

Deterrence in cyberspace

Spare the costs, spoil the bad state actor:
Deterrence in cyberspace requires consequences

Chris Painter



Policy Brief
Report No.4/2018

About the author

Chris Painter is a distinguished non-resident fellow at ASPI's International Cyber Policy Centre. He is a globally recognized leader and an expert on cybersecurity, cyber policy, cyber diplomacy and combatting cybercrime. He has been on the vanguard of US and international cyber issues for over twenty-five years—first as a leading federal prosecutor of some of the most high-profile cybercrime cases in the country, then as a senior official at the Department of Justice, the FBI, the National Security Council and finally as the world's first top cyber diplomat at the State Department. He has helped drive, initiated or been involved in virtually every major US cyber policy for over a decade and has created innovative new organizations and approaches to deal with threats and take advantage of opportunities in cyberspace. Among other things, he currently serves as a Commissioner on the Global Commission for the Stability of Cyberspace and is a member of the Board of Directors for the Center for Internet Security.

What is ASPI?

The Australian Strategic Policy Institute (ASPI) was formed in 2001 as an independent, non-partisan think tank. Its core aim is to provide the Australian Government with fresh ideas on Australia's defence, security and strategic policy choices. ASPI is responsible for informing the public on a range of strategic issues, generating new thinking for government and harnessing strategic thinking internationally.

ASPI International Cyber Policy Centre

The ASPI International Cyber Policy Centre's mission is to shape debate, policy and understanding on cyber issues, informed by original research and close consultation with government, business and civil society.

It seeks to improve debate, policy and understanding on cyber issues by:

1. conducting applied, original empirical research
2. linking government, business and civil society
3. leading debates and influencing policy in Australia and the Asia-Pacific.

We thank all of those who contribute to the ICPC with their time, intellect and passion for the subject matter. The work of the ICPC would be impossible without the financial support of our various sponsors but special mention in this case should go to the Australian Computer Society (ACS), which has supported this research.

Chris Painter's distinguished visiting fellowship at ASPI's International Cyber Policy Centre was made possible through the generous support of DFAT through its Special Visits Program. All views expressed in this policy brief are the authors.

Important disclaimer

This publication is designed to provide accurate and authoritative information in relation to the subject matter covered. It is provided with the understanding that the publisher is not engaged in rendering any form of professional or other advice or services. No person should rely on the contents of this publication without first obtaining advice from a qualified professional person.

ASPI

Tel +61 2 6270 5100

Fax + 61 2 6273 9566

Email enquiries@aspi.org.au

www.aspi.org.au

www.aspistrategist.org.au

 facebook.com/ASPI.org

 [@ASPI_ICPC](https://twitter.com/ASPI_ICPC)

www.aspi.org.au/icpc/home

© The Australian Strategic Policy Institute Limited 2018

This publication is subject to copyright. Except as permitted under the *Copyright Act 1968*, no part of it may in any form or by any means (electronic, mechanical, microcopying, photocopying, recording or otherwise) be reproduced, stored in a retrieval system or transmitted without prior written permission. Enquiries should be addressed to the publishers.

First published June 2018.

Cover image: Global international connectivity graphic © spainter_vfx/iStock.



Australian Government
Department of Foreign Affairs and Trade

Deterrence in cyberspace

Spare the costs, spoil the bad state actor:
Deterrence in cyberspace requires consequences

Chris Painter

Contents

Foreword	03
What's the problem?	04
What's the solution?	05
The challenge of attribution	05
Naming and shaming	06
General considerations	06
Collective action	07
Using and expanding the tools in the toolkit	07
Recommendations	09

Foreword



In the past three years, barely a week has gone by without a report of a critical cyberattack on a business or government institution. We are constantly bombarded by revelations of new ransomware strains, new botnets executing denial of service attacks, and the rapidly expanding use of social media as a disinformation and propaganda platform.

Perhaps most alarmingly, a great many of these attacks have their origin in the governments of nation states.

In the past decade we have moved well beyond business as usual signals intelligence operations. Some of the largest malware outbreaks in recent years, such as NotPetya and WannaCry, had their origins in state-run skunkworks.

Cyberattacks initiated by nation states have become the new normal, and countries including Australia have struggled with the challenge of how to respond to them. Far too often they're considered a low priority and met with a shrug of the shoulders and a "What can you do?"

In this paper, Chris Painter offers us a way forward. Chris presents a reasonable framework for deterrence, a way that we as a nation can help limit the deployment of cyberwarfare tools.

His recommendations are designed to properly punish bad actors in a way that discourages future bad behaviour. They're modelled on actions that have worked in the past, and serve, if not as a final solution, at least as a starting point for us to scale back on the increasing number of state-sponsored cyber attacks.

Most importantly, these actions aren't just to the benefit of the state—they will allow us to better protect private citizens and companies that all too often get caught in the cyberwarfare crossfire. To put it simply, if we can ensure there are costs and consequences for those who wrongly use these tools to wreak damage, bad actors might start thinking twice before engaging in this destructive behaviour.

Yohan Ramasundara
President, Australian Computer Society

What's the problem?

Over the past few years, there's been a substantial increase in state attacks on, and intrusions into, critical information systems around the globe—some causing widespread financial and other damage.¹ They have included:

- attacks by North Korea on Sony Pictures in 2014
- widespread Chinese theft of trade secrets and intellectual property
- Russian state-sponsored interference in the US elections
- North Korea's sponsorship of the WannaCry ransomware worm that caused, among other things, a meltdown of the UK's National Health System
- the Russian-sponsored NotPetya worm that caused tens of millions of dollars of damage and disruption around the world.

The pace and severity of these attacks show no sign of declining. Indeed, because there have usually been little or no consequences or costs imposed on the states that have taken these actions, they and others have little reason not to engage in such acts in the future.

The US, Australia and many other countries have spent years advancing a framework for global stability in cyberspace. This framework comprises:

- the application of international law to cyberspace
- acceptance of certain voluntary norms of state behaviour in cyberspace (essentially, voluntary rules of the road)
- the adoption of confidence and transparency building measures.

Although much progress has been achieved in advancing this framework, the tenets of international law and norms of state behaviour mean little if there are no consequences for those states that violate them. This is as true in the cyber world as in the physical one. Inaction creates its own norm, or at least an expectation on the part of bad state actors that their activity is acceptable because there are no costs for their actions and no likely costs for future bad acts.

Individually as countries and as a global community, we haven't done a very effective job of punishing and thereby deterring bad state actors in cyberspace. Part of an effective deterrence strategy is a timely and a credible response that has the effect of changing the behaviour of an adversary who commits unacceptable actions. Although there are some recent signs of change, in the vast majority of cases the response to malicious state actions has been neither timely nor particularly effective. This serves only to embolden bad actors, not deter them. We must do better if we're to achieve a more stable and safe cyber environment.

What's the solution?

It is a well-worn and almost axiomatic expression that deterrence is hard in cyberspace. Some even assert that deterrence in this realm is impossible. Although I don't agree with that fatalistic outlook, it's true that deterrence in cyberspace is a complex issue. Among other things, an effective deterrence framework involves strengthening defences (deterrence by denial); building and expanding the consensus for expectations of appropriate state behaviour in cyberspace (norms and the application of international law); crafting and communicating—to potential adversaries, like-minded partners and the public—a strong declaratory policy; timely consequences, or the credible threat thereof, for transgressors; and building partnerships to enable flexible collective action against those transgressors. Although I'll touch on a couple of those issues, I'll focus here on imposing timely and credible consequences.

The challenge of attribution

One of the most widely cited reasons for the lack of action is the actual and perceived difficulty in attributing malicious cyber activity. Unlike in the physical world, there are no launch plumes to give warning of the location of the origin of a cyberattack, and sophisticated nation-states are adept at hiding their digital trail by using proxies and routing their attacks through often innocent third parties. But, as recent events illustrate, attribution, though a challenge, is not impossible. Moreover, attribution involves more than following the digital footprints; other forms of intelligence, motive and other factors all contribute to attribution. And, ultimately, attribution of state conduct is a political decision. There's no accepted standard for when a state may attribute a cyberattack, although, as a practical, political and prudential matter, they're unlikely to do so unless they have a relatively high degree of confidence. Importantly, this is also true of physical world attacks. Certainly, a state doesn't require 100% certainty before attribution can be made or action taken (as some states have suggested). Whether in the physical or the cyber world, such a standard would practically result in attribution never being made and response actions never being taken.

Although attribution is often achievable, even if difficult, it still seems to take far too long—at least for public announcements of state attribution. Announcing blame, even if coupled with some responsive actions, six months to a year after the event isn't particularly timely. Often by that point the impact of the original event has faded from public consciousness and so, too, has the will to impose consequences. Part of this delay is likely to be due to technical difficulties in gathering and assembling the requisite evidence and the natural desire to be on solid ground; part is likely to be due to balancing public attribution against the possible compromise of sources and methods used to observe or detect future malicious activity; but part of it's probably due to the need to summon the political will to announce blame and take action—particularly when more than one country is joining in the attribution. All of these cycles need to be shortened.

Naming and shaming

Public attribution of state conduct is one tool of deterrence and also helps legitimise concurrent or later responses. The US, the UK, Australia and other countries came together recently to attribute the damaging NotPetya worm to Russia and, a few months ago, publicly attributed the WannaCry ransomware to North Korea. This recent trend to attribute unacceptable state conduct is a welcome development and should be applauded.² It helps cut through the myth that attribution is impossible and that bad state actors can hide behind the internet's seeming anonymity.

However, public attribution has its limits. Naming and shaming has little effect on states that don't care if they're publicly outed and has the opposite effect if the actor thinks their power is enhanced by having actions attributed to them. In the above two cases, it's doubtful that naming and shaming alone will change either North Korea's or Russia's conduct. Public attribution in these cases, however, still serves as a valuable first step to taking further action. Indeed, in both cases, further actions were promised when public attribution was made. That raises a couple of issues. First, those actions need to happen and they need to be effective. President Obama stated after the public attribution to North Korea in relation to the Sony Pictures attack that some of the response actions 'would be seen and others unseen'. A fair point, but at least some need to be seen to reinforce a deterrent message with the adversary, other potential adversaries and the public at large. The other issue is timing. The public attribution of both WannaCry and NotPetya came six months after the respective attacks. That delay may well have been necessary either for technical reasons or because of the work required to build a coalition of countries to announce the same conclusion, but attribution that long after the cyber event should be coupled with declared consequences—not just the promise that they're to come. Some action did in fact come in the NotPetya case about a month after public attribution, when the US sanctioned several Russian actors for election interference, NotPetya and other matters. That was a very good start but would be even more effective in the future if done when the public attribution occurs. Action speaks louder than attribution alone, and they must be closely coupled to be effective.

General considerations

A few general considerations apply to any contemplated response action to a cyber event. First, when measures are taken against bad actors, they can't just be symbolic but must have the potential to change that actor's behaviour. That means that one size does not fit all. Different regimes hold different things dear and will respond only if something they prioritise or care about is affected. Tailored deterrence strategies are therefore required for different states.³ For example, many have opined that Russia is more likely to respond if sanctions are targeted at Putin's financial infrastructure and that of his close elites than if simply levied in a more general way. Second, the best response to a cyberattack is seldom a cyber response. Developing cyber tools and having those tools as one arrow in the quiver is important, but other responses will often be more effective. Third, the response to a cyber event shouldn't be approached in a cyber silo but

take into account and leverage the overall relationship with the country involved. The agreement that the US reached with China that neither should use cyber means to steal the trade secrets and intellectual property of the other to benefit its commercial sectors wouldn't have come about if widespread cyber-enabled intellectual property theft was seen only as a cyber issue. Only when this problem was seen as a core national and economic security issue, and only when President Obama said that the US was willing to bear friction in the overall US–China relationship, was progress really possible. Fourth, a responsive action and accompanying messaging needs to be appropriately sustained and not a one-off that can be easily ignored. Fifth, potential escalation needs to be considered. This is a particularly difficult issue when escalation paths aren't well defined for an event that originates in cyberspace, whether the response is a cyber or a physical one, and the chance of misperception is high. And finally, any response should comport with international law.

Collective action

Collective action against a bad actor is almost always more effective than a response by just one state and garners more legitimacy on the world stage. Of course, if the 'fiery ball of cyber death' is hurtling towards you, every country has the right to act to defend itself, but, if possible, acting together, with each country leveraging its capabilities as appropriate, is better. Collective action doesn't require any particular organised group or even the same countries acting together in each instance. Flexibility is the key here and will lead to swifter results. The recent attribution of NotPetya by a number of countries is a good example of collective action to a point. It will be interesting to see, following the US sanctioning of Russia, whether other states join in imposing collective consequences.

One challenge for both collective attribution and collective action is information sharing. Naturally, every state will want to satisfy itself before taking the political step of public attribution, and that's even more the case if it's taking further action against another transgressing state. Sharing sensitive attribution information among states with different levels of capability and ability to protect that information is a tough issue even in the best of times. But, if collective action is to happen, and happen on anything approaching a quick timeline, enhancing and even rethinking information sharing among partner countries is foundational.

Using and expanding the tools in the toolkit

The current tools that can be used in any instance to impose consequences are diplomatic, economic (including sanctions), law enforcement, cyber responses and kinetic responses. Some of them have been used in the past to varying degrees and with varying levels of effectiveness but not in a consistent and strategic way. Some, like kinetic responses, are highly unlikely to be used unless a cyber event causes death and physical injury similar to a physical attack. Others admittedly take a while to develop and deploy, but we have to have the political willingness to use them decisively in the appropriate circumstances and in a timely manner. For example, the US has

had a cyber-specific sanctions order available since April 2015 and, before its recent use against Russian actors in March, it had only been used once in December 2017 against Russian actors for election interference. For the threat of sanctions to be taken seriously, they must be used in a more regular and timely manner, and their targets should be chosen to have a real effect on the violating state's decision-making.

Our standard tools are somewhat limited, so we must also work to creatively expand the tool set so that we can better affect the unique interests of each adversarial state actor (identified in a tailored deterrence strategy), so that they'll change course or think twice before committing additional malicious acts in the future. That is likely to need collaboration not just within governments but between them and the private sector, academia, civil society and other stakeholders in order to identify and develop new tools.

Recommendations

Of course, foundational work on the application of international law and norms of voluntary state behaviour should continue. That work helps set the expectation of what conduct is permissible. In addition, states should articulate and communicate strong declaratory policies. Declaratory statements put potential adversaries on notice about what's unacceptable⁴ and can contain some detail about potential responses. In addition, a number of other things can aid in creating an environment where the threat of consequences is credible:

1. Shorten the attribution cycle.

Making progress on speeding technical attribution will take time, but delays caused by equity reviews, interagency coordination, political willingness, and securing agreement among several countries to share in making attribution are all areas that can be streamlined. Often the best way to streamline these kinds of processes is to simply exercise them by doing more public attribution while building a stronger political commitment to call bad actors out. The WannaCry and NotPetya public attributions are a great foundation for exercising the process, identifying impediments and speeding the process in the future. Even when attribution is done privately, practice can help shorten interagency delays and equity reviews.

2. If attribution can't be made or announced in a fairly brief period, couple any later public attribution with at least one visible responsive action.

Attribution six months or a year after the fact with the vague promise of future consequences will often ring hollow, particularly given the poor track record of imposing consequences in the past. When attribution can be made quickly, the promise of a future response is understandable, but delaying the announcement until it can be married with a response may be more effective.

3. Mainstream and treat cybersecurity as a core national and economic security concern and not a boutique technical issue.

If cyberattacks really pose a significant threat, governments need to start thinking of them like they think of other incidents in the physical world. It is telling that Prime Minister Theresa May made public attribution of the Salisbury poisonings in a matter of days and followed up with consequences shortly thereafter. Her decisive action also helped galvanise an international coalition in a very short time frame. Obviously that was a serious matter that required a speedy response, but the speed was also possible because government leaders are more used to dealing with physical world incidents. They still don't understand the impact or importance of cyber events or have established processes to deal with them. Mainstreaming also expands and makes existing response options more effective. As noted above, a prime reason for the US–China accord on intellectual property theft was the fact that it was considered a core economic and national security issue that was worth creating friction in the overall US–China relationship.

4. Build flexible alliances of like-minded countries to impose costs on bad actors.

A foundational element of this is improving information sharing, both in speed and substance, to enable better collective attribution and action. Given classification and trust issues, improving tactical information sharing is a difficult issue in any domain. However, a first step is to discuss with partners what information is required well in advance of any particular incident and to create the right channels to quickly share that information when needed. It may also require a re-evaluation of what information must absolutely be classified and restricted and what can be shared through appropriately sensitive channels. If there's greater joint attribution and action, this practice will presumably also help build mechanisms to share information and build trust and confidence in the future with a greater number of partners.

5. Improve diplomatic messaging to both partners and adversaries.

Improved messaging allows for better coordinated action and serves to link consequences to the actions to which they're meant to respond. Messaging and communication with the bad actor while consequences are being imposed can also help with escalation control. Of course, effective messaging must be high-level, sustained and consistent if the bad actor is to take it seriously. Sending mixed messages only serves to undercut any responsive actions that are taken.

6. Collaborate to expand the toolkit.

Work with like-minded states and other stakeholders to expand the toolkit of potential consequences that states can use, or threaten to use, to change and deter bad state actors.

7. Work out potential adversary-specific deterrence strategies.

Actual or threatened responsive actions are effective only if the target of those actions is something that matters to the state in question, and that target will differ according to the particular state involved. Of course, potential responses should be in accord with international law.

8. Most importantly, use the tools we already have to respond to serious malicious cyber activity by states in a timely manner.

Imposing consequences for bad action not only addresses whatever the current bad actions may be but creates a credible threat that those consequences (or others) will be imposed in the future.

None of this is easy or will be accomplished overnight, and there are certainly complexities in escalation, proportionality and other difficult issues, but a lot comes down to a willingness to act—and the current situation isn't sustainable. The recent US imposition of sanctions is a step in the right direction, but imposing tailored costs when appropriate needs to be part of a practice, not an aberration, and it must be accompanied by high-level messaging that supports rather than undercuts its use.


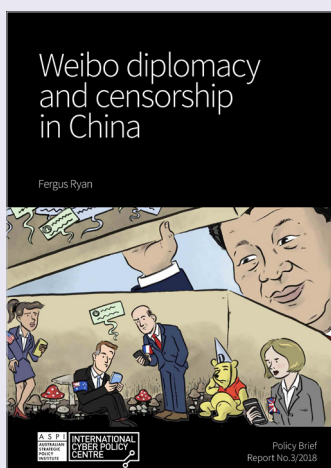
The 2017 US National Security Strategy promises ‘swift and costly consequences’ for those who target the US with cyberattacks. Australia’s International Cyber Engagement Strategy states that ‘[h]aving established a firm foundation of international law and norms, the international community must now ensure there are effective consequences for those who act contrary to this consensus.’ On the other hand, Admiral Rogers, the head of US Cyber Command and the National Security Agency, recently told US lawmakers that President Putin has clearly come to the conclusion that there’s ‘little price to pay here’ for Russia’s hacking provocations, and Putin has therefore concluded that he ‘can continue this activity’.

We must change the calculus of those who believe this is a costless enterprise. Imposing effective and timely consequences for state-sponsored cyberattacks is a key part of that change.

Notes

- 1 Of course, there are an ever-increasing number of attacks and intrusions by criminals, including transnational criminal groups, as well. Deterring this activity is a little more straightforward—the consequences for criminals are prosecution and punishment and, in particular, a heightened expectation that they'll be caught and brought to justice. I don't address deterring criminal actors in this paper, although there have been advances in ensuring that countries have the laws and capacity to tackle these crimes and there have been a number of high-profile prosecutions, including transnational cases. Much more needs to be done to deter these actors, however, as many cybercriminals still view the possibility that they'll be caught and punished as minimal.
- 2 One downside of a practice of publicly attributing state conduct is that it creates an expectation that victim states will do this in every case and leads to the perception that when they don't it means they don't know who is responsible—even if they do. For that reason, states, including the US, have often said in the past that they'll make public attribution when it serves their deterrent or other interests. There are also cases in which a state or states may want to privately challenge a transgressor state to change its behaviour or in which calling out bad conduct publicly risks sources and methods that may have a greater value in thwarting future malicious conduct. Nevertheless, the seeming trend to more cases of public attribution is a good one, and these concerns and expectations can be mitigated in a state's public messaging or by delaying public attribution when necessary.
- 3 Defence Sciences Board, Task Force on Cyber Deterrence, February 2017.
- 4 Such statements should be relatively specific but need not be over-precise about exact 'red lines', which might encourage an adversary to act just below that red line to escape a response.

Some previous ASPI publications



WHAT'S YOUR STRATEGY?

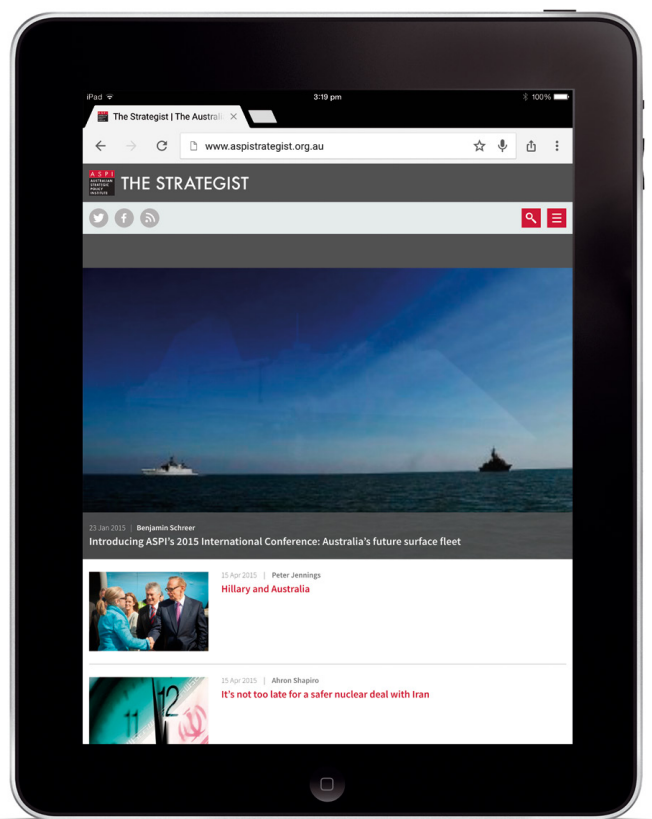


Stay informed via the field's leading think tank, the Australian Strategic Policy Institute.

The Strategist, ASPI's commentary and analysis website, delivers fresh ideas on Australia's defence and strategic policy choices as well as encouraging discussion and debate among interested stakeholders in the online strategy community. Visit and subscribe to an email digest at www.aspistrategist.org.au.

 facebook.com/ASPI.org

 [@ASPI_org](https://twitter.com/ASPI_org)



Supported by



To find out more about ASPI go to www.aspi.org.au
or contact us on 02 6270 5100 and enquiries@aspi.org.au.

A S P I
AUSTRALIAN
STRATEGIC
POLICY
INSTITUTE

INTERNATIONAL
CYBER POLICY
CENTRE

