Pipeline
Knowledge Is Power

www.pipelinepub.com Volume 6, Issue 6

Exposing the Dark Side of the Cyber World

By Dr. Antonio Nucci

The Internet has become the central nervous system for our networked life. As a global network of loosely connected IP-based networks, it reaches into every country and provides governments, businesses and consumers worldwide with a common platform for communication. And now, a new kind of criminal has emerged.

As the 21st century criminal has moved into new realms and dimensions, law enforcement agencies and government organizations are in hot pursuit. The pervasive nature of cyber crime ranges from loss of proprietary corporate information to the loss of life, from national security to cyber warfare. From predators exchanging child porn and scammers stealing identities to countries attacking countries, cyber crime does not discriminate.

Quantifying the Spread and Impact of Cyber Crime and Cyber Terrorism

The FBI estimates that all types of computer crime in the U.S. now cost industry about $400 billion, while officials in the Department of Trade and Industry in Britain say computer crime has risen by 50 percent from 2005 to 2006. It is estimated that only 5 percent of cybercriminals are ever arrested or convicted because the anonymity associated with Web activity makes them hard to catch, and the trail of evidence needed to link them to a cyber crime is hard to unravel. CERT/CC estimates that as much as 80 percent of all computer security incidents remain unreported (according to Marcia Savage of SearchSecurity.com).

There are certain steps to be taken before we can successfully combat cyber crime.

First, and foremost, it is time to increase our understanding of the language and the many dialects (i.e. protocols, applications and services) being spoken in the cyber world. Network traffic monitoring and measurement is increasingly regarded as an essential function for understanding and improving the performance and security of our cyber infrastructure. With networking technologies and services evolving rapidly, as witnessed by the explosive growth of the Web, peer-to-peer networks, multimedia, gaming, etc., accurate network traffic monitoring is required to ensure the security and optimize the efficiency of our cyber world.

Second, it is time to promptly identify cyber users and communities of cyber users whose activity and content may harm the safety and transparency of the cyber world.

Third, it is important to gain visibility into who is the real person behind an alias or cyber-identifier used to enter the cyber world. A critical problem in this digital world is knowing with whom you are interacting.

Current Approaches – and Their Weaknesses

Network traffic monitoring and measurement is increasingly regarded as an essential function for understanding and improving the performance and security of our cyber infrastructure. With networking technologies and services evolving rapidly, as seen with the explosive growth of the Web, peer-to-peer (P2P) networks and the GRID, accurate network traffic monitoring is required to ensure the security and optimize the efficiency of our cyber world.

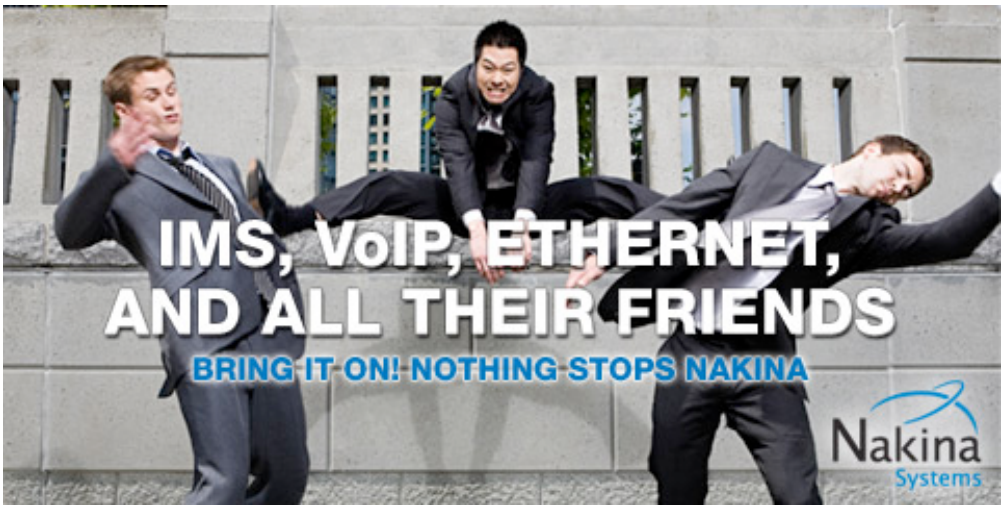

Critical to the success of any such tool is its ability to accurately -- and in real time -- identify and categorize each flow (i.e., sequence of packets associated with the same cyber world transaction/connection) by the application responsible for it. Identifying network traffic using port numbers was the standard in the recent past. This approach was successful because many traditional applications use port numbers assigned by or registered with the Internet Assigned Numbers Authority (IANA). The accuracy of this approach, however, has been questioned because of the evolution of applications that do not communicate on standardized ports. Many current-generation P2P applications use ephemeral ports, and in some cases, use ports of well-known services such as Web and FTP to make them indistinguishable to the port-based classifier.

Techniques that rely on inspection of packet contents have been proposed to address the diminished effectiveness of port-based classification. These approaches attempt to determine whether or not a flow contains a characteristic signature of a known application. Studies show that these approaches work very well for today's Internet traffic, including P2P flows. In fact, commercial bandwidth management tools and network security appliances use application signature matching to enhance robustness of classification and deep inspection of packet content even in the case of encapsulated protocols within each other (i.e., x in HTTP).

Indeed, very recently several threats appeared to use this technique to hide their presence and break through firewalls and other security devices. The progress in hardware acceleration has allowed

packet content inspection techniques to run at speeds as high as 40 Gbps and made them the most commonly used approach to gain visibility into any Internet stream.

Nevertheless, packet-inspection approaches face two severe limitations. First, these techniques only identify traffic for which signatures are available. Maintaining an up-to-date list of signatures is a daunting task. Information is rarely available, up to date or complete. Furthermore, the traditional ad-hoc growth of IP networks, the continuing rapid proliferation of applications of different kinds, and the relative ease with which almost any user can design and infiltrate a new application to the traffic mix in the network with no centralized registration, contribute to this "knowledge gap." Second, packet inspection techniques only work if full packets (i.e., header and payload) are available as an input and are completely harmless whenever coarser information is provided (i.e., traffic flows). Unfortunately, only a few service providers today have equipped their networks with packet inspection appliances, while the majority of them has access only to traffic flows extracted directly from the routers, either sampled or unsampled.

To overcome these two fundamental problems of packet content inspection appliances, the research community has focused on a new family of techniques called "flow-features-based analysis." The common goal of these techniques is to identify which application class a traffic flow belongs to when using traffic flow information only. These techniques achieve the flow-application class mapping by extracting and analyzing hidden properties of the flow, either in terms of "social interaction" of hosts engaged in such a flow or the spatial-temporal behavior of several flow features such as flow duration, number and size of packets per flow, inter-packet arrival time, and so on. A variety of more sophisticated data mining algorithms have been proposed on top of such framework, such as supervised and un-supervised machine learning, clustering and graph-theoretical approaches, to increase the detection rate while decreasing the false-positive rate.

These techniques all lack a fundamental attribute that make them impractical from an operational perspective (i.e., the precision identification of the application responsible for the observed flow in contrast to packet content inspection techniques). This is a fundamental question to answer, as today's network operators must know the nature -- legitimate or malicious -- of any single bit of information flowing through their pipes. Furthermore, as the application classification process might still be prone to classification errors, these techniques are not reliable for content billing or for robust application security.

Exploring the feasibility of bringing together the benefits of the two families has not attracted much attention in the research community. The only framework available is called ACAS, aimed at automatically extracting application-specific signatures by processing the first 200 bytes of the first few packets. Although this work is novel from a pure conceptual perspective, the practicality of such a framework is still questionable in many aspects. First, it has been tested on only a very few well-known applications such as FTP, POP3, IMAP, HTTPS, HTTP, SMTP and SSH. Thus, it is not clear how well it will perform in a more application-enriched environment. Second, its underlying algorithms require offline training on the set of applications that the operator is interested in detecting. Thus, it is not capable of recognizing "zero day" applications but it is still based on the network operator's knowledge of which applications are on the wire. Most importantly, the network operator is still required to go over this manual and tedious process of generating traffic with the set of applications he is interested in to properly train ACAS. The ultimate training of ACAS on these "never seen" applications must be executed in a controlled and clean lab environment. ACAS may suffer high false-positive rates for these new applications due to the discrepancy in environments, (i.e., a clean and controlled lab for offline training and an enriched and more complicated real network for online application classification).

The ideal solution would leverage the merits of the packet content inspection techniques by guaranteeing the high-accuracy in classifying application-specific traffic, while providing the robustness to detect zero-day applications and couple them with ability to work with both packets and flow characteristics of the flow-based behavioral analysis techniques.

A New Way to Think About the Cyber World

While the cyber world is seen as a "dark" space and governments have increasingly expressed their concern about the cyber world's role in public safety and national security, we still have not done enough to shed light on the cyber world and its users.

To do so, we must first understand it. The cyber infrastructure must not be thought of as just the physical infrastructure made of optical fibers, servers and routers. Rather, the cyber infrastructure is also about protocols, applications and services being used to enable communications among any number of end points (users). We must discover who is behind the nickname, Mac or IP address, or VoIP number – perhaps by using novel biometric techniques to profile users' communication as they access the cyber world. Reconstructing today's missing links between the cyber ID and the real person would make the cyber world a safer place to visit.