1 The Battleground

My commander used to tell me that to defend against the enemy, you have to first know who your enemy is: their methods of attack, tools and tactics, and objective. This military doctrine readily applies to network security just as it did in the Army. The blackhat community is the adversary; we must defend against this threat. However, to be successful, we must first know our enemy.

When I first entered the field of network security, I was frustrated by the lack of information about the blackhat community. It was easy to find technical information on exploits, scanners, and various other attacker tools. But these play only a small part in the overall picture. I wanted to know more. What are the attackers' goals? What are attackers trying to achieve? Why? How do they identify vulnerable systems and then compromise them? What happens once attackers control a system? How do they communicate among themselves? Are we dealing with a single threat or a variety of threats?

Many of these questions were ones we asked in the military; in the military, however, we had answers. Specific organizations, commonly called military intelligence, or S2, were dedicated to obtaining and disseminating information on the enemy. The more we knew about the enemy, the better we could defend ourselves. As a tank officer, for example, I was expected to have an intimate knowledge of Soviet armored tactics and capabilities. I was expected to know the technical makeup of a single Soviet tank company. We were trained on the tank's

range, speed, and performance capabilities. We read books on the history and political structure of our threat. We conducted hands-on training of captured equipment. This information is critical for defending against this threat. By knowing the tank's range, I can estimate when the enemy will begin opening fire on me and thus when I should begin firing back. By knowing the speed of the enemy's tanks, I will know how much lead time I should have when I call in for artillery fire. By knowing the performance of the enemy's tanks—its rate of fire—I can estimate how many rounds the enemy can fire at me in a minute and the probability of hits. By crawling around inside a captured T-72, I had a better understanding of the visual capabilities of the crew inside. All this information is critical in defending against my opponent. The more information I have, the better I can stop and defeat that enemy.

What amazed me in the field of network security was the lack of this type of intelligence. I found little information on who the enemy was, how it attacked, what the motivations or the tactics involved were. The security community was focused on the specific technical tools used by the blackhat community and the tools used in defense but not in the tactics or motives involved. What I wanted to learn was how the blackhat community was identifying and probing for vulnerable systems. What happened once a system was compromised? What activity was going on that I did not know about? I had a lot of questions but very few answers, and this scared me. It was my job to protect against a threat, an enemy. But I did not even know who my enemy was, let alone its tools and techniques. I wanted to learn more, but how could I do so?

It took several years to develop a solution. The plan is simple: Have the enemy teach us its own tools, tactics, and motivations. Why attempt to develop theory when you can have the blackhats show you step-by-step how they operate? No other source is more reliable or more complete. In the military, some consider this battlefield intelligence, whereby you gather information from the enemy. In network security, we can attempt to do the same. We will let the blackhats teach us how they operate. Now the question is, How do you gather battlefield intelligence when you don't even know where the battlefield is?

For me, the battlefield landed in my wife's dining room in 1998. In the beginning of that year, I received my first dedicated connection to the Internet. Anyone in

the world had access to my network at home, at any time. At first, I had no idea what security implications this meant; I did not realize just how aggressive a war is going on in cyberspace. Fortunately, I was researching firewall logs at the time and detected a great deal of suspicious traffic probing my network. I decided to learn more about this traffic, so I researched a variety of papers on the Internet. Although I found a wealth of technical information, most of it focused on specific exploits or the tools used in the exploits. I found little in the way of intelligence on the bad guys. I wanted to learn more but was not sure how. I decided to place a production system on my network, closely monitor this system, and then wait and see what happened. My intent was to have the blackhat community show me how it operates by probing, attacking, and exploiting the system. I used a default installation of Linux Red Hat 5.0, a version of the UNIX operating system, and connected it to the exposed network. I had no idea what to expect. Would anyone even find the system? If so, how long would it take? Would the system be attacked, and what would happen once the system was compromised? All these were questions I was hoping to answer, but would the solution work? On February 25, 1999, I connected the system to my network. Within 15 minutes, my system had been identified, probed, and exploited. Little did I know at the time, but an idea was born.

I learned a lot from that experience, mainly how not to set up such an environment. After compromising the system, the blackhat quickly figured that something was not right, erased the hard drive, and never returned. I lost most of the valuable data that could have been gained, such as the blackhat's keystrokes, toolkits, and system activities. Little was learned, but I had proof that this could be done. By placing production systems on a network and then monitoring all the activity to and from that system, it is possible to learn more about the enemy.

Over time, this concept grew into the Honeynet Project, 30 security professionals dedicated to learning the tools, tactics, and motives of blackhats and sharing those lessons learned. The group learns by building production systems and then monitoring all activity to and from the systems. We capture and analyze data as these systems are probed, attacked, and exploited. Everyone volunteers time and unique skills in the research and development of the project. By combining our skills and knowledge, we can exponentially increase our learning about the blackhat community. We then share this information with the security

community. The end goal is to improve our understanding of the enemy. Armed with this knowledge, we and the security community can better defend against the blackhat community. What makes us unique is that we share as much as possible with the security community; we want everyone to benefit from our research. The more people who understand how the enemy works, the more secure systems will be, which indirectly benefits everyone.

The project began informally in April 1999. I needed help in developing methods to capture blackhat activity. The system compromised in February demonstrated the need to develop more comprehensive and sophisticated methods for data capture. Once the data was captured, I also needed help in analyzing it. I just did not understand a great deal of network and system activity, such as decoding a specific exploit captured from the network. I asked certain members of the community to assist me. Fortunately, the security community is made up of many dedicated and helpful individuals. For example, Marty Roesch, developer of Snort, coded new functionality into the IDS (intrusion detection system) just to help our research—in this case, keystroke logging, called session breakout. Max Vision stepped up to help with sophisticated exploit attacks, decoding exploits based on their network signature. Without the help of these and other members, the project would not have been possible.

Honeynets capture all sorts of unusual network and blackhat activity. No single person can understand all the issues involved. Our small group continued to grow as we realized that we needed the expertise of a larger group of people. Over the next year, the project informally grew as more people were willing to help. Each individual had unique skills, experiences, and backgrounds that contributed to the project. However, we all shared a common motivation: to learn about the blackhat community and to share those lessons learned. We were not a highly organized group; many of us had never met in person. We infrequently shared information via e-mail in attempting to improve the Honeynet concept or to decode a specific signature or attack.

This all dramatically changed in June 2000, when a Solaris 2.6 honeypot was compromised by an organized blackhat group that used our honeypot to communicate among themselves. For a three-week period, we captured all their conversations. Tracking all this activity required the skills of the entire group, from decoding spe-

cific IRC (Internet relay chat) configurations to translating Urdu into English. This event helped galvanize our informal group into an organized project.

We had never even considered ourselves an organized group until then. In fact, the name Honeynet Project was created at the last minute, as we had to call ourselves and our research something when our findings were released. Since then, the group has attracted additional members, such as psychologist Max Kilger, Ph.D., who focuses on blackhat behavior. We have also established relationships with various national and international organizations. We continue to develop our techniques and research, always sharing with the security community our lessons learned. This book represents another step in sharing that information.

The key tool that the team uses is called a Honeynet, a network designed to be compromised. We can then learn who our adversary is and how it operates. Every packet that enters and leaves the Honeynet is captured and analyzed. Every action on the systems is logged and secured. The beauty of the project is that there is no theory. The blackhats show us step-by-step how they operate in the real world. Once we have gathered this information, we can then review the data and better establish who the enemy is and understand its goals, motives, and methods of operation.

Throughout this book, we use the term *blackhat* to represent the enemy, the attacker. Many people have used the term hacker, cracker, or a variety of other labels. We prefer not to get involved in the political debate of what words define which users. We standardize on using the term blackhat to mean the bad guys, the enemy. The enemy can be male or female, a disgruntled company employee, a teenager in South East Asia, or a highly trained former KGB agent. In many cases, you will not know the identity of the enemy. In some cases, we have been able to identify the individual(s) and have noted that here whenever possible. Often, however, the only identity you can assign is the term blackhat. Regardless, this is the individual or entity attempting unauthorized activity with one of your resources.

The common theme throughout this book is learning about our adversary, the blackhat community. In Chapters 2, 3, and 4, we introduce you to the Honeynet, the primary learning tool of the Honeynet Project. We discuss what these

production systems are; their value; how we build, use, and maintain them; and the risks/issues involved. In Chapters 5–8, we cover how we use Honeynets to capture the blackhat activity and then analyze the captured data. Based on this analysis, we are able to learn the tools, tactics, and motives of the blackhat community. Our analysis includes system forensics, packet analysis, and log review. In Chapters 9–12, we review what we have learned about the blackhat community from some well-documented compromises. This will show you step-by-step how the enemy thinks and acts. We attempt to discuss as little theory as possible and focus instead on what we have learned. The end goal of the book is to teach you

- **The Honeynet:** What a Honeynet is, its value to the security community, how a Honeynet works, and the risks and issues involved
- **The Analysis:** How to analyze captured data and from that learn the tools, tactics, and motives of the blackhat community
- **The Enemy:** What we have learned about the blackhat community

We hope that you learn and have as much fun with this book as we have had in the past several years.