## DNS Blocking: A Viable Strategy in Malware Defense

### Electronic Countermeasures

During the wars in Iraq and Afghanistan, insurgents' use of improvised explosive devices (IEDs) proliferated. The United States ramped up its development of counter-IED equipment to improve standoff detection of explosives and explosive precursor components and to defeat IEDs themselves as part of a broader defense capability. One effective strategy was jamming or interrupting radio frequency (RF) communications to counter radio-controlled IEDs (RCIEDs). This approach disrupts critical parts of RF communications, making the RCIED's communication to activate ineffective, saving both warfighter and civilian lives and property. For some time now, the cyber world has also been under attack by a diffuse set of enemies who improvise their own tools in many different varieties and hide them where they can do much damage. This analogy has its limitations; however, here I want to explore the idea of disrupting communications from malicious code such as ransomware that is used to lock up your digital assets or data-exfiltration software that is used to steal your digital data.

The IED analogy has been leveraged in the Lockheed Martin Corporation Cyber Kill Chain (CKC), which depicts the stages of cyber attack in a process-based approach (again, this imagery has its limitations). Although the CKC has been used mostly on data exfiltration as the "action on objectives," a similar approach can be depicted for ransomware, with lockup of your digital assets as the intended final outcome. Much like the counter-IED initiatives developed by the U.S. military, cybersecurity capabilities are available to an enterprise in various phases of this malicious software's "lifecycle" and can be deployed to counter these attacks. The malicious code or malware depends heavily on its network communications media to perform various nefarious activities and implement many stages of its exploit. In one sense, the CKC demonstrates that if the communication chain can be broken, you can neutralize the effect of malicious code and disable its ability to achieve its action on objective. The figure below depicts a cybersecurity portfolio with various cybersecurity capabilities that can break the chain of malicious activity against an enterprise.

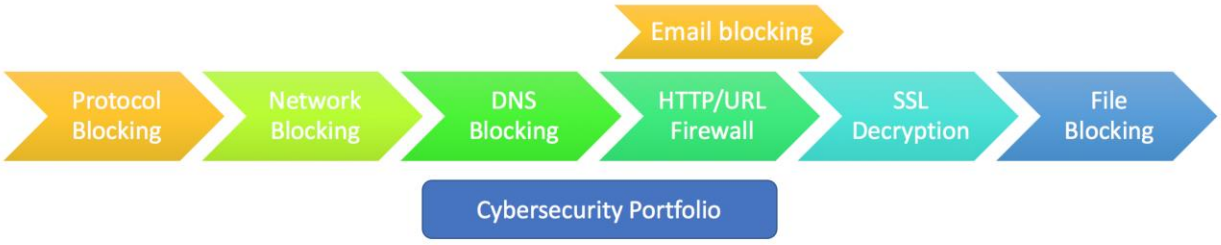


These capabilities provide varying levels of signal interruption or communications jamming to the malicious code, blocking its ability to complete its action on objective. The left-most capability—protocol blocking—blocks risky network communications and can be performed at the network level. The right-most capability—file blocking—usually requires a presence at the end point, or the host. As you move to the right of this portfolio of

defense, you can ensure more accuracy in your defense against malicious code. However, it comes with more complications in implementation and the expectation to have visibility of every end-point device in the enterprise. Domain Name Services (DNS) blocking falls in the middle of these types of capability, ensuring a wide impact while avoiding the complexity of having to install or instrument every device in your enterprise.

DNS provides a phonebook-like lookup of Internet resources. DNS blocking denies the phonebook lookup or responds in a way that disables communication for a particular internet resource. In this sense, DNS blocking provides a valuable defense against multiple stages of the CKC and can be compared to Duke V3, an electronic countermeasure system for RCIEDs developed by SRC, Inc., and used in Afghanistan. The Duke V3's success was attributed to its ability to interrupt the critical control channel communications that activate the IED. Duke V3 was touted for its low cost and high protection radius. DNS blocking, in a similar sense, can be deployed at an enterprise's perimeter with reasonable cost and have a wide impact in blocking malicious communications. This blog steps through the process of creating and enabling a DNS blocking capability in your enterprise.

**Choosing What to Block Using DNS**

Today's enterprise-level recursive DNS services provide a variety of options for blocking domain names. Here we will explore some categories and their benefits.

1. Top-level domain blacklist
   Instead of choosing domain names to block uniquely, you choose to block entire top-level domains such as .XXX (pornography), .BIT (BitTorrent file exchange servers), or .TOR2WEB (TOR proxies). This approach can be extremely beneficial in reducing the risky categories of domain names from being resolved by your enterprise.

2. New domains
   Services such as Newly Observed Domains (NOD) can provide quick protection from domain names that were registered recently for malicious purposes. This type of block list can make it difficult for malware authors to create new domains and use domain-generation algorithms to keep their malicious communications untraceable.

3. Response Policy Zone (RPZ) subscription
   Instead of building elaborate lists of bad domain names, DNS RPZ can be used for on-demand lookup of domain names through a service provider to obtain either their reputations or their categories and use them to block the domain names.

4. By response or requesting IP/network block
   Domain names that resolve a network block that is either in a country or in an untrusted network zone can be blocked entirely. This approach can also be used to block an internal device or network that is requesting access to an Internet-based resource. Using this method, you could restrict entire internal network segments of your organizations from access to external resources. You can also restrict other

risky communications, such as DNS rebinding, in which an external domain name pretends to be a local resource.

5. By anomalous data types and answers
   Domain name responses to unknown types of requests (non-RFC complaint, not relevant) and nonsensical responses (NULL) can be used to block DNS tunneling and abuse of the DNS protocol to build covert channels. This method is very effective in eliminating bogus communications or covert channel communications attempts using DNS.

**Choosing How to Block**

DNS blocking is performed for malicious domains at the recursive boundary of the enterprise using three broad name-response categories:

1. Nonexistent domain (NXDOMAIN): This method is used to provide a response that the requested domain or the domain entity itself does not exist.
2. Domain redirected: In this case, a malicious domain request is redirected to a local resource to quarantine it and provide support for mitigation and recovery.
3. Request denied: The server refuses to answer a specific query (Response Code 5: Query Refused) to disrupt anomalous DNS payloads.

In many cases, the NXDOMAIN response is simpler to implement and provides a way to deny requests to malicious domain name resources. However, NXDOMAIN makes it difficult to provide feedback to users who might click on malicious links or attempt to work around the block, not knowing it is a security violation. These three options give you a variety of choices for planning your "jamming" of malicious communications so that you are able not only to limit risk but also to recover devices that are likely infected in your enterprise. For example, choosing domain redirect to quarantine a set of high-risk domains allows you to collect information on your enterprise computers or systems that are infected with malware.

**Challenges to DNS Blocking**

It is important to note that DNS blocking poses some challenges to the enterprise. Here are a few practices that your organization can use to reduce their impact on your business or mission and bolster your solution by enabling DNS blocking at your perimeter.

1. Whitelist before blacklist
   Trusted domain names and domain extensions that should be allowed in your enterprise to ensure critical communications such as VOIP, email, and communication to trusted partners are not negatively impacted by this capability. It is important that you build a whitelist and ensure that it is maintained

with clear documentation, tracking, and lifecycle management.

2. Using DNSSEC
   DNS Security Extensions (DNSSEC) enables digital signatures of domain names, making them tamper-proof. Digital signatures can be a challenge to implementing DNS blocking because the enterprise's internal resources may ignore the answer provided by the DNS blocking system. One mitigation technique is to have the DNS blocking system act as an authoritative server for blocked DNSSEC domains. You can pursue this once you have measured the usage of DNSSEC in your enterprise and assessed the need to prevent any bottlenecks that it introduces.

3. Managing the blacklist (volume, lifecycle, private-vs.-public, classified blacklist)
   DNS blacklists that are locally managed can quickly become unmanageable in size and complexity. In most cases, it is best to partner with a provider who can manage these blacklisted "zones" and provide responses via RPZ on-demand. This reduces the size of the blacklist while providing an on-demand lookup to valid unknown domains. Enterprises should also manage the local blacklist with clear lifecycle activities and ensure that domains are expired and removed from the blacklist. The slew of blacklist providers can also be difficult to assess, but evaluating them for your industry will help you choose the most appropriate blacklist service provider. I recommend that you reserve classified (government) or proprietary (internal threat intelligence) blacklists for the final lookup, after investigating all the other blacklists, to ensure that your critical threat analysis resources are focused on advanced threat actors.

4. Logging and analysis
   DNS queries for an enterprise can be voluminous, ranging from 30,000 to 80,000 cumulative queries per second in a large-scale enterprise. These can be very hard to log for the entire DNS recursive system. However, logging blocked DNS queries and analyzing them are essential for understanding the effectiveness of your DNS blocking capability. Logging and analysis also enables systematic requests from users to whitelist certain domain names. In addition to applying the general best practices for your cybersecurity tools, DNS blocking logs should maintain minimal information to sufficiently understand (a) what was blocked, (b) when the blocking was triggered, (c) which downstream device or network requested the blocked resource, and (d) what policy was triggered in the blocking event.

**Way Forward in DNS Blocking**

DNS blocking will continue to play a critical role in the enterprise cybersecurity capabilities value chain. Every enterprise should explore its role and its pertinent approach to enable DNS blocking. As your organization matures in DNS blocking, here are a few forward-thinking ideas in this area to explore to ensure that the service does not become stale but dynamically changes to address new challenges in cybersecurity:

1. Explore options for Internet Service Provider (ISP) or managed services.
   An ISP-provided or -managed DNS blocking service can be both cost-effective and simple for your enterprise to pursue. But typically you should consider an ISP-managed service only after you have a

DNS blocking capability at your enterprise because first you should understand your organization's needs in this area so that you can effectively communicate these as requirements to an external service provider. An upstream ISP that provides DNS blocking as a service may also be exploring the use of various RPZ providers to ensure security of the data communications. An ISP-managed blocking service could be a viable option, especially if it also provides logging for your analysis.

2. Augment your blacklist with timeliness and context.
   Many blacklist providers do not offer sufficient context or understanding of domains blocked for malicious usage or malicious intent. This context is critical for an organization that is building its own blacklist and protecting against some advanced threat actors. Wherever possible, augment your blacklist with as much context as possible to ensure that you have sufficient information to analyze blocking events and understand specific threats and risks experienced by the enterprise.

3. Move beyond indicators to understand adversarial techniques.
   Newer modalities in computing provide the ability to analyze data at scale and use patterns to recognize adversarial techniques. Methods such as machine learning and predictive analysis can help glean knowledge from the data obtained using capabilities such as DNS blocking. Once you have matured your capability and have voluminous data, you can apply techniques to enlarge your blacklist to detect malicious usage patterns and expand your blacklist to capture adversarial techniques (like domain-generation algorithms or domain-registration techniques) instead of specific indicators (like domain names).

Techniques used by cyber adversaries continue to evolve, using more application layer attacks backed by a very sophisticated set of tools. It is necessary for an enterprise defense strategy to be timely, cost-effective, and active to continue to protect its systems and data. DNS blocking is clearly one such capability to activate and mitigate the risks associated with cyber threats.

**References and Further Reading**

(References to specific commercial products, processes, or services do not necessarily imply endorsement by Carnegie Mellon University or the Software Engineering Institute.)


Counter IED: https://en.wikipedia.org/wiki/Counter-IED_equipment

Mietzner, Jan, et al. "Responsive communications jamming against radio-controlled improvised explosive devices." IEEE Communications Magazine 50.10 (2012): 38-46.

A threat-driven approach to cybersecurity: http://lockheedmartin.com/content/dam/lockheed/data/isgs/documents/Threat-Driven%20Approach%20whitepaper.pdf

DNS RPZ introduction: https://dnsrpz.info/

RPZ configuration techniques: http://www.zytrax.com/books/dns/ch7/rpz.html

RPZ at enterprise scales: http://blogs.cisco.com/security/using-dns-rpz-to-block-malicious-dns-requests

Spamhaus botnet RPZ service: https://www.spamhaus.org/news/article/669/spamhaus-dbl-as-a-response-policy-zone-rpz

Farsight NOD service: https://www.farsightsecurity.com/solutions/threat-intelligence-team/newly-observed-domains/

Cisco Umbrella service overview of DNS security: https://learn-umbrella.cisco.com/solution-briefs/dns-layer-network-security

DNS rebinding: https://en.wikipedia.org/wiki/DNS_rebinding

Copyright 2017 Carnegie Mellon University. All Rights Reserved.
This material is based upon work funded and supported by the Department of Defense under Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.
The view, opinions, and/or findings contained in this material are those of the author(s) and should not be construed as an official Government position, policy, or decision, unless designated by other documentation.
NO WARRANTY. THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.
[DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited distribution. Please see Copyright notice for non-US Government use and distribution.
Internal use:* Permission to reproduce this material and to prepare derivative works from this material for internal use is granted, provided the copyright and “No Warranty” statements are included with all reproductions and derivative works.
External use:* This material may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other external and/or commercial use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.
* These restrictions do not apply to U.S. government entities.
Carnegie Mellon® and CERT® are registered in the U.S. Patent and Trademark Office by Carnegie Mellon University.
DM17-0298