised the importance of having accountable skilled technically competitive people, including testers, who could implement the Code of Practice.

She added that SMEs for instance, don’t necessarily have internal auditors, inhouse experts or an accountable person in the company when it comes to technology. There could be a risk too that SMEs won’t comply with the code, if it is voluntary, because as far as they are concerned, to do so essentially slows down innovation.

Nicola and Steve reiterated the point that it was important that the Code of Practice had text about the impact of the AI on wider society, and Nicola added there needed to be clarity around who the testers of the AI systems were.

Adam said this highlighted the importance of standards that will define suitably technical professionals and the skills required. He said the proposed standard should consider providing a conceptual mapping to the stakeholders used in the EU's AI Act.

The next questions are going to ask you specifically about the Code of Practice that has been designed and proposed by DSIT. There will be a question on whether you support the inclusion of each principle in the Code of Practice and whether you have any feedback on the provisions in each principle.

Q.10 Do you support the inclusion of Principle 1: “Raise staff awareness of threats and risks within the Code of Practice?”

Yes

If the government set up an **ongoing government-led awareness-raising campaign** on cyber-security and cyber resilience, backed by industry partners, then wording could be amended to:

- ***Ensure staff keep up-to-date about the latest threats, advice and support available on gov.uk cyber security site.***

Ideally, we’d like to see an **easily accessible one-stop** shop on gov.uk, in particular to help SMEs meet all their cybersecurity information and advice needs. Currently there is a proliferation of sites which can be confusing to smaller organisations who might not have the staff capacity to monitor the latest developments.

Best practices should already incorporate cybersecurity, but publicising these practices would be beneficial. Setting the most secure defaults from the start is essential. Government departments could share information on recommended default settings.

Q11. Do you support the inclusion of Principle 2: “Design your system for security as well as functionality and performance” within the Code of Practice?

Yes

ISSG member Andrew Wright, Head of Cybersecurity at two hospital trusts in northwest London, felt the text could be expanded today:

Design your system for security as well as functionality and performance and **ensure they are informed by AI specific risks**

Q12. Do you support the inclusion of Principle 3: "Model the threats to your system" within the Code of Practice?

Yes

Suggested changes - Model the threats to your system **throughout its life-cycle**

Q13. Do you support the inclusion of Principle 4: "Ensure decisions on user interactions are informed by AI-specific risks" within the Code of Practice?

Yes

Q14. Do you support the inclusion of Principle 5: "Identify, track and protect your assets" within the Code of Practice?

Yes

Q15. Do you support the inclusion of Principle 6: "Secure your infrastructure" within the Code of Practice?

Yes

The panel felt it might be a duplication of 14, although the point was made this could be expanded to refer to external resources.

The panel thought that additional words that could be added to that sentence could be **Secure your infrastructure, assets and third party suppliers**.

Q16. Do you support the inclusion of Principle 7 "Secure your supply chain" within the Code of Practice?

Yes

Q17. Do you support the inclusion of Principle 8: "Document your data, models and prompts" within the Code of Practice?

Yes

Q18. Do you support the inclusion of Principle 9: "Conduct appropriate testing and evaluation" within the Code of Practice?

Yes

It was suggested by the panel there could be a clearer definition of whom the suitable testers were. Adam said such work should be undertaken by 'suitably skilled testers', but Nicola said the testing could be done by lots of different people within an organisation so how do we define who the tester is?

Q19. Do you support the inclusion of Principle 10: "Communication and processes associated with end-users" within the Code of Practice?

No

Steve Sands said: "I'm concerned that we're keeping this to end-users and it needs to be wider than that and include other stakeholders and maybe regulators."

Q20. Do you support the inclusion of Principle 11: "Maintain regular security updates for AI models and systems" within the Code of Practice?

Yes

Q21. Do you support the inclusion of Principle 12: "Monitor your system's behaviour and inputs" within the Code of Practice?

Yes

Q22. Are there any principles and/or provisions that are currently not in the proposed Code of practice that should be included?

Yes

Adam Leon Smith proposed that what is missing is a "secure your model" principle. He said: "There's a gap in state-of-the-art technical requirements to detect and avoid specific AI-specific vulnerabilities. As well as technology guidance, this could include model hardening levels, for example.

"There's a very complex value chain in AI, and vulnerabilities can be injected by adding data to a file share that's not properly secured. So having a very life-cycle based approach is necessary to really secure AI systems end-to-end, at least where there's a significant risk to consumers."

Q23. [If you are responding on behalf of an organisation] Where applicable, would there be any financial implications, as well as other impacts, for your organisation to implement the baseline requirements?

Don't know

Q24. Do you agree with DSIT's analysis of alternative actions the Government could take to address the cyber security of AI, which is set out in Annex E within the Call for Views document?

Don't know

Q25. Are there any other policy interventions not included in the list in Annex E of the Call for Views document that the Government should take forward to address the cyber security risks to AI?

Don't know

Q26. Are there any other initiatives or forums, such as in the standards or multilateral landscape, that that the Government should be engaging with as part of its programme of work on the cyber security of AI?

Yes

Adam Leon Smith: Yes, the standard should look at EN 303 645 for some inspiration. This is a good example of a horizontal IoT standard that deals with the specific cybersecurity aspects of a technology class.

Sarish Chandra, ISSG member and Director of Security, Risk & Compliance GE Health Care, suggested there could be a certification component such as the Cyber Essentials programme where staff could get trained, possibly using a modular approach, which could be geared to specific instances. He said:

“Then you know you're using AI securely or you're developing AI securely, or you're offering a secure AI service as a vendor to your customers. And there are obviously different tiers of how AI could be used, developed and consumed, and those could be the tiers and a certification model that would be valuable in a commercial setting.”

Q27. Are there any additional cyber security risks to AI, such as those linked to Frontier AI, that you would like to raise separate from those in the Call for Views publication document and DSIT’s commissioned risk assessment. Risk is defined here as “The potential for harm or adverse consequences arising from cyber security threats and vulnerabilities associated with AI systems”.

Yes

Q28. Thank you for taking the time to complete the survey. We really appreciate your time. Is there any other feedback that you wish to share?

Yes

Additional comment on standards:

Adam Leon Smith BCS Fellow, Chair of the BCS Fellows Technical Advisory Group previously set up information security systems for start-ups through to banks. His primary focus now is on AI standardisation.

He said: “Developing it in ETSI means that there's a global set of stakeholders inputting into it, which is good. It also fits very well with the Product Safety thinking mindset that's being used in Europe.

“There’s a market-need for a standard in this space. There isn't really a state-of-the-art baseline for securing a machine-learning model in comparison to a Windows Server or a Java application. We just don't have that baseline level of what's an acceptable configuration, what's an acceptable level of risk, and that applies even to models that are decades old. It's not just generative AI, which I think a lot of the focus has been on.”

Rob Wilson, Cyber Security Consultant Specialist, Telecoms and ISSG member: “We need standards to be in place so that we know the quality of the products that we're buying meet certain criteria.

Patrick Burgess Cyber Security Consultant, SMEs’ and ISSG member: “It feels like we need to be setting framework rules for innovation. It feels like AI is innovating on an hourly basis. So, the people who innovate should understand that there are rules you have to adhere to. There should be a set of standards that include getting audits, penetration tests, third party checks, whatever you need to be doing to demonstrate that you are doing things in a way that is both ethical, responsible, reasonable and safe.”

Additional recommendations:

- Introduce a **Cybersecurity Code of Practice with mandatory breach reporting** and quarterly reporting on risk (including third-party risk) rather than a voluntary code, as currently
- Develop an **ongoing government-led awareness-raising campaign** on cyber-security and cyber resilience, backed by industry partners.
- Set up an **easily accessible one-stop** shop on gov.uk for SMEs need to meet all their cybersecurity information and advice needs.

On a mandatory Code of Practice Steve Sands said: “We've been advocating that any code of practice around cyber should be mandatory rather than voluntary because the absolute truth is that for as long as these things are voluntary, then a substantial part of the market will choose not to do it and cyber, whether it's underpinning AI or whether it's underpinning anything else.

On an ongoing government-led awareness-raising campaign Steve added**: “**The best practice guidance is distributed around various bits of government and non government around the UK and and really and truly I would always want to go to one single version of the truth, a single point of contact for anything to do with cyber.”

Thanks to all those who took part in the consultation:

Adam Leon Smith Chair of the BCS Fellows Technical Advisory Group, AI standardisation expert

Steve Sands, Chair of the BCS Information Security Specialist Group (ISSG)

Nicola Martin, Chair of the Software Testing specialist group Fellow of the BCS and Women’s Engineering Society.

Patrick Burgess, who has a focus on SMEs (ISSG)

Sarith Chandra Director of Security, Risk & Compliance in Health Care (ISSG)

Rob Wilson, Cyber expert with a Telecoms focus

Matt Mason Nottingham Trent University, Head of Infrastructure and Operational Security (ISSG)

Andrew Wright NHS Digital Leader specialising in Cyber and Information Security in NW London (ISSG)

Immo Huneke an Expert Software Engineer with Zuhlke Engineering Ltd., member of South London BCS Branch