

Know Your Enemy: Honeynets in Universities

Deploying a Honeynet at an Academic Institution.

[Honeynet Project](#)

<http://www.honeynet.org>

Last Modified: 26 April, 2004

Honeynets have demonstrated their value as a research tool in the area of Information Assurance (IA). Many researchers and organizations in the security community, both public and private, are currently employing honeynets to continue to gather knowledge concerning the tactics, techniques and procedures of the hacker community. Since the summer of 2002, Honeynet Alliance members at [The Georgia Institute of Technology \(Georgia Tech\)](#), successfully deployed a honeynet on the internal network to collect information on hackers and to help secure their campus enterprise network. The purpose of this paper is to help academic organizations deploy honeynets in .edu environments by sharing with you their experiences and lessons learned. We assume that you have already read and are familiar with the concepts of a honeynet as discussed in the [KYE: Honeynets](#) paper.

The deployment of a honeynet on a large enterprise network such as that found on a major college or university can offer numerous benefits to an institution. Based on our experience, we identified two primary benefits. The first is the ability to use the data collected as a teaching and research tool for any type of computer security related course or research that is being offered. Professors and students can potentially use the honeynet as a testing ground for classes or research. In fact, one student recently received his Ph.D based on our honeynet. The second, and based on our experience the more significant benefit of a honeynet, is it can serve as a network security tool to dramatically increase the overall security posture of that institution's network. For example, our honeynet identified over 165 compromised systems on the GA Tech networks, providing extensive information what was compromised, how, and potentially by whom. Later on in the paper we cover in greater detail the value our honeynet provided GA Tech and its faculty, staff, and students.

Getting That Bad Boy Approved

Lets say you are interested in the idea of deploying a honeynet on your internal, academic networks, for either research or detection purposes. Where does one start? The first step you need to take is getting approval. A common concern for academic honeynets is permission; how do you get institution authorization to deploy this on internal networks? One thing we learned is it is necessary to receive permission from several organizations before establishing a honeynet at a University. This usually involves two things; demonstrating the value of honeynets and addressing issues concerning legality, security and privacy. These areas must be addressed since you will not own the network in which you will be using to establish your honeynet. You must coordinate with your network administrators and your university administration to ensure that issues previously mentioned are addressed. You do not want to set up a honeynet without your network administrators being aware of it.

For value, we initially requested to set up our honeynet for research purposes. Honeynets provide fertile ground for research topics including: databases, distributed agents, data analysis, agent technology, network fundamentals and advanced topics. Additionally, honeynets provide a wealth of data collections for research. However, the honeynet also quickly demonstrated its value in detection, we soon realized its role as part of the university security infrastructure. The value of honeynets extend beyond the research lab and into the system/security administrators toolkit. As detailed later in the paper, the deployment of honeynets at Georgia Tech has significantly increased the cyber-security posture of the university.

The second area to be addressed are the concerns of legality, security and privacy. These are all potential topics for the network administrators and the university administration to use as justification for denying the honeynet deployment and must be addressed up front. It will most likely be necessary to receive permission from the Network Administrators of your campus network in order to establish a honeynet at your institution. Some network administrators may be reluctant to do so fearing that a honeynet will increase the chance of a system compromise occurring. We have found the opposite to be true here at Georgia Tech. The honeynet can provide additional analysts to examine suspicious traffic that is occurring within the campus network. If the principle of data control is followed, the danger posed by a honeynet machine being utilized to compromise any other campus machines is greatly mitigated. Frequent monitoring of the honeynet data and providing frequent reports of suspicious incidents to the network administrators should put them at ease concerning the establishment of a honeynet. In our case, a network administrator (Technical Project Director from the Network Security Branch of the Georgia Tech Office of Information Technology (OIT)) actually took an active role in our honeynet initiatives, co-authoring a recent paper titled, [The use of a Honeynet to Detect Exploited Systems across Large Enterprise Networks](#). This paper is a detailed discussion of the usefulness of honeynets in networks that can be used to help convince your network administrators of the value of a honeynet in helping to secure the campus network. The campus network administrators should provide you with IP addresses located within the production campus network.

In addition to the technical security buy-in needed from the network staff, the university administration will have to approve the deployment of honeynets. The administration will primarily be concerned with issues of legality and privacy. The best practice is to obtain a policy letter that identifies the legal and procedural guidelines on the use of honeynets. One of the most frequently asked legal question is whether a honeynet violates the Wiretap act. This question has been answered by the the United States Department of Justice recognizing that a honeynet can qualify under an Exception to Wiretap Act, the Provider Exception (System Protection) clause. The second legal concern frequently raised is the question of entrapment. Given that honeynets are deployed on networks and no advertisement enticing people to scan and/or break into the honeynet, the case of entrapment has little merit. A final administrative legal concern is the privacy of data on the network. Information such a credit card numbers, user IDs/passwords, social security numbers, and a myriad of other data could easily be captured by a honeynet. Proper handling of this data is very important.

Your network administrators will most likely be aware of issues concerning the monitoring of traffic on the campus network as well as what steps can be legally taken to secure the campus network. If this is not the case, you, or the network administrators, may need to approach your campus legal department. The Georgia Tech campus legal department has decided the use of a honeynet on the

Georgia Tech campus networks provides the university with a legal method to unobtrusively observe anomalous and misuse traffic directed on the campus network. The Georgia Tech Office of Information Technology (OIT) has cleared the establishment of the Georgia Tech HoneyNet with the Georgia Tech Legal Department. OIT authorized us to establish and monitor a honeynet on the Georgia Tech campus network in order to assist in protecting this network. These are the guidelines used at Georgia Tech, and may need to be modified to meet the specific concerns of both the network and the university administrators.

- The Georgia Tech HoneyNet is located within the production address space of the Georgia Tech Campus Network.
- The Georgia Tech HoneyNet uses computer operating systems that are representative of computers found on the Georgia Tech network.
- No IRC traffic is collected from our honeynet. If an IRC server is established on our honeynet that machine is immediately taken off line. This is to stay in compliance with the US Wiretap Act since a client utilizing that IRC may be unaware that the channel is on a compromised machine. A user such as this may have an expectation of privacy. We decided that it is best not to take the chance of violating the US Wiretap Act. Other nations will most likely have similar statutes concerning means that can be utilized to secure networks.

Technology We Used

Once you get approval to deploy a honeynet, the next step is to decide what type you want to deploy. There are currently two types of honeynets that can be employed on a network. These are GEN I, or first generation, and GEN II, or second generation. The type of honeynet depends on many factors including availability of resources, types of hackers and attacks that you are trying to detect, and overall experience with the honeynet methodology.

- GEN I HoneyNets - a simpler methodology to employ. This is the technology we choose to initially deploy when we first set up our honeynet. Its simplicity allowed us to learn how to run our honeynet. Although this is a simpler version of the honeynet, it still offers the opportunity to learn a lot about the basics concerning honeynets.
- GEN II HoneyNets - a more sophisticated version of the honeynet. This is the current configuration of the Georgia Tech HoneyNet after learning the basics by utilizing a GEN I technology. A [Honeywall CDRom](#) will soon be available that greatly simplifies the deployment of a GEN II honeynet. For additional information on honeynet technologies, refer to [KYE: HoneyNets](#).

We had chosen to initially deploy a GEN I HoneyNet on our enterprise network. Our initial objective was to detect machines within our enterprise network that had been compromised by automated script type attacks in addition to collecting rootkit research. We also wanted to start simple. As we became more comfortable with our data analysis methods we then employed a GEN II HoneyNet in our research efforts. Regardless of which type of honeynet you deploy, you will need resources (computers, switches, hubs, etc). Fortunately, we found honeynets to be pretty cheap hardware wise. The only people using your honeynet are attackers, so you don't need high performance systems. We used surplus machines and available research lab space here at Georgia Tech to deploy our honeynet. Surplus machines should be available at most academic institutions. Addition tools might

include a KVM type switch box to provide the ability to connect the honeynet machines to a single monitor and keyboard. If possible, the honeynet machines should be configured with removable hard drives to allow for compromised systems to be analyzed off-line. Given the high probability of available existing resources, you should not have to spend a lot of money for hardware to set up a honeynet. Time, on the other hand, can become a high demand resource concerning the honeynet (hint, grad students are cheap). We used surplus available resources here at Georgia Tech (to include student time, which is free). Were we to purchase the current configuration we estimate it would cost:

- 8 port KVM switch OSU21032 \$355.00
- Dell Dimension 2400 Computer x 8 @\$675.00 \$5400.00
- Data Castle BT-27 removable drives 2 per computer @ \$20.00 \$160.00
- Computer Rack \$1058.00
- **Total \$6973.00**

Our honeynet, in its current configuration, could be set up for less than \$7000. All computer resources, to include a KVM switch box, surplus computers, and removable drives with a disk duplicator were also available. It did not cost anything, besides time, to establish and maintain our honeynet. If you had to buy the hardware we used, we estimate the cost to be \$7000.00. You may find similar circumstances at your university.

Maintaining and Monitoring our Honeynet

While honeynets may be relatively simple and cheap to deploy, they can be a time intensive endeavor to maintain. Here at Georgia Tech we spend on average one hour per day to analyze traffic that was collected from the previous day. Keep in mind, the advantage with a honeynet is that you are not analyzing the data to determine what is hostile traffic and what is benign. Instead, with honeynets you operate under the assumption that all data it collects is bad, the reason you are analyzing it is to derive value from it. This makes honeynets a very effective tool for learning what unauthorized activity is happening on your network. When we initially established the honeynet, we used a session of snort running with the default rule set employed and then correlated these alerts against the tcpdump data via ethereal in order to learn what to focus on in ethereal. When we began to get more comfortable with snort we began to write our own snort rules to alert on specific suspicious traffic that we were observing on our campus network. Later on, using the tcpdump data we were able to go directly to ethereal and write our own filters to look for specific data. The capability to automate this analysis process does not exist; however, someone will have to do the manual analysis to determine what to automate on.

When a compromise occurs we expect to spend up to forty hours of analysis for each hour of attack traffic that has been collected from the honeynet. This process can be very time consuming but very rewarding. We now split this analysis up here at Georgia Tech and everyone who is involved in the investigation contributes to the compromise report. A report is produced for every compromise that occurs on the honeynet. Any suspicious behavior from campus machines directed to the honeynet is also reported. These reports are sent to OIT personnel, professors, individual network administrators, and other parties of interest here at Georgia Tech.

Our Honeynet Proves its Worth

The Georgia Tech campus is typical of many university environments, we face many of the same challenges and problems. We are very large environment required to give open access to many of our resources, making it extremely difficult to secure. Some statistics of Georgia Tech.

1. 15000 Students, 5000 Staff, 69 Departments
2. Mix of many different computers and operating systems
3. Not uncommon to have systems that can boot up in four different operating systems
4. 30000-35000 networked computers on campus
5. Academic, administrative, resident (REZNET), and research networks

Like many academic institution, we also find ourselves a large target, for several reasons.

1. High bandwidth capability to/from Internet (600Mbps/4 Terabytes per day here at Georgia Tech)
2. Lots of student machines with large storage capability
3. Can hide in "Academic Freedom" traffic. The concept of academic freedom concerning network traffic is that since Georgia Tech is a research institution, network traffic should not be restricted so that on-going research efforts are not interrupted. Because of the requirement for academic freedom, Georgia Tech chooses not to implement a firewall between the campus network and the Internet. Individual enclaves within the Georgia Tech network do utilize firewalls. An IDS is also run at the campus gateway with Out of band monitoring and follow-on investigation. These network characteristics make Georgia Tech well suited as a network to employ a honeynet as an additional IDS tool. Other academic institutions may have similar characteristics and benefit from the employment of a honeynet.

The original purpose, and one of the benefits of our honeynet, is research. The honeynet network is used by students conducting research in the areas of operating system and network security here at Georgia Tech. The honeynet has assisted in research efforts to include devising a new methodology for characterizing rootkits to aid in their subsequent detection as part of Ph.D. research. Other Ph.D. students are looking to incorporate analysis of the two years of honeynet data that we have collected here into their Ph.D. research efforts. One research effort produced a detailed analysis of new and existing rootkits collected from the honeynet and examination of inadequacies in the existing state of the art rootkit detection methodologies. We have produced one paper concerning the analysis of a previously unseen rootkit that was collected from the honeynet. We were able to characterize this rootkit and produce signatures in order to detect subsequent incidents. We have also provided suggestions to authors of GPL tools for rootkit detection in order address shortcomings and improve their tools. There is a current on-going research effort incorporating the honeynet with a Darknet (a research effort to track traffic to unused IP addresses) in order to establish a linkage between these two types of IDS tools. As previously mentioned, the use of a honeynet allows for operating systems of interest to be deployed in order to collect any attack traffic for subsequent analysis. The honeynet has provided for multiple research papers to be written here at Georgia Tech as well serving as a platform for on-going computer and network security research.

We have also found the honeynet to be an outstanding teaching tool - working with the honeynet provides hands-on experience to students who monitor the honeynet as part of an independent study program. This hands-on work provides a number of opportunities for these students. Studying the attacks and the root-kits allows these students to produce documentation and notes on how and when

the attacks occurred as well as analysis of the root-kits: how to remove them, how they work, and how they could be improved. As a member of the HoneyNet Alliance, Georgia Tech is expected to send a report each quarter to the Alliance detailing what we have seen on our Net. The students who monitor the honeynet as part of their independent study are expected to produce this report as well as a report of any compromise that occurred during this time period. HoneyNet data is also used in network security classes in order to teach students how to use tools such as ethereal and tcpdump in order to analyze attack traffic. Presentations on the honeynet are also given to classes and other groups of interest. During the 2003 'Capture the Flag' network security exercise hosted by the University of California Santa Barbara (UCSB) technology from the GEN II HoneyNet was employed. A session of Snort-Inline was run in order to analyze attack traffic. One Georgia Tech team placed third in the exercise after two UCSB teams.

What surprised us was the tremendous value of a honeynet for detection. Our honeynet helped to accurately detect over 165 compromised systems on our internal network. As our honeynet had very few false positives, it was in many ways easier to use, and more effective, than traditional detection technologies. Its greatest detection benefit was the ability to give detailed information, we were able to collect all data concerning these actual compromises to include passwords, remote IP addresses, and the methods of compromises used by the hackers. Very early on in its deployment our honeynet enabled us to identify an internal Georgia Tech system that had its passwords compromised by a hacker. This system was then used by the attacker to connect to other internal systems, including our honeynet. Our honeynet quickly identified this attacker and the tools and tactics he was using to infiltrate systems. The attacker set up a back door port on the honeynet in order to connect at a later time. We knew that this system had been compromised but kept it up and running in order to track the hacker's behavior. Several days after this system was compromised the hacker connected to the back door port established on this computer using another computer from within the Georgia Tech Enterprise Network. We immediately notified the Georgia Tech Office of Information Technology (OIT) personnel of this other potential compromised computer on campus. The OIT personnel took this computer off-line for analysis. Upon conducting analysis the OIT personnel could not find any indication that this other machine had been compromised. We learned this machine was not compromised by an exploit; instead it appeared the password to the system had been compromised. OIT personnel speculated that the hacker used some method to get the password of this machine. The hacker could have used a brute force technique to guess this password. He could have harvested this password from a dummy web site set up to harvest usernames and passwords from unwitting users (social engineering). The OIT personnel instructed the user to change his password by selecting a password that was more secure and to not use this password when establishing accounts at other web sites. The Georgia Tech OIT personnel have stated that it would have been very difficult for them to detect that this system had been compromised using the existing security measures that they have available. Our GEN I HoneyNet allowed us to detect a system that most likely had been compromised by a hacker with some skill. At this point the Georgia Tech OIT personnel became supporters of the honeynet after having seen its worth in detecting compromised systems on campus. Since its establishment in the summer of 2002 the honeynet has detected over 165 compromised machines on campus.

Lessons Learned

Out of all of this were a variety of lessons learned things to do and NOT to do. Hopefully this short list can help you avoid some common mistakes.

1. **Start Small** - If you are going to install a honeynet within your enterprise, start small. Begin initially with two machines (in order to detect sweep scans of your honeynet) with operating systems that you are familiar with installed behind the reverse firewall. This will allow you to begin to understand how to analyze the data that you will receive on the honeynet. You will also be able to fine tune your configuration. The more machines that you have, the more data you will most likely receive going to and from the honeynet.
2. **Maintain good relations with your enterprise administrators.** THIS IS CRITICAL! Inform your network administrators of the types of exploits that you are seeing. In some cases, they will already be aware of these exploits, but in other cases, you will have been the first person to notice them. The enterprise administrators should benefit from your efforts since they most likely provided you with the range of IP addresses that you are using for the honeynet.
3. **Focus on attacks and exploits originating from within your enterprise network.** These are the attacks that can do the most damage to your enterprise. Inform your enterprise administrators immediately of these types of attacks since they indicate machines that have already been compromised within the enterprise.
4. **Don't publish the IP address range of the honeynet.** There is no need to do this. Hackers and worms are constantly scanning across the Internet for machines to exploit. Your honeynet will be found and attacked.
5. **Don't underestimate the amount of time required to analyze the data collected from the honeynet.** This data must be analyzed every day. You will be collecting lots of information and it must be analyzed to provide any benefit. Most attacks take seconds to compromise and take over a vulnerable system. It can take weeks to analyze and document such an attack. Once again, we think it's well worth the effort.
6. **Powerful machines are not necessary to establish the honeynet.** The Georgia Tech Honeynet did not use state of the art machines and it functioned as intended. Everything we needed to establish our honeynet was already available on campus.

Conclusion

Honeynets can be a very powerful tool when deployed in an academic environment. Not only can it be used for securing your production environment by acting as a reliable and in-depth detection solution, but can also be used for a variety of extensive research projects. GA Tech has had great success applying honeynets to both areas.

The Honeynet Project