New Blacklisting Technique Improves Network Security

Researchers have developed an algorithm that generates useful blacklists for networks by taking information from victims of past network attacks and predicting which hacker sites are likely to target specific networks in the future.

Computer scientists from SRI International—a nonprofit research institute—and the SANS Institute—which provides information-security training, certification, and research—created the highly predictive blacklist (HPB) algorithm.

Blacklists, which contain IP addresses previously involved in malicious activity, are an increasingly popular security technique. However, there are problems with the two main blacklisting approaches.

Large blacklists that security companies compile and sell to customers contain the most prolific attack sources they have found. But they can load up firewall filters with many IP addresses that will never attack the host organization.

Specific networks can create their own local blacklists based on the IP addresses that have most frequently attacked them, but this won't predict new potential sources.

The SRI and SANS researchers have taken a different approach, developing a strategy for constructing blacklists by correlating attackers' previous preferences for victims' networks as a way to predict other networks they might target.

The new approach works with information about harmful online activity that SANS collects via its DShield system, which receives firewall logs from participating organizations and uses them to analyze attack trends.

HPB uses two analysis engines to create a blacklist for each network it protects. One engine ranks attack sources based on their relevance to the network for which it is developing a blacklist. The other determines the severity of potential attacks.

The system uses a relevancy-ranking scheme that calculates how likely it is that certain IP addresses involved in previous attacks will attack a specific network in the future and thus should be included on a blacklist, explained SANS chief research officer Johannes Ullrich. The scheme, which SRI developed, is mathematically similar to the one Google uses to determine search results' relevancy to queries.

HPB looks at commonalities among prior attacks by a site, such as the nature of the assault, noted SRI program director Phillip Porras.

For example, the technique determines relevancy by identifying the types of networks that certain types of attacks favor. If an attack works via a specific TCP/IP port, they would not be considered relevant to networks that typically block that port.

The DShield service uses this type of information to construct daily customized blacklists—with potential threats ranked by relevancy—for networks that share their firewall logs with DShield. A wide range of organizations participate, including ISPs, universities, and corporations, noted Porras.

Each participant chooses how to use the information, Ullrich explained.

According to Porras, 90 percent of HPB users have reported improvements in blacklisting performance.

Blacklisting's key problem is the number of false positives it can generate, said Jaime Yaneza, threat researcher manager for Internet-security vendor TrendMicro.

SRI has applied for a patent for HPB, Porras added, and may license it for third-party use. The researchers hope to apply the approach to other types of security intelligence, such as that used to fight spam.

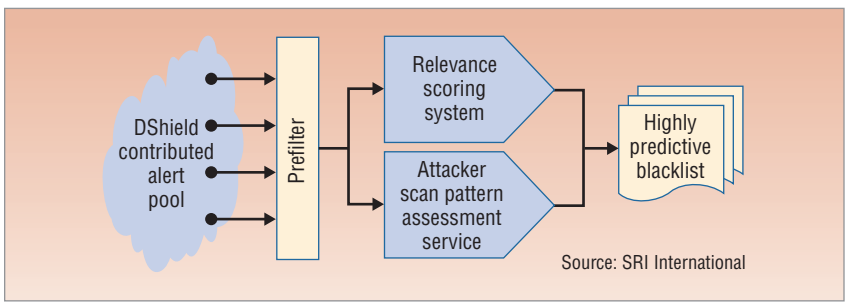

Scientists have developed an algorithm that creates blacklists by using information from past network attacks to predict which hacker sites are likely to target specific networks in the future. The highly predictive blacklist approach works with information about harmful online activity that the SANS Institute collects via its DShield system. After filtering out unnecessary information, HPB runs the data through one system that ranks attack sources based on their relevance to a network being protected and one that determines potential attack severity.

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

Fat-Tree Network Approach Could Cut Data-Center Costs

A new technique could provide data-center owners with an alternative to buying faster, more expensive switches as a way to improve performance.

University of California, San Diego, scientists have applied cluster-computing techniques to data-center network architectures. This is designed to let network administrators employ large numbers of widely available Gigabit Ethernet switches and routers to deliver better performance at lower cost than using high-priced 10-Gigabit-Ethernet equipment in the network, explained professor Amin Vahdat.

The approach lets data-center operators increase performance without significantly raising expenditures.

A 20,000-node system using 10-Gigabit Ethernet switches at higher network levels could cost as much as $28 million to construct and would cost even more when 40- and 100-Gigabit Ethernet technologies hit the market, Vahdat explained.

A 20,000-node network using only Gigabit Ethernet switches would cost just $4 million, he noted.

He said his team's approach utilizes a version of the fat-tree network architecture based on a standard three-tier architecture: each machine on the network works with an edge switch, the edge switches work with aggregation switches, and the aggregation switches connect to a core switching infrastructure.

This technique is a type of cluster computing, in which large groups of less powerful devices are organized so that they outperform more powerful devices.

The approach's primary obstacle is that its greater number of switches introduces operational complexity and requires more cabling, said Jonathan Eunice, president and principal analyst of Illuminata, a market-research firm.

To solve this problem, Vahdat noted, his team is looking at implementing its technique via wireless technology or free-space optics.

In addition, Eunice said, the researchers haven't adequately addressed potential bandwidth and latency problems caused by using slower switches.

According to Vahdat, his team also still must develop ways to get TCP/IP to work effectively on top of the new topology.

The researchers are working with several companies and hope to have a demonstration data center using the technology ready within 12 to 18 months.

Virtual Worlds Use Virtual Punishments to Rein in Participant Behavior

Parts of the virtual-world community have a reputation as wide-open places where participants can expect many types of behavior, including those that offend some. But the operators of many of those virtual worlds are now beginning to try to limit such behavior.

Most virtual worlds are environments with stringent rules, noted Danny O'Brien, international outreach coordinator for the Electronic Frontier Foundation, a digital-rights-advocacy organization.

Linden Lab created *Second Life*, the biggest virtual world, as an open environment but has always had user-behavior guidelines, according to Ken Dreifach, the company's deputy general counsel.

However, O'Brien said, some virtual worlds started with few rules and little or no supervision. And some, O'Brien noted, did not limit specific types of conduct.

This has caused concern about some users sending profanity, sexually explicit content, or links to such material, as well as threats or harassing messages.

Untoward behavior has been a concern of VZones' members and staff, said Justine Reichman, CEO of the company, which operates virtual worlds.

In response, she said, offenders using prohibited behavior can be given warnings, temporary service restrictions, and other punishments, which can escalate if problems continue. To pursue punishment, though, she noted, the company requires proof of wrongdoing, such as screen captures of offensive content.

VZones also provides an "ignore" function that lets users block communications from designated individuals and thereby stop unwelcome communications and contact from them.

According to Reichman, VZones' terms of service offer participants firm expectations for behavior that will be enforced and are updated to ensure they stay current. And ultimately, she said, the company has the legal authority to bar misbehaving users from the service.

Other virtual worlds are following this trend and are enacting rules and enforcing them, according to O'Brien.

Tongue Computing System Could Help the Disabled

Georgia Tech University researchers have created a system that could let severely disabled people, such as those who have suffered a stroke or paralysis, use the inside of their mouth to issue commands to computers, wheelchairs, home appliances, or other devices.

Their prototype Georgia Tech Tongue Drive System lets users control a power wheelchair by moving their tongue around their mouth, said assistant professor Maysam Ghovanloo.

The system maps the motion of the user's tongue to specific computer mouse or keyboard commands. So far, Ghovanloo noted, his team has mapped six different commands: left, right, forward, backward, single-click, and double-click. The system uses the computer as an interface to control a wheelchair or other device.

In the future, he added, the technology could work with many additional commands because it could detect and map even the smallest individual and combination tongue movements, including touching each of the 32 teeth in the typical adult mouth.

The system, which consists of off-the-shelf technologies, currently works via a 3-millimeter-wide magnet placed on the tip of the user's tongue. Highly sensitive magnetic sensors attached to a headset with extensions that reach the wearer's cheeks monitor changes in the magnetic field as the tongue moves around the mouth. The sensors send data to a receiver on the headset, which transmits the information wirelessly to a nearby laptop for processing and translation into commands.

The researchers plan to adapt their code-translation software so that it could run on small devices such as a PDA or smart phone that users could wear or attach to a wheelchair.

Experiments on able-bodied volunteers have used a magnet adhered to the tongue using human-tissue glue, which lasts a few hours. People could use a lentil-sized magnet implanted under their tongue's surface or a magnetized tongue piercing for a longer-lasting approach, Ghovanloo said.

The Christopher and Dana Reeve Foundation, which funds research into curing spinal-cord injury and provides information and advocacy services for victims, has given the project a $150,000 grant. The US National Science Foundation has awarded it $120,000 and plans to provide another $70,000 to $80,000 annually for three years.

Douglas S. Landsman, director of the Reeve Foundation's Individual Research Grants Program, said the Georgia Tech project is encouraging because the system promises to be fast, functional, and minimally invasive.

In addition, Landsman said, the tongue is a good controller because it moves quickly, offers fine motor control, and communicates directly to the brain and thus isn't affected by spinal-cord injuries.

The Georgia Tech approach could be an improvement over the puff-sip system, which lets disabled users issue commands to devices by inhaling or exhaling into a tube. This technique offers only four commands. And eye-tracking control systems can be costly, slow, and inaccurate, and can obscure the user's field of vision.

The researchers soon plan to test their approach on disabled individuals, with the assistance of a local rehabilitation hospital.

Ghovanloo said his team is trying to make its system lighter, more functional, and reasonably priced before licensing it to vendors for commercial distribution, perhaps by 2010.

Researchers have created the Georgia Tech Tongue Drive system, which could let severely disabled people use the inside of their mouth to issue commands to computers, wheelchairs, home appliances, or other devices. The system works via a magnet placed on the tip of the user's tongue. Magnetic sensors attached to a headset monitor changes in the magnetic field as the tongue moves around the mouth. The sensors send data to a receiver on the headset, which transmits the information wirelessly to a nearby computer for processing.

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