Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: September 7, 2020

K. Paine
UK National Cyber Security Centre
O. Whitehouse
NCC Group

March 6, 2020

Indicators of Compromise (IoCs) and Their Role in Attack Defence
draft-paine-smart-indicators-of-compromise-00

Abstract

Indicators of Compromise (IoCs) are an important technique in attack defence (often called cyber defence). This document outlines the different types of IoC, their associated benefits and limitations, and discusses their effective use. It also contextualises the role of IoCs in defending against attacks through describing a recent case study. This draft does not pre-suppose where IoCs can be found or should be detected - as they can be discovered and deployed in networks, endpoints or elsewhere - rather, engineers should be aware that they need to be detectable (either by endpoint security appliances or network-based defences, or ideally both) to be effective.

Table of Contents

1. Introduction

This draft aims to describe, and illustrate the purpose of, Indicators of Compromise (IoCs), which are widely used in attack defence (often called cyber defence). The concept of the 'Pyramid of Pain' [PoP] will also be introduced to show the properties of the

broad range of defences that IoCs can provide. Furthermore, this draft will describe a real intrusion set, APT33, for which IoCs were identified and used for defence. This document is not a comprehensive report of APT33 and is intended to be read alongside APT33 open source material (for example, [Symantec]).

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. What are IoCs?

Indicators of Compromise (IoCs) are artefacts observed about an attacker; their techniques, tactics, procedures or associated tooling and infrastructure. These indicators can be observed at a combination of network or host levels and can, with varying degrees of confidence, help to identify an occurrence of an intrusion or associated activity to a known intrusion set. These IoCs are used by network defenders (blue teams) to protect their networks. Examples of IoCs can include:

- IP addresses
- domain names
- TLS Server Name Indicator values
- certificate information
- signatures such as binary code patterns and strings
- hashes of malicious binaries or scripts
- attack tools, such as mimikatz [Mimikatz]
- attack techniques, such as Kerberos golden tickets [GoldenTicket]

IoCs are often found initially through manual investigation and then shared at scale so a variety of individuals and organisations can defend themselves. They can be found in a range of locations, including in networks and at endpoints, but wherever they exist, they need to be made available to security appliances to ensure that they can be deployed quickly and widely. For IoCs to provide defence-in-depth (see Section 5), which is one of their key strengths, they should be deployed on both the network and on endpoints through

solutions that have sufficient privilege to act on them, to cope with different points of failure.

IoCs can be of varying quality. An IoC without context is not much use for network defence - a defender could do different things with an IoC (e.g. monitor it, block it, log it) depending on this context. Without the associated context, for example the threat actor it relates to, the last time it was seen in use, its expected lifetime or other related IoCs, the usefulness of an IoC is greatly reduced. On the other hand, an IoC delivered with context is much more useful to a network defender, who can then make an informed choice on how to use it to protect their network.

3. Why use IoCs?

3.1. IoCs can be used even with limited resource

IoCs are scalable and easy to deploy which makes them a really valuable asset for smaller entities. IoCs are also inexpensive to use. For example, take a small manufacturing subcontractor in a supply chain that produces a critical, highly specialised, component. The small manufacturer represents an attractive target because there would be disproportionate impact on both the supply chain and the prime contractor if it were compromised. In addition, it is likely to have comparatively smaller resource to manage the risks it faces. It is reasonable to assume that this small manufacturer will have only basic security (in the form of firewalls, similar network protection devices, or an outsource agreement), however it can still leverage IoCs to great effect. IoCs can be deployed to give a baseline protection against known threats by small entities without access to a significant defensive team or the threat intelligence relationships necessary to perform resource-intensive investigation. In addition, as detailed further in Section 3.2, the prime contractor can also supply IoCs to the subcontractor to provide an uplift in defensive capability in order to protect the prime contractor. Affording some level of protection to organisations across a spectrum of resource, maturity, and sophistication is a major part of the appeal for IoCs.

3.2. IoCs have a multiplier effect on attack defence effort

The correspondence is one-to-many: simply blocking one IoC may protect thousands of users within an organisation. Discovering one IoC can be intensive, but sharing IoCs via well-established routes such as the Malware Information Sharing Platform (MISP) [MISP] will protect thousands of organisations and end users. The shareability and reproducibility of IoCs is a huge advantage; it allows a threat defender to look for things consistently and automate the process of

defending their networks. It doesn't require intensive training (as needed to, for example, manually analyse tipped machine learning events), nor does it require time-intensive resource to deploy IoCs.

In the case of an ongoing email phishing campaign, IoCs can be monitored, discovered and deployed quickly and easily. If they are deployed quickly via a mechanism such as a protective DNS filtering service, they can be more effective still, as the same email campaign is mitigated before a recipient clicks the link or before malicious payloads can call out for instructions. While this approach can therefore be faster than some traditional defences, the most important benefit is that other parties can be protected without additional effort.

3.3. IoCs are easily shareable

This is due to two major factors: firstly, because lists of identifiers are easy to distribute, and secondly, due to standards such as Structured Threat Information Expression (STIX) [STIX] that provide a well-defined format for sharing. This allows IoCs to give blanket coverage for organisations and allow widespread mitigation in a timely fashion. They can be shared with systems administrators - from small to large organisations, from large teams to a single individual - allowing them to implement defences on their network.

3.4. IoCs can be attributed to specific threat actors

Deployment of various modern system security services, such as endpoint detection and response or firewall filtering, comes with an inherent trade-off between breadth of protection and risk of false positives (see Section 4). An organisation can examine their own risk, impact and threat - they can perform their own information assurance and threat modelling - and work to manage those threats they wish to. This means an organisation can prioritise or accept trade-offs against a subset of malicious actors; tying IoCs to threat actors allows organisations to focus their defences against particular risks. Organisations should have the technical freedom and the capability to choose their risk posture and defence methods.

3.5. IoCs can provide significant time savings

Not only are there time savings from sharing IoCs, saving duplication of investigation effort, but deploying them automatically at scale is seamless for many enterprises. Where automatic deployment of IoCs is working well, organisations and users get blanket protection with minimal human intervention and minimal effort, a key goal of attack defence. Conversely, protecting a complex network without automatic deployment of IoCs could mean updating every single endpoint or

network device consistently and reliably to the same security level. The work this entails (including polling for logs, locating assets and devices, and manually checking patch levels) introduces complexity and a need for skilled analysts. While it is still necessary to invest effort to eliminate false positives when widely deploying IoCs, the cost and effort involved can be far smaller than the work entailed in reliably patching all endpoint and network devices - for example, particularly on legacy systems that may be particularly complicated, or even impossible, to update.

3.6. IoCs allow for discovery of historic attacks

A network defender can use recently acquired IoCs in conjunction with historic data, such as logged DNS queries or email attachment hashes, to hunt for signs of past compromise. Not only can this technique help to build up a clear picture of past attacks, but it also allows for retrospective mitigation of the effects of any previous intrusion. This use case is reliant on historic data not having been compromised itself, by a technique such as Timestomp [Timestomp], or being incomplete due to third party policies, but is nonetheless valuable for detecting past attacks.

3.7. IoCs underpin and enable multiple of the layers of the modern defence-in-depth strategy

Firewalls, Intrusion Detection Systems and Security Incident Event Management platforms all employ IoCs to identify and mitigate threats. Anti-Virus (AV) products, as part of a multi-faceted approach, deploy IoCs via catalogues or libraries to all supported client endpoints. Of course, IoCs do not address all attack defence challenges - but they form a vital tier of any organisation's layered defence.

As an example, open source malware can be deployed by many different actors, each with their own "Tactics, Techniques and Procedures" (TTPs) and infrastructure. However, if the same executable is used, the hash remains the same - and this IoC can be deployed in endpoints through AV to protect regardless of TTP and infrastructure. Should this defence fail however, other defences can prevent malicious actors progressing further through the attack chain - for instance, by blocking known malicious domain name look-ups and thereby preventing the malware calling out to its command and control infrastructure.

A different malicious actor changes the intrusion set deployed across different campaigns, but their access vector remains consistent and well-known. Therefore, this TTP and pattern of activity can be recognised and proactively defended against. For example, if their

access vector consistently exploits a vulnerability in software, regular and estate-wide patching can prevent the attack from taking place. Should these pre-emptive measures fail however, other IoCs observed across multiple campaigns can be used to prevent the attack at different stages in the attack chain. Hence, IoCs can underpin multiple layers of any modern defence-in-depth strategy.

4. Pain, Fragility and Precision

The variety of IoC types inherently embody a set of trade-offs between the risk of false positives (misidentifying non-malicious traffic as malicious) and the risk of failing to identify attacks. The pain of modifying attacks to subvert known IoCs inversely correlates with both the fragility of various IoCs as a tool for attack defence, and the precision with which IoCs identify an attack. Research is needed to elucidate the exact nature of these trade-offs between pain, fragility and precision.

4.1. Pyramid of Pain

IoCs form a "Pyramid of Pain" [PoP] that can be used for prevention, detection and mitigation. This represents how much pain it is: to an adversary to change and for the defender to gather. The layers of the PoP range from hashes to TTPs and the pain ranges from recompiling code to creating a new attack strategy.

On the lowest (and least painful) level are hashes of malicious files. These are easy for a defender to gather and can be given to firewalls, for example, to block malicious downloads. To subvert this defence, an adversary need only recompile code with some trivial changes, thereby changing the hash. IoCs aren't the only route for doing this blocking but are a quick, less intrusive and more convenient method.

The next two levels are IPs and domain names. These are blockable, with varying false positive rates (see Section 4.4), and often cause more pain to an adversary to subvert; they may have to change IP ranges, find a new provider, and change their code if the IP address is hard-coded (rather than resolved). Domain names are more granular than IP addresses and are more painful for an adversary to change.

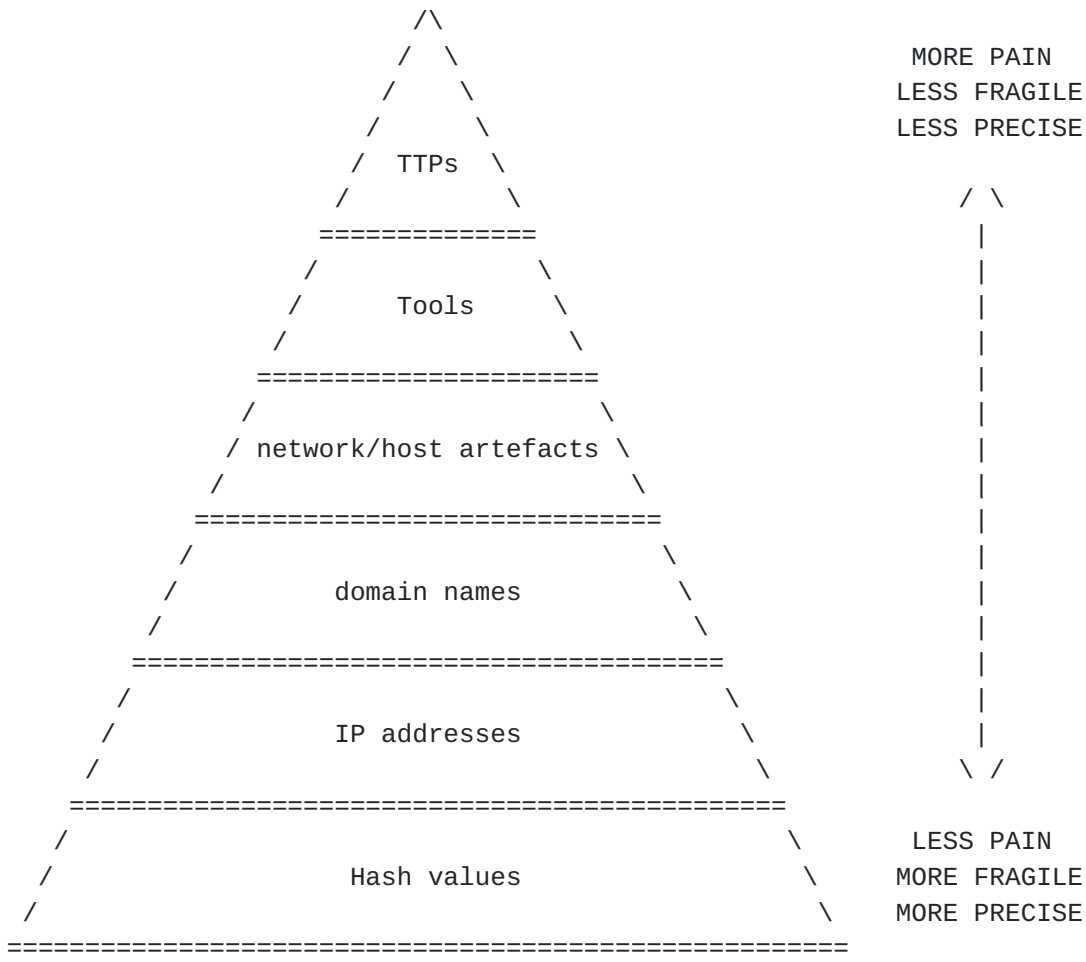

Figure 1

Network and host artefacts, such as changed timestamps of files left on the endpoint (see [Timestomp]) or a beaconing pattern on the network, are harder still to change, as they relate specifically to the attack taking place and may not be under the direct control of the attacker.

Tools and TTPs form the top two layers of the pyramid; these layers describe a threat actor's methodology - the way they perform the attack. An example could be deployment of malicious code to perform reconnaissance of a victim's network, which pivots laterally to a valuable endpoint, then downloads a ransomware payload. Tools refer to the software used to conduct the attack, whereas TTPs relate to the broader attack strategy being used. Information on TTPs and Tools take intensive effort to diagnose on the part of the defender, but they are fundamental to the attacker and campaign and hence incredibly painful for the adversary to change.

4.2. Fragility

For a network defender, the PoP can also be thought of in terms of fragility. The less painful it is for the attacker to change the IoC, the more fragile that IoC is as an attack defence tool. It should be relatively simple to get a hash of the various file attachments (and then deploy this through AV or other means) or to get an email subject for a particular campaign. However, when thinking in terms of fragility, the hash IoCs or email subjects are fragile and can be changed; in reality, they will be changed easily between campaigns. IPs and domain names can also be changed between campaigns, but it is harder - and if the IoCs didn't change but weren't blocked, that's a missed opportunity.

4.3. Precision

The PoP can be also considered in terms of how precise the defence can be, with the false positive rate roughly increasing as we move up the pyramid. A hash (e.g. MD5, SHA1 or SHA2) can specify a particular executable, so the false positive is nil. On the other hand, TTPs or fuzzier rules may apply to various binaries, and even benign software may share the same identifying methodology. This corresponds with the consequences for fragility mentioned above, as the more precise IoCs, such as hashes, are also the most fragile.

4.4. Comprehensive Coverage

IoCs provide the defender with a range of options across the PoP layers, balancing between precision and fragility to give high confidence events that are practical and useful. Broad coverage of the PoP is important as it allows the defender to cycle between high precision but high fragility options and more robust but less precise indicators. As fragile indicators are changed, the more robust IoCs allow for continued detection and faster rediscovery. This is why it's important to collect as many IoCs as possible across the whole PoP.

At the top of the PoP, TTPs identified through anomaly detection and machine learning are more likely to have false positives, which gives lower confidence and, vitally, requires better trained analysts to understand and implement the defences. Hashes, at the bottom, are precise and easy to deploy but are fragile and easily changed within and across campaigns by malicious actors.

In the middle of the pyramid, IoCs related to network information (such as domains and IP addresses) can be particularly useful. They allow for broad coverage, without requiring each and every endpoint security solution to be updated. This means they can shine in

contexts where ensuring endpoint security isn't possible such as "Bring Your Own Device" (BYOD), IoT and legacy environments. Using these network level IoCs can also protect against a compromised endpoint as, even if the compromise passes unnoticed, the IoCs can still be checked against network traffic, allowing detection of the attack. For example, in a BYOD environment, enforcing security policies on the device can be difficult, so non-endpoint IoCs and solutions are needed to allow detection of compromise even with no endpoint coverage.

Covering a broad range of IoCs gives defenders a wide range of benefits: easy to deploy, high confidence enough to be effective, painful enough to change and disruptive to bad actors. The combination of these factors cements IoCs as a particularly valuable tool for defenders with limited resources.

5. Defence in Depth

Endpoint Detection and Response (EDR) or Anti-Virus (AV) are often the first port of call for protection from intrusion but aren't a panacea. One issue is that there are many environments where it is not possible to keep them updated, or in some cases, deploy them at all. For example, the Owari botnet, a Mirai variant [Owari], exploited Internet of Things (IoT) devices where such solutions could not be deployed. It is because of such gaps, where endpoint solutions can't be relied on (see [CLESS]), that a defence-in-depth approach is commonly advocated, using a blended approach that includes both network and endpoint defences.

If an attack happens, then you hope an endpoint solution will pick it up. If it doesn't, it could be for many good reasons: the endpoint solution could be quite conservative and aim for a low false-positive rate, it might not have ubiquitous coverage or it might only be able to defend the initial step of the kill chain. In the worst cases, the attack specifically disables the endpoint solution or the malware is brand new and so won't be recognised. Going up the pyramid, IP addresses are next, and here we have ACLs (access control lists) that can go on firewalls - or your favourite DNS filtering service for protection. Using IPs will blanket-defend a range of endpoints, from printers [IoT] to "Bring Your Own Device" (BYOD) to capable endpoints. Going further through the pyramid, domains are next - these are more granular.

One example of how IoCs provide a layer of a defence-in-depth solution is Protective DNS (PDNS), a free and voluntary DNS filtering service provided by the UK NCSC for UK public sector organisations [PDNS]. In 2018, this service blocked access to 57.4 million DNS queries for 118,527 unique reasons (out of 68.7 billion total

queries) for the organisations signed up to the service [ACD2019]. 28 million of them were for domain generation algorithms (DGAs), including 15 known DGAs. IoCs such as malicious domains can be put on PDNS straight away and can then be used to prevent access to those known malicious domains across the entire estate of over 460 separate public sector entities that use NCSC's PDNS [Annual2019]. Coverage can be patchy with endpoints, as the roll-out of protections isn't uniform or necessarily fast - but if the IoC is on PDNS, a consistent defence is maintained. This offers protection, regardless of whether the context is a BYOD environment or a managed enterprise system. Other IoCs, like Server Name Indicator values in TLS or the server certificate information, also provide IoC protections.

Similar to the AV scenario, large scale services face risk decisions around balancing threat against business impact from false positives. Organisations need to be able to retain the ability to be more conservative with their own defences, while still benefiting from them. For instance, a commercial DNS filtering service is intended for broad deployment, so will have a risk tolerance similar to AV products; whereas DNS filtering intended for government users (e.g. PDNS) can be more conservative, but will still have a relatively broad deployment if intended for the whole of government. A government department or specific company, on the other hand, might accept the risk of disruption and arrange firewalls or other network protection devices to completely block anything related to particular threats, regardless of the confidence, but rely on a DNS filtering service for everything else.

Other network defences can make use of this blanket coverage from IoCs, like middlebox mitigation, proxy defences, and application layer firewalls, but they're out of scope for this draft. Note too that DNS goes through firewalls, proxies and possibly to a DNS filtering service; it doesn't have to be unencrypted, but these appliances must be able to decrypt it to do anything useful with it, like blocking queries for known bad URIs.

6. Case Study: APT33

To contextualise IoCs, we describe a real world case study: a current campaign by the threat actor APT33, also known as Elfin and Refined Kitten (see [Symantec]). APT33 has been assessed by industry to be a state-sponsored group [FireEye], yet in this case study, IoCs still gave defenders an effective tool against such a sophisticated and powerful adversary. The group has been active since at least 2015 and is known to target a range of sectors including petrochemical, government, engineering and manufacturing. Activity has been seen in countries across the globe, but predominantly in the USA and Saudi Arabia.

6.1. Overall TTP

The techniques employed by this actor exhibit a relatively low level of sophistication: typically, spear phishing is used with document lures that imitate legitimate publications. Once inside a target network, the actor will attempt to pivot to other machines to gather documents and gain access to administrative credentials. In some cases, users are tricked into providing credentials that are then used to enable the use of RULER, a freely available tool that allows exploitation of an email client. The attacker, in possession of a target's password, uses RULER to access the target's mail account, and embed a malicious script which will be triggered when the mail client is next opened, resulting in the execution of malicious code (often additional malware retrieved from the Internet) (see [FireEye2]).

When a destructive tool is deployed, it relies on overwriting the master boot record (MBR) of the hard drives in as many PCs as possible. This type of tool, known as a wiper, results in data loss and renders devices unusable until the operating system is reinstalled. In some cases, the actor is able to use administrator credentials to invoke execution across a large swathe of a company's IT estate at once; where this isn't possible the actor may attempt to spread the wiper to as many PCs as possible manually, or by using wormlike capabilities against unpatched vulnerabilities on the internal network.

6.2. IoCs

As a result of investigations by both industry and NCSC in partnership, a set of IoCs were compiled that could then be shared out with government and industry to enable network defenders to search for these indicators in their networks. Detection of these IoCs is likely indicative of APT33 targeting and could indicate potential compromise and subsequent use of destructive malware. Network defenders could also initiate processes to block these IoCs and foil future attacks. This set of IoCs comprised:

- 9 fragile indicators including hashes and email subject lines
- 5 IP addresses
- 7 domains

These IoCs mostly cover the bottom few levels of the PoP, with the network level IoCs giving resilience not provided by the fragile indicators. Not only can these IoCs be used to check historical data for evidence of past compromise, but they can also be deployed to

block further infection and/or to detect infection in a timely manner, thereby contributing to preventing the loss of user and system data.

7. Conclusions

IoCs are versatile and powerful. IoCs underpin and enable multiple of the layers of the modern defence-in-depth strategy. IoCs are easy to share, providing a multiplier effect on attack defence effort and they save vital time. Network-level IoCs offer protection, especially valuable when an endpoint-only solution isn't sufficient. These properties, along with their ease of use, make IoCs a key component of any attack defence strategy and particularly valuable for defenders with limited resources.

For IoCs to be useful, they don't have to be unencrypted or visible in networks - but crucially they do need to be made available, along with their context, to entities that need them. It is also important that this availability and eventual usage copes with multiple points of failure, as per the defence-in-depth strategy, of which IoCs are a key part.

8. Acknowledgements

Thanks to all those who have been involved with improving cyber defence in the IETF and IRTF communities.

9. IANA Considerations

This memo includes no request to IANA.

10. Security Considerations

This draft is all about system security.

11. Informative References

[ACD2019] Levy, I. and M. S, "Active Cyber Defence - The Second Year", 2019, <https://www.ncsc.gov.uk/report/active-cyber-defence-report-2019>.

[Annual2019] NCSC, "Annual Review 2019", 2019, <https://www.ncsc.gov.uk/annual-review/2019/ncsc/docs/ncsc_2019-annual-review.pdf>.

[CLESS] Taddei, A., Wueest, C., Roundy, K., and D. Lazanski, "Capabilities and Limitations of an Endpoint-only Security Solution", 2019, <https://datatracker.ietf.org/doc/draft-taddei-smart-cless-introduction/>.

[FireEye] O'Leary, J., Kimble, J., Vanderlee, K., and N. Fraser, "Insights into Iranian Cyber Espionage: APT33 Targets Aerospace and Energy Sectors and has Ties to Destructive Malware", 2017, <https://www.fireeye.com/blog/threat-research/2017/09/apt33-insights-into-iranian-cyber-espionage.html>.

[FireEye2] FireEye, "OVERRULED: Containing a Potentially Destructive Adversary", 2018, <https://www.fireeye.com/blog/threat-research/2018/12/overruled-containing-a-potentially-destructive-adversary.html>.

[GoldenTicket] Soria-Machado, M., Abolins, D., Boldea, C., and K. Socha, "Kerberos Golden Ticket Protection", 2014, <https://cert.europa.eu/static/WhitePapers/UPDATED - CERT-EU_Security_Whitepaper_2014-007_Kerberos_Golden_Ticket_Protection_v1_4.pdf>.

[IoT] NCC Group, "Security Impact of IoT on the Enterprise", 2019, <https://www.nccgroup.trust/globalassets/our-research/uk/whitepapers/2019/11/iot-whitepaper-matt.pdf>.

[Mimikatz] Mulder, J., "Mimikatz Overview, Defenses and Detection", 2016, <https://www.sans.org/reading-room/whitepapers/detection/mimikatz-overview-defenses-detection-36780>.

[MISP] MISP, "MISP", 2019, <https://www.misp-project.org/>.

[Owari] NCSC, "Owari botnet own-goal takeover", 2018, <https://www.ncsc.gov.uk/report/weekly-threat-report-8th-june-2018>.

[PDNS] NCSC, "Protective DNS", 2019, <https://www.ncsc.gov.uk/information/pdns>.

[PoP] Bianco, D., "The Pyramid of Pain", 2014, <https://detect-respond.blogspot.com/2013/03/the-pyramid-of-pain.html>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[STIX] OASIS Cyber Threat Intelligence, "STIX", 2019, <https://oasis-open.github.io/cti-documentation/stix/intro>.

[Symantec] Symantec, "Elfin: Relentless", 2019, <https://www.symantec.com/blogs/threat-intelligence/elfin-apt33-espionage>.

[Timestomp] OASIS Cyber Threat Intelligence, "Timestomp", 2019, <https://attack.mitre.org/techniques/T1099/>.

Authors' Addresses

Kirsty Paine
UK National Cyber Security Centre

Email: kirsty.p@ncsc.gov.uk

Ollie Whitehouse
NCC Group

Email: ollie.whitehouse@nccgroup.com