SUBDOMAIN INFRINGEMENT

AN UNSEEN THREAT

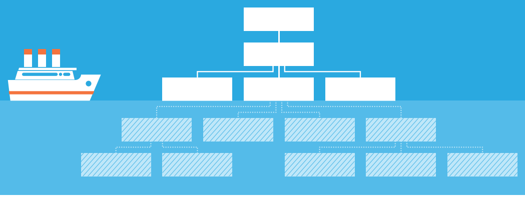

Domain infringement is when threat actors use brand names within illegitimate web domains to imply affiliation with a brand to deceive end users about who's behind the content they see on a site. They use this exploitation of trust as a lure to phish for sensitive data, distribute malware, promote scams, generate revenue from ads on parked domains, and drive monetizable traffic to other sites.

Most brands are well aware of these risks and have an internal program in place to identify and prevent domain infringement, but typically, these only cover searching for infringement in parent domain names, i.e., example.com, leaving all fraudulent subdomains, i.e., somethingelse.example.com, in use by threat actors undetected.

But therein lies a critical problem: infringing subdomains are just as dangerous and destructive to a brand and an organization's security posture in the hands of threat actors as infringing parent domains. Ignoring instances where a brand is being abused in subdomains can be severely detrimental to the organization and its employees and customers.

In this white paper, we'll show you why subdomain infringement detection is critical for information security and fraud protection, as well as:

- Show why subdomain infringement is even more dangerous than parent domain infringement
- List the reasons why subdomain infringement slips through the cracks on an organizational level
- Outline the techniques used by threat actors to efficiently acquire domains from which to launch large-scale subdomain infringement campaigns
- Explore how prevalent the problem is, using five leading financial service companies as real examples that show how dangerous subdomain infringement can be
- Explain the steps you can take to turn domain infringement detection into comprehensive domain threat detection

More Dangerous than Parent Domain Infringement

Brand impersonation and infringement is a powerful tool for third parties seeking to enrich themselves at the financial and reputational expense of the brands they target. Threat actors can leverage the names of trusted entities to drive monetizable traffic to their sites or deceive users into sharing sensitive information, falling for scams, or downloading malware. While the relationship between parent domain infringement and brand impersonation may seem more direct—and thus more effective—in reality, subdomain infringement is arguably more dangerous to brands for three key reasons:

1. There is an essentially infinite number of subdomains: Because multiple, non-unique subdomains can be added to any parent domain, the scale of domain infringement can be much larger than that of parent domains. Though perhaps still reasonably large, the number of infringing parent domains in existence that relate to a given brand is finite, whereas the number of potential infringers in subdomains is practically enormous.

2. Threat actors don't need to typosquat: Because parent domain names are unique and relatively limited in number, the fact that subdomains can use a verbatim brand name without having to typosquat increases the strength of deception. Because spelling mistakes are easier to identify than where the breaks are in a long domain name, they can easily clue in users that a site is fake. For example, "brandname.com.tk" may be harder for users to spot as a fake at a glance than "br4nddname.tk".

3. Infringing subdomains can hide: A single parent domain can have multiple subdomain levels, which hide the true parent domain name by making it appear as part of a long URL path. For example, the parent domain "abc.xyz" is obscured by multiple subdomains in the example domain name "brandname.com.secure-login-auth-user.abc.xyy".

Who's Dropping the Ball?

If subdomain infringement carries the same risk as parent domain infringement—and on a larger scale—then why do organizations continue spending time and resources detecting parent domain infringement while essentially ignoring subdomain abuse? Why haven't they adapted their defenses to address this increasingly popular method of attack?

The answer is subdomain infringement often falls in a gray zone of organizational responsibility, where no one is explicitly tasked with looking for it. There are two likely reasons why:

1. Detecting subdomains is really hard: The technical hurdles involved in detecting subdomain infringement are greater than those for detecting parent domain infringement. The WHOIS system, which stores information on all registered parent domain names on the Internet,

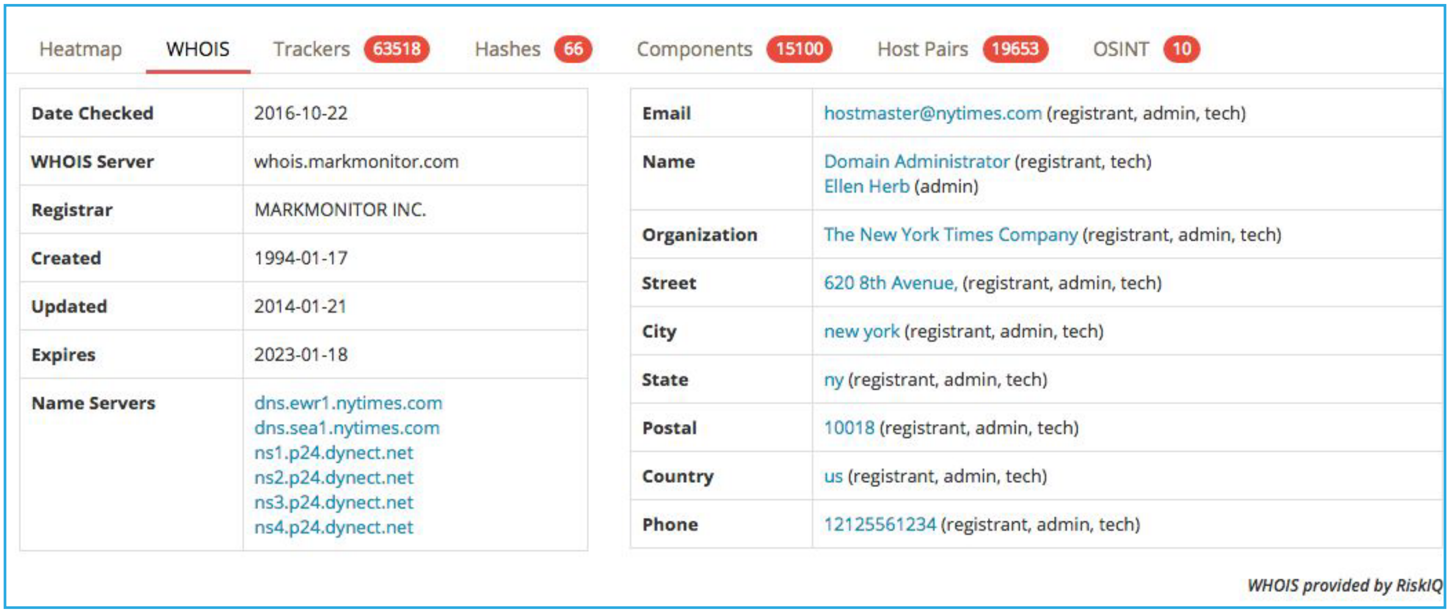

Date Checked	2016-10-22
WHOIS Server	whois.markmonitor.com
Registrar	MARKMONITOR INC.
Created	1994-01-17
Updated	2014-01-21
Expires	2023-01-18
Name Servers	dns.ewr1.nytimes.com dns.sea1.nytimes.com ns1.p24.dynect.net ns2.p24.dynect.net ns3.p24.dynect.net ns4.p24.dynect.net

Email	hostmaster@nytimes.com (registrant, admin, tech)
Name	Domain Administrator (registrant, tech) Ellen Herb (admin)
Organization	The New York Times Company (registrant, admin, tech)
Street	620 8th Avenue, (registrant, admin, tech)
City	new york (registrant, admin, tech)
State	ny (registrant, admin, tech)
Postal	10018 (registrant, admin, tech)
Country	us (registrant, admin, tech)
Phone	12125561234 (registrant, admin, tech)

WHOIS provided by RiskIQ

Fig-1 WHOIS record for the domain name nytimes.com on passivetotal.org, RiskIQ's free threat research tool

provides a simple, cost-effective way to find brand-infringing parent domains. Below, you can see the WHOIS information for Nytimes.com, which provides basic information about that domain, including who owns it, when it was created, and when it expires.

However, there's no similar universal system for searching subdomains. Passive DNS data offers a way to detect infringing subdomains, but it requires vast infrastructure and technical investment to collect a large enough repository of data to be useful for this purpose. To collect PDNS, a sensor is installed on the local network and set up to receive and record DNS requests as they happen. However, the sensor will only record DNS traffic that occurs on that network, and not the entire Internet. So, to collect the requisite amount of PDNS coverage, you need to have sensors on a huge number of networks and aggregate all that data.

Until relatively recently, no vendors were able to offer this type of data, and consequently, there was no need to ask anyone to solve a problem for which the necessary tools to do so did not exist. Even now, many organizations have not yet begun to take advantage of newly available data sources and adopt subdomain infringement detection as a best practice of their domain infringement programs.

2. It's viewed as a legal problem, not a security problem: Despite the potential for brand infringement to be leveraged by threat actors to carry out cyber attacks, traditionally, domain infringement detection and mitigation has been a function of either brand protection or legal teams, not security teams. As a result, most anti-infringement programs tend to focus on the detection of infringement in parent domain names only, especially because infringement in parent domains is weighted more heavily than infringement in subdomains in trademark infringement claims.

Security and Threat Intelligence teams are far more likely to have access to and be familiar with PDNS data than legal or brand protection teams, but most often focus on using it exclusively to investigate other issues and may not be aware of its value in detecting subdomain infringement. This means the people who care about subdomain infringement don't have—or may not be aware—of the tools needed to solve that problem.

A Landscape Ripe for Subdomain Infringement

There are several reasons why using branded subdomains vs. parent domains has additional advantages for threat actors, most of which have to do with the landscape of today's web. They've identified an array of lucrative techniques for leveraging malicious or fraudulent subdomains with which they can launch low-risk, large-scale attacks.

Cheap and Free Registration of Parent Domains: Before 2013, the number of top-level domains (TLDs) in the world was manageable. But, in 2012, ICANN began accepting applications for creating new TLDs in response to wild demand. Since then, the amount of total domain names in existence has exploded. Now, there are more than 24 million domain names and counting, registered across roughly 1,200

Freenom: A tool for Threat Actors

Freenom, the world's first provider (registrar) to give free domain names away, is frequently leveraged by threat actors. Freenom may fancy itself the Robin Hood of domains, helping expand access to custom domains and email addresses, but the more unsavory set may be getting the greatest benefit from its service. The price chart below shows how anyone can register a variety of top-level domain names for free as long as they have a valid email address:

TOP LEVEL DOMAINS	REGISTRATION	RENEWAL	TRANSFER
TK	**$0.00**	**$0.00**	**$0.00**
ML	**$0.00**	**$0.00**	**$0.00**
GA	**$0.00**	**$0.00**	**$0.00**
CF	**$0.00**	**$0.00**	**$0.00**
GQ	**$0.00**	**$0.00**	**$0.00**

Fig-2 Freenom pricing list

There are upwards of 5,000 services offering free domains (and growing rapidly), which is yet another reason why we expect this sort of abuse to increase in the future. The graph below shows one example of a campaign tracked by the RiskIQ Research team where threat actors used domains registered on Freenom to ultimately direct users to the Neutrino exploit kit.

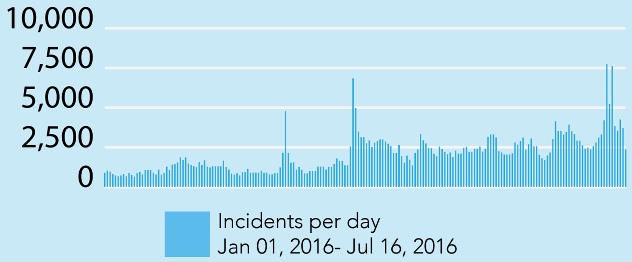

Fig-3 A threat campaign using free domains registered on Freenom to direct traffic to exploit kits new TLD extensions, with new ones opening up all the time.

By editing the DNS record of the domain name, the owner of a parent domain name can create new subdomains, hosting unique web content on each. And, since subdomains can be more than one level deep, a single domain name could theoretically have over a hundred different subdomain levels, all completely undetectable via WHOIS records. This creates an endless infrastructure with which threat actors can carry out attacks, at the cost of about what you can find between your couch cushions.

Easy to scale: There is a limited number of infringing parent domain names available for any given brand, so these domains are extremely valuable and may cost more to purchase. By contrast, brand-related subdomains can be appended to any parent domain, regardless of what it is. This means threat actors can launch a greater number of attacks in a shorter period and don't have to worry about finding another viable domain afterward in order to attain the same level of deception through brand impersonation. This also means a single parent domain can be used to target multiple brands, which further cuts down operating costs and ease of attack.

Advanced tactics: Acquiring a parent domain and appending infringing subdomains to it may be easy, but why use your own domain if you can add subdomains to someone else's?

Domain shadowing is an advanced method that threat actors use to leverage subdomains for their campaigns. In this technique, threat actors steal credentials to domain registrant accounts associated with existing, legitimate websites. This is somewhat aided by the availability of cheap and free domain registrations since, with more domains in existence overall, there are more targets for threat actors to choose from, and the owners of free and cheap domains may be less inclined to protect them.

Once they gain access, the threat actors add a small number of unauthorized subdomains to parent domains owned by the registrant, and use them to host their own malicious content. Whereas when the threat actor owns his or her own domains they can add large numbers of subdomains and create a huge number of malicious sites, with domain shadowing, it is essential for the threat actor to create only a small number of subdomains in order to prevent ruining the good reputation of the parent domain as a whole or the domain owner from finding out in order to prolong the lifetime of the attack. Though a smaller number can be added per parent domain, these bad subdomains are extremely difficult to detect because they are mixed in with good subdomains on otherwise reputable parent domains with no discernable connections to any known threat infrastructure, and thus, may stay active for a longer period of time.

Little risk of getting caught: Brands are proactively monitoring for registrations with similar spelling, "branddname.tk" for example, even before any website content is associated with that domain. However, registering something like "abc.xyz" is unlikely to raise the alarm from anyone. Even after "brandname" is added as a subdomain to "abc.xyz", most organizations have no way to detect it until it's too late and they start receiving complaints from users about deceptive sites that have been scamming them or distributing malware.

Using domain shadowing, threat actors can even stay anonymous and avoid using infrastructure of their own by hijacking a legitimate site and creating hidden subdomains underneath it to avoid detection and leverage the parent domain's good reputation.

Financial Services: A Lucrative Target

Any security programs that aren't looking for subdomain infringement are missing a significant piece of the puzzle—the piece most likely to be associated with the high-risk group of infringements that pose the greatest threats to an organization and its users.

RiskIQ recently took a sample of nearly 4,000 recent infringements, both parent and subdomain, across five financial services brands. We found that subdomains make up a sizeable chunk of domain infringements (about 25% of the total). And, even though subdomain infringements make up just a quarter of the total, they make up an overwhelming majority of the bad stuff—75% of malware and phishing instances identified within infringing domains were found when the infringement took place in the subdomain rather than the parent domain name.

Given this data, any given subdomain infringement that you find is significantly more likely to be a severe security risk than a given parent domain infringement. In our sample, the risk increase was ten-fold. The chance of a domain with brand infringement in the subdomain being associated with malware or phishing was 1000% that of domains with brand infringement in the parent domain only.

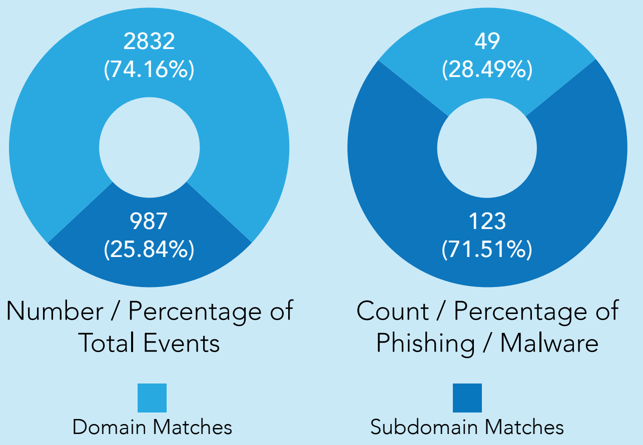

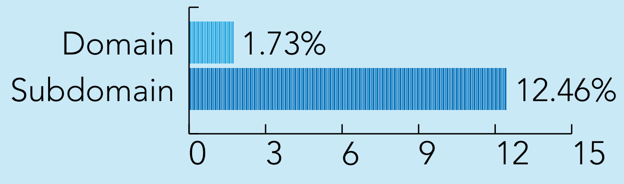

Proportion of events associated to phishing or malware within the set of parent domain infringements vs. subdomain infringements

Going Beyond WHOIS with Passive DNS

Layering Passive DNS data over WHOIS lookups tells a much more detailed story about the threat actors behind domain infringement and allows brands to consider subdomains as well as parent domains.

Passive DNS is a system of record that stores DNS resolution data for a given location, record, and period. This historical resolution data set allows analysts to view which domains resolved to an IP address and vice versa, as well as time-based correlation of domain or IP overlap.

That means Passive DNS can help determine when a new subdomain is first observed and analyze all newly observed hosts containing a brand-infringing subdomain in much the same way that monitoring new WHOIS registrations can provide this information for parent domains.

RiskIQ has one of the largest repositories of PDNS data in the world. The DNSIQ program allows organizations who contribute data from their own network to gain access to the central repository of data shared by all contributing members.

Domain Threat Detection and Monitoring

Adding crawling infrastructure and advanced security analytics to monitoring WHOIS and PDNS for brand-use in parent domains and subdomains turns domain infringement detection into domain threat detection. Crawling any websites associated with infringing domains gives security teams context around how an infringing domain is or isn't being used by a threat actor, and what level of threat it poses to the organization.

Crawls provide data about a site's content and behavior, including whether the site is malicious, phishing, contains brand names or logos in the page, or whether the site redirects visitors somewhere else. RiskIQ's sophisticated network of virtual users work as an emulated human user with a fully instrumented browser and algorithms to simulate human-like mouse movements and click behavior, storing the entire chain of events as forensic evidence for analysis during each crawl. This gives security teams access to updated threat data and connects infringement to other pieces in an attack chain, so researchers know precisely when an infringement became active in an attack, which infringements associate with live content, and where an infringing site is redirecting users. This way, they know how severe the threat is so they can prioritize their organization's response to incoming infringements.

Monitoring sites' content and behavior over time provides security teams with a full view of the threat lifecycle, showing the chain of events from an infringing subdomain getting initially created up through a malicious site getting hosted on that subdomain and then showing when the threat has been resolved and the malicious content taken down.

Domain Threat Correlation

Detecting and monitoring domain infringement and associated web content can be just the beginning. Correlating this brand infringement data with other external threat data sets such as social media and mobile app data can reveal a larger story about threats facing an organization. For example:

- Let's say that a fake social media account is claiming to represent your brand's customer support team. They might then fool user traffic by posting links like mybrand.secure-login-auth.abc.xyz, which drives them to a phishing page.

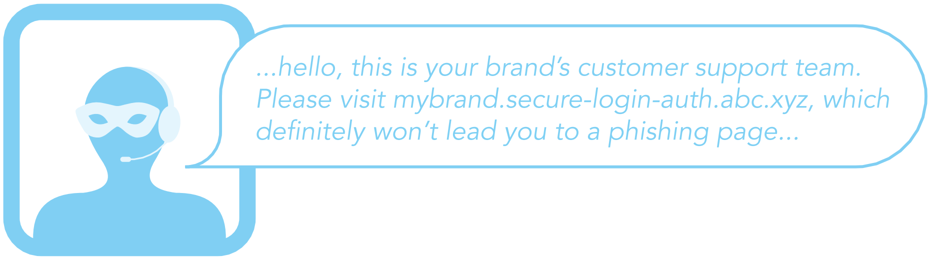
...hello, this is your brand's customer support team. Please visit mybrand.secure-login-auth.abc.xyz, which definitely won't lead you to a phishing page...

- Or, a threat actor might use an online ad to drive users to a website deceiving users into thinking `they're downloading your brand's official mobile app. However, what they're really downloading is mobile malware.

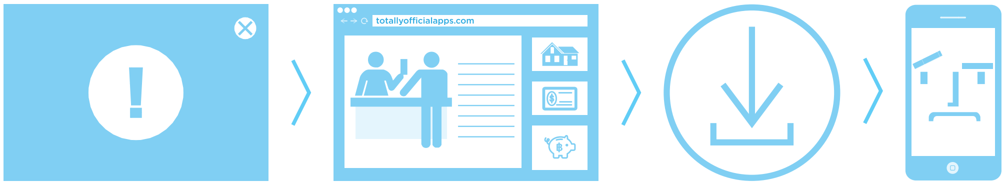

By continuously monitoring what, if any, web content is being served by sites containing subdomains that infringe on your brand—and to where they're sending traffic—RiskIQ can automatically correlate that information against threat data we collect elsewhere, such as mobile apps and social media channels, to uncover more details about how threat actors are targeting your organization and the level of risk posed by all the separate pieces of infrastructure based on their role in the larger threat.

Domain Threat Investigation and Incident Response

PassiveTotal is RiskIQ's free threat research tool with which you can conduct queries against RiskIQ's Passive DNS data set, as well as many others, to pivot off an indicator like an infringing domain to quickly uncover and preemptively block all threat infrastructure owned by that same actor—stopping potential threats from reaching your organization before they even happen. Additionally, you can use domain threat intelligence for proactive defense by ingesting RiskIQ's Security Intelligence

Services feeds to block suspicious traffic from all newly created domains or hosts in your firewall automatically until they reach a certain age.

Proactive Defense Against Subdomain Infringement

The age of a parent domain name or a specific subdomain under it is often one of the strongest indicators of reputation, and brand new subdomains on older parent domains are extremely unlikely to have legitimate business reasons to show up in corporate network traffic. As a result, they should always be viewed by your security team with a high degree of suspicion. To implement this best practice at your organization, RiskIQ offers feeds of newly observed domains and newly observed hosts to integrate directly into your firewall.

With RiskIQ's PassiveTotal and Security Intelligence Services, your team can use threat intelligence gleaned from domain infringement data to proactively identify and block threats as well as find trends over time and potentially find campaigns targeting specific regions or industries from certain brands showing up in subdomains, which can enhance other investigation activities. For instance, monitoring for new hosts created on your own company domains is a way to identify a potential compromise.

Why RiskIQ

RiskIQ's industry-leading External Threat Management platform combines huge stores of threat data with a sophisticated and powerful crawling infrastructure and a vast global proxy network to give organizations visibility into their security posture outside the firewall.

As the only company that monitors digital risk across all channels, Forrester Research recognized RiskIQ as a Leader in the category-defining Forrester Wave™: Digital Risk Monitoring Q3 2016 report. The independent research firm, which evaluated nine different companies based on 27 criteria such as the ability to detect and mitigate corporate risk and gather data to monitor for risk, gave our External Threat Management Platform the highest scores for Current Offering and Market Presence categories.

With the broadest coverage of channels, RiskIQ is uniquely positioned to help companies understand their digital footprint and to detect and respond to threats across the web, mobile, and social threat landscapes.