

Blacklists & Dynamic Reputation

Understanding Why the Evolving Threat Eludes Blacklists

by Gunter Ollmann, VP Research, Damballa

As cybercriminals have increased the sophistication and diversity of the threat they pose to enterprise businesses, reputation systems have come under growing pressure to deal with the agile nature of the threat.

Most approaches to applying protection against Internet threats depend on lists of what are essentially "good" or "bad" for those in need of defense. These lists are effectively the minimum set of information that a particular protection technology needs to be supplied with in order to achieve its stated security objectives. For example, Anti-Spam technologies rely on lists of "bad" IP addresses - where each IP address represents a previously identified spamming server. Meanwhile URL content-filtering services offer a categorized view of Internet content that enables organizations to allow or block access to certain Web pages.

The lists used by today's protection technologies and filtering strategies are effectively the aggregated observations and analysis of past threats. In essence, these lists are employed as "reputation" sources from which automated decisions are made to allow or block traffic.

The reputation systems used by the security technologies charged with protecting enterprise networks on a daily bases can be divided into two key classes:

- **Static Reputation** - The traditional approach to employing reputation for defense is through the use of static list-based reputation systems. They contain information of known or previously encountered and classified threats and are typically distributed in the form of blacklists or whitelists. Static reputation systems provide a binary perspective of the threat - the threat is either on the list or it isn't.
- **Dynamic Reputation** - Dynamic reputation systems were developed to counter more agile and faster moving threats than were capable of being addressed by their static reputation system counterparts. They are typically formula-based and rely on a mix of historical information and "real-time" intelligence - making use of sliding windows of observations and aggregate scoring systems. The dynamic nature of the system means that non-binary reputation "scores" are supplied on a query-answer basis. Instead of lists, dynamic reputation systems often employ queryable API's to ensure the most accurate score.

Dynamic reputation systems can be sub-divided into two additional types:

- **Cascading** - Cascading dynamic reputation systems rely on the input of multiple related data sources - each feature of which is weighted and scored. The more information related to the threat being queried, the more distinct the score.
- **Predictive** - Predictive dynamic reputation systems employ similar data sources and feature weighting systems to that of cascading dynamic reputation systems, but they also utilize models and trained classifiers in order to project reputation scores on threats for which there is little confirmed information available. As such, they are able to predict the reputation score of brand new threats before other confirming information can be cascaded down.

Reputation Over Time

As cybercriminals have increased the sophistication and diversity of the threat they pose to enterprise businesses, reputation systems have come under growing pressure to deal with the agile nature of the threat. Dynamic reputation systems have come to the forefront in dealing with the dynamic threat landscape and keeping pace with the changing face of cybercrime.

In theory, static reputation systems such as blacklists and whitelists should be more accurate (when they are being constructed) than their dynamic reputation counterparts. Static lists by definition are validated by hand and are generally much more limited in scope (e.g. containing only a few hundred or thousand items). In reality, the actionable output from static reputation systems is typically only more accurate at the time of the list's construction

and, over time, the list's accuracy decreases markedly. Dynamic reputation systems provide a more accurate depiction of the threat over time, since they are constantly processing new intelligence and providing real-time scores.

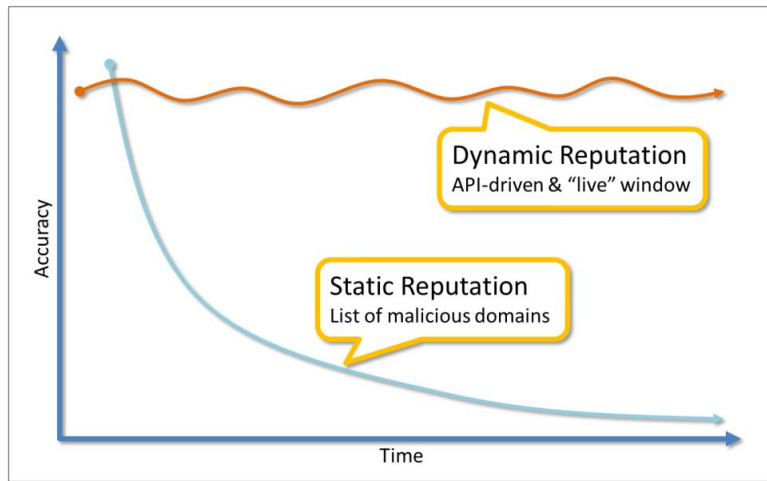


Figure 1: Threat reputation accuracy over time. Fluctuations in dynamic reputation accuracy reflect the changing time-window of observations and new data inputs. Decreases in static reputation accuracy are attributed to aging of a static list.

The actionable intelligence list produced by a static reputation system depreciates in accuracy over time after the list has been generated (assuming that it was accurate to begin with). As the threat changes, the disparity between the static reputation list and the cybercriminals hosting infrastructure grows. Dynamic reputation systems fluctuate a few percent over time due to the use of sliding-window threat observations and the addition of new data.

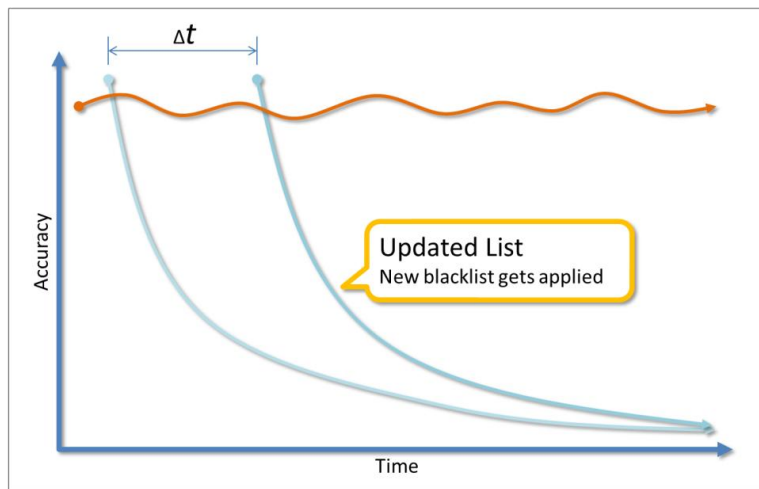


Figure 2: The depreciating accuracy of static reputation derived lists can be partially compensated through the periodic release of new "fresh" lists.

To partially compensate for the degradation of static reputation lists over time, it is important to continuously update the particular blacklist or whitelist and release periodic updates. Depending on the dynamism of the threat to which the reputation system is being applied and the frequency of list updates, a saw-tooth pattern to accuracy can often be observed.

Know What You Know

Reputation systems primarily deal with the known; that is to say they tend to focus on the classification of threats that have been observed sometime in the past. Static reputation systems are locked to qualifying only those threats that have been previously observed. Meanwhile, dynamic reputation systems offer more flexibility and some implementations can extend to predictive assignments of reputational score.

An important aspect of reputation systems when discussing accuracy is that of true positives (TP) and false positives (FP). The binary nature of static reputation systems (i.e. it's either in the list and real, or it's not in the list and is therefore unknown) means that, in theory, it should only contain TP listed threats. However, as previously discussed, over time, the accuracy of the list depreciates and FP results are introduced (especially in the cases of IP-based blacklists).

Dynamic reputation systems offer more flexibility in identifying a threat. By definition, they cover an entire spectrum of threats that are visible during the training phase and assign a probability (or confidence) score for a particular domain, IP address, URL, etc. This means that they can also cover both known and unknown threats - where "unknown" represents threats that have not been directly observed by the reputation system and/or are derived from cascaded threat data.

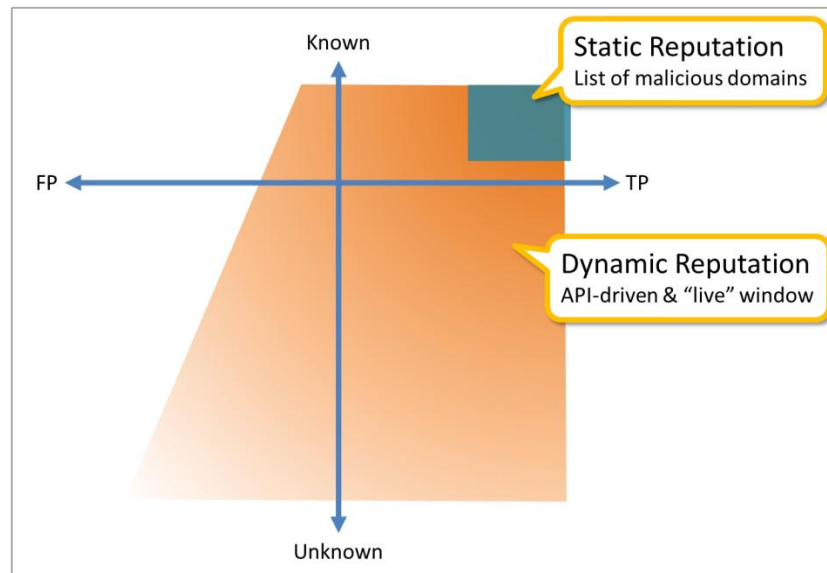


Figure 3: Static and dynamic reputation systems compared against FP vs. TP threat determination and the ability to identify known and unknown threats.

Static reputation systems are prone to aging effects. As the reputation data contained within the blacklist ages, it gets stale and includes a growing number of FP results - but is still limited to the handling of known threats.

Because of the non-Boolean nature of dynamic reputation systems, a FP is actually a misnomer. Each domain, IP address, URL, etc. can be assigned a score (representing a probability or confidence) of being affiliated with a particular threat. Ranging from 0 to 1 (or 0 to 100% for some systems), a dynamic reputation score is employed by protection technologies based on the tolerance for FP indicators. For example, a score of 0.95 and above may be the functional equivalent of a static reputation "True" state (i.e. a TP), while a score less than 0.25 may be interpreted as a "False" (i.e. it is unlikely related to the threat category). In practice, protection technologies may set a threshold of 0.75 for alerting and manual investigation and a threshold of 0.90 for automated response.

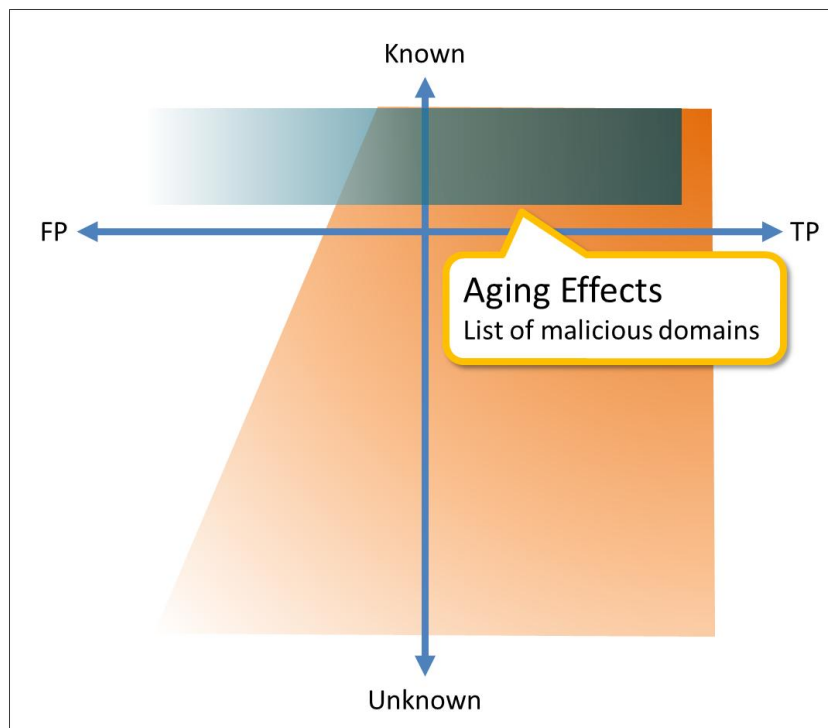


Figure 4: The aging effect on static reputation lists leads to FP threat identification.

Using Static Reputation Services

The majority of reputation systems employed today within the security industry utilize static reputation systems - or, more specifically, blacklists. Many of the legacy protection technologies dependent on blacklists continue to provide the bulk, first-pass, filtering of network traffic associated with the most common threat categories (e.g. Spam, phishing, drive-by-download, botnet command and control, malicious download sites, copyright infringing material, etc.).

Blacklists are easily compiled and there is no shortage of vendors offering lists for various threat categories. In some cases the data used to populate the static reputation system is gathered from their own threat observation capability, or mixed with commercially available threat feeds. The better the vendor is in acquiring the raw data and the more capable they are in analyzing and classifying that data themselves, the more accurate their blacklist will be (subject to data aging and the pace at which updates to the blacklist are distributed to subscribers).

If an organization relies on protection technologies that are reliant on static reputation lists (blacklist or whitelist), they should consider the following:

- **Who is the original source of the reputation data?**
Not all blacklists are equal. The breadth at which the vendor monitors the global threat and their ability to correctly label the data is very important. Vendors need be transparent in their data collection processes and articulate the level of global coverage in order for comparisons between vendors to be made. The majority of vendors supplying "proprietary" reputation data are often aggregating other third-party blacklist information - making it difficult to ascertain the accuracy of the data and a natural propensity towards FP alerts.
- **Will aggregating multiple blacklists increase protection?**
Combining multiple blacklists from various sources can cause more problems than it conceptually solves. Many blacklist vendors rely on public and widely available sources of information for compiling their lists. Purchasing multiple lists from vendors that use the same sources of data will not yield any incremental threat coverage. For vendors that produce blacklists exclusively from their own vetted intelligence sources,

there is often little overlap between blacklists. However concerns arise over what other threats are being missed and what was the vendor's logic for including or rejecting a threat in the published blacklist.

- **Why is a particular threat listed?**

False positives are a critical problem for static reputation systems since the protection technology relying on them is generally inflexible in the application of that threat intelligence. When a suspected FP is encountered, it is vital that the enterprise customer quickly uncover the details behind its inclusion in the blacklist (aggregating blacklists from multiple suppliers can exacerbate the problem). The blacklist vendor needs to be transparent in the logic behind the inclusion and removal of threats within the lists they provide and provide the tools and facilities to answer questions as to a particular threat's inclusion.

- **Is bigger better?**

Several list vendors extol the virtues of "bigger is better" when it comes to static reputation services. In practice, this tends to not be the case. In any static reputation service system there will inherently be FP inclusions. Smaller lists tend to be highly vetted and more accurate, meanwhile long lists tend to be less qualified and have a higher percentage of FP inclusions. Accuracy is more important than size. It is impossible to preemptively label as good or evil, every day, the tens of millions of unique domain names and IP addresses that will be contacted by corporate devices.

- **Who has access to the blacklist?**

The static nature of blacklists is an obvious weakness when dealing with dynamic threats and sophisticated cybercriminal opponents. Cybercriminals often monitor the public blacklists and purchase many of the cheap threat feeds or blacklists in order to identify whether their criminal infrastructure has been listed. If it has been listed, they will simply change DNS settings, IP addresses or hosting facilities in order to evade the protection technologies deployed by their targets. The cybercriminal doesn't need to be fast, just faster than the updates to the blacklists.

- **How frequent are the blacklist updates?**

When comparing blacklists from various vendors, it is important to understand the frequency of updates to the blacklist and the ramifications on the protection technology that will be utilizing them. Big blacklists can take a long time to download and distribute, therefore managing updates is more important. Scheduled releases of blacklist updates can be managed more easily than ad hoc updates - but ad hoc updates may enable more timely mitigation against new and particularly virulent threats.

Pressure on Static Reputation Systems

Despite the agility of the cybercriminals, there are other legitimate changes to the Internet and commercial server hosting capabilities that have a dramatic effect on static reputation systems - and are key factors for the growth of dynamic reputation systems. Some of the most known areas of weakness in static reputation systems are due to the following:

1. **Running out of IPv4 ranges**

Most of the IPv4 netblocks have been allocated. All /8 ranges were allocated by February 2011 and smaller ranges are similarly being exhausted. Because of this, there is an open market for unregulated "trading" of IPv4 IP addresses. A netblock of IP addresses that may have belonged to one organization today may belong to another tomorrow.

2. **The transition of IPv4 to IPv6**

As the limitations of IPv4 become more restrictive, businesses are switching to IPv6. IPv6 includes a tremendously larger number of IP addresses and routing features that make it impractical for list-based or scanner-based technologies to function.

3. Commercial cloud computing

As businesses move to the cloud, so has cybercrime. Cloud-based computing offers numerous advantages to cybercriminals - in particular the ability to rapidly initiate a malicious server and to present a dynamically allocated IP address to their victims that, if blocked or filtered, would likely impact multiple "good" commercial services hosted with the same cloud provider.

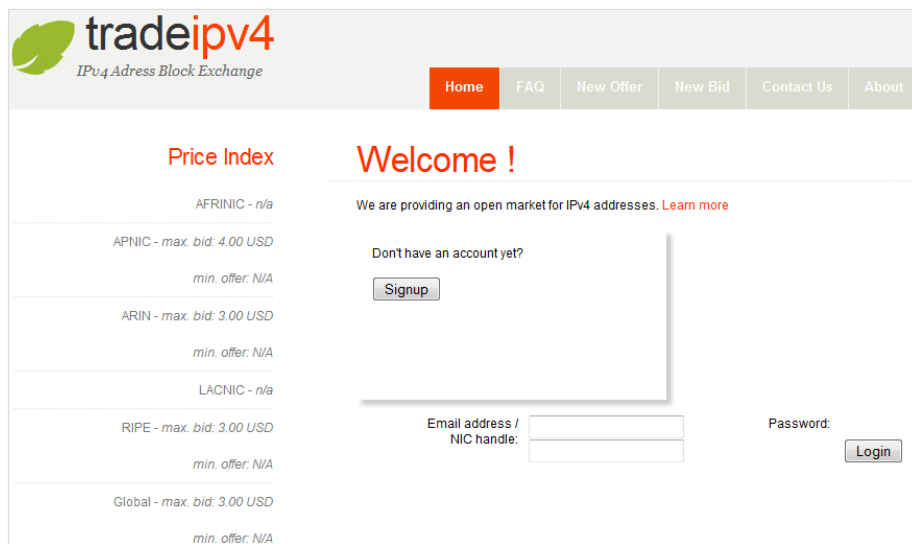


Figure 5: Trading IPv4 address blocks and IP addresses.

Using Dynamic Reputation Services

In response to the increased dynamism of the threat and the Internet at large, dynamic reputation systems are replacing their less agile static reputation counterparts. Unfortunately, it isn't a one-for-one replacement or integration process. Legacy protection systems require a definitive list of good and bad information from which to alert and take action. If a threat is on the list (e.g. blacklist) then it should do this, and if it isn't on the list then it should do that.

Dynamic reputation systems have come to the forefront over the last few years as new approaches to handling big data and the analysis of high-speed stream data sets have been invented. They have been designed to provide "live" scores associated with any suspected threat. Instead of the Boolean response to "is it on the list or not," a feature of the suspected threat (e.g. IP address, domain name, URL, etc.) is passed to the dynamic reputation system and a score is calculated at the time of the query - ensuring that a real-time evaluation of the threat is performed using the most timely and accurate information available. This score must be interpreted by the protection technology - allowing for multiple decisions to be made depending on the specific value returned.

If required, some dynamic reputation systems can be forced to output a blacklist. This is done by agreeing on a minimum scoring threshold (e.g. 0.975 for ISP's that are focusing on the most significant abuses or 0.75 for organizations being targeted by state sponsored attackers) and including all domains, IP addresses, etc. above the threshold as members of the blacklist. While this is sometimes possible to do given the query-answer nature of the system, a list of candidate threats (with a small set of data points per threat) will often need to be supplied for scoring. The process of selecting the candidate threats will obviously be a limiting factor in the accuracy and length of the static blacklist that happens to be generated.

The key features of dynamic reputation systems can be summarized as follows:

- Purposefully designed to handle a dynamic Internet and agile threats
- API query-answer driven framework to ensure the most timely and accurate intelligence
- Flexible scoring that encompasses both accuracy and confidence

- Automated analysis and dissection of big data record sets
- Threshold-based response to threat alerts and corresponding actions
- Inherent ability to predict and score newly formulated threats

Conclusions

Reputation systems have long held a critical role in corporate protection strategy. They have traditionally performed the role of the bastion defense in the form of blacklists - rapidly blocking or filtering the worst of the worst.

As the threats have become more dynamic and more diverse, blacklists (and the static reputation system technologies they are derived from) have been struggling to keep pace. Today's threats are agile and employ a growing number of techniques specifically designed to evade or subvert protection technologies that depend on static reputation systems.

Dynamic reputation systems, by their very nature, have been designed to handle agile threats and the morphing topology of the Internet. Techniques employed by attackers to subvert static reputation systems hold no sway against the latest generation of dynamic reputation systems. While gaming of dynamic reputation systems is possible, the process of doing so is not trivial and is generally cost prohibitive.

To effectively utilize dynamic reputation systems and replace redundant static reputation systems, changes are necessary within the protection technologies themselves. Instead of periodic updates to lists of reputation data, they must utilize query-answer systems (often in the form of queryable API's) to obtain the latest threat reputation information.

Agile threats require agile responses. Static reputation systems and the protection technologies that depend on them have failed against today's dynamically changing threat landscape. When presented as a first-pass filtering technology they can perform adequately against well-known and well-studied threats. They will however continue to be marginalized as a core component of modern protection technology when deployed against modern agile threats.

About Damballa, Inc.

Damballa is a pioneer in the fight against cybercrime. Damballa provides the only network security solution that detects the remote control communication that criminals use to breach networks to steal personal and intellectual information, and conduct espionage or other fraudulent transactions. Patent-pending solutions from Damballa are platform and system-agnostic, protecting networks with any type of device including PCs, Macs, smart phones, as well as mobile and embedded systems. Damballa customers include Fortune 1000 companies, Internet and telecommunications service providers, government agencies and educational organizations.
<http://www.damballa.com>

Copyright © 2011, Damballa, Inc. All rights reserved worldwide