Application Note

TrustedSource™ in McAfee® Firewall Enterprise

McAfee® Firewall Enterprise

version 8.1.0 and earlier

This document uses a question and answer format to explain the TrustedSource reputation service in McAfee Firewall Enterprise version 8.1.0 and earlier to system administrators, giving information on scope and usage and addressing performance and security concerns.

In this document ...

About TrustedSource

What is TrustedSource?

TrustedSource™ is a global Internet reputation intelligence system that determines what is good and bad behavior on the Internet by using real-time analysis of worldwide behavioral and sending patterns for email, web activity, malware, and machine-to-machine behavior. Using data obtained from the analysis, TrustedSource dynamically calculates reputation scores that represent the level of risk posed to your network when you visit a web page. The result is a database of reputation scores for IP addresses, domains, specific messages, URLs, and images.

I use email and web reputation already. Why do I need reputation for other protocols?

Because hackers don’t limit themselves to email and web. For example, a common blended attack used to involve using phishing email to get a user to click on a link to a web site, which would then install malware on the user's computer. As web reputation systems have become more common, some attackers have started to bypass them entirely by using FTP links to deliver the malware. Botnets don't just send spam and host malware, they also launch scans and mass attacks against a broad variety of vulnerabilities. Reputation can help protect your network from this activity.

How does it work?

When a connection is attempted through the firewall, it is compared against each policy rule in order until a match is found. Each rule can be configured with a reputation threshold. If the connection's reputation does not meet the threshold, the rule does not match. Thus, for example, if a rule allows only connections with Neutral (Low Risk in version 8.x) reputation or better, a connection with Suspicious (Medium Risk in version 8.x) reputation will not match the rule. It might match a later, less stringent rule, or it might only match the “deny all” rule at the end of the policy. Thus, reputation can be used not only for simple blacklisting, but also to create a graduated security policy for riskier connections.

Detailed configuration instructions are in the latest *McAfee Firewall Enterprise Product Guide*.

What is “connection reputation”?

TrustedSource reputation is the result of sophisticated behavioral analysis algorithms. A reputation query includes not just the remote IP address, but information about what ports and protocol are being used. This helps the TrustedSource servers detect subtler patterns of anomalous behavior.

Why would one address have different reputation for different protocols?

TrustedSource reputation data is derived from many sources. One of the principal goals of TrustedSource is to track malicious botnets such as the Storm worm. An infected botnet host might have a bad reputation in all protocols.

However, some risks are more targeted. For example, a commercial mail server might have a poor mail reputation from sending unsolicited messages while still hosting a perfectly benign web site. As another example, if a new SSH bug is being broadly exploited, then SSH connections from hosts that have never been seen to use the protocol before are good candidates for a reputation adjustment.

What new things can I do with TrustedSource on McAfee Firewall Enterprise?

Creative use of reputation can give you finer control over risk on your network. For example:

- **Managing active web content** — JavaScript, Flash, and other active web content present many risks to your enterprise security. Despite this, it is usually impractical to disable it for all web sites, or even to maintain a whitelist of "trusted" web sites. Using TrustedSource on a McAfee Firewall Enterprise, however, you could allow active content from well-reputed sites, strip it from suspicious sites while still allowing access, and block malicious sites.
- **Detecting "pharming" attacks** — Pharming attacks trick a computer into using a hostile DNS server, which lets the attacker mount man-in-the-middle attacks against trusted sites, potentially inject malware, or even scan your internal network. TrustedSource reputation data can identify connections to rogue DNS servers used in such attacks. Blocking DNS queries to malicious servers stops attackers in their tracks, and can alert you to compromised internal systems.

Why wouldn't I just reject all traffic with a bad reputation?

Reputation is a useful tool in your security policy, but it is no substitute for risk analysis. In many cases, a host will have a bad reputation because it has been compromised without the knowledge of its owner. If consumers with a compromised home computer want to visit your public web site to read about your products, it makes little sense to protect them. On the other hand, if they want to liquidate their bank account, you might want to seek additional authentication. As is frequently the case in security, using reputation wisely is about balancing risk with availability.

Deployment concerns

Will it introduce latency? How much?

When a connection requires a TrustedSource reputation lookup, some latency is inevitable. We've done everything we can to minimize this.

First, we check reputation only when it is required to decide how to handle a connection. If the security policy allows a decision without a reputation check, none is made. Second, there is an intelligent caching architecture. In normal network usage patterns, most desired connections will be resolved by the cache without a live reputation query.

Finally, there are several TrustedSource data centers located around the world. When a query is made, your firewall will automatically direct it to the server that can give you the fastest response, usually in under 100 milliseconds. Altogether, this means that there is little to no latency added for most applications.

Will this processing require more firewall CPU?

No. In "worst case" system tests, reputation queries never consumed more than two percent of system CPU.

What if I can't reach the TrustedSource servers? Does my traffic stop?

If your firewall cannot reach any of the TrustedSource servers, it automatically assigns all applicable connections a default reputation, which is configurable in the Admin Console. The firewall checks periodically for server availability. As a result, your users don't suffer any latency waiting for reputation queries to time out, and you retain control over how traffic is handled in this situation.

Can a bottleneck occur because too much traffic is being TrustedSource validated?

Every effort has been made to prevent this. Reputation values are cached, and query responses are processed out of order to avoid bottlenecks. The policy subsystem on the firewall monitors latency in TrustedSource queries, and if delays are experienced for any reason your firewall will proactively fall back to using default reputation until a reliable connection can be made.

Can an attacker forge reputation information?

All queries from your firewall to the TrustedSource data centers occur through an encrypted tunnel secured with bidirectional certificate authentication. If your firewall detects a man-in-the-middle attack, it will alert you and fall back to using default reputation.

Does enabling TrustedSource expose information about my company?

A reputation query implicitly tells the reputation server that a connection was attempted. This data is statistically mingled with other queries, and provides a key input to the behavioral analysis algorithms that produce TrustedSource reputations. The reputation database is too large and dynamic to practically distribute to individual systems. While we do not individually inspect this data, and make every effort to protect it, if the traffic patterns passing through your firewall are themselves highly confidential, you should avoid using reputation in your policy.

In the future, McAfee Firewall Enterprise might be enhanced to provide additional behavioral data (indicating attacks detected by IPS, for example) to the TrustedSource data centers. Any such feedback will be optional, and will need to be enabled by a system administrator.

How can I check my own network's reputation?

Visit the TrustedSource web site at www.trustedsource.org to check the reputation of your domain or specific IP addresses.