Testimony of

Dr. Doug Jacobson

September 9, 2003

Testimony
United State Senate Committee on the Judiciary
Pornography, Technology, and Process: Problems and Solutions on Peer-to-Peer Networks

September 9, 2003

Dr. Doug Jacobson
Associate Professor, Electrical and Computer Engineering, Iowa State University
President & Chief Technology Officer, Palisade Systems, Inc.

Mr. Chairman and members of the committee I would like to thank you for the opportunity to appear before you today to discuss the issues surrounding peer-to-peer networks.

By way of introduction, I wear several hats. I'm currently an associate professor of Electrical and Computer Engineering at Iowa State University. My duties include teaching advanced network protocols and a course on information warfare. I also serve as the Director of Iowa State's Information Assurance Center, which has over 30 faculty associates on campus. Iowa State University was one of the first seven universities designated a Center of Academic Excellence in Information Assurance Education by the National Security Agency. For the past seven years I have assisted local law enforcement agencies in investigating numerous computer crimes including software theft, harassment, and child pornography.

I also started a security and Internet filtering company in 1996 and currently serve as its president and chief technology officer. Palisade Systems was formed to address the issue of pornography on school networks. The patented technology licensed from Iowa State University enabled Palisade Systems to develop its first product called ScreenDoor, which is a web filter. In summer of 2000 we introduced PacketHound which is designed to detect, monitor, and block unauthorized peer-to-peer applications. This product was so unique that it won an R&D 100 award in 2001, an honor reserved for the top technology innovations of a given year. Since that time Palisade Systems has become a leader in peer-to-peer technology, conducting research and publishing studies on the subject.

During my testimony today I will discuss the use of peer-to-peer applications for the distribution of pornography. I will talk about how peer-to-peer applications have evolved to challenge the efforts of security professionals to stop their misuse for criminal purposes. Finally I will address the current state of technology used to monitor or block peer-to-peer applications.

First, I would like to address the issue of pornography and peer-to-peer networks. There are several issues that make pornography on peer-to-peer networks more problematic than web or FTP-hosted pornography. You don't have to look for pornography on peer-to-peer networks; it

will find you. There are no effective controls regarding content provided on a peer-to-peer network, the only information you are given is a file name. A good example of this problem hit home this spring when I was teaching my information warfare class. To give students an opportunity to study the security problems associated with peer-to-peer networks, I set up a peer-to-peer node. I searched for a file that I had created and placed on the peer-to-peer network. I received several matches to my search request, but when I downloaded and viewed the files, they contained embedded links to pornography sites. It turns out that people have designed custom applications that will answer any query on a peer-to-peer network and return a match to the search string. When the unsuspecting user requests the file, the custom software will return a file containing the embedded links.

Another issue is that peer-to-peer networks seem safer to someone who does not want to be discovered. A user simply types in a search string and receives a list of files that match words in the search string. The user then clicks on the file name and that file is transferred to his or her hard drive. This gives the user a feeling of anonymity. Moreover, this feeling of anonymity works both ways in that the both the provider and the requester of the file are not easily identified and therefore they may believe that their act cannot be monitored.

It is easier to search for pornographic files whose titles contain key words on a peer-to-peer network than to use a standard web search engine. For example, I entered the search string "Teen Sex" into Google and received over 5 million hits, ranging from websites that indicated they contained graphic images to others devoted to discussions of teen pregnancy and other legitimate health and social issues. By contrast, upon entering the same search string into a Gnutella peer-to-peer network, after one minute I received a far more narrow listing of more than 1300 matches containing only movies and pictures whose names implied specifically pornographic content, including a number of sources offering access to child pornography. A quick review of both lists showed that the peer-to-peer network provided a quicker and easier way to obtain pornography of all varieties.

Palisade Systems conducted a study of searches on a Gnutella network. Acting as a node on the network, we gathered 22 million searches between February 6 and February 23, 2003. Details identifying the individual such as the IP addresses were removed to maintain the privacy of the users. In addition, none of the requested files were downloaded.

Over the nearly three weeks of monitoring, the study found that:

? 42% of all requests were for all types of pornography (168,000 of 400,000 search requests)
? 6% of all search requests were specifically for child pornography (24,000 of 400,000 search requests)

While this study focused on search requests you could reasonably assume that the number of hits would be proportional to the number of requests, since people don't constantly search for things that they can't find. Studies published by other organizations have come to similar conclusions all backing the claim that peer-to-peer networks have become a widely accepted vehicle for the distribution of pornography.

Palisade Systems has been monitoring peer-to-peer networks on an on-going basis since the conclusion of the first study. Although I am not able to release statistics, we have gathered significant examples of pornography being downloaded at federal and local government agencies and elementary schools across the nation. There are two or three other companies that can also provide similar information. We will be reporting our findings in a report within the next month.

Many articles have been published about the more dangerous uses of peer-to-peer applications, ranging from security violations to child pornography. I would like to focus briefly on the methodology used by peer-to-peer applications to avoid detection and to hide both the content and the source of the information. An argument could be made that, other than for governmental security purposes or to protect the transfer of confidential or proprietary information between commercial interests, legitimate peer-to-peer applications would not need to hide from detection or evade monitoring.

Information sharing on the Internet has moved from a centralized approach with FTP servers, WAREZ sites, and web servers to a distributed approach. For those of you unfamiliar with them, a WAREZ site is used to store and distribute pirated software or other illegal information. In some measure, this transition has been a reaction to the successful efforts of system administrators and law enforcement agencies in stopping centralized approaches to illegal file sharing. With a centralized approach there needs to be someone in charge of the system and therefore, an identifiable person who can be prosecuted. However, several new peer-to-peer protocols have been designed to evade detection and to circumvent standard security mechanisms developed to address centralized approaches. While you can argue that peer-to-peer represents a new paradigm in networking and has legitimate uses, a case can be made that certain protocols have been designed expressly to aid in the covert procurement and/or distribution of information. Indeed the reason Palisade Systems developed PacketHound in the first place was because many of these new applications claimed they could not be blocked, thus suggesting a high likelihood of criminal intent.

First a little background. In order for applications to communicate with each other, they need to know two things; they need to know the address of the computer that hosts the application and the address of the specific application. The address of the computer is known as the "IP address" and the address of the application is called the "port number". To use a familiar model, the address of a computer is analogous to the address of house, and the address of the application is analogous to the name of a particular person in that house. Of course, this is how mail can be delivered to the right person at the right house. Similarly, with computer applications you need to know the port number of the application you wish to communicate with. In most cases the port number has been assigned to a given application; for example the World Wide Web is port 80 and email is port 25. From a security standpoint, such an assignment is also helpful insofar as you can stop all traffic directed toward a given set of applications. To use out postal model once more, this would be like tossing all mail addressed to "Occupant."

Peer-to-peer has become the latest battlefield in which developers of peer-to-peer applications are constantly working to outsmart the defenders of the network. I would now like to talk briefly about the state of technology used to monitor, control, or block Peer-to-peer networks. I have chosen to divide the technology into four categories.

The first category is port blocking, which is a common method used at the front end of most corporate networks Early peer-to-peer networks evolved to employ techniques that bypass port blocking mechanisms. The first method was "port hopping", in which the port number is not fixed but can be changed. This approach prevents network administrators from blocking the particular ports associated with the peer-to-peer applications. As administrators adapted and started filtering out all but a few critical ports, peer-to-peer protocols adopted a new technique called "tunneling". In this technique, peer-to-peer protocols used the port number assigned to other applications like web traffic so that their traffic is allowed to pass through the defenses.

Another method is called "signature based," which is the concept behind the Palisade Systems PacketHound product. This technology examines all of the packets on the network to determine if certain patterns (called signatures) can be found on the network. For example, Gnutella always starts a new connection with a certain set of messages. Signature-based technologies require updating whenever a new protocol is found or a current protocol changes. PacketHound can currently detect over 20 Peer-to-peer protocols and has signatures for over 100 protocols.

The "content-based" method involves trying to determine the content of the file being transferred by looking at the data. This technology has had limited success in the web filtering market. We have all heard about web filters blocking access to a website dedicated to cancer research because the word "breast" is contained on some of the pages. This method can work to identify known data like copyrighted music, but does not work when the data is not known or cataloged. For example, there is not a catalog of all known pornographic images therefore a content-based identification could not be used. Palisade Systems has teamed up with a company called Audible Magic to create a product that indicates when peer-to-peer networks are exchanging copyrighted music. Another drawback to content-based filtering is the use of encryption by some peer-to-peer networks.

The final technology is to allow only known traffic types on a network and to block all unknown traffic. This type is often called "white listing" where only the know traffic is allow to pass. This technology can work well in a controlled environment where the traffic is known. Palisade Systems has introduced a product called FireBlock designed to allow only know traffic, which also won the R&D 100 award this year.

The newest steps in the evolution of peer-too-peer networks are "encryption" and "anonymous access". Anonymous access is designed to hide the source of the information. First generation peer-to-peer protocols provided the address of the computer originally containing the file that was downloaded, thus allowing network administrators and law enforcement agencies to trace the source of the files. The latest step in the evolution of peer-to-peer applications uses encryption techniques to hide both the source of the data and the data itself. The best example of this evolution is the newest application called Earthstation 5. This protocol uses encryption, anonymous access, and tunneling. The web site for Earthstation 5 makes it clear they are working to stop any efforts at filtering.

A couple of observations can be made from a review of these technologies. First, while each technology has certain limitations, using multiple technologies in a layered approach seems like the best defense in a corporate environment. However, this method often required knowledgeable staff and constant monitoring of the networks. Second, most technologies are focused on a

corporate market. Many of the solutions are cost prohibitive for small organizations like schools, small business, etc. Furthermore, these technologies are not designed for the home users. This leaves individuals on their own to solve the problems of peer-to-peer networking, which naturally leads us to the question "What's the home owner to do?"

Peer-to-peer applications are easy to find and install. If a home user allows these applications to be installed, little can be done to prevent downloading of pornography or other material. Unlike web filtering where certain sites can be blocked and web access can be monitored, peer-to-peer traffic cannot be filtered based on its content. This leaves a home user no choice but to either allow peer-to-peer activity and all of its associated risks or not allow any peer-to-peer applications on their machines.

It would be possible for Internet service providers to offer a service that blocks peer-to-peer traffic similar to the web filtering provided. The bottom line is that the home user needs to be educated about the potential dangers of peer-to-peer networking.

In summary, I've outlined how peer-to-peer networking has evolved to avoid detection and filtering. I see no signs of this evolutionary path slowing down in fact with the advent of the newest protocols like Earthstation 5, we will be facing increasing challenges in the years ahead. Also, given the inherent distributed nature of peer-to-peer protocols and the difficulty in identifying these networks, I predict that peer-to-peer networks will become a method of choice to distribute illegal materials across the internet. Companies like Palisade Systems in conjunction with research universities like Iowa State University will continue to develop new technologies to combat the evolution of peer-to-peer networks.

I would like to thank you for this opportunity to testify today. I would be pleased to answer any questions you might have.