



Best Practices Guide

Federal Government

SIGNAL STRENGTH

135.09



Welcome to this Best Practices Guide from DomainTools. This reference offers insights into the cyber threats facing the Government sector, what the landscape looks like for defenders, and how security teams are making effective use of adversary infrastructure analysis to gain an edge. The Guide consists of four sections:



The current threat landscape



Successes and limitations of common defensive strategies



Ĥ

The value of DNS and DNS-adjacent data in Zero Trust initiatives and in adversary analysis, and why DomainTools is a leader in this space

How government-sector security teams are solving important security problems with DomainTools

At the end, we offer links to various resources to help you learn more about DomainTools offerings.

The Cyber Threat Landscape

It is a truism almost to the point of cliché that the threat landscape is evolving, and going into the mid-2020s, this is as true as ever. **Certain trends, however, suggest that this evolution may be quickening relative to earlier times.** In particular, the introduction and widespread adoption of large language models (LLMs) such as ChatGPT and its peers seem to be accelerating the pace of change, with these tools offering phishers more convincing lures, Business Email Compromise (BEC) actors the possibility of deepfake voice impersonations of trusted colleagues, and malware authors new ways to craft variants that can bypass many detection technologies. The Government sector is in no way immune to these threats; CISA has published advisories on LLMs, deepfakes, and malware authoring.

And machine learning and its derivatives are not the only rapidly changing threats. Ransomware and APT groups are also proving agile and polymorphous, either with changes to their tactics, techniques, and procedures (TTPs), changes to their organizational structures, or both (see DomainTools <u>reporting</u> on the most prolific ransomware families). Moreover, the discovery and exploitation of new vulnerabilities, some of them critical, is also occurring at a brisk pace.

But no matter how sophisticated or unique the cyber threat, **something virtually** all of them have in common is that they rely on the use—or abuse—of Internet infrastructure that is observable, comparatively static, and often rich in contextual information that defenders can, and do, use to considerable effect in aligning defenses with confirmed or suspected adversaries. This Guide will show you how.



Current State— What's Working and What's Not

Security technologies and practices have not stood still while the threats evolved—they have evolved right along with them. And security innovations have not always lagged threats, with seemingly daily advances that improve detection, defense, visibility, and remediation. Nevertheless, breaches and compromises roll on. A reasonable assessment of a security organization that is doing things "right" might be necessary but not sufficient: the steps today's defenders take are generally good and prudent ones, and in some cases, truly stellar work is being done and shared. Yet, the teams that have the strongest postures and best track records will be the first to admit that they are anything but invincible. So the goal is not perfection; it is to make reasonable, cost-effective advances that make measurable positive differences in outcomes.

Technology or approach	Gaps that remain
Reputation lists and observation-based threat intelligence feeds	Newly-registered domains are generally a blind spot to these technologies because reputation feeds are built on observed harm in the wild, or analysis of traffic already in the environment
Deep packet inspection, sandboxing, heuristics-based rules	Susceptible to novel techniques or malware or non- technical social engineering techniques
Forensic analysis of domains or IP addresses that touched the protected environment	Threat actors usually control more infrastructure than what is initially observed. If forensics do not account for this additional infrastructure, it may cause future harm.

Many government agencies deploy Zero Trust policies along with a good set of defenses, which typically include relatively late-model products in the realms of network defense, host defense, identity and access management, visibility and situational awareness tools, cyber threat intelligence, and orchestration and automation of some or all of these—and this is an incomplete list. But it's worth pointing out two truths about technology stacks that incorporate these tools:

- 1. Malicious activity is still proceeding and frequently succeeding
- 2. Almost every technology mentioned operates, at some level, within the framework of DNS

This second point is what we will explore next.





The Importance of Adversary Infrastructure Analysis

Because today's SOCs, fusion centers, intelligence teams, law enforcement centers, and any other entities entrusted with cyber defense are moving at such a rapid pace, with constrained staffing, it is fair to ask why resources should be expended on infrastructure analysis. After all, that time has an opportunity cost; each minute or hour spent on such analysis cannot be spent on other tasks.



While there is no individual "right" answer, our work with practitioners around the world has led us to certain well-tested axioms:

- Everything that happens on the Internet uses domains and/or IP addresses. Malware families come and go; network- or protocol-based attacks have their moments in the sun; online social engineering attacks become prevalent and supplant earlier ones; but amid all of these cycles and evolutions, the fundamental infrastructure on which the vast majority of them rely remains relatively familiar: domains and IPs. This is a weak point for adversaries, because it means that...
- There are almost always clues available; It is very difficult for adversaries to cover their tracks completely. Some are more adept than others, of course; but staying all the way in the shadows of the Internet is challenging, time-consuming, and often works against the scale and speed that bad actors depend on to make crime pay.
- You can tell a lot about a domain by the company it keeps. Malicious domains tend not to be "lone wolves." Any malicious campaign designed to have a significant impact will almost universally rely on multiple objects (domains, IPs, certificates, etc). Moreover, these components almost always have some features in common with each other, either for technical reasons or because the actors controlling them re-use certain patterns—or both. These relationships often shed a lot of light on the nature of any individual constituent part, in much the same way that an individual tile might not mean much until the observer can see the larger mosaic in which it sits.

Adversaries make mistakes. Defenders know that they have to be right 100% of the time, and the attacker only has to be right once. However, adversaries face this same asymmetry. If they want to ensure that they can't be identified or blocked, they have to avoid leaking identifying or connecting information. That is not particularly hard when the actor is running a single domain, but when they scale that to dozens, hundreds, or thousands, the odds of a leak become much greater. Defenders can and do use actor OPSEC (operational security) slip-ups to their advantage.

Top-performing security teams around the world operate on these axioms daily.



7

DomainTools

DomainTools provides comprehensive Internet intelligence to security practitioners and advanced security teams. The solutions are used to identify external risks, investigate threats, and proactively protect organizations in a constantly evolving threat landscape. DomainTools constantly monitors the Internet and brings together the most comprehensive and trusted domain, website, and DNS data to deliver context and machine learning-driven risk analytics in near-real time, providing critical tools and services for the following use cases:

Threat

Intelligence

Detect relevant indicators earlier in their lifecycle to identify and disrupt incipient attacks.

Zero Trust Internet policies

Detect or deny connections to newly-created and/or high-risk domains.

Forensics and Incident Response

Respond to and triage potential incidents with confidence and speed.

Phishing and Fraud Prevention

Know if and when malicious domains and infrastructure are spoofing your assets before they can cause damage.

Threat Hunting

Discover indicators of compromise (IOCs) and malicious infrastructure that may be targeting your network.

Enrichment

Enrich homegrown or third-party security applications with effective Internet intelligence.

The DomainTools **Iris Internet Intelligence Platform** is made up of three components. Iris **Detect** provides a near real-time Internet infrastructure detection, monitoring, and enforcement platform and API; **Iris Enrich** is a robust API that includes Whois, DNS, SSL certificate, and risk scoring elements to enrich indicators at scale; and **Iris Investigate** provides a platform and API that supplies and maps domain intelligence, risk scoring, and industry-leading passive DNS data.

Farsight Newly Observed Domains and **Newly Observed Hostnames:** feeds taken from the Farsight worldwide passive DNS sensor array for the earliest possible detection of emerging threat infrastructure. DomainTools also provides **Threat Intelligence Feeds** that can be integrated into threat intelligence platforms and other tools to provide predictive domain risk scoring, hotlists, newly discovered hostnames and domains, and more.

DomainTools Monitors can provide alerts to security teams that signal early warnings of when adversaries are preparing to attack or when known campaigns are evolving.

Farsight DNSDB is a comprehensive passive DNS near real-time and historical database of global internet infrastructure data, that can be accessed and queried by DomainTools customers and integrated into tools through an API to help reduce risk.

Zero Trust and DomainTools

Increasingly, security and risk management personnel are implementing various initiatives centered on the concept of Zero Trust, whether because of regulatory guidelines such as <u>Executive Order 14028</u> Sec. 3, CISO guidance, or both. Internet intelligence from DomainTools can play an important role in such initiatives, because connections from the protected environment to unknown (and therefore untrusted) infrastructure represent a genuine and pervasive risk. DomainTools offers a variety of tools and data that help security teams:

- Identify and/or block connections to newly-created domains
- Build context around adversary-controlled infrastructure

- Identify clusters of malicious activity based infrastructure patterns
- Monitor emerging attack campaigns as the adversary develops them

Rather than implicitly trusting domains unless or until there is a reason to block them, it is becoming increasingly popular to block all domains younger than a certain threshold by default. This is effective because many web proxies, SMTP proxies, and other controls have no way to categorize domains when they are first registered; such controls often rely on reputation scoring or analysis of served content. When a domain is initially registered, some time may pass before it is provisioned and thus able to be assessed by traffic or application analysis tools.

DomainTools data, as enrichment in SIEM, SOAR, or TIP, or proprietary tools, can enable newly-created domain alerting or blocking. Similarly, DomainTools Risk Scores provide a means of identifying domains that may not be on a typical observation-based denylist as yet, but which may represent a threat and thus should not be trusted. DomainTools is already trusted by US Federal/ Government agencies, as well as many private sector enterprises, to help meet Zero Trust architecture needs around threat intelligence. The foundational objective of Zero Trust is to prevent trouble before it occurs. However, full prevention of incursions or dangerous connections is not achievable in real-world environments. When trusted assets have connected to malicious infrastructure, DomainTools enrichment and investigative tools can be applied by IR (incident response) or forensic teams. In such scenarios, DomainTools data may provide insights that cause the IR team to "retroactively revoke trust" of a given domain—that is, a domain that was not flagged or blocked previously but which, thanks to DomainTools data, is now seen to be dangerous. Several principles of Zero Trust are addressed by DomainTools products and data in the following specific ways:

Do not trust unknown resources

- Use Iris Enrich, or Farsight Newly Observed Domains or Newly Observed Hostnames to flag newly observed domains as seen in event logs.
- Use Domain Risk Score to flag high-risk domains.
- Use I**ris Detect** or **Newly Observed Domains** or **Newly Observed Hostnames** to identify domains that spoof particular keywords such as the organization's name, or its close associates or vendors, in order to flag risky infrastructure before it is weaponized (i.e. before it has a chance to appear in the protected environment).
- In any of the above scenarios, domains flagged by DomainTools can then be incorporated into custom denylists or other security controls. SOAR or proprietary scripting can automate these processes.

Apply Least Privilege principles to user access of Internet-based resources

- Apply Least Privilege principles to user access of Internet-based resources.
- For security controls with a spectrum of dispositions available, DomainTools enrichment and/or risk scores may be referenced to calibrate the level of control.
 - For example, in an email filter, deny any connections from domains younger than a given value, or with risk scores above a given threshold; allow connections but disable attachments and links for domains with ages or risk scores within a designated band of age/score values, etc.
 - In a web filter, deny connections to domains younger than a given value or with high risk scores; place an interstitial warning for domains in a slightly lower risk/age band, etc.

Monitor the environment in real time

 Use machine-scale Iris Enrich and/or Farsight DNSDB enrichment to identify young and/or high-risk domains (or domains meeting more customized criteria such as hosting geography, registrar, ASN, etc).

Proximity OVERALL SCORE 63 62 8444 91 91	TO Proximity OVERALL SCORE OVERALL SCORE 63 62 20 70 PREMER Hall WARP NEAR NETWORK TY	amain Profile			E , 11		
TO Proximity OVERALL SCORE OVERALL SCORE 63 62 20 PRRMKN HIZ WARP NEAH	OVERALL SCORE OVERALL SCORE 63 62 20 70 PREMERT Had Water TEAM TEAM FOR THE SCORE		Rai Sco				
OVERALL SCORE OVERALL SCORE COVERALL SCORE COVERAL COVERALL SCORE COVERALL SCORE COVERAL	OVERALL SCORE	70		Prox	imity		
63 62 20 70 PREMIER NEAR NEAR NEAR	63 62 20 70 HINKET HILIWARE THE HINKETY	OVERALLECON		OVERAL	LISTOPRE		
		63	62	20	70		
			10 M IPP	974H			

Three Common SOC Use Cases and Where DomainTools Fits In

Each of the following is a summarized sequence giving an example of how security teams use certain DomainTools <u>products in common workflows</u>. The exact use case will have slight variations for every organization.



Threat Hunting:

- Ingest IOCs from trust group, threat actor report, or other source
 - Alternatively or in addition, ingest newly discovered domains from **Iris Detect** or **Farsight Newly Observed Domains or Newly Observed Hostnames.**
- Run a query on the identified IOCs in Iris Investigate; pivot and expand to uncover additional.

connected infrastructure; save query for expanded set as an Iris Investigate hash (saved query).

- Retro-hunt for presence of any of the expanded indicator set in earlier logs or alerts.
- Set SIEM or security control alerts for traffic involving any of the expanded indicator set.
- Re-run Iris Investigate hash (a form of stored query) daily to pick up new indicators matching the established pattern.
- Identify and investigate hits on any of the indicators; hand off to analyst or IR teams as appropriate.

Incident Response:

Upon firing of any alert with a high enough severity that the team decides to investigate:

- Identify any external domains or IP addresses associated with the alert.
- Search on the domain(s) or IP(s) in Iris Investigate; pivot and expand to uncover additional indicators; save query for expanded set as an Iris Investigate hash.
 Some teams will also pivot in **DNSDB** for additional connections.
- Retro-hunt for presence of any of the extended indicator set in earlier logs or alerts.
- Any traffic flows to any of the extended indicator set are now immediately suspicious.
- The full scope of traffic to any of the extended indicator set may be considered part of the incident.
- Set SIEM or security control alerts for traffic involving any of the extended indicator set.



Security Analysis:

- Ingest IOCs from trust group, threat actor report, or other source
- Search on IOCs in Iris Investigate; pivot and expand to uncover additional indicators; save query for expanded set as an Iris Investigate hash
- Analyze extended infrastructure (in other tools such as Censys, Shodan, etc) for clues about additional TTPs that may be telegraphed by it
- Re-run Iris Investigate hash daily to pick up new indicators matching the established pattern; detect and investigate hits on any of the indicators
- Use Iris Detect to monitor names and brands of vendors for potential imitations, e.g. Microsoft365, Salesforce, etc;
- When spoof domains are discovered, work with Detection Engineering to set up monitoring of any outbound connections to the spoof domains (Or, set up blocking rules ahead of time for the spoof domains)

Additionally, many SOC personnel use the Iris Enrich API to decorate domains appearing in popular SIEM platforms such as Splunk, Chronicle, Microsoft Sentinel, and others. Such enrichment allows analysts to quickly assess connections made to any domains identified as high-risk (according to the DomainTools Risk Score), newly-created, or both. Armed with this information, analysts can then make informed decisions about which domains might merit further investigation.



Example:

Government and Banking Spoof Campaign

Registering spoofs of legitimate domains is often one of an adversary's first moves in creating a phishing or watering hole attack, or creating later-stage servers for command and control (C2) or data exfiltration; in any of these activities, the domain names are intended to deceive end-users or security personnel. And while ransomware regularly makes headlines and features in CISA advisories, phishing and Business Email Compromise still account for larger financial losses on an annual basis, <u>according to the FBI</u>. Thus, it is important for organizations to keep tabs on infrastructure intended for such activities. A recent (as of this writing) cluster of activity involved a number of domains that spoofed U.S. government sites as well as financial institutions.



Monitoring the string "ustreasury" for spoofs with Iris Detect (or filtering for these terms within Newly Observed Domains) would have turned up the following domains, which have been exported to Iris Investigate for further examination in this example:

L			new reactions		And the second of		
1	weeks the proof	[section for]	- 802-001 <u>100</u>	The second second second	10.5(5NP)	a sussion.	300 B.M.
	K mar A service set of	1210125800		an a	nanoon on an tari ku din too	97 17 0000	process que transférier de la conference (C. 1997) Maria Com process que constituir de la conference process que constituir de la conference
						Les anno 1 Les anno 1 Les anno 1 Les anno 1 Les anno 1	saandadadan waxa
		inter y transmitter. y annual del y -	ina o in canada de la contra com	6337 F235			
						UIVOV 1. Kali ilas 5. C. Oliv 1. Kali ilas 5.	
1	armoninde mes com	State Street	A classify and	IF BECHYTER ALE CO	Bonis-a	B SPORMATCH	/2004.001
	Алине 🛛 🖬 В малениян Списток на на	CALC: NO.	4 - 1 - 1 - 1 - 1 - 1 - 1 - 1 - 1 - 1 -	4.1.14 A.1.14	account and but fully a set of the	errow a worker i	ne si a Mageri i ant A DDTTTege i Ione
						E POR E	A STAR PAGE - Lane - 1987 Pages Low
					nord of the Charles Inc.	Kass, Raill - c cover 0 Kass, Raill - - c cover 0 Kass, Raill - - c cover 0 c cover 0	eran sekaran kerana

Since we know that these are not legitimate domains, we have tagged them as spoofs. But to carry out a further investigation, we examined what other domains might be closely connected to these. Iris Investigate shows us that around 20 domains all shared the same street address (which the registrant has invented, as the given street does not match any known street name in Brooklyn or anywhere else):

	Contact Stre	ret:
	435 stella berkly rd.	brooklyn ny
	99 Arg Rei 192	Art Alle
	DOMAIN	RISK SCORE -
	warldbankwdc-us.com	-60
	worldbankefa.com	100
	whitehouseus-gov.com	100
	ustreasurydept us.com	100
	uschbankilmited.com	. 100
	us-depthomelandsec gov.com	
	ubaamericaa.com	100
	stgxx.com	160
	headoffice.com	
	finimasyfinitiosupauthority.co	em 🗰
	chreig-gev.com	*
	al-rayaribank-sk.com	100
	1shé-tx.com	
	laba-us.com	100
	usdepartmentafdefease-gow	oom 🖷
	nitctioni-ni.com	
	federalineserveboard.com	
	fbi-usgor.com	Ť.
	usbancorprivatebanking.com	
0 *	© corporateleadqtrs.com	

• Example: Government and Banking Spoof Campaign

We can see that the registrant demonstrates a clear pattern of activity, repeatedly registering new domains with this fictitious address in the Whois records, so we tag these domains with "Activity Cluster." We can also use the Visualization panel in Iris Investigate to quickly see patterns or clustering within the set of domains.



We now have some options available to act on the information we have just developed. We can:

- Use the Iris Investigate API to create a recurring query for any new domain registrations matching this street address, since this actor appears to adhere to this pattern
- Create alerts for any traffic from the protected environment to any of the domains
- Create blocking rules for the domains and/or the IP addresses associated with them
- Share the domains and/or IP addresses with a trust group and/or law enforcement

Finding the domains tied to the original two spoofs of US Treasury sites would not have been possible without connected-domain data, and the additional context provided by **Domain Risk Scoring** helps increase our confidence that the domains in question are malicious.





Organizational Benefits

Because of the benefits of techniques such as the one illustrated above, DomainTools customers consistently <u>report</u> significant organizational wins in the form of cost savings, improved detection rates, analytical efficacy, and more. According to Enterprise Strategy Group, DomainTools customers identified as many as 83% more malicious domains with DomainTools than with alternatives, and detected malicious domains up to 96% earlier than with industry-standard blocklist sources.



Figure 3. Blended Customer-reported Metrics for DomainTools versus Alternative Methods and Industry-standard Blocklists





Source: Enterprise Strategy Group, a division of TechTarget, Inc.

DomainTools customers also report that their teams were more efficient, with some reporting savings of between 1.5 and 2 hours per day per employee.

"DomainTools gives us the earliest and most updated feed of newly created and updated domain and DNS infrastructure—so the second someone creates a domain, within five minutes, we know about it."

"Out of 1,000 domains determined to be malicious by Iris Detect, 68% did not appear in any other industrystandard blocklist. Of those that were detected elsewhere, Iris Detect and Investigate detected three days earlier on average, with most being detected within a three- hour period."



Conclusion and Additional Resources

The great majority of cyber threats today use DNS, and leave traces that can be exploited for forensic and predictive purposes. DomainTools has amassed the world's largest datasets around Internet infrastructure, and for many years has leveraged the data to produce detection, enrichment, and investigative tools deeply informed by close work with practitioners in many of the world's most sophisticated security organizations. We believe that the data, tools, and methods described here have the potential to aid in the implementation of security initiatives and, more broadly, to make a meaningful contribution to the protection of organizations around the world.

Recommended Resources:

- Schedule a personalized <u>demo</u> of DomainTools products
- Zero Trust and DomainTools
- Formulating a Robust Pivoting Methodology
- Valuable Datasets to Analyze Network Infrastructure
- Using Infrastructure Analysis to Get Ahead of Attacks in Cyber Defense