[1] P. Barford and V. Yegneswaran. An inside look at botnets. In *Special Workshop on Malware Detection, Advances in Information Security*. Springer Verlag, 2006. [bib | .pdf]

The continued growth and diversification of the Internet has been accompanied by an increasing prevalence of attacks and intrusions [20]. It can be argued, however, that a significant change in motivation for malicious activity has taken place over the past several years: from vandalism and recognition in the hacker community, to attacks and intrusions for financial gain. This shift has been marked by a growing sophistication in the tools and methods used to conduct attacks, thereby escalating the network security arms race.

Our thesis is that the *reactive* methods for network security that are predominant today are ultimately insufficient and that more proactive methods are required. One such approach is to develop a foundational understanding of the mechanisms employed by malicious software (malware) which is often readily available in source form on the Internet. While it is well known that large IT security companies maintain detailed databases of this information, these are not openly available and we are not aware of any such open repository. In this paper we begin the process of codifying the capabilities of malware by dissecting four widely-used Internet Relay Chat (IRC) botnet codebases. Each codebase is classified along seven key dimensions including botnet control mechanisms, host control mechanisms, propagation mechanisms, exploits, delivery mechanisms, obfuscation and deception mechanisms. Our study reveals the complexity of botnet software, and we discuss implications for defense strategies based on our analysis.

[2] J. R. Binkley and S. Singh. An algorithm for anomaly-based botnet detection. In *Proceedings of USENIX Steps to Reducing Unwanted Traffic on the Internet Workshop (SRUTI)*, pages 43-48, July 2006. [bib | .html]

We present an anomaly-based algorithm for detecting IRC-based botnet meshes. The algorithm combines an IRC mesh detection component with a TCP scan detection heuristic called the TCP work weight. The IRC component produces two tuples, one for determining the IRC mesh based on IP channel names, and a sub-tuple which collects statistics (including the TCP work weight) on individual IRC hosts in channels. We sort the channels by the number of scanners producing a sorted list of potential botnets. This algorithm has been deployed in PSU's DMZ for over a year and has proven effective in reducing the number of botnet clients.

[3] BotHunter distribution page, 2007. http://www.cyber-ta.org/releases/botHunter/. [bib | http]

[4] E. Cooke, F. Jahanian, and D. McPherson. The zombie roundup: Understanding, detecting, and disrupting botnets. In *Proceedings of USENIX Steps to Reducing Unwanted Traffic on the Internet Workshop (SRUTI)*, pages 39-44, July 2005. [bib | .html | .pdf]

Global Internet threats are undergoing a profound transformation from attacks designed solely to disable infrastructure to those that also target people and organizations. Behind these new attacks is a large pool of compromised hosts sitting in homes, schools, businesses, and governments around the world. These systems are infected with a bot that communicates with a bot controller and other bots to form what is commonly referred to as a zombie army or botnet. Botnets are a very real and quickly evolving problem that is still not well understood or studied. In this paper we outline the origins and structure of bots and botnets and use data from the operator community, the Internet Motion Sensor project, and a honeypot experiment to illustrate the botnet problem today. We then study the effectiveness of detecting botnets by directly monitoring IRC communication or other command and control activity and show a more comprehensive approach is required. We conclude by describing a system to detect botnets that utilize advanced command and control systems by correlating secondary detection data from multiple sources.

[5] D. Dash, B. Kveton, J. M. Agosta, E. Schooler, J. Chandrashekar, A. Bachrach, and A. Newman. When gossip is good: Distributed probabilistic inference for detection of slow network intrusions. In *Proceedings of the 21st National Conference on Artificial Intelligence*, pages 1115-1122, July 2006. [bib | .pdf]

Intrusion attempts due to self-propagating code are becoming an increasingly urgent problem, in part due to the homogeneous makeup of the internet. Recent advances in anomalybased intrusion

detection systems (IDSs) have made use of the quickly spreading nature of these attacks to identify them with high sensitivity and at low false positive (FP) rates. However, slowly propagating attacks are much more difficult to detect because they are cloaked under the veil of normal network traffic, yet can be just as dangerous due to their exponential spread pattern. We extend the idea of using collaborative IDSs to corroborate the likelihood of attack by imbuing end hosts with probabilistic graphical models and using random messaging to gossip state among peer detectors. We show that such a system is able to boost a weak anomaly detector D to detect an order-of-magnitude slower worm, at false positive rates less than a few per week, than would be possible using D alone at the end-host or on a network aggregation point. We show that this general architecture is scalable in the sense that a fixed absolute false positive rate can be achieved as the network size grows, spreads communication bandwidth uniformly throughout the network, and makes use of the increased computation power of a distributed system. We argue that using probabilistic models provides more robust detections than previous collaborative counting schemes and allows the system to account for heterogeneous detectors in a principled fashion.

[6] F. C. Freiling, T. Holz, and G. Wicherski. Botnet tracking: Exploring a root-cause methodology to prevent distributed denial-of-service attacks. In *Proceedings of the 10th European Symposium on Research in Computer Security (ESORICS)*, volume 3679 of *Lecture Notes in Computer Science*, pages 319-335, September 2005. [bib | DOI]

Denial-of-Service (DoS) attacks pose a significant threat to the Internet today especially if they are distributed, i.e., launched simultaneously at a large number of systems. *Reactive* techniques that try to detect such an attack and throttle down malicious traffic prevail today but usually require an additional infrastructure to be really effective. In this paper we show that *preventive* mechanisms can be as effective with much less effort: We present an approach to (distributed) DoS attack prevention that is based on the observation that coordinated automated activity by many hosts needs a mechanism to remotely control them. To prevent such attacks, it is therefore possible to identify, infiltrate and analyze this remote control mechanism and to stop it in an automated fashion. We show that this method can be realized in the Internet by describing how we infiltrated and tracked IRC-based *botnets* which are the main DoS technology used by attackers today.

[7] J. Göbel, J. Hektor, and T. Holz. Advanced honeypot-based intrusion detection. *USENIX ;login: magazine*, 31(6):17-25, December 2006. [bib]

[8] G. Gu, P. Porras, V. Yegneswaran, M. Fong, and W. Lee. BotHunter: Detecting malware infection through IDS-driven dialog correlation. In submission, 2007. [bib]

[9] S. T. King, Z. M. Mao, D. G. Lucchetti, and P. M. Chen. Enriching intrusion alerts through multi-host causality. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, 2005. [bib | .pdf]

Current intrusion detection systems point out suspicious states or events but do not show how the suspicious state or events relate to other states or events in the system. We show how to enrich an IDS alert with information about how those alerts causally lead to or result from other events in the system. By enriching IDS alerts with this type of causal information, we can leverage existing IDS alerts to learn more about the suspected attack. Backward causal graphs can be used to find which host allowed a multi-hop attack (such as a worm) to enter a local network; forward causal graphs can be used to find the other hosts that were affected by the multi-hop attack. We demonstrate this use of causality on a local network by tracking the Slapper worm, a manual attack that spreads via several attack vectors, and an e-mail virus. Causality can also be used to correlate distinct network and host IDS alerts. We demonstrate this use of causality by correlating Snort and host IDS alerts to reduce false positives on a testbed system connected to the Internet.

[10] Z. Li, Y. Chen, and A. Beach. Towards scalable and robust distributed intrusion alert fusion with good load balancing. In *Proceedings of the SIGCOMM Workshop on Large-Scale Attack Defense (LSAD)*, pages 115-122, 2006. [bib | DOI]

Traffic anomalies and distributed attacks are commonplace in today's networks. Single point detection is often insufficient to determine the causes, patterns and prevalence of such events. Most existing distributed intrusion detection systems (DIDS) rely on centralized fusion, or distributed fusion with unscalable communication mechanisms. In this paper, we propose to build a DIDS based on the

emerging decentralized location and routing infrastructure: *distributed hash table (DHT)*. We embed the intrusion symptoms into the DHT dimensions so that alarms related to the same intrusion (thus with similar symptoms) will be routed to the same sensor fusion center (SFC) while evenly distributing unrelated alarms to different SFCs. This is achieved through careful routing key design based on: 1) analysis of essential characteristics of four common types of intrusions: DoS attacks, port scanning, virus/worm infection and botnets; and 2) distribution and stability analysis of the popular port numbers and those of the popular source IP subnets in scans. We further propose several schemes to distribute the alarms more evenly across the SFCs, and improve the resiliency against the failures or attacks. Evaluation based on one month of DShield firewall logs (600 million scan records) collected from over 2200 worldwide providers show that the resulting system, termed *Cyber Disease DHT (CDDHT)*, can effectively fuse related alarms while distributing unrelated ones evenly among the SFCs. It significantly outperforms the traditional hierarchical approach when facing large amounts of *diverse* intrusion alerts.

[11] D. J. Malan and M. D. Smith. Host-based detection of worms through peer-to-peer cooperation. In *Proceedings of the 2005 ACM Workshop on Rapid Malcode (WORM)*, pages 72-80, November 2005. [bib | DOI]

We propose a host-based, runtime defense against worms that achieves negligible risk of false positives through peer-to-peer cooperation. We view correlation among otherwise independent peers' behavior as anomalous behavior, indication of a fast-spreading worm. We detect correlation by exploiting worms' *temporal consistency*, similarity (low temporal variance) in worms' invocations of system calls. We evaluate our ideas on Windows XP with Service Pack 2 using traces of nine variants of worms and twenty-five non-worms, including ten commercial applications and fifteen processes native to the platform. We find that two peers, upon exchanging snapshots of their internal behavior, defined with frequency distributions of system calls, can decide that they are, more likely than not, executing a worm between 76% and 97% of the time. More importantly, we find that the probability that peers might err, judging a non-worm a worm, is negligible.

[12] D. J. Malan and M. D. Smith. Exploiting temporal consistency to reduce false positives in host-based, collaborative detection of worms. In *Proceedings of the 4th ACM Workshop on Recurring Malcode (WORM)*, pages 25-32, 2006. [bib | DOI]

The speed of today's worms demands automated detection, but the risk of false positives poses a difficult problem. In prior work, we proposed a host-based intrusion-detection system for worms that leveraged collaboration among peers to lower its risk of false positives, and we simulated this approach for a system with two peers. In this paper, we build upon that work and evaluate our ideas “in the wild.” We implement Wormboy 2.0, a prototype of our vision that allows us to quantify and compare worms' and non-worms' temporal consistency, similarity over time in worms' and non-worms' invocations of system calls. We deploy our prototype to a network of 30 hosts running Windows XP with Service Pack 2 to monitor and analyze 10,776 processes, inclusive of 511 unique non-worms (873 if we consider unique versions to be unique non-worms). We identify properties with which we can distinguish non-worms from worms 99% of the time. We find that our collaborative architecture, using patterns of system calls and simple heuristics, can detect worms running on multiple peers. And we find that collaboration among peers significantly reduces our probability of false positives because of the unlikely appearance on many peers simultaneously of non-worm processes with worm-like properties.

[13] J. J. Parekh, K. Wang, and S. J. Stolfo. Privacy-preserving payload-based correlation for accurate malicious traffic detection. In *Proceedings of the SIGCOMM Workshop on Large-Scale Attack Defense (LSAD)*, pages 99-106, 2006. [bib | DOI]

With the increased use of botnets and other techniques to obfuscate attackers' command-and-control centers, Distributed Intrusion Detection Systems (DIDS) that focus on attack source IP addresses or other header information can only portray a limited view of distributed scans and attacks. Packet payload sharing techniques hold far more promise, as they can convey exploit vectors and/or malcode used upon successful exploit of a target system, irrespective of obfuscated source addresses. However, payload sharing has had minimal success due to regulatory or business-based privacy concerns of transmitting raw or even sanitized payloads. The currently accepted form of content exchange has been limited to the exchange of known-suspicious content, e.g., packets captured by honeypots; however, signature generation assumes that each site receives enough traffic

in order to correlate a meaningful set of payloads from which common content can be derived, and places fundamental and computationally stressful requirements on signature generators that may miss particularly stealthy or carefully-crafted polymorphic malcode.Instead, we propose a new approach to enable the sharing of suspicious payloads via *privacy-preserving* technologies. We detail the work we have done with two example payload anomaly detectors, PAYL and Anagram, to support generalized payload correlation and signature generation *without* releasing identifiable payload data and without relying on single-site signature generation. We present preliminary results of our approaches and suggest how such deployments may practically be used for not only cross-site, but also *cross-domain* alert sharing and its implications for profiling threats.

[14] M. A. Rajab, J. Zarfoss, F. Monrose, and A. Terzis. A multifaceted approach to understanding the botnet phenomenon. In *Proceedings of the 6th ACM SIGCOMM on Internet Measurement (IMC)*, pages 41-52, 2006. [bib | DOI]

The academic community has long acknowledged the existence of malicious botnets, however to date, very little is known about the behavior of these distributed computing platforms. To the best of our knowledge, botnet behavior has never been methodically studied, botnet prevalence on the Internet is mostly a mystery, and the botnet life cycle has yet to be modeled. Uncertainty abounds. In this paper, we attempt to clear the fog surrounding botnets by constructing a multifaceted and distributed measurement infrastructure. Throughout a period of more than three months, we used this infrastructure to track 192 unique IRC botnets of size ranging from a few hundred to several thousand infected end-hosts. Our results show that botnets represent a major contributor to unwanted Internet traffic - 27% of all malicious connection attempts observed from our distributed darknet can be directly attributed to botnet-related spreading activity. Furthermore, we discovered evidence of botnet infections in 11% of the 800,000 DNS domains we examined, indicating a high diversity among botnet victims. Taken as a whole, these results not only highlight the prominence of botnets, but also provide deep insights that may facilitate further research to curtail this phenomenon.

[15] snort web site, 2007. http://snort.org/. [bib | http]

[16] S. Stafford, J. Li, and T. Ehrenkranz. On the performance of SWORD in detecting zero-day-worm-infected hosts. In *Proceedings of the International Symposium on Performance Evaluation of Computer and Telecommunication Systems (SPECTS)*, 2006. [bib | .pdf]

[17] E. Stinson and J. C. Mitchell. Characterizing the remote control behavior of bots. In submission, 2007. [bib]

[18] R. Vogt and J. Aycock. Attack of the 50 foot botnet. Technical Report 2006-840-33, Department of Computer Science, University of Calgary, August 2006. [bib | http | .pdf]

The trend toward smaller botnets may be *more* dangerous in terms of large-scale attacks like distributed denials of service. We examine the possibility of “super-botnets,” networks of independent botnets that can be coordinated for attacks of unprecedented scale. For an adversary, super-botnets would also be extremely versatile and resistant to countermeasures. Our simulation results shed light on the feasibility and structure of super-botnets and some properties of their command-and-control mechanism. Possible defenses against the threat of super-botnets are suggested.

[19] J. Yang, P. Ning, X. S. Wang, and S. Jajodia. CARDS: A distributed system for detecting coordinated attacks. In *Proceedings of the IFIP TC11 Fifteenth Annual Working Conference on Information Security for Global Information Infrastructures*, pages 171-180, 2000. [bib | .pdf]

A major research problem in intrusion detection is the efficient Detection of coordinated attacks over large networks. Issues to be resolved include determining what data should be collected, which portion of the data should be analyzed, where the analysis of the data should take place, and howto correlate multi-source information. This paper proposes the architecture of a Coordinated Attack Response & Detection System (CARDS). CARDS uses a signature-based model for resolving these issues. It consists of signature managers, monitors, and directory services. The system collects data in a flexible, distributed manner, and the detection process is decentralized among various monitors and is event-driven. The paper also discusses related implementation issues.

[20] V. Yegneswaran, P. Barford, and J. Ullrich. Internet intrusions: global characteristics and

prevalence. In *Proceedings of the ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, pages 138-147, June 2003. [bib | DOI]

Network intrusions have been a fact of life in the Internet for many years. However, as is the case with many other types of Internet-wide phenomena, gaining insight into the *global* characteristics of intrusions is challenging. In this paper we address this problem by systematically analyzing a set of firewall logs collected over four months from over 1600 different networks world wide. The first part of our study is a general analysis focused on the issues of distribution, categorization and prevalence of intrusions. Our data shows both a large quantity and wide variety of intrusion attempts on a daily basis. We also find that worms like CodeRed, Nimda and SQL Snake persist long after their original release. By projecting intrusion activity as seen in our data sets to the entire Internet we determine that there are typically on the order of 25B intrusion attempts per day and that there is an increasing trend over our measurement period. We further find that sources of intrusions are uniformly spread across the Autonomous System space. However, deeper investigation reveals that a very small collection of sources are responsible for a significant fraction of intrusion attempts in any given month and their on/off patterns exhibit cliques of correlated behavior. We show that the distribution of source IP addresses of the non-worm intrusions as a function of the number of attempts follows Zipf's law. We also find that at daily timescales, intrusion targets often depict significant spatial trends that blur patterns observed from individual “IP telescopes”; this underscores the necessity for a more global approach to intrusion detection. Finally, we investigate the benefits of shared information, and the potential for using this as a foundation for an automated, global intrusion detection framework that would identify and isolate intrusions with greater precision and robustness than systems with limited perspective.

[21] C. C. Zou and R. Cunningham. Honeypot-aware advanced botnet construction and maintenance. In *Proceedings of the International Conference on Dependable Systems and Networks (DSN)*, pages 199-208, June 2006. [bib | DOI | .pdf]

Because “botnets” can be used for illicit financial gain, they have become quite popular in recent Internet attacks. “Honeypots” have been successfully deployed in many defense systems. Thus, attackers constructing and maintaining botnets will be forced to find ways to avoid honeypot traps. In this paper, we present a hardware and software independent honeypot detection methodology based on the following assumption: security professionals deploying honeypots have liability constraints such that they cannot allow their honeypots to participate in real (or too many real) attacks. Based on this assumption, attackers can detect honeypots in their botnet by checking whether the compromised machines in the botnet can successfully send out unmodified malicious traffic to attackers' sensors or whether the bot controller in their botnet can successfully relay potential attack commands. In addition, we present a novel “two-stage reconnaissance” worm that can automatically construct a peer-to-peer structured botnet and detect and remove infected honeypots during its propagation stage. Finally, we discuss some guidelines for defending against the general honeypot-aware attacks.

[22] C. C. Zou, D. Towsley, and W. Gong. A firewall network system for worm defense in enterprise networks. Technical Report TR-04-CSE-01, University of Massachusetts Amherst, College of Engineering, February 2004. [bib | .pdf]
