ARP-based Detection of Scanning Worms Within an Enterprise Network

David Whyte*   P.C. van Oorschot*   Evangelos Kranakis*

Abstract

Rapidly propagating worms are arguably the greatest security threat currently facing the Internet. To date, worm writers have been successful in penetrating most security countermeasures. Signature-based detection schemes often fail to detect zero-day worms, and their ability to rapidly react to new threats is limited as they typically require some form of human involvement to formulate updated attack signatures. We propose an anomaly-based detection technique designed to protect internal networks from scanning worm infections. This is the first publication in the open literature (to our knowledge) proposing and providing a detailed description of a method to detect propagation of scanning worms within individual network cells. We show that this technique is both accurate and rapid enough to enable automatic containment and suppression of worm propagation within a network cell. Implemented in software, our detection approach relies on an aggregate anomaly score, derived from the correlation of Address Resolution Protocol (ARP) activity from individual network attached devices. Our preliminary analysis and prototype indicate that this technique can be used to rapidly detect zero-day worms within a very small number of scans, e.g. three scans with a false positive rate of five over a two week period in our test environment. The necessary individual ARP activity system profiles are automatically generated during a training period and thus the software can be rapidly deployed with minimal tuning and administration.

1 Introduction

Scanning worms are rapidly evolving. Despite some very promising recent proposals (e.g. see [17]), most signature-based detection techniques are limited in their ability to detect newly emerging worms. Compounding the scanning worm detection challenge are the numerous worm variants typically launched after an initial outbreak. Slight code modifications enable worms to evade many

---

signature-based countermeasures. Worms can also employ a variety of anti-detection techniques such as: anti-forensics, dynamic behavior, and modularity of attack tools [13]. Unfortunately, even the best network security posture can fail due to newly discovered vulnerabilities, careless users, and improper system configuration. Clearly the advantage in this digital arms race lies with the worm writer. Network defenders must close all possible security exposures, whereas an attacker only needs to find one in order to penetrate the network.

Worm infected hosts must be rapidly detected in order to contain worm propagation. Infected hosts can potentially initiate thousands of infection scans per second [11] making automated suppression and containment strategies necessary. The detection methods used within these automated schemes must be both fast and accurate. Additionally, if the detection technique is prone to false positives, it can mean legitimate users and services are impacted. Many worms now spread so quickly that traditional intrusion detection methods (i.e. generation and deployment of attack signatures) are too slow [12]. Most current worm propagation detection methods are limited in: (1) their speed of detection, (2) their inability to accurately detect *zero-day* worms, (3) their inability to detect slow scanning worms, and (4) their high false positive rate.

Worm containment systems can contain an outbreak if the number of infected systems does not exceed an *epidemic threshold* [12, 22], depending on the properties of a worm. Containment within an enterprise network typically involves dividing the network into cells [4]. Although these cells can be as small as a single system (i.e. host-based containment) in a large network, the individual cells could be comprised of hundreds of hosts. The goal of such containment is to limit infection to individual cells; infection between hosts within a cell remains unchecked, with the potential loss of the hosts within a network cell considered an acceptable cost given the alternative of the entire network becoming infected. Topological worms focus their scans on a particular subnet (e.g. harvested from the infected hosts) of logically or physically connected systems [18], and communicate targeting information among infected hosts to bypass security countermeasures and increase the propagation rate. In this situation, limiting infection to a single network cell still al-

lows considerable damage within the cell. In this paper, we propose a detection technique that detects scanning worms attempting to propagate between hosts within an individual network cell. For the purposes of our approach, we define a *cell* as any portion of a network within a common broadcast domain.

Typically, scanning worms use a pseudo-random generator to generate 32-bit random numbers that correspond to the IPv4 addresses of their intended victims. The use of a numeric IP address instead of qualified domain name to locate the system obviates the need for a domain name system (DNS) query. It has been shown that this characteristic can be used as a behavioral signature[1] to detect scanning worms within an enterprise environment [20]. We propose another behavioral signature based on the observation that a scanning worm targeting systems within its own network cell exhibits anomalous behavior distinct from normal Address Resolution Protocol (ARP) [14] activity; infected systems trigger unusual ARP request activity as they try to infect susceptible systems within their respective network cells. More specifically, intra-cell worm-initiated scans result in discernible behavioral changes in the amount and pattern of ARP request activity of infected hosts, because a scanning worm targeting same-cell systems triggers the broadcast of anomalous ARP requests.

Our Contributions. We present a technique, which we have implemented and tested in a software prototype, to both rapidly and accurately detect worm propagation within enterprise network cells. Based on the following three factors, we derive an anomaly score for each individual device[2], and use this as an infection indicator for each device within a cell.

1. *Peer list*: connections to devices outside the set of internal devices a host normally interacts with.
2. *ARP activity*: increases in the average number of ARP requests each host issues per unit time.
3. *Internal network dark space* [1, 2]: connections to vacant internal IP addresses (i.e. addresses not bound to any active devices).

The greater the anomaly score attributed to a network device, the more likely it is infected with a scanning worm.

In our test environment, our technique detected simulated scanning worm activity with a minimum sustained scanning rate of one scan per minute, within three scans, with a very low false positive rate (see next paragraph). The precision of this first-mile detection enables the use of automated containment and suppression strategies [12] to stop scanning worms before they infect other devices

within a network cell. This scheme relies solely on the observation and correlation of activities related to a single protocol (i.e. ARP). ARP was designed as the intermediary between IP and MAC addresses within local subnets. Other technologies such as Windows Internet Naming Service (WINS), Internet Protocol (IP), and Domain Name Service (DNS) that enable communication outside of local subnets all rely on ARP.

Our anomaly-based detection approach is appealing for a number of reasons including the following:

1. *Speed*: the possibility to detect an infected system within e.g. three scans.
2. *Detection of zero-day worms*: possible because the approach does not rely on the matching of existing worm signatures to identify suspicious traffic.
3. *Scanning rate independence*: the approach can detect both fast and slow scanning worms (assuming a sustained scanning rate of at least one scan per minute - although this is somewhat configurable).
4. *Intra-cell protection*: our approach addresses infection of systems within a network cell.
5. *Low-false positive rate*: our prototype, albeit on a small network (see Section 4.1), experienced only five false positives within a two week period and this number could be further reduced (see Table 2).
6. *Ease of deployment*: the approach is network-based, runs on commodity hardware, and relies on the observation of a protocol found in every network (i.e. ARP).

Enterprise-wide scanning worm detection is a crucial component of any effective *network security posture* and thus has been the subject of considerable research. Although certain marketing-related documentation offers the possibility that intra-cell worm propagation is addressed in at least one available commercial product [2] (albeit not involving ARP), we are not able to find any details from available published materials. To the best of our knowledge, our paper is the first in the open literature to propose a method and detailed description for, and to report on an implementation of, a detection technique that enables scanning-worm detection *within* a network cell (i.e. intra-cell propagation).[3] Used in conjunction with other internal scanning worm containment schemes that address propagation between network cells [19, 20, 21], our technique enables a complete enterprise-wide scanning worm detection capability (i.e. both intra-cell and inter-cell).

Although our analysis confirms the validity of this technique, in networks of differing characteristics, its effectiveness may vary. Specifically, our technique may have limited applicability in networks where typically devices

---

[1]Behavioral signatures [6] are used to describe common aspects of worm behavior particular to a given worm.

[2]Hereafter by *device*, we typically mean a network addressable host within the broadcast network.

[3]Here, we assume a network cell contains more than a single device.

interact with a large percentage of other devices within the network cell (e.g. distributed or peer-to-peer networks (P2P)). In such instances, our technique could be integrated as a secondary input to a more sophisticated suite of anomaly detectors.

The sequel is structured as follows. Section 2 outlines the basic approach of the detection technique. Section 3 discusses our prototype, which is analyzed in Section 4 with details for a particular dataset and small testbed. Section 5 presents limitations. Section 6 discusses future work. Section 7 reviews related work. We conclude in Section 8.

2 Basic Methodology and Approach

In this section we give a high-level overview of our ARP-based scanning worm anomaly detection approach.

In larger enterprise networks, it is not unusual for network segments to be either logically or physically separated. In fact, an enterprise network may be comprised of several distinct subnets for a variety of reasons including security, ease of administration, and geographical location. This natural separation of networks can be leveraged to contain worm propagation within distinct network segments. Silicone Defense's *CounterMalice* solution [4] divides the enterprise network into segments called *cells* with the specific intention that each cell contains a worm containment device to confine and contain worm infection. As mentioned earlier, by our definition a cell is any portion of a network within a common broadcast domain. Figure 1 illustrates how an enterprise network can be divided into cells. The enterprise network is connected to the Internet through a single egress point (i.e. firewall).

As has been noted elsewhere (e.g. see [20]), the propagation of scanning worms can be characterized as: local to local (L2L), local to remote (L2R), or remote to local (R2L). In L2L propagation, the worm targets systems within the boundaries of the enterprise network in which it resides e.g. scanning within network cells (i.e. *intra-cell*) or between them (i.e. *inter-cell*)[4]. L2R propagation refers to a scanning worm within an enterprise network targeting systems outside of its network boundary. Finally, R2L propagation refers to worm scanning from the Internet into an enterprise network. Our approach addresses L2L worm propagation within a network cell (i.e. L2L intra-cell); it does not address L2L inter-cell, L2R, or R2L worm propagation.

2.1 ARP Anomaly Detection Approach

Devices that reside within the same network cell use ARP rather than the Domain Name Service (DNS) to communicate. ARP is a layer 2 protocol (i.e. data link layer)

---

[4]Topological scanning worms (e.g. Nimda) employ this strategy [18].
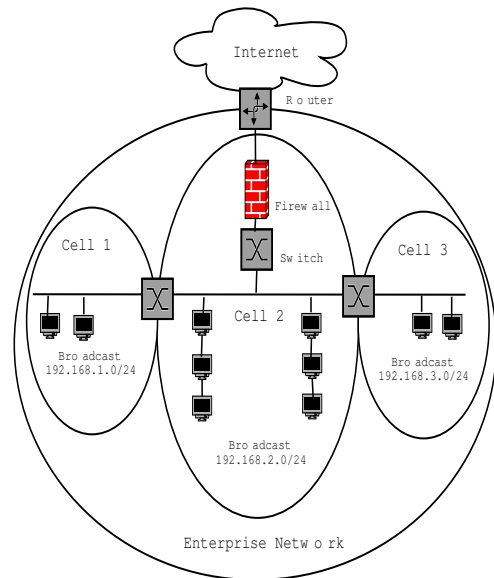
Figure 1. Dividing an Enterprise Network into Cells. Three network cells are separated by switches. Each cell has its own Class C network and distinct broadcast domain.

used by the IP protocol to map IP addresses to the physical hardware (MAC) addresses of network devices. When a device needs to resolve a given IP address to a MAC address, it broadcasts an ARP request. The ARP request packet contains the sender's IP address (source protocol address), the sender's MAC address (source hardware address), and the destination IP address (target protocol address). Each device within the common broadcast domain receives this packet. The protocol specifies that only the device with the specified destination IP address will respond with an ARP reply. An ARP reply contains both the IP address and MAC address of the device that responds. ARP activity is a necessary precursor to communications between devices as it provides the data link layer with the necessary mappings between the IP and MAC addresses of communicating devices.

For the purposes of our approach, we define a *scan* as an ARP request. An ARP request indicates that a system is trying to resolve an IP address to a MAC address for some type of connection. However, it is possible for a host to connect to a device without an immediately preceding ARP request. Once a device performs an ARP request, the MAC-to-IP address mapping within an ARP reply is maintained locally in the device that receives it in a table called an ARP cache. Only the device that made the ARP request receives the ARP reply. The ARP cache entries also have an associated time to live (ttl) and are dynamically entered and removed. If an ARP reply is re-

3

ceived by a device and the MAC address already appears within its cache, it is overwritten by the update. As long as the ARP reply remains in the local cache, subsequent connections to the same device will result in the MAC address being obtained from the cache rather than through an ARP request. The affect of ARP caches on our approach is discussed in Section 2.1.2. Our technique can be deployed on any device within the broadcast domain, even in a switched network fabric (see Section 6.1), and as it only processes ARP requests, it is extremely efficient.

L2L worm propagation can occur within a particular cell or span multiple cells depending on the scanning strategy of the worm. L2L scanning activity results in unusual ARP activity, namely: (1) an infected device will use ARP to try to connect to some devices within the internal network with which it had no previous history of connecting to pre-infection; (2) the number of ARP requests generated per fixed unit time (e.g. every 60 seconds) will increase; and (3) in those networks where not all IP addresses within a netblock have been allocated, ARP requests will be generated for nonexistent systems (i.e. for *internal network dark space*). We now discuss in turn how we use each of these behaviors to derive an aggregate anomaly score for each device within a cell.

2.1.1 Peer List (Customary ARP request targets)

A training period is used to determine normal device interactions within a cell. ARP requests observed during the training period allow us to characterize normal device interactions through the use of a peer list (see Figure 2). A peer list is indexed by all the devices which made ARP requests (i.e. served as a source protocol address within an ARP request) during a training period. Each ARP chain contains entries of the devices being queried in requests by the indexing device (i.e. target protocol address within the ARP request).

The source protocol address of an ARP request is the IP address of the device trying to resolve an IP address to its corresponding MAC address. The target protocol address, encapsulated within an ARP request, is the IP address of the system being queried which will be the eventual target of a connection. Each time an ARP request is generated, any observed new source protocol address is recorded as an index entry within the peer list. Figure 2 contains four entries within its index (i.e. 192.168.1.1, 192.168.1.2, 192.168.1.5, and 192.168.1.9). The corresponding target protocol addresses of the respective queries are added as entries indexed by the corresponding source protocol address. Over a training period, we build an index of active systems within the network cell (i.e. ARP requestors) and the list of devices (*ARP chains*) they are trying to connect with. The individual elements within each ARP chain are
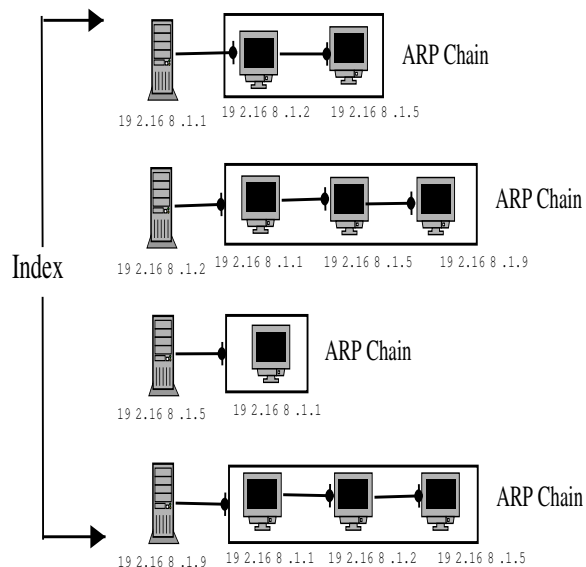
Figure 2. Peer List for a small network with four active devices. There are four ARP chains of 2, 3, 1 and 3 elements. Each ARP chain is comprised of an index value equal to an ARP requestor's IP address. The index and the associated collection of ARP chains is known collectively as the peer list.

derived from the set of IP addresses queried by the ARP requestor. If an ARP requestor queries the same device more than once, this activity is ignored (i.e. no duplicate entries exist within an ARP chain). Typically, individual devices will only communicate with a small subset of other internal devices that offer some sort of service (e.g. DNS, file, router, etc.).[5]

For each device i in a cell, we assign an anomaly score for the *peer list factor* (a_1) as: $a_{1,i} = x$ where x is the number of ARP requests as made by device i (in the current sample interval) which are outside of device i's ARP chain. In our testbed, the device identifier i corresponds to the last octet in the device's IP address. For instance, if two ARP requests are made by device 192.168.1.9 to two different IP addresses not contained within its ARP chain within the current sample interval, then $a_{1,9} = 2$. Subsequent distinct connection attempts outside a device's ARP chain within the detection window (see Section 2.2) result in a linear growth for this anomaly factor.

The running total of the a_1 factor for device i is the sum of the $a_{1,i}$ values over all sample intervals in the current detection window. More specifically, at sample interval j, let $a_{1,i}$ as defined above be shorthand for $a_{1,i}^{(j)}$; then the

---

[5] Of course, this assumption is violated e.g. in P2P networks and highly distributed cooperative network environments. We discuss this further in Section 5.

running total for $a_{1,i}$ for the current detection window of width w is $A_{1,i} = \sum_{k=0}^{w-1} a_{1,i}^{(j-k)}$.

If a device tries to connect to another device outside of both its ARP chain and the peer list, it is assigned an anomaly score based on the factor described in Section 2.1.3. Note that a possible source of false positives results from the fact that although perhaps suspicious, a device may legitimately contact another device outside of its ARP chain (e.g. a legitimate infrequently occurring interaction between two devices may not be captured during the training period thus not all such occurrences are necessarily an indication of a scanning worm).

2.1.2 ARP Activity (Number of ARP requests)

The number of ARP requests is recorded for each active device within the network over discrete sample intervals (e.g. 60 seconds) during the training period. Once the training period is complete, the mean ($\overline{x}$) and standard deviation (σ) of the observed ARP request activity are calculated for each individual device. The standard deviation reveals the normal fluctuation in ARP activity that can be expected during sample intervals.

We set a (somewhat arbitrary) upper bound and call it the *expected maximum ARP request activity* (E_i) *for device i* within a sample interval:

$$E_i = \overline{x} + 2\sigma \qquad (1)$$

A primary factor in choosing this value is that in a normal (i.e. Gaussian) distribution, 95% of the data values will fall within two standard deviations of the mean value; however, it should be clear that other selections may be equally or more useful. Once the training period has ended, the observed (i.e. subsequently monitored) ARP request activity O_i for each device i, is compared to E_i. $O_i > E_i$ may indicate anomalous scanning activity.

We assign an anomaly score for the *ARP activity factor* (a_2) for device i as follows:

$$a_{2,i} = \left\{ \begin{array}{ll} O_i - E_i & \text{; if } O_i > E_i \\ 0 & \text{; if } O_i \leq E_i \end{array} \right. \qquad (2)$$

Similar to the a_1 factor, this calculation is performed during each sample interval to determine a *running* total within the *detection window* (see Section 2.2) for each device. As device i's ARP request activity (O_i) increases beyond E_i, $a_{2,i}$ increases linearly. For instance, server 192.168.1.11 has $\overline{x} = 1.57974$, $\sigma = 1.03620$, and $E_{11} = 3.65214$. If $O_{11} = 5$ then $a_{2,11} = 1.34186$. This factor is particularly useful in addressing the affect of local ARP caches and large ARP chains (as discussed in Section 4.5).

2.1.3 Internal Network Dark Space

Internal network dark space is defined during the training period. Looking at the peer list in its entirety, we derive a set of internal system addresses that comprise the active systems within the cell during the training period. ARP requests for IP addresses not contained within this set we consider to be anomalous, and refer to as *internal network dark space*.

Each network consists of a block of network addresses. Although it is not unusual for an internal network to consist of IPv4 Internet addresses, most internal networks use non-routable IP address schemes. If RFC 1918 [15] is followed, the internal address space can be as large as a Class A network. This allows internal networks to consist of potentially millions of hosts. Regardless of whether the IP addresses within a network are non-routable or Internet-accessible, there may exist IP addresses that are not bound to any device within the cell. Connection attempts to these vacant IP addresses are anomalous. We assign an anomaly score for the *internal network dark space factor* (a_3) during a given sample interval as:

$$a_{3,i} = \left\{ \begin{array}{ll} 0 & \text{; if no dark space scans} \\ y & \text{; if device i scans dark space} \end{array} \right. \qquad (3)$$

We suggest that the value y be assigned such that it is the same for all devices (i) and a single observed connection to an internal network dark space address should generate a value $a_{3,i}$ sufficient on its own to meet the *alert threshold r* and generate an alarm (i.e. in our prototype, we set $y = r$; see Section 2.2).

2.2 Setting Alert Thresholds

Our implementation requires that a scanning worm exhibit a minimum sustained scanning threshold of one scan per minute. Therefore, we define a sample interval as 60 seconds (i.e. $t = 60$ seconds). The choice of sample interval directly affects the amount of state information that must be maintained by the prototype. The *detection window of width w* (= number of sample intervals) is the period of time in which observed anomaly scores for factors a_1 and a_2 must be maintained in state. In our implementation, we set $w = r$ (see definition of r below).

For example, for $r = 1$ the detection window is 60 seconds and an alert is generated upon a single anomalous scan observed within a one minute period. If we set $r = 2$, the detection window is 120 seconds and two anomalous scans must occur within two minutes to trigger an alarm. Anomalous scans get aged out over time; scans which slide out as the current detection window moves no longer contribute to the anomaly score. Anomaly scores are attributed to devices as ARP requests are processed by the prototype. Therefore, alarms can be generated at

any time regardless of the size of the detection window or when they are observed within a sample interval.

For each sample interval and each device, we derive the total anomaly score for device i as:

$$a_{T,i} = a_{1,i} + a_{2,i} + a_{3,i} \qquad (4)$$

If the current sample interval is denoted as sample interval j, and $a_{T,i}$ above is shorthand for $a_{T,i}^{(j)}$, then the total anomaly score for a window of width w ending at sample interval j is:

$$A_i = \sum_{k=0}^{w-1} a_{T,i}^{(j-k)} \qquad (5)$$

An alarm is generated when $A_i \geq r$ for any device i. With respect to factor a_1 and a_2 (or a combination thereof), the configurable alert threshold (r) for A_i is the minimum number of anomalous scans that must be made by device i within the detection window before an alarm is generated. As usual, a balance must be struck between incurring false positives and negatives (refer to Section 4.3). The lower the alert threshold, the more sensitive to ARP fluctuations the detection system becomes and the greater the possibility of a false positive. r can be manually set before the training period or automatically determined by the detection system. Our prototype automatically sets r to the floor of the highest E_i value it has calculated over the training period. For instance (cf. Table 1), 192.168.1.11 has the highest E_i score (3.65214) therefore $r = 3$. We discuss how the three anomaly factors interact next to produce A_i.

The first two factors (a_1, a_2) reflect the number of anomalous scans observed within the detection window. An anomalous scan for device i is considered to be any of: (1) each ARP request outside of a device's ARP chain (a_1 factor); (2) each scan by which O_i is in excess of E_i (a_2 factor); or (3) the detection of a single scan to internal network dark space (a_3 factor). The third factor is unique in that per our parameter settings, an alarm was triggered after a single observation of a scan to internal network dark space regardless of the value of r and the time of observation. For instance, if $r = 2$ an alert will be generated after detection of two anomalous scans within the detection window (i.e. $A_{1,i} + A_{2,i} = r$) or the single observation of a scan to internal network dark space (i.e. $A_{3,i} = y = r$; see Section 2.1.3).

3 High-Level System Design

Our software implementation uses the *libpcap* [3] library and is comprised of two logical components: PPE and ACE. The Packet Processing Engine (PPE) is responsible for extracting the relevant features from the live network activity or saved network trace files (see Section
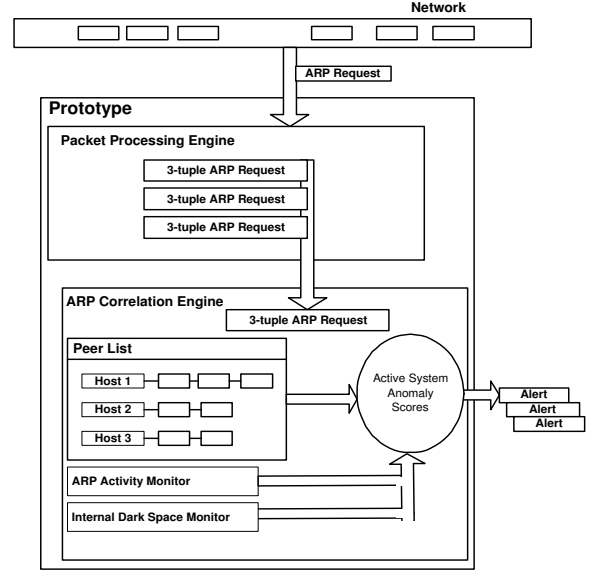

Figure 3. High Level Design of Prototype Implementation

3.1). The ARP Correlation Engine (ACE) includes a dynamically generated peer list and the list of IP addresses it considers to be internal network dark space (see Section 3.2). The ACE maintains in state all relevant ARP information extracted by the PPE. To detect worm propagation between network cells, an instance of such a prototype would need to be deployed in each broadcast domain. Figure 3 shows the high-level design.

3.1 Packet Processing Engine (PPE)

The PPE is responsible for extracting all ARP request packets from network capture files or live off the network. Due to the transmission mechanism of ARP requests (i.e. broadcast), the prototype can be deployed on any device within the network cell. Other forms of ARP activity (e.g. ARP replies) are ignored making this scheme stateless. Feature extraction from the ARP request packets includes 3-tuple tokens (source IP address, target IP address, timestamp) which are passed to the ACE for processing.

3.2 ARP Correlation Engine (ACE)

The ACE processes all network features passed to it by the PPE. The ACE is responsible for four major functions. During the training period, the ACE: (1) creates individual-device specific ARP request statistics, and (2) creates the peer list. Once the training period is complete the ACE: (3) uses ARP request activity to generate a three-factor anomaly score, and (4) generates alarms when the alert threshold has been met or exceeded.

Network ARP Statistic Extraction. During the training period, all tokens passed by the PPE are processed to determine system specific ARP request activity. The ACE maintains ARP request statistics for each active device (i) within the network. ARP requests, encapsulated in tokens from the PPE, are processed in sampling intervals of duration t. For our implementation, we chose a value of t such that it matched the default ARP cache time to live (ttl) of our devices (i.e. Linux operating systems). If within a sampling interval there is no ARP request activity, this observation is excluded from final mean and standard deviation calculations (see Section 2.1.2). This is to compensate for frequent periods of inactivity (e.g. nights and weekends) that would skew the ARP request statistics giving them lower values than in peak usage times.

At the end of the training period, the mean and standard deviation of ARP request activity is calculated for each device (see for example Table 1). These values comprise the expected maximum (E_i) ARP request activity for each individual device within the cell.

Peer List. The peer list, constructed during the training period, contains a listing of all *live* devices and the IP addresses of the internal devices they were in communication with. In a typical network environment (i.e. client server model), internal devices will try to access only a subset of the devices within their local subnet, as recorded by their respective ARP chains. For any given device, connecting to a device within its respective ARP chain should occur at a higher probability than other devices within the peer list. If a worm performs unrestricted scanning[6] of a network cell, anomalous ARP activity will result. First, ARP requests between system pairs (i.e. infected host and a subset of victims) that have not been observed communicating during the training period will occur. Second, as an infected host tries to connect to internal devices the number of ARP requests it issues per unit time will increase.

Anomaly Score and Generating Alerts. Once the training period is complete, an anomaly score for each individual device within the network cell is maintained (see Section 2.1). An alert is generated when $A_i \geq r$ for any device i. The timestamp from the triggering ARP request is used as the timestamp for the alert, which also indicates the alert triggering source and destination address.

4 Prototype and Analysis

In this section, we describe the network and data set (network traffic) we used with our software prototype as a proof-of-concept to validate our proposal and refine our system design, and discuss how our prototype performed. Four weeks of network traffic was collected in one of our university research labs. The first two weeks of the net-

---

[6]Recall that our definition of a scan is an ARP request.

---

work data set was used as the training period. We then tested the prototype on the remaining two weeks of data, to determine both the validity of our detection technique, and the affect of the configured alert thresholds on false positive rates. We tested two different approaches to setting alert thresholds:

1. Common threshold approach (Section 4.1.1): give every device within the network cell the same alert threshold r.
2. Function-specific threshold approach (Section 4.1.2): partition the network cell to give devices that perform different functions (e.g. server, workstation) different alert thresholds r_j where j is the function used to partition the network cell.

Finally, we will describe our scanning worm simulations and report on the performance of our detection software in detecting these scans.

4.1 Data Set for Prototype Evaluation

To validate our approach, we developed and tested a fully functional software prototype with all features discussed in Sections 3.1 and 3.2. The software was installed on a commodity PC running Linux with a 10/100 network interface card. The lab network consisted of a one quarter Class C network of Internet-reachable IPv4 addresses (i.e. 63). Using the cell definition from Section 2, the lab contained one cell. The lab analysis allowed us to test the prototype's ability to detect L2L intra-cell worm propagation. Network traffic was collected from November 11 to December 11, 2005.

From the two week training period the prototype automatically determined each device's peer list size, mean number of ARP requests per minute, standard deviation of ARP requests per minute, and the largest number of ARP requests observed by each device within the sampling interval (we used a 60 second period). The last characteristic is not used in the determination of the anomaly score but as an input to analyze the effectiveness of the approach as discussed in Section 4.3.1. After the training period we recorded, for analysis, ARP activity within the network cell in a single *pcap* file for the next two weeks. During this analysis period, we monitored the internal network independently with an intrusion detection system (i.e. snort [16]) to ensure no worm activity was included within the data set. Finally, we simulated scanning worm propagation within the test network using the *Nmap* [7] security scanner to test our detection software.

The respective ARP request activity for each active system within the internal network is included in Table 1. Network infrastructure devices (i.e. firewall and switches) were excluded from analysis and thus do not appear as index entries within the peer list. The total size of ARP

Table 1. ARP Statistics for Prototype System on Lab Network. Mean and Standard Deviation for number of ARP requests made per sample interval of 60 seconds during the training period. Max ARP Requests is the maximum observed ARP requests over any sample interval during the training period.

Servers					
IP Address	ARP Requests	ARP Chain Size	Mean	Standard Deviation	Max ARP Requests
192.168.1.11	13 532	21	1.57974	1.03620	8
192.168.1.12	13 335	17	1.20305	0.61021	9
192.168.1.13	11 318	16	1.17641	0.42872	4
192.168.1.14	2 701	2	1.01310	0.11376	2
Workstations					
IP Address	ARP Requests	ARP Chain Size	Mean	Standard Deviation	Max ARP Requests
192.168.1.15	10 216	8	1.29761	0.51830	4
192.168.1.16	8 254	10	1.24309	0.46756	4
192.168.1.20	9 080	9	1.21419	0.46745	5
192.168.1.21	3 335	8	1.17489	0.47421	5
192.168.1.22	6 016	9	1.21649	0.44508	3
192.168.1.23	5 513	10	1.20426	0.44432	4
192.168.1.24	6 863	8	1.17125	0.42354	4
192.168.1.25	5 546	9	1.20074	0.43677	5
192.168.1.26	5 642	11	1.06825	0.26134	3
192.168.1.27	2 311	9	1.19744	0.46979	5
192.168.1.30	7 708	5	1.72427	0.67429	4
192.168.1.31	5 746	4	1.52241	0.63850	3
192.168.1.33	17 782	6	1.23520	0.44930	3
192.168.1.45	57	4	1.07547	0.26667	2
192.168.1.51	99	4	1.42307	0.75006	4
192.168.1.52	10	5	1.45454	0.93419	4

request traffic for the entire four week period was approximately 9 MB.

4.1.1 Approach 1: Common Threshold

Table 1 separates servers and workstations in our testbed. Note that, with the exception of one server (a secure log server), the servers within the network have the largest ARP chains. The device with the largest peer list (i.e. 192.168.1.11) was the DNS/mail server for the network. This is not unexpected in a typical client-server model. Likewise, the servers within the network also had the largest observed ARP requests within the sampling intervals.

By applying our technique on all the devices within the network, we determined the number of false alarms generated as a function of our alert threshold. Applying a common alert threshold to all devices (*common threshold approach*), we ran the prototype on the second two weeks of archived ARP request data, varying this threshold to observe the affect on false positive rates. Each trial *run* of the prototype (i.e. processing a two week data file) took less than one minute to complete. A subset of our results are captured in Table 2.

Setting $r = 1$ resulted in 99 false positives over the two week dataset. Recall that a scan to an IP address considered to be internal network dark space was set to immediately generate an alarm regardless of r by our suggested configuration of $y = r$. As expected, we observed no scans from internal devices to internal network dark spaces.

Setting the alert threshold at $r = 2$ causes an alarm to be generated after observing two anomalous scans within two time intervals (i.e. 2 minutes). With $r = 2$, 22 false positives resulted over the two-week period. For contrast, for $r = 3$ (the value automatically selected by the prototype, see Section 2.2), only 5 false positives resulted over the two-week period. For use in an automated response system, the occurrence of 5 false positives within a two-week period may be too great. In our test network, as we

increased r, the number of false positives decreased. If we manually set $r = 6$ only 1 false positive from all devices is generated within a two-week period.

4.1.2 Approach 2: Function-Specific Thresholds

One method to refine our approach is to use different alert thresholds, for different categories of devices based on the system function (not currently implemented by our prototype). For instance, we observed that most servers (with one exception) have a higher ARP chain counts than workstations. Additionally, two servers have the two highest observed per minute ARP request counts during the training period (servers 192.168.1.11 and 192.168.1.12, see Table 1). Not surprisingly, a server must be able to handle bursts of requests from other network devices. However, legitimate bursts in ARP request activity may cause the a_2 factor to exceed its alert threshold causing false positives. All 5 false positives generated at the $r = 3$ threshold were caused by the two servers with the highest per-minute ARP request count. We refer to distinguishing of devices within the network cell to allow differing alert thresholds based on the function of a device (e.g. server or workstation) as the *function-specific thresholds* approach. We could allow the workstations alarm threshold to be set at $r = 3$ reducing their false positive rate to zero (i.e. for our test network). Increasing the server alarm threshold to $r = 5$ reduced not only their false positive rate but the overall false positive rate for the two-week period to 2.

Varying alarm thresholds could be extended to other classes of systems within the network cell if required, depending on the nature of applications running on the network. An example would be relaxing r on devices that use distributed or P2P applications. We expect that these devices would repeatedly initiate a greater number of connections (ARP requests) than other devices during the sampling intervals. This increased activity would typically give these devices larger mean and standard deviation scores causing higher E_i values. Devices with larger E_i values will require the observation of more ARP requests than less active devices (i.e. lower E_i) before r will be met or exceeded. In this scenario, the ability to assign differing r values to devices with similar E_i would minimize both false positives and negatives. However, as r increases so does the number of worm scans before an alarm is generated.

4.2 Simulating Scanning Worm Activity

To simulate scanning worm propagation within a network cell we used the port scan option of the Nmap security scanner. Just like a worm, the kernel and networking components of the workstation performing Nmap scans use ARP in order to make contact with the devices within

Table 2. Alarm Threshold Analysis

Alerts			
Threshold r	Server	Workstation	Total
1 scan	37	62	99
2 scans	19	3	22
3 scans	5	0	5
4 scans	3	0	3
5 scans	2	0	2
6 scans	1	0	1
Anomalous Connection Activity			
	Server	Workstation	Total
Outside ARP chain	181	36	217
Dark space	0	0	0

the network cell. We configured Nmap to scan a single port (port 80) on all the devices within the network cell. Ignoring the two broadcast addresses left 61 usable IP addresses. These port scans simulated a scanning worm trying to find versions of vulnerable HTTP servers within the network cell. The device we used to scan the network cell was a workstation that had the highest ARP chain count (192.168.1.26 had 11 ARP chain entries; see Table 1). We set our alert threshold to $r = 3$. Based on Section 4.1.2, this was the minimum threshold that incurred no false positives from workstations during testing and it was also the value automatically selected by our prototype. To fully exercise our detection software, Nmap was run in two modes, each with two types of scanning strategies as follows.

The first two tests consisted of scanning port 80 on every device within the network cell using Nmap's *normal mode* (scanning the network as quickly as possible) employing both the sequential and random scanning strategies. The last two tests consisted of scanning port 80 on every system within the network cell using *sneaky mode* (waiting 15 seconds between scans in an effort to become stealthy) employing both the sequential and random scanning strategies.

Table 3. Network ARP Statistics

Number of Scans Before Detection		
	Normal	Sneaky
Sequential	2	2
Random	1	3

4.2.1 Worm Simulation Results

Nmap Sequential Scanning Strategy. The sequential scans for both normal and sneaky mode were detected within two port scans (see Table 3). Nmap was config-

ured to sequentially scan the host range from 192.168.1.1 until 192.168.1.62 (thus omitting network broadcast addresses). 192.168.1.2 was the target of the second scan in both sequential scans. IP address 192.168.1.2 was assigned to a network switch and does not appear within the peer list and therefore is considered internal network dark space. In these cases, sequential scanning detection was triggered by the $a_{3,26}$ factor within the aggregate anomaly score for the system.

Nmap Random Scanning Strategy. In normal mode, the random scan was detected within one scan. Of the 61 usable addresses within the network, the total peer list size of servers and workstations was only 21 (i.e. approximately 66% of our network was defined during the training period to be internal network dark space). The first random scan in normal mode was to an internal network dark space. Again, the a_3 factor dominated the aggregate anomaly score and caused an alarm to be generated after the first scan.

In sneaky mode, the random scan was detected within three scans. In this case, although statistically improbable, no internal dark space addresses were scanned. The $a_{2,26}$ factor became the dominant factor and triggered after detecting three ARP requests above E_i for the device within a three minute period. However, it is statistically probable (see Section 4.2.2) that subsequent tests using the same parameters would be detected by the a_3 factor before three scans.

4.2.2 The Affect of Dark Space on Sequential and Random Scanning Detection

Overall, our detection testbed implementation benefited from the sparse internal IP addressing scheme within the network cell. Internal network dark space comprised approximately 66% of the network cell's useable IP addresses. If p is the probability that a random scan will be to internal network dark space then $1 - p$ is the probability that a random scan will not be to internal network dark space (e.g. 0.3387). Random scans are independent events. The probability that the $a_{3,i}$ factor will trigger causing an alarm after the occurrence of three random scans is $1 - (1 - p)^r = 1 - .3387^3 = 0.9611$.

The large amount of internal network dark space also aided our prototype in detecting sequential scanning. Topological worms typically harvest network configuration information from their victims for new targets [18]. In our testbed, any sequential scanning strategy that started from the lowest IP value within a device's network subnet configuration value (i.e. 192.168.1.1) would be detected within the second scan (see Section 4.2.1).

4.3 Discussion of False Positives and Negatives

The following two sections discuss the impact and causes of false positives and negatives on our detection technique. Since the analysis is valid for both the common threshold and function specific thresholds approaches, we discuss only the specific results of the approach with the greatest number of false positives (i.e. common threshold approach).

4.3.1 False Positives

All five false positives which arose when the alert threshold was set to $r = 3$ (triggering at r or more scans) were caused by servers, and specifically by bursts in server activity. A typical scenario for normal network activity involves users logging onto their workstations and requesting network services (e.g. DNS, mail, etc.) that allow them to execute desired tasks. As a workstation generates an ARP request to determine the MAC address of the server, the server also typically generates ARP requests to determine the MAC address of other servers that assist them in performing their tasks. When a number of users access services simultaneously, this will cause a burst in ARP requests from the servers. The two most active servers (i.e. 192.168.1.11 and 192.168.1.12) in our testbed have the highest observed maximum ARP requests in the sampling interval. ARP request bursts caused by servers answering legitimate service requests can produce false positives. The occurrence of false positives could be reduced by raising r. However, each increment of r allows another scanning worm infection attempt to occur before an alarm is raised.

Automated attack and scanning tools share the same searching strategies as scanning worms. These tools can perform random or sequential scanning at differing speeds to either exploit or identify vulnerable systems. Our ARP-based detection technique discovers intra-cell scans caused by such tools (e.g. if the minimum sustained scanning rate exceeds one scan per minute), but does not distinguish them from scans resulting from a scanning worm.

4.3.2 False Negatives

A false negative occurs when malicious activity occurs without triggering an alarm. In this section, we discuss the affect of ARP caches and large ARP chains on possible false negatives for our detection technique.

ARP Cache. If a device happens to have the MAC address of the device it wants to communicate with within its local cache, no ARP request is generated. For this reason, typical scanning worms exploit an infected device's local ARP cache just as any legitimate network application would. However, our prototype is network-based and

does not have access to the local ARP caches of the devices within the network cell. For this reason, with respect to the $a_{2,i}$ factor, this activity is not reflected within the O_i activity count.

However, our calculation of E_i (see equation (1), Section 2.2.1) does offer some insight on the affect of ARP caches. For example, system 192.168.1.30 (statistically the most active workstation in Table 2) has $\overline{x} = 1.72427$ and $\sigma = 0.67429$, making $E_{30} = 3.07285$. This represents the maximum *expected* number of ARP requests a device can make within a one minute period without causing a positive $a_{1,30}$ score. Recall (see Section 3.2) that the default ttl for the entries within the ARP caches of the devices in our network cell is 60 seconds, which matches our $a_{2,i}$ sampling interval. ARP replies (i.e. MAC and IP address pairs) are cached locally on the devices that generated the associated ARP requests.

According to E_i, we expect that three is a reasonable upper bound for the number of entries within this device's ARP cache. If a scanning worm happened to select any of the IP addresses within the cache, no network ARP request would be sent and our prototype would not detect the scan (i.e. O_i would not be incremented by 1). ARP caches can be a source of false negatives. However, in our calculations of $\overline{x}$ and σ for each device, we ignored the affect of long periods of inactivity. This was done to better approximate the ARP request activity during active usage (i.e. ignoring user and system inactivity) to ensure that our false positive rates would be minimized for this factor. E_i represents a reasonable upper bound to the number of entries within the respective local ARP caches. In practice, the ARP caches will typically contain fewer entries than the E_i values for each system since by design, E_i exceeds the corresponding mean value. Therefore, we expect that ARP caches have a minimal affect on factor $a_{2,1}$.

Additionally, the $a_{2,i}$ factor is not applied in isolation. In order to avoid any anomaly score contribution from the remaining two anomaly score factors (i.e. $a_{1,i}$ and $a_{3,i}$) all these scans would have to be limited to the device's ARP chain. The probability of detection avoidance through limiting scans to the respective ARP chain is discussed in the next section.

Large ARP Chain. We use ARP chains to characterize normal network interactions. A system with a large ARP chain will be able to connect to the devices within the chain without contributing to the $a_{1,i}$ factor. The largest ARP chain in our testbed belonged to the server 192.168.1.11 with 21 entries. The largest workstation ARP chain belongs to 192.168.1.26 with 11 entries. Table 4 shows the probabilities of the two systems scanning within their respective ARP chains, internal network dark space, and their peer lists.

As discussed in Section 4.2.1, internal network dark

Table 4. Anomaly Factor Triggering Probabilities in Testbed

System Specific ARP Request Statistics			
	ARP Chain Size	$\overline{x}$	σ
192.168.1.11	11	1.57974	1.03620
192.168.1.26	21	1.06825	0.26134
Scan Location Probability			
	ARP Chain	Dark Space	Peer List
192.168.1.11	0.3387	0.6613	0
192.168.1.26	0.1774	0.6613	0.1613

space dominates the overall scanning possibilities for worms within our testbed network cell. Even the device with the largest ARP chain (192.168.1.11) had only a 0.3387 probability of selecting an IP address within its ARP chain. During a random scan, the chance of 3 successive scans all targeting IP addresses in the ARP chain is approximately $(1-p)^3 = 0.3387^3 = 0.03885$ (see Section 4.2.2). If a sequential scanning strategy was used, the probability that 3 successive scans would all be to devices within an ARP chain would depend on the IP address composition of the ARP chain. Using our network as a practical example, a sequential scan from 192.168.1.11 would deviate from its ARP chain on the second scan (see Section 4.2.1) and be detected by $a_{1,11}$ and $a_{3,11}$. Regardless, if a scanning worm managed to only scan IP addresses within the device's ARP chain, it would have to do so with a sustained scanning rate of less than 1 scan per minute or it would be detected by $a_{2,11}$. Therefore, the application of a_2 ensures large ARP chains have a minimal affect on the detection technique.

5   Limitations

Limitations. Our detection technique cannot detect L2L inter-cell, R2L and L2R worm propagation (see Section 2.1). It would be useful as a complementary technique used in conjunction with other approaches, e.g. [20] to detect L2L intra-cell scanning worm propagation. Furthermore, our approach relies solely on the observation of ARP requests. We do not try to match the associated ARP replies to determine if the subject of the ARP requests are actually active on the network. In the event that a system broadcasts an ARP request to a device currently not active on the network (e.g. powered down or disconnected from the network) due to scheduled maintenance or some unscheduled failure, this will not be considered a scan to internal network dark space as long as its address was observed during the training period. In this scenario,

the amount of internal network dark space would be understated as inactive devices would be considered *live*. If we extended our approach to correlate ARP requests and replies we could determine which devices are live or actually internal network dark space.

ARP request and reply correlation would also address another potential limitation that arises in networks that use the Dynamic Host Configuration Protocol (DHCP) [5]. DHCP allows network devices to determine their IP addresses from a central server rather than from a static configuration file. When a device becomes active on the network, it contacts its DHCP server to retrieve an IP address that it can use on the network. The MAC address of the requesting device is then associated with an IP address assigned by the DHCP server. DHCP-assigned IP addresses are *leased* to the devices that request them. A DHCP lease is the amount of time that the DHCP server grants permission for a device to use a particular IP address. The devices in our test network used static IP addresses and thus the MAC and IP pairing was constant. In a DHCP-enabled network, the MAC address and IP pairing is not guaranteed to be constant (e.g. when the lease expires a device may receive a different IP address and thus the IP address MAC pairing is different). Our prototype uses IP addresses to identify devices and thus would be adversely affected by allowing a device to have different IP addresses. ARP request and reply correlation would enable us to use the MAC addresses (which are fixed and never change) of devices for identification which DHCP has no affect on.

Another limitation is that network dark space addresses are determined by observing ARP requests on the network and building a peer list (see Section 2.1.1). Once the training period is completed, in our description thus far there is no mechanism to add to the peer list or determine if previously active devices have been taken off the network. This may provide an inaccurate accounting of internal network dark space. To address this limitation, we could correlate ARP requests and replies to determine the emergence of new devices. Currently, if we observe an ARP request to an IP address outside the peer list an alarm is generated. If we modified our approach to consider ARP replies, we could dynamically determine internal network dark space.

In a P2P or distributed computing environment, network devices may interact with a large number of other devices. The ARP chains for the devices could be quite large and homogeneous. In this scenario, the a_1 factor would be affected as a device could interact with a large percentage of the network cell and still remain within its ARP chain. Furthermore, if the device was involved in performing tasks that required frequent interaction with multiple devices over long periods of time its $\overline{x}$ and σ would be large. This would require our prototype to observe a greater number of worm scans before the a_2 factor would trigger.

Attempted Circumvention. A possible worm infection strategy would be to only perform infection attempts after the device had initiated a connection through legitimate use. In this scenario, a scanning attempt could be initiated when the IP address is in the local cache thus obviating the need for an ARP request. A slight modification to this strategy would be for the worm to install itself on a host and monitor ARP request activity before propagation. Worm propagation could then be restricted to those devices that were the subject of ARP requests. In these scenarios, the a_1 and a_3 factors would not be affected by this activity. However, once propagation begins the worm's sustained scanning rate would still trigger an alert by a_2 unless the scanning rate was less than 1 scan per minute.

6 Suppression and Containment Extensions

To suppress and contain scanning worms within a network cell apparently requires that one employ an automated active response. Two scanning worm containment strategies are: (1) stop the scans as they traverse the network before they reach their intended victims; and (2) stop the scans before they leave the infected host. The first strategy could be achieved through integrating a containment capability directly into the network fabric (e.g. switches). The second strategy could be achieved through integrating a containment capability on each host itself. We believe that our ARP-based detection technique could be used in either strategy to enable automatic containment. Each strategy has its benefits and limitations; we discuss these in turn.

6.1 Integration into Network Fabric

A network switch provides a channel for incoming data from any one of its input ports to a specific output port that connects a device to the network. Most network switches have the capability to maintain statistics on the type and amount of network data passing through its ports. We believe our ARP-based technique could be integrated into switches to allow for monitoring of ARP activity on all the devices they connect to the network. Switch integration would enable us to address a current limitation of our approach by allowing ARP requests and replies to be correlated (see Section 5). ARP replies are not broadcast and must be observed at each individual device at which they are received or on the switches that forward them. Integration of our detection technique into the switching fabric would enable us to perform a more complex correlation for better detection and better provide the ability to take active response. For instance, if worm activity was detected on a specific port, a switch could simply turn the

port off isolating the device connected to that port from the rest of the network.

A network-based detection approach such as this does have some advantages over a host-based approach. A network-based device can be hardened against attack and is typically administered by security conscious personnel. In a host-based detection approach, the host is the subject of the attack and there is a risk that a successful attack may include turning detection software off or instructing it to not alarm. Furthermore, administration and reporting of alarms is typically easier in network-based detection approaches as fewer devices need to be maintained and monitored.

6.2 Host-based Worm Containment

The success or failure of a worm infection is more easily determined at the host than at the network layer. In a network-based approach, one may see an attack go by without knowing if it succeeded, unless one observes subsequent anomalous activity (i.e. attacks) from the suspected victim. Our ARP-based technique could be integrated into a host's networking software or hardware to enable containment. As in Section 6.1, a host-based integration would also enable us to address a current limitation of our approach by allowing us to correlate ARP requests and ARP replies (see Section 5). ARP replies received at the individual host would enable us to perform a more complex ARP correlation for better detection accuracy. Again, as in Section 6.1, a host-based approach would also give us the capability of an active response on the intended victim. For instance, if worm activity is detected on the individual host, the host could remove itself from the network by not forwarding or receiving data.

7 Related Work

The ARP protocol has been the subject of so-called *ARP cache poisoning* attacks [10], in which a device's ARP cache is updated with forged ARP request and reply packets in an effort to change the MAC address of the device to one in which an attacker can monitor or perform a denial of service attack. Other than the involvement of ARP, our work is unrelated to ARP cache poisoning.

A host-based solution to the scanning worm problem was discussed by Williamson [21]. His approach is to limit or throttle the rate of malicious code at the host by determining if a system is trying to connect to new addresses. If so, the connection is delayed in order to delay malicious code propagation. The decision to delay and not drop suspicious connections reduces the impact of false alarms while still limiting the spread of malicious activity in the network. Granger et al. [8] developed a software architecture to enable self-securing network interfaces. Packets are examined as they are processed by

host software where malicious activity can be detected and potentially blocked.

A variety of network-based solutions to the scanning problem exist including Silicon Defense's *CounterMalice* [4] which is a worm defense solution that proactively identifies and automatically blocks worm activity within a network. The solution partitions the network in cells and prevents worms from spreading between the cells.

Jung et al. [9] developed an algorithm called Threshold Random Walk (TRW), to identify malicious remote hosts, based on the observation that scanners are more likely to attempt to access hosts and services that do not exist than legitimate remote hosts. If the connection is determined to succeed the random walk is driven upward; failure to succeed drives the random walk down. By giving legitimate network traffic a higher probability to succeed than attack traffic, a determination can be made on whether a series of connection attempts is a scan.

Weaver et al. [19] developed a scan detection and suppression algorithm based on a simplification of TRW. They use caches to track the activity of both new connections and IP addresses to reduce the random walk calculation, making the algorithm suitable for implementation in both hardware and software. Their technique allows a scanning host to be detected and stopped in fewer than 10 scans with a low false positive rate.

Zou et al. [22] model requirements for the dynamic quarantine of infected hosts. They demonstrate that epidemic thresholds exist for differing detection and response times. This work provides a benchmark that any new proposed detection algorithm could use to measure its efficiency.

Singh et al. [17] propose a method to automatically detect worm propagation based on network traffic characteristics such as highly repetitive packet content, and IP source and destination address distributions. They have shown that this method has the ability, with no human intervention, to generate worm content signatures with a low false positive rate.

Two commercial scanning detectors offer an alternate approach to detect worm propagation. Forescout [1] and Mirage [2] networks use a technique called *dark-address detection*. These detectors either have knowledge of, or route unoccupied address spaces within an internal network and detect when systems attempt to connect to these unused spaces.

8 Concluding Remarks

The main objectives of this paper have been to describe our ARP-based anomaly detection approach, and confirm that the general idea works. To do so, we implemented the approach in a prototype and tested it in a small lab environment (nonetheless with a reasonable amount of normal

network traffic). Obviously tests in larger and more diverse networks are required to fully exercise the approach; however, we believe our testbed was sufficient in establishing the practical viability of this approach, given that it is designed to be deployed in a single *cell*. Furthermore, regarding the scalability of the approach, our prototype has been used to successfully conduct preliminary tests on a much larger cell of 591 network devices giving us some reassurance. While we made specific recommendations on the anomaly score parameters and settings, these were based on our network topology. Different network environments will most likely require different parameter settings that can be determined during the training period.

The minimum sustained scanning rate constraint of one scan per minute was a limitation of our prototype and not the overall approach. The same approach could be used to detect worms that scan beneath this threshold at the expense of more memory and poorer results in terms of accuracy and false positives. The amount of internal network dark space in our testbed significantly influenced (positively) the performance of the detection technique. In the full version of this paper, a more rigorous analysis of the affect of internal dark space will be undertaken.

Our detection system is anomaly-based and therefore has the ability to detect emerging worms. The prototype automatically calculates the required individual device statistics and can determine an appropriate network-specific alert threshold (r). We have developed a full implementation of our ARP-based approach in a software prototype that runs on commodity hardware. We plan to make this software available to the public.

Acknowledgements

References

[1] Forescout. Wormscout. http://www.forescout.com/wormscout.html.

[2] Mirage Networks. http://www.miragenetworks.com.

[3] tcpdump/libpcap public repository. http://www.tcpdump.org.

[4] Silicon Defense. Worm containment in the internal network. Technical report, 2003.

[5] R. Droms. Dynamic Host Resolution Protocol. RFC, March 1997. http://www.ietf.org/rfc/rfc2131.txt?number=2131; accessed on January 24, 2005.

[6] D. Ellis, J. Aiken, K. Attwood, and S. Tenaglia. A behavioral approach to worm detection. In *Proceedings of The Workshop on Rapid Malcode*, 2003.

[7] Fyodor. Remote OS detection via TCP/IP stack fingerprinting. *Phrack*, 54, December 1998.

[8] G. Granger, G. Economou, and S. Bielski. Self-securing network interfaces: What, why and how. Technical report, Carnegie Mellon University, CMU-CS-02-144, May 2002.

[9] J. Jung, V. Paxson, A. Berger, and H. Balakrishman. Fast portscan detection using sequential hypothesis testing. In *IEEE Symposium on Security and Privacy*, 2004.

[10] S. Manwani. ARP cache poisoning prevention and detection. Technical report, Faculty of Computer Science, San Jose State University, December 2003.

[11] D. Moore, V. Paxson, S. Savage, C. Shannon, S. Staniford, and N. Weaver. Inside the Slammer Worm. In *IEEE Security and Privacy Magazine*, pages 33–39, July/August 2003.

[12] D. Moore, C. Shannon, G. Voelker, and S. Savage. Internet quarantine: Requirements for containing self-propagating code. In *Proc. of the 2003 IEEE Infocom Conference, San Francisco, CA*, April 2003.

[13] R. Pethia. Attacks on the Internet 2003. *Congressional Testimony, Subcommittee on Telecommunications and the Internet*, November 2003.

[14] D. Plummer. An Ethernet Address Resolution Protocol. RFC, November 1982. http://www.ietf.org/rfc/rfc0826.txt?number= 826; accessed on January 24, 2005.

[15] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, and E. Lear. Address Allocation for Private Internets. http://www.ietf.org/rfc/rfc1918.txt; accessed on Jan 12, 2005.

[16] M. Roesch. Snort - lightweight intrusion detection for networks. In *LISA*, 1999.

[17] S. Singh, C. Estan, G. Varghese, and S. Savage. Automated worm fingerprinting. In *ACM/USENIX Symposium on Operating System Design and Implementation*, Dec. 2004.

[18] N. Weaver, V. Paxson, S. Staniford, and R. Cunningham. A taxonomy of computer worms. In *The First ACM Workshop on Rapid Malcode*, Oct. 2003.

[19] N. Weaver, S. Staniford, and V. Paxson. Very fast containment of scanning worms. In *Proceedings of the 13th USENIX Security Symposium*, 2004.

[20] D. Whyte, E. Kranakis, and P. van Oorschot. DNS-based detection of scanning worms in an enterprise network. In *Proc. of the 12th Annual Network and Distributed System Security Symposium*, Feb. 2005.

[21] M. Williamson. Throttling viruses: Restricting propagation to defeat malicious mobile code. In *Annual Computer Security Applications Conference*, 2002.

[22] C. Zou, L. Gao, W. Gong, and D. Towsley. Monitoring and early warning for Internet worms. In *Proceedings of the 10th ACM Conference on Computer and Communications Security*, 2003.