

Sorting Gray Alerts Using Domain Reputation Scores

Posted on January 30, 2020



The job of managed detection and response (MDR) teams, as their name suggests, is not limited to detecting cybersecurity threats. They are also responsible for carrying out the right actions in response to specific threat alerts.

If there were less than a hundred alerts, and they were all black or white, everything would go smoothly; at least when it comes to following up with the appropriate responses. Alerts with a definite malicious component would then be processed easily to quarantining and blocking stages, while benign alerts are ignored. But the cybersecurity landscape has become more complicated than that, for several reasons, including the facts that:

- Threats are no longer black or white. MDR teams are seeing more and more gray alerts or those that need to be investigated further to assess if they are malicious or not.
- We are not talking about a hundred alerts. Some 55% of IT professionals said they receive more than [10,000 alerts](#) per day, while 27% get more than 1 million of them daily.

With this complexity in mind, we further explore the challenge of grey alerts and how domain reputation scores gleaned from our [Domain Reputation API](#) can help make the job of MDR teams and IT security professionals more comfortable.

Alert Fatigue: When There Are Too Many Gray Alerts

Alert fatigue happens when cybersecurity professionals are frequently bombarded with large volumes of alerts until they become desensitized. Alarm fatigue can lead to longer response times, and can even result in overlooking critical alerts. In both instances, an organization becomes vulnerable to costly cyberattacks, such as ransomware infections, data breaches, and denial-of-service (DoS) attacks.

With the sheer volume of alerts that security teams receive daily, alert fatigue is quite common. Around **31.9% of IT security professionals**, for instance, ignore alerts since many turn out to be false positives.

A white paper by **Enterprise Management Associates (EMA)** detailed the following statistics about security alerts:

- Some 74% of IT security professionals are overwhelmed by the volume of security alerts and the work required to address them.
- Around 79% revealed that their alert assessment process remains manual.
- Security analysts spend 24-30 minutes to investigate each alert.
- Some 52% of security alerts are not prioritized correctly by the tools employed, so they have to be manually prioritized.
- Around 46% of threat alerts are incorrectly tagged as “critical,” while 31% turn out to be false positives.

The Target data breach that occurred in 2014 was tagged as “the greatest retail hack in U.S. history.” And it may be a classic case of alert fatigue wreaking havoc on an organization. Six months before the attack, Target employed malware detection software from an MDR service provider. As such, it did receive alerts labeled “malware.binary,” but these **were ignored**.

The malware infection eventually led to the theft of 40 million payment card information from its network, costing Target more than \$300 million in settlement costs, which experts said could even reach \$1 billion. The announcement that the malware alerts were indeed missed also caused the retailer’s stock price to dip at that time.

Gray Alert Prioritization Using Domain Reputation Scores

The Target data breach taught us that failure to assess gray alerts could be very costly. And since cybersecurity staff members receive thousands of alerts daily, the key lies in proper prioritization. They need to identify which alerts need to be addressed first, and one way to do that is by using **domain reputation scores**.

Domains remain a critical attack vector, with millions of malicious domains used to distribute malware, send spam, and enable threat actors to infiltrate networks overall. From 2017 to 2018, the number of phishing sites increased by 220%. It's therefore essential to prioritize threat alerts from domains or IP addresses trying to access the organization's network based on **domain reputation scores**.

Consider the scenario where a company receives thousands of gray alerts, including those from three domains: laticivue[.]com, yntscp[.]com, and twitter.com. The first two domains are associated with the Hancitor malware, according to [Pastebin](#), while the last may trigger an alert in organizations where access to social media sites is generally prohibited.

We ran the three domains on our [Domain Reputation API](#) and found the following:

- Laticivue[.]com has a risk score of 66.69.



laticivue.com



Search by IPv4, domain name

Warnings detected

Score: 66.99

WHOIS Domain check

- Registered 0 day ago
- Owner details are publicly available

SSL certificate validity

- Recently obtained certificate, valid from 2019-12-02 08:50:17

SSL vulnerabilities

- HPKP headers not set
- HTTP Strict Transport Security not set
- Heartbeat extension disabled

- Yntscp[.]com has a risk score of 78.76.



www.yntscp.com



Search by IPv4, domain name

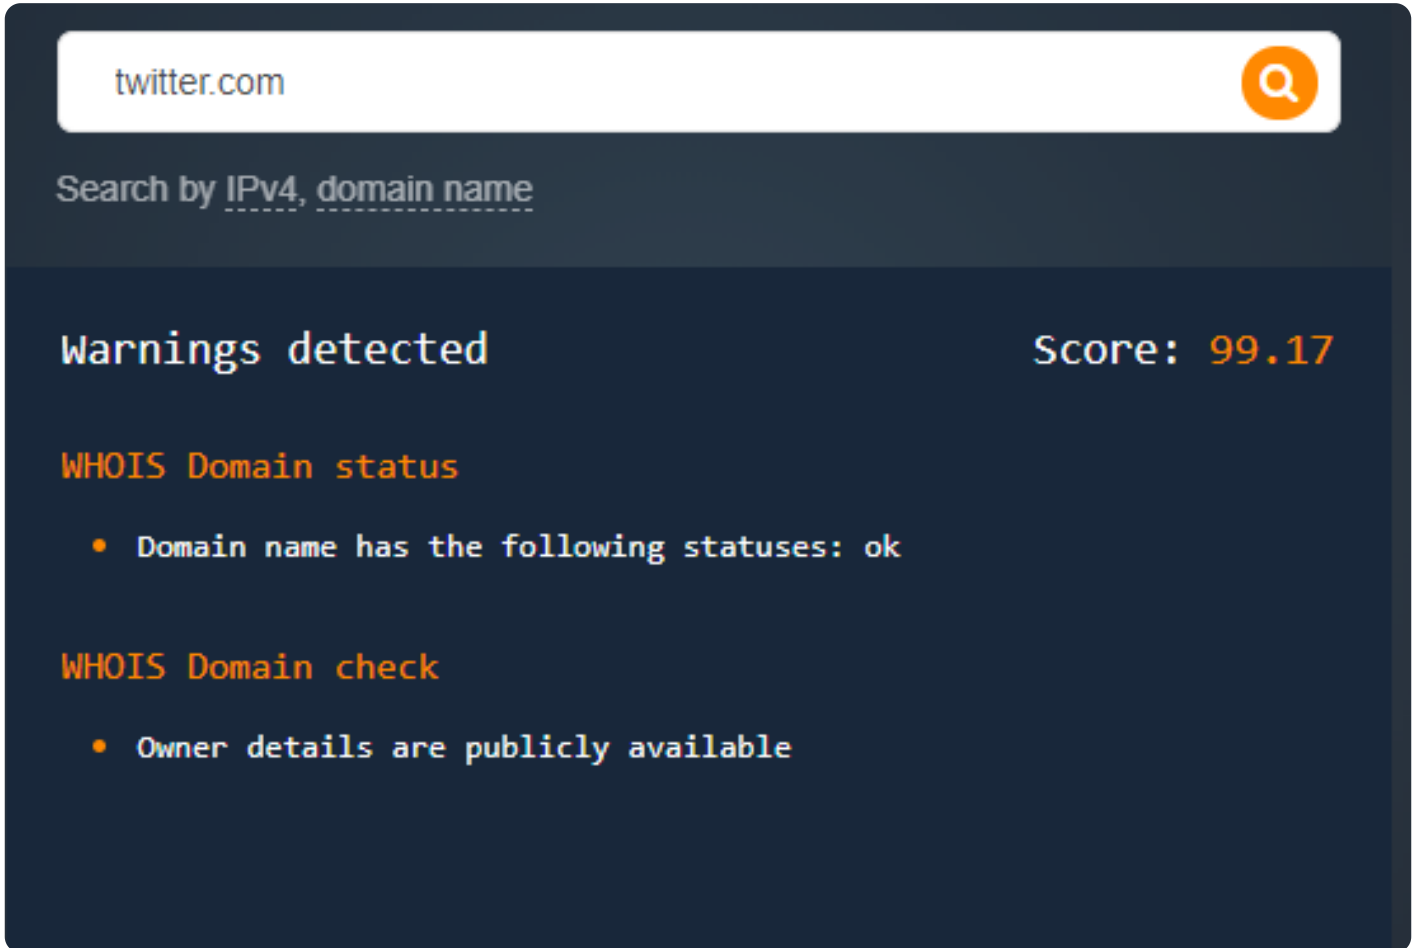
Warnings detected

Score: 78.76

WHOIS Domain check

- Registered 1 month and 24 days ago
- Owner details are publicly available

- Twitter.com has a risk score of 99.17.



The screenshot shows a search interface for the domain 'twitter.com'. The search bar contains 'twitter.com' and a magnifying glass icon. Below the search bar, it says 'Search by IPv4, domain name'. The results section is titled 'Warnings detected' and shows a 'Score: 99.17'. Underneath, there are two sections: 'WHOIS Domain status' with a bullet point 'Domain name has the following statuses: ok', and 'WHOIS Domain check' with a bullet point 'Owner details are publicly available'.

MDR teams can thus prioritize the alerts based on the **domain reputation scores** returned by the API. In the sample scenario, the first alert should be assessed first, followed by the second alert and then the third one. The higher the score (0 = high risk; 100 = safe), the safer the domain is to access.

The **domain reputation scores** given out by [Domain Reputation API](#) is a reliable basis for prioritizing security alerts. Its algorithm takes account of several factors, including the following:

- Website content and host configuration

- The domain's Secure Sockets Layer (SSL) certificates, connection, and configuration
- The domain's WHOIS record
- Inputs from different malware data feeds
- Reverse IP lookup to see all associated domains

In short, domain reputation scoring can help MDR teams and IT security professionals sift through thousands of gray alerts faster and gives organizations the ability to better ward off alert fatigue.