;login:
THE MAGAZINE OF USENIX & SAGE
December 2002 • volume 27 • number 6

Focus Issue: Security

Guest Editor: Rik Farrow

inside:

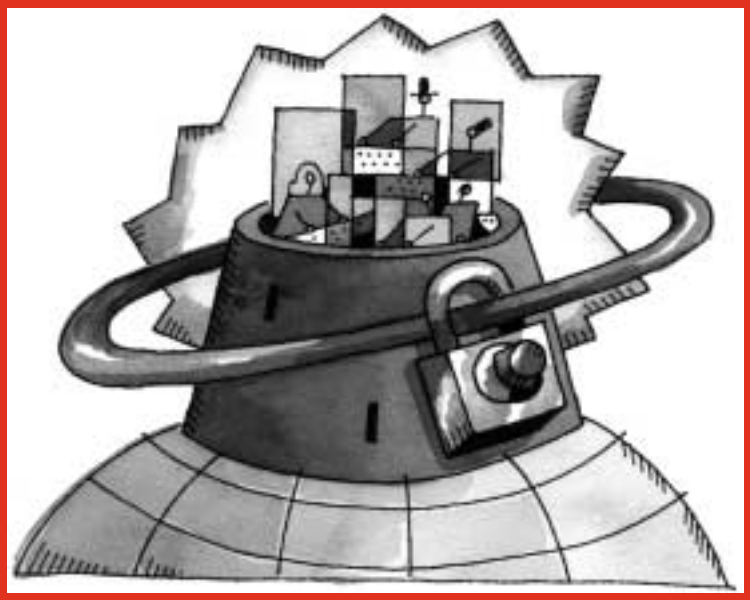

USENIX & SAGE
The Advanced Computing Systems Association &
The System Administrators Guild

HOSUS (honeypot surveillance system)

by Lance Spitzner

Lance is a security geek whose passion is using honeypots to study blackhats. This love for tactics began as a tanker in the US Army's Rapid Deployment Force. He is a senior security architect with Sun Microsystems.

lance@honeynet.org

Within the past several years, the information security community has increasingly recognized the value of honeypots. First discussed in 1989 and 1990 by Clifford Stoll[1] and Bill Cheswick,[2] honeypots are a unique security technology; they are resources designed to be attacked. Many people have different interpretations of what a honeypot is. For the purposes of this paper, I will use the following definition for honeypots: a security resource whose value lies in being probed, attacked, or compromised.[3]

This is a highly flexible definition, but then again, honeypots are a highly flexible technology, able to satisfy a variety of different goals and objectives. Commercial vendors, such as ManTrap, Smoke Detector, or Specter, have developed honeypots that can be used to detect attacks.[4] Other organizations, such as the Honeynet Project, have deployed honeypots for research purposes.[5] This paper attempts to describe one possible deployment of honeypots, called HOSUS (HOneypot SUrveillance System), a concept based on the Navy's SOSUS (SOund SUrveillance System) program.[6]

SOSUS

During the Cold War, one of the greatest threats facing the United States was the Soviet nuclear submarine threat. Submarines could move virtually undetected through the oceans. Used as platforms to launch nuclear attacks or gather intelligence based on intercepted signals, submarines could covertly collect sensitive information on or obliterate a country. One of the greatest assets of a submarine is its ability to operate clandestinely. Nothing is more frightening or more dangerous than an enemy you cannot find or track.

To counter this threat, the United States Navy in the 1950s developed and first deployed the SOSUS program. The intent was to monitor and track enemy submarines, thus neutralizing their greatest advantage and diminishing the threat. SOSUS consists of hydrophones placed along the bottom of various oceans. These hydrophones are linked together to passively capture activity-generated sounds, which are then used to identify, better understand, and track enemy threats.

HOSUS

Much as the United States faced hidden threats in the vastness of the oceans during the Cold War, organizations now face an even greater magnitude of hidden threats in the vastness of cyberspace. Just like the oceans, the Internet is an international domain, where threats come and go, allowing the enemy to strike at a time and target of their choosing. It is extremely difficult to identify and track this threat. HOSUS can provide a solution similar to the one provided by SOSUS. Like hydrophones that passively collect data from the ocean's depths, honeypots deployed throughout the Internet can passively capture attacker activity.

As an information collection and detection technology, honeypots have several advantages. First, they have no real production value, so any activity sent to them is most likely a probe, scan, or attack. This dramatically reduces false positives, one of the greatest challenges faced by most detection technologies. In addition, honeypots can also capture unknown attacks, reducing false negatives, as demonstrated with the Solaris dtspcd exploit captured in the wild in 2002.[7] Last, unlike most detection technologies, honeypots can interact with the attacker, giving more information on the

1. Clifford Stoll, *The Cuckoo's Egg* (New York: Doubleday), 1989.

2. Bill Cheswick, "An Evening with Berferd," *USENIX Proceedings*, January 20, 1990.

3. Lance Spitzner, *Honeypots: Tracking Hackers* (Boston: Addison-Wesley, 2002).

4. ManTrap: *http://www.mantrap.com*; Specter: *http://www.specter.com*; Smoke Detector: *http://palisadesys.com/products/smokedetector/prod_smokedet.shtml*.

5. The Honeynet Project: *http://www.honeynet.org*.

6. SOSUS: *http://www.pmel.noaa.gov/vents/acoustics/sosus.html*.

7. Solaris dtspcd exploit: *http://www.cert.org/advisories/CA-2002-01.html*.

attacker's activities and intent. Examples of this capability can be found in the series of Know Your Enemy white papers and challenges sponsored by the Honeynet Project.[8]

Threats using active measures, such as probes, attacks, or exploits, would be captured by these devices. Once correlated at a central point, this information can give organizations a much better understanding of the threats they face within cyberspace. More importantly, this information can potentially detect activity and predict attacks before they happen. However, HOSUS has the potential of capturing far more information then SOSUS ever could.

A concept similar to this has already been employed, though in a limited fashion. An organization known as the Honeynet Research Alliance,[9] an extension of the Honeynet Project, has passive surveillance devices, known as Honeynets, deployed throughout the world (as of July 1, 2002, they currently have 10 Honeynets). Data is passively collected on threats and attacks, then forwarded to a central point for data correlation and analysis (see Figure 1). This data has proven extremely valuable, resulting in analysis and publication of the characteristics and methodology of many different threats within cyberspace. The HOSUS concept could be employed on a much larger scale.

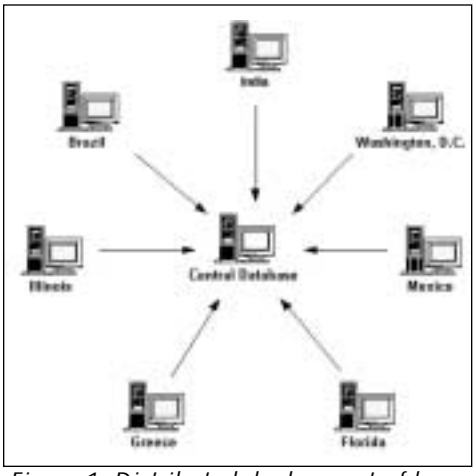

Figure 1: Distributed deployment of honeypots (in this case honeynets) passively collecting and then forwarding data to a central location. Source: Honeynet Research Alliance.

Deployment

There are two approaches to deploying the HOSUS concept: low interaction and high interaction. Honeypots are categorized by the level of interaction they provide to the attacker.[10] The greater the interaction, the more functionality honeypots have. For example, a low-interaction honeypot would emulate a Linux server running the wu-ftp service, limiting the amount of interaction the attacker would have with the system. A high-interaction honeypot would be a real Linux server running a real version of the wu-ftp service; there would be no limitation, since the attacker would have access to a real system. The attacker could exploit the service, take over and reprogram the computer, and then use it as a base for communication. The greater the level of interaction, the more we can learn about the attacker. However, the greater the interaction, the more work involved and the greater the risk the system could be subverted to attack or harm other non-honeypot systems.

Both low- and high-interaction solutions have their advantages with a HOSUS deployment. Low-interaction solutions are much simpler to deploy and maintain. But they are limited primarily to the transactional information of an attack, such as IP addresses, port numbers, and the time/date of the attack. Depending on the level of emulation with the low-interaction solution, some of the attacker's activities, such as login attempts, could be captured. This data can be extremely useful for detection, early warning, and prediction of activity, or statistical analysis of attack behavior.

High-interaction honeypots have the advantage of capturing far greater levels of information. They provide real operating systems and applications for attackers to interact with, just as they exist in the real world. One example of high-interaction honeypots, Honeynets, could be used to capture detailed information on the enemy, including their communications, latest tools and techniques, motives, and organization. Additional measures could be taken to create realistic Honeynets, perhaps even solutions that contain false information designed to mislead attackers. These Honeynets could be customized to appear as different targets, such as a university, government , or hospital site.

8. Know Your Enemy papers: *http://www.honeynet.org/papers/*; challenges: *http://www.honeynet.org/misc/chall.html.*

9. The Honeynet Research Alliance: *http://www.honeynet.org/alliance/*.

10. Spitzner, *Honeypots: Tracking Hackers*.

The quantity of passive listening devices deployed has a direct correlation to the amount of data collected and the value and statistical significance of the data analysis. The more sensors (honeypots) you can deploy, the more data you can obtain. To facilitate this wide deployment, it's possible to create rapidly deployable honeypot sensors. One idea is to create a simple appliance, such as a bootable CD-ROM. The CD-ROM would contain all the required software for the establishment and maintenance of the honeypot. It would be preconfigured to remotely and securely log all captured information to a central collection point. To facilitate ease of deployment, honeypots could also be pre-configured to passively monitor any IP address that is not specifically assigned to a system. This allows for easy and rapid deployment within most organizations. Whenever the honeypot sees activity for unassigned IPs, it simply assumes the identity of the victim, interacts with the attackers, prevents outbound attacks, captures the information, then securely forwards that information to the central data collection point.

The idea of monitoring unused IP space is not new, having been demonstrated by organizations such as CAIDA[11] and Arbor Networks, Inc.[12] However, in addition to monitoring IPs in unused networks, HOSUS monitors unused IPs within production networks of valid organizations as well. And honeypots take this concept one step further by not only monitoring but also interacting with attacks.

For a low-interaction deployment, technology like this already exists, such as the open source solution honeyd,[13] developed by Niels Provos. Honeyd is a low-interaction solution that emulates a variety of operating systems and services. When combined with a technology called arpd, honeyd can dynamically monitor unused IP space, then interact with any activity or attacks bound for those systems. A high-interaction solution would be more difficult to automate the deployment process, but would also have far greater information-gathering capabilities. It may be possible to build a Honeynet solution that also boots off a single CD-ROM, creating the Honeynet architecture, fulfilling data control, data capture, and data collection requirements. Then the Honeynet would only need to be populated with target systems. This process could even be streamlined further by creating virtual Honeynets,[14] multiple systems running off a single physical computer. Similar to honeyd, virtual Honeynets already exist and have been successfully deployed.

Risk

Just like any technology, HOSUS has inherent risks. The greatest risk is identification by the enemy. If the enemy can identify the existence and location of the deployed honeypots, he can neutralize their effectiveness. In the case of the low-interaction honeypots, the attacker can merely avoid the devices, avoiding detection. With high-interaction solutions, the attackers could not only avoid the systems but, if they so chose, feed it bad information, establishing, for example, a false IRC channel with bogus communications. A second risk exists: the honeypots can potentially be compromised and then be used to attack or harm other non-honeypot systems. This risk is especially prevalent with high-interaction honeypots, as we provide actual operating systems for the attackers to interact with.

These risks can be mitigated. By making the deployment of low-interaction honeypots simple and efficient, their identity and location can quickly be changed with minimal impact. Honeypots can be rotated to new locations on a weekly or monthly basis. In cyberspace, unlike the ocean, it is extremely easy to reconfigure and redeploy assets.

11. CAIDA, "Inferring Internet Denial-of-Service Activity," *http://www.caida.org/outreach/papers/2001/BackScatter/*.

12. Arbor Networks, "A Snapshot of Global Worm Activity," *http://research.arbor.net/up_media/up_files/snapshot_worm_activity.pdf*.

13. Honeyd: *http://www.citi.umich.edu/u/provos/honeyd/*.

14. Virtual Honeynets: *http://www.honeynet.org/papers/virtual/*.

This capability also exists with Honeynet technology. Honeynets can mitigate this risk even further by creating a highly realistic environment, running services and applications just as they would be found in a production environment.

For the risk of compromise, measures can be taken to control the attacker. For low-interaction honeypots, the emulated services are created to limit the attacker's interaction. These emulated services are designed not to be compromised; they do not provide enough functionality for the attacker to exploit. At most, they merely appear as vulnerable or exploited services. For high-interaction solutions, data-control mechanisms can be used to control outbound connections, mitigating the risk of the honeypot harming others. One example is Hogwash,[15] a solution that allows hostile egress from the honeypot but alters a bit in the malware to negate the outbound attack. This provides the maximum realism for an intruder inside the honeypot, leading the miscreant to believe the attack tool to be flawed. Other examples of data control include data throttling and counting outbound connections.

Conclusion

The purpose of this paper is to highlight the value of the honeypot deployment concept HOSUS. Similar to the hydrophones deployed during the Cold War, distributed honeypots could be used to passively collect information on threats and hostile activity within cyberspace. Once centrally correlated, this information could then be analyzed to better understand the threats that exist, detect indications of hostile activity, and prevent or, if required, defend against the cyberattack. The types of deployment, low interaction and high interaction, each has its advantages and disadvantages, depending on the data to be captured. Most likely, a successful deployment would require a combination of both technologies. However, both technologies share the same risks: detection and compromise. HOSUS is one possible method to better understand and protect against cyberthreats. If you are interested in learning more about honeypot technologies, *http://www.tracking-hackers.com* is an excellent place to start.

15. Hogwash: *http://hogwash.sourceforge.net*.