Using Good Botnets to Beat Bad Ones

A university team has developed a technique that uses large groups of computers, such as those in many existing networks, to defeat botnets, which are large groups of computers that hackers infect and commandeer to launch spam, distributed denial-of-service, and other attacks.

The University of Washington researchers say their Phalanx system would be relatively inexpensive to implement and could scale to fight botnets of any size.

The approach would utilize any large group of voluntarily participating computers as a shield against botnet traffic, explained Colin Dixon, a doctoral student who worked on the project. This could include existing systems such as content distribution networks (CDNs), which companies use to send video, audio, and other material to subscribers.

Via vectors such as infected e-mail attachments and web pages, hackers surreptitiously gain access to and plant malicious software on multiple computers. The software lets the hackers issue commands to the computers and thereby launch massive attacks.

Current approaches to fighting botnets generally have had limited success, according to Dixon.

For example, heuristic-based filtering of attack-related traffic is ineffective because the filters must distinguish between real and malicious traffic in real time, which is increasingly difficult as many attackers can make their traffic appear legitimate, Dixon said.

Phalanx uses a large network of computers acting as nodes to shield a server. Information coming to protected servers first passes through the shield.

The system breaks incoming traffic into pieces and distributes them randomly to the subset of nodes protecting a particular transmission. The random nature of the distribution makes it difficult for botnets to target the specific nodes that are shielding a server. And attacking every potential node would be difficult because there are so many of them.

To pass through the system, packets must have a cryptographic token that the Phalanx administrator issued or the sending computer must solve a cryptographic puzzle. Hackers won't have a token, and their computers are generally set to send out as many messages as possible and not spend a lot of time on a single transmission, which would be necessary to solve the puzzle. Legitimate senders' computers, on the other hand, will generally take the necessary time.

Phalanx uses a filtering ring that won't admit packets if they don't have the necessary token or don't contain information that shows that the sending computer solved the puzzle.

The shielded server uses the information contained in packet headers to reassemble them in the proper order.

The protective network's size lets it handle traffic without excessive delays.

The University of Washington researchers simulated an attack by a million-machine botnet on a server protected by 7,200 shielding computers, and the server functioned almost normally, Dixon said.

Adding more nodes and expanding the filtering ring would enable a Phalanx system to scale.

David Farber, distinguished career professor at Carnegie Mellon University, said that Phalanx is an interesting idea but that he is unsure whether it will work in practice. Also, he said, the technology will have to be able to adapt to changing approaches by hackers.

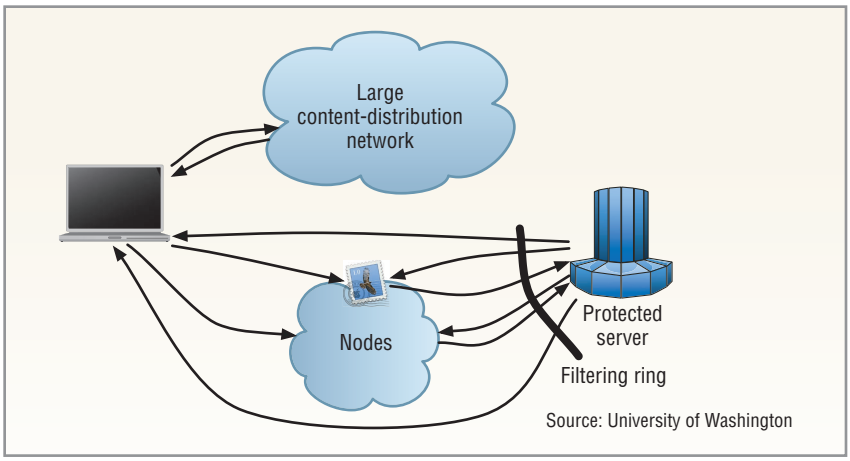

The Phalanx system uses a large group of computers, like those in many networks, as a shield to combat botnets. Phalanx confounds hackers by breaking incoming traffic into pieces and distributing them randomly to the nodes protecting a server. Moreover, the system uses a filtering ring. To pass through, either packets must have a Phalanx cryptographic token or the sending computer must solve a cryptographic puzzle. Neither is likely to be the case for botnet traffic.

News Briefs written by ***Linda Dailey Paulson****, a freelance technology writer based in Ventura, California. Contact her at ldpaulson@yahoo.com.*

Published by the IEEE Computer Society

New Circuit Element Could Yield Better Memory Chips

Hewlett-Packard scientists have developed a new circuit type—previously the subject only of theoretical work—that will enable construction of very dense, energy-efficient memory chips that could be used in small, powerful computers.

The circuit is called a *memristor* (memory resistor). It controls electrical current by providing flow resistance, and this gives circuits memory properties. Moreover, the memristor doesn't lose stored data when the host device shuts down.

According to HP Senior Fellow R. Stanley Williams, his research team has used the technology to build ultrahigh-density, 15 × 15 nanometer prototype crossbar switches, which channel data between any two devices attached to its ports. The technique could be used to construct even smaller circuits, he added.

Williams said memristors will be particularly useful for nonvolatile random-access memory. Unlike today's RAM, NVRAM would retain information after a computer shuts down, giving workers faster data access.

University of California, Berkeley, professor Leon Chua predicted the memristor in a 1971 paper and gave the element its name. Chua argued that a fourth type of basic passive circuit element—in addition to resistors (which provide static resistance to the flow of electrical charge), capacitors (which store charges), and inductors (which resist changes to the flow of charge)—is necessary to provide an important but missing functionality.

Unlike a resistor, a memristor increases or otherwise changes its resistance based on the electrical current flowing through it. The circuits could thus record and rewrite data by interpreting high and low

Robot Conducts the Detroit Symphony Orchestra

Honda Motor Co.'s well-known Asimo robot conducted the 51-member Detroit Symphony Orchestra in a special sold-out performance promoting the organization's The Power of Dreams Music Education Fund for Detroit youth.

Charles Burke, director of education for the Detroit Symphony Orchestra (DSO) and artistic director for the group's Civic Youth Ensembles, explained that the organization was interested in letting Asimo conduct its musicians as part of an effort to call attention to the need for support for local music education.

Honda Motor Co.'s Asimo robot recently "conducted" the Detroit Symphony Orchestra. Engineers programmed Asimo to carry out the tasks necessary to lead the orchestra in playing "The Impossible Dream."

This is crucial because budget cuts have hurt music programs throughout the state of Michigan, noted Honda spokesperson Alicia Jones. She said the company also wanted to determine how advanced robotics might work with musical applications.

Asimo conducted "The Impossible Dream" from the musical *Man of La Mancha*. Jones said the piece was chosen based on its theme of following one's dreams.

Burke said that the experience was fun for the musicians and that the robotic conductor was easy to follow.

Asimo (Advanced Step in Innovative Mobility) is the result of research that Honda began in 1986 into building a walking, humanoid robot. Honda built the first Asimo in 2002 and introduced the current version in 2005.

The robot moves via 34 servomotors powered for an hour between charges by a 51.8-volt lithium-ion battery in the machine's backpack. Two stereoscopic cameras in Asimo's head provide the robot and its operator with the ability to "see" its surroundings.

Honda engineers programmed Asimo to mimic the motions a human would make if conducting "The Impossible Dream." Jones said the company wanted to ensure the movements would be accurate enough for the musicians to follow. The robot used its onboard computer to maintain a steady rhythm as well as to signal changes in the music accurately. ■

resistance as binary ones and zeros, respectively.

Memristors can also read a full range of resistance levels and thus can function as analog devices, which lets them perform more complex computing operations.

HP researchers built their memristor when they were trying to develop molecule-sized switches. These switches could replace transistors once photolithographic manufacturing techniques can no longer continue making conventional circuits smaller, an effort many scientists are working on. Smaller circuits are desirable because they let manufacturers either pack more of them onto processors, thereby increasing power; or pack the same number onto smaller chips, thereby reducing their size.

During their research, the HP scientists discovered Chua's paper. They also experimented with titanium dioxide and found that its oxygen atoms are negatively charged ions and its electrical field is huge. This lets oxygen ions move and change the material's conductivity, a necessity for memristors.

The researchers then sandwiched two thin titanium dioxide layers between two electrodes. One titanium dioxide layer has 1 percent fewer oxygen atoms than the other. Applying a small electrical current causes the atoms to move around and quickly switch the material from conductive to resistive, which enables memristor functionality.

In the process, the device uses little energy and generates only small amounts of heat. Also, when the device shuts down, the oxygen atoms stay put, retaining their state and the data they represent.

Chua said it appears HP has made the perfect example of the memristor he wrote about 37 years ago.

Manufacturers could make memristors in the same chip fabrication plants used now, so companies would not have to undertake expensive retooling or new construction.

The memristor's major challenges are its relatively low speeds and the need for designers to learn how to build circuits with the new element.

The devices may begin appearing commercially in about five years, Williams predicted.

Getting Ready for SuperSpeed USB

A standards organization is finalizing an ultrafast version of the popular universal serial bus plug-and-play hardware interface designed to better meet the needs of devices requiring high-speed communications.

The USB Implementers Forum—which supports and promotes universal serial bus technology—plans to complete work on the USB 3.0 specification during the next few months, according to organization president Jeff Ravencraft, who is also a technology strategist in Intel's Communication Technology Lab. The USB-IF expects the first USB 3.0 products to appear in late 2009 or early 2010.

Ravencraft said the technology's data rate will be 5 Gbits per second, more than 10 times USB 2.0's 480 Mbps.

The new technology, also known as SuperSpeed USB, would connect products handling large data volumes and requiring fast performance—such as consumer electronics, communications devices that store and transfer large amounts of data, external storage, and flash-based peripherals—to PCs or servers. These devices could include digital cameras and camcorders, USB drives, and media players.

Today's products require faster technologies to quickly transfer the multimedia and other large data files that have become so popular, Ravencraft explained. A 25-gigabyte file would transfer in about 15 minutes with USB 2.0 but in only about 70 seconds with SuperSpeed USB.

The USB 1.0 Promoter Group released the short-lived USB 1.0 in 1996 and then the popular 12-Mbps USB 1.1 in 1998. The USB 2.0 Promoter Group released its standard in 2000. There is also a wireless version of the technology.

In addition to being fast, SuperSpeed USB provides full-duplex communications. It can thus send transmissions and acknowledgments simultaneously, making the process even more efficient.

Moreover, USB 3.0 utilizes two additional high-speed pairs of wires that can transmit data simultaneously. This enables a very fast signal rate, explained Rahman Ismail, Intel software architect, chair of the USB-IF's Compliance Committee, and the USB 3.0 Promoter Group's chief technical officer.

USB 3.0 works with copper wiring or optical fiber, which lets it flexibly interact with different transmission media, noted Ravencraft.

Unlike its predecessor, SuperSpeed USB will not continuously poll devices to determine whether there is traffic to transmit. When there is none, the new technology will move into a low-power state. It will thus be more energy efficient, Ismail said.

SuperSpeed USB also adds quality-of-service capabilities for devices requiring guaranteed bandwidth.

Because of USB 2.0's wide implementation and suitability for most of today's applications, USB 3.0 probably won't be widely adopted for a while, said Brian O'Rourke, principal analyst with market research firm In-Stat.

Editor: Lee Garber, *Computer*, l.garber@computer.org