Poster Proposal: Detecting Zero-Day Self-Propagating Internet Worms Based on Their Fundamental Behavior

Shad Stafford, Toby Ehrenkranz, and Jun Li
Department of Computer and Information Science
University of Oregon
{staffors, tehrenkr, lijun}@cs.uoregon.edu

Self-propagating worms pose a significant threat to the health of the Internet and rapid detection of them is of paramount importance. There are many existing worm detection mechanisms but all suffer from significant drawbacks: signature-based detection techniques are vulnerable to polymorphic worms, honeypots will not detect worms that do not scan random addresses, and systems that require installation on each end-host are cumbersome to deploy. In our research, we propose and evaluate a novel worm detection mechanism that detects worms based on their fundamental behaviors without which they can hardly propagate. It is installed at a network gateway, is entirely transparent to end-hosts, does not examine the payload of connections or do byte-stream comparisons, and does not have to assume that worms scan potential victims randomly in order to detect them.

Although there are many proposed approaches to worm detection, they fall short when attempting to detect zero-day worms due to various drawbacks. A number of these systems have focused on content-based detection which relies on using byte-pattern-based worm signatures to detect worm traffic. When the byte pattern of a given traffic flow matches the byte pattern defined by a worm signature, that traffic is identified as being worm traffic. In order to create these signatures, systems have been proposed to look for common byte stream patterns [1], [2], or to use honeypots to capture only illegitimate traffic sent towards unused addresses [3]. It is clearly computationally intensive to analyze the payloads of all traffic in order to compare against the byte pattern from a signature, but the greatest hurdle these systems face is generally considered to be polymorphic worms—worms which change the byte patterns in their payload. Without similar payloads between worm connections, it is very difficult to create a signature that is able to match all of the worm connections. More recent advances such as [4] attempt to address this issue, but research in [5] shows that such schemes will likely never be sufficient for detecting all polymorphic worms.

We propose a behavior-based solution that avoids the problem entirely. Instead of attempting to fix the polymorphism issue in signatures and content-based systems, our solution looks for the essential characteristics of all worms—characteristics which do not require the examination of payload byte patterns.

Our solution works as follows at a high level. First, we apply two novel heuristics to capture essential characteristics of worms at the connection level. Because a host will only begin sending out worm traffic after itself being infected by worm traffic, the first heuristic looks for connections which were possibly caused by earlier *similar* connections. Furthermore, since worm-infected hosts will send worm traffic towards many unique hosts that would not normally be contacted, a second heuristic looks for connections which deviate from a host's normal destination visiting pattern. Connections which are flagged by both heuristics are considered *wormlike*. Second, we look at the number of wormlike connections during a sliding window. There may be some legitimate traffic which appears wormlike, and having a wormlike connection does not automatically trigger a worm alert. However, when a worm is indeed occurring, we will *consistently* see wormlike connections. When the number of wormlike connections observed over a sliding window reaches a threshold value, we can then detect that there is a worm active.

We have also thoroughly evaluated our worm detection system. We use a mixture of real traffic traces and simulated worm traffic. The simulation of worm traffic allows us to evaluate the system against a variety of worm speeds, scanning algorithms, and vulnerability levels. Through these experiments and measurements, we have shown that our system is both fast and accurate and is able to detect worms with a broad range of scanning algorithms, speeds, and vulnerability levels. In our experiments, the system was able to detect the presence of a worm in a network with zero false-positives or false-negatives, and furthermore was able to correctly identify individual infected hosts with high speed and accuracy.

References

[1] S. Singh, C. Estan, G. Varghese, and S. Savage, "Automated worm fingerprinting," in *Proceedings of the 6th Symposium on Operating System Design and Implementation (OSDI)*, 2004.

[2] H.-A. Kim and B. Karp, "Autograph: Toward automated, distributed worm signature detection," in *USENIX Security Symposium*, August 2004, pp. 271–286.

[3] C. Kreibich and J. Crowcroft, "Honeycomb: Creating intrusion detection signatures using honeypots," in *Workshop on Hot Topics in Networks*, 2003.

[4] J. Newsome, B. Karp, and D. Song, "Polygraph: Automatically generating signatures for polymorphic worms," in *Symposium on Security and Privacy*, 2005.

[5] J. R. Crandall, Z. Su, S. F. Wu, and F. T. Chong, "On deriving unknown vulnerabilities from zero-day polymorphic and metamorphic worm exploits," in *Conference on Computer and Communications Security*, 2005.