Editors: Elias Levy, aleph1@securityfocus.com
Iván Arce, ivan.arce@coresecurity.com

The Land of the Blind

And a thousand thousand slimy things lived on; and so did I.
— Samuel Taylor Coleridge, *The Rime of the Ancient Mariner*

We are in a tunnel. It's been rather dark, and will stay that way for quite a while.
— Hugh Hefner

Iván Arce
Core Security Technologies

A quick review of the articles in *IEEE Security & Privacy* since its launch in January 2003 reveals a wide range of software security topics. Several articles and departments delved deep into security technologies and software examining the security-oriented features and capabilities of the information security models or the fundamental building blocks used to create them. (A good example of this is "When Hashes Collide," the first installment of the magazine's newest department, Crypto Corner.[1]

Yet, we rarely read about security technology's strengths and weaknesses in the specific hardware and software products used in real-world environments. Isn't it curious that almost no published security product review or comparison explicitly assesses or provides any more than superficial details about products' security? How could this be—especially considering that so many current technologies are deeply rooted in cryptography, a field that derives its evolution from an iterative attack-and-defend, prove-and-refute process and whose practitioners accept and embrace the idea of breaking and reinventing the very systems, algorithms, and tools that result from long hours of hard work?

We can assume that cryptographers apply logic that's not only fundamental to modern scientific development but that's also been used for decades or, arguably, centuries. If that's true, we could justify the lack of scientific rigor in security products by noting the information security industry's immaturity or by including in the analysis other variables that drive its evolution. For such an analysis, we should at least consider the economic, cultural, political, and ideological backgrounds, as well as goals and motivations, of the organizations and individuals that shape the information security microcosm.

Bearing in mind that humans drive the discipline's evolution, and adding the weight of the author's own subjectivity to the analysis, you'll find this Attack Trends installment looks at the security vendor space—of which I am part—in search of new vulnerability types and attack trends.

Software security and security software

I've been told repeatedly since I was a boy that money doesn't grow on trees; although this idea originally seemed absurd and alien, I eventu-

ally came to grips with the realities of the modern economy and the value of currency—often in hard ways. A similarly outrageous theory would indicate that new vulnerabilities and attack trends don't grow on trees either; instead, they're the outcome of an evolving field driven by scientists, technologists, business people, and various other forms of intelligent (yet error-prone) humans with individual and collective motivations.

We can assess the status of the information security discipline by identifying, quantifying, and analyzing such relevant events as vulnerability discovery and disclosure, security incident reports, research and development advances in security technology, security service and product deployment, and the relevant regulatory and legal events that affect the field. Substantial variation in the quantity or quality of these monitored events over a period of time indicates a new trend. A notable qualitative exception to the overall set of monitored events can signal a new type of event.

We can efficiently process these cold, hard facts with a well-defined methodology, but the trend analysis is heavily influenced by a single variable: humans. The people who collect and process the information and draw conclusions throughout that process often include their environments, cultural backgrounds, experiences, motivations, preferences, goals, and values. To paraphrase the 20th century Spanish philosopher José Ortega y Gasset, "we are we, plus our circumstances." Therefore, observers might find it difficult to identify trends or new vulnerabilities in the information security microcosm if they're part of those phenomena: observations and analysis are necessarily tinted with human subjectivity.

Observers might find it difficult to identify trends or new vulnerabilities in the information security microcosm if they're part of those phenomena.

To avoid further epistemological diatribes, let's look at some interesting events observed in the information security vendor space that exemplify the mechanics of an ongoing search for new vulnerabilities and attack trends.

Hook

The outbreak of the Witty worm and the subsequent analysis of its distinguishing characteristics is a good example of a closely observed security event that indicates an exceptional and qualitatively unique data point among the set of monitored security variables.

The Witty worm began spreading on 19 March 2004 at 8:45 p.m. and reached its propagation peak of 12,000 infected hosts within 75 minutes. Researchers Colleen Shannon and David Moore at the Cooperative Association for Internet Data Analysis (CAIDA) provided an in-depth account and analysis of Witty's spread (www.caida.org/analysis/security/witty/) and outlined some of its unique features:

- Witty exploited a vulnerability in a widely deployed information security software product from a well-established and known security vendor.
- Witty started to propagate from a total number of compromised hosts that was an order of magnitude (around 100 hosts) greater than previous worms. This led analysts to conclude that Witty was deployed on those hosts prior to the start of its propagation.
- The worm began propagating within 48 hours of public disclosure of the vulnerability it exploited. To date, this represents the shortest interval between a vulnerability's public disclosure and its widespread automated exploitation.

Several other interesting characteristics of Witty are present in Shannon and Moore's report, but this list satisfies our selection criteria for a qualitative exception to our set of security events. Is this a new type of attack? Is it a new trend?

In the June 2004 issue of the Usenix magazine *;login:*, Nicholas Weaver and Dan Ellis provide their hypothesis of the worm author's goals and motives based on their analysis of the worm's code and the monitored events during its propagation (www.usenix.org/publications/login/200406/pdfs/weaver.pdf). They characterize the worm as a dangerous new trend that combines both skill and malice on the part of the attacker.

In a more recently published paper (www.icir.org/vern/papers/witty-draft.pdf), Abishek Kumar, Vern Paxson, and Nicholas Weaver combine a time-spatial correlation of monitored Witty worm events with a meticulous dissection of the code to reveal a seemingly extraordinary number of findings about the worm and the mechanics of its propagation. One of them is that a specific host among the entire IP address space (2^{32} unique addresses) could be identified as the worm's initial spreading point (*patient zero*). By analyzing the disassembled code of the worm's binary, identifying design and implementation flaws in the use of a pseudorandom number generator for its target selection, and matching the expected behavior of the code to the events observed during worm

propagation, the authors singled out a specific IP address that couldn't possibly belong to the set of IP addresses generated by the worm's code but nonetheless appeared to be spreading the worm with a traffic pattern distinctly different from the traffic behavior observed during the incident.

Beyond the great technical details and plausible motives attributed to the Witty worm attacker in these research papers, another interesting meta-analysis is worth noting: The characterization of the attacker's techniques, practices, and skills is quite similar to those of the researchers' own. The worm's author seems to have learned from previous worm incidents and refined his or her skills. The author followed a somewhat sophisticated development process, tried to avoid design and implementation problems, and used (possibly independently discovered) information about an exploitable vulnerability in a third party's security software package.

In turn, the researchers, armed with their own arsenal of skills, applied the same techniques and practices—analyzing a third party's software package (the worm's code), identifying an "exploitable" flaw in the program's code, and using it to fulfill their stated goal (enhancing our understanding of the security event, expanding our information security knowledgebase, and improving the chances of achieving a more robust security posture).

In this context, the extraordinary characteristics of the Witty worm outbreak and its analysis, along with the qualitative differences that singled it out among myriad monitored security events, are the natural but hardly surprising outcome of an evolving discipline in which all participants—the good, the bad, and the ugly—have demonstrated growth in practice.

Line

An eagle's eye view of certain monitored information security events within a certain time period can provide the raw material to extrapolate trends. Technology industry analyst firm, Yankee Group, recently presented such a report to the public (www.yankeegroup.com/public/products/decision_note.jsp?ID=13157). The report presents statistical data about the vulnerabilities disclosed to the public since 2001. The report then compared the resulting figures to the number of disclosed vulnerabilities in Microsoft's products during the same time period. By correlating this data, the report identifies a relative increase in disclosed vulnerabilities in security vendor products over the past year and the first quarter of 2005, thus extrapolating a new attack trend: an apparent shift in vulnerability research activity toward finding bugs in security vendor products.

The report also states that third parties, usually researchers affiliated with other security vendors or independent discoverers, found most of the disclosed flaws in security vendor products. Presumably, these attributions come from the finder "credits" section in the vendor's security advisories and patch notifications. This data helps explain the proposed trend as a result of two key factors:

- The depletion of easily identifiable vulnerabilities in Microsoft products due to the company's ongoing security push and the deployment of Windows XP Service Pack 2 (XP SP2). The easily identifiable vulnerabilities (the low-hanging fruit) remain present in widely deployed security products.
- The prevalence of individual and organizational self-interest as a guiding principle for vulnerability research. According to this premise, security researchers are eager to find vulnerabilities in security products due to an "adolescent enthusiasm" to one-up their peers. Similarly, security vendor organizations are economically motivated to find bugs from competing organizations, which would provide competitive advantages and enable them to obtain increased market share at the expense of vulnerable vendors.

Despite the possible inaccuracies, fallacies, and logical pitfalls of this simplified analysis, and in the absence of methodological details to support the research results, the final conclusions and recommendations seem a lot more reasonable and well within the realm of common sense. Simply put, vendors should adopt information security best practices in their product development life cycle.

Security vendors are prone to errors, yet their products are widely deployed by security-conscious organizations. In the presence of bugs, the strength of the increasingly ubiquitous nature of information security technology could become a serious weakness in organizations' security postures—a reality that won't go unnoticed by today's security researchers and practitioners, whatever their goals and intentions. Accordingly, organizations deploying security products should exercise caution, require due diligence from their security providers, and mitigate risk by diversifying and mixing the sources of the security technologies they deploy.

Security vendors are prone to errors, yet their products are widely deployed by security-conscious organizations.

And sinker?

So far, I've presented two interrelated yet distinct information security events that we can use to extrapolate an emerging attack trend:

- The Witty worm's appearance and its subsequent study signals a singular, qualitative exception to all previous intelligence on malware dissemination (it's been viewed from multiple angles, such as the attacker's sophistication, payload semantics, elapsed time from disclosure to massive exploitation, and propagation techniques).
- A cross-section of vulnerabilities disclosed in the past few years has elicited a new real or perceived trend of targeting security products for attack.

The discovery and disclosure of vulnerabilities in security vendor products isn't a new practice.

Does the combination of these two events provide sufficient evidence to single out a new attack trend? Or is the perceived new trend the natural consequence of an ongoing process in the information security discipline that can be reasonably explained with additional information from a wider time span?

To delve a little deeper, let's examine a few more security events that seem similarly relevant. Although their statistical value is negligible, their qualitative characteristics could prove them to be more than just anecdotal evidence.

The revenge of the nerds

The discovery and disclosure of vulnerabilities in security vendor products isn't a new practice. In 1998, Thomas Ptacek and Timothy Newsham discovered and disclosed fundamental design flaws in network intrusion detection systems ("Insertion, Evasion, and Denial of Service: Eluding Network Intrusion Detection;" www.insecure.org/stf/secnet_ids/secnet_ids.html). In 1996, Peiter "Mudge" Zatko, a researcher at BBN, reported similar flaws in SecureID, the industry-leading token-based secure authentication product from Security Dynamics (www.tux.org/pub/security/secnet/papers/secureid.pdf). Adam Shostack, another security researcher and entrepreneur, exposed these flaws further at the 1996 Network Threats workshop at the Center for Discrete Mathematics and Theoretical Computer Science (DIMACS; http://dimacs.rutgers.edu/Workshops/Threats/Shostack.ps).

Simpler software implementation bugs in security vendor products aren't news either, as Ptacek and his fellow researchers demonstrated in 1998 (http://cert.uni-stuttgart.de/archive/bugtraq/1999/02/msg00315.html). Thanks to Thomas Lopatic, Dug Song, and John McDonald's presentation at the Las Vegas Black Hats Briefings Conference in 2000, firewall vendors quickly discovered that security vulnerabilities weren't necessarily prevalent in single segments of the security vendor space (http://blackhat.com/presentations/bh-usa-00/Song-McDonald-Lopatic/Song_McDonald_lopatic.ppt). The vulnerabilities they found in one firewall's stateful packet inspection technology later proved to also apply to other security vendors (www.securityfocus.com/bid/979).

The SSH history

The secure shell (SSH) program was originally conceived as a replacement for a client-server application used to login remotely to Unix systems. The standard Unix programs used for this purpose, `rlogin` and `telnet`, transmitted user authentication credentials and all session data in the clear over the network. An attacker with access to the network path between the client and server could capture and reuse the authentication credential to obtain unauthorized access to the remote systems or to eavesdrop on the session data and compromise the communication's privacy.

SSH provided encrypted network communications and cryptographically strong authentication mechanisms to replace insecure Unix programs and quickly became the de facto standard for secure remote login among system administrators and security-conscious users. It's debatable if SSH is a security application or simply an administrative application with the necessary security technology built-in, but both software and security vendors rapidly adopted it. Even giant software and networking vendors such as Sun Microsystems and Cisco adopted SSH as a default administrative interface.

Finnish developer Tatu Ylonen, founder of SSH Communications Security, made the first SSH implementation available for free in 1995 (www.ssh.com/company/newsroom/article/650/). Security vendor firm F-Secure (known then as Data Fellows) also sold the application as a commercial package through a partnership with SSH Communications. In 1999, OpenSSH—a new implementation of the program based on the original `ssh` version 1.21 release—appeared as an open-source project sponsored by developers from the OpenBSD operating system project (www.openssh.org/history.html).

In 1998, Ariel Futoransky and Emiliano Kargieman, security researchers and cofounders of Core Security Technologies (a security vendor firm), discovered a serious design flaw known as the "CRC insertion attack" (www.coresecurity.

com/files/files/11/CRC32.pdf) in the SSH protocol and notified Tatu Ylonen, the program's lead developer at SSH Communications (www.coresecurity.com/common/showdoc.php?idx=131&idxseccion=10). Fortunately, Futoransky and Kargieman also produced a patch to prevent exploitation of the flaw and submitted it to Ylonen. The patch became part of the SSH package and continued to be used as an SSH version 1 protocol fix by all vendors in almost all existing `ssh` implementations. Two years later in February 2001, the Polish security expert Michal Zalewski identified an exploitable integer-overflow vulnerability in Futoransky and Kargieman's original `ssh` patch (www.bindview.com/Services/RAZOR/Advisories/2001/adv_ssh1crc.cfm). Shortly afterwards, a patch to Futoransky and Kargieman's patch was produced and included in the SSH distribution to prevent exploitation of the newly discovered bug.

Almost seven months later, David Dittrich, senior security engineer at the University of Washington, discovered an exploit (SSH x2, attributed to the TESO hacker group) being used in the wild to compromise systems running vulnerable SSH implementations (http://staff.washington.edu/dittrich/misc/ssh-analysis.txt). The inclusion of the exploit as a component of an automated compromise tool (a mass-rooter) seems to be a case of automated widespread exploitation of a bug in security vendor code that predates the Witty worm's spread.

Scientists, technologists, and businesspeople

The SSH story and other assorted anecdotes help describe the information security discipline's development over the past 10 years, but the individuals and organizations who contributed to this evolution acted in varied roles and adhered to multiple guiding principles, as most humans do. As scientists, technologists, and businesspeople, they've strived to fulfill their goals using their skills, values, and inspiration. They most likely faced many technological, scientific, and business pitfalls, hopefully learning from their own mistakes and various historical lessons along the way. But even if they did, the results of their collective effort are certainly far from perfect, and it's increasingly evident to technology users.

Within the same basic 10-year time interval, the information security industry has exploded at a frantic pace. Through mergers, acquisitions, spin-offs, and the cyclic emergence and disappearance of tenths or hundredths of information security start-up companies, close to 700 security vendors have had and still struggle to get a grip on a moving target estimated to be worth US$16 billion by 2008 according to Morgan Stanley Research analysts Peter Kuper and Brian Essex.[2]

To complicate matters further, the importance of addressing information security problems in a timely and effective manner is an increasingly primary concern for many organizations. The adoption of information security technologies and the deployment of security vendor products have also increased significantly, thanks to several disastrous incidents, a growing general concern for security, and a complementing set of compliance requirements and regulations.

Perhaps the Witty worm incident and the statistical analysis of the disclosure of security vendor vulnerabilities in the past three years are just symptoms of the existing phenomena of an industry's rapid growth with foundations in a discipline that is still far from maturity, rather than the hint of an emerging attack trend.

According to Greek mythology, the Cyclopes were strong, powerful, one-eyed beings who built well-crafted weapons and indulged in other blacksmith arts. After their brother Kronos imprisoned them, they invented the thunderbolt for Zeus so that he could overthrow Kronos and rule heaven and earth. In exchange, the Cyclopes were freed. Their story ended sadly, though: Apollo killed them after learning that Zeus had used a Cyclopes-forged thunderbolt to kill his son. Greek mythology is quite allegorical, seriously convoluted, and certainly off-topic for a department in *IEEE Security & Privacy* magazine, except for that distinguishing Cyclopes characteristic: one eye.

As information security practitioners, we like to think we've been imprisoned for years in a dark land populated with seemingly blind people. By using just one eye, we've managed to build ingenious devices, but now we've found our way into the open, and this vast new land requires us to use both eyes, right now. If we are to avoid a tragic ending, we (as well as our discipline and industry) must step up and mature very rapidly. □

References

1. P. Gutmann, D. Naccache, and C.C. Palmer, "When Hashes Collide," *IEEE Security & Privacy*, vol. 3, no. 3, 2005, pp. 68–71.
2. P. Kuper and B. Essex, *Data—The Next Perimeter of Defense*, Morgan Stanley Research, Jan. 2005; http://investext.com.

Iván Arce *is chief technology officer and cofounder of Core Security Technologies, an information security company based in Boston. Previously, he worked as vice president of research and development for a computer telephony integration company and as information security consultant and software developer for various government agencies and financial and telecommunications companies. Contact him at ivan.arce@coresecurity.com.*