

Internet Engineering Task Force (IETF)
Request for Comments: 9424
Category: Informational
ISSN: 2070-1721

K. Paine
Splunk Inc.
O. Whitehouse
Binary Firefly
J. Sellwood

A. Shaw
UK National Cyber Security Centre
August 2023

Indicators of Compromise (IoCs) and Their Role in Attack Defence

Abstract

Cyber defenders frequently rely on Indicators of Compromise (IoCs) to identify, trace, and block malicious activity in networks or on endpoints. This document reviews the fundamentals, opportunities, operational limitations, and recommendations for IoC use. It highlights the need for IoCs to be detectable in implementations of Internet protocols, tools, and technologies -- both for the IoCs' initial discovery and their use in detection -- and provides a foundation for approaches to operational challenges in network security.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9424>.

Copyright Notice

Copyright (c) 2023 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction
2. Terminology
3. IoC Fundamentals
 - 3.1. IoC Types and the Pyramid of Pain

3.2.	IoC Lifecycle	
3.2.1.	Discovery	
3.2.2.	Assessment	
3.2.3.	Sharing	
3.2.4.	Deployment	
3.2.5.	Detection	
3.2.6.	Reaction	
3.2.7.	End of Life	
4.	Using IoCs Effectively	
4.1.	Opportunities	
4.1.1.	IoCs underpin and enable multiple layers of the modern defence-in-depth strategy.	
4.1.2.	IoCs can be used even with limited resources.	
4.1.3.	IoCs have a multiplier effect on attack defence efforts within an organisation.	
4.1.4.	IoCs are easily shared between organisations.	
4.1.5.	IoCs can provide significant time savings.	
4.1.6.	IoCs allow for discovery of historic attacks.	
4.1.7.	IoCs can be attributed to specific threats.	
4.2.	Case Studies	
4.2.1.	Cobalt Strike	
4.2.1.1.	Overall TTP	
4.2.1.2.	IoCs	
4.2.2.	APT33	
4.2.2.1.	Overall TTP	
4.2.2.2.	IoCs	
5.	Operational Limitations	
5.1.	Time and Effort	
5.1.1.	Fragility	
5.1.2.	Discoverability	
5.1.3.	Completeness	
5.2.	Precision	
5.2.1.	Specificity	
5.2.2.	Dual and Compromised Use	
5.2.3.	Changing Use	
5.3.	Privacy	
5.4.	Automation	
6.	Comprehensive Coverage and Defence-in-Depth	
7.	IANA Considerations	
8.	Security Considerations	
9.	Conclusions	
10.	Informative References	
	Acknowledgements	
	Authors' Addresses	

1. Introduction

This document describes the various types of IoCs and how they are used effectively in attack defence (often called "cyber defence"). It introduces concepts such as the Pyramid of Pain [PoP] and the IoC lifecycle to highlight how IoCs may be used to provide a broad range of defences. This document provides suggestions for implementers of controls based on IoCs as well as potential operational limitations. Two case studies that demonstrate the usefulness of IoCs for detecting and defending against real-world attacks are included. One case study involves an intrusion set (a set of malicious activity and behaviours attributed to one threat actor) known as "APT33", and the other involves an attack tool called "Cobalt Strike". This document is not a comprehensive report of APT33 or Cobalt Strike and is intended to be read alongside publicly published reports (referred to as "open-source material" among cyber intelligence practitioners) on these threats (for example, [Symantec] and [NCCGroup], respectively).

2. Terminology

Attack defence:

The activity of providing cyber security to an environment through the prevention of, detection of, and response to attempted and successful cyber intrusions. A successful defence can be achieved through blocking, monitoring, and responding to adversarial activity at the network, endpoint, or application levels.

Command and control (C2) server:

An attacker-controlled server used to communicate with, send commands to, and receive data from compromised machines. Communication between a C2 server and compromised hosts is called "command and control traffic".

Domain Generation Algorithm (DGA):

The algorithm used in malware strains to periodically generate domain names (via algorithm). Malware may use DGAs to compute a destination for C2 traffic rather than relying on a pre-assigned list of static IP addresses or domains that can be blocked more easily when extracted from, or otherwise linked to, the malware.

Kill chain:

A model for conceptually breaking down a cyber intrusion into stages of the attack from reconnaissance through to actioning the attacker's objectives. This model allows defenders to think about, discuss, plan for, and implement controls to defend against discrete phases of an attacker's activity [KillChain].

Tactics, Techniques, and Procedures (TTPs):

The way an adversary undertakes activities in the kill chain -- the choices made, methods followed, tools and infrastructure used, protocols employed, and commands executed. If they are distinct enough, aspects of an attacker's TTPs can form specific IoCs as if they were a fingerprint.

Control (as defined by US NIST):

A safeguard or countermeasure prescribed for an information system or an organisation designed to protect the confidentiality, integrity, and availability of its information and to meet a set of defined security requirements [NIST].

3. IoC Fundamentals

3.1. IoC Types and the Pyramid of Pain

IoCs are observable artefacts relating to an attacker or their activities, such as their tactics, techniques, procedures, and associated tooling and infrastructure. These indicators can be observed at the network or endpoint (host) levels and can, with varying degrees of confidence, help network defenders to proactively block malicious traffic or code execution, determine a cyber intrusion occurred, or associate discovered activity to a known intrusion set and thereby potentially identify additional avenues for investigation. IoCs are deployed to firewalls and other security control points by adding them to the list of indicators that the control point is searching for in the traffic that it is monitoring. When associated with malicious activity, the following are some examples of protocol-related IoCs:

- * IPv4 and IPv6 addresses in network traffic
- * Fully Qualified Domain Names (FQDNs) in network traffic, DNS resolver caches, or logs
- * TLS Server Name Indication values in network traffic
- * Code-signing certificates in binaries

- * TLS certificate information (such as SHA256 hashes) in network traffic
- * Cryptographic hashes (e.g., MD5, SHA1, or SHA256) of malicious binaries or scripts when calculated from network traffic or file system artefacts
- * Attack tools (such as Mimikatz [Mimikatz]) and their code structure and execution characteristics
- * Attack techniques, such as Kerberos Golden Tickets [GoldenTicket], that can be observed in network traffic or system artefacts

The common types of IoC form a Pyramid of Pain [PoP] that informs prevention, detection, and mitigation strategies. The position of each IoC type in the pyramid represents how much "pain" a typical adversary experiences as part of changing the activity that produces that artefact. The greater pain an adversary experiences (towards the top), the less likely they are to change those aspects of their activity and the longer the IoC is likely to reflect the attacker's intrusion set (i.e., the less fragile those IoCs will be from a defender's perspective). The layers of the PoP commonly range from hashes up to TTPs, with the pain ranging from simply recompiling code to creating a whole new attack strategy. Other types of IoC do exist and could be included in an extended version of the PoP should that assist the defender in understanding and discussing intrusion sets most relevant to them.

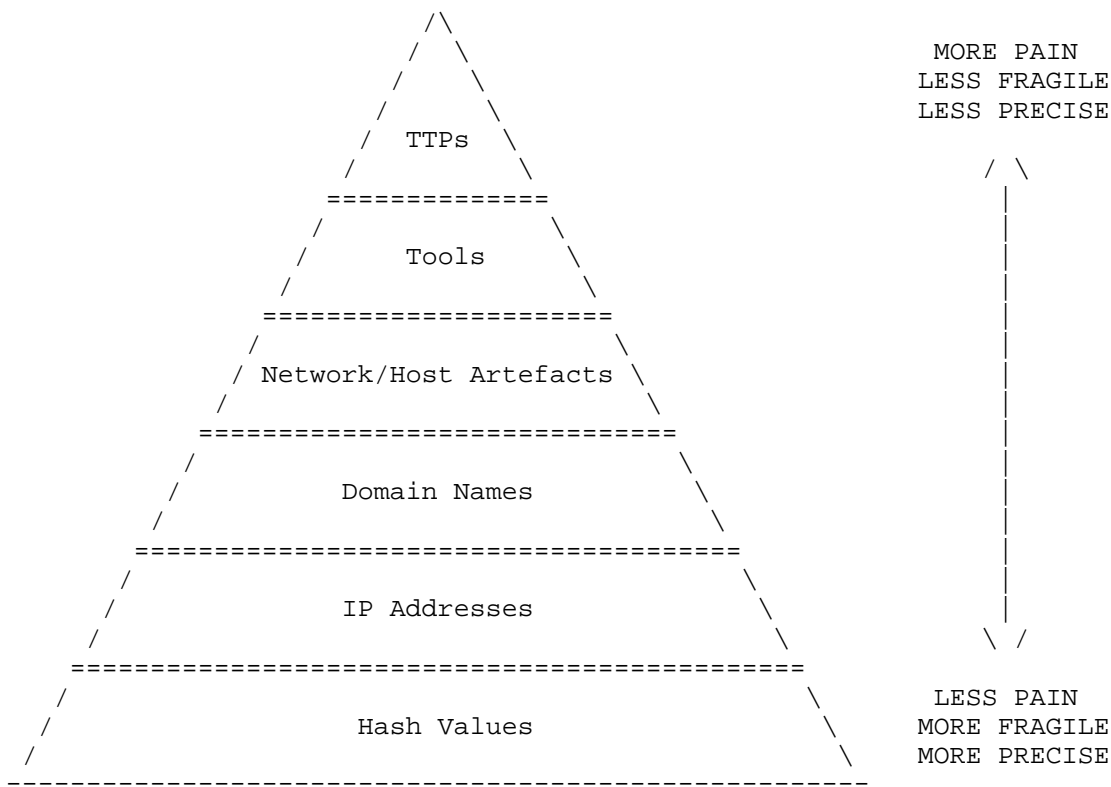


Figure 1

On the lowest (and least painful) level are hashes of malicious files. These are easy for a defender to gather and can be deployed to firewalls or endpoint protection to block malicious downloads or prevent code execution. While IoCs aren't the only way for defenders to do this kind of blocking, they are a quick, convenient, and nonintrusive method. Hashes are precise detections for individual files based on their binary content. To subvert this defence, however, an adversary need only recompile code, or otherwise modify

the file content with some trivial changes, to modify the hash value.

The next two levels are IP addresses and domain names. Interactions with these may be blocked, with varying false positive rates (misidentifying non-malicious traffic as malicious; see Section 5), and often cause more pain to an adversary to subvert than file hashes. The adversary may have to change IP ranges, find a new provider, and change their code (e.g., if the IP address is hard-coded rather than resolved). A similar situation applies to domain names, but in some cases, threat actors have specifically registered these to masquerade as a particular organisation or to otherwise falsely imply or claim an association that will be convincing or misleading to those they are attacking. While the process and cost of registering new domain names are now unlikely to be prohibitive or distracting to many attackers, there is slightly greater pain in selecting unregistered, but appropriate, domain names for such purposes.

Network and endpoint artefacts, such as a malware's beaconing pattern on the network or the modified timestamps of files touched on an endpoint, are harder still to change as they relate specifically to the attack taking place and, in some cases, may not be under the direct control of the attacker. However, more sophisticated attackers use TTPs or tooling that provides flexibility at this level (such as Cobalt Strike's malleable command and control [COBALT]) or a means by which some artefacts can be masked (see [Timestamp]).

Tools and TTPs form the top two levels of the pyramid; these levels describe a threat actor's methodology -- the way they perform the attack. The tools level refers specifically to the software (and less frequently, hardware) used to conduct the attack, whereas the TTPs level picks up on all the other aspects of the attack strategy. IoCs at these levels are more complicated and complex -- for example, they can include the details of how an attacker deploys malicious code to perform reconnaissance of a victim's network, pivots laterally to a valuable endpoint, and then downloads a ransomware payload. 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.

The variation in discoverability of IoCs is indicated by the numbers of IoCs in AlienVault, an open threat intelligence community [ALIENVAULT]. As of January 2023, AlienVault contained:

- * Groups (i.e., combinations of TTPs): 631
- * Malware families (i.e., tools): ~27,000
- * URL: 2,854,918
- * Domain names: 64,769,363
- * IPv4 addresses: 5,427,762
- * IPv6 addresses: 12,009
- * SHA256 hash values: 5,452,442

The number of domain names appears out of sync with the other counts, which reduce on the way up the PoP. This discrepancy warrants further research; however, contributing factors may be the use of DGAs and the fact that threat actors use domain names to masquerade as legitimate organisations and so have added incentive for creating new domain names as they are identified and confiscated.

3.2. IoC Lifecycle

To be of use to defenders, IoCs must first be discovered, assessed, shared, and deployed. When a logged activity is identified and correlated to an IoC, this detection triggers a reaction by the defender, which may include an investigation, potentially leading to more IoCs being discovered, assessed, shared, and deployed. This cycle continues until the IoC is determined to no longer be relevant, at which point it is removed from the control space.

3.2.1. Discovery

IoCs are discovered initially through manual investigation or automated analysis. They can be discovered in a range of sources, including at endpoints and in the network (on the wire). They must either be extracted from logs monitoring protocol packet captures, code execution, or system activity (in the case of hashes, IP addresses, domain names, and network or endpoint artefacts) or be determined through analysis of attack activity or tooling. In some cases, discovery may be a reactive process, where IoCs from past or current attacks are identified from the traces left behind. However, discovery may also result from proactive hunting for potential future IoCs extrapolated from knowledge of past events (such as from identifying attacker infrastructure by monitoring domain name registration patterns).

Crucially, for an IoC to be discovered, the indicator must be extractable from the Internet protocol, tool, or technology it is associated with. Identifying a particular exchange (or sequence of exchanged messages) related to an attack is of limited benefit if indicators cannot be extracted or, once they are extracted, cannot be subsequently associated with a later related exchange of messages or artefacts in the same, or in a different, protocol. If it is not possible to determine the source or destination of malicious attack traffic, it will not be possible to identify and block subsequent attack traffic either.

3.2.2. Assessment

Defenders may treat different IoCs differently, depending on the IoCs' quality and the defender's needs and capabilities. Defenders may, for example, place differing trust in IoCs depending on their source, freshness, confidence level, or the associated threat. These decisions rely on associated contextual information recovered at the point of discovery or provided when the IoC was shared.

An IoC without context is not much use for network defence. On the other hand, an IoC delivered with context (for example, the threat actor it relates to, its role in an attack, the last time it was seen in use, its expected lifetime, or other related IoCs) allows a network defender to make an informed choice on how to use it to protect their network (for example, simply log it, actively monitor it, or outright block it).

3.2.3. Sharing

Once discovered and assessed, IoCs are most helpful when deployed in such a way to have a broad impact on the detection or disruption of threats or shared at scale so many individuals and organisations can defend themselves. An IoC may be shared individually (with appropriate context) in an unstructured manner or may be packaged alongside many other IoCs in a standardised format, such as Structured Threat Information Expression [STIX], Malware Information Sharing Platform (MISP) core [MISPCORE], OpenIOC [OPENIOC], and Incident Object Description Exchange Format (IODEF) [RFC7970]. This enables distribution via a structured feed, such as one implementing Trusted Automated Exchange of Intelligence Information [TAXII], or

through a Malware Information Sharing Platform [MISP].

While some security companies and some membership-based groups (often dubbed "Information Sharing and Analysis Centres (ISACs)" or "Information Sharing and Analysis Organizations (ISAOs)") provide paid intelligence feeds containing IoCs, there are various free IoC sources available from individual security researchers up through small trust groups to national governmental cyber security organisations and international Computer Emergency Response Teams (CERTs). Whoever they are, sharers commonly indicate the extent to which receivers may further distribute IoCs using frameworks like the Traffic Light Protocol [TLP]. At its simplest, this indicates that the receiver may share with anyone (TLP:CLEAR), share within the defined sharing community (TLP:GREEN), share within their organisation and their clients (TLP:AMBER+STRICT), share just within their organisation (TLP:AMBER), or not share with anyone outside the original specific IoC exchange (TLP:RED).

3.2.4. Deployment

For IoCs to provide defence-in-depth (see Section 6) and so cope with different points of failure, correct deployment is important. Different IoCs will detect malicious activity at different layers of the network stack and at different stages of an attack, so deploying a range of IoCs enables layers of defence at each security control, reinforcing the benefits of using multiple security controls as part of a defence-in-depth solution. The network security controls and endpoint solutions where they are deployed need to have sufficient privilege, and sufficient visibility, to detect IoCs and to act on them. Wherever IoCs exist, they need to be made available to security controls and associated apparatus to ensure they can be deployed quickly and widely. While IoCs may be manually assessed after discovery or receipt, significant advantage may be gained by automatically ingesting, processing, assessing, and deploying IoCs from logs or intelligence feeds to the appropriate security controls. As not all IoCs are of the same quality, confidence in IoCs drawn from each threat intelligence feed should be considered when deciding whether to deploy IoCs automatically in this way.

IoCs can be particularly effective at mitigating malicious activity when deployed in security controls with the broadest impact. This could be achieved by developers of security products or firewalls adding support for the distribution and consumption of IoCs directly to their products, without each user having to do it, thus addressing the threat for the whole user base at once in a machine-scalable and automated manner. This could also be achieved within an enterprise by ensuring those control points with the widest aperture (for example, enterprise-wide DNS resolvers) are able to act automatically based on IoC feeds.

3.2.5. Detection

Security controls with deployed IoCs monitor their relevant control space and trigger a generic or specific reaction upon detection of the IoC in monitored logs or on network interfaces.

3.2.6. Reaction

The reaction to an IoC's detection may differ depending on factors such as the capabilities and configuration of the control it is deployed in, the assessment of the IoC, and the properties of the log source in which it was detected. For example, a connection to a known botnet C2 server may indicate a problem but does not guarantee it, particularly if the server is a compromised host still performing some other legitimate functions. Common reactions include event logging, triggering alerts, and blocking or terminating the source of

the activity.

3.2.7. End of Life

How long an IoC remains useful varies and is dependent on factors including initial confidence level, fragility, and precision of the IoC (discussed further in Section 5). In some cases, IoCs may be automatically "aged" based on their initial characteristics and so will reach end of life at a predetermined time. In other cases, IoCs may become invalidated due to a shift in the threat actor's TTPs (e.g., resulting from a new development or their discovery) or due to remediation action taken by a defender. End of life may also come about due to an activity unrelated to attack or defence, such as when a third-party service used by the attacker changes or goes offline. Whatever the cause, IoCs should be removed from detection at the end of their life to reduce the likelihood of false positives.

4. Using IoCs Effectively

4.1. Opportunities

IoCs offer a variety of opportunities to cyber defenders as part of a modern defence-in-depth strategy. No matter the size of an organisation, IoCs can provide an effective, scalable, and efficient defence mechanism against classes of attack from the latest threats or specific intrusion sets that may have struck in the past.

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

Firewalls, Intrusion Detection Systems (IDSs), and Intrusion Prevention Systems (IPs) all employ IoCs to identify and mitigate threats across networks. Antivirus (AV) and Endpoint Detection and Response (EDR) products deploy IoCs via catalogues or libraries to supported client endpoints. Security Incident Event Management (SIEM) platforms compare IoCs against aggregated logs from various sources -- network, endpoint, and application. Of course, IoCs do not address all attack defence challenges, but they form a vital tier of any organisation's layered defence. Some types of IoC may be present across all those controls while others may be deployed only in certain layers of a defence-in-depth solution. Further, IoCs relevant to a specific kill chain may only reflect activity performed during a certain phase and so need to be combined with other IoCs or mechanisms for complete coverage of the kill chain as part of an intrusion set.

As an example, open-source malware can be deployed by many different actors, each using their own TTPs and infrastructure. However, if the actors use the same executable, the hash of the executable file remains the same, and this hash can be deployed as an IoC in endpoint protection to block execution regardless of individual actor, infrastructure, or other TTPs. Should this defence fail in a specific case, for example, if an actor recompiles the executable binary producing a unique hash, other defences can prevent them progressing further through their attack, for instance, by blocking known malicious domain name lookups and thereby preventing the malware calling out to its C2 infrastructure.

Alternatively, another malicious actor may regularly change their tools and infrastructure (and thus the indicators associated with the intrusion set) deployed across different campaigns, but their access vectors may remain consistent and well-known. In this case, this access TTP can be recognised and proactively defended against, even while there is uncertainty of the intended subsequent activity. For example, if their access vector consistently exploits a vulnerability in software, regular and estate-wide patching can prevent the attack

from taking place. However, should these preemptive measures fail, other IoCs observed across multiple campaigns may be able to prevent the attack at later stages in the kill chain.

4.1.2. IoCs can be used even with limited resources.

IoCs are inexpensive, scalable, and easy to deploy, making their use particularly beneficial for smaller entities, especially where they are exposed to a significant threat. For example, a small manufacturing subcontractor in a supply chain producing a critical, highly specialised component may represent an attractive target because there would be disproportionate impact on both the supply chain and the prime contractor if it were compromised. It may be reasonable to assume that this small manufacturer will have only basic security (whether internal or outsourced), and while it is likely to have comparatively fewer resources to manage the risks that it faces compared to larger partners, it can still leverage IoCs to great effect. Small entities like this can deploy IoCs to give a baseline protection against known threats without having access to a well-resourced, mature defensive team and the threat intelligence relationships necessary to perform resource-intensive investigations. While some level of expertise on the part of such a small company would be needed to successfully deploy IoCs, use of IoCs does not require the same intensive training as needed for more subjective controls, such as those using machine learning, which require further manual analysis of identified events to verify if they are indeed malicious. In this way, a major part of the appeal of IoCs is that they can afford some level of protection to organisations across spectrums of resource capability, maturity, and sophistication.

4.1.3. IoCs have a multiplier effect on attack defence efforts within an organisation.

Individual IoCs can provide widespread protection that scales effectively for defenders across an organisation or ecosystem. Within a single organisation, simply blocking one IoC may protect thousands of users, and that blocking may be performed (depending on the IoC type) across multiple security controls monitoring numerous different types of activity within networks, endpoints, and applications. The prime contractor from our earlier example can supply IoCs to the small subcontractor and thus further uplift that smaller entity's defensive capability while protecting itself and its interests at the same time.

Multiple organisations may benefit from directly receiving shared IoCs (see Section 4.1.4), but they may also benefit from the IoCs' application in services they utilise. In the case of an ongoing email-phishing campaign, IoCs can be monitored, discovered, and deployed quickly and easily by individual organisations. However, if they are deployed quickly via a mechanism such as a protective DNS filtering service, they can be more effective still -- an email campaign may be mitigated before some organisations' recipients ever click the link or before some malicious payloads can call out for instructions. Through such approaches, other parties can be protected without direct sharing of IoCs with those organisations or additional effort.

4.1.4. IoCs are easily shared between organisations.

IoCs can also be very easily shared between individuals and organisations. First, IoCs are easy to distribute as they can be represented concisely as text (possibly in hexadecimal) and so are frequently exchanged in small numbers in emails, blog posts, or technical reports. Second, standards, such as those mentioned in Section 3.2.3, exist to provide well-defined formats for sharing large collections or regular sets of IoCs along with all the

associated context. While discovering one IoC can be intensive, once shared via well-established routes, that individual IoC may protect thousands of organisations and thus all of the users in those organisations. Quick and easy sharing of IoCs gives blanket coverage for organisations and allows widespread mitigation in a timely fashion -- they can be shared with systems administrators, from small to large organisations and from large teams to single individuals, allowing them all to implement defences on their networks.

4.1.5. 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. The ability to do this at scale and at pace is often vital when responding to agile threat actors that may change their intrusion set frequently and hence change the relevant IoCs. Conversely, protecting a complex network without automatic deployment of IoCs could mean manually updating every single endpoint or network device consistently and reliably to the same security state. The work this entails (including locating assets and devices, polling for logs and system information, and manually checking patch levels) introduces complexity and a need for skilled analysts and engineers. While it is still necessary to invest effort both to enable efficient IoC deployment and to eliminate false positives when widely deploying IoCs, the cost and effort involved can be far smaller than the work entailed in reliably manually updating all endpoint and network devices. For example, legacy systems may be particularly complicated, or even impossible, to update.

4.1.6. 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 a clear picture of past attacks, but it also allows for retrospective mitigation of the effects of any previous intrusion. This opportunity is reliant on historic data not having been compromised itself, by a technique such as Timestamp [Timestamp], and not being incomplete due to data retention policies, but it is nonetheless valuable for detecting and remediating past attacks.

4.1.7. IoCs can be attributed to specific threats.

Deployment of various modern security controls, such as firewall filtering or EDR, come with an inherent trade-off between breadth of protection and various costs, including the risk of false positives (see Section 5.2), staff time, and pure financial costs. Organisations can use threat modelling and information assurance to assess and prioritise risk from identified threats and to determine how they will mitigate or accept each of them. Contextual information tying IoCs to specific threats or actors and shared alongside the IoCs enables organisations to focus their defences against particular risks. This contextual information is generally expected by those receiving IoCs as it allows them the technical freedom and capability to choose their risk appetite, security posture, and defence methods. The ease of sharing this contextual information alongside IoCs, in part due to the formats outlined in Section 3.2.3, makes it easier to track malicious actors across campaigns and targets. Producing this contextual information before sharing IoCs can take intensive analytical effort as well as specialist tools and training. At its simplest, it can involve documenting sets of IoCs from multiple instances of the same attack campaign, for example, from multiple unique payloads (and therefore

with distinct file hashes) from the same source and connecting to the same C2 server. A more complicated approach is to cluster similar combinations of TTPs seen across multiple campaigns over a period of time. This can be used alongside detailed malware reverse engineering and target profiling, overlaid on a geopolitical and criminal backdrop, to infer attribution to a single threat actor.

4.2. Case Studies

The following two case studies illustrate how IoCs may be identified in relation to threat actor tooling (in the first) and a threat actor campaign (in the second). The case studies further highlight how these IoCs may be used by cyber defenders.

4.2.1. Cobalt Strike

Cobalt Strike [COBALT] is a commercial attack framework used for penetration testing that consists of an implant framework (beacon), a network protocol, and a C2 server. The beacon and network protocol are highly malleable, meaning the protocol representation "on the wire" can be easily changed by an attacker to blend in with legitimate traffic by ensuring the traffic conforms to the protocol specification, e.g., HTTP. The proprietary beacon supports TLS encryption overlaid with a custom encryption scheme based on a public-private keypair. The product also supports other techniques, such as domain fronting [DFRONT], in an attempt to avoid obvious passive detection by static network signatures of domain names or IP addresses. Domain fronting is used to blend traffic to a malicious domain with traffic originating from a network that is already communicating with a non-malicious domain regularly over HTTPS.

4.2.1.1. Overall TTP

A beacon configuration describes how the implant should operate and communicate with its C2 server. This configuration also provides ancillary information such as the Cobalt Strike user licence watermark.

4.2.1.2. IoCs

Tradecraft has been developed that allows the fingerprinting of C2 servers based on their responses to specific requests. This allows the servers to be identified, their beacon configurations to be downloaded, and the associated infrastructure addresses to be extracted as IoCs.

The resulting mass IoCs for Cobalt Strike are:

- * IP addresses of the C2 servers
- * domain names used

Whilst these IoCs need to be refreshed regularly (due to the ease of which they can be changed), the authors' experience of protecting public sector organisations shows that these IoCs are effective for disrupting threat actor operations that use Cobalt Strike.

These IoCs can be used to check historical data for evidence of past compromise and deployed to detect or block future infection in a timely manner, thereby contributing to preventing the loss of user and system data.

4.2.2. APT33

In contrast to the first case study, this describes a current campaign by the threat actor APT33, also known as Elfin and Refined

Kitten (see [Symantec]). APT33 has been assessed by the industry to be a state-sponsored group [FireEye2]; yet, in this case study, IoCs still gave defenders an effective tool against such a 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.

4.2.2.1. Overall TTP

The techniques employed by this actor exhibit a relatively low level of sophistication, considering it is a state-sponsored group. Typically, APT33 performs spear phishing (sending targeted malicious emails to a limited number of pre-selected recipients) with document lures that imitate legitimate publications. User interaction with these lures executes the initial payload and enables APT33 to gain initial access. Once inside a target network, APT33 attempts 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 with Ruler [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 embeds a malicious script that 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 [FireEye]).

APT33 sometimes deploys a destructive tool that overwrites 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 uses 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 first attempt to spread the wiper manually or use worm-like capabilities against unpatched vulnerabilities on the networked computers.

4.2.2.2. IoCs

As a result of investigations by a partnership of the industry and the UK's National Cyber Security Centre (NCSC), a set of IoCs were compiled and shared with both public and private sector organisations so network defenders could search for them 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 to foil future attacks. This set of IoCs comprised:

- * 9 hashes and email subject lines
- * 5 IP addresses
- * 7 domain names

In November 2021, a joint advisory concerning APT33 [CISA] was issued by the Federal Bureau of Investigation (FBI), the Cybersecurity and Infrastructure Security Agency (CISA), the Australian Cyber Security Centre (ACSC), and NCSC. This outlined recent exploitation of vulnerabilities by APT33, providing a thorough overview of observed TTPs and sharing further IoCs:

- * 8 hashes of malicious executables
- * 3 IP addresses

5. Operational Limitations

The different IoC types inherently embody a set of trade-offs for defenders between the risk of false positives (misidentifying non-malicious traffic as malicious) and the risk of failing to identify attacks. The attacker's relative pain of modifying attacks to subvert known IoCs, as discussed using the PoP in Section 3.1, inversely correlates with the fragility of the IoC and with the precision with which the IoC identifies an attack. Research is needed to elucidate the exact nature of these trade-offs between pain, fragility, and precision.

5.1. Time and Effort

5.1.1. Fragility

As alluded to in Section 3.1, the PoP can be thought of in terms of fragility for the defender as well as pain for the attacker. The less painful it is for the attacker to change an IoC, the more fragile that IoC is as a defence tool. It is relatively simple to determine the hash value for various malicious file attachments observed as lures in a phishing campaign and to deploy these through AV or an email gateway security control. However, those hashes are fragile and can (and often will) be changed between campaigns. Malicious IP addresses and domain names can also be changed between campaigns, but this may happen less frequently due to the greater pain of managing infrastructure compared to altering files, and so IP addresses and domain names may provide a less fragile detection capability.

This does not mean the more fragile IoC types are worthless. First, there is no guarantee a fragile IoC will change, and if a known IoC isn't changed by the attacker but wasn't blocked, then the defender missed an opportunity to halt an attack in its tracks. Second, even within one IoC type, there is variation in the fragility depending on the context of the IoC. The file hash of a phishing lure document (with a particular theme and containing a specific staging server link) may be more fragile than the file hash of a remote access trojan payload the attacker uses after initial access. That in turn may be more fragile than the file hash of an attacker-controlled post-exploitation reconnaissance tool that doesn't connect directly to the attacker's infrastructure. Third, some threats and actors are more capable or inclined to change than others, and so the fragility of an IoC for one may be very different to an IoC of the same type for another actor.

Ultimately, fragility is a defender's concern that impacts the ongoing efficacy of each IoC and will factor into decisions about end of life. However, it should not prevent adoption of individual IoCs unless there are significantly strict resource constraints that demand down-selection of IoCs for deployment. More usually, defenders researching threats will attempt to identify IoCs of varying fragilities for a particular kill chain to provide the greatest chances of ongoing detection given available investigative effort (see Section 5.1.2) and while still maintaining precision (see Section 5.2).

5.1.2. Discoverability

To be used in attack defence, IoCs must first be discovered through proactive hunting or reactive investigation. As noted in Section 3.1, IoCs in the tools and TTPs levels of the PoP require intensive effort and research to discover. However, it is not just an IoC's type that impacts its discoverability. The sophistication of the actor, their TTPs, and their tooling play a significant role, as does whether the IoC is retrieved from logs after the attack or

extracted from samples or infected systems earlier.

For example, on an infected endpoint, it may be possible to identify a malicious payload and then extract relevant IoCs, such as the file hash and its C2 server address. If the attacker used the same static payload throughout the attack, this single file hash value will cover all instances. However, if the attacker diversified their payloads, that hash can be more fragile, and other hashes may need to be discovered from other samples used on other infected endpoints. Concurrently, the attacker may have simply hard-coded configuration data into the payload, in which case the C2 server address can be easy to recover. Alternatively, the address can be stored in an obfuscated persistent configuration within either the payload (e.g., within its source code or associated resource) or the infected endpoint's file system (e.g., using alternative data streams [ADS]), thus requiring more effort to discover. Further, the attacker may be storing the configuration in memory only or relying on a DGA to generate C2 server addresses on demand. In this case, extracting the C2 server address can require a memory dump or the execution or reverse engineering of the DGA, all of which increase the effort still further.

If the malicious payload has already communicated with its C2 server, then it may be possible to discover that C2 server address IoC from network traffic logs more easily. However, once again, multiple factors can make discoverability more challenging, such as the increasing adoption of HTTPS for malicious traffic, meaning C2 communications blend in with legitimate traffic and can be complicated to identify. Further, some malwares obfuscate their intended destinations by using alternative DNS resolution services (e.g., OpenNIC [OPENNIC]), by using encrypted DNS protocols such as DNS-over-HTTPS [OILRIG], or by performing transformation operations on resolved IP addresses to determine the real C2 server address encoded in the DNS response [LAZARUS].

5.1.3. Completeness

In many cases, the list of indicators resulting from an activity or discovered in a malware sample is relatively short and so only adds to the total set of all indicators in a limited and finite manner. A clear example of this is when static indicators for C2 servers are discovered in a malware strain. Sharing, deployment, and detection will often not be greatly impacted by the addition of such indicators for one more incident or one more sample. However, in the case of discovery of a DGA, this requires a reimplementing of the algorithm and then execution to generate a possible list of domains. Depending on the algorithm, this can result in very large lists of indicators, which may cause performance degradation, particularly during detection. In some cases, such sources of indicators can lead to a pragmatic decision being made between obtaining reasonable coverage of the possible indicator values and theoretical completeness of a list of all possible indicator values.

5.2. Precision

5.2.1. Specificity

Alongside pain and fragility, the PoP's levels can also be considered in terms of how precise the defence can be, with the false positive rate usually increasing as we move up the pyramid to less specific IoCs. A hash value identifies a particular file, such as an executable binary, and given a suitable cryptographic hash function, the false positives are effectively nil (by "suitable", we mean one with preimage resistance and strong collision resistance). In comparison, IoCs in the upper levels (such as some network artefacts or tool fingerprints) may apply to various malicious binaries, and

even benign software may share the same identifying characteristics. For example, threat actor tools making web requests may be identified by the user-agent string specified in the request header. However, this value may be the same as that used by legitimate software, either by the attacker's choice or through use of a common library.

It should come as no surprise that the more specific an IoC, the more fragile it is; as things change, they move outside of that specific focus. While less fragile IoCs may be desirable for their robustness and longevity, this must be balanced with the increased chance of false positives from their broadness. One way in which this balance is achieved is by grouping indicators and using them in combination. While two low-specificity IoCs for a particular attack may each have chances of false positives, when observed together, they may provide greater confidence of an accurate detection of the relevant kill chain.

5.2.2. Dual and Compromised Use

As noted in Section 3.2.2, the context of an IoC, such as the way in which the attacker uses it, may equally impact the precision with which that IoC detects an attack. An IP address representing an attacker's staging server, from which their attack chain downloads subsequent payloads, offers a precise IP address for attacker-owned infrastructure. However, it will be less precise if that IP address is associated with a cloud-hosting provider and is regularly reassigned from one user to another; it will be less precise still if the attacker compromised a legitimate web server and is abusing the IP address alongside the ongoing legitimate use.

Similarly, a file hash representing an attacker's custom remote access trojan will be very precise; however, a file hash representing a common enterprise remote administration tool will be less precise, depending on whether or not the defender organisation usually uses that tool for legitimate system administration. Notably, such dual-use indicators are context specific, considering both whether they are usually used legitimately and how they are used in a particular circumstance. Use of the remote administration tool may be legitimate for support staff during working hours but not generally by non-support staff, particularly if observed outside of that employee's usual working hours.

For reasons like these, context is very important when sharing and using IoCs.

5.2.3. Changing Use

In the case of IP addresses, the growing adoption of cloud services, proxies, virtual private networks (VPNs), and carrier-grade Network Address Translation (NAT) are increasing the number of systems associated with any one IP address at the same moment in time. This ongoing change to the use of IP addresses is somewhat reducing the specificity of IP addresses (at least for specific subnets or individual addresses) while also "side-stepping" the pain that threat actors would otherwise incur if they needed to change IP address.

5.3. Privacy

As noted in Section 3.2.2, context is critical to effective detection using IoCs. However, at times, defenders may feel there are privacy concerns with how much and with whom to share about a cyber intrusion. For example, defenders may generalise the IoCs' description of the attack by removing context to facilitate sharing. This generalisation can result in an incomplete set of IoCs being shared or IoCs being shared without clear indication of what they represent and how they are involved in an attack. The sharer will

consider the privacy trade-off when generalising the IoC and should bear in mind that the loss of context can greatly reduce the utility of the IoC for those they share with.

In the authors' experiences, self-censoring by sharers appears more prevalent and more extensive when sharing IoCs into groups with more members, into groups with a broader range of perceived member expertise (particularly, the further the lower bound extends below the sharer's perceived own expertise), and into groups that do not maintain strong intermember trust. Trust within such groups often appears strongest where members interact regularly; have common backgrounds, expertise, or challenges; conform to behavioural expectations (such as by following defined handling requirements and not misrepresenting material they share); and reciprocate the sharing and support they receive. [LITREVIEW] highlights that many of these factors are associated with the human role in Cyber Threat Intelligence (CTI) sharing.

5.4. Automation

While IoCs can be effectively utilised by organisations of various sizes and resource constraints, as discussed in Section 4.1.2, automation of IoC ingestion, processing, assessment, and deployment is critical for managing them at scale. Manual oversight and investigation may be necessary intermittently, but a reliance on manual processing and searching only works at small scale or for occasional cases.

The adoption of automation can also enable faster and easier correlation of IoC detections across different log sources and network monitoring interfaces across different times and physical locations. Thus, the response can be tailored to reflect the number and overlap of detections from a particular intrusion set, and the necessary context can be presented alongside the detection when generating any alerts for defender review. While manual processing and searching may be no less accurate (although IoC transcription errors are a common problem during busy incidents in the experience of the authors), the correlation and cross-referencing necessary to provide the same degree of situational awareness is much more time-consuming.

A third important consideration when performing manual processing is the longer phase monitoring and adjustment necessary to effectively age out IoCs as they become irrelevant or, more crucially, inaccurate. Manual implementations must often simply include or exclude an IoC, as anything more granular is time-consuming and complicated to manage. In contrast, automations can support a gradual reduction in confidence scoring, enabling IoCs to contribute but not individually disrupt a detection as their specificity reduces.

6. Comprehensive Coverage and Defence-in-Depth

IoCs provide the defender with a range of options across the PoP's layers, enabling them to balance precision and fragility to give high confidence detections that are practical and useful. Broad coverage of the PoP is important as it allows the defender to choose between high precision but high fragility options and more robust but less precise indicators depending on availability. As fragile indicators are changed, the more robust IoCs allow for continued detection and faster rediscovery. For this reason, it's important to collect as many IoCs as possible across the whole PoP to provide options for defenders.

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. However, these are very painful for attackers to change, so when tuned appropriately, they provide a robust detection. Hashes, at the bottom, are precise and easy to deploy but are fragile and easily changed within and across campaigns by malicious actors.

Endpoint Detection and Response (EDR) or Antivirus (AV) are often the first port of call for protection from intrusion, but endpoint solutions 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, that a defence-in-depth approach is commonly advised, using a blended approach that includes both network and endpoint defences.

If an attack happens, then the best situation is that an endpoint solution will detect and prevent it. 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 [KillChain]. In the worst cases, the attack specifically disables the endpoint solution, or the malware is brand new and so won't be recognised.

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, as they may be detected and enforced in a more centralised manner at network choke points (such as proxies and gateways). This makes them particularly useful in contexts where ensuring endpoint security isn't possible, such as Bring Your Own Device (BYOD), Internet of Things (IoT), and legacy environments. It's important to note that these network-level IoCs can also protect users of a network against compromised endpoints when these IoCs are used to detect the attack in network traffic, even if the compromise itself passes unnoticed. 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.

One example of how network-level IoCs provide a layer of a defence-in-depth solution is Protective DNS (PDNS) [Annual2021], a free and voluntary DNS filtering service provided by the UK NCSC for UK public sector organisations [PDNS]. In 2021, this service blocked access to more than 160 million DNS queries (out of 602 billion total queries) for the organisations signed up to the service [ACD2021]. This included hundreds of thousands of queries for domains associated with Flubot, Android malware that uses DGAs to generate 25,000 candidate command and control domains each month (these DGAs [DGAs] are a type of TTP).

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 925 separate public sector entities that use NCSC's PDNS. Coverage can be patchy with endpoints, as the roll-out of protections isn't uniform or necessarily fast. However, if the IoC is on PDNS, a consistent defence is maintained for devices using PDNS, even if the device itself is not immediately updated. This offers protection, regardless of whether the context is a BYOD environment or a managed enterprise system. PDNS provides the most front-facing layer of defence-in-depth solutions for its users, but other IoCs, like Server Name Indication values in TLS or the server certificate information, also provide IoC protections at other

layers.

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 it 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 are out of scope for this document. Large enterprise networks are likely to deploy their own DNS resolution architecture and possibly TLS inspection proxies and can deploy IoCs in these locations. However, in networks that choose not to, or don't have the resources to, deploy these sorts of mitigations, DNS goes through firewalls, proxies, and possibly 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.

Covering a broad range of IoCs gives defenders a wide range of benefits: they are easy to deploy; they provide a high enough confidence to be effective; at least some will be painful for attackers to change; and their distribution around the infrastructure allows for different points of failure, and so overall they enable the defenders to disrupt bad actors. The combination of these factors cements IoCs as a particularly valuable tool for defenders with limited resources.

7. IANA Considerations

This document has no IANA actions.

8. Security Considerations

This document is all about system security. However, when poorly deployed, IoCs can lead to over-blocking, which may present an availability concern for some systems. While IoCs preserve privacy on a macro scale (by preventing data breaches), research could be done to investigate the impact on privacy from sharing IoCs, and improvements could be made to minimise any impact found. The creation of a privacy-preserving method of sharing IoCs that still allows both network and endpoint defences to provide security and layered defences would be an interesting proposal.

9. Conclusions

IoCs are versatile and powerful. IoCs underpin and enable multiple layers of the modern defence-in-depth strategy. IoCs are easy to share, providing a multiplier effect on attack defence efforts, and they save vital time. Network-level IoCs offer protection, which is 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 it is crucial that they be made available, along with their context, to entities that need them. It is also important that this availability and eventual usage cope with multiple points of failure, as per the defence-in-depth strategy, of which IoCs are a key part.

10. Informative References

- [ACD2021] UK NCSC, "Active Cyber Defence - The Fifth Year", May 2022, <<https://www.ncsc.gov.uk/files/ACD-The-Fifth-Year-full-report.pdf>>.
- [ADS] Microsoft, "File Streams (Local File Systems)", January 2021, <<https://docs.microsoft.com/en-us/windows/win32/fileio/file-streams>>.
- [ALIENVAULT] AlienVault, "AlienVault: The World's First Truly Open Threat Intelligence Community", <<https://otx.alienvault.com/>>.
- [Annual2021] UK NCSC, "NCSC Annual Review 2021: Making the UK the safest place to live and work online", 2021, <<https://www.ncsc.gov.uk/files/NCSC%20Annual%20Review%202021.pdf>>.
- [CISA] CISA, "Iranian Government-Sponsored APT Cyber Actors Exploiting Microsoft Exchange and Fortinet Vulnerabilities in Furtherance of Malicious Activities", November 2021, <<https://www.cisa.gov/uscert/ncas/alerts/aa21-321a>>.
- [COBALT] "Cobalt Strike", <<https://www.cobaltstrike.com/>>.
- [DFRONT] Infosec, "Domain Fronting", April 2017, <<https://resources.infosecinstitute.com/topic/domain-fronting/>>.
- [DGAs] MITRE, "Dynamic Resolution: Domain Generation Algorithms", 2020, <<https://attack.mitre.org/techniques/T1483/>>.
- [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", September 2017, <<https://www.mandiant.com/resources/blog/apt33-insights-into-iranian-cyber-espionage>>.
- [FireEye2] Ackerman, G., Cole, R., Thompson, A., Orleans, A., and N. Carr, "OVERRULED: Containing a Potentially Destructive Adversary", December 2018, <<https://www.mandiant.com/resources/blog/overruled-containing-a-potentially-destructive-adversary>>.
- [GoldenTicket] Mizrahi, I. and Cymptom, "Steal or Forge Kerberos Tickets: Golden Ticket", 2020, <<https://attack.mitre.org/techniques/T1558/001/>>.
- [KillChain] Lockheed Martin, "The Cyber Kill Chain", <<https://www.lockheedmartin.com/en-us/capabilities/cyber/cyber-kill-chain.html>>.
- [LAZARUS] Kaspersky Lab, "Lazarus Under The Hood", <<https://media.kasperskycontenthub.com/wp->

content/uploads/sites/43/2018/03/07180244/
Lazarus_Under_The_Hood_PDF_final.pdf>.

[LITREVIEW]

Wagner, T., Mahbub, K., Palomar, E., and A. Abdallah, "Cyber Threat Intelligence Sharing: Survey and Research Directions", January 2019, <<https://www.open-access.bcu.ac.uk/7852/1/Cyber%20Threat%20Intelligence%20Sharing%20Survey%20and%20Research%20Directions.pdf>>.

[Mimikatz] Mulder, J., "Mimikatz Overview, Defenses and Detection", February 2016, <<https://www.sans.org/white-papers/36780/>>.

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

[MISPCORE] Dulaunoy, A. and A. Iklody, "MISP core format", Work in Progress, Internet-Draft, draft-dulaunoy-misp-core-format-16, 26 February 2023, <<https://datatracker.ietf.org/doc/html/draft-dulaunoy-misp-core-format-16>>.

[NCCGroup] Jansen, W., "Abusing cloud services to fly under the radar", January 2021, <<https://research.nccgroup.com/2021/01/12/abusing-cloud-services-to-fly-under-the-radar/>>.

[NIST] NIST, "Glossary - security control", <https://csrc.nist.gov/glossary/term/security_control>.

[OILRIG] Cimpanu, C., "Iranian hacker group becomes first known APT to weaponize DNS-over-HTTPS (DoH)", August 2020, <<https://www.zdnet.com/article/iranian-hacker-group-becomes-first-known-apt-to-weaponize-dns-over-https-doh/>>.

[OPENIOC] Gibb, W. and D. Kerr, "OpenIOC: Back to the Basics", October 2013, <<https://www.fireeye.com/blog/threat-research/2013/10/openioc-basics.html>>.

[OPENNIC] "OpenNIC", <<https://www.opennic.org/>>.

[Owari] UK NCSC, "Owari botnet own-goal takeover", 2018, <<https://webarchive.nationalarchives.gov.uk/ukgwa/20220301141030/https://www.ncsc.gov.uk/report/weekly-threat-report-8th-june-2018>>.

[PDNS] UK NCSC, "Protective Domain Name Service (PDNS)", August 2017, <<https://www.ncsc.gov.uk/information/pdns>>.

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

[RFC7970] Danyliw, R., "The Incident Object Description Exchange Format Version 2", RFC 7970, DOI 10.17487/RFC7970, November 2016, <<https://www.rfc-editor.org/info/rfc7970>>.

[RULER] MITRE, "Ruler", <<https://attack.mitre.org/software/S0358/>>.

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

[Symantec] Symantec, "Elfin: Relentless Espionage Group Targets Multiple Organizations in Saudi Arabia and U.S.", March 2019, <<https://www.symantec.com/blogs/threat-intelligence/>>.

elfin-apt33-espionage>.

[TAXII] OASIS Cyber Threat Intelligence (CTI), "Introduction to TAXII", <<https://oasis-open.github.io/cti-documentation/taxii/intro.html>>.

[Timestamp] MITRE, "Indicator Removal: Timestamp", January 2020, <<https://attack.mitre.org/techniques/T1099/>>.

[TLP] FIRST, "Traffic Light Protocol (TLP)", <<https://www.first.org/tlp/>>.

Acknowledgements

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

Authors' Addresses

Kirsty Paine
Splunk Inc.
Email: kirsty.ietf@gmail.com

Ollie Whitehouse
Binary Firefly
Email: ollie@binaryfirefly.com

James Sellwood
Email: james.sellwood.ietf@gmail.com

Andrew Shaw
UK National Cyber Security Centre
Email: andrew.s2@ncsc.gov.uk