Intrusion Detection… Or Prevention?

David Piscitello

Vendors are refining security products that "sense" attacks. But what if software were made less vulnerable in the first place?

Attacks against computers and networks have escalated steadily in number and sophistication, to the point where many security experts believe the situation is beyond "out of control." Statistics compiled by the Computer Emergency Response Team Coordination Center (CERT/CC) show a more than five-fold increase in security incidents reported between 1999 and 2001 (9,859 to 52,658).

Some of this increase can be attributed to a welcome appreciation among organizations that reporting incidents does more good than harm—CERT, among others, assures the confidentiality of victims. Nevertheless, the figures are staggering.

Attackers have become more efficient in recent years. Freely available tools automate the processes of information gathering, target acquisition and vulnerability isolation, and they are becoming increasingly stealthy as well: While many security administrators will quickly identify novices or "script kiddies" by the clumsy methods they use to probe computers and networks for vulnerabilities, it will take talented administrators to identify the new-millenium attacker.

This more advanced attacker goes to elaborate lengths to study, then circumvent intrusion detection systems (a tactic called intrusion detection "detection") and to defeat other security countermeasures. The current state of security has been characterized by one self-described hacker—"Simple Nomad," organizer of the Nomad Mobile Research Centre (www.nmrc.org)—as a game of Network Cat and Mouse.

David Piscitello, president of Core Competence, Inc. is an internationally recognized expert in security technology and founder of the Internet Security Conference. He has chaired the Networld+Interop Conference Program committee since 1996.

Intrusion Detection

Intrusion detection systems (IDSs) rank among the most highly publicized elements of network security. But until recently, IDSs have failed to deliver the lofty results vendors promised. Organizations that have deployed IDS complain most about complexity, completeness (attack coverage) and inaccurate detection.

Many organizations find they do not have sufficient expertise in-house to configure the IDS they've purchased, or to interpret and take appropriate actions based on the alerts the IDS generates. Too often, a misconfigured or inappropriately placed IDS results in waves of false positives and undetected actual attacks. IDSs also have been criticized for poor performance.

Perhaps the biggest disappointment lies in the inability of most IDSs to do more than report attacks. The information contained in IDS alerts in and of itself is useless to all but serious security experts, and organizations quickly find that incident correlation and real-time response are beyond their expertise.

Organizations that believe they can't do without IDS but can't manage the systems themselves are turning to managed IDS services from companies like RipTech, Counterpane and SecureWorks. These and other managed service providers use the same IDS products, but have the core competencies to deploy them more effectively.

What's more, attackers have successfully thwarted certain IDSs by using slow, stealthy probes, and by coordinating an attack across multiple sources. "Slow, stealthy" attacks exploit the fact that signature-based IDSs look at a sequence of packets to see if that sequence matches a pattern of a known attack. In a "slow" attack, the packets aren't sent all at once, but over a longer time period, so that the IDS doesn't detect the attack pattern. Similarly, a "stealthy" attack may inject legitimate traffic between the attack packets, hiding their signature from the IDS.

In some cases, the IDS itself is attacked, either as part of a deception or to force it offline. Attackers have been known to download evaluation copies of IDS software, then reverse-engineer or study the software to learn how the IDS reacts to traffic patterns. Armed with this information, they can alter their attack patterns to circumvent the IDS (another example of Network Cat and Mouse).

Next-Generation IDS

The latest generation of network intrusion detection systems, from companies like OneSecure, Tippingpoint, NFR Security, MazuNetworks and IntruVert, promises measurable improvement. IntruVert, for example, applies stateful inspection to its signature-based attack detection, and complements this with traffic anomaly detection, a real-time comparison of network traffic against a baseline of "routine" or normal traffic to detect unusual and potentially harmful traffic.

IntruVert's detection techniques use state information within the data packets and seek out multiple token matches to identify attack signatures that span multiple packets or arrive as an out-of-order packet stream. Likewise, OneSecure also incorporates stateful analysis with its signature detection, and anomaly detection to identify new attacks as well as attacks that span multiple sessions. The IDS systems from IntruVert and OneSecure also incorporate denial-of-service detection, i.e., the identification of an unusually large volume of seemingly normal traffic designed to sap network resources from legitimate traffic and thus deny service to rightful users.

Increasingly, the IDS sensors—i.e., appliances that tap into networks to examine and in some cases block traffic—are hardware accelerated. NFR Security and OneSecure appliances operate at 100 Mbps, while TippingPoint's appliance and IntruVert sensors promise 1 Gbps support.

When next-generation IDSs are deployed as in-line instead of passive monitoring appliances, they will drop traffic identified as intrusive. This is often labeled intrusion prevention, but a more accurate term would be attack blocking or intrusion rejection: The vulnerability still exists, but the in-line IDS is able to block the attack.

That's no small accomplishment, but an even more secure system would be one in which the vulnerability didn't exist in the first place. Achieving this higher level of security requires tighter integration and greater scalability.

Intrusion Prevention

The most frustrating thing about network security today is that the canonical security posture is entirely defensive and largely reactive. IDS is more a testimony to how poorly we have written software and how feebly we manage systems and networks than how clever we are at repelling intruders. We've created an entire industry to protect us from the bad code we passively accept from software vendors and the shoddy practices our budgets and "Internet pace" foster as well.

And the problem is getting worse. CERT/CC statistics reveal nearly six times as many vulnerabilities (flaws in software that an attacker exploits to gain control of or misuse a computer) reported in 2001 than 1999 (2,437 to 417). In the first two months of 2002 alone, CERT reported multiple vulnerabilities in SNMP, a buffer overflow vulnerability in Microsoft Windows UPnP, vulnerabilities in SSH implementations and the W32/Bad-Trans Worm. Thumb through CERT's Vulnerability summary reports or search the Bug-Traq mailing list archive (http://msgs.securepoint.com/bugtraq/) and you'll see that it's not simply Microsoft: Cisco, Oracle, IBM, CheckPoint, Novell, Netscape, commercial versions of Linux—all have been proven exploitable.

In practice, intrusion prevention won't be a popular notion

We need to move from reactive intrusion detection to a proactive stance of intrusion prevention. That might sound like a concept everyone can agree with, but in practice, intrusion prevention—in the form of demanding software that's not easily exploited—won't be a popular notion. It flies in the face of the "as is" software licenses we all accept with a mouse click, daily. Intrusion prevention in its most basic form means vendor quality-assurance programs and third-party, independent source code review performed to assure that code is not exploitable due to logic or coding errors.

To illustrate the power of this approach to intrusion prevention, imagine that a router vendor has prevented buffer overflow attacks from occurring in the first place, because it has eliminated the very possibility of the exploit by the thorough analysis of the source code *before* shipping it to a user/customer. This is of course more expensive for the vendor than employing the user base as beta testers. You do get what you pay for.

Intrusion prevention also involves careful configuration of operating and file systems. Regrettably, the cardinal rule of security, *that which is not expressly permitted is prohibited*, is largely neglected today. Products ship with far too many default settings that are excessively permissive.

Under pressure to get software and equipment in production, harried administrators are unlikely to read documentation thoroughly enough to appreciate the pitfalls that lie ahead—like the fact that it's possible to execute arbitrary shell commands (DOS or *NIX) from SAP/R3. An administrator is equally unlikely to ferret out every default Oracle account and password before putting a sensitive database online. But these defaults are easily acquired from hacking websites. Clearly, if administrators don't have time—and training—to be diligent with configurations, it's hardly a surprise that less technical employees rarely consider such subtleties when they install new software.

Even with substantial improvement in secure code development, vulnerabilities will be revealed. Vendors will provide patches, hot fixes and new releases, but organizations *must* improve their cumulative track record for applying them. The vast majority of successful attacks exploit known vulnerabilities, for which the vendor has already released a patch.

Had A Security Physical Lately?

Curtis E. Dalton

Intrusion prevention is vital at the networking level, but it's equally important for companies to prevent physical encroachments on their network assets. Unfortunately it took September 11 to remind many organizations of the need for adequate physical security.

If you conduct a physical security assessment of your organization, you will almost certainly be surprised to learn how incredibly vulnerable you are. While many security practitioners were busy tightening down electronic access over the last few years, many continued to leave the front door, back door and first floor window open.

The purpose of physical security is to extend the umbrella of security over those resources that cannot be protected by purely logical means. Resources to protect might include anything from laptops to personnel.

While the protection mechanisms may vary, the security fundamentals remain the same. Physical security begins just outside your facility, and never stops from there inward. If your facility's business is to guard something confidential, don't advertise it along the rooftop over the main entrance (e.g., "BIG ASP Data Center: Home of many financial institution computer systems"). Should potential customers wish to visit your facility, there are better ways to make sure they can find it.

The importance of protecting entry and exit points is of no news to anyone; *how* they are protected is crucial. Are employees able to simply walk in the front entrance, or must they present or wear identification badges? How closely do guards inspect these badges—from 40 or 50 feet away and partially obstructed by the employee's overcoat through a thick crowd of people entering and exiting?

Or maybe employees must present their identification badges at the guard desk for entry and exit into and out of the facility—but does anyone actually read these logs to see if the same person has entered twice in the last two hours and hasn't departed in over a week? Also, what level of difficulty is required to counterfeit these identification badges? Typically, the answer is: Not much. It's a widely known fact that millions of 19- and 20-year-olds are able to counterfeit driver's licenses—or more appropriately, drinking licenses. So it's likely that if a teenager with $50 can get a fake ID, a would-be intruder to your facility can, too.

TABLE A: Basic Physical Security Requirements

- All assets have been identified and classified as to security level with a list of personnel authorized to access those assets.
- Proper change control practices exist (adding, altering or removing personnel access to facilities).
- There is adequate guard presence at key locations and revolving roving patrol schedules and routes that take guards throughout the facility.
- Redundant power exists (with separate facility ingress and egress points).
- Redundant communications links exist (with separate facility ingress and egress points, and separate communications offices for each link).
- Restricted access to sensitive equipment (sensitive computer systems, UPS and generator power, communications equipment, alarm systems).
- Adequate fire and smoke detection and suppression.
- Adequate leak or water detection and suppression.
- Adequate interior and exterior surveillance.

Many companies could realize an immediate improvement in security if they were to more prudently assign access privileges to users, and make use of stronger authentication methods. Too many users have administrative privileges on their own systems, which on poorly secured LANs readily extends to other desktops and servers. File and printer sharing is commonly permitted without authentication. Web servers permit directory browsing. Password composition and expiration rules are largely ignored, and the persistent reliance on passwords is itself a major problem, but one easily corrected with token-based security products, certificates, biometrics or combinations of such authentication methods.

Preventative Measures

Organizations with significant purchasing power *can* influence the "software as is" situation. Just as these organizations negotiate service level agreements with telephone, cellular and Internet service providers, why should they overlook the opportunity to negotiate software confidence (i.e., security/reliability) agreements?

Vendors should rethink the tradeoff between convenience/ease of use and security, and should minimize default settings that favor the former at the expense of the latter. A first small step might be to ensure that any software configuration is deemed incomplete until all default passwords have been changed.

In addition, organizations should invest in training. Software and network configuration are complex tasks; secure configuration of software and networks is doubly demanding. It's unreasonable to expect an administrator to absorb user and administration guides of 800–1,200 pages prior to deploying, say, a Certificate Authority (and indeed, one major PKI vendor's documentation is this size).

Training not only provides a fast start for the administrator, but connects him/her with a community of individuals with similar job descriptions, creating a resource to counterbalance the community that the attackers maintain. A week of training will cost an organization under $10,000; Cahners In-Stat Group reports that an average network security breach costs $259,000. Do the math.

The security you require of entrants at the front door should differ from that of those wishing to enter more critical areas of your facility, such as the datacenter or network operations center. Obviously, you should restrict unauthorized access to sensitive areas within or about the facility that provide power, communications, data access or storage; monitoring, access and authorization controls (keys, access logs, etc.), alarm notification and surveillance systems should likewise be restricted. A basic list of physical security checklist requirements is detailed in Table A.

Even with all the high-tech electronic gadgetry, physical security mechanisms are useless without adequate support from humans. This support requires diligence and training. In most cases, a diligent eye towards security will return far more than electronics and padlocks. This isn't to say that high-tech access control and surveillance electronics, sturdy computer cabinetry and padlocks are not needed, but rather that these things should be used to enhance your security, rather than representing your sole solution.

For security staff training to be beneficial, it should be as realistic as possible. If you are going to test procedures and protection mechanisms against a perimeter breach (unauthorized person entering the facility), stage a simulated breach with staff acting in various roles and record the events that follow. Management should then review what transpired during the exercise, noting both strengths and weaknesses so that future efforts can be tailored to strengthen deficient infrastructure, practices or staff training in particular areas. Physical security must be closely supported by solid operations practices to be effective.

Of course, security will almost always take a back seat to business productivity, since there's no need for corporate security if the corporation goes out of business. Also, there will always be a trade-off between security and productivity. As you increase one, the other will be decreased. Balancing the two in a manner that best suits your organization is the task at hand□

Curtis Dalton is a managing partner of Principal Security Group, LLC, an information security consulting firm based in the Boston area. He can be reached at cdalton@psgsite.com

- Use of signage only for required purposes.
- Adequate locking computer cabinetry.
- Adequate backup power (diesel generator) that is regularly load-tested.
- Adequate perimeter access controls (walls, locked doors, fences).
- Adequate interior access controls for access to secured floors, network or data storage closets and datacenter.
- Use of employee picture IDs for entrance.
- Use of biometric devices plus password (PIN) for access to highly secured areas.
- Visitor restriction policies such as escort required and visitor badges that expire.
- Known local police and fire response times for contingency planning.
- No opening or unlocked doors (other than guarded entrances) or windows on the first floor.
- Adequate Radio Frequency blocking within datacenters and network operations centers to prevent unauthorized outbound or inbound cellular or wireless LAN communications.
- Secure on- and offsite data storage and retrieval practices.

Organizations also should assign responsibility for vulnerability awareness. This individual(s) should be responsible for monitoring CERT, BugTraq and other security lists for advisories regarding newly-uncovered vulnerabilities in software that the organization runs. Vendors should willingly disclose vulnerabilities to registered users as soon as they are revealed (another clause that belongs in the software confidence agreement). The vulnerability awareness individual or team should rapidly respond to new vulnerabilities by acquiring the necessary patch or work around, testing it, then distributing or installing it.

Conclusion

Intrusion prevention is more like a vaccine than an antibiotic (i.e., anti-virus). Vaccines make the human body resistant to viral attacks by enhancing the immune system. Similarly, intrusion prevention measures make your systems and networks resistant (immune) to certain attacks. Intrusion detection is sexy and absolutely necessary for certain kinds of attacks, but over the long haul, intrusion prevention will better serve an organization□

Companies Mentioned In This Article

CERT (www.cert.org)

CheckPoint (www.checkpoint.com)

Cisco (www.cisco.com)

Counterpane (www.counterpane.com)

IBM (www.ibm.com)

IntruVert (www.intruvert.com)

Mazu Networks (www.mazunetworks.com

Microsoft (www.microsoft.com)

NFR Security (www.nfr.com)

Novell (www.novell.com)

OneSecure (www.onesecure.com)

Oracle (www.oracle.com)

RipTech (www.riptech.com)

SecureWorks (www.secureworks.com)

Tippingpoint (www.tippingpoint.com)