DOMAINTOOLS

Predictive Domain Reputation Scoring

10 Billion+
Current and Historical
Whois Records

4.5 Billion+
IP Address
Change Events

1.8 Billion+
Registrar
Change Events

3 Billion+
Name Server
Change Events

580 Million+
Screenshots

Assessing Threat Before It Hurts

Everyone knows that appearances can be deceiving. From people, to books, to businesses, a first impression can turn out to be dangerously inaccurate Internet domains can be even harder to judge. They can look innocent, yet harbor dangerous malware. It can be very hard to sort out the good from the bad.

Today's threat environment demands a fast, reliable, automated way to assess and report the risk levels of domains so that appropriate security measures can be enacted. Reputation scoring is an effective family of technologies to accomplish this.

What is Reputation Scoring?

The concept of reputation, as applied to domains, IP addresses, and URLs, enables humans or automated security technologies to make decisions on whether to allow, deny or conditionally allow various types of connections. In forensic investigations, it helps guide the investigator toward those domains most likely to represent risk. Unlike simple allow/deny blacklisting, reputation scoring enables tailoring the security posture to the level of risk.

Why It Matters

Securing people and data: In an age where millions of dollars of data are sneaking out to unknown domains through sophisticated attacks, companies need tighter controls and automated systems for checking traffic to and from the internet.

Pre-connection checking: Many businesses conduct automated transactions from domain to domain, such as e-commerce systems that connect to other domains to enable financial transactions. Such a business has a vested interest in screening out dangerous domains--but it is not practical for human security personnel to vet each and every domain that seeks to connect.

The Problem With Reactive Reputation Scoring

Most reputation feeds are reactive rather than predictive. They assign risk based on observed behavior—for example, hosting malicious software (malware), or hosting command-and-control infrastructure for botnets. Because of this, traditional reputation feeds require one or more victims to be hurt before the domain lands on a reputation list. This creates a lag in time between when the malware is actually in operation to when it is caught and its presence distributed across blacklists. **This lag time represents a significant risk.**

A Predictive Reputation Feed

The DomainTools Reputation Engine (DTRE) assigns risk scores to domains as they are registered or even, in effect, before they are registered (by flagging malicious serial domain registrants). This means that users are protected from malicious domains before they have fired their first shot.

The DTRE is based on an algorithm called Proximity to Known Danger.

Proximity to Known Danger

Malicious or dangerous domains are usually controlled by organizations with many domains in their holdings—"lone wolves" are uncommon. Therefore, if "Domain A" is observed to be malicious, then other domains controlled by the same organization ("Domain B," Domain C," etc) automatically inherit an elevated risk profile. Registrant email addresses are analyzed for their connection to known-evil domains, and any other domain registered with that email address receives an elevated risk score.

The same principle applies to IP addresses and name servers. An IP address that hosts a high concentration of malicious domains—that is, the percentage of known-bad domains out of all domains on the IP address—is a risky "neighborhood," and any domain hosted there inherits a degree of elevated risk. Likewise, Domain Name System (DNS) servers that have high concentrations of known-dangerous domains served by them impart a higher risk score to other domains on them. As the concentration of bad domains approaches 100%, the risk scores for any domains associated with those IPs or name servers become very high.

A Service You Can Count On

DomainTools is uniquely positioned to deliver next-generation predictive reputation technology, having not only the broadest domain coverage, but also the most categories of data on which to build reputation scoring. The databases powering the DTRE are the world's largest repository of current and historical domain profile information.

The DomainTools Reputation Engine has been designed and built to be accurate, reliable, efficient, and scalable. The performance of the system reflects this rigor: accuracy rates are above 98% and continue to rise as the algorithm is fine tuned.

Want to Learn More?

The DTRE is a leading-edge project of DomainTools' research and development team. It is available for early-access evaluation on a limited basis. To stay informed about the DTRE release dates, or to discuss strategic technology partnerships, please contact us at **product@domaintools.com**.

> DomainTools' reputation scoring technologies slam shut the window of vulnerability left open by reactive reputation systems.