ASSURANCE BY PRINCIPLE:

Preparing for the next generation of product security assurance

MATT SPENCER
Centre for Interdisciplinary Methodologies
University of Warwick
CV4 7AL
✉ *m.spencer.1@warwick.ac.uk*

POLICY REPORT
V1.0 15/12/23

WARWICK
THE UNIVERSITY OF WARWICK

UK Research and Innovation

RISCS
Research Institute for Sociotechnical Cyber Security

Executive Summary

Product security assurance policy has wide ranging potential impacts on markets for secure technology, on the security profession, and on the security posture of organisations of many kinds. The UK's new Principles Based Assurance (PBA) model is a bold new vision for the future of product security assurance, learning from the past and anticipating a future of robust cyber security decisions based on explicit, goal-based assurance arguments.

This report aims to support the PBA implementation. It recommends that to ensure that the rollout of PBA is successful, the NCSC should maintain a socio-technical capability specialising in assurance communication, focused on:

- Supporting design-led policy implementation for PBA;
- conducting or commissioning ongoing research into attitudes and understandings of PBA among vendors, customer and evaluator groups;
- developing tools, workshops and models to help stakeholder groups (particularly customer organisations) to understand and use the new methodology effectively, and;
- ensuring that PBA is integrated into long-term strategic thinking.

Over the coming months and years, many stakeholders will make choices about whether and when they embrace the PBA approach: customer organisations will decide whether to request or use PBA reports in their procurement processes, evaluators will decide whether to seek 'Trusted Lab' status, and product developers will decide whether to commission PBA reports for the products they make. In all these cases, adoption entails upfront investment, in time, money and upskilling of staff. Expectations of future returns, however, will depend on expectation about how widely and how rapidly PBA is embraced by other firms (generating economies of scale and network effects). Whether justified or not, narratives within the professional community about whether PBA is likely to be a long-term commitment of NCSC, and whether it will succeed or fail, have the potential to become self-fulfilling prophesies. By informing expectations about potential future returns, such narratives can dampen (or enhance) engagement, generating self-reinforcing feedback loops. This is a particularly salient policy risk due to the relatively short shelf-life of previous schemes. Responsive, engaged, and strategic communication will therefore be vital.

This report summarises recommendations for a policy audience arising from the UKRI-funded 'Scaling Trust' project, which examined the early design and pilot phases of PBA as a case study. The text focuses primarily on contextual considerations relevant to the current stage of implementation. Recommendations from the text are picked out below.

Summary of Recommendations

i. It is recommended that the NCSC develops and maintains an 'assurance communication' capability to support PBA rollout. Such a capability would be highly attentive to emerging narratives about the scheme's success or failure within the professional community and would engage with emerging concerns in an adaptive and responsive manner. 13

ii. Wherever PBA arguments depend upon assertions about users following correct procedure, the need to avoid a 'blame culture' should also be explicitly stressed (Dekker 2018). Customer organisations must be advised to adapt the assurance argument to their actual working practices, rather than attempting to turn assumptions about practices into imposed rules. 18

iii. It is recommended that the NCSC conducts usability testing on PBA reports, and examines real procurement decisions that use them, to gain confidence that customer organisations are understanding reports well and weighing evidence appropriately. This should include decisions where some candidate products have PBA reports available, and some do not. 18

iv. Once PBA has been rolled out more extensively, it is recommended that the NCSC conducts a survey of stakeholders' perceptions of Trusted Labs and uses this data to inform interventions such as communication campaigns or greater transparency around how evaluators are regulated. 19

v. It is recommended that the NCSC explore alternative approaches for setting the scope and scale of PBA evaluations, to understand how different ways of constraining these decisions affect the profile of benefits to different stakeholders. 30

vi. It is recommended that PBA evaluations have an explicit overall goal of evaluating whether a product 'functions as expected', as this will provide evaluators with the flexibility to address any security-related factors of potential relevance to customers. 31

vii. It is recommended that PBA reports include a section of the summary detailing a prioritised list of areas in which further testing during implementation would provide the greatest 'added value' to the assurance case. 32

viii. It is recommended that Trusted Labs be encouraged to develop an internal capability with expertise in the CAE framework, to provide oversight of ongoing projects and to offer training to vendors/customers as part of evaluation engagements. 32

ix. It is recommended that PBA contracts include a standard framework for vendors to commission 'add on' evaluation work where this would enhance the assurance case. NCSC should monitor PBA evaluations during the early rollout and take steps to address any opportunistic behaviour (or perceptions of it). 33

x. It is recommended that re-evaluations refer explicitly to any earlier reports conducted on that product, or product line, by any evaluator, so that changes in confidence level are clearly traceable, and to ensure that product vendors cannot game the system by seeking multiple evaluations. Some form of centralised registry of reports will likely be needed. 34

xi. It is recommended that the NCSC puts in place a mechanism through which customer organisations can receive updates about the availability of newer PBA reports on products they use or that are under consideration. 35

xii. It may be helpful for the NCSC to commission a study examining a sample of PBA reports after a time period has elapsed (one year, say), to determine whether new information that has come to light in that time (principally, new vulnerabilities) makes a material difference to the validity of the arguments. 35

xiii. It is recommended that the NCSC ensures that PBA contracts include a built-in cost framework for future minor updates, such as impacting a newly discovered vulnerability against the CAE structure, with standardised costs and response time. 36

xiv. Whether product developers should be obliged to 'maintain' their PBA documentation is a difficult question. At least in the initial phase of PBA roll-out, any such requirement is likely to discourage adoption, but it should be a priority topic for longer term consideration. 36

xv. It is recommended that the NCSC ensures that PBA is integrated into the strategic vision. One way to do this is to identify adjacent problem spaces and strategic objectives beyond improving procurement, where there are synergies with PBA. This could include a vision of assurance-enabled inclusive security based on the explicit articulation of security reasoning, relating evidence to claims and claims to the security goals of diverse stakeholders. 39

VERSION CONTROL

Date	Version	Comment
Aug 23	V0.1	Early draft produced for feedback on content and structure
Oct 23	V0.2	Draft for general feedback and review
Dec 23	V1.0	Final version

DATA AVAILABILITY STATEMENT

In order to protect participant confidentiality, supporting data cannot be made openly available. Further information about the data and conditions for access are available from WRAP at http://wrap.warwick.ac.uk/166796/

KEY TERMINOLOGY

Accreditor	A term used within the UK Ministry of Defence for the individual responsible for approving technical products for use
APC	Assurance Principles and Claims
CAE	Claims, Argument, Evidence
Certifier	The individual or organisation responsible for approving an evaluation, so that certification of the evaluated product is authorised
CESG	Computer Electronic Security Group
CPA	Commercial Product Assurance
Common Criteria	Common Criteria for Information Technology Security Evaluation
Customer organisation	The organisation with an interest in buying and/or using an evaluated product
Developer	The organisation that has developed an evaluated product
Evaluator	The organisation that conducts a product security assurance evaluation
ITSEC	Information Technology Security Evaluation Criteria
NCSC	National Cyber Security Centre
PBA	Principles Based Assurance
TCSEC	Trusted Computer Security Evaluation Criteria
Trusted Labs	The name given to the evaluators authorised to perform PBA evaluations

DISCLAIMER

The views expressed in this document are the author's alone and do not represent the views of the NCSC or any other official body.

ACKNOWLEDGEMENTS

This work was funded by UK Research and Innovation, grant MR/S037373/1

LICENSE

Table of Contents

Background

PROBLEM DEFINITION

Over the past 3 years, the UK's National Cyber Security Centre (NCSC) has been developing a new approach to assurance: Principles Based Assurance (PBA). PBA is expected to play a significant role in ensuring the cyber security of the UK in the future. It is a new approach for technology security evaluation with wide-ranging implications for product development across a wide range of product categories, for procurement and risk management practices in customer organisations, and for the role of third-party evaluation in organisational security. As a policy intervention into the institutions of security assurance, PBA has the potential to impact a very large number of organisations in the UK. Implementation will be gradual and will need to be evidence-based and responsive to ensure that the NCSC's goals are met.

MOTIVATIONS FOR PRINCIPLES BASED ASSURANCE

The PBA approach is motivated by a desire to ensure that organisations in the UK can have appropriate confidence in the security of the technologies they use. It will entail the construction of an institutional framework for product developers and customer organisations to commission third party evaluations of technical products (hardware-based and software-based). With PBA, the NCSC seeks to put into practice lessons learned from the successes and failures of previous approaches to assurance, and aspires to create an enduring and flexible scheme capable of balancing risk, cost and complexity.

The main elements of PBA are:

- **Principles:** formal texts authored by the NCSC specifying normative security goals for distinct categories of products. Instead of prescribing details of product design (as previous schemes have done), principles set out overarching desired outcomes.
- **Evaluations:** assurance activities will be carried out by 'Trusted Labs' authorised by NCSC to produce PBA reports.
- **Reports:** the write-up of assurance activities created by the evaluator. Reports communicate the evaluator's degree of confidence in the product, key findings and risks, supporting potential customer organisations to make informed decisions. Instead of communicating isolated test results, PBA reports will construct an assurance argument, explicitly relating evidence to the overarching claims about security being made.

Image by Michael Dziedzic / Unsplash

The underlying philosophy of principles based assurance was set out in a 2021 White Paper[1], with later updates clarifying the anticipated operating model[2]. While the overarching approach is largely settled, there remains some areas in which the details of implementation remain to be worked out. These include:

- Details of institutional arrangements for PBA evaluations, including governance of evaluations and 'maintenance' of assurance arguments.
- Details of the business model and patterns of interaction that will emerge within this new institutional arrangement.
- Details of the style of communication to be used within assurance reporting, including templates, and conventions for re-use.

To gain an initial understanding of PBA, it is helpful to consider how it departs from traditional forms of product security assurance. Traditional product assurance can be divided into two main categories:

1. Evaluations of commodity technology under Common Criteria (ISO 15408) or Commercial Product Assurance (CPA). These exercises, which rely on prescriptive criteria, specifying details of the design for each product category, produce a relatively simple 'pass or fail' certificate. They are best suited to contexts where decisions are simple, and a customer organisation wants (or needs due to regulation) to use a product with a recognised certificate.
2. Specialist security evaluations commissioned as one-off projects from test labs, or through NCSC schemes like CTAS (CESG Tailored Assurance Scheme). These exercises provide extensive, context-specific reports, tailored to that customer organisation's use case. These are expensive and typically used where security is a critical consideration (for example in military applications).

[1] https://www.ncsc.gov.uk/collection/technology-assurance/future-technology-assurance

[2] https://www.ncsc.gov.uk/blog-post/making-principles-based-assurance-a-reality

Figure 1: Schematic depiction of product security assurance prior to principles based assurance

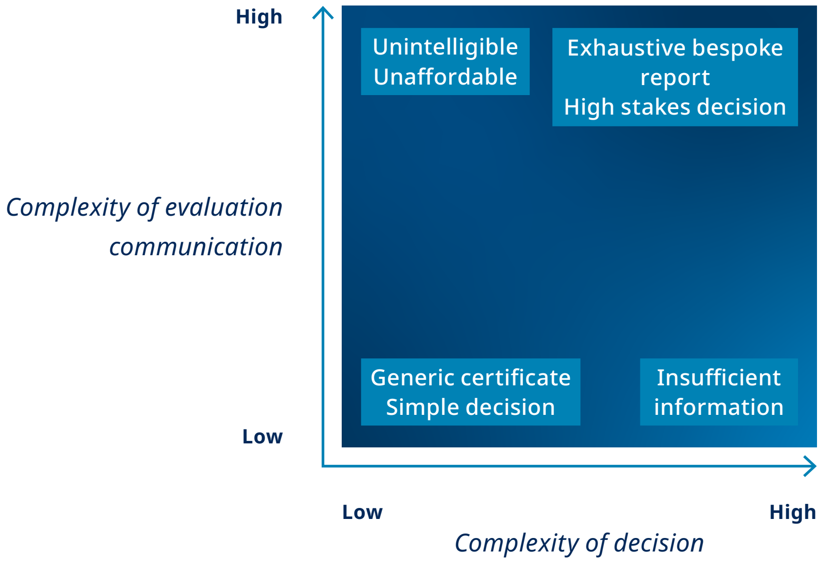

These two categories represent a general matching of the degree of complexity of the evaluation (and, crucially of the evaluation communication) against the degree of complexity of the decision needing to be made (see Figure 1). Principles based assurance can be understood as introducing a new 'middle range' option, aiming to produce evaluation reports of intermediate complexity. Instead of giving a simple 'pass/fail' result, the goal is to communicate the evaluator's assessment in a way that supports a customer to make a risk-based assessment of product suitability for their context. The goal, put simply, is to support organisations to make better (more informed) cyber security decisions, with a level of complexity and cost that is appropriate to their setting.

Figure 2: Schematic depiction of technology assurance, showing principles based assurance as intermediary

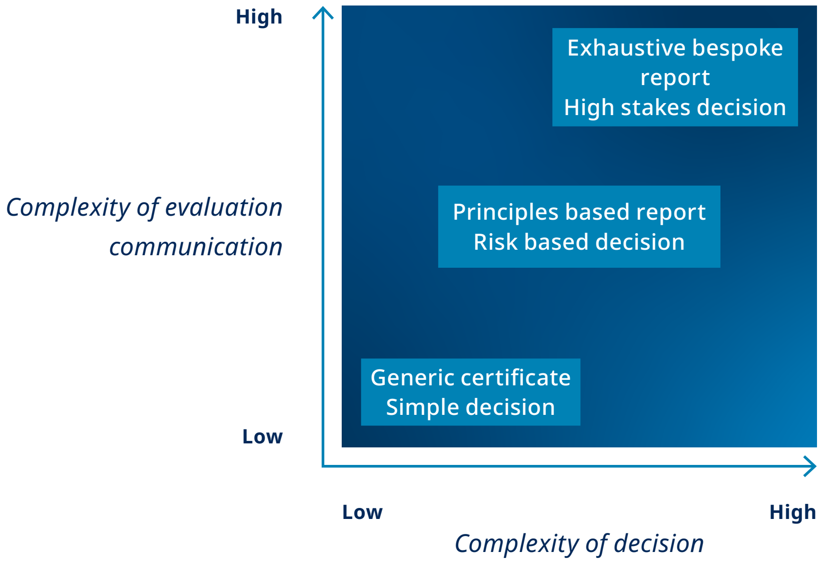

The development of PBA has followed from the recognition within NCSC that many decisions about technology use sit in this relatively under-served zone of intermediate complexity, where organisations would benefit from having more information than is conveyed by a simple certificate, and where the risk is not high enough to warrant commissioning a costly bespoke security evaluation. This intermediate space has historically been partially addressed by security consultancies (firms and independent consultants) that offer advisory services about technology choices. The distinct value proposition of PBA lies in standardisation, trust, and reuse. Principles provide a common framework for assurance, a common language with the potential improve communication within and between organisations, Trusted Labs offer the benefits of third-party advice with additional NSCS oversight ensuring independence and impartiality, while evaluations promise economies of scale, as a PBA report can be shared with many of a product developer's potential customers.

AIMS FOR THIS POLICY REPORT

This policy report summarises insights with a bearing on the PBA implementation that have arisen from the research conducted under the first phase of the Scaling Trust project (October 2019 – April 2024). This is UKRI funded Future Leaders Fellowship based at the University of Warwick.

In this document, I approach the problem holistically, contextualising PBA theoretically and empirically, and then unpacking the main elements of the approach in light of these observations. Recommendations are highlighted throughout the text, as well as being compiled in a dedicated section at the top of the document.

The brief draws on the following sources:

- **Literature reviews:** The historical and theoretical background to principles-based assurance was studied through reviews of academic literature on the history and theory of computer security, and approaches to governance, policy and regulation.
- **Consultation with NCSC specialists:** Discussions with NCSC representatives shaped the research from the outset and inform the content.
- **Interviews:** This report draws on over 30 expert interviews conducted as part of Scaling Trust. These included interviews with cyber security professionals, risk managers, evaluators, technology developers, and policymakers.
- **Documentary analysis:** Reports from a pilot of principles-based assurance, from CPA evaluations and from CTAS evaluations were studied, excerpts extracted, sanitised and fictionalised, to support the analysis of communication.
- **Workshops:** This report draws on insights from a set of four 'Trust Mapping' workshops, conducted to examine technology assurance within organisations, after implementation.

The Context of Principles Based Assurance

This section examines five areas that have influenced PBA, and/or will influence its reception: certification schemes, bespoke security evaluations, audit and accountability more generally, cloud security principles, and safety cases — and draws out potential lessons for implementation.

CERTIFICATION SCHEMES

Product security assurance is a well-established area of concern that played an important role in the history of computer security. In the 1960s and 1970s, much of the pioneering work of figures like Willis Ware, James Anderson and Grace Nibaldi was responding to practical questions about how the security of computer systems could be evaluated. Governments around the world, including in the US where much influential early research was funded, wanted to be able to use commercially developed computing systems to handle sensitive data. To do this, they needed assurance that these systems met appropriate security requirements. This meant pinning down what 'secure' meant when applied in this context, and designing evaluation methodologies capable of establishing whether the evidence supported the claim that a given system met such a definition.

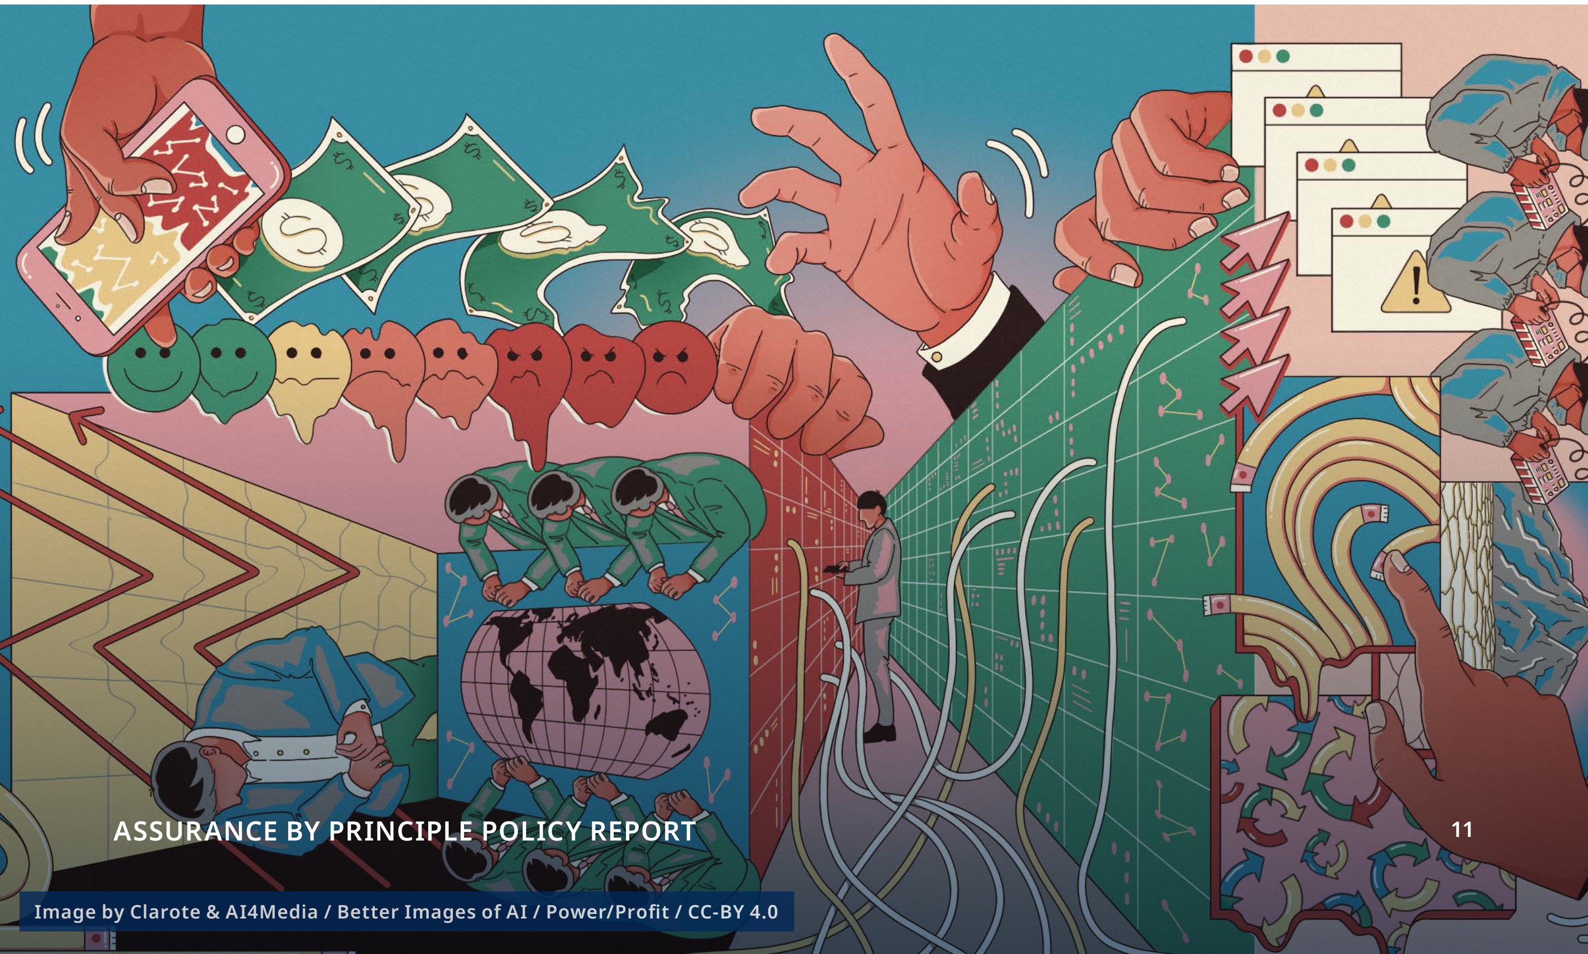

Image by Clarote & AI4Media / Better Images of AI / Power/Profit / CC-BY 4.0

The US TCSEC (Trusted Computer Security Evaluation Criteria), launched in 1985, was the first standardised scheme for security assurance evaluations. A set of pain points soon emerged, however, as evaluations were costly and time consuming. These issues would continue to motivate further development of product security assurance approaches right through to this day, proving persistent and hard to solve. Developed in the 1990s, the Common Criteria for Information Technology Security Evaluation was another milestone. Common Criteria replaced schemes like TCSEC and the European equivalent ITSEC and offered greater economies of scale by enabling recognition across multiple international markets. Common Criteria also responded to important developments in the market, by allowing for evaluations to be tailored according to product categories, recognising that security assurance no longer applied simply to 'computers' but to more specific categories of products (such as routers, firewalls, operating systems, and so on).

Like the earlier schemes, Common Criteria in turn gained a reputation for being costly and time consuming, and the discovery of vulnerabilities in certified products led to doubts about its effectiveness, and about the perils of reliance on certified products. In 2011, in response to these issues, CESG (the part of GCHQ responsible for information security) launched the UK Commercial Product Assurance (CPA) scheme. This domestic scheme was intended to provide a lighter-weight route to product certification, initially to satisfy government procurement rules, but later extended to other product categories (electricity and gas smart meters, for instance). However, concerns emerged in turn about CPA encouraging 'tick box' security, about time and cost, as well as patchy participation by product developers. This led to a growing consensus among policymakers that a more radical overhaul of product security assurance was needed. Paving the way for PBA, in 2019 the UK gave up its 'certificate producer' status under Common Criteria, and in 2020 it was announced that CPA would also be phased out. Work on PBA was underway and would be publicly announced in 2021.

It is clear from this context that it is vital for PBA to be, and to be understood to be, a decisive break from the past. Ensuring that PBA will endure longer than predecessor schemes will require a responsive and agile approach, attuned to emerging challenges and attentive to lessons learned. If PBA is perceived to be repeating the mistakes of the past, it will be very hard to gain the kind of stakeholder buy-in necessary for enduring success. This is the claim I would attach most significance to: the conditions are present for narratives about PBA failing to live up to its promises to become 'self-fulfilling prophesies'. Crucial to the success of the PBA rollout, then, is the ability of the NCSC to be highly attentive to emerging narratives among the professional community, and to maintain an adaptive and responsive approach to community engagement. This is what I describe here as an 'assurance communication' capability.

i. It is recommended that the NCSC develops and maintains an 'assurance communication' capability to support PBA rollout. Such a capability would be highly attentive to emerging narratives about the scheme's success or failure within the professional community and would engage with emerging concerns in an adaptive and responsive manner.

Like traditional schemes, PBA evaluations are conducted in reference to formal documentation of expectations (the principles) and the evaluation is carried out by a third-party private sector evaluator (indeed many of these labs are likely to be the very same organisations that carried out CPA evaluations under the previous regime). It also involves the same broad types of evidence gathering activities, including testing, inspection of hardware, review of documentation, interviewing and witnessing. The reports themselves will have a different style, but nevertheless retain some similarities in overall structure, with a summary followed by a breakdown of evidence. The main points where PBA departs from convention are as follows:

- Traditional security assurance schemes use prescriptive standards that define details of product design. This is regarded to be problematic because it constrains innovation and because it means whoever sets the standards shares responsibility for secure design with the product developer. PBA's principles are outcome-based, and this allows much greater flexibility for innovation, and ensures that responsibility for secure design sits clearly with the product developer. It is also more flexible, allowing for the evaluation of products that partially meet principle-sets. This is reflected in PBA report structure, which is likely to place the emphasis on communicating areas in which claims are unsupported (in traditional sciences, products which do not meet all requirements would not be certified).
- The prescriptive standards used in traditional assurance schemes define evaluation activities to a high level of detail, so that evaluations are typically fairly formulaic. PBA offers more opportunity for evaluators to add value by explicitly developing assurance activities and arguments based on considerations of cost effectiveness.
- Traditional schemes generate simple certifications that create the conditions for 'blame avoidance' (Hood 2010), where security problems that do arise can be blamed on the evaluation. PBA reports present arguments concerning security that support explicit risk-based decisions, designed to support traceability and clearer apportionment of responsibility.
- Traditional schemes failed to gain comprehensive participation within the market, with some vendors choosing not to obtain certifications due to expense, and newer products generally being uncertified due to time delays in the process. PBA is designed to be more flexible, and its qualitative reports should enable easier comparison of evaluated and non-evaluated options than the simpler 'all or nothing' certifications. It is also possible that PBA evaluations can be performed more rapidly than those of traditional schemes, although this will need to be established during implementation.

BESPOKE SECURITY EVALUATIONS

Instead of relying on generic certifications alone, organisations that operate in high security contexts also commission bespoke evaluations of key products as part of the procurement process. Such exercises may also be conducted to support ongoing reviews and audits of security risk. It is common for companies that offer evaluation services under certification schemes to take on this kind of custom project work as well. The requirements that form the backbone of the exercise are provided by the customer organisation rather than being generic standards (though standards may be used as a part of the requirements). The report on the product is also addressed to the customer's specific needs. While PBA is most directly a replacement for certification schemes like CPA, the approach also draws on lessons learned from the NCSC's experience conducting or commissioning risk-based bespoke evaluations.

An important example is CTAS (CESG Tailored Assurance Scheme), under which bespoke security evaluations are commissioned by NCSC on behalf of government customers. These assessments are carried out by private labs, but the NCSC acts as an intermediary, helping to specify requirements, conducting quality assurance on the reports and providing a summary of findings and recommendations. Each exercise begins with the design of an 'Evaluation Work Programme' based around a set of low-level declarative claims, such as:

> ***'Changes to the configuration shall be logged.'***

The work programme specifies activities to be performed that are deemed capable of giving confidence in whether such claims are true. The NCSC acts as an intermediary, reviewing and approving the work programme prior to the evaluation taking place, with a focus on the coherence of the underlying logic. An example of such feedback:

> ***'Consider changing "Ensure" to Demonstrate" as the Accreditor should be provided with evidence that something is taking place. The evaluator cannot ensure things are happening but can demonstrate whether or not they are.'***

The development of PBA has relied on pilot exercises, which have focused on understanding the ability of evaluators to produce work programmes and arguments of the expected kind. In these pilots, the NCSC has acted as mediator in a way reminiscent of CTAS, reviewing and amending plans and reports. However, as the implementation of PBA goes forward, the NCSC will need to step back, and ensure that the right framework is in place to ensure that the evaluations conducted under PBA conform to expectations, without the need for NCSC to act as an intermediary.

The CTAS report structure follows the familiar format of certification schemes, with a short introductory section followed by tables documenting claims against activities undertaken and results. Unlike in CPA evaluations, where the process must end in a pass or fail decision, CTAS evaluations may also conclude with a ‘partial pass’. This is made possible by the fact that (unlike in CPA) the report itself is made available to the customer organisation. It also means that areas of incomplete evidence are more easily accommodated, for instance where pragmatic factors impede the availability of certain items of evidence. In allowing for ‘unsupported’ as well as ‘supported’ claims, PBA has a further resemblance to this form of reporting.

The CTAS ‘assessment statement’ is drafted by specialists within the NCSC rather than by the evaluator. A fictional example of such a statement is provided below.

Changes to the configuration shall be logged

Findings

There is no comprehensive logging capability present for [the system]. Whilst there is some very limited event logging, it is not considered to be adequate. Furthermore, the lack of a syslog collection / SIEM platform means any logs created are stored locally on devices, are offloaded periodically and only ever reviewed on an ad-hoc basis if, or when, there is an occurrence which necessitates it. Proactive monitoring of logs to establish patterns and highlight problems is not undertaken for the system.

The [role] can make changes to the [system] to address operational issues… but changes made in this manner will not necessarily generate any events. This lack of logging could potentially allow for unauthorised modifications to be made to the system with no record captured or retained for review or evidential purpose. The environment did feature a fault log that tied into endpoint monitoring; however, this was provided from an operational and solution status monitoring perspective rather than being used to audit and log changes that may reflect an impact to asset security.

There is a formal operational defect process for notifying the developer of problems, this process follows a defined workflow to ensure the problem is fully understood, replicated and a solution defined. Any updates / upgrades required to address operational defects raised go through a formal review and release process. Scheduled changes must go through a full and comprehensive change management process which is defined in accordance with the plan – Design Change Procedure before being authorised for implementation.

Recommendation

Look at introducing a Security Information Event Management (SIEM) system within the [system] which captures any details of any changes made to the system as a whole.

In PBA, the evaluator will be expected to author the report, including the high level summary. The fictionalised CTAS assessment statement provides an insight into some of the key requirements of PBA reports, given that the latter will need to communicate the implications where principles are partially met. Each paragraph explains a finding in relation to the overarching claim, unpacking it in terms of lower-level details, and/or implications, and/or limitations of existing measures. There is no simple formula for how this should be done: in each case the author must judge what the reader will understand, what clarifications are needed, and how that understanding may develop over the steps of the argument. In the example above, the recommendation at the end is very short and simple, and note that there is an implied argument that taking this recommended step will eliminate the issues discussed in the findings.

One reason that CTAS assessment statements are written by NCSC specialists is that they have greater knowledge of the customer context. In PBA, the evaluator will be expected to construct these kinds of summaries without necessarily having specific knowledge of the customer context. Indeed, a PBA report may need to address a wide range of possible customers. Given the novelty of this distribution of responsibilities, it is likely that support will be needed in the early days of the PBA implementation to ensure that such summaries are written in a format that is widely understandable.

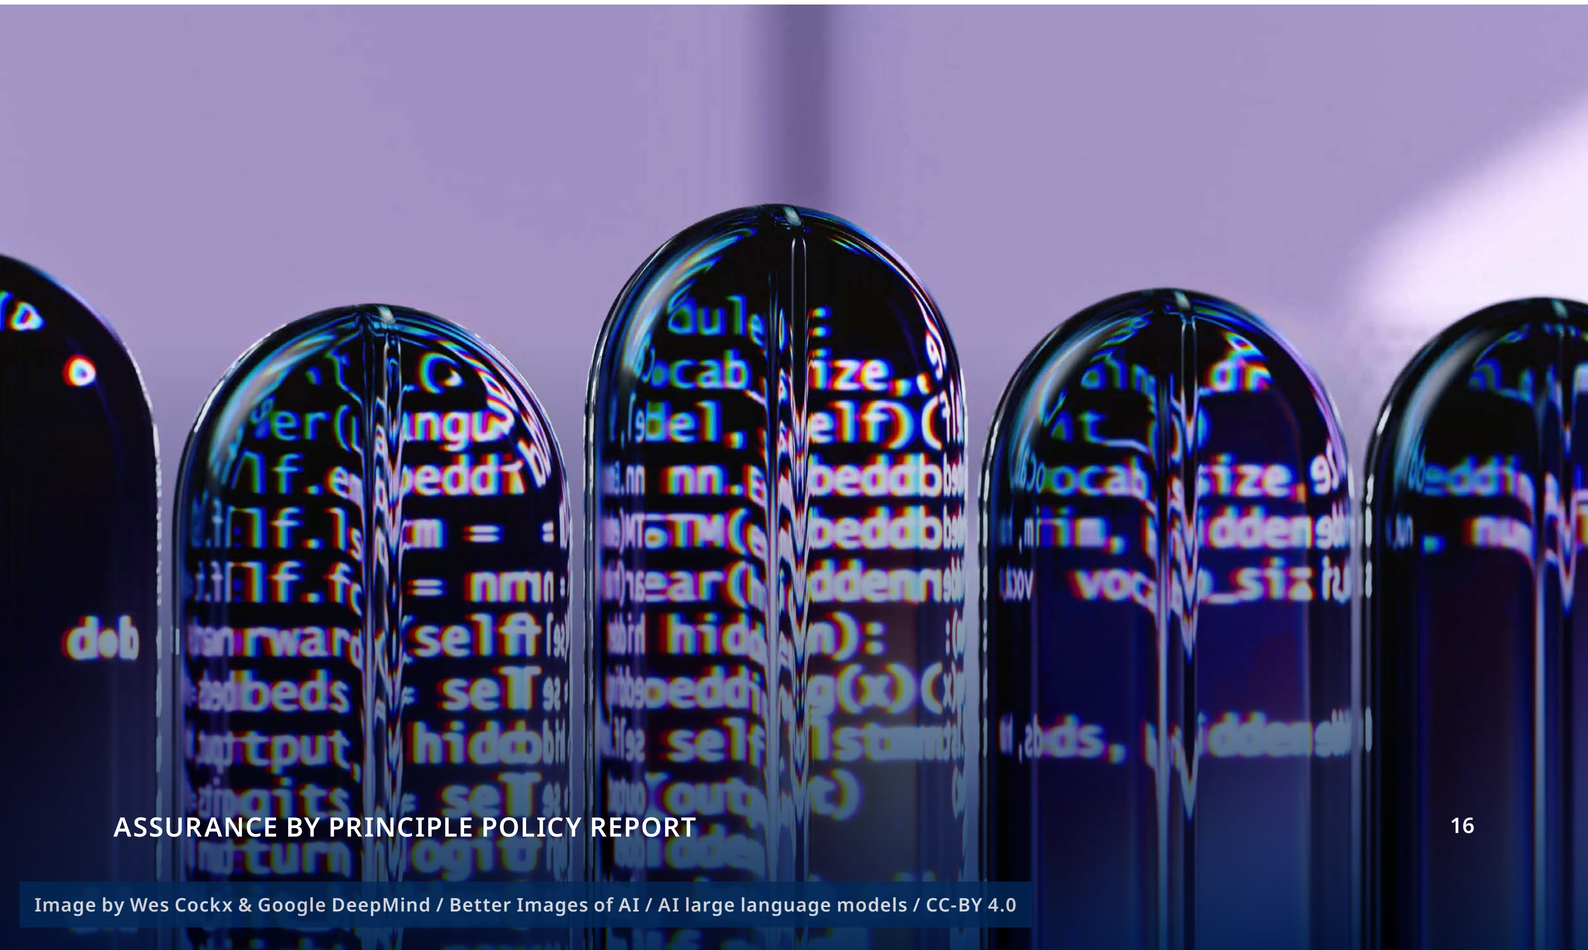

Image by Wes Cockx & Google DeepMind / Better Images of AI / AI large language models / CC-BY 4.0

AUDIT AND ACCOUNTABILITY

Taking a step or two back in our contextual analysis, we can put PBA in the context of the great expansion of audit and accountability that has spread across many different subject matters over the past century and a half.

Financial audit as we know it today was developed in the 19th Century in response to the emergence of shareholder capitalism. Across the 20th Century, these techniques expanded into areas as varied as environmental accounting, quality management, and equality, diversity, and inclusion (EDI). In the 1990s the 'New Public Management' saw the techniques of audit spread to an array of public institutions. Councils were required to report on metrics of care and service quality, universities on research excellence, hospitals on waiting times, and so on. In each case, specific forms of value are subject to measurement, qualities made commensurate and quantified, enabling the creation of durable records of evaluation that go on to enable accountability and decision making at many levels.

But with the growth of audit and measurement, a series of important critiques have emerged, and any new evaluation policy ought to be constructed in a way that is mindful of the kinds of unintended negative outcomes that have been observed in other areas.

1. Critics argue that accountability can undermine safety where it incentivises people to keep quiet about problems, and thus obstructs organisational learning and contributes to 'disaster incubation' (Dekker 2018).
2. Critics note that metrics are often only partially aligned with desired outcomes, and accountability can create the conditions in which organisations chase metrics instead of focusing on improving outcomes, leading to potentially absurd results (Tsoukas 1997).
3. The aura of impartiality associated with the figure of the auditor has become significantly weakened, given that today activities such as consultancy represent the major source of revenue for most of the large audit firms (CMA 2018). Given 'vulnerable' public trust, the Financial Reporting Council recently emphasised the need for auditors to 'live by the spirit' of the process, not by compliance to rules and standards alone (FRC 2018: 5).

While these concerns have been expressed in very different domains, all three represent key threats to the future perceived legitimacy of product security assurance schemes:

1. As well as depending on facts about the intrinsic characteristics of the evaluated product, assurance arguments also depend upon assumptions about working practices (for instance relating to use, updates and maintenance). Under PBA, decisions about procurement would be made based on such assumptions. There is a risk that the approach incentivises users/maintainers to keep quiet about areas in which they transgress specified operating procedures. Assurance arguments must be adapted to the idiosyncrasies of real working practices. The big challenge for PBA in this respect is that the PBA report is generic and cannot be adapted to the practices of any particular organisation. It is therefore important that the interpretation of assurance arguments, and their translation into an organisation's own internal assurance processes, is informed by considerations of usability.

ii. Wherever PBA arguments depend upon assertions about users following correct procedure, the need to avoid a 'blame culture' should also be explicitly stressed in the text (Dekker 2018). Customer organisations must be advised to adapt the assurance argument to their actual working practices, rather than attempting to turn assumptions about practices into imposed rules.

2. Principles based assurance is a direct response to policymakers' concern that schemes like CPA may incentivise customer organisations to pick certified products over other, potentially better, options, because certification can readily deflect blame. The judgement that customer organisations must exercise in interpreting a PBA report is supposed to overcome this 'short circuiting' of responsibility.

iii. It is recommended that the NCSC conducts usability testing on PBA reports, and examines real procurement decisions that use them, to gain confidence that customer organisations are understanding reports well and weighing evidence appropriately. This should include decisions where some candidate products have PBA reports available, and some do not.

3. The rather turbulent history of product security assurance has led to widespread cynicism about the figure of the security evaluator, and whose interests they really serve (see Spencer 2022 for a more detailed look at this issue). Interviewees recanted stories of "gamed" evaluations. One interviewee suggested that governments' interests in running these schemes was less about improved outcomes for customer organisations, and more about creating a reliable stream of funding to support a reserve of expert security professionals within the country. Ensuring that the Trusted Labs are perceived as working in the interests of customer organisations is important, but a challenging area for policymakers to have direct influence on.

iv. Once PBA has been rolled out more extensively, it is recommended that the NCSC conducts a survey of stakeholders' perceptions of Trusted Labs and uses this data to inform interventions such as communication campaigns or greater transparency around how evaluators are regulated.

CLOUD SECURITY

An influential precedent for PBA was the CESG cloud security principles, published in 2014. This framework was designed to support customers (with a particular focus on government departments) in making the transition from traditional datacentres to cloud service providers. Doing so required a change in mindset around assurance, as organisations that had been used to being able to physically inspect datacentres or commission third party audits were not generally able to do so when it came to cloud. The cloud security guidance that CESG produced set out principles for how to assure for cloud architectures, given that this would need to be a collaborative exercise between customer organisations and cloud service providers. In some areas, independent inspections may be possible, but the principles also set out guidance for circumstances in which assurance would have to rely on the assertions of service providers.

Although the main focus of principles based assurance is on products that can be inspected, the success of the cloud security guidance demonstrated the potential value of principles. Principles describe the kinds of claims that a customer organisation needs to gain assurance in, and guidance is provided on how that may be done, instead of depending on creating a specification for the underlying technology. This approach also opened the way for greater engagement with safety assurance.

GOAL-BASED REGULATION AND SAFETY CASES

Intersections between safety and security have been proliferating in recent years. The principles around which PBA is constructed are inspired by goal-based regulation, an approach with roots in safety regulation. In his report on the 1988 Piper Alpha disaster (an oil rig explosion that killed 165 people), Lord Cullen found fault with prescriptive safety regulations, and remarked that:

> ***'[m]any regulations are unduly restrictive in that they are of the type which impose 'solutions' rather than 'objectives' and are out-of-date in relation to technological advances… There is a danger that compliance takes precedence over wider safety considerations… The principal regulations should take the form of requiring stated objectives to be met.' (1990: 4).***

Cullen also recommended that organisations in safety critical industries should maintain 'safety cases'. The British safety case approach thus coalesced around three principles (adapted from Leveson 2011):

- Responsibility for controlling risks sits with those who create the risk;
- safety should be based on setting and achieving goals instead of focusing on compliance with prescriptive standards, and;
- risks that remain should be understood and acceptable.

Specialists in safety science have developed a number of practical methods for compiling safety cases. Goal Structured Notation, for instance, is an influential approach, providing a formal framework for organising the safety argument. A related approach is Claims, Arguments, Evidence (CAE), developed at Adelard in the 1990s by Peter Bishop and Robin Bloomfield. CAE will be adopted as the framework for organising and communicating PBA evaluations, with each set of principles developed into an 'Assurance Principles and Claims' (APC) document which provides a starting point for evaluation[3].

[3] https://www.ncsc.gov.uk/blog-post/making-principles-based-assurance-a-reality

Inspired by the philosopher Steven Toulmin's naturalistic account of argument, CAE aims to provide a framework for an evaluator to give an explicit account of how various kinds of evidence gathered in the evaluation process warrant confidence in the overarching claim being asserted, through deterministic, probabilistic, and qualitative forms of argument. Toulmin suggested that argument be studied on analogy with jurisprudence. 'A sound argument, a well-grounded or firmly-backed claim, is one which will stand up to criticism, one for which a case can be presented coming up to the standard required if it is to deserve a favourable verdict' (2003 [1958]). Given that safety cases do at times find their way into the courtroom, the analogy is particularly apt. From Toulmin, the CAE approach inherits an emphasis on the pragmatics of reception: a safety case is one that can be *demonstrated*.

The contrast with traditional product security assurance is informative: while CPA or Common Criteria evaluations may be highly rigorous and extensive, the resulting tabular portfolio of evidence cross references evidence only with low-level claims. Exactly how this set of many low-level claims adds up to warranted confidence in the overarching argument about product security is often opaque and indeed treated as a responsibility of the standard-setter, not of the evaluator. CAE, in contrast, generates tree-like diagrams that visually represent the hierarchical decomposition of the argument. In the PBA reports, which refer to and extend decompositions set out in APC documents, the evaluator and customer organisation will be able to 'see' how the evaluation 'adds up' in a way that has rarely previously been possible in cyber security.

CAE has itself evolved and PBA will need to remain attuned to new developments in the future. The original description of CAE focused on making clear the positive argument that a system is safe. Faced with concerns about confirmation bias, an updated approach to CAE as described within 'Assurance 2.0', also places emphasis on complementary 'negative' techniques for the analysis and documentation of 'defeaters', conditions that would undermine the argument (and thus potentially the safety of the system) (Bloomfield & Rushby 2020). CAE has thus moved beyond Toulmin's legal analogy (in court, making a case means presenting only one side of the argument) and ensures that the focus is on confidence in the veracity of the claim being analysed, and not merely on confidence in the persuasiveness of an argument to be made about it, a subtle but crucial distinction.

The Theory of Principles Based Assurance

This section briefly summarises three theoretical lenses through which PBA can be usefully analysed: firstly a functional analysis, secondly an analysis of judgement devices and finally an analysis of policy instruments and the changing nature of governmental social control.

ORGANISATION: JUDGEMENT AS A SOURCE OF REQUISITE VARIETY

Principles based assurance represents an attempt to exploit the potential of human contextual judgement as a source of cyber resilience.

At the most general level, the problems of cyber security arise from the potential for computer and network technologies to be turned to malicious use. Organisations use a variety of security control measures to constrain such possibilities, but that constraint is always incomplete, and knowledge of the true robustness of control measures is always limited and fallible.

Under a simple functional analysis, PBA, like alternative assurance approaches, can be understood as a response to this general control problem. Organisations face an operational environment that may throw both expected and unexpected inputs at their systems' interfaces. Security controls limit the ability of unexpected inputs to product unwanted outcomes. PBA provides the organisation with information about control measures that reduces uncertainty and allows for better optimisation.

In his work on regulation and control, the cybernetician Ross Ashby set out what he called the 'law of requisite variety': the more variety a system is exposed to from its environment, the wider the channels of communication it needs, and the more variety in internal structure it needs, if it is going to maintain the same degree of control over outcome (Ashby 1956, chapter 11). Traditional security assurance schemes provide an organisation with information in a binary 'yes/no' 'certified/uncertified' schema. In Ashby's terms, this is a narrow channel of communication, and requires correspondingly modest internal structure to interpret this information.

Principles based assurance generates reports on evaluated products which do not simply recommend them for use: instead, they provide an assessment with much more nuance, and thus constitute a wider communication channel. But for this to amount to an ability to cope with greater external variety, correspondingly greater internal structure is also required. This 'internal variety' is found in the ability of decision makers to exercise contextual judgement about the appropriateness of evaluated products, something that requires familiarity with the format, and with the contextual environment of the organisation. In short, with PBA, the capability of people to make judgements about what is a good choice in their context is employed as a resource for cyber resilience.

MARKET: PBA AS A JUDGEMENT DEVICE

Principles based assurance represents a re-organisation of judgement devices in the market for secure digital technology.

Markets for security-related products are affected by quality uncertainty: it is hard for buyers to judge the quality of the products they may wish to purchase (i.e. how secure they are). The largest customer organisations may be able to persuade a product developer to provide technical briefings, even facility tours or formal audits. But many organisations have little access to information, and in any case it may be difficult to ascertain even which qualities they should be looking for.

Information economics has focused on homogeneous markets in which products have intrinsic quality. Economic sociology, in contrast, has developed our understanding of differentiated markets where what counts as a reasonable choice depends upon collective understandings as much as on intrinsic properties. Following Karpik (2010), we can distinguish a number of types of 'judgement devices' that help buyers to determine what is a reasonable choice.

- **Personal Networks** are an important source of credible insight about products, and primarily operate through informal spoken communication.
- **Critics and guides** exert a soft form of authority over the space of products, through public communication, reviews and evaluations.
- **Confluences** are techniques that guide buyers towards certain products. In cyber security, we can observe the construction of confluences at professional conferences, through the tradition of company stalls, employing backdrops, welcoming sales representatives, free gifts and demonstrations.
- **Rankings** are constructed from comparative metrics that enable products to be sorted onto a common scale. Cyber security examples include Gartner Peer Insights.
- **Appellations** include brands and certifications and are labels that associate products with specific origins or that have met a certain standard of quality. All cyber security products are branded in this way, and this is also the space in which traditional assurance schemes have intervened, generating a kind of 'label' that marks certified products as having passed the evaluation.

Principles based assurance can be understood as creating a new kind of judgement device that is somewhere between an appellation and a guide. To pass through the process, it is necessary that evaluated products meet some minimum standard of quality, but instead of producing a simple label, it produces a guide reflecting on the qualities of that particular product. Furthermore, the market shaping potential of such a device goes beyond enabling the claims of a seller to be trusted: the new PBA process may improve buyers' abilities, in general, to what counts as security-relevant quality.

GOVERNMENT: PBA AS TOOL OF SOCIAL CONTROL

The development of principles based assurance represents a subtle shift in how government acts in relation to cyber security.

In a classic analysis of governmental social control, Christopher Hood (1980) categorised the resources available to government into four categories, with associated forms of control:

Resource	Type of control	Examples in cyber security policy
Monetary	The government controls the printing of money, as well as aspects of its flow via financial policy.	Monetary policy has little direct relevance to cyber security. However, governmental budgets for cyber security technology can be large and for countries such as the USA, this purchasing power has historically exerted a significant influence on the market.
Nodality	The government occupies a highly connected position in networks of communication, which gives it the ability to intervene through public information campaigns and propaganda.	A great deal of cyber security policy operates through communication. This includes: • The organisation of sector-specific cyber security interest groups, forums, and conferences. • The publication of advice, guidance and information campaigns aimed at citizens or organisations of various types. • Press releases from government security experts, engagement with the media, and so on.
Organisation	Government is itself an organisation of significance size and capability, and can directly act in its own right.	The government necessarily invests considerable resources in securing its own digital infrastructures. In some priority sectors, government resources are directly involved in aspects of security work. This includes: • Incident response in contexts of national importance. • Government intervention, for instance taking down Botnets.
Authority	Government has powers to determine conduct in many ways through law and regulation.	Legislation in the UK addressing aspects of cyber security includes the Computer Misuse Act (1990), the Data Protection Act (2018) and the Network and Information Systems Regulations (2018). Many sectors of the UK economy are governed by independent public bodies that are granted regulatory powers by government, but which are funded through fees from regulated firms. Examples include the Financial Conduct Authority (FCA) and the Solicitors Regulation Authority (SRA).

Product security assurance has historically involved a combination of resources: authority in establishing or giving recognition to formal schemes, and in requiring the use of schemes in certain sectors. Electricity and gas smart meters, for example, require certification in the UK through the Commercial Product Assurance scheme. The most extensive mandatory use of certified technology has been in the public sector, and assurance schemes thus have wielded power through a combination of authority and organisation, exerting influence on how and what products are developed via control of government procurement. Underlying the influence government exerts on the market is the use of nodality, the ability to set, or at least influence, the standards to which products are developed and evaluations assured, and thus produce the key documents that communicate product expectations to the market.

Principles based assurance weakens the use of authority and heightens the use of nodality. It represents a subtle but important shift in the orientation of governmental social control. Because principles based assurance does not produce categorical 'yes/no' answers, it cannot so readily be used in strict regulatory regimes and is more suited to softer regulatory approaches. But in making greater demands on the ability of customer organisations to exercise good judgement, the use of nodality is enhanced. The language of the principles and of the reports is intended to be used beyond the evaluation process, and to some extent will need to be adopted by consumer organisations in their internal risk evaluation processes. In this respect, PBA can be understood as part of a general effort by the NCSC to improve the quality of risk communication as applied to matters of cyber security.

Hood's typology of government resources can be used to examine the ways in which government can gather information as well as how it can exert influence. Intervention typically requires some form of feedback loop. It is likely to be relatively easy for government to observe labs and evaluations, and to examine the reports that are output. More challenging (but nevertheless highly important) will be the observation of stakeholder perceptions and of how customer organisations are using PBA reports.

Image by Google DeepMind / Unsplash

The Components of Principles Based Assurance

This section provides a high-level overview of the main components of PBA: the principles, evaluation approach, communication approach, evaluation governance, operational governance and business models. In all these areas previous policy development has established the foundations, but there remain important details to be explored through implementation.

PRINCIPLES

Principles say ‘what’ is to be achieved, not ‘how’. In addition to creating the institutional architecture for PBA, the NCSC also plays a key role in writing and maintaining a comprehensive repository of principles. Such principle sets exist in a number of areas, although some rework may be required to ensure that those drafted earlier are well aligned with the overarching PBA approach. Further sets of principles will need to be drafted, and existing sets updated in response to changes in technologies and threat environments. At the time of writing, existing principle sets cover:

- Cloud security[4]
- Connected places[5]
- Cross domain solutions[6]
- Machine learning[7]
- Protocol design[8]
- Secure communications[9]
- Secure design[10]
- Software as a service[11]
- Supply chain[12]
- Zero trust architecture[13]

```
You, 7 months ago | 1 author (You)
import VueRouter from "vue-router";
import routes from "./routes/routes";
import store from "./store/index";
import vuexI18n from "vuex-i18n";
import enLangFile from "./lang/en";

// Set config file into the global variable
window.config = require("./vue.config");

    // Import bootstrap file
    require("./bootstrap");

    // Set Vue globally
    window.Vue = require("vue");

    require("./directives");
```

Image by Mohammad Rahmani / Unsplash

It is important to recognise that principles as they have been used in NCSC guidance cover two distinct areas: products and design. The object of product principles is a product or solution, while design principles in contrast apply to a design process. The former are of the type 'X should' while the latter are of the type 'you should'.

The principles serve a number of functions:

- They provide product developers with an understanding of the NCSC's expectations, informing product development and marketing.
- They provide a channel for dissemination of expertise from the intelligence community to organisations across all sectors.
- They provide a framework for assurance, whether this be a PBA evaluation, or a consumer.

An example provides clarity on what 'goal-based' standards look like.

> ***'A well-engineered CDS [Cross Domain Solution], with export functionality, should protect against the unauthorised export of information.'***

The principle states the outcome to be avoided: 'the unauthorised export of information'. What is expected of the CDS is that this outcome should be 'protected against'. Note that the principle does not say that 'unauthorised export of information should be impossible'. The semantics of 'protection' are crucial to the specification of the goal, as protection is relative to some particular type or level of threat. The principle thus implies a set of contextual judgements about the threat, and how protection may be assured. To provide the basis for the PBA evaluation, these judgements will need to be made, anticipating what is reasonable for a customer organisation (and in turn a reader will need to judge whether their context is typical).

[4] https://www.ncsc.gov.uk/collection/cloud/the-cloud-security-principles

[5] https://www.ncsc.gov.uk/collection/connected-places-security-principles/about-the-principles

[6] https://www.ncsc.gov.uk/collection/cross-domain-solutions

[7] https://www.ncsc.gov.uk/blog-post/introducing-our-new-machine-learning-security-principles

[8] https://www.ncsc.gov.uk/whitepaper/protocol-design-principles

[9] https://www.ncsc.gov.uk/guidance/secure-communication-principles

[10] https://www.ncsc.gov.uk/collection/cyber-security-design-principles

[11] https://www.ncsc.gov.uk/collection/saas-security/saas-security-principles

[12] https://www.ncsc.gov.uk/collection/supply-chain-security/principles-supply-chain-security

[13] https://www.ncsc.gov.uk/collection/zero-trust-architecture

[14] https://www.ncsc.gov.uk/collection/cross-domain-solutions/using-the-principles/protect-against-unauthorised-export

The formal text of the principles also contains a series of subsidiary normative claims, listed as 'defensive techniques.' Like the main principle, defensive techniques are normative claims about the product in question. Here we find a degree of variability. Some anticipate a qualitative argument, for instance:

> ***'Authorisation to release information (or 'release authorisation') should take into account business processes required to sanction information release.'***

Among the defensive techniques, we also find areas in which the NCSC does use prescriptive language. For instance:

> ***'Systems that need to enforce synchronous bi-directional communication (transactional query-response traffic such as HTTP requests or SQL queries) should correlate requests and responses, so that a response cannot pass without a corresponding request of the correct type.'***

Such prescriptions are best understood to represent areas in which details of the solution are understood to be so obvious, that for the sake of clarity it is better to state them explicitly.

An important point to make is that implementing the full set of defensive techniques, so that each of these subsidiary claims can be justified with confidence, does not, on its own, imply a strong argument for the overarching principle. The Assurance Principles and Claims document (APCs) provides an initial decomposition for each claim and are likely to do some of the work in showing how a demonstration of the implementation of the defensive techniques 'adds up' to a demonstration of the overarching claim. In some cases the nature of this overarching argument may depend on design specifics of the product in question, and in these cases the evaluator will need to adapt the argument (and hence the structure of the evaluation) to the product in question.

As PBA enters into operation, we can expect the growth of a penumbra of informal agreements and secondary texts, including the APCs, a 'case law' of previous assessments, templates for reports, as well as notes or memos from meetings and discussions with the NCSC. Evaluators will rely on such sources to inform their expectations of how the principles ought to be interpreted. Even in CPA, where standards were highly prescriptive, informal agreements about interpretations (for instance verbal discussions in meetings) can be found referenced as part of the report agreement process and could in certain situations override a literal interpretation of the security characteristics. With PBA, we would expect that similarly the practice of interpreting principles will be gradually developed, and evaluators will need to invest in building up familiarity with the process in order to effectively participate.

EVALUATION

The NCSC will oversee a group of ‘Trusted Labs’ which are authorised to conduct PBA evaluations.

In dialogue with the product developer (and potentially also with a customer organisation, depending on how the evaluation is commissioned), the evaluator will need to define how the evaluated product satisfies the relevant principles, and from this derive a plan detailing the suite of evidence required to construct a compelling assurance argument.

Technology security evaluation often requires collaboration between evaluator and product developer. The ‘Security Target’ document that defines the target of evaluation for Common Criteria evaluations is typically written by the product developer. PBA intensifies this relationship as there is additional room for interpretation. Interviews with people involved in a pilot of PBA evaluations of Cross Domain Solutions emphasised this point: a product developer for instance reported confusion around whether the operating system on which their product was running should have been in scope for the evaluation or not. Clear communication between developer and evaluator is essential to ensure the resulting evaluation meets expectations.

PBA evaluations will typically involve a combination of:

- Targeted testing of the product
- Reviews of documentation
(design documentation, manuals, operating procedures, development records)
- Reviews of source code
- Inspection of physical hardware
- Interviews with developers

It is important to recognise that the activities that are conducted under a PBA evaluation will be limited by time and cost, and evaluators will need to develop ways to scope evaluations that give good coverage. Scoping out what evidence is required to make a security assurance case also means scoping out what activities are unnecessary.

Where an evaluation is 'time-boxed' with a set budget, it becomes crucial to focus on the activities that provide the greatest contribution to overall confidence, and it would be beneficial to produce a clear record of areas in which additional evidence (beyond what could be gathered within the scope of the evaluation) would enhance the assurance case, for the benefit of customers. Given that the results of evaluation activities are not known in advance, the evaluator faces uncertainty in designing an optimal profile of activities.

A wider issue is determining 'how big' an evaluation ought to be. A lighter-weight, cheaper evaluation may result in an output that has a different profile of costs and benefits between developer, evaluator and customer compared with a more involved, more expensive assessment. It may be hard to standardise with a time-boxed approach because different products, even within the same general category, may have very different levels of complexity in design and build, may be targeted at use cases involving different levels of threat, and may be best evaluated with a different profile of activities. As a result, what counts as an 'equivalent' level of evaluation for different products may be hard to strictly specify and may remain fuzzy.

These issues about how decisions about scope and scale of PBA evaluations are made are likely to be an important area for future policy development. It is likely that some flexibility is given to vendors to define how exhaustive they wish the evaluation to be (it would be hard to justify constraining them in this respect), but with minimum levels defined in terms of effort, time, and/or coverage. In any case, all PBA reports will need to clearly state how scope and scale was determined.

v. It is recommended that the NCSC explore alternative approaches for setting the scope and scale of PBA evaluations, to understand how different ways of constraining these decisions affect the profile of benefits to different stakeholders.

Image by Google DeepMind / Unsplash

COMMUNICATION

There are two key components of the evaluation report in principles based assurance: 1) the overall assessment or 'executive summary', 2) the documentation of evidence. Evaluators will need to adopt a new style for both components.

The executive summary will need to anticipate the context of use. This may be challenging for products that can be used in several different ways, in different styles of architecture. An abridged, fictionalised PBA-style summary is given in Appendix 1, based on reports from the Cross Domain Solutions PBA pilot. These reports were constructed in close dialogue with NCSC stakeholders, and the example in the appendix shows some of the high-level features of reporting that were devised in that context. One notable feature is that the highest-level claim being reported about the product is not one of the principles at all. Indeed, it serves as a higher-level claim, a 'master' claim for the whole exercise. This is, roughly, that 'the product was found to function as expected'. The principles partially define those expectations. But this claim also allows space for the evaluator's discretion, for instance around other functional expectations that may interest a customer, or 'security promises' made by the developer that are not covered by the principles. In addition to knowing that the product can meet the goals the NCSC has specified, it is also important, as one interviewee put it, to know if the product 'does what it says on the tin'.

vi. It is recommended that PBA evaluations have an explicit overall goal of evaluating whether a product 'functions as expected', as this will provide evaluators with the flexibility to address any security-related factors of potential relevance to customers.

The summary also gives an overview of high-level risks, areas in which security problems could be introduced in the implementation of the product. In a CAE structure these would amount to areas in which the argument would depend on contextual details of how the product is implemented and cannot be supported based on evidence about the intrinsic characteristics of the product.

A further advantage of the CAE structure, to be adopted in future reports, will be in providing a format in which it is straightforward to highlight key areas in which further evidence could strengthen the assurance case, with particular attention to areas where this further evidence could be gathered by the customer organisation as part of implementation. For this to be effective it will be necessary to find appropriate ways to summarise quantitively or qualitatively 'how much' confidence is provided, and 'how much' more could be gained via further activities.

vii. It is recommended that PBA reports include a section of the summary detailing a prioritised list of areas in which further testing during implementation would provide the greatest 'added value' to the assurance case.

The reporting of evidence, which forms the main body of most assurance reports, will differ in PBA compared with previous report formats. The CAE framework will provide an organising structure. But a major departure from precent is a focus on unsupported claims rather than on the documentation of all areas of evidence (though it is possible that evaluators will choose to include the latter in an appendix). As evidence need not be reported in areas where claims are deemed to be met, there will be a greater emphasis on the evaluator's ability to assemble a tight assurance argument. Some evaluators may gain a reputation for being more cautious than others in their treatment of evidence, with knock-on effects on the interpretability and equivalence of reports.

EVALUATION GOVERNANCE

Adequate evaluation governance processes will be needed for the smooth running of PBA. This is because there will be flexibility in how evaluators run their activities and because each evaluation will be tailored to a product. There are risks of additional costs, for example if the product developer, the NCSC, or the customer organisation raises a problem with the acceptability of the output and requests additional work or rework. In CTAS, which has a similar flexibility, this risk is reduced by the fact that the NCSC approves the Evaluation Work Programme in advance of evidence gathering activities taking place. It is not expected that the NCSC will play a similar role in PBA. Contracts will likely set a milestone requiring approval of a similar document from the customer organisation and/or the product developer. However, success in managing this approval process will depend on these parties' understanding of the PBA philosophy and the CAE framework. Particularly in the early days, a lack of consensus on the goals of the process (for example, if a product developer were to believe that the process is oriented around ticking off the 'defensive techniques' for each principle) could prove a major obstacle to successful evaluations, and if unchecked such issues could hinder the long term adoption of the approach.

viii. It is recommended that Trusted Labs be encouraged to develop an internal capability with expertise in the CAE framework, to provide oversight of ongoing projects and to offer training to vendors/customers as part of evaluation engagements.

Governance arrangements, including contract management, relationship management, and NCSC engagement, will need to provide straightforward mechanisms to resolve the scenario where additional work is needed to fill 'gaps' in the evidence. These may occur due to oversight in planning. But they also may arise where evidence turns out to be weaker or less conclusive than expected in ways that could not be anticipated in advance. This kind of situation will account for some of the 'unsupported claims' in the report (other kinds of unsupported claims can arise because the product simply does not satisfy the principles on its own and requires additional security measures wrapped around it). In designing governance arrangements, it will be important for NCSC to put mechanisms in place to prevent evaluators being perceived as indulging in opportunistic behaviour, such as categorising claims as unsupported in order to gain additional paid evaluation work.

ix. It is recommended that PBA contracts include a standard framework for vendors to commission 'add on' evaluation work where this would enhance the assurance case. NCSC should monitor PBA evaluations during the early rollout and take steps to address any opportunistic behaviour (or perceptions of it).

Image by Google DeepMind / Unsplash

OPERATIONAL GOVERNANCE

It will be important to consider the 'afterlife' of principles based assurance, in other words what happens after a procurement decision is made, after a contract is signed for a product or service, and then throughout its operational lifetime. Note that a safety case, on which PBA is partially modelled, is expected to be a 'live' document, not a one-off 'point in time' exercise. A criticism of traditional product security assurance, voiced by several interviewees, is that 'point in time' evaluations become rapidly out of date. In an ideal world, there would be open sharing of information that impacts upon confidence in the security of products, so that all stakeholders can update their assurance cases and take appropriate action, but there are a number of practical barriers.

There may be particular technology categories (such as electricity and gas smart meters) where a PBA evaluation is mandatory for these products to be sold and used. In these cases, it is likely to also be mandatory that such products face ongoing 'renewal' or 'assurance maintenance' through periodic follow up exercises. But it is unlikely that all PBA evaluations will require mandatory follow-ups. A product vendor may seek a PBA evaluation as a one-off exercise to help substantiate claims made in marketing their product. A customer organisation may seek a PBA evaluation to provide input to a project assurance process or procurement decision. In these cases, making follow-up evaluations mandatory could discourage uptake of the scheme as it implies commitments to bearing future costs.

For some products, a later PBA evaluation will replace an earlier one. These might be performed by different Trusted Labs, but should retain (though potentially modifying) the established CAE structure. One issue to be determined is whether future reports may rely on evidence gathered in a previous assessment, or whether it is necessary to produce evidence afresh. Allowing justified reuse would imply that subsequent PBA evaluations would have the potential to progressively strengthen the assurance argument by focusing on areas de-prioritised in previous time-boxed exercises.

x. It is recommended that re-evaluations refer explicitly to any earlier reports conducted on that product, or product line, by any evaluator, so that evidence and changes in confidence level are clearly traceable, and to ensure that product vendors cannot game the system by seeking multiple evaluations. Some form of centralised registry of reports will likely be needed.

Where a newer PBA report is available, customer organisations already using that product will need to assess how any changes translate into changes in operational risk. To make this as straightforward as possible, and build the wider value of PBA, the NCSC could produce tools to help customer organisations to adopt CAE for their internal project assurance/security lifecycles.

xi. It is recommended that the NCSC puts in place a mechanism through which customer organisations can receive updates about the availability of newer PBA reports on products they use or that are under consideration.

In addition to the need for consideration of how updated PBA reports will be disseminated and impacted is the need for consideration of who, if anyone, is obliged to update PBA reports when new information is available. When digital technology is procured, contracts commonly include a service element, with the developer providing, at a minimum, security patches and updates, but also often some level of technical support, and extending all the way to handling elements of system integration and configuration. Whether the evaluator, or product developer, is responsible for 'patching' a PBA report if, for instance, a new vulnerability is found to affect it, is a challenging question to answer. Requiring developers to pay for evaluators to 'patch' the PBA reports for their products would, in theory, create the conditions for better ongoing assurance. It is not clear what the best solution would be, as this 'patching' of reports would increase uncertainty about future costs, which could discourage participation, while the alternative, leaving 'patching' of reports to the developers' discretion, could undermine customer confidence.

xii. It may be helpful for the NCSC to commission a study examining a sample of PBA reports after a time period has elapsed (one year, say), to determine whether new information that has come to light in that time (principally, new vulnerabilities) makes a material difference to the validity of the arguments.

Although the solution may need to be determined based on further evidence, one step that could be taken to reduce uncertainty for developers is to mandate that all PBA contracts provide a cost framework for future minor work, including impacting new information against the argument and issuing a minor update version of the report.

xiii. It is recommended that the NCSC ensures that PBA contracts include a built-in cost framework for future minor updates, such as impacting a newly discovered vulnerability against the CAE structure, with standardised costs and response time.

Another reason for uncertainty in this area of operational governance is the fact that the legal situation may change. A judge presiding on a well-publicised civil litigation case brought by a group of Postmasters against the Post Office for the harms inflicted by the Horizon IT system scandal made observations that may prove consequential in the years ahead (the case was settled in advance of a final judgement, so these observations do not have the legal weight they could have had). Standard UK contract law generally follows the principle of 'buyer beware', i.e., a seller has no obligation to make the buyer aware of deficiencies in a product (unless this is explicitly included in the contract), but in certain circumstances judges have interpreted contracts as 'relational contracts' in which obligations of good faith do apply (even if not explicitly stated). The judge in the abovementioned Horizon case found that the contract between Postmasters and the Post Office, under which the software services of Horizon were provided, was a relational contract, and among the obligations the Post Office should have met, was an obligation to candidly disclose problems (see Lloyd 2022 for a discussion). Given that most products assured under PBA will be sold along with a service element included, it may well be that an obligation to correct the record when it comes to assurance arguments becomes standard in future. In such an event, far from being a source of additional costs, PBA could provide a framework for easy dissemination of new information to stakeholders.

xiv. Whether product developers should be obliged to 'maintain' their PBA documentation is a difficult question. At least in the initial phase of PBA roll-out, any such requirement is likely to discourage adoption, but it should be a priority topic for longer term consideration.

Image by Google DeepMind / Unsplash

BUSINESS MODELS

A great many details of policy will affect the kinds of business model operated by Trusted Labs and product developers. The NCSC will need to be responsive in the early years to ensure that emerging issues are detected and corrected in a timely fashion.

The cost of performing PBA evaluations will depend on a number of factors. The first few evaluations are likely to be more costly given the increased risk of rework where the interpretation of the principles and the process is still uncertain. This could discourage smaller test lab outfits from taking part, and could present an ongoing barrier to entry. Early entrants are likely to benefit from additional NCSC support while the scheme is being implemented, and provision of onboarding support for new labs should be considered further down the line.

In search of efficiencies, evaluators will inevitably develop quasi-standard inventories of tests, and fragments of CAE structures, that they can reuse across evaluations. While this is not necessarily a problem, and to some extent it will be essential for achieving the kind of economies of scale necessary for widespread adoption, some degree of oversight will be necessary at a minimum to avoid overly templated approaches. One strategy that may be useful is in having some (or all) evaluations culminate in a 'live' presentation of results to ensure that evaluators are able to articulate how the argument they are making holds together, and to act as a soft mitigation against the risk of evaluators relying overly on existing template material for constructing new assurance arguments.

Apart from where PBA is mandatory, all firms using or delivering PBA evaluations will make the decision to take part based on assessments of cost and benefit. Up-front investment will be needed, in skills and in participation in evaluations, while many of the benefits will arise at a later date and will be subject to uncertainty, at least in the initial stages of implementation. Evaluators will need to train staff in CAE, develop processes and materials to support PBA evaluations, and engage with the NCSC. Vendors will need to pay for initial evaluations, to engage with evaluators in the evaluation process, and learn to interpret and communicate the results. Customer organisations will need to learn how to interpret and integrate the reports into assurance processes. Many benefits are deferred downstream: for the evaluators, a future stream of PBA evaluations, for the vendors, better communication with customers, and for customers, lower operational cyber security risk. However, whether those benefits will be realised depends upon other players in the market, the extent of their uptake and whether the approach is successful.

Because of this, the shape of the market and the success of the approach is likely to be strongly influenced by perceptions of the success or failure of PBA, by the prevailing narratives and sentiment in the professional community. Negative perceptions of PBA, that regard it (like CPA or Common Criteria) as having a limited lifespan, could easily act as 'self-fulfilling prophesies', discouraging participation and thus further entrenching those negative perceptions and their effects. The flip side of this is that the same positive feedback loop can be exploited by the NCSC to drive engagement: if the vision is clear to all stakeholders, if criticisms are met with appropriate and timely responses, and if the NCSC's long-term commitment to the approach is made clear, the business case for adoption, for all stakeholders, will likewise be clear, and if benefits are realised sooner, this would reinforce the positive narrative. One tool for achieving this is to integrate PBA more comprehensively with the NCSC's overarching strategic narrative. I close with a brief comment on how this could be done.

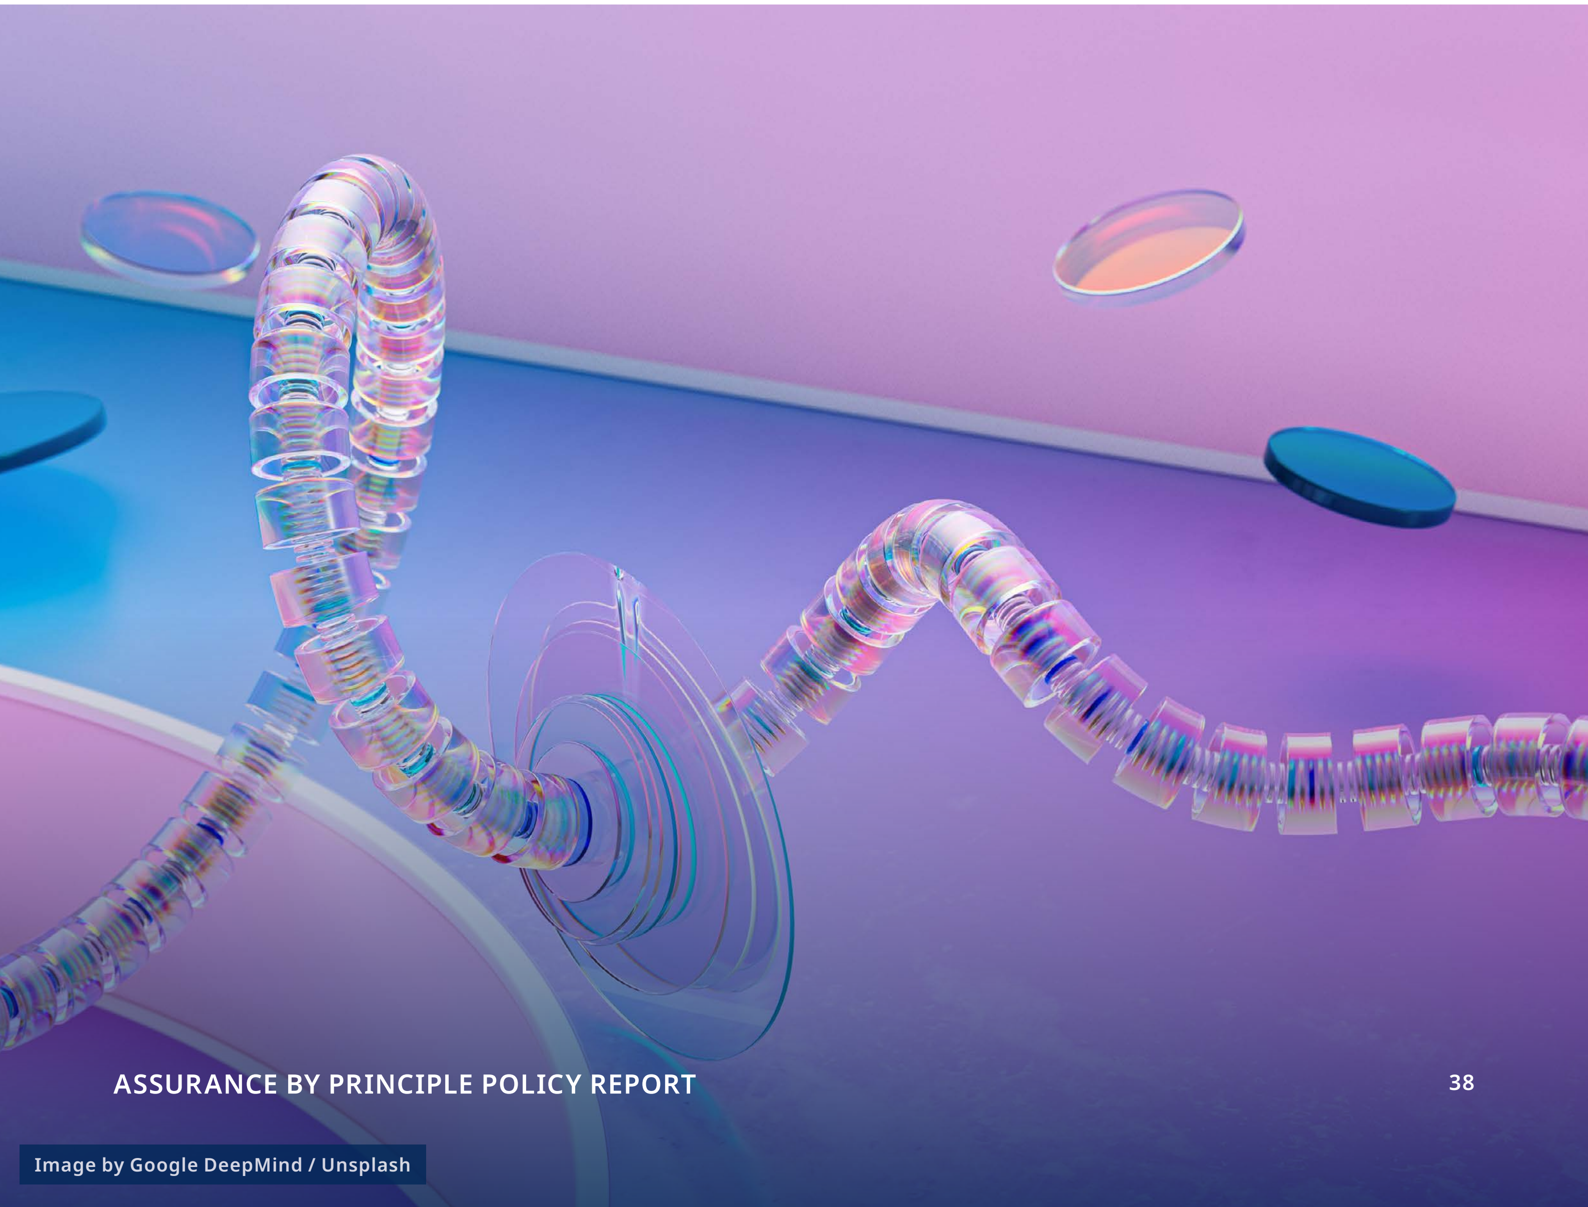

Image by Google DeepMind / Unsplash

The Future of PBA (towards inclusive security)

The development of the PBA approach has been driven by the goal of creating the conditions for better (more informed) technology choices. This is a specific case within what can be understood as a broader NCSC agenda to enhance the quality of reasoning and communication in cyber security in the UK. An adjacent challenge within this problem space is that of enabling 'inclusive security,' an umbrella term for cyber security practices with a focus on addressing a broader set of stakeholder needs. Just like PBA, inclusive security entails a communication challenge: we are in need of a common language, widely understood and easy to apply, in which to express how security objectives relate to the needs of different stakeholders. While PBA provides a framework with which to articulate whether and how particular products or systems satisfy key principles, inclusive security requires a framework to justify how those principles satisfy different groups' needs and goals.

One of the ways in which the NCSC could influence the narrative within the professional community, to ensure that the PBA approach is understood as a long-term policy commitment, is to identify adjacent problem spaces like that of inclusive security, and to emphasise synergies in key strategy texts. This would also help organisations to see the wider value of adopting PBA. In short, organisations that are able to construct explicit security arguments about confidence in technology choices are in a good position to take the next step, to make explicit how those choices relate to broader security goals, determinations of whose security is being prioritised. Integrating PBA as broadly as possible into the strategic narrative should be a priority for assurance communication going forward.

xv. To maximise stakeholder engagement with PBA, it is recommended that the NCSC integrate PBA into the wider strategic vision as comprehensively as possible. One way to do this is to identify adjacent problem spaces and strategic objectives, and synergies between them. This could include a vision of assurance-enabled inclusive security based on the explicit articulation of security reasoning, relating evidence to claims and claims to the security goals of diverse stakeholders.

Works Cited

Ashby, W. Ross. *"An introduction to cybernetics."* Chapman and Hall. 1956.

Bloomfield, Robin, and John Rushby. *"Assurance 2.0: A manifesto."* arXiv preprint arXiv:2004.10474. 2020.

Bishop, Peter, and Robin Bloomfield. *"A methodology for safety case development."* In Safety and Reliability, vol. 20, no. 1, pp. 34-42. Taylor & Francis, 2000.

CMA (Competition and Markets Authority). *"Statutory audit services market study."* Update paper. 2018.

Cullen, W.D. *"The public inquiry into the Piper Alpha disaster."* UK Department of Energy 1990.

Dekker, Sidney. *"Just culture: restoring trust and accountability in your organization."* Crc Press, 2018.

FRC (Financial Regulatory Authority). Audit Culture Thematic Review. 2018.

Hood Christopher. *"The tools of governments."* Macmillan. 1983.

Hood, Christopher. *"The Blame Game."* Princeton University Press, 2010.

Karpik, Lucien. *"Valuing the unique: The economics of singularities."* Princeton University Press, 2010.

Leveson, Nancy G. *"The use of safety cases in certification and regulation."* MIT White Paper, 2011 http://sunnyday.mit.edu/SafetyCases.pdf

Lloyd, Ian. *"Lessons on Robustness and Reliability of Software Solutions from the Horizon System for UK Post Offices."* Computer Law Review International 23, no. 1. 2022: 6-12.

Spencer, Matt. *"Characterising assurance: scepticism and mistrust in cyber security."* Journal of Cultural Economy 2022: 1-16.

Toulmin, S. *"The Uses of Argument"* (2nd ed.). Cambridge: Cambridge University Press. 2003 [1958].

Tsoukas, Haridimos. *"The tyranny of light: The temptations and the paradoxes of the information society."* Futures 29, no. 9. 1997: 827-843.

Image by Google DeepMind / Unsplash

Appendix 1: Fictionalised PBA Report Executive Summary

EXECUTIVE SUMMARY

This report summarises findings of the evaluation of the Example Systems' File Transfer Pro. This product provides a means for secure file transfer across network domains of different classifications. It is available as a physical appliance or as a virtual application.

The product was tested against the NCSC's 13 Cross Domain Solutions Security Principles by Cyber Labs, an ISO 17025 certified test lab. Following an initial review with the developer, a test approach was drawn up that could cover each of the principles. The testing that was carried out included:

- *interviews with developers*
- *reviews of technical documentation*
- *reviews of process documentation and records of development activities*
- *penetration testing of the product in a test rig designed specifically for this purpose*
- *physical examination of the product*

The evaluation was time-limited, and the test plan and scope have been designed in order to maximise coverage within this constraint. Details of assumptions and coverage are included in section X.X below.

When considering cyber security threats, note that advanced threat actors may be capable of investing very extensive resources into the discovery and exploitation of vulnerabilities; time-limited testing cannot therefore provide comprehensive assurance against this category of threat.

Overall, the product was found to function as expected, and to provide good security controls around the transfer of files between domains. It was not possible during the evaluation to circumvent the system, and pass data inappropriately.

This is a very flexible and configurable product and certain aspects of the technical documentation were found to be ambiguous, creating the risk of misunderstandings or misconfigurations during the implementation of the end-to-end solution.

It was also found to be possible to create a Denial of Service issue by sending large volumes of messages, leading to unavailability of the solution. Rate limiting should be implemented as part of the end-to-end solution.

By design the system does not check incoming files for malware. Security architects should consider including this as part of the end-to-end solution.

There are trade-offs to be made between usability and security. With default configuration, several avenues exist for an 'insider threat' attack, including data exfiltration. While the system can be configured to reduce these avenues, there is a trade off in terms of usability, and many organisations will find it more appropriate to rely on additional controls to mitigate this residual risk.

Maintaining the system over the long term will largely be a matter of updating configuration files and applying update patches. Due to the nature of the system, configuration is complex, and the security of the end-to-end solution will depend upon modifications to the configuration being set appropriately. Systems administrators will need training, and maintenance of configuration must be subject to a rigorous review, test and release process.

[fictionalised, but based on combining various observations made in the reports]