



Lost in the Noise:
4 Weak Attack Signals
your IDS will Miss

TABLE OF CONTENTS

INTRODUCTION AND OVERVIEW	3
THE SIGNALS MAY BE SUBTLE BUT THEIR IMPACT IS NOT	3
HOMOGLYPHS	4
Why your legacy IDS will not detect homoglyphs	4
How a modern NDR can detect homoglyphs	5
UNAUTHORIZED USER ACTIVITY	6
Why your legacy IDS will miss unauthorized activity	7
How a modern NDR can detect unauthorized activity	8
MALWARE C2 BEACONS	9
Why your legacy IDS will not detect malware C2 beacons	9
How a modern NDR can detect malware C2 beacons	10
ANOMALOUS NETWORK ACTIVITY	11
Why your legacy IDS will not detect anomalous activity	12
How a modern NDR can detect anomalous activity	12
SUMMARY	14

INTRODUCTION AND OVERVIEW

Legacy intrusion detection systems (IDS) can be effective at making signature-based detections and spotting explicit attack signals, but what happens when the attack signal is weak, low-amplitude, or subtle? These indicators can easily be missed by the IDS, because the IDS lacks the fundamental mechanisms needed to analyze the activities that generate these signals.

Any threat that sneaks past your legacy IDS leaves your organization open to risk, so employing more advanced detection methods to identify them is vital to the safety of your organization. Don't be fooled by the weakness or subtlety of these attack signals – their impact can be anything but weak or subtle.

In this paper we explore four types of network activity that your legacy IDS will likely miss which – if detected – can provide early warning of a cyber attack. For each, we describe the mechanisms used by modern IDS alternatives to detect them.

THE SIGNALS MAY BE SUBTLE BUT THEIR IMPACT IS NOT

The team at Stamus Labs has identified four types of weak attack signals that are commonly missed by IDS detection. Later in this paper, we describe the ways these attack signals can be detected effectively using a modern IDS alternative.

For each, we describe the signals, explain why an IDS will have difficulty detecting them, and highlight some of the ways these can be effectively detected on your network.



Homoglyphs



Unauthorized User Activity



Malware C2 Beacons



Anomalous Network Activity



HOMOGLYPHS

Homoglyphs (sometimes known as homographs) are a common method of deception used primarily in phishing attempts. In this type of attack, the attacker disguises their malicious domain, URL, or TLS certificate by using characters that appear identical to those that are used by the spoofed domain, URL, or TLS server name indication (SNI).

Recent Purchase

Dear Mark,

Your Apple ID, dmarkdurrett@gmail.com, was just used to purchase 50GB of iCloud storage on a computer or device that has not previously been used to make purchases. You may also be receiving this email if you reset your password since your last purchase.

If you made this purchase, or have reset your Apple ID password since your last purchase, you can disregard this email.

If you did not make this purchase, or believe someone may have accessed your account, go to <https://appleid.apple.com> and change your password as soon as possible.

Regards,
Apple

Apple ID Summary • Terms of Sale • Privacy Policy

Copyright © 2020 Apple Inc.
All rights reserved.

Did you mean apple.com?



The site you just tried to visit looks fake. Attackers sometimes mimic sites by making small, hard-to-see changes to the URL.

Ignore [Go to apple.com](https://apple.com)

<https://appleid.apple.com>

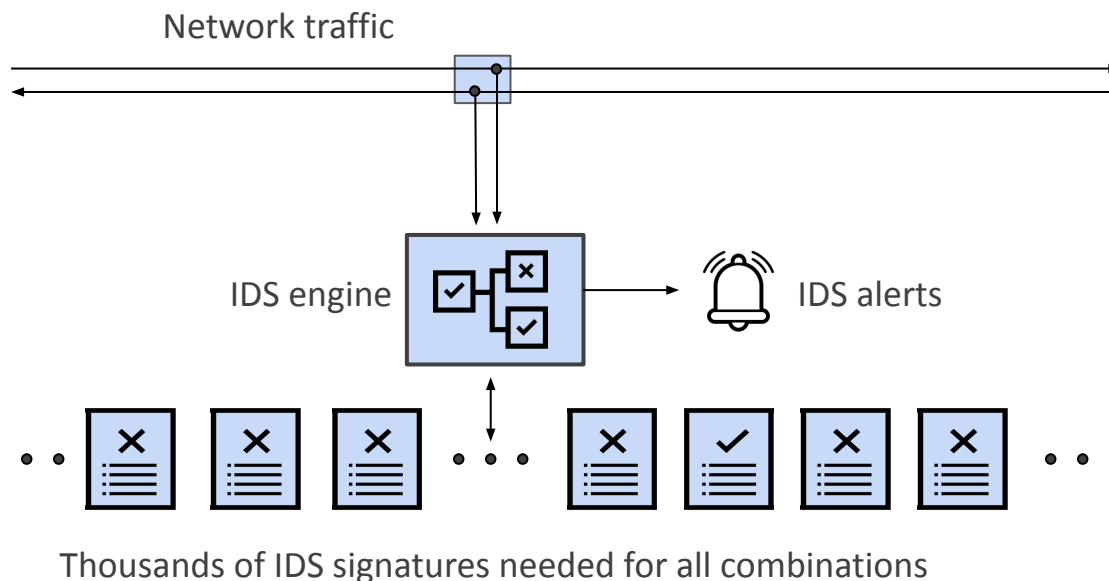
Because of the vast number of Unicode combinations and potential spoofs, the number of possible homoglyph combinations is essentially infinite. This makes detecting homoglyphs incredibly difficult without the right technology.

Homoglyphs – Can you spot the difference?



Why your legacy IDS will not detect homoglyphs

IDS functions by comparing a stream of packets to an explicit rule. To trigger an alert, an IDS must see a match between network traffic and the pre-defined indicator of compromise, known malicious IP address, untrusted domain name, or any other explicitly identifiable characteristic.



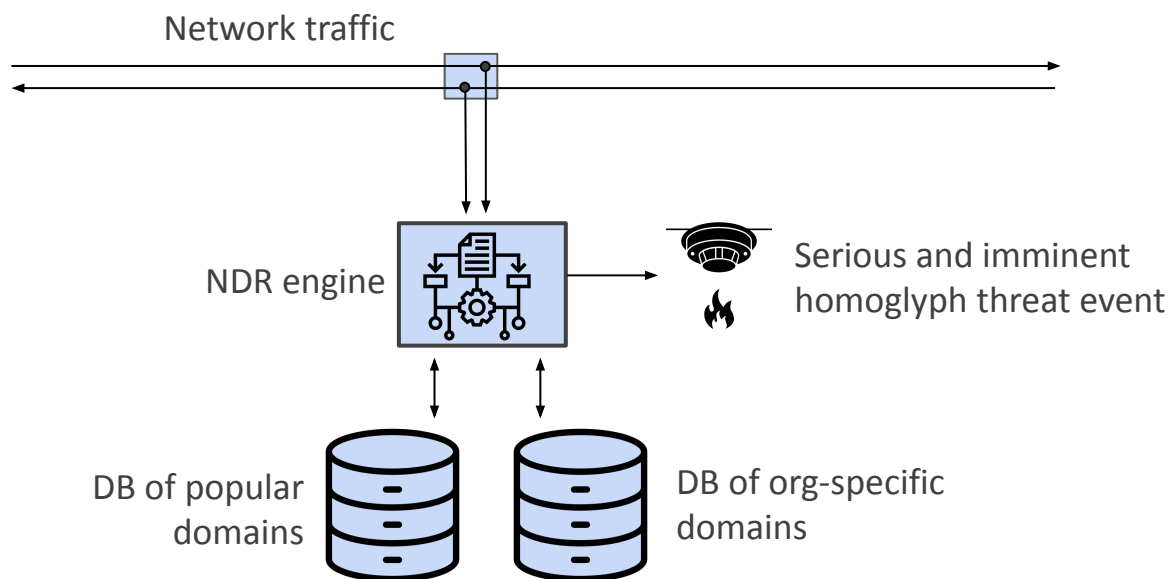
While, technically speaking, rules could be written to detect certain commonly known instances of domain spoofing homoglyphs, it is impractical to assume that a rule could be written for every possible instance of homoglyph usage. If it was even possible to write that nearly infinite number of rules, the IDS would still fail to be capable of storing that many rules while also effectively checking traffic against them.

How a modern NDR can detect homoglyphs

Homoglyph detection requires advanced functionality that the IDS simply does not possess – advanced logic on Unicode decoding. In this method, there must be a database of commonly spoofed domains (such as the Alexa top 100 domain list paired with a custom list of known domains specific to the monitored network).

When traffic moves through the network, it is checked against this list for similarity and an alert is triggered if the estimated similarity is below a given threshold. Essentially, known and trusted domains are defined, and then an engine is used to perform computational logic which compares the domains seen in incoming traffic against those known and trusted or regularly spoofed domains.

This analysis must be conducted by a post-processing engine. By dedicating computing power to the inspection of key pieces of metadata (like URLs, Domain Names, and SNI Certificates) the logic engine can analyze the Unicodes present in the serving domain and trigger alerts as needed. By using post-processing to do Unicode decoding and data analysis, the detection engine does not need to store countless rules the way an IDS would.



It is important to note that modern web browsers (like Firefox and chromium-based browsers) only show the non-Punycode version of the domain when all characters are the same language. Other browsers convert all Unicode URLs to Punycode or use optical character recognition (OCR) to determine if a URL could be interpreted differently. These are great first lines of defense, but links sent by text message, email, or other methods still pose a phishing risk.



UNAUTHORIZED USER ACTIVITY

There are several different types of network activity that can be classified as unauthorized user activity. Essentially, any type of activity that isn't explicitly approved by the organization's security team and IT department can be considered "unauthorized". This can vary across each organization. Within this umbrella also falls shadow IT (the use of unapproved software, systems, or devices) and policy violations (when a user breaks a defined rule or a tool is misconfigured).

Unauthorized user activity doesn't necessarily signal the presence of a malware actor on your network, nor does it mean that your users are purposefully trying to violate your policies. Regardless, maintaining oversight into these instances is still an incredibly important part of defending the organization. While this kind of activity does not always indicate that you are under attack, unauthorized user activity can leave your organization vulnerable.

Why your legacy IDS will miss unauthorized activity

In order to trigger an alert, the IDS must match between specific pieces of data in network traffic and predefined indicators of compromise, known malicious IP addresses, untrusted domain names, or other explicitly identifiable characteristics.

Unfortunately, this type of detection does not help uncover unauthorized user activity which must be detected by monitoring host activities and actively hunting for known violations.




Organizations should establish a baseline for what is authorized and what is not; however, the responsibility for monitoring user activity and auditing these policies often falls on the security team.

An IDS alone cannot maintain the host state needed to view the relevant data which comprises a user's history and activity.

Examples of Unauthorized user activity



Unauthorized proxy servers




Off limits file sharing apps



Forbidden cloud service



Clear text passwords



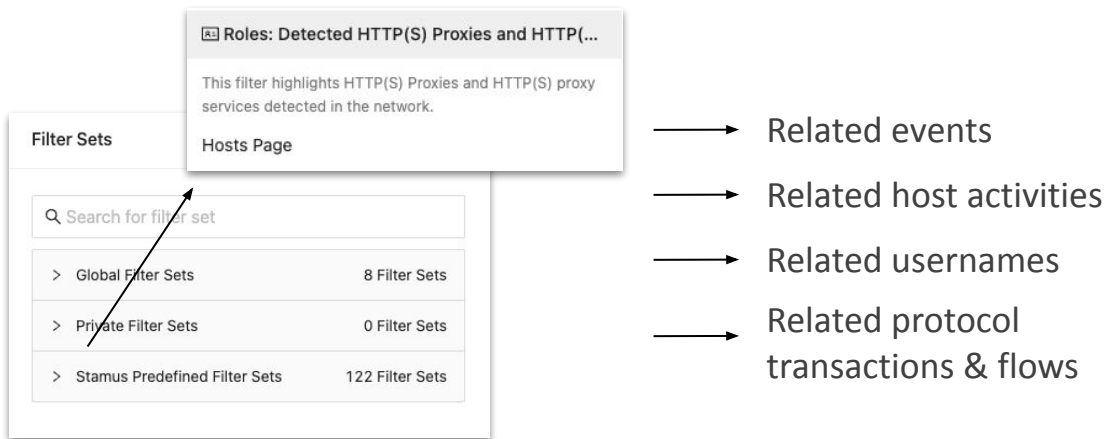
ChatGPT usage



Other shadow IT

How a modern NDR can detect unauthorized activity





Thankfully, integrated threat hunting tools provided by more modern systems like network detection and response (NDR) can mitigate the amount of work an analyst would need to do to identify unauthorized user activity by providing a panel of insights from the host that can quickly and easily be filtered to look for violations in policy.



Example of threat hunting for rogue proxy servers

While legacy IDS systems generate most of the data needed to do this (related logs and NSM data help complete the picture), there is not typically an automated process to trigger alerts based on user activity. Experts recommend a proactive approach to finding unauthorized activity using threat hunting tools which can query all the relevant host data for a specific time window.

Other Hunting filter examples


-  Clear Text Passwords
-  Possible Tor Traffic
-  FTP Network Services
-  FTP Network Services



MALWARE C2 BEACONS

Malware beaconing is when malware communicates with an attacker's command-and-control (C2) server to receive new instructions or tasks to complete on a target machine. Attackers configure the frequency and method of these communications with the goal of hiding them in seemingly normal network traffic.

Periodic signals between infected systems and malware command and control (C2)



Basic malware beacons will transmit data at regular intervals, which is not overly difficult for most systems, but sophisticated evasion techniques like low frequency, randomized communications or varied communication channels can cause beacons to be missed.

Beacons themselves are not actually harmful to a system, but the instructions they contain that are passed on to malware present in the target machine can lead to data breaches, stolen information, or ransomware attacks.

Why your legacy IDS will not detect malware C2 beacons

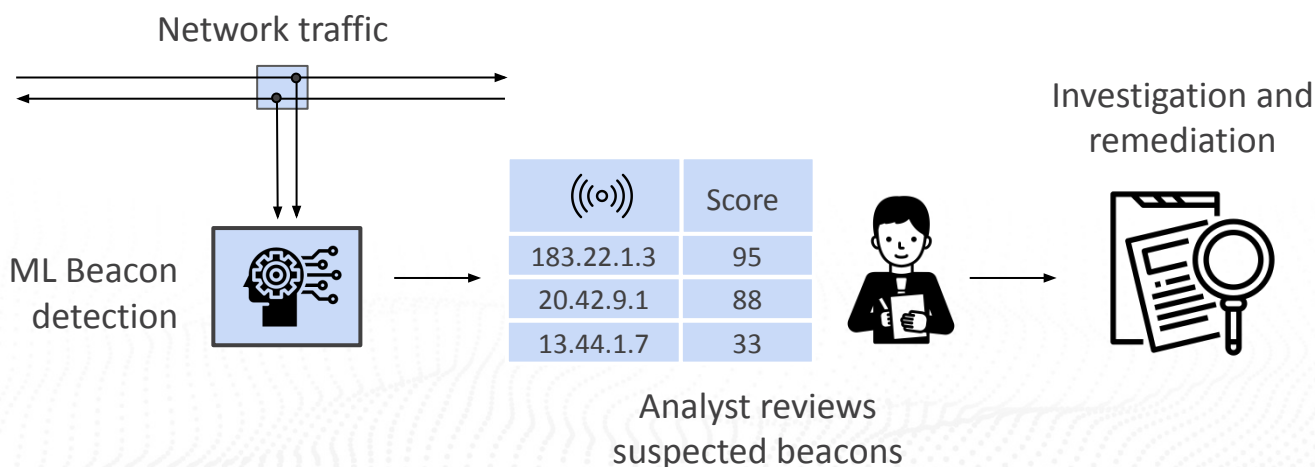
IDS can detect specific types of malware or command-and-control servers, assuming that those sources are already known, and the corresponding rules have already been written. But when the C2 server has never been seen before or the malware has already found access into the target system, IDS has no way of detecting its presence. When this happens, the best way to locate the threat and block the servers access is to identify their communications and then trace the source and destination.

The main reason IDS cannot detect malware beaoning communications is because they happen over time. IDS signatures happen on a single packet in a single moment in time. Detecting beacons requires aggregate data which must then be analyzed to look for regular frequencies or suspicious behaviors. IDS just doesn't have the ability to track these changes and conduct the analysis needed to identify these low volume attack signals.

How a modern NDR can detect malware C2 beacons

The amount of data being transmitted in every beacon request and response is often consistent, and the intervals at which the malware calls home is regular regardless of the frequency. Beacons follow a pattern, and no matter how randomized that pattern might be it can be identified using the right technology.


By continually analyzing various pieces of flow data (packet size, jitters, standard deviation, repetition, etc) a machine learning algorithm will be able to identify patterns that signal possible beacons and then aggregate that information towards a specific IP address or JA4S fingerprint for further analysis.



A ML based detection system generates a confidence score (a beacon metric) that helps the security team quickly assess the likelihood that a communication is a malware beacon based on several behavioral factors. The beacon metric is a weighted score prioritizing TLS servers exhibiting behavior patterns typically associated with beaoning traffic. In other words, communications with clear periodicity (even if the periodicity fluctuates) and specific packet profiles are highlighted by this mechanism.

The beacon metric can range from 0 to 100. The higher the beacon score, the higher the possibility that communication is a potential beacon. By investigating the assets impacted on these beacon profiles, a security practitioner can quickly identify command and control server activity based on IP address communications or JA4S value, a fingerprint of the server side of a TLS handshake.

When an analyst confirms the presence of a beaconing system, they can quickly investigate, gather evidence, and remediate.



ANOMALOUS NETWORK ACTIVITY

Anomalous network activity is any change in the established standard communication happening on a network. An anomaly could signal malware or another type of cyberattack. Further investigation could uncover network problems or equipment failure. Regardless, anomaly detection is important because it helps identify early attack signals that could be missed elsewhere while also giving greater visibility into the health and efficiency of your network.

Anomalous behavior has likely never been seen on the network. There are any number of behaviors that could be considered anomalous, and it just depends on the baseline that has been set for what is considered “normal”.


Essentially, when something or someone on the network is not behaving in the way you expect it to behave or when something or someone is present on the network that you have never seen before, that should be considered an anomaly.



New network clients and devices



Unusual network connections



Changed command structure



Unusual data packets



Previously unseen communications



Unusual user behavior patterns

Why your legacy IDS will not detect anomalous activity

There are five main reasons why your legacy IDS cannot detect anomalous behavior.

- Anomalous activity must be observed and tracked over time
- Signature-based IDS has no concept of time or state
- Must create a baseline and track deviations against that baseline
- Anomaly detection requires details about each host, unavailable to an IDS
- Activity monitoring requires metadata not observed and maintained by IDS, such as protocol transactions, flow data, files, etc

Traditional IDS uses signature-based detection. Traffic on the network must be compared against a library of explicit, predefined rules. When a traffic pattern matches a rule, an IDS alert is triggered. This type of detection does not work for anomalous behavior because it cannot maintain the host state and view all the relevant pieces of metadata over time.

Maintaining state requires keeping track of the combination of original data plus any changes seen in that data over time. Your legacy IDS simply does not maintain the state of the hosts and their related metadata, preventing it from seeing the changes which could signal anomalous behavior.

How a modern NDR can detect anomalous activity

To detect anomalies, the detection system must have some way of maintaining the host state and then provide a way for the analyst to see a full panel of the host's activities over time. Your legacy IDS does generate a lot of this data, and in addition, it must be paired with the related logs and NSM data in order to get the full picture. Change cannot be tracked unless the detection engine provides a way to see how the host has behaved over a period of time.

There are three primary mechanisms that may be used to detect anomalous network activity using information gathered from the host state. The first is machine learning, which is becoming a popular method of anomaly detection. With machine learning, the system analyzes host data to learn what is “normal” activity. Deviations from “normal” are considered anomalous.


The second detection method is the use of statistical algorithms. These detection engines can locate previously unseen or otherwise unfamiliar network activity which could help signal an anomalous behavior.

Finally, proactive threat hunting is a common way analysts locate anomalous network activity. Using analysis tools, guided filters, or custom filters, a threat hunter can search through host data to find specific types of anomalies (such as users from non-IT departments performing advanced administrative processes).



Manual Hunting - skilled analysts using guided threat hunting tools to spot anomalies

Can be used to identify suspicious activity for further investigation



Statistical Algorithms - outliers and previously unseen activity based on statistical probabilities

Can be used to identify suspicious activity and can be used to build evidence



Unsupervised ML - machine learning detection outliers

Can be used to identify suspicious activity and can be used to build evidence

Each of these mechanisms have their strengths and weaknesses. A mature security team uses all three to help their organizations uncover anomalous network activity.

SUMMARY

In order to expand visibility into increasingly subtle attack signals that will routinely be missed by a legacy IDS, organizations should consider modern alternative options.

However, it is ultimately up to the organization to decide whether the risk of these attack signals is worth making the switch from legacy IDS to a more modern solution. While a legacy IDS can be a reasonably capable network threat security system, IDS falls short in many detection scenarios.

Forward looking organizations will look for a modern network detection and response (NDR) system that preserves the evidentiary value of an IDS while dramatically improving threat detection and response capabilities.

Clear NDR™ is that solution. Built on Suricata — a highly effective open-source legacy IDS — but empowered by multiple modern detection methods like machine learning, stateful logic, anomaly detection, and more, Clear NDR gives users the best features of IDS without the limitations.

To learn more about how an Clear NDR can provide important benefits for organizations, visit www.Stamus-Networks.com.

ABOUT STAMUS NETWORKS

Stamus Networks believes that cyber defense is bigger than any single person, platform, company, or technology. That's why we leverage the power of community to deliver the next generation of open and transparent network defense. Trusted by security teams at the world's most targeted organizations, our flagship offering – Clear NDR™ – empowers cyber defenders to uncover and stop serious threats and unauthorized network activity before they harm their organizations. Clear NDR helps defenders see more clearly and act more confidently through detection they can trust with results they can explain.



5 Avenue Ingres
75016 Paris
France

450 E 96th St. Suite 500
Indianapolis, IN 46240
United States

✉ contact@stamus-networks.com

🌐 www.stamus-networks.com