

Identifying and Tracking Emerging and Subversive Worms Using Distributed Intrusion Detection Systems

Nathan Einwechter 2002-10-16

Worms continually become more sophisticated, as new propagation methods and stealth techniques are developed and implemented. As worms continue to evolve, so must our ability to detect and track them. One solution is the use of distributed intrusion detection systems (dIDS) to identify new and emerging worms that utilize new subversive propagation techniques. This paper will discuss how and why the dIDS design is able to identify, detect, and track worms even as they implement more advanced propagation methods.

Defining Emerging and Subversive Worms

To understand the solution, we must first understand the problem by defining two terms. For the purpose of this paper, we will define emerging worms as those that are previously unknown, or those that are beginning to emerge as a high-priority threat due to an accelerated or high rate of infection. Subversive worms will be further defined as those worms that are typically more difficult to detect or identify, particularly as they use new and or unusual (which is to say previously unseen) techniques to propagate and communicate. Thus, a subversive worm would be any worm that utilizes covert channels or stealth scanning techniques to disguise its activities.

So, according to these definitions, at time of its discovery, [Code Red](#) would be an emergent but not a subversive worm because it was not particularly difficult to identify or detect. Emergent worms that are not subversive are usually easily picked up by traditional methods, such as anti-virus programs. Emergent worms that ARE subversive and that have not been previously identified are not likely to be detected by traditional methods; wherein lies the problem that dIDS is able to solve. An example of such a worm is the [W32.Leave.Worm](#), which will be discussed in detail and used for the case study in this paper.

Overview of dIDS

To understand how distributed intrusion detection systems (dIDS) are able to assist in the identification and tracking of emerging and subversive worms, it will be helpful to understand what a dIDS is and how it works.

Taken from [An Introduction to Distributed Intrusion Detection Systems](#);

"A distributed IDS (dIDS) consists of multiple intrusion detection systems (IDS) [or any other network sensor] over a large network, all of which communicate with each other, or with a central

server that facilitates advanced network monitoring, incident analysis, and instant attack data. By having these co-operative agents distributed across a network, incident analysts, network operations, and security personnel are able to get a broader view of what is occurring on their network as a whole.

A dIDS also allows a company to efficiently manage its incident analysis resources by centralizing its attack records and by giving the analyst a quick and easy way to spot new trends and patterns and to identify threats to the network across multiple network segments."

There are several examples of commercially available dIDS. One implementation example of the dIDS system is the [Symantec Threat Management System \(TMS\)](#), which gathers data from a large number of sensors (IDS systems, firewalls, etc.) This system is a perfect example of how a dIDS works in that it collects the information from a wide variety of distributed geographic and virtual locations. It can almost be imagined as a spider web or series of spider webs, where the centre of the web is the server that collects the data and allows further analysis, and each point working outward toward the edges of the web represent a single IDS or set of IDSs. dIDSs then often take this data and perform further basic and predictive analysis on the data, as is the case with TMS.

Detection/Identification of Worms Using dIDS

Now that it is clear how distributed intrusion detection systems (dIDS) work, we can examine the ways in which they are ideally suited to identifying emerging and subversive worms. We can also discuss, then, exactly how such worms can be identified, and why they would be difficult to identify by other conventional IDS means such as the more common stand-alone type of IDS. Examples of this will be demonstrated.

dIDS systems are unique in that they provide a broad scope of what is occurring on a large number of computers or devices across multiple segments and geographical areas. This provides a perfect system to allow for the detection of emerging and subversive worms and attack methods by giving security analysts more information from a diversity of geographically dispersed sources that can be correlated and viewed than can a stand-alone IDS. Essentially, the detection of one or two events in the logs of a single stand-alone IDS doesn't show scanning patterns, whereas the correlation of events across multiple IDS does. For example, if a dIDS system is not in place and several IDS sensors detected a specific type of attack or probe, analysts reading the attack logs at each IDS sensor would most likely dismiss the event as "normal attack activity". If, however, these attack logs were aggregated and analyzed, the analyst would be more likely to perceive any unusual attack patterns or trends. The analyst would then be able to correlate the attack data and demonstrate that a specific pattern of events is occurring across a large network area.

Once this pattern of events is identified, the analyst would then be able to identify characteristics of the attacks/probes and thus begin a well-informed investigation, which may ultimately lead to the discovery of a worm. The correlating data provides evidence that the perception of small attack patterns is not mere coincidence: it backs up the data to indicate that there may be something more dubious than what be concluded from a single, small set of logs. It is this ability to correlate data across numerous IDSs that allows the dIDS to be able to detect emerging subversive worms. What may look like relatively normal activity for a given IDS can quickly turn into an obvious pattern of scans and communications to the dIDS due to its broad scope.

Tracking Worms Using dIDS

Once the new worm is detected, the analyst is then faced with the task of analyzing specific characteristics or traits of the worm, recovering the files that the worm consists of, and tracking the worm to determine infection rates and methods. It is also important to track the new worm to discover where the worm has and has not gone so it can be more quickly destroyed.

Due the distributed nature of the dIDS, again, the analyst is able to track the movement of the worm across the Internet or a given network using much the same aggregation and correlation techniques utilized to detect the worm in the first place. These techniques can give the analyst a map of how and where the worm is spreading. Depending on the size of the dIDS, the worm's movements can be tracked across one or many subnets on the Internet. Knowing when and where the attacks occur can allow the analyst to then project infection rates and identify certain characteristics of the worms scan patterns, and its movements. With a large dIDS network, the approximate number of infected hosts can also be determined, as can exactly who the infected hosts are. The dIDS system can also allow the analyst to track the author (or near the author) of the worm by viewing the initial stages of the worm's outbreak. It can also help analysts to detect the author of the worm attempting to connect to infected machines using covert methods.

Case Study: the W32.Leave.Worm

The W32.Leave.Worm scanned large Net blocks at an extremely high rate, looking for hosts that were previously infected with the [Sub7 Trojan, also known as Backdoor.SubSeven](#). Once an infected host was found, the worm files were uploaded to the target host then executed to infect the computer with the worm. The worm then disabled the Sub7 Trojan to prevent further reinfection by the worm as well as prevent any other attackers from gaining access to the system. The infected system then synchronized the host systems clock with the time server at the U.S. Naval Observatory then downloaded an encrypted Web page, which was believed to contain instructions for the worm. Due to the fact it was reinfecting hosts previously infected by Sub7, the worm would have normally been

difficult to detect by conventional means and most likely would have gone undetected until a serious attack was carried out. Once information regarding the worm was discerned and circulated by various members of the security industry, the propagation methods turned into less subversive methods by utilizing e-mail propagation, with a somewhat subversive twist in and of itself.

To further demonstrate the abilities of the dIDS to accomplish the task of identifying and tracking emerging worms, this case study will cover the two major areas of identification/discovery and tracking (or initial analysis) that were discussed earlier in this article.

Identification/Discovery

The W32.Leave.Worm was first noticed within the dIDS system called [myNetWatchman](#) for whom the author of this article volunteered with incident analysis and back tracing. The worm was first noticed because an unusually high number of Sub7 scans began to take place occurring from net blocks owned by AOL Inc. In fact, there were such a high number of attacks originating from the AOL net blocks with varying IPs that it was initially believed to be a RST DoS attack against the AOL networks. This was also believed to be the case because AOL was the only provider to see so many Sub7 scans.

This theory was shattered, however, as we decided to graph Sub7 scans by provider over time to look for any correlations. The first graph we created (Figure 1) using data from the database compared the aggregate number of Sub7 probes for every hour for the previous eight days.

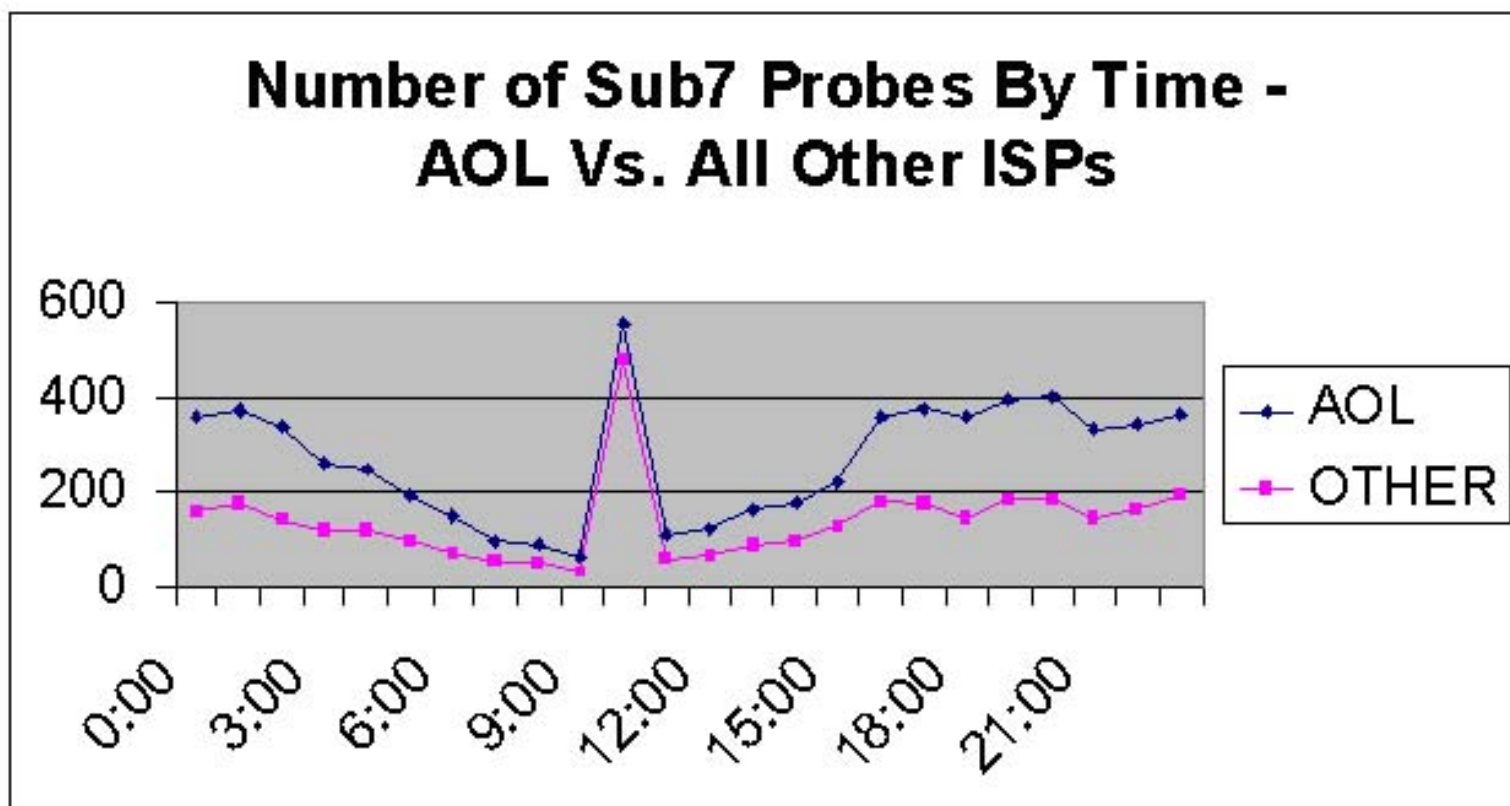


Figure 1

Once the results of this graph were viewed, it was apparent that there was a strong correlation between the probes on AOL versus those originating from all other ISPs in the database. This demonstrated that the event was more widely distributed than originally thought. It also demonstrated a curious correlation across the hours, in which both lines seem to follow the same pattern throughout the hours of the day, suggesting probes are triggered by the hour of the day, perhaps in an attempt to hide the probes among other traffic during busy times of the day.

As a result of Figure 1, another graph was constructed that shows all the ISPs in the database at the time from which Sub7 scans had originated during a period of time. This graph can be seen in Figure 2, and shows the correlation again; only this time the correlation is across days instead of aggregated by the hour over several days. It also shows the worm's evolution as it begins to spread. It can also be seen that the worm was first released into the wild some time on or around June 7.

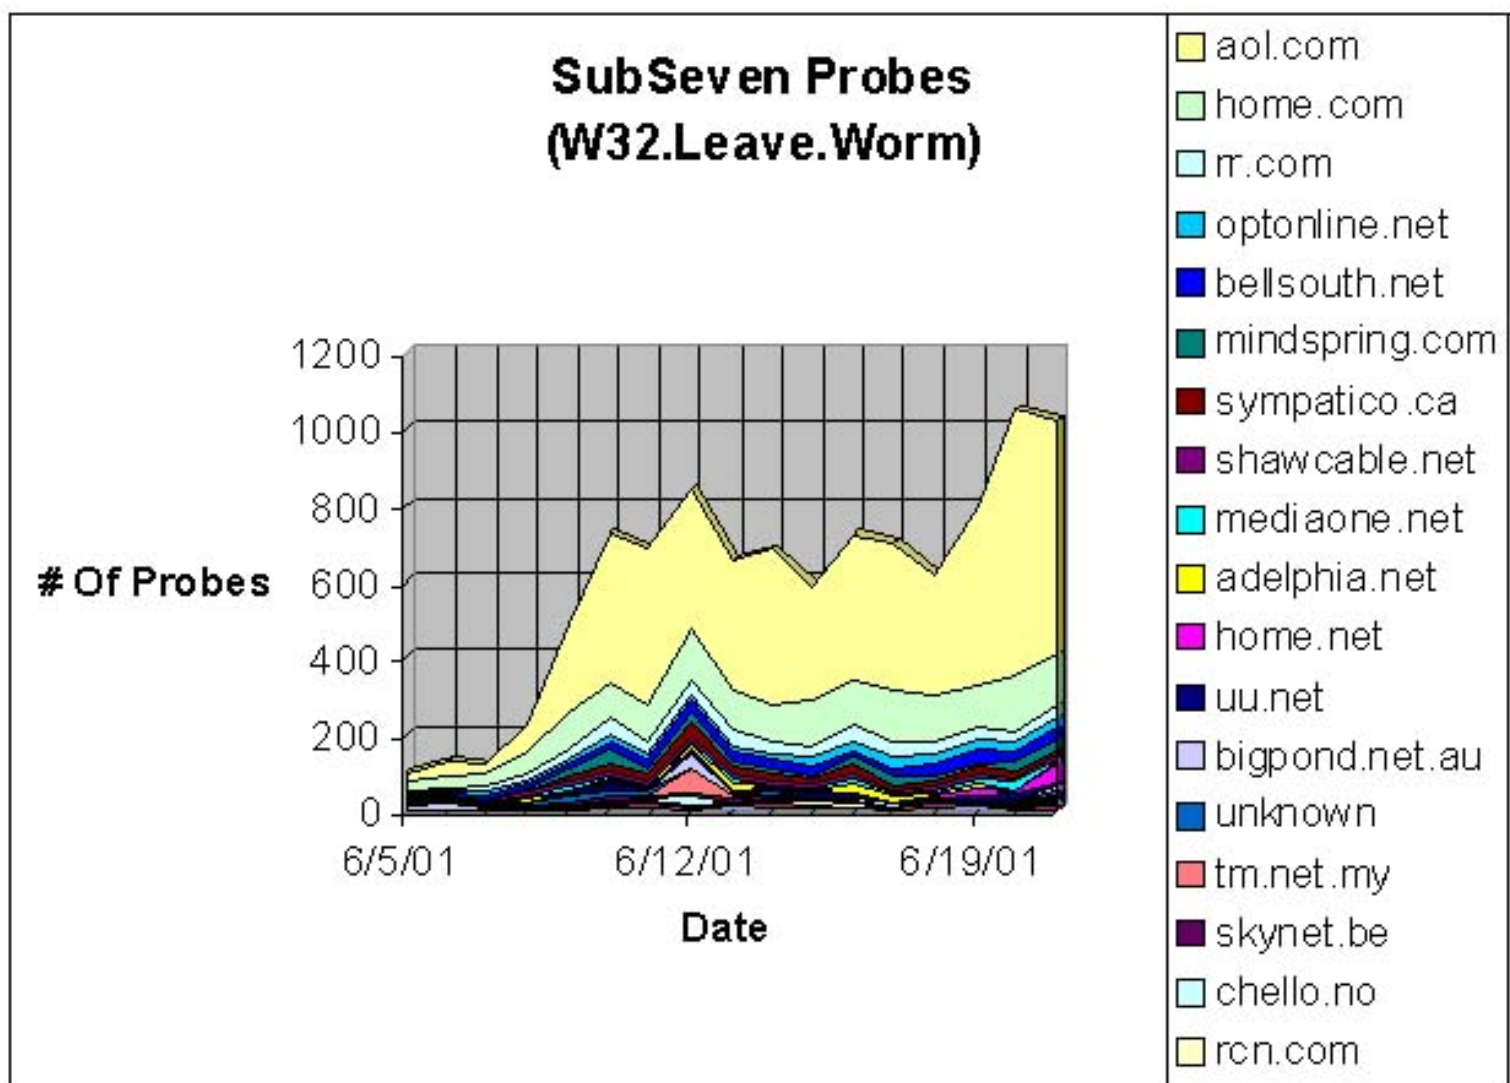


Figure 2

With these two graphs, it was easily discernable that this was not a simple RST DoS attack using spoofed AOL addresses, or even mere co-incidence. It is now highly evident that we are dealing with a more sinister problem such as a worm.

Initial Analysis and Tracking

The data collected and available through the dIDS system then enabled us to continue with initial analysis and do real-time tracking of the worm's progress across the Net. This was possible due to the dIDS sensor network consisting of approximately 1,500 sensors across the Internet and distributed across many net blocks and geographical locations.

Some of the initial analysis afforded to us by the dIDS system included the ability to project infection rates based on scan rates, as well as identifying infected hosts. By identifying infected hosts, we were able to quickly notify the owners of the system or the owners of the net block the system resides on, to prevent the problem from getting out of hand in that way. By allowing us to identify infected hosts, the dIDS also allowed us to keep track of the worm in a new way using the data collected. That is, we were then able to create reports to see what other attack activity had been detected from the infected hosts, to detect any coordinated attacks they may be carrying out, or new propagation methods being used.

Case Study Conclusion

Although the dIDS was not used in the end to identify the author of the worm, it was able to detect and track the worm's progress throughout the Internet. This is particularly important when considering the very quiet infection method it uses. Had such a dIDS system not been in place, the worm could have continued to spread for some time before detection. Eventually, the worm was caught in a honeypot we set up on a network and was then analyzed. Eventually the [author of the attack was arrested](#) in violation of the "Computer Misuse Act of 1990" in the United Kingdom, as a result of a joint investigation.

Conclusion

The dIDS design as a whole is easily able to identify, detect and track new and emerging worms that use subversive techniques to attempt to go undetected for longer periods of time. By using the dIDS system, we are more easily able to prevent major zombie networks from being surreptitiously set up and thus prevent major distributed denial of service (DDoS) attacks from occurring. By early detection and analysis, subversive worms can be dealt with before they become a major problem or

create widespread damage. It is by this early detection that we can protect ourselves, as well as everyone else, on the Internet.

[Nathan Einwechter](#) is currently a Senior Research Scientist with Fate Research Labs, as well as a System Developer/ Incident Analyst with myNetWatchman.Com. While working with myNetWatchman, Nathan assisted in the discovery and analysis of the "W32.Leave.Worm" along with SANS, the FBI, and NIPC. The discovery of which was made possible by analysis of data collected by a dIDS system.

[Privacy Statement](#)

Copyright 2006, SecurityFocus