Protecting legitimate cyber security activity

A proposal for a principles-based framework for the application of a statutory defence under a reformed Computer Misuse Act

October 2021

Introduction

Section 1 of the UK's Computer Misuse Act 1990 currently prohibits unauthorised access to computer material.

The CyberUp Campaign argues that this binary way of distinguishing legal and illegal activity is outdated, because it restricts UK-based cyber security professionals' ability to undertake essential vulnerability and threat intelligence research that, at times, necessarily takes place without prior, or sufficiently clear, authorisation.

To tackle this, the CyberUp Campaign has proposed reforming the Computer Misuse Act to include a statutory defence that would allow cyber security professionals to defend those activities that constitute unauthorised access, but which are performed in the public interest, to improve cyber security or detect or prevent crime.

While the concept of statutory defences is a well-established legal principle, we have, perhaps unsurprisingly, received a number of queries seeking clarification on how such a defence would work in practice, minimise the risk of abuse and strike the right balance between protecting the cyber security ecosystem and prosecuting criminals effectively.

In the spirit of facilitating constructive policy debate, and to support the more concrete discussion of the practicalities of a statutory defence in the cyber context, the CyberUp Campaign proposes a principles-based framework for applying the statutory defence which we set out in detail in this document.

In developing the framework, we consulted with a number of cyber security professionals, senior executives at cyber security companies, in-house legal professionals and legal academics, all of whose insights have proved invaluable.

We reflect on their feedback in the Discussion section of this document and explain clearly which changes we accepted, and why some suggestions are not included in the final iteration of the framework.

To avoid confusion, misunderstanding and crossed wires, we believe that it is important that we also explain clearly how we envisage the principles-based framework to be used, and make explicit any other assumptions we have made in developing it:

How the principles-based framework should be used in practice and other assumptions:

1. We aim to clarify how a statutory defence would operate in the cyber context. The principles-based framework seeks to demonstrate that courts are capable of successfully and consistently applying an assessment of whether an act of unauthorised access was defensible, and thereby inform an evolving understanding of what constitutes legitimate conduct in cyber space.

2. We do not intend for the details of the framework to be included in primary legislation as part of a reformed Computer Misuse Act. Instead, we advocate for updated legislation to mandate the courts to "have regard to" Home Office or Department for Digital, Culture, Media and Sport (DCMS) guidance on applying a statutory defence that would, ideally, be based on the framework we propose. We would also want courts to be upskilled on cyber matters over time. As is standard practice in criminal law, courts could seek evidence from independent expert bodies, such as, in this case, the UK Cyber Security Council, to understand technical details before them in the course of their work.

3. We deliberately focus on principles rather than a list of defined activities to future-proof our approach from the outset. Our aim is to avoid being too prescriptive in a way that will date any guidance too quickly, and so to provide a framework that is able to be used consistently as technology, capabilities, threats, working practices and societal norms evolve.

4. We focus on the reform of defences rather than narrowing offences to apply only to 'bad actors'. We do so because we believe that a statutory defence offers an essential degree of flexibility to adapt to a changing world.

5. At present, the development of a principles-based framework is necessarily abstract. However, in reality, where the framework guides the application of a statutory defence, we can assume that a concrete act of unauthorised access has taken place and is being assessed by the courts. That means that:

    a. We assume the defence to be relevant only in those instances where gaining authorisation has been impossible to obtain, which has necessarily informed our view of which principles are appropriate to include. Of course we take the view that it is always desirable for a cyber security professional to seek authorisation where there is no reason to believe it might be unreasonably withheld. However, there will be instances where a system owner cannot be (accurately) identified, or doing so would unjustifiably delay any relevant acts in question.

    b. We take the view that, on balance, it is right that there is a reversed legal burden, so that whoever undertook the act of unauthorised access has to prove that doing so was defensible. The principles-based framework will guide the defendant on the quality of evidence they need to provide, and offer objective standards against which the courts will judge any such evidence and, ultimately, conduct.

The proposed principles-based framework

We propose four principles against which an act of unauthorised access should be assessed to determine if a new statutory defence should apply to render the act defensible under a reformed Computer Misuse Act.

As will become apparent, there is necessary overlap across these principles. We argue that they should be applied comprehensively to any act in question to arrive at a fair and meaningful determination of its defensibility.

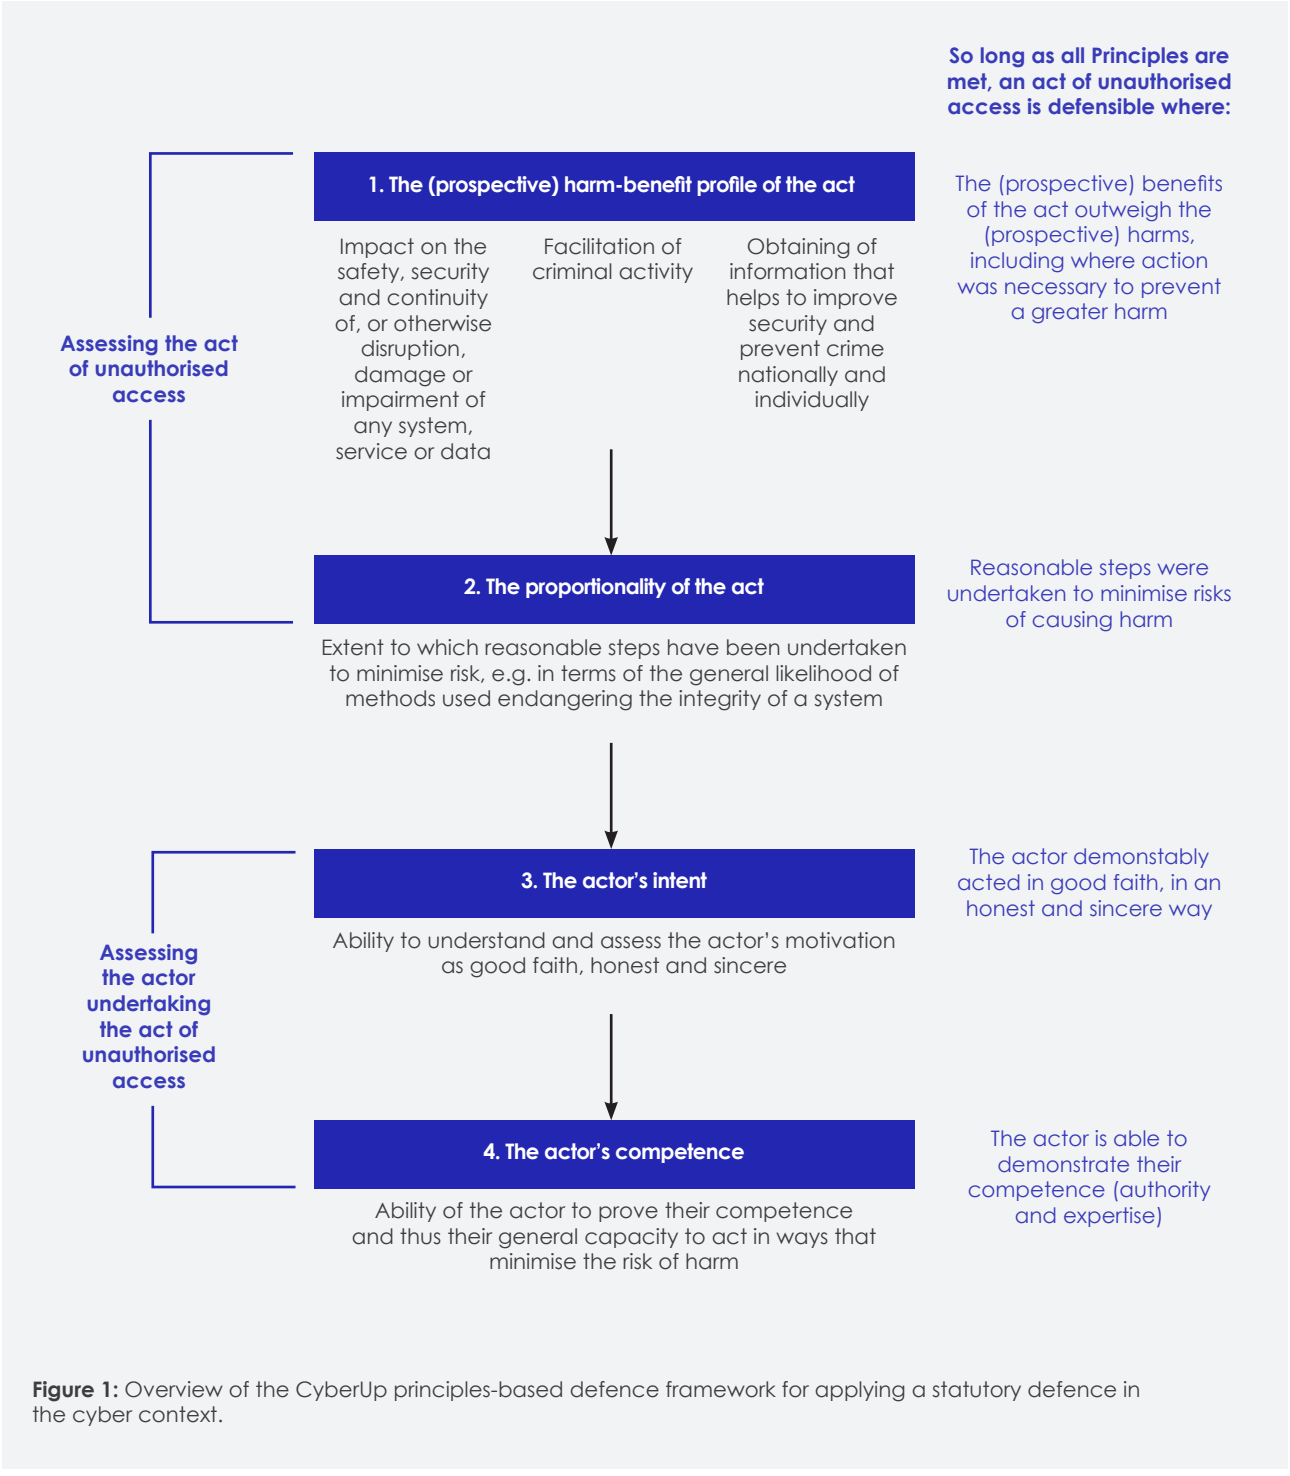

Figure 1: Overview of the CyberUp principles-based defence framework for applying a statutory defence in the cyber context.

Principle 1: The (prospective) harm-benefit profile of the act

"Principle 1 states that an act of unauthorised access is defensible where the (prospective) benefits of the act outweigh the (prospective) harms, including where action was necessary to prevent a greater harm, so long as the other Principles are also met."

One of the most important factors in considering whether an act is defensible is considering the (prospective) harms of the act in question versus its (prospective) benefits.

While cyber security activities arguably often inherently involve risk and uncertainty, we believe they can be placed on a (prospective) harm-benefit spectrum, as illustrated below.

Figure 2: The CyberUp (prospective) harm-benefit spectrum of acts of unauthorised access.

Harms include:

- The infringement of the rights of another person/organisation, such as impacting the safety, security and continuity of systems, services or data
- Economic/financial harms and/or harms to public health, such as facilitating criminal activity
- Geo-political risks

Benefits include:

- Improved cyber resilience and reduced cyber crime, for individual system owners, and at national scale, such as through obtaining information or actionable intelligence that helps improve security or prevent crime

We understand that, in practice, where a statutory defence is applied, an act will have been committed that has resulted in real-life outcomes which might have caused actual harms and benefits.

However, it is right to state that criminal defences should be judged by focusing on the point of action rather than on their consequences, or on "what happens next". Ultimately, any act in question need not have resulted in a quantifiable or qualified outcome to be defensible. For example, it would still be defensible to act even where that act, unforeseeably, causes financial damage, whereas an act might be indefensible where a benefit arises, purely and unforeseeably, by chance.

Nevertheless, we do also believe that the successful application of a defence to a criminal charge would not necessarily undermine potential civil liabilities.

Principle 2: The proportionality of the act

"Principle 2 states that an act of unauthorised access is defensible where reasonable steps were undertaken to minimise risks of causing harm, so long as the other Principles are also met."

The proportionality of the act considers whether the same outcome could have been achieved by an act that involved less risk, and whether the actor took all reasonable steps to minimise the risk of harm involved in the act.

If it is deemed that the same objective could have been achieved in a way that had a lower chance of producing the harms set out above, then this would count against a researcher's eligibility for a defence.

We include this principle as we believe that it will encourage cyber security professionals to act in a way that minimises unnecessary risk of harm.

Principle 3: The actor's intent

"Principle 3 states that an act of unauthorised access is defensible where the actor demonstrably acted in good faith, in an honest and sincere way, so long as the other Principles are also met."

We have long argued that one of the problems with the Computer Misuse Act is that Section 1 prohibits unauthorised access to computer material irrespective of whether this step was taken by a cyber security researcher with noble intentions, or by an internationally renowned gang of cyber criminals for nefarious ends.

We understand concerns that proving intent will always be difficult but believe it is possible to demonstrate an actor's good faith motivations, for example where they acted in the honest belief that their conduct was reasonable to prevent crime, or protect a system, and where they can show that no personal profit or gains were made as a result of their conduct.

In addition, the CyberUp Campaign has published research[1] that shows an accepted consensus of what does and does not constitute 'good faith' activity which we propose to build upon further. In our proposed framework, intent would be a necessary but insufficient condition of defensible acts of unauthorised access. For example, while an actor might have the 'right' intent, their actions might still be deemed disproportionate, or foreseeably have caused damage to a system, making it ultimately indefensible.

For the purposes of establishing whether an actor had the correct intent, we believe that an actor's conduct following the action should also be assessed. This includes, for example, whether an actor disclosed a vulnerability, shared information with law enforcement, or cooperated fully with relevant organisations in a timely manner. An actor's subsequent actions are generally not relevant to criminal liability (though they may be relevant at sentencing, or in decisions to prosecute) but they can be an important factor in establishing an actor's intentions at the time they performed the act in question.

[1] https://www.cyberupcampaign.com/news/new-research-public-bodies-are-already-defining-good-faith

Principle 4: The actor's competence

"Principle 4 states that an act of unauthorised access is defensible where the actor is able to demonstrate their competence (authority and expertise), so long as the other Principles are also met."

We argue that there is a series of factors that ought to serve as a proxy to determining an actor's competence, and thus their general capability of acting in a way that minimises the risk of harm to the greatest extent possible. These include:

- An actor's level of qualification, certification or accreditation
- An actor's membership of a professional organisation and compliance with a code of ethics
- An actor's professional capacity during the act in question – whether an actor was acting under commercial, academic research or other contracts, or participating in a bug bounty or other kind of product attack challenge programme
- An actor's prior track record of work, research and investigations – self-taught ethical hackers may not have any qualifications or be affiliated with any accrediting body, but this doesn't necessarily mean that the defence shouldn't apply to them
- An actor's previous associations – similarly, successful participation in schemes like bug bounty programmes should count towards competence

We do not propose that qualification, certification or accreditation act as a precondition for eligibility for a statutory defence but argue that the existence of these proficiencies, alongside the other factors set out here, could all act as supporting factors in the actor's favour when determining the defensibility of their acts.

Indeed, some cyber criminals may be highly qualified, and may even have a track record of successful research that led to actionable intelligence or otherwise improved general cyber resilience, but it will be for the courts to determine, on the basis of the remaining principles, whether the defence should ultimately apply.

Discussion

We mention in the Introduction that we consulted on the principles-based framework with cyber security and legal stakeholders.

At a high level, while there was general consensus on the usefulness and appropriateness of the framework we propose, stakeholders' views on individual principles ranged from concerns they might offer "too much freedom for abuse" and would have to be stricter, to worries they did "not go far enough in giving a good enough shield [and] safe haven as needed."

We believe that this demonstrates that what we propose ultimately strikes the right balance. Nevertheless, we have incorporated stakeholders' comments and made changes to our initial draft framework:

- We removed a reference to the cyber security sector being advantaged economically as a benefit from an act of unauthorised access. While this has been a key argument for the CyberUp Campaign in proposing the introduction of a statutory defence, we agree that it will be a desirable consequence, but not appropriate for courts to weigh this as a consideration when applying a statutory defence.

- The initial draft of the document framed Principle 1 in terms of the risk and reward of the act in question. We received feedback that the language of harm and benefit is more appropriate in the legal context. We do believe that there is reasonable consensus around the harms and benefits arising from certain acts and techniques and where, therefore, they sit on the prospective harm-benefit spectrum.

- We initially labelled Principle 4 as relating to an actor's qualifications, but have revised this into the broader category of competence. This responds to valuable comments that some of the most skilled cyber security professionals are entirely self-taught, and this fact shouldn't impede their eligibility for a defence, and that, as one consultation response reminded us, "some criminals do have certificates, they just do crime in their spare time."

- Our initial draft included an actor's 'subsequent conduct' as a 5th Principle. We received feedback that, in criminal law, an actor's subsequent conduct is not relevant to criminal liability, though may be for sentencing and a decision being taken on whether to pursue a prosecution, and can be an important factor in establishing intent. For this reason, we have included this as part of Principle 2. This also reflects similar feedback we received from cyber security professionals.

There were comments and feedback that we have chosen not to include in the final iteration of the proposed framework in this document. Here we set out those concerns and our rationale for not including them:

- Many respondents commented on the overlap between the principles, or argued that certain principles should be prioritised over, or collapsed into others. Some argued, for example, that intent and proportionality should be given priority, because they could be proven to the criminal standard of proof. Moreover, there were comments that an actor's prior track record of good faith security research should count towards intent. There is also clearly a relationship between proportionality and competence.

    - Our view is that the fact many of these principles overlap, and that information relevant to one principle will also be a relevant factor in one or more of the other principles, is a strength rather than a weakness of the approach we have set out. Having four principles offers an extra safeguard against someone who has carried out a truly criminal act successfully having a defence applied to them, because it is possible to more easily bring the criminal character of the act in question to light which might be obscured by a reduced or more limited set of principles.

- We resisted the suggestion to include a list of typical activities or techniques that would be deemed as broadly acceptable under a reformed Act. We set out in the Introduction that we believe that our approach will allow a defence to be applied consistently over time, providing a degree of future-proofing to the legislation and its underpinning guidance. That said, we do believe that at any one time – be it a case in, say, 2022 under a reformed Act or in 15 years of time – there will always exist a degree of consensus about the prospective 'benefit-harm' profiles of different acts and techniques, and their broader defensibility, and a court will be able to draw upon this expert consensus in the application of Principle 1. We plan to undertake further research work to establish with more precision where that objective expert consensus currently lies, but for the purposes of this document we have sought only to set out future-proofed principles for the application of a statutory defence.

- There were also many comments about the need to look at the offences of the Computer Misuse Act. This is understandable. The natural corollary of adding weight to the benefit or harm of an action is that reference to these factors ought to be captured in some way by the offences the Act establishes. Nevertheless, we have chosen with this piece of work to focus specifically on a framework for the implementation of a statutory defence, because:

    - We must acknowledge the current legal position we are working from. Rather than creating a new offence to target new conduct, we are recommending that conduct currently caught within the law should be excused. This leads to a natural concern that any change should not risk 'bad actors' escaping liability – a concern we meet regularly in meetings with stakeholders, and we have seen in several responses to this consultation.

    - Focusing on the reform of defences (rather than offences) provides additional legal protections to reassure against such worries. As we outline in the Introduction, these include a reversed legal burden for the defendant, and the use of objective standards in assessing the defendant's conduct, rather than their subjective perceptions alone, informed by the principles discussed in our consultation.

- There were also some comments about how the defence would apply to actors in different jurisdictions, who might have different views of what constitutes an accepted interest. The argument is that a Russian criminal might well argue that their act of unauthorised access offers a benefit to their criminal enterprise, or even to the Russian state. We believe this concern is misplaced. In UK criminal law, a court will apply the standard to what a 'reasonable observer' in the UK would deem a benefit of the act. This would very likely result in the benefit being determined in terms like those we have set out in this document – namely, the UK's, or UK entities' improved cyber resilience and reduced cyber crime through actionable intelligence and/or improved security systems.

- One respondent suggested they would prefer to see a system in which cyber security professionals took requests to either law enforcement agencies or intelligence services to have proposed work pre-cleared with set boundaries and periodic reporting. Our concern is that a system of pre-clearance would have the effect of stifling many of the upsides that reforming the Act, according to the approach we have set out, would bring – specifically, significantly improved cyber resilience. One of the main reasons we are proposing a defence is to allow cyber security professionals to have unauthorised access in some scenarios where that authorisation cannot be secured. Removing one requirement for an authorisation process and replacing it with another in the form of a pre-clearance requirement, which will still be limited by the time it will take to jump through the various hurdles, is not the right approach. Our view is that, over time with case law, and ideally with clear guidance from prosecutors, the boundaries of legal conduct will be clear enough that the high degree of oversight that is sought by those who prefer a system more tightly regulated by government bodies is not necessary. That said, the CyberUp Campaign has previously mentioned that it does believe that anyone undertaking an act of unauthorised access should keep documents and logs of the act and any related activities, and we still believe that this is appropriate.

Conclusion and next steps

This document has set out a principles-based framework for the application of a statutory defence in the cyber context. The purpose of this project has been to demonstrate that courts are capable of successfully and consistently applying an assessment of whether an act of unauthorised access was defensible, and thereby inform an evolving understanding of what constitutes legitimate conduct in cyber space.

The concerns that some still harbour about the potential for a reformed Computer Misuse Act being open for abuse are understandable, and indeed need to be aired and revisited as the reform process continues to move forward to ensure that any legislation that makes its way on to the statute books cannot be manipulated for nefarious ends. But, in establishing a set of principles that would guide the application of a statutory defence, this document is a starting point in positing that it is possible for this well-established legal concept to be made to work in a cyber context. Worries about the inclusion of a statutory defence leading to abuse are not a reason to press the pause button on efforts to reform this outdated legislation, but rather, we argue, an important part of scrutiny on the path towards reform.

In putting forward these principles, we have been clear to emphasise the need for the legislation and guidance to be able to evolve over time, and to resist setting out which activities are, according to present expert understanding of their harm-benefit profile, legitimate and therefore, assuming the other principles are met, ought to be legal instances of unauthorised access. We understand, though, that it would be useful for policy makers to have an understanding of what that current consensus looks like. As part of its upcoming work programme, the CyberUp Campaign intends to produce another piece of work drawing out that consensus using existing academic literature and further consultation with cyber security stakeholders.