White Paper

Reputation: The Foundation Of Effective Threat Protection

By Jamie Barnett

Table of Contents

Executive Summary

Reputation systems have been used for years across many disciplines—from doctors diagnosing illnesses to mathematical experts rating financial instruments—to assess situations and make decisions. Since the early days of online communities and ecommerce, providers and consumers of goods, services, and information via the web have sought ways to gauge the reputation of the parties involved in transactions. Reputation calculation tools are more critical today to cybersecurity than ever before, as more users access more online tools via more devices and interact with colleagues, friends, and strangers in more online venues. Reputation provides a comforting level of assurance around identity and integrity in critical Internet-based personal and professional transactions, for which physical-world verification is impossible.

In this paper, we arm security decision makers with information about what makes electronic security reputation systems effective; they, in turn, can apply that knowledge to both near-term security policy and long-term strategy. This paper:

- Addresses the dynamic nature of electronic threats and the need for an entity's reputation to reflect its current state at any given moment
- Contrasts the notion that reputation exists in the gray area between "absolute bad" and "absolute good" with more static blacklist or whitelist protection methodologies
- Discusses the four elements that ensure high confidence in calculating reputation: data volume, data longevity, data trustworthiness and, most important, correlation across a broad set of data

> *Reputation services identify the source of the threat and determine if it is legitimate. In other words, does your company do business with firms located at particular IP addresses, websites, or email names? If yes, are they trusted sources of transactions and interactions? Large-scale threat intelligence systems analyze data from tens of millions of endpoints and millions of servers.*
>
> —Chris Christiansen, IDC

> Since the early days of online communities and ecommerce, providers and consumers of goods, services, and information via the web have sought ways to gauge the reputation of the parties involved in transactions. This desire has spawned a variety of third-party "trust" models, from lightweight community voting to heavyweight certification authorities and seal-of-approval programs. Those models, in one way or another, have relied on reputation. Today, perhaps the most urgent need for reputation-based systems is in the world of cybersecurity to identify and prevent online threats such as network intrusions and malware.

Background

*On April 30 at 9:56am, www.multimedia***.com, a newly registered website that allowed users to post, search, and view amateur videos, came online. The website was part of a group of 160 new domain registrations, and was identified by the network of sensors and data feeds that contribute to McAfee Global Threat Intelligence. Seemingly legitimate, many of these domains had all of the trappings of media-sharing sites, except for a clue that prompted us to adjust their reputations to "high risk" in our system. What tipped us off?*

Wikipedia defines reputation as "the opinion (more technically, a social evaluation) of the group of entities toward a person, a group of people, or an organization on a certain criterion."[1] At McAfee, we have been dealing with the reputation of electronic entities—from files to senders to websites—for years, and our definition has broadened to include further elements.

First, reputations are dynamic and temporal. For example, a previously legitimate website can become infected with malware and then be cleaned up in a short time. A reputation must be refreshed as quickly as content is refreshed. Second, an entity's reputation is seldom "absolutely good" or "absolutely bad," but rather lies somewhere in the vast gray area in between, making the intersection of reputation with policy an empowering thing to security decision makers. Finally, confidence is a critical consideration in calculating reputation. By confidence we mean the confidence interval, or reliability of our estimate. The more data points and evaluation criteria we consider in our analysis, the more accurate the reputation we calculate is likely to be at that moment. The four things that contribute to increasing reliability are data volume, data longevity, data trustworthiness, and broad data correlation.

[1] "Reputation," Wikipedia. http://en.wikipedia.org/wiki/Reputation

At McAfee, we calculate the reputations of hundreds of millions of electronic entities—files, websites, web domains, messages, DNS servers, and network connections—using a highly granular scoring system based on a variety of information about the entity's behaviors, characteristics, and our own experience of how comparable entities behave. Among other inputs, we rely on telemetry data, billions of queries per day from tens of millions of McAfee products—ranging from anti-malware clients to web and email gateways to firewalls—that we have deployed around the globe and that act as sensors for our cloud-based analysis engine. For example, in dynamically calculating the reputation score of a network connection, we look at thousands of attributes and behaviors, including the IP address' life span, ports and protocols it uses, network activity vis-à-vis a baseline of expected behavior, attack history, and associations with other known IPs.

Catching Risky Behavior

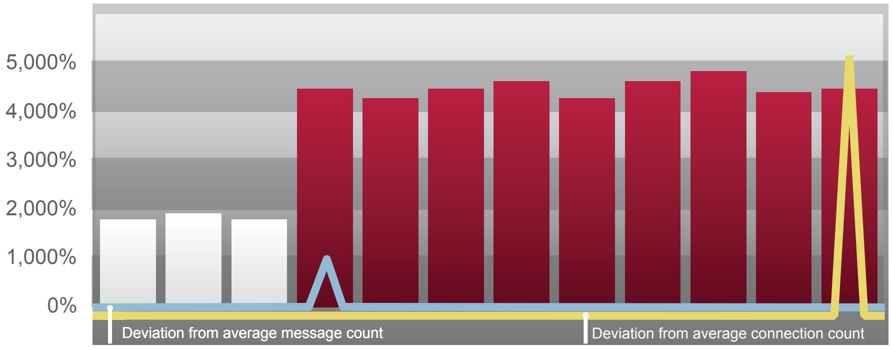

Figure 1. McAfee's reputation system observes potentially risky behavior in an IP address, for example, and would proactively block its messages from reaching our customers.

Figure 1, above, illustrates an instance in which our systems detect anomalous behavior leading up to a distributed denial of service attack and predictively adjust its reputation. The blue line shows the deviation from an IP address' average message count. The message deviation (blue spike) that occurs in the first-third of the graph prompts the connection's reputation score (measured by the vertical bars) to rise and change from "unverified" (gray) to "high risk" (red). When the actual attack is carried out, shown by the deviation from the IP address' average connection count (yellow line), that "high risk" reputation score tells McAfee products to block the messages to protect customers.

Reputation is not just an important component of any security system; it is essential. Threats move too quickly or too stealthily to rely on traditional techniques such as signature-based protection and blacklists. If a threat's intention is to hit as many computers as possible, it can propagate much more quickly than a signature can be written and deployed, and blacklist solutions don't capture the nuances that a reputation score does. On the other end of the spectrum, we see razor-targeted threats whose goal is not to spread quickly but instead to avoid detection, cause minimal impact, and achieve a very subtle, directed objective. To combat each of these extremes (and everything in between), security professionals and their vendors realize that today's threat landscape requires a system that calculates an entity's reputation in real time based on collective intelligence about that entity, and then takes action based on that reputation.

Operation Aurora, the attack against Google and more than 20 other companies in late 2009 and early 2010, used a directed effort to zero in on a specific set of individuals. The attackers used sophisticated, evasive technologies to gain access to those users' machines and, from there, to companies' valuable information and intellectual property. Despite their subtlety and efforts to avoid detection, threats such as Operation Aurora have a small number of associated entities—emails emanating from temporarily bad IP addresses luring unsuspecting users to malware-infected websites, for example—whose reputations can change from one moment to the next. McAfee uses changes in these reputations to automatically detect and prevent malicious activity, protecting essential people, assets, and information.

Threat Dynamics

The fast-moving and ever-changing nature of electronic threats requires a reputation system that takes this dynamism into account. As systems and researchers learn more information about an entity, that information should be used to adjust the entity's reputation on a continuous basis. For example, a legitimate computer becomes infected with Trojan malware that causes it to be part of a spam-sending botnet. Within a short time the computer is cleaned and safe. In a matter of minutes, the computer has come full circle—from low risk to high risk and back again. An effective reputation system should be sensitive enough to reflect the accurate state of the computer at any time. (See Figure 2.)

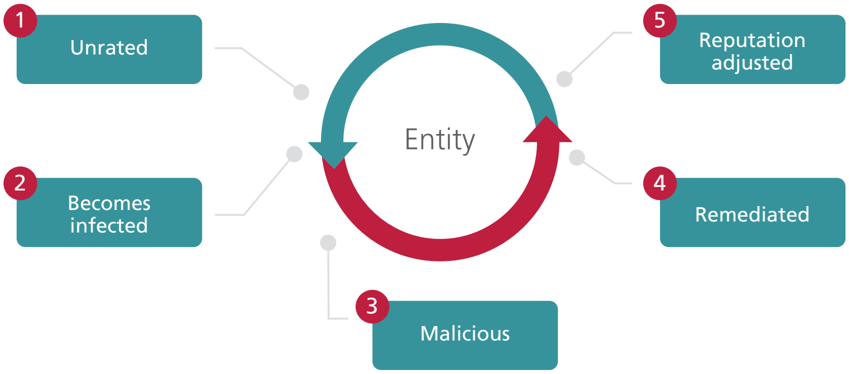

Figure 2. A strong reputation system should recognize even short-term changes to the entities it tracks.

Queries and their responses are important inputs into a reputation score. Robust reputation systems, those with millions of products deployed in real-world settings, maintain a feedback loop in which those products query the reputation system upon some local product trigger, and draw on the system's response to determine local action. Even the queries and responses themselves—based on their volume and frequency—can trigger the system to change an entity's reputation. For example, if a new IP address goes online and email gateways deployed around the world quickly query enough messages sent from it, the reputation system may increase its confidence that the IP is sending spam, and would accordingly adjust the connection's reputation. Similarly, when an infected entity is cleaned up, the reputation system would reflect the improved status.

A valuable benefit of having a reputation system that incorporates and adjusts itself with each additional data point is that cybercriminals who attempt to test malware or do a "dry run" of a network attack may inadvertently alert the system of their activities. An effective reputation system forces cybercriminals to choose between testing their ware in the real world, and being thwarted in the process, or forging ahead with an attack with an arsenal of untested tools. Either way, the system puts the criminals at a disadvantage.

Because we are dealing with both ends of the cybersecurity spectrum, from rapidly propagating to under-the-radar threats, an effective reputation system must collect data efficiently, rapidly analyze large data sets, and nearly instantaneously distribute the results to computers across the globe. The world is full of examples of computer systems that are optimized for one or perhaps two of those dimensions, but few exist that perform all three well.

Gray Matters

*What tipped us off that something was amiss about www.multimedia***.com and the other domains? It was the domains' behavior compared with our expectations. For example, some of the media-sharing domains didn't follow the traffic pattern usually exhibited by media-sharing sites. Although some acted normally, others acted like hosting domains, a model often used for redirecting traffic to obfuscate spam-sending IP addresses or for hosting malware executable files, botnet control instructions, or phishing credentials. We also noticed that several of the domains were rapidly shutting down and showing up on different IP addresses, a practice called fast flux that is used to avoid detection. And we also saw that many of the behaviors, such as domains moving from one IP address to another at the same time, were happening at the same instant, indicating to us that the domains were controlled by a single entity. As this story unfolded, our systems downgraded the domains reputations to "high risk."*

Reputation systems differ from blacklist and whitelist technologies: The former deal with the gray area in between "good" and "bad"; the latter are static, with administrators adding to and subtracting from them at regular intervals. Reputations, on the other hand, change each time the system "learns" something new; so they are inherently fluid. More important, the high number of online entities, combined with their dynamic reputations, makes it nearly impossible to say that an entity is 100 percent anything. By the time we could finish scanning a system to determine whether it is fully good or bad, it might already have changed. Because reputation systems help security products make split-second decisions, they need to provide the best possible answer at any time. Given this uncertainty, the absence of a reputation system means that an organization necessarily ends up under- or over-blocking threats. A robust reputation system can derive dynamic reputations with fewer false-positives, for example, than a blacklist, giving administrators a higher degree of accuracy. So the gray area really does matter, and that's where dynamic reputation scores meet policy. (See Figure 3.)

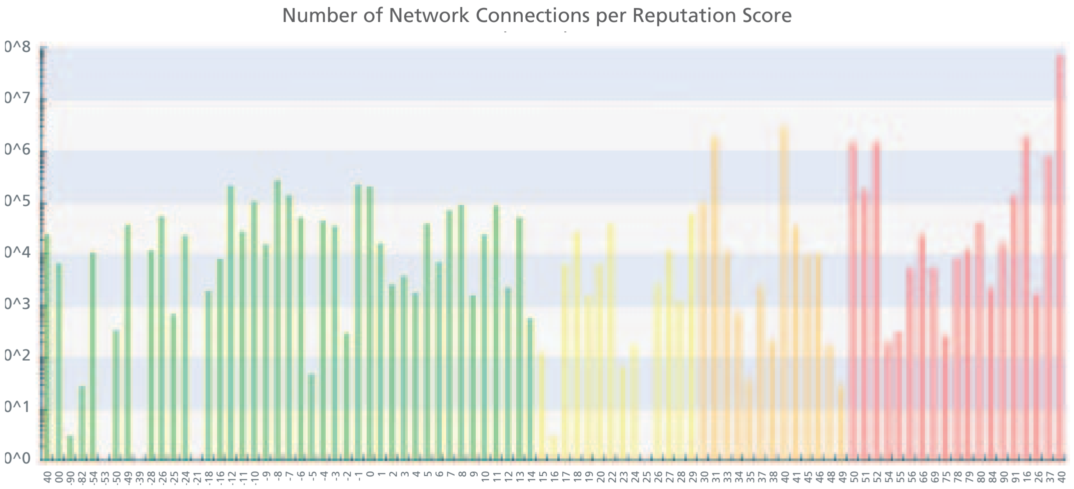

Figure 3. This graph illustrates the granularity of McAfee's reputation scoring system. We plot the network connections we track against each of the reputation scores along a continuum (the x-axis). The y-axis, reflected in a logarithmic scale, shows how many connections have each reputation score at a point in time. The colors indicate the scores at which we consider the risk levels of the connections to be low (green), unverified (yellow), medium (orange), or high (red).

Several elements determine whether a company should take security-related actions such as blocking a file or curtailing network communications: an organization's risk profile, productivity requirements, tolerance for false-positives, asset criticality, alternative security measures, and a variety of other factors. Reputation systems arm organizations with consistent, objective scoring and empower decision makers to set policies based on the risks and tradeoffs specific to their own organizations. Thus their security infrastructure can automatically take action in accordance with those policies.

Confidence Building

Because we are dealing with shades of gray rather than simple black and white, boosting the confidence level of our assessment of an entity's reputation is essential. Security professionals and their organizations rely on reputation scores to make dynamic policy decisions based on known probabilities, and it is the job of the reputation system to ensure that those scores reflect the highest possible confidence level. The more dimensions we take into consideration when calculating a score, the higher our confidence.

A useful analogy is a medical diagnosis. How does a doctor ascertain a patient's illness? She follows a series of steps, in which the goal of each step is either to zero in on a hypothesis or to gain increasing confidence that the hypothesis is correct. The doctor may start by asking the patient what's wrong, and then take the patient's temperature. Based on these two activities, she may have some indication of what the problem is. But it isn't until she correlates that information with a new piece of data, blood pressure, for example, that she increases her confidence in her hypothesis. To reach a high confidence level, she may need to look at a dozen dimensions that, on their own, don't tell her much but, correlated with each other, offer her high confidence that her diagnosis is accurate. Taking the analogy one step further, the doctor can now take a predictive stance. Following her diagnosis of one patient, she may find that ten people with the same symptoms have just walked into her office in a 10-minute period, allowing her to draw from one diagnosis to diagnose others. From there, the doctor and her medical colleagues can take action to block the illness where it is likely to strike next.

Similarly, cybersecurity reputation systems rely on the correlation of data across a large number of dimensions. To illustrate the increase in confidence that results from multidimensional correlation, in Figure 4 we plot known bad (red) and known good (green) IP addresses side-by-side in three graphs. The first shows the IPs along a single dimension, message volume. In this graph, it is difficult to ascertain which are bad and good, and the line of delineation between the two. Adding a second dimension, IP address persistence (or an IP address' continuous existence), in the second graph allows greater separation of the data, giving the viewer more certainty in determining to which group a data point is likely to belong. Adding a third dimension, the breadth of an IP address' recipient base, in the third graph gives a much more nuanced view of the data, showing the IPs in easily discernable clusters that make predictions much more accurate. In our actual analyses, we evaluate data across more than a thousand dimensions to boost our reputation-score confidence levels.

Multidimensional Reputation

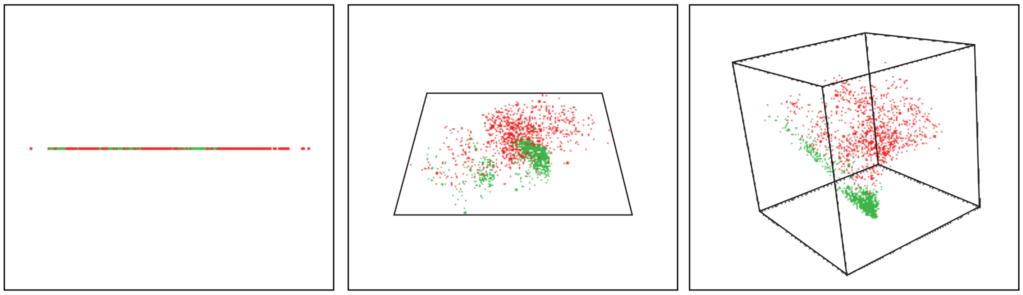

Figure 4. Adding dimensions to reputation scores increases the confidence level of the scores.

The ability to carry out robust reputation analysis across many dimensions depends on the ability to bring together enough relevant data from a broad set of sources. At McAfee, we gather data from our products deployed around the world as the foundation of our cloud-based intelligence, which feeds our reputation system. Besides serving as the basis for a robust reputation system, telemetry data are useful in building confidence levels in the following ways:

- **Data volume.** If each received query also serves as a data point influencing the overall system, more data translate to a higher degree of confidence in the reputation score. Think of this as the aperture on a telescope: the more data volume (light) taken in, the deeper the viewer can see into space. At McAfee, we receive billions of queries per day in our cloud-based intelligence systems; this enables us to see threat activity quickly and identify it with greater accuracy.
- **Data longevity.** Collecting data over a long period contributes to system maturity. It ensures that the reputation system will have a solid baseline for how entities are expected to behave based on their and their peers past behavior. This helps not only detect anomalies when they occur but also identify attacks based on recognized patterns.
- **Data trustworthiness.** A serious consideration when dealing with reputation systems is trustworthy data. Both the volume of data collected and the automated methodology for analyzing the data open the door for pollution of the results, such as by user collusion. Robust reputation systems must have mechanisms for authenticating the data they receive as well as for adjusting for the credibility of the source. Factors such as the location, configuration, and past behavior of the data source can affect how heavily that source's data will be weighed in the overall reputation calculation.
- **Data correlation.** The most critical factor in a robust reputation system is the ability to collect and correlate telemetry data from a broad range of sources representing *all threat vectors*. At McAfee, we leverage our span of products, including anti-malware on endpoints, web and email gateways and firewalls at the perimeter, and intrusion prevention systems to examine threats from all facets—file, web, email, network, and even application. Being able to correlate data representing a 360-degree view of a threat is like having all of the edge pieces of a puzzle.

The Power of Reputation

Pulling together telemetry data from all vectors helps us understand a threat and gives us far greater precision in calculating the reputation of any entities involved with the threat. Figure 5 represents one way McAfee uses telemetry data gathered from one threat vector to identify the threat in other vectors and update its reputation systems to protect our customers. Our reputation system receives queries from sensors around the world requesting information about potential malware, web threats, network connections, and email messages, among other things. In this example, our anti-malware client technology sends us file-reputation queries based on a hash, or fingerprint, of the file. The number, frequency, and geographical distribution of queries lead us to establish with a high level of confidence—even without the file—that it is malware. Our reputation system returns a score that causes the anti-malware client software to block or quarantine the file. Separately, an email sender at an IP address that we have not seen before attempts to send an email with a file with the same hash to a user behind one of our email gateways. The gateway queries our cloud, learns that it is malware, and blocks the message. Our system scans our database for websites hosting a file with the hash, adjusting their web reputations and the websites' associated network connections, to "high risk," and retrieves the malware file for further processing. Finally, because our anti-malware clients, email and web gateways, and firewalls all query our reputation system, the threat is blocked regardless of the vector over which it arrives.

Global Threat Intelligence at Work

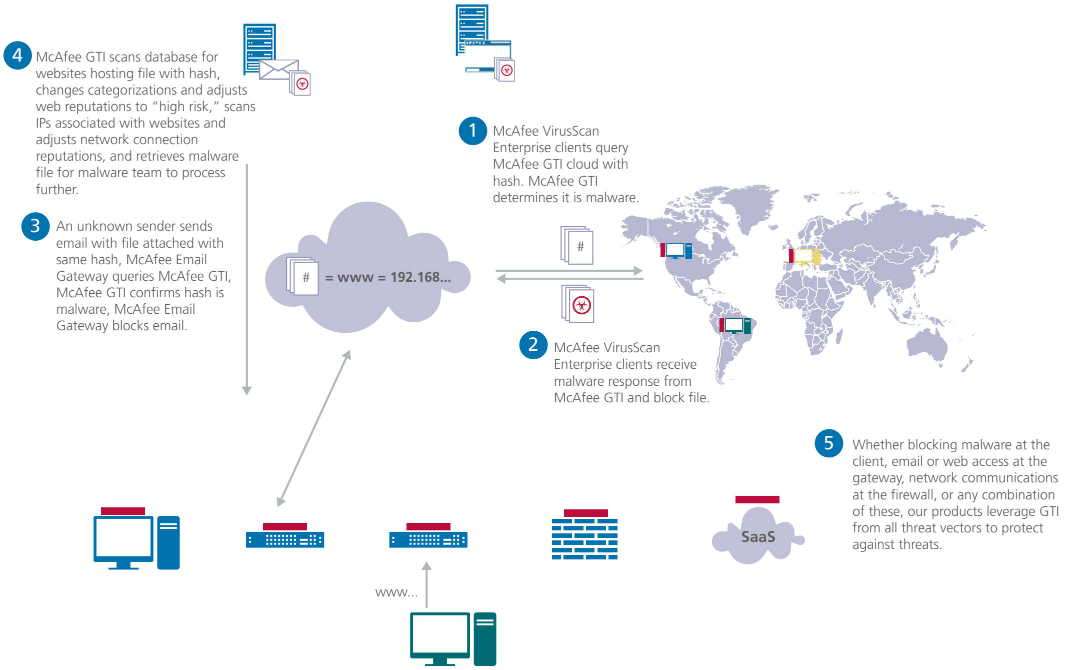

4 McAfee GTI scans database for websites hosting file with hash, changes categorizations and adjusts web reputations to "high risk," scans IPs associated with websites and adjusts network connection reputations, and retrieves malware file for malware team to process further.

3 An unknown sender sends email with file attached with same hash, McAfee Email Gateway queries McAfee GTI, McAfee GTI confirms hash is malware, McAfee Email Gateway blocks email.

1 McAfee VirusScan Enterprise clients query McAfee GTI cloud with hash. McAfee GTI determines it is malware.

2 McAfee VirusScan Enterprise clients receive malware response from McAfee GTI and block file.

5 Whether blocking malware at the client, email or web access at the gateway, network communications at the firewall, or any combination of these, our products leverage GTI from all threat vectors to protect against threats.

Figure 5. No matter the threat vector, McAfee products can query our cloud-based intelligence and stop new threats from causing damage.

Conclusion

*Our story ends as we learn www.multimedia***.com and the other 159 suspicious domains were an attempted widespread phishing attack based on Zeus, a well-known builder application for password-stealing Trojan malware. The malicious domains were actually phishing sites for stealing login credentials. We examined the many entities that "touched" these domains: IPs hosting the domains, emails with embedded links to URLs within the domains, malware files hosted on the domains, etc. We analyzed these entities and, when appropriate, adjusted their reputations to "high risk" in our systems so that McAfee products, deployed locally in corporate and consumer environments, could protect against the threats, regardless of their delivery mechanisms.*

The nature of today's cybersecurity landscape calls for smart and sophisticated defenses, and a robust reputation system is a critical component. The notion of reputation-based security has been around for years, but today we must deal with a rapidly growing number of threats ranging from fast-spreading viruses to narrowly targeted and evasive IP heists to everything in between. This challenge requires a consistent, objective security framework for understanding and calculating the status of an incredibly dynamic set of entities. Knowing an entity's status with a high degree of confidence—derived from a trustworthy set of correlated telemetry data—is the keystone for providing comprehensive protection.

Earning a Reputation

In our cloud-based reputation systems, McAfee calculates reputations for electronic entities in the following ways. These reputations intersect with our products' policies to enable security professionals to make the right decisions based on their organizations' risk profiles and business needs.

File reputation. McAfee's cloud-based system receives daily nearly 50 million file reputation queries (based on a file hash) and responds with a score that reflects the likelihood that the file in question is malware. The score is based not only on the collective intelligence from sensors querying the McAfee cloud and the analysis performed by the researchers and automated tools of McAfee Labs™, but also on the correlation of cross-vector intelligence from web, email, and network threat data. The local McAfee anti-malware engine—whether deployed as part of an endpoint anti-malware, gateway, or other solution—uses the score to determine action (block, quarantine, let pass, etc.) based on local policy.

Web reputation. McAfee's cloud-based system receives daily 2.5 billion web reputation queries and responds with a score that reflects the likelihood that the URL, web domain, or DNS server in question is malicious (phishing site, infected with malware, etc.). The score is based not only on the collective intelligence from sensors querying the McAfee cloud and the analysis performed by the researchers and automated tools of McAfee Labs, but also on the correlation of cross-vector intelligence from file, email, and network threat data. The local McAfee product—such as the McAfee Web Gateway—uses the score in combination with its local engine to determine action based on local policy. McAfee not only calculates reputations for URLs, but also for domains, their associated IP addresses, and DNS servers.

Message reputation. McAfee receives hundreds of millions of email queries daily, takes a fingerprint of the message content (versus the content itself, for privacy reasons), and analyzes it along many dimensions. Message reputation combines with factors such as spam-sending patterns and IP behavior to determine the likelihood that the message in question is malicious (spam, malware, etc.). The score is based not only on the collective intelligence from sensors querying the McAfee cloud and the analysis performed by the researchers and automated tools of McAfee Labs, but also on the correlation of cross-vector intelligence from file, web, and network threat data. The local McAfee product—such as an email gateway—uses the score to determine action based on local policy.

Network connection reputation. McAfee collects information from billions of IP addresses and network ports, providing hundreds of trillions of unique views, and calculates a reputation score based on data about network traffic, including port, destination, protocol, and inbound and outbound connection requests. The score reflects the likelihood that a network connection poses a threat (for example, is associated with botnet control). The score is based not only on the collective intelligence from sensors querying the McAfee cloud and the analysis performed by the researchers and automated tools of McAfee Labs, but also on the correlation of cross-vector intelligence from file, web, and network threat data. The local McAfee product—whether a firewall or intrusion prevention system—uses the score to determine action based on local policy.

About the Author

Jamie Barnett is Director of Product Marketing for McAfee Global Threat Intelligence. Prior to working at McAfee, she served as VP Product Management and Marketing for software company Blue Vector, and prior to that cofounded EMC's security initiative, leading the charge on the data management company's acquisition of RSA Security. A closet strategist and wannabe technologist, Barnett pretends she's not in marketing. When she's not writing white papers, she watches reruns of Monty Python's Flying Circus, which she knows by heart.

About McAfee Labs™

McAfee Labs is the global research team of McAfee, Inc. With the only research organization devoted to all threat vectors—malware, web, email, network, and vulnerabilities—McAfee Labs gathers intelligence from its millions of sensors and its cloud-based reputation technologies such as McAfee® Artemis™ and TrustedSource.™ The McAfee Labs team of 350 multidisciplinary researchers in 30 countries follows the complete range of threats in real time, identifying application vulnerabilities, analyzing and correlating risks, and enabling instant remediation to protect enterprises and the public.

About McAfee, Inc.

McAfee, Inc., headquartered in Santa Clara, California, is the world's largest dedicated security technology company. McAfee is relentlessly committed to tackling the world's toughest security challenges. The company delivers proactive and proven solutions and services that help secure systems and networks around the world, allowing users to safely connect to the Internet, browse, and shop the web more securely. Backed by an award-winning research team, McAfee creates innovative products that empower home users, businesses, the public sector, and service providers by enabling them to prove compliance with regulations, protect data, prevent disruptions, identify vulnerabilities, and continuously monitor and improve their security. www.mcafee.com.

McAfee, Inc.
3965 Freedom Circle
Santa Clara, CA 95054
888 847 8766
www.mcafee.com

McAfee, the McAfee logo, McAfee Labs, and TrustedSource are registered trademarks or trademarks of McAfee, Inc. or its subsidiaries in the United States and other countries. Other marks and brands may be claimed as the property of others. The product plans, specifications, and descriptions herein are provided only for information and are subject to change without notice. They are provided without warranty of any kind, expressed or implied. Copyright © 2010 McAfee, Inc.