HOUSE OF COMMONS

Science, Innovation and Technology Committee

Oral evidence: Cyber resilience of the UK's critical national infrastructure, HC 559

Wednesday 24 April 2024

Ordered by the House of Commons to be published on 24 April 2024.

Watch the meeting

Members present: Greg Clark (Chair); Dawn Butler, Dame Tracey Crouch; Dr James Davies; Rebecca Long Bailey; Stephen Metcalfe; Graham Stringer.

Questions 57 - 125

Witnesses

I: Professor John Goodacre, Professor of Computer Architecture, University of Manchester, and Challenge Director, Digital Security by Design, Innovate UK; Richard Grisenthwaite, Executive Vice President and Chief Architect, Arm.

II: Henry Harrison, Co-founder and Chief Scientist, Garrison Technology; Dr Vasilios Mavroudis, Principal Research Scientist and Co-Theme Lead, AI for Cyber Defence Research Centre, Alan Turing Institute.

Examination of witnesses

Witnesses: Professor John Goodacre and Richard Grisenthwaite.

Q57 **Chair:** This morning we continue our inquiry into the cyber resilience of the UK's critical national infrastructure. To help us with that, I am pleased to welcome our first pair of witnesses. We have with us here today Professor John Goodacre, professor of computer architecture at the University of Manchester and challenge director of the digital security by design programme at Innovate UK, the Government innovation agency. Joining him is Mr Richard Grisenthwaite, who is the executive vice president and chief architect at Arm.

Thank you very much indeed, both of you, for joining us today. Perhaps I can start with Professor Goodacre. Could you just brief the Committee on what the key principles are of secure by design and default?

Professor Goodacre: Yes, I can. What I might do is expand that question just slightly to bring the definition of cyber security into that context as well. If we look at what cyber security is today, it is primarily the risk management and the response—how do you maintain your systems if and when you are attacked? That is done by the person who is responsible for that system.

If we start looking at by default, that is where the responsibility is moving earlier in the supply chain, to the people building the systems, where they are trying to reduce the likelihood of there being vulnerabilities that can be exploited. There are fewer manufacturers than there are users in that regard, so it is good to start pushing the responsibility for security down into the supply chain.

By default is making sure that the product does not have as many likely vulnerabilities by default—so we do not turn on features if the user is not going to use them, because all that does is make it more likely that there is something that they do not know about that can be attacked.

When we get to by design, this is where security is implemented within the design process of the products that those companies use or build. This is where they select components that are being architected or built in a way that can protect the user from the effects of vulnerabilities. Basically, in the design process, you are, by design, designing security principles into those products.

Q58 **Chair:** Thank you very much indeed. Mr Grisenthwaite, you are the chief architect of Arm and, therefore, involved in planning and designing new innovations. How prevalent is this in your thinking in the company?

Richard Grisenthwaite: Security is one of the fundamental principles that we have to obey, because the whole premise of computing is that you have to be able to trust it. The issue is not so much whether we care about security. We absolutely do and have always put in lots of features that, for well-written software, will inherently be secure.

The challenge is that there are increasingly complex techniques that people use to find minor flaws in software and then exploit security. The question then is what we do about that and how we make it secure in the face of improperly designed software, even with small amounts of error. There are millions of lines of code in software, and the idea that you can have perfect software is a fallacy. The reality is that everybody makes mistakes, which do not always get picked up, and so there will always be these security vulnerabilities.

For the last 20 years, we at Arm have been incrementally adding features to try to address security issues as we become aware of them. There is a particular class of problem that seems to be incredibly prevalent, which is called memory safety. I will not go into the details of exactly what that means, but some 70% of the reported vulnerabilities in software for many years—it has been an almost-consistent run rate—find their root cause in memory safety.

We have been looking for years at ways in which we can enhance the architecture to provide that resilience when there are these issues in software so that, essentially, either they can be found when you try running them or, in the event that somebody finds an exploit, you can limit the damage that can be caused as a result of that.

Q59 **Chair:** Within the company, how do you take steps to make sure that there are not people who, as it were, smuggle themselves into the company and embed things that might be useful to outsiders later on?

Richard Grisenthwaite: It is a good question. We have a very strong culture of peer review of our hardware as we design it. The way in which you design hardware is that you write it in what is called a hardware description language, or HDL. That stuff is regularly reviewed by peers, so it would need collusion between multiple people for somebody to be able to smuggle something in like that.

Q60 **Graham Stringer:** Can you tell us what the state of play is at the present time in using secure by design and default in the national infrastructure?

Professor Goodacre: "By design" and "by default" are two fairly new terms. They are terms that have evolved in the last five years. The definition that I gave you is one that is evolving as the preferred definition, primarily through an agency called CISA in the US, which is cementing that terminology.

Therefore, in terms of those technologies specifically being deployed in CNI today, I do not believe that they are going beyond the attack reduction aspect. By default is something that they can do today, so that, when you deploy a system, you consider its configuration and minimise its attack.

There is a challenge for by design. In the digital security by design programme, which I am director for, we have been looking at a

technology that Richard alluded to, called CHERI, which protects against the exploitation of these memory safety issues. That is not commercially available today, so that clearly is not part of the CNI deployment today. However, there are other by design approaches.

For example, there is one called root of trust. You may be familiar with the fact that, in the Windows 11 release, they said that you have to have a certain version of a PC to be able to run Windows 11. That was to implement a root of trust. This is something that can start stopping ransomware, for example, so you can always get your system back by pressing the reset button if you have that version of Windows 11. That kind of technology is optionally available to the CNI sector. I do not have the information to provide on whether it is actively being selected to be used.

Richard Grisenthwaite: In terms of the digital security by design programme that Professor Goodacre referred to, which was partially funded by UKRI, Arm created a prototype whereby we took a high-end process that you might, for example, have in a high-end server and added these new technologies there. We have made available several hundred of the boards associated with this to many universities and companies in the UK to allow them to evaluate the technology.

That programme is still ongoing. We are getting a lot of positive feedback that it can help alleviate quite a number of these issues. There was a very impressive report from Microsoft Research a few years ago that showed that, in principle, this technology could address some 70% of the issues.

People are looking at how this could be commercially deployed, but, at the moment, we are not at that stage. We are still at that transition of research towards development.

Q61 **Graham Stringer:** Talking about security is difficult, but are there any particular problems in protecting our national infrastructure by improving cyber security?

Professor Goodacre: There are two ways to look at the protection of CNI. The first is what they can do today that they are not doing. As part of the digital security by design programme, I have been hosting a number of workshop panels with members of the critical national infrastructure communities, particularly in energy security, and communications and telecoms. We find that there are probably a lot of hygiene issues that could be addressed today.

For example, we were given information that a lot of the control of our infrastructures is in cleartext. There is no encryption. It is, basically, open and non-authenticated. You do not even need passwords to be able to control what the infrastructure is doing. Those issues should and could be promoted for addressing today.

The second stage to consider is digitalisation, and the benefits of bringing in digital systems and of the economy, if you like, of digitisation, and making sure that those are not also then going to be troubled by the 70% of memory safety issues that we talk about here. There is the immediate need to do something to clean up the system today, but, as we do that, let us make sure that we do it in a by design and by default manner for the future as well.

Richard Grisenthwaite: I would observe that security is an arms race. Every time you deal with any problem, the attackers get more sophisticated. We have seen that progressively at Arm, with many forms of attack. It is very important to also recognise that there is no one magic solution to security.

The hygiene factors that John talked about at the start are very prevalent. The National Cyber Security Centre put out strong principles for just building in the most basic security. We put out something called the platform security architecture, which has multiple levels that allow people to not have to reinvent the wheel with security: "Here is state-of-the-art best practice for that basic hygiene."

The by design principles that we are talking about build on top of that, because there is no point in worrying about these very sophisticated attacks if you do not have that basic foundation right—for example, if somebody is not even using passwords or is passing stuff in the clear. There is a tremendous amount of building on top of some very basic stuff that needs to be done.

The industry is getting better. The critical national infrastructure is getting better. What we are trying to do with programmes like digital security by design is make sure that, when those basics are done, the next steps are also ready for the technology to be deployed.

Q62 **Graham Stringer:** Do we have a definition of "national infrastructure" that we are using in the security world? Is there co-operation between the different parts of national infrastructure?

Richard Grisenthwaite: There are the agencies that regulate our critical infrastructures and their services. You have to look primarily at the service that they are delivering in terms of that definition. They co-ordinate through the NCSC. They have a programme called the CAF—cyber assessment framework—which provides what I would say is a description of what security should look like, and they are assessed against that description.

Something that the Committee could consider is whether those descriptions from Government can move into more prescriptive, "What you should be doing," as opposed to, "Have you done this?" That could be, for example, "Have you ensured that there is no cleartext or unauthenticated access to your systems?" as opposed to asking, "What have you done about access to your systems?"

It is very much about moving the language that we see in cyber security from, "Have you understood your risks? Have you understood what your monitoring mechanisms are?" through to, "Have you done something to protect yourself in the first place?" That is what is missing in a lot of the language that we use, in both legislation and recommendation, in the frameworks that the honourable MP was describing there.

Something that we can do proactively to help against the arms race, as Richard introduced it, is to ask, "What have you done? What can you do? Have you addressed this issue?" That can help with those frameworks.

Richard Grisenthwaite: With the greatest respect, the question was, "What is critical national infrastructure?" It is something that I have discussed with the National Cyber Security Centre. I would not say that they have been vague, but they have said that the exact scope of it is growing. More and more of our lives are dependent on technology. Are things like the mobile phone network part of the national infrastructure? To what extent is having a reliable internet connection in your home part of your own infrastructure? The way in which we now interact with the entirety of society is based on this.

One of the things that the NCSC was telling me is, "Do not think just in terms of those critical switches in power stations or stuff like that. That is the easy stuff." Increasingly, we have to recognise that the entirety of society is interacting digitally and online. Everything is becoming part of the national infrastructure; otherwise, we will be back in the 1840s and communicating by waving at each other. That means that we should not think that these technologies for security apply only to a particular small niche. Increasingly, we need security and secure design principles in all our technologies.

Q63 **Graham Stringer:** That must be right, because, if you ask, "What is not part of the infrastructure?" you suddenly find that it is. That takes me to my final question. Security covers everything, basically. You can have excellent cyber security, but, if something else in the organisation falls down, you do not have security.

Professor Goodacre, you are from Manchester University, which I wish well as a Manchester MP. Is there an overall consideration of security at Manchester University? You might make your computers secure, but one of the bad actors that we are talking about here is China. If China withdrew all its students and funding from Manchester University, you would be in a terrible mess, would you not? Are those factors considered as well?

Professor Goodacre: I would not want to comment on Manchester University's student employment preferences on that.

Q64 **Graham Stringer:** It is not employment. It is them being students paid for by China.

Professor Goodacre: Clearly, within the UKRI funding regime, there is a trusted research agenda in terms of recommendations and guidelines on what our academic communities should be doing. I can forward the references to that if necessary. I would not want to comment further on the choice of students at the university.

Q65 **Graham Stringer:** You are not going to say anything. It is not really about the choice. It is about the vulnerability to funding being withdrawn.

Professor Goodacre: Pulling that back into the by design, by default mechanisms, you could view that the systems that are being run should be run in a way that are secured to be by default, not dependent on any external influence, whether that is student numbers within the university or the number of connections to computers in the university. The systems themselves should be resilient and integral to the operation that they are intended for, without those external influences.

Graham Stringer: Thank you.

Q66 **Dawn Butler:** Professor Goodacre, you said that the Government could pivot from developers saying, "You should do something," to asking, "Have you done it?" in regard to secure by design. How can we tell whether a developer has considered secure by design?

Professor Goodacre: This is a very good question, because it is very difficult to know if they have done what they say they have done. You have a certification for your PSA, haven't you, Richard?

Richard Grisenthwaite: Yes.

Professor Goodacre: Richard can cover that later. It is the idea that you can put trigger points within an implementation to say, "Let us test it then." In that regard, you can see it.

The other thing is that it can be down to what tools or components they have used within the system. If it is by design, those components will have delivered that level of security. For example, if we were to say, "Please give me a new energy meter," and you said, "Which chip does it have in it? Does that have an implementation approach that is PSA-certified with by design characteristics, or is it an off-the-shelf chip that does not?"

Clearly, as the vendor, they just have to tell you what chip it is and you can look it up in a catalogue to know whether it has those design features within it, so it would be possible to have an audit level there. Whether they turn features on or off is a configurational issue.

Just coming back to the joint of those two questions, security underpins all the technology priority areas. In Government language, we talk about the five priority areas, so AI, semiconductors, biomedical, telecoms and so on. They all require this level of security and it is, in many cases, common across those as well. We should see about the language

becoming a consistent underpinning, if you like, of security across the critical aspects, so the network and the infrastructures themselves, but also the technology priority areas.

Richard Grisenthwaite: If I could add, I have regular contacts with the NCSC. One thing that we have been discussing there is about having clearer provenance of the materials inside your chips. Many people think that a chip is a nice small thing, but, in reality, there are literally billions of transistors and component blocks being brought in from multiple places. In some cases, they are being brought in from an open source community, where the provenance may not be clear and there are some concerns that, for example, people might bring in a piece of a design for a chip that has been made in China without anyone knowing what is in it.

One very important area, particularly in the CNI space, is to be able to have a clear provenance of where all the individual components inside a chip come from. You can look up a catalogue and say, "This has been built by a reputable US or UK company," but that might have been assembled from a whole bunch of different collections of intellectual property that have been brought together. Security always depends on the weakest link, so, if there is one little part of that that is wrong, it could compromise the overall security.

That provenance story for critical national infrastructure is a really important part of it. For the last couple of years, Arm has been talking with the NCSC about creating a scheme to try to make that clearer. People talk in software about a bill of materials. Nowadays, chips are so complicated and made of so many components that there needs to be a bill of materials for what is inside your chips as well.

Q67 **Dawn Butler:** How far away are we from systematically making sure that that happens, so that we can trace the provenance of each part?

Richard Grisenthwaite: There is still work to be done.

Dawn Butler: That is very diplomatic.

Richard Grisenthwaite: The first step in any journey is to understand what the potential problems are. Again, I am very aware that, when I say to people that I design chips, they ask, "What, those tiny little things? That cannot be very difficult, can it?" It takes 300 people a year to design a modern processor. There is an awful lot of complexity. It takes tens to hundreds of millions of pounds-worth of activity to create such a thing.

Because this is not particularly well recognised, and because they are being assembled from components from other places, it is about making people aware of the scope for a problem. Arm has been doing a lot of work with the NCSC and with agencies in the US to raise some of the concerns here. We are starting to see Governments go, "Yes, we need to put in place standards as to how you report this stuff." We can write down exactly what we have done, but you need standards so that people can say, "I bought this. I want to compare it with that." That

standardisation process requires standards bodies such as NIST in the US and the NCSC in the UK to push those standards through. Those conversations are happening, but we are not yet at that point.

Q68 **Dawn Butler:** Professor Goodacre, what is the digital security by design challenge fund, and what are its goals?

Professor Goodacre: This was something that Richard and I both kicked off in 2018 when I worked at Arm. Richard had identified some fairly good research at Cambridge University—this is where we are seeing the UK's excellence in academic research coming through—that could change the way that one of these chips runs the software, so that it can protect the user against memory safety errors. As I am sure you have seen in the press recently, there was a supply chain attack that could have done some really bad things to a lot of systems. It also has the technology that could have addressed that.

Q69 **Dawn Butler:** Tell us a bit more about that.

Professor Goodacre: The CHERI technology that I mentioned earlier is a research technology that Arm took to prototype with a research vehicle called Morello, so that we could take it out to the broader community—those hundreds of boards that we have been sharing. We are finding that that technology, as described, will block those memory safety issues. These are the things that you are probably seeing a lot of reports coming out of America on, saying, "We cannot accept any more memory safety issues." This is something that we can clearly allow our Government to do as well. Is it socially acceptable still to allow that 70% of vulnerabilities, or allow ransomware into our systems, if we know now that you can address that?

Another aspect of that technology is that it can put aspects of the design into what we call isolation or boxes, so that it cannot get out. Even if the code in that box has a supply chain vulnerability in it, it cannot affect the rest of your system. It is a very technical aspect of doing it. As an integrator or designer of a system, you can have more confidence. You are de-risking the cyber security of your product by using what is called compartmentalisation.

It is a long word, but, basically, it is putting bits of code in boxes where you are not, or do not want to be, 100% sure whether the provenance of that code has had a supply chain attack. You put it in a box and talk to it very precisely. It cannot then go and read your PIN codes and passwords, which is what happens today. If you have just one vulnerability in your application, it has access to all your data. This is a mechanism in the technology that stops that.

This programme was funded by the then industrial strategy. The key aspect to it was that it allowed a challenge to exist. Innovate UK ran about 18 of them. These were major investments of probably £50 million to £200 million in size.

The DSbD one was a £70-million investment from the Government. That was backed at the time by £117 million of industry money. We have already recognised £144 million and are on track to recognising around £220 million of industrial co-investment against the Government's initial £70 million. In essence, it is a very interesting programme to many companies, and I am sure that we will get into why it is not taking off commercially at this point. The programme worked with Arm through a project to create this prototype.

We then ensured that we broke the market failure. It takes years to build new hardware. Software churns every six months. If software does not have a piece of hardware, they ain't going to write the software. If the hardware does not have some software available, they ain't going to build the chip. The Government's challenge fund co-ordinated those two communities and made them large enough that we could then evaluate whether this technology addresses the questions that we think it does. I think we can say that it has. We have had some fairly large demonstrators from various industries. We have 34 projects across the community, but another 40 SMEs have been involved there as well.

We have built a skills group from half a dozen scientists up to probably around 500 skilled engineers, who, critically, know how by design technologies work. It is a significant uplift in what we now have in the community in the UK, and that is now having international effects on the demand. That is what the programme has been.

At the moment, we are looking at the end of the ISCF funding challenge in March 2025. With the uncertainty of the spending review and things, we have a challenge to make sure that we can still support it through to the adoption aspects that everybody agrees need to happen. It is going to require recommendations from committees such as yours that say, "We have to understand that this is an underpinning across our critical priority sectors that needs to be addressed in its own manner." In that way, all our digital infrastructures can be secured, not just, "How do we make an AI secure? How do we make telecoms secure?" This is beneath that and is across all the underpinnings.

Basically, anything that you consider digital today has a computer chip in it, which may be vulnerable to the exploitation of mistakes that programmers make in software. I can tell you that there are no tools or humans that are perfect. They are there, so let us contain those errors. It is not a magic bullet, but it is a significant part of the risk and de-risking of cyber security.

Dawn Butler: That is very comprehensive. Thank you both very much.

Q70 **Chair:** To follow that up in terms of our recommendations—I should declare that I set up the industrial strategy challenge to which this pertains—explain why it requires public funding. Why is it not, in a very commercial world, achievable through private investment?

Professor Goodacre: The DSbD challenge was proposed, basically, to overcome a market failure, where no single industry had either the co-ordinating or financial interest, or the ability to address that market failure. In that regard, it was Arm, Microsoft, Google and Amazon.

Richard Grisenthwaite: Those were the public ones.

Professor Goodacre: They were the public ones that we had in there, supported by NCSC and so on. Basically, they were saying, "Government, help us here. It is a big problem."

Q71 **Chair:** Does that market failure, as you put it, persist?

Professor Goodacre: The one that we addressed, which was the hardware/software market failure, has been overcome. The one that we are now looking at and struggling with is the commercial viability of the supply chain to adopt this technology. Arm sold 30 billion processors last year.

Richard Grisenthwaite: It was something of that order.

Professor Goodacre: That is the breadth of the digital ecosystem. Which company has an influence on 30 billion? There are probably two or three in the world that have a significant chunk of that. In essence, if you want to get the 30 billion, it is a market failure, because there is not the commercial return on investment and the supply chain to allow that trigger to cause it. Whether it is policy, regulation or social responsibility, or whatever that trigger or financial incentive, you have to look at the market as a whole. It is not a telecoms problem. It is not a biomed problem. It is not an AI problem. It is a problem to all of them.

Q72 **Chair:** Rather than through investing public money, as we did in that challenge fund, could you do it through regulation requiring this if you are to participate in the UK market?

Professor Goodacre: You would still need a co-ordinating function. Clearly, there are still aspects of innovation required. It is not a done deal at this stage. It is a concept that has been proven to answer the demands. It is not something that is commercially ready so you could just say, "Please buy it."

Q73 **Chair:** Mr Grisenthwaite, could you comment briefly on that from a commercial point of view?

Richard Grisenthwaite: From a commercial point of view, the reality is that, as I have alluded to, the costs of designing a new processor are very substantial. Because of that, we, as a commercial company, have to direct our resources to the most commercially lucrative opportunities. That is the nature of a business.

The challenge is that, to a certain extent, security is something that people see as a hygiene factor, not as a reason to buy something new. Certainly in consumer devices, I will predict that nobody around here has

bought a new phone because they thought, "Mine is not secure enough, so I must get a new one." You may well have bought yourself a new phone because you want higher performance or a better camera, or things like that, but security is a very difficult thing to go and sell.

We are working on how the major software ecosystems—things such as the Google Android operating system in your phones, or the Linux operating system in big servers—can really make this such an attractive feature that people will say, "I want to have this device. It will be worth spending some more on it." Adding these new features will cost more money because they take more transistors. That is just the way that it works. It is more area on the silicon. There needs to be a reason for people to say, "I will pay extra for this feature." The market failure is that having something that is secure by design is not necessarily seen as particularly attractive for a lot of commercial devices.

Professor Goodacre: If we pull that specifically into the critical national infrastructure scenario, you could legislate, "You must use it," but they would not be able to buy it. That is the problem that has to be overcome.

Q74 **Chair:** Mr Grisenthwaite, just to follow up on the exchange that you had with Dawn Butler on the provenance of components, you described the complexity and the sheer number of different components. Is it possible at this stage to exclude, for example, components from China for most of these modern applications?

Richard Grisenthwaite: At the moment, it is a design decision by the companies. If somebody is building a chip for, let us say, a telecoms application, they will design some parts themselves and take off-the-shelf components from various sources. At the moment, there is no particular guidance or requirement that the source of those is reported. Because all of these are companies, people will tend to access materials from the cheapest possible source. That is just how business works.

As a result, you are left with people pulling together these components in order to make, say, an advanced telecoms chip—to take the example that I am using. At the end of the day, that will get sold under the name of whichever company has made that chip, and people will not necessarily know where those things came from. There is no obligation whatever to report that. There is no knowledge about where things come from. Essentially, this stuff can almost get hidden.

The fundamental premise is, "Should you report where you get your IP from, so that at least people can see what is happening?" That will be a start to it. To a certain extent, you can have certification schemes on top of that.

Q75 **Chair:** That is not in place at the moment.

Richard Grisenthwaite: None of that exists now.

Q76 **Chair:** This would be a huge undertaking. It is the equivalent in

agriculture of being able to trace a piece of meat back to the animal on the farm where it was reared. That has taken quite a long time in agriculture, and it is fairly simple. In this case, it is a vast undertaking.

Richard Grisenthwaite: It is tractable, because you can identify various principles as to which things are more or less security vulnerable in terms of the various components, and make a start in that direction. There is a technology called firewalling, where you can, essentially, segment off some of the less trustworthy materials and provide principles. We have put in some suggestions to the NCSC about how you could start a scheme in this direction. It may take a lot of time to do it, but that should not be a reason not to do it. It is just an exercise that will take time.

Q77 **Chair:** Is the way that you described of identifying some components that are more security vulnerable robust enough? Could you not have a small, non-critical component that could be embedded with capability to go beyond what its purpose is and interact with all the other components?

Richard Grisenthwaite: If you surround that with technologies that limit what it is able to do, you can constrain its ability to cause harm. That compartmentalisation is something that hardware can do relatively straightforwardly, much as you can do in software with the technologies that John was talking about.

There are techniques and there are principles that you can write down to say, "These are the areas that have to be from a known provenance. These are the areas that, if you firewall them off using this sort of technique, could be from a less trustworthy source." The fundamental first step is to document what you have done, so that people can then audit it—so that if you are a Government agency saying, "I want to buy something for my secure communication system in the MOD," I know exactly what is in that. At the moment, the lack of transparency is a significant challenge.

Chair: To be clear, this is about the future at the moment. This Committee had an inquiry into 5G, and the Government, as you know, took a decision to exclude Huawei from that.

Richard Grisenthwaite: I know that only too well.

Q78 **Chair:** Other countries such as Australia have been more comprehensive in, as they claim and say, excluding Huawei, but the point is that, without that knowledge of the provenance at the moment, it is impossible to say whether there is componentry somewhere that is from this source.

Richard Grisenthwaite: There is a risk there. I would not want to overstate it, but it is a risk that is only going to get worse. Our geopolitical times are getting more complicated. The opportunities are there. We are seeing people submarining technologies into open source software, which is something that Professor Goodacre alluded to a little

earlier, and there have been some reports of that. The same concepts could exist in hardware.

We are trying to think ahead to how we can start protecting against this. This ties in with what I said earlier, which is that there is an awful lot of poor hygiene in security and we have to think of this whole thing as a set of principles built on top of each other. A lot of people are fixing those basics. We are thinking about what is next and what the next possible attacks are.

Q79 **Graham Stringer:** We have decided to be wary enough to exclude Huawei, and to be wary about Volvo cars or electric cars made in China, but I have no idea of the market in chips. Are the Chinese feeding chips into an overall market, so that you do not know what they are? In my ignorance, I thought that people would not be buying Chinese chips, but can you tell me what percentage of the market they comprise?

Richard Grisenthwaite: We can give a written response as to our understanding of the market, because I do not have that data today.

Graham Stringer: That would be very interesting.

Richard Grisenthwaite: At a more general level, it is a very globalised industry and there are an awful lot of different types of chips. There are tiny power controllers. There are chips that do the big processing. What I am highlighting is that the bit that does the big processing—the thing that makes your phone do all the clever stuff that it does—is a combination of lots of what we call IP blocks. They are little subcomponents of design. The basis of Arm's business is selling these, but other people sell them. You can get them through open source as well. The provenance of not even chips but those parts of chips is what I am concerned about.

The reality is that an awful lot of semiconductors and chips are made in China. That is part of the global supply chain that exists and part of the reason why you are seeing things like the CHIPS Act in the US and in Europe, which is looking to bring more capability into more friendly nations. That is the basis of the UK strategy.

I sit on the semiconductor advisory panel and, as part of that, we are thinking in terms of working with like-minded nations to ensure better supply chains, given that the global system is breaking down because of the change in geopolitics.

Q80 **Dawn Butler:** We have taken evidence that we were about 20 years behind China in the development of chips and microprocessors. I know that you have over 100 patents on microprocessors. I just wondered how close we are to catching up.

Richard Grisenthwaite: I am not sure that I recognise that term "20 years". It depends on what exactly you are looking at. In terms of the designs of microprocessors, which are the brains of a phone or a server,

Arm's implementations are being demonstrated as being absolutely leading edge. The manufacturing side is different. The UK does not have its own manufacturing base, and I am not necessarily sure that it needs to, but we do need to make sure that we have access to manufacturing in trustworthy countries.

Q81 **Dawn Butler:** I thought that the Five Eyes were working together to have a manufacturing base.

Richard Grisenthwaite: That is the direction of travel.

Q82 **Dawn Butler:** So we are not there yet.

Richard Grisenthwaite: The leading-edge chips, the so-called 2 nanometre or 3 nanometre chips, which are the smallest geometries and the most advanced chips, are built in three places in the world—some in America, some in South Korea, and the vast majority in Taiwan. Taiwan presents a strategic risk going forward because of the geopolitical tensions there.

Professor Goodacre: That is the printer. The artist is in the UK. We draw the pictures. They print them. Clearly, they can change the printing, and that is where Richard was talking about knowing the provenance of the components. In essence, they have the leading printer, but we have the leading artists.

Chair: That is a good analogy.

Q83 **Stephen Metcalfe:** Just to clarify that last point about Taiwan, Taiwan is, at the moment, a friendly nation.

Richard Grisenthwaite: It is, but there is risk there.

Q84 **Stephen Metcalfe:** It is the risk of where it might be in the future, which could disrupt the supply chain.

A lot of what I wanted to discuss has already been answered. The first question was really around the challenge fund, which you have discussed quite clearly, but it does sound as though, when the challenge fund was set up, there was a clear path that you wanted to follow. Was that a constraint, or was it agreed among the wider community: "This is the challenge. This is what we want to explore. Please go away and do that"?

Professor Goodacre: The challenge bid itself asked for a bit more money than it received, so that we could look more at the adoption side. As a group of stakeholders and an advisory group that included a lot of Government players as well, we prioritised the essential bits that must be fixed, without which adoption would not work. That was what we addressed.

We focused on breaking the market failure between availability of hardware and software to evaluate whether the technology answers the problems. We then, in the vision, if you like, said, "We will test it in two or three sectors." We have actually tested it in six. We have done VPN in

telecoms. We have done two automotives. We have done e-commerce through a very large UK company called the Hut Group, which has been looking at how secure e-commerce is on the internet. There are a couple of others that I cannot remember. They then said, "Yes, this is finding things that we did not anticipate at the beginning."

We also constrained the funds to make sure that we did that really well on the Morello scenario. More recently, we have extended that to include the RISC-V architecture as well. The technology that we are describing is patent-free, so any architecture. There are, basically, three in the world that you should care about—the x86 one, the Arm architecture, and the RISC-V one. The programme was able to extend to include RISC-V in terms of, "What does this look like there and can we do anything on that as well?"

Q85 **Stephen Metcalfe:** Thank you very much. That is very helpful. I think the subtext of that was, if you had had a bit more money, you might have been able to do a bit more, or the scope might have been a bit wider, but it did not hamper you.

Professor Goodacre: It did not hamper us proving that the technology works. Thanks to Arm for working with us on that. We are the co-ordinators of this, not the doers.

Richard Grisenthwaite: I would stress that the Morello chips and boards that are out there are a resource that researchers can and will carry on using. Funding researchers in order to do that, particularly in universities, would be a good thing, because it is not like, after five years, everything is solved. There are more software investigations that can be done.

We are seeing growing interest from within the US Government to look at it. I have a meeting coming up with the ONCD, which is part of the White House, looking at national cyber issues. They are asking about access to the Morello system, simply so that they can understand the values of this technology and what they could do. This is a completely global industry. There is no point in Britain trying to secure its bit and not having the US involved. A lot of the software on your phones and systems comes from the US as well as the UK.

Professor Goodacre: The challenge is not so much what we did with the funding that we had in that period, but more the long-term commitment and now extending that to move beyond what we have. How do we influence DSIT, in this case, to prioritise within UKRI, for example, the focusing of security principles across the technology areas for the next phase? That is now the challenge, so that we can move it beyond, "Okay, it works," and get it into a critical infrastructure.

Q86 **Stephen Metcalfe:** Do you want to put a figure on that?

Professor Goodacre: What I can say is that the current programme was £70 million. If we take the example of energy resilience of the new EV

meters or the constraint management of things that they are going through, led by DESNZ, and the implementation of the security around that, why do we not pre-product procure an end-to-end scenario for that? That is probably in the high tens. If you want to be able to expand that to multiple markets, it is probably also in the high tens. Around the £100 million level would not be unusable.

Q87 **Stephen Metcalfe:** I think Arm supplied 30 billion chips last year.

Richard Grisenthwaite: We do not provide any chips. We sell designs that other companies use.

Stephen Metcalfe: You created 30 billion chips.

Richard Grisenthwaite: I do not have last year's numbers. Over our lifetime, we have shipped just under 300 billion chips, and the majority of that has been in the last few years, so that number is not inconsistent with that.

Q88 **Stephen Metcalfe:** Just so that we can all understand the scale of the chip market, is there a figure for how many chips are produced globally?

Richard Grisenthwaite: We have already undertaken to make a certain amount of written responses. You have to be very careful about chips. There are memory chips; there are lots of different types. If we talk about processing chips for high-end mobile phones and so on, we know that several billion of those are made every year, and so that is several billion processing chips. That is one of the biggest markets for high-end processing. Servers are a much smaller market, but they have more chips in them.

That 300 billion number is also inflated by the fact that there are lots of very small microcontrollers in things that control your electric windows in your car, or your windscreen wipers, so it is an incredibly wide market. McKinsey has a figure at the moment indicating that it is about a half a trillion dollar industry overall. I expect it to go up to about a trillion by 2030. That is the semiconductor industry as a whole. That probably gives you an idea of the scope of it.

Q89 **Stephen Metcalfe:** When there was a disruption following covid to global supply chains, you could not get a car for love nor money and sales of second-hand cars went through the roof; part of that was due to the lack of supply of chips. I was quite surprised to be informed at the time that an average high-end car would use 5,000 chips, and some would use 10,000.

Richard Grisenthwaite: Those numbers are going up.

Stephen Metcalfe: Hence the doubling by 2030.

Richard Grisenthwaite: Yes, and the complexity is also going up, because there is clearly a lot of interest in things like AI at the moment. Cars are starting to be able to steer and drive themselves. We are

providing technologies that do all that processing. They take a lot of work and chips.

Stephen Metcalfe: As the cars start to drive themselves, we need to make sure that they are secure—beyond that, yes please.

Q90 **Rebecca Long Bailey:** Mr Grisenthwaite, can you tell us a bit more about the Morello board approach and how it makes systems more resilient to cyber attacks?

Richard Grisenthwaite: The Morello board is a board that contains a chip that has been augmented with technologies based on this fundamental research done by Cambridge University. They call that CHERI, which stands for capability hardware enhanced RISC instructions. Alongside each reference to memory, it carries some additional information to, essentially, restrict it to being able to refer only to the things that you want it to refer to.

The memory safety issues that are at the heart of some 70% of vulnerabilities really come into classes such as when I have a set of data that I want to pick something out of—for example, look up your name in a database—malicious programmers find ways of exploiting that so it goes out of range and picks up some other piece of data that it is not meant to get hold of. That can then be used, if you are not only reading data but writing it, to corrupt the way that it is executing, to take over your computer and to do bad things. That is the basis of these attacks.

By carrying constraints alongside references to memory in a single package, you are then able to, essentially, defeat that class of attack. That then becomes a building block to deal with a great many of the memory safety issues that are there.

In addition, because of that very tight containment of information, you are then able to use this as a way to do a better compartmentalisation of your software, which is what John was talking about earlier. Much as I was talking about with hardware before, software is made up of lots of different libraries of material that come from lots of different people. People will write something to handle your pictures or your email, or whatever. If something goes wrong in the thing that is displaying your pictures, you do not want that to be able to access your email.

By having stronger compartment boundaries, which is what this technology allows you to do, it then means that, even if there are exploits that are found, there is less damage and less ability to take over the whole device. A single exploit being the basis by which you can take over the whole device is, unfortunately, the situation that we are in today, so better compartmentalisation is very important.

Professor Goodacre: Richard has very nicely articulated the technical merits of Morello. You should also understand that it is a demonstrational platform. It is not something that anybody can go and buy. It is not something that a consumer can use to protect themselves in the ways

that Richard described. That is the status. We have used the demonstrational platform with all these great features to prove to and show people that these benefits can be realised.

The challenge that we all have now is how we get together, through regulations, policy, recommendations, guidelines, frameworks, interventions or procurement, to de-risk as much of the cyber security risk of the UK and, in some cases, the national security. How do we get the commercial answers to be put forward in terms of taking Morello into product?

Richard Grisenthwaite: I can explain why that challenge exists. Essentially, we provide a building block. There are lots of software engineers who spend their lives writing software. They need to be able to use these new building blocks and tools to create that more secure world that John is talking about. To a certain extent, it takes time for software engineers to work out the right way of using these building block tools to come up with a compelling story that says, "This will really still give everything that people want from their phone or from their server"—or whatever else—"and can be more secure." That is the step that we are still going through.

There are some great demonstrations being pulled together by Cambridge University and by us that show this, but it is taking more time to port software, because software is so complicated, in order to make use of this technology.

Q91 **Rebecca Long Bailey:** Professor Goodacre, how does the potential uncertainty of Innovate UK's digital security by design challenge funds post 2025 impact long-term planning, both for investment strategies and research, and in improving cyber security for critical national infrastructure?

Professor Goodacre: The DSbD programme through Innovate UK has been a co-ordinating function. In other words, we are the ones who pooled the software, the hardware and the researchers. As part of the industrial strategy fund, we also included social scientists and economic people, so it has been a very cross-disciplinary approach to that.

It is not clear how those approaches could pick up the ecosystem, those 500 skilled staff, and the trajectory and maturity of that software. How we sustain that over the spending review is a very pertinent question, as is how or whether we can execute on the adoption and diffusion of the technology beyond that. That is the question that I am sure you are all very familiar with for the latter half of this year and the start of next year.

Q92 **Chair:** Professor Goodacre, you are a professor at the University of Manchester. Not far down the road from you in Cheshire is the Hartree facility run by the Science and Technology Facilities Council. There has been an announcement that a new Lenovo, Chinese-made supercomputer

is to be installed there to be used by, among other people, Rolls-Royce and the UK Atomic Energy Authority.

Ciaran Martin, former head of the National Cyber Security Centre, has given evidence to this inquiry and was quoted as saying that this raises important strategic and national security questions. What is your instinct about your neighbouring facility?

Professor Goodacre: Unfortunately, I do not think that I am equipped to answer. I am not an expert in the selection of HPC systems. I have worked with Hartree previously. In fact, they have one of my prototype exascale supercomputers sat in their car park that I was using as part of my research as well. I know that they work diligently towards what is acceptable for their use and market, but, apart from that, I cannot comment on the system. It is using standard components in processing capability terms, because there are only a limited number of vendors that provide those components.

Q93 **Chair:** Mr Grisenthwaite, we have had a conversation about provenance. If it is a computer that is made in China, we can be pretty clear about its provenance. What is the point of doing this great exercise in finding out where things come from if such procurement happens anyway?

Richard Grisenthwaite: I cannot talk about procurement policies or what due diligence was done on that. In terms of why provenance matters, we need to be aware of where things come from so that people can make informed choices and decisions. Ignorance and lack of transparency is the point where you wake up and discover, "My goodness, things are much worse than we realised." That transparency is desperately important if we are to have the ability to do a proper audit of our systems going forward.

Q94 **Chair:** Thank you. You have made that point very clearly and compellingly. Thank you both for your evidence this morning. I am now going to invite the next pair of witnesses to join us at the table. Thank you very much indeed.

Examination of witnesses

Witnesses: Henry Harrison and Dr Vasilios Mavroudis.

Q95 **Chair:** As they take their seats, I will introduce them. I am pleased to welcome Henry Harrison, co-founder and chief scientist of Garrison Technology Ltd, which is a British cyber security company specialising in translating security solutions developed for national security purposes to commercial applications. It will be very interesting to hear from you, Mr Harrison. Dr Vasilios Mavroudis is the principal research scientist and co-theme lead for AI for cyber defence at the Alan Turing Institute. Thank you both very much indeed for joining us.

Perhaps I could start with a question to Dr Mavroudis. Tell us about how

artificial intelligence is being used both defensively and offensively in cyber security. While that is relevant to this inquiry, we also have a parallel inquiry into AI, so we would be very interested in your thoughts on that.

Dr Mavroudis: Good morning and thank you for having me today. Machine learning has been traditionally used in cyber security in the context of not only critical national infrastructure but protecting retail and commercial systems that are not part of this infrastructure. The security industry has typically been quite conservative with regard to how complex the machine learning models that it employs are. This is for a good reason. We need those systems to be explainable and easy to audit and understand. As we move towards more complicated models, which are those that we see in newspaper headlines, all those processes become harder.

On the defence side, we have been using machine learning for anomaly detection, intrusion detection and malicious software detection in systems. This is usually in combination with even more basic systems—for example, detecting malicious signatures or behaviours.

Besides that, there is the more offensive side, which is a concern. Phishing attacks, which are social engineering attacks, usually on high-value individuals, have been on the rise. It seems that the latest advancements with language models will propel this further. These are two use cases, let us say, for defence and offence.

There is a third area, which is more on the research side, but it seems to be gaining traction lately. It has to do with automated vulnerability discovery and patching, which is, essentially, fixing. This has not yet made it to commercial products to the extent that we would expect to see, but the latest research work suggests that, in the next few years, we will have something in that area.

Q96 **Stephen Metcalfe:** Someone on our previous panel mentioned that, effectively, this is an arms race. If you take one step, bad actors will try to combat that. Part of the problem, though, is that the pace of change is ever increasing as well, and we have lots of emerging technologies that are now taking hold. Do you feel that the operators of our critical national infrastructure understand the cyber security risks that these new technologies present?

Henry Harrison: I do not seek to represent the entirety of the UK's critical infrastructure sector, but, in looking at this, there is a useful analogy that I have relied on through my career in security, which is to think of the problem of you and your friend out for a walk being chased by a bear. You have two strategies available to you. One is to try to outrun the bear, and the other is to try to outrun your friend. Those two strategies are, indeed, employed in different contexts, but a large part of the mainstream commercial world is really focused on trying to outrun the friend rather than the bear.

My world spans both of these, because I deal with commercial companies as well as the national security sector. If our agencies that are trying to protect our secret and top secret information took the view that they are only trying to outrun their friend, we, as citizens, would not be terribly pleased with them. They clearly are trying to outrun the bear rather than the friend. I may have drawn this metaphor to its final conclusion here.

Does everybody understand everything? No, but then not everybody is trying to understand everything. If you wanted to ask the question about the real understanding—who has that information and who really understands—you need to look to the agencies in the UK and in our allied partner nations in terms of the deepest understanding of what works and what does not work in these areas.

Q97 **Stephen Metcalfe:** As it stands, do you think that there is enough understanding of the risk?

Henry Harrison: The risk is phrased very differently. I had a conversation with the chairman of a major financial organisation sometime back, who said that their ambition in terms of cyber security risk management was to be in the upper quartile. In other words, that is very clearly about trying to avoid being the laggards that are caught out by people going after the weakest. That is very different from those who say, "No, we have very clearly defined adversaries—foreign nation states—and we need to understand what it takes to defend against them when they are coming after us specifically."

Q98 **Stephen Metcalfe:** Do you agree with that entirely?

Dr Mavroudis: I do. I have something small to add, in that, through interactions and even interviews that we have conducted with CNI operators and people at NCSC who are close to national infrastructure, the overall stance that they have been keeping seems to be very healthy. They are curious. They are learning about what is coming out without necessarily jumping quickly into using the latest and greatest, which we often do not understand very well.

Q99 **Stephen Metcalfe:** Are they working adequately with their supply chain as well to ensure that the whole ecosystem is knowledgeable and working to reduce potential risk?

Dr Mavroudis: The supply chain risks for machine learning come from the hardware, which I consider relatively safe. The bigger risk comes with adopting pre-trained models from third parties that might not have the security that you might expect, or those models getting intercepted on the way to your premises by a state-level adversary, potentially, and embedding a hidden functionality in them. The problem is that those models, if we are talking about advanced AI models, are hard to audit, and so, once this functionality is embedded in them, it is usually hard to discover it.

Henry Harrison: If I might just add a little bit about the supply chain, the cyber security industry is a very large, very well-funded industry, with a lot of capital being constantly invested in new technologies that are being developed.

It is also an industry where the vast majority of buyers do not have the skills, the information or the time to assess whether those technologies are effective at protecting against cyber security risk. Consequently, I have read studies that suggest that the cyber security industry spends more on marketing as a percentage of revenue than any other industry on the planet.

The previous technical director of NCSC described many of the offerings in the cyber security technology market as no better than magic amulets. The current chief technology officer of the NCSC has described the market as broken.

As I have said, there is a very big difference between the mainstream market and the market where we are looking to protect the things that the UK has determined really matter, which have historically been classified secrets, and where very significant work goes on between the NCSC or its partners in the NSA or CSE and other allies to look at the technology and work with the suppliers in order to assess whether those are really secure. That does not apply to the vast majority of the security market.

Q100 **Stephen Metcalfe:** How do we address that? Do the Government need to invest in this? Is it the market?

Henry Harrison: Yes, they do, but the problem is that this is a very large global problem. You heard from your previous witnesses that there are a great many initiatives around secure by design. Speaking as the representative of a private sector company, we are under extreme commercial pressure. If nobody is marking your homework, there is an overwhelming incentive to cheat. It is very easy to stand up and say, "I do fantastic technology. I am secure by design. This will protect you against anything," but to evidence that is very difficult. Because this is a global market, the UK has to see this as something it has to work with its allies to address, rather than expecting to resolve the problem unilaterally.

Q101 **Stephen Metcalfe:** Thinking about that challenge further, the UK has a history of having the ability to create quite good standards, regulations and quality marks. Is that an area where we could, while working with our global partners, lead the way with them and find some sort of accreditation system?

Henry Harrison: There is activity under way. The area that I work in is cross domain solutions. We are a type of solution that, historically, is designed specifically for protecting things that really matter. The UK is the recognised world leader in this space. On the second day of last

year's two-day US Government conference on cross domain solutions, all but one of the speakers was British. The NCSC has been doing some fantastic work on looking at how homework marking can be done for that sector, but there is a great deal of work to do in terms of publicising and helping people to understand this, because, at the moment, the knowledge, while not classified, is very restricted in terms of the distribution that it has found.

Q102 **Stephen Metcalfe:** If you were to encourage us to make a recommendation to the Government about where they should spend some money, is that where you think it should go, or should it be spent on a greater understanding of the relationship between emerging technologies and cyber security?

Henry Harrison: I am a salesman, so you must take everything that I say, as with any vendor, on the understanding that I have an interest in this. We are in a situation where the UK is absolutely the recognised world leader in the area of how to protect against high-end nation state cyber attacks, and we are doing insufficient work to shout about that on the global stage and to advertise our world-leading-ness.

Dr Mavroudis: I want to bring attention to another area that is often overlooked and has to do with start-ups working on something that is not necessarily critical national infrastructure in the strictest definition. For example, we have start-ups in Oxford working on nuclear fusion reactors. I am not a nuclear physicist. From what I know, there is no risk of leaking radiation if something goes wrong, but this comes with a lot of innovation. It would be great for the nation if we had a breakthrough there and perhaps we need some support for those start-ups to protect their secret from, first, espionage and, secondly, potentially malicious intrusions to their systems for destruction of property, which would definitely set their research and processes back several years.

Q103 **Dawn Butler:** We often say that we are world leaders in everything in the UK. Mr Harrison, we heard an analogy earlier. Should we be investing more in not just the creation of the art but the printing of it?

Henry Harrison: That was a very good analogy at the time, but it is, of course, more complex than that. There are, in fact, multiple levels in this process. For example, my company manufactures hardware, so we have UK-based manufacturing and we ship physical units of a certain size to our customers in the UK and overseas, but we do not, of course, manufacture every element inside that box. Some of those are chips that we purchase from elsewhere. I cannot comment on the UK as a whole, but I can say that it is possible to build businesses in the UK that not only design but manufacture technology, and we do that.

Q104 **Dawn Butler:** That is important, is it not, if we are going to be confident about the provenance of all the components. How is artificial intelligence being created with secure by design as part of it? What I like about the Alan Turing Institute is that it always has an ethical approach at its core

whenever I read the briefings. Is artificial intelligence being delivered with a secure by design approach?

Dr Mavroudis: That is a fantastic question and a really difficult one to answer thoroughly. I will make my best effort, though. The problem with machine learning right now is that we are aiming for performance, so we are optimising for that, to get to the best model that can deliver the best performance. Alignment with humanity's goals, perhaps, or security is almost an afterthought.

However, in the last few years, because the community realised that we were lagging behind on that, there has been a significant effort to make it part of the training of the model to ensure that it matches certain other properties, one of them being security. Is it perfect? It is not, but we have seen substantial progress towards this goal and the hope is that we can reach something that is considered acceptable in the real world soon.

Another area that we have been making lots of progress in is explainability, which has to do more with reasoning as to why the model suggested or did something based on some data. The research in this area is lagging behind a little, particularly due to the complexity of the latest models.

Q105 **Dawn Butler:** Is there a global agreement on the standards of security? Because it is a global problem, is there a global agreement on what people would consider is secure at level 1 or level 2?

Dr Mavroudis: Although different countries have different standards, with regard to embedding a machine learning model into a system, we do not have to necessarily come up with a whole new set of rules for how we can safely do that. The existing best practices from the cyber security industry will suffice to a very large extent. However, in many cases, especially when it comes to using AI models to make decisions, cyber security is at odds with usability. There is some friction there.

Henry Harrison: You have identified a very large problem in the security industry, which is that some sort of calibration scale for security is lacking.

When we talk about security, I would like to see us differentiate between something that is secure enough to repel a teenager in their bedroom who is willing to spend an afternoon trying to get into it, and something that is secure against a foreign nation state willing to spend $100 million and a year trying to get into it. That would be a useful scale, to my mind, but historically nobody has really been willing to opine on that because of the risk of getting it wrong, and also because the sorts of people who are capable of making those sorts of technical evaluations are not typically well connected with the process of describing things in terms of money and time that is most useful to decision-makers.

Q106 **Dawn Butler:** That is an area that we really need to pursue as a committee. We will probably do more work on that. Should Government

and operators of critical national infrastructure use only artificial intelligence systems that are designed and deployed securely, to avoid harms to individuals and systems? Is that currently possible?

Henry Harrison: I do not really have any comment. I am not an artificial intelligence expert, and I cannot comment on that.

Dr Mavroudis: Can you expand a little bit?

Dawn Butler: Basically, should the national infrastructures be designed to avoid harms? Should that be embedded in the system? Is that currently possible?

Dr Mavroudis: Is that with regard to machine learning or AI, or in general?

Dawn Butler: It is with regard to AI.

Dr Mavroudis: Absolutely, and this is the end goal of all the efforts in machine learning. However, as I touched on earlier, it is usually pretty difficult if we are talking about advanced models to know exactly all the information encoded in the neurons of that model and the parameters.

To reduce harm, there are processes we can add on top of the model. For example, there is a standardised process called levels of automation, by which you could have an AI model as part of a more complex system. It defines how this model gets to influence decisions. For example, you do not have to go directly to autonomous decision making where the model decides everything on its own, but you could have a human in the loop, either approving the model's suggestions or receiving recommendations from the model without them being binding.

Q107 **Dawn Butler:** Do you think that having human in the loop is enough, or should it be more human-based? Is having a human in the loop enough to mitigate any risk?

Dr Mavroudis: For use cases where humans are sufficient, I would say that we do not have an immediate pressure to switch to using the model, except perhaps to reduce costs, but besides that, no. However, there are many cases where adding a machine learning component into an existing system would help us scale it much further and process a lot more data. In those cases, there is some pressure to move towards systems that are more autonomous than they are right now.

Q108 **Dawn Butler:** My last question is to Mr Harrison. Does the UK have all the talent that we need to continue to be world leaders?

Henry Harrison: Of course we do not have all the talent that we need. I run a company, so I always want more engineers who are better and cheaper, but we have been very successful over the past 10 years in hiring excellent talent into our engineering organisation. There have been ups and downs based on the swings in the market, but generally

speaking, I have no complaints about the UK as a skills base for what I have been doing in the last 10 years.

Q109 **Dawn Butler:** What is the industry doing to encourage and support the next generation of workers?

Henry Harrison: We are particularly interested in supporting the diversity of the next generation of engineers. We still have a very significant issue, particularly with gender diversity in the engineering base. This is a global issue, not uniquely a UK issue, but it is certainly one of the things that we as an organisation have tried to support in the small way that we can by encouraging females at school and university to consider careers in science and engineering.

Q110 **Dawn Butler:** What does that look like? How do you encourage them in schools?

Henry Harrison: I am afraid that I am probably not the right person to answer that question, but I am happy to arrange for us to provide some written evidence to the Committee on that.

Q111 **Chair:** Just to follow up on Dawn's point, there are programmes like the STEM ambassadors, for example, where people who are working in science, technology, engineering and maths go into schools and mentor people. Do you participate in that as a company?

Henry Harrison: We have participated in some of those things, but again, on the details, I would prefer to supply those as written evidence.

Q112 **Chair:** As someone who has worked and practised in the field, you have observed this. You have, as it were, grown up in it.

Henry Harrison: My role, above all, has been to support the female engineers within my organisation to allow them to go and support these types of things. I really would like to source the information from them rather than try to represent them.

Q113 **Chair:** I do not want to be difficult, but you are the chief scientist for your company and you expressed an aspiration to have more people coming through. What can you do in your own job as chief scientist to help and encourage that?

Henry Harrison: I am a middle-aged, white male. There is not a lot I can do about that. Consequently, my role, as I say, is above all to support others within my organisation who are better placed than I am to do that sort of outreach work.

Q114 **Chair:** Dr Mavroudis, you work in a related field, which is similarly not as diverse as it could be and therefore missing people of talent. What do you observe both in the organisations that you work in and in your personal capacity?

Dr Mavroudis: Turing has been very active in working towards improving EDI within our teams and processes. We currently have an

internship programme specifically targeted to address this problem for undergraduate students. We have been engaging with students in high school, essentially giving them some insight into what it means to be a researcher and to work in cyber security. Besides that, we do our best when hiring to advertise to places that will attract a diverse set of candidates.

However, I still feel that we come a little late into the chain of events to completely overturn some trends that are there in hiring as part of our team. For example, most of our researchers have a PhD. The percentages in terms of gender balance of the candidates with a PhD in cyber security are already quite skewed. It is hard for us to intervene at that point, although we are doing our best.

Another issue we are facing as a team, which we have been creative in addressing, is that hiring people at the intersection of machine learning and security is especially hard. As you can imagine, hiring someone right now who is an expert in AI is very competitive. Not only that, but we are looking for someone who has expertise on both this and cyber security. There are very few people in the UK and in the world who have this set of skills. Instead, we hire people with expertise in one of the two areas, typically machine learning, and we upskill them and train them in the other one, which is usually cyber security.

Q115 **Chair:** Do you personally go to visit universities and younger people, perhaps even in schools? Is that something you do?

Dr Mavroudis: Yes, typically Turing organises and makes it easy for us to engage with students. It is usually a day or two off for them in some nice environment where they can engage with many of us from different teams and different groups at different stages in our career, and get an insight into what that looks like.

Q116 **Dr Davies:** Mr Harrison, you have already touched on cross domain solutions. I know you have a particular interest in them. Could you explain in basic terms what they are and how they work?

Henry Harrison: Sure: 20 years ago, if you had computer systems that contained highly sensitive material—typically secret or above secret information—the way you protected them was simply not to connect them to the internet. We have not done a scientific study of this, but what is interesting is that, anecdotally, the majority of professionals in the cyber security industry believe that still to be the case today. It is, largely speaking, not the case today because it is not practical to try to run defence or intelligence matters in that way in the 21st century.

Consequently, in fact, today these sorts of highly sensitive systems are connected to the internet or to other dangerous systems, but they are connected via a class of technology called cross domain solutions. That is what the technology is designed to do. It is designed to allow connectivity

between highly sensitive systems and highly dangerous systems, while protecting the highly sensitive systems.

Q117 **Dr Davies:** It is a bit like a very powerful firewall and encryption.

Henry Harrison: Obviously, firewalls and encryption are also technologies that have a role in cyber security, but there is not necessarily very much in common between them and cross domain solutions. Interestingly, although very few people in the industry have heard of cross domain solutions, you can in fact go on to the NCSC's website, where it has published its principles for what constitutes building a good cross domain solution. Rather than try to represent the NCSC, I would refer anybody to see what it has to say about this, as the recognised world authority on it.

Q118 **Dr Davies:** Where are they currently deployed?

Henry Harrison: Typically, today where cross domain solutions are deployed is precisely for the use case I described, which is around the protection of systems containing classified information. The really interesting thing is that, 20 years ago, those were very clearly the most sensitive things that the country had—classified military and intelligence secrets. Today, we also worry very much about critical infrastructure, where the issue is not secrecy but the continued and correct operation of those things.

What we see is that there has been very little knowledge about the existence or the role that cross domain solutions could play among those who are discussing the security of critical infrastructure, as opposed to the protection of secrets.

Q119 **Dr Davies:** Yes, why is that the case? Why is there such poor awareness?

Henry Harrison: Again, if we go back 20 years, there was poor awareness because the information was classified, so there was meant to be poor awareness. That has changed recently. About three or four years ago, the NCSC took the decision to publish detailed guidance about cross domain solutions on its website. While we have seen bottom-up technical information about cross domain solutions, we have not seen any top-down communication that emphasises the importance of understanding about this. That is something that I would heartily commend the country to do.

Q120 **Dr Davies:** Dr Mavroudis, do you have any additional points to add to what I have raised?

Dr Mavroudis: I want to add something on the point that Mr Harrison made about secrecy. Secrecy is perhaps necessary in some cases, and no one is going to argue against that, but it is hurting research in another way. The machine learning community made huge advancements in the last few years, just by being very open with their datasets, their

challenges and their benchmarks, even putting things up on the internet and allowing individuals around the world to take a chance, try their approach, see if it works and share it with everyone.

We have the problem that it is becoming increasingly hard, or it is typically hard for us, to do the same because a lot of those things are classified. There is a way forward, though, which is identifying the underpinning problems and, after stripping away all the technical details that might be confidential, sharing with the world what the actual scientific challenge is there, so that we can attract more interest and more people working on those.

Q121 **Rebecca Long Bailey:** I have some very brief questions. Mr Harrison, we have briefly discussed cross domain solutions, but can you speak specifically to how they would improve the cyber security and resilience of critical national infrastructure?

Henry Harrison: I will divide my answer into two parts. First, the UK as a country has a great deal of intelligence about what works and what does not work in the context of cyber security. We know that both because of our agencies' intelligence gathering and because we, like all other countries, are engaged in the business of trying to break into other people's systems as well. Obviously, that is not something that we can talk about, or that I am authorised or qualified to talk about.

Clearly, the UK, through its appropriate authorities, is very well informed about what works and what does not work, and has chosen a methodology for protecting secret and above secret information, presumably on the basis that it believes that that is what is most effective. If that is what is most effective there, it seems equally likely to be the case that that would be most effective in the area of protecting critical infrastructure.

Whether we wish to be that effective in protecting our critical infrastructure is a decision that I cannot opine on, because that is a risk decision that we need to make. There are costs to be incurred and so on. None the less, it seems to me a logical inevitability that, if that is the best way of protecting one set of things, it would largely be the best for protecting another.

If you look at the regulatory framework in the UK, the sector that is the bellwether for UK cyber security regulation is the telecommunications industry, with the Telecommunications (Security) Act and the code of practice that has been published for telecommunications. If you read the code of practice carefully, you will see that it does indeed specifically refer to cross domain ways of working—I think that is the phrase used—and several patterns that come directly from the principles about cross domain.

However, among those who are working with and interpreting the code of practice from a player and a regulatory perspective, there is still a

general lack of understanding of the role that cross domain solutions and cross domain ways of working play within that classified space, because historically the people who have known that have been those who have been active in that area, rather than in the telecommunications industry. None the less, it signposts a direction of travel that is very encouraging.

Q122 **Rebecca Long Bailey:** My next question follows on from your response, Mr Harrison. How do you think the Government could improve the uptake and visibility of cross domain solutions across critical national infrastructure?

Henry Harrison: I have said before that the cyber security industry is very noisy, and it spends a lot of money on marketing. The net result of that, of course, is that cyber security vendors are not terribly well trusted, because all cyber security vendors will stand up and say, "We have the best thing since sliced bread. This is absolutely the thing you need to protect yourself."

There is, therefore, a critical role for Government in promoting what they believe to be the most effective method of protection. NCSC has done a good job of assembling some of those materials, but we have to be conscious that NCSC is not a marketing agency; it is not set up to be a marketing agency. If you are going to try to cut through the noise that is generated by the billions of dollars that are spent on marketing by predominantly US and Israeli cyber security companies, there is a role for marketing-like visibility within the UK.

What I do not know, because it is not my area of speciality, is how the UK Government would best organise themselves to do that kind of communication. It seems to me that a good starting point—or maybe not a starting point but something we should do—would be to have within the next iteration of the national cyber security strategy a reference, for example, to cross domain solutions.

In the United States, there has been a public statement by the White House on the importance of cross domain solutions in something called national security memorandum 8, which was published in January 2022. That called out the particular importance of cross domain solutions as a defensive technique within the US national security sector. Because we are different, we have not had to do a similar thing in the UK, but it means that there is something of a void at the top in terms of expressing the importance of this sector from a security perspective.

Q123 **Graham Stringer:** I am bowled over, really, by the expertise of the people we have seen this morning. This is a very general question. Many parts of science and technology are bedevilled by specialisms that cannot talk to each other. I possibly do not know enough about your world to even ask sensible questions, but you all appear to know everything about everything. I am sure that is not true. Is there a problem in specialisation and when things go wrong of finding the right people who really do understand the nuts and bolts? I could tell you how an internal

combustion engine works, but could I repair one now? No, I would not have the faintest idea.

Henry Harrison: I would hate to give the impression that I know everything about everything, because I really do not. For example, immediately next to me is something I know very little about, which is the artificial intelligence and machine learning work that Dr Mavroudis works on. I also know very little about what works in cyber security, because in order to know what works you need to have access to the intelligence about what has and what has not been successful, both when you are trying to do it and when your adversaries are. You are absolutely right that people do not know this. The real challenge is to get those authoritative voices heard about what does and what does not work in cyber security, because they are few and far between.

Q124 **Graham Stringer:** How do you deal with that compartmentalisation and specialisation problem?

Henry Harrison: I deal with it in a very specific way, which is not generalisable to other people. The way I deal with it is that I have a particular set of customers in the UK and its allies, who I believe are probably better qualified than anybody else, at least outside of Russia and China, to opine on what does and what does not work. Broadly speaking, if they tell me what I am doing works, then I am going to trust them. That is my approach, but it does not really scale very effectively to anybody else in the industry.

Dr Mavroudis: I agree. Working with interdisciplinarity is a problem but also a necessity. I totally agree with Mr Harrison. I know things in a very narrow spectrum. That is what I am talking about today.

I gave the example earlier of the nuclear fusion reactor start-ups. As part of our research, I realised that some of those start-ups were using machine learning models to control the plasma in the reactor. Indeed, this delivered much better results than the manual way with human operators that they used previously. I wanted to investigate if this could lead to problems with embedded hidden functionality that does not do exactly what you expect it to do at certain times, but I did not have the expertise to investigate that. The way we tackle this in the institute is that we have a very wide network. I was able through that to reach out to nuclear experts in universities. I agree; it took a little bit of effort. We do not speak the same language, but after a few interactions and meetings we got to a really good point.

Q125 **Dawn Butler:** I just wondered, Dr Mavroudis, what excites you about AI and what scares you about AI.

Dr Mavroudis: That is a difficult question. The way I see it, AI right now is a hammer and we are looking for nails. It can do everything and not that much. The examples we see on the news are primarily around chatbots. That is the tip of the iceberg. Overall in the community, there is

a sense that we can do so much more with it. We just need to find the use cases that we can solve right now with the tools we have.

The problem we are facing, perhaps for the first time from a cyber security perspective with machine learning, is that we have a truly black box component that is very hard to analyse and study from first principles and using analytical techniques we have developed in the past decades. They do not necessarily apply to machine learning models. We can compartmentalise it, which in some cases is sufficient, but what I find scary is that we might be pressured to deploy machine learning systems or AI that we do not fully understand. It is important to keep this in mind and follow standard and basic cyber security practices, which will dictate that having such a system as part of a more complicated machinery is not always a good option.

Chair: Mr Harrison and Dr Mavroudis, thank you for your evidence on cyber security and your reflections on some broader questions this morning. That concludes this meeting of the Committee.