

Accelerare le Security Operation con il DNS Infoblox

Gianluca De Risi gderisi@Infoblox.com

5 Giugno 2019



Dal Rapporto Clusit 2019 –Alert su Azure-

- Più di un terzo degli alert generati a fronte di attacchi su Azure faccia riferimento a **comunicazioni DNS malevole**. Si tratta in particolare di rilevazioni di client che tentano di comunicare con domini malevoli e canali di command&control. Il DNS viene spesso usato inoltre come canale per esfiltrare dati da sistemi compromessi, attraverso ad esempio il tunneling del traffico TCP attraverso l'infrastruttura DNS, oppure tramite server DNS custom in grado di interpretare messaggi DNS opportunamente codificati.

Alert Display Name	Numero di attacchi Q4 2017	Numero di attacchi Q4 2018	Numero di attacchi Q4 2019
Malicious DNS communication	0	5886	14957
Incoming RDP Brute Force Attacks – Generic Algorithm	4662	2251	6264
Unfamiliar sign-in properties	0	0	3840
Random process name detected	0	0	2959
Suspicious authentication activity	0	0	1889
Anonymous IP address	0	0	1629




DNS

DNS Firewall Response Policy Zone

Response Policy Zones

- First universal standard for DNS firewall policy
 - Vixie & Schryver, 2009
- We carry DNS firewall rules inside DNS zones
 - ...thus the name, RPZ
- Rules are published, subscribed, shared, by normal DNS zone transfer protocol
 - Including IXFR, NOTIFY, TSIG
 - So, propagation is timely, efficient, and authentic
- Add only 5% of delay to process millions of entries

DNS Data Path with DNS Firewall Added



RPZ Policy Triggers

- RPZs support five different categories of triggers
 1. The domain name owner of a resource record in an answer
 2. Response IP address
 3. Client IP address
 4. NSDNAME
 5. NSIP
 6. Machine Learning Algorithms (Infoblox)

RPZ Policy Actions

- RPZs support six different categories of actions
 1. NXDOMAIN
 2. NODATA
 3. PASSTHRU
 4. DROP
 5. TCP-Only
 6. Local Data

RPZ “Feeds”

- An RPZ “feed” is a Response Policy Zone that an organization maintains and makes available for transfer
 - Often keeping it updated with rules addressing the latest threats
- You can
 - Subscribe to one or more internal/external feeds
 - Maybe build your own feed for internal use
 - Maybe offer your feed to external subscribers

What Can I Do with RPZ?

- Easy stuff:
 - Block access to DGA C&C's
 - Block access to known phish/drive-by
 - Block e-mail if envelope/header is spammy
 - Block data data exfiltration via DNS
- More interesting stuff:
 - Block DNS A/AAAA records in bad address space
 - E.g., import Cymru Bogons or Spamhaus DROP list
 - Block DNS records in your own address space
 - After allowing your own domains to do so, of course
 - Block client on the network with no domain suffix search list queries “wpad”
 - RPZ allows monitoring of specific records (useful also for marketing), these records may or may not exist, RPZ can catch both
 - RPZs for User Education, don't just block redirect!

TLDR: DomainTools Top Risky TLDs

TLD	Domain Count
.top	4,106,700
.tk	3,158,181
.xyz	2,526,182
.tw	2,476,020
.loan	2,203,017
.ga	1,950,194
.ml	1,682,226
.cf	1,667,317
.gq	1,534,950
.club	1,407,012
.online	1,210,913
.site	1,182,432
.ltd	652,897
.work	597,322
.vip	566,351

An example of Domain Tools intelligence use with RPZ

- Add the TLD to RPZ
- Enable passthrough
- Enable log
- Analyze on a 30 days timeframe


<https://blog.domaintools.com/2019/05/using-domaintools-threat-profile-to-identify-risky-tlds/>

Exfiltrating Data via DNS Tunneling

- Uses DNS as a covert communication channel to bypass firewalls
- Attacker tunnels other protocols like SSH, or web within DNS
- Enables attackers to easily insert malware, pass stolen data or tunnel IP traffic without detection
- A DNS tunnel can be used as a full remote-control channel for a compromised internal host

Examples

- Iodine
- OzymanDNS
- SplitBrain
- DNS2TCP




Data Exfiltration over DNS Queries

- Sophisticated (zero-day)
- Infected endpoint gets access to file containing sensitive data
- It encrypts and converts info into encoded format
- Text broken into chunks and sent via DNS using hostname.subdomain or TXT records
- Exfiltrated data reconstructed at the other end
- Can use spoofed addresses to avoid detection

Data Exfiltration via host/subdomain Simplified/unencrypted example:

MarySmith.foo.thief.com
SSN-543112197.foo.thief.com
DOB-04-10-1999.foo.thief.com
MRN100045429886.foo.thief.com



Areas of Security DDI Enables

#1

Infrastructure Protection

Better Application and Service Availability



#2

Data Protection and Malware Mitigation

Protect Users and Data



#3

Threat Containment and Operations

Efficiency & Optimization of Security Operations



Areas of Security DDI Enables

#1

Infrastructure Protection

Better Application and Service Availability



#2

Data Protection and Malware Mitigation

Protect Users and Data




#3

Threat Containment and Operations

Efficiency & Optimization of Security Operations




Current Security Siloed Architecture



Problem

- Security ops is spending money on siloed solutions and costly, isolated threat feeds
- This is that we named “Accidental Security Tower”
- The SIEM is tracking events based upon separate uncorrelated vendor threat feeds
- The Ops team doesn’t know what feed is trustworthy and what not
- With the digital landscape it becomes an issue

NextGen Threat Intel Architecture



an Infoblox company



Solution

- Infoblox Introduced TIDE (Threat Intel Data Exchange) Platform in order to solve that
- Take all of these feeds and turn them into TIDE for curating, organization, conflict resolution and normalization
- We simplify all this complexity it down into one single feed, optimized and curated
- We natively speak with each policy enforcement point, unifying the policy across all the Security Ecosystem
- Now the SIEM has fewer conflicting events to remediate. This is a big cost reduction
- Our DDI Repository in an hybrid architecture, can enrich the SIEM events with all the added info we have

Bulk Correlation within TIDE

ActiveTrust TIDE is how Infoblox extends its Threat Intelligence beyond The DNS Firewall Application, to cross-pollinate our known threat indicators within your existing security environment. We can pull in timely, comprehensive indicator sets to cross-compare against recent traffic or internal events. For example, via API, you can pull:

- Indicators from a source or set of sources over a given time.
- Indicators related to a specific Threat Class or property, across available sources.
- Indicators in an active state from a source.

The Info I need	The Call to Make
I need IID-sourced hostnames for the past 30 minutes.	https://platform.activetrust.net:8000/api/data/threats?type=host&profile=IID&period=30min&data_format=json
I need iSight Partners and DHS AIS IPs for the past week, in CSV format.	https://platform.activetrust.net:8000/api/data/threats?profile=AIS-FEDGOV,iSIGHTPARTNERS&period=1w&data_format=csv
I need all Malware C2 hosts across all my available sources, for the past day, limited to 50,000 records in CEF format.	https://platform.activetrust.net:8000/api/data/threats/host?&class=MalwareC2&period=1d&data_format=cef&rlimit=50000
I need all active hostname threats from IID.	https://platform.activetrust.net:8000/api/data/threats/state/host?Profile=IID&data_format=json

Single Correlation within TIDE

In some cases, there is a clear need to programmatically ask about a specific indicator. As examples, this might be in response to:

- In response to a SIEM rule trigger.
- A result of traffic analysis on web proxy logs where an indicator raises concern.
- A mail gateway application generated a new set of suspect indicators to verify.

Both TIDE and Dossier offer the capability to query a single indicator and instantly derive additional related intelligence. TIDE will tell you correlation with large indicator sets and Dossier will provide deeper context.

The Info I need	The Call to Make
I need to search against IID for a given hostname.	<code>https://platform.activetrust.net:8000/api/data/threats?host=jnnuvinskycatteredifg.com&profile=IID&data_format=json</code>
I need to search for an IP address across all my available sources, in CEF format.	<code>https://platform.activetrust.net:8000/api/data/threats?type=ip&ip=81.196.20.134&data_format=cef&rlimit=100</code>

Now The Network Context....

“What’s on my Network?”

You cannot protect or defend what you cannot see...

CIS Critical Security Controls - Version 6.0

To learn more about the CIS Critical Security Controls and download a free detailed version please visit:
<http://www.cisecurity.org/critical-controls/>

CSC #1

- CSC 1: Inventory of Authorized and Unauthorized Devices
- CSC 2: Inventory of Authorized and Unauthorized Software
- CSC 3: Secure Configurations for Hardware and Software on Mobile Device Laptops, Workstations, and Servers
- CSC 4: Continuous Vulnerability Assessment and Remediation
- CSC 5: Controlled Use of Administrative Privileges
- CSC 6: Maintenance, Monitoring, and Analysis of Audit Logs
- CSC 7: Email and Web Browser Protections
- CSC 8: Malware Defenses
- CSC 9: Limitation and Control of Network Ports, Protocols, and Services
- CSC 10: Data Recovery Capability
- CSC 11: Secure Configurations for Network Devices such as Firewall Routers, and Switches
- CSC 12: Boundary Defense
- CSC 13: Data Protection
- CSC 14: Controlled Access Based on the Need to Know
- CSC 15: Wireless Access Control
- CSC 16: Account Monitoring and Control
- CSC 17: Security Skills Assessment and Appropriate Training to Fill Gaps
- CSC 18: Application Software Security
- CSC 19: Incident Response and Management
- CSC 20: Penetration Tests and Red Team Exercises

CSC #1 – “Actively manage (inventory, track, and correct) all devices on the network so that only authorized devices are given access, and unauthorized and unmanaged devices are found and prevented from gaining “access.”

Source: <https://www.sans.org/critical-security-controls>

Know what is where - this does not work...

B4 fx cbd-rtr-hsrp


	A	B	C	D	E
1	Network	192.168.0.1/24			
2	VLAN	85			
3	IP	Host	Description		
4		1 cbd-rtr-hsrp	HSRP for VLAN 85		
5		2 cbd-rtr01-fe03	FE03 cbd-rtr01-fe03		
6		3 cbd-rtr02-fe03	FE03 cbd-rtr02-fe03		
7		4 Leave for network managers			
8		5 Reserved			
9		6 Reserved			
10		7 Reserved			
11		8 Reserved			
12		9 Reserved			
13		10 cdb-prt-flr01-01	First floor printer 1 aka marketing		
14		11 cdb-prt-flr01-02	First floor printer 2 aka finance		
15		12 Free			
16		13 Free			
17		14 Free			
18		15 Start DHCP RANGE			
19		16 DHCP RANGE			
20		17 DHCP RANGE			
21		18 DHCP RANGE			
22		19 DHCP RANGE			
23		20 DHCP RANGE			

192.168.1.0 192.168.2.0 192.168.3.0 192.1

```
jim — -bash — 80x24
PING 192.168.0.13 (192.168.0.13): 56 data bytes
64 bytes from 192.168.0.13: icmp_seq=0 ttl=64 time=96.012 ms
64 bytes from 192.168.0.13: icmp_seq=1 ttl=64 time=82.121 ms
64 bytes from 192.168.0.13: icmp_seq=2 ttl=64 time=33.851 ms
64 bytes from 192.168.0.13: icmp_seq=3 ttl=64 time=2.855 ms
64 bytes from 192.168.0.13: icmp_seq=4 ttl=64 time=81.400 ms
--- 192.168.0.13 ping statistics ---
5 packets transmitted, 5 packets received, 0.0% packet loss
round-trip min/avg/max/stddev = 2.855/59.248/96.012/35.187 ms
swampy-mac:~ jim$
```


... and does not scale
for complex environments...

What does work?



Security Orchestration

Accelerating Incident Handling and Response with Automation



Context to Prioritize Remediation



DHCP

Device Audit Trail and Fingerprinting

- Device info, MAC, lease history



IPAM

Application and Business Context

- "Metadata" via Extended Attributes: Owner, app, security level, location, ticket number
- Context for accurate risk assessment and event prioritization



DNS

- Malicious activity *inside the security perimeter*
- Includes BYOD and IoT devices
- Profile device & user activity

**Protect Users Everywhere - On-Premises,
Roaming and in Remote Office/Branch Office**

**Using an On-Premises,
or an hybrid On-Premises+Cloud
or a pure Cloud approach**

Infoblox Secure DDI the Switzerland of IT Security





Grazie

gderisi@Infoblox.com

