

Protecting the Crown Jewels

An Enterprise-Class Approach to Application-Level Security

INTRODUCTION

Hackers tend to go where the targets are the most attractive, and the defenses are the weakest. Therefore, it shouldn't be surprising that enterprise applications and databases are increasingly coming under attack from the kind of threats once associated mostly with operating systems and desktop applications.

Most large organizations have already installed antivirus software, firewalls and even intrusion detection systems (IDSs) to protect their networks and host operating systems. But by comparison, enterprise-class applications have received relatively little attention, on the assumption that they are protected by firewalls and other defenses at the network perimeter. Yet these applications and databases are the major reason enterprises invest in IT in the first place, and the data they contain are often the enterprise's most valuable assets. Indeed, an enterprise without database security is like a bank with locks on the doors and armed guards by every entrance, but no vault.

Though a critical component of a layered defense, firewalls cannot detect and stop the new class of threats now being directed at applications and databases. Another widely deployed tool, intrusion detection systems, perform only passive monitoring and after-the-fact forensics rather than preventing attacks. Indeed—the Gartner Group recently highlighted the limitations of these measures:

"...most organizations have learned that perimeter firewalls, antivirus software, and intrusion detection systems are not enough to protect them from cyber attack. Attacks have moved to the application level, circumventing network-based firewalls. Worms propagate so quickly that signature-based antivirus protection is useless. Intrusion detection systems do not provide protection, only faster notification that your security has failed. The ideal form of protection requires hardened, locked-down server and desktop configurations!..."

Organizations need to bring the same level of protection to applications and databases as they apply to servers and networks, with solutions that can automatically detect and respond to application-level threats in real time, and that are granular enough to provide access for customers and business partners while keeping attackers out.

THE NEW TARGETS: APPLICATIONS AND DATABASES

In the last several years, there has been a substantial growth in potential vulnerabilities as well as actual attacks on applications and databases.

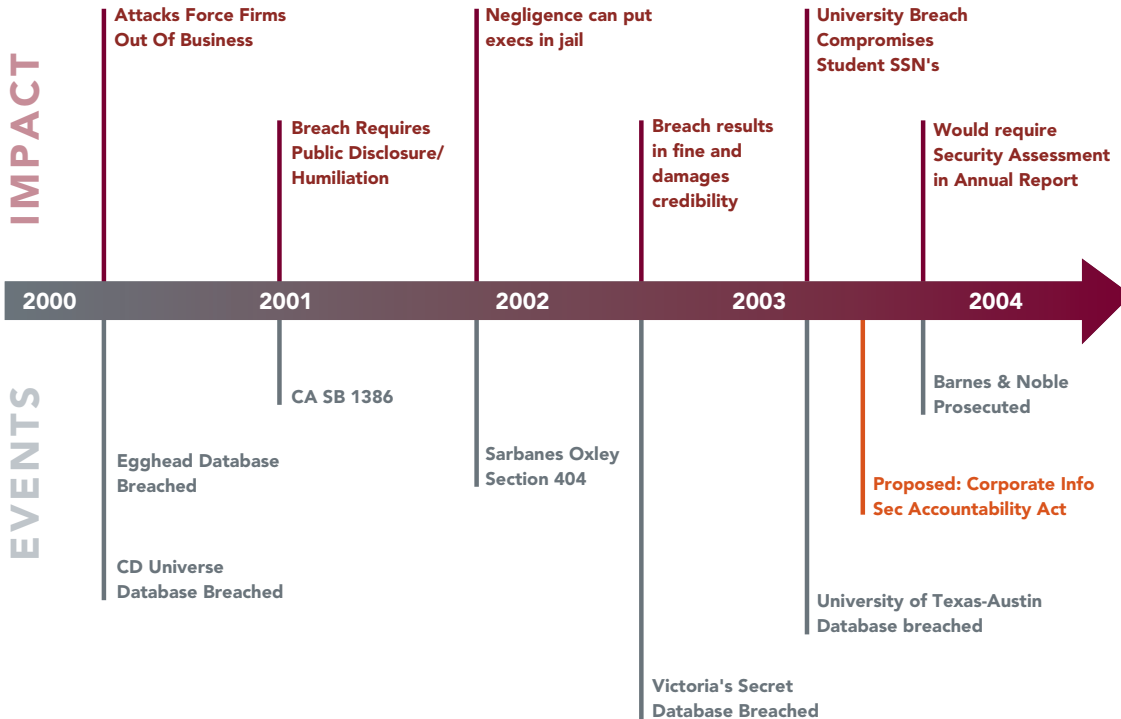
Many of these threats are a result of the changing nature of enterprise applications and databases. A decade or more ago, databases were usually kept physically secure in a central data center and accessed mostly by applications within the corporate borders. Today, though, applications and databases may be distributed in business units to meet local needs. Even more critical is the fact that these applications and databases are increasingly made available to suppliers, customers and business partners in order to conduct business over the Web. It is a business imperative today that, for example, suppliers can check their customers' production schedules to coordinate just-in-time deliveries of raw materials, and consumers can pay bills and check their cell phone usage on-line. But with this increased access comes increased risk.

In 2002 alone, Symantec documented more than 65 vulnerabilities affecting database products from Microsoft, Oracle and

IBM. Microsoft issued 11 security bulletins for SQL Server 2000 and 7.0 in 2002, while Oracle published 20 security alerts. This pace of security alerts shows the dynamic and ever-changing nature of application and database threats, which requires organizations to respond with more than just annual or semi-annual audits.

Many of the new threats take advantage of the fact that today's databases are not mere repositories for information, but robust development environments that allow developers—and hackers—to carry out complex functions within the database. In one common form of attack, such as a SQL injection, a hacker uses the SQL database-access language to insert into a database not legitimate business logic, but malware designed to infiltrate, corrupt or gain illegitimate access to the database. Another common attack is the buffer overflow, in which a hacker inserts more data into a buffer (temporary storage area) than the buffer was designed to hold. The extra data can corrupt the legitimate data in the buffer, or in adjoining areas of memory, or contain instructions that allow illegitimate access to information.

DATABASES ARE UNDER ATTACK



This time line shows a handful of well published breaches, which have resulted in a variety of legal measures, not to mention the dire consequences suffered by the firms targeted by the attack.

Another threat comes from application worms, which are automated, self-propagating attacks on the custom code written for many Web applications. Application worms take advantage of publicly available Web indexes to find sites to attack, and to determine how best to attack them.

Examples of application and database attacks are not hard to find. In March of 2004, hackers downloaded 55,200 names and Social Security numbers from databases at the University of Texas at Austin—the second such incident in six months. The same month, San Diego State University reported hackers had broken into a server and had access to the names and Social Security numbers of more than 178,000 former and current students, applicants and employees. Also within the same month, BJ’s Wholesale Club reported the theft of tens of thousands of customer credit card numbers from a corporate database. This followed reports in the fall of 2002 that personal information about customers at the Victoria’s Secret Web site, including their name, address and the items they ordered, could be accessed through the company Web site—a breach that led to a \$50,000 fine from the state of New York and an agreement with the state that it improve its information security policies.

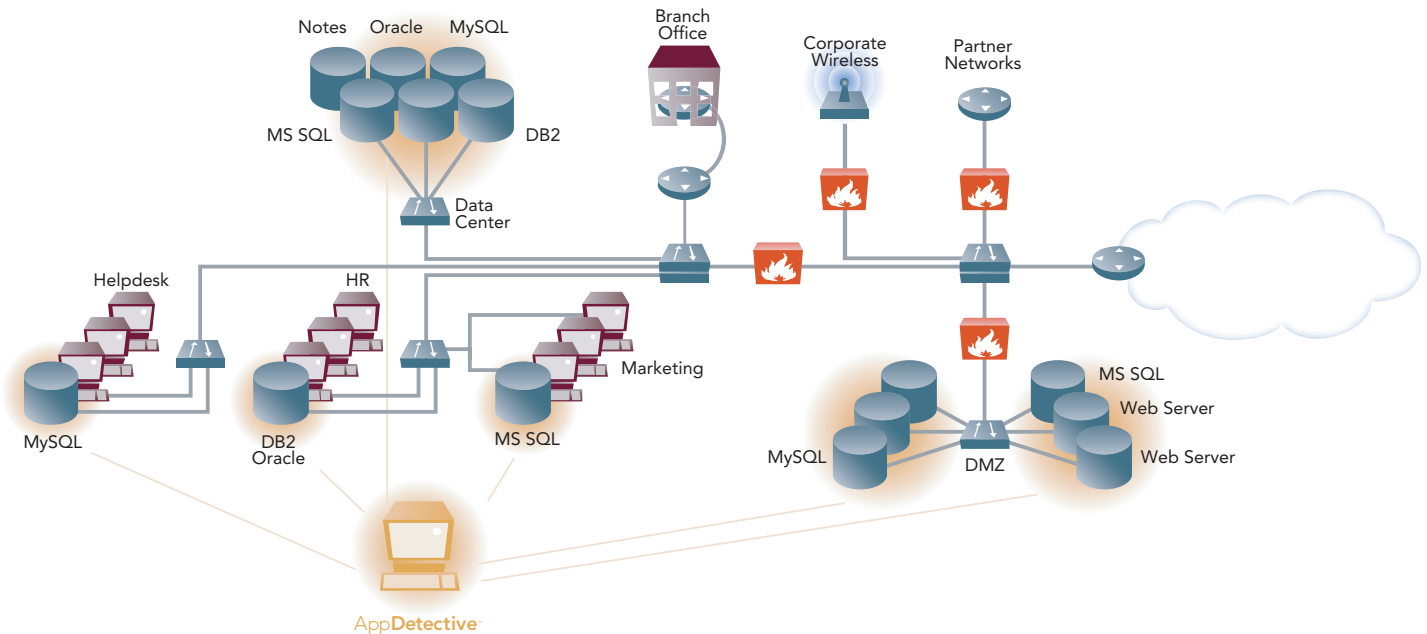
The threat of such costly and embarrassing penalties is becoming greater as state and federal governments enact new information privacy and security laws that touch virtually every industry and market sector. California SB 1386, passed in the fall of

2002, requires companies or individuals that store any personal information about California residents to notify those residents if any information about them has been stolen or accessed by a hacker. At the federal level, the Sarbanes-Oxley Act requires (among other things) that officers of publicly traded companies vouch for the processes behind the creation of their financial reports, which includes an assurance that financial information was not corrupted and did not fall into the wrong hands before its release to the public. HIPAA (the Health Insurance Portability and Accountability Act) requires health care providers to safeguard patient information, while the North American Electric Reliability Council (NERC) at the behest of federal regulators has adopted strict rules to protect electric utilities’ computers, software, and networks from intrusion.

SHORTCOMINGS OF NETWORK/HOST PROTECTION

The existing defenses most organizations have in place for their servers and networks are not designed to detect application-level attacks, nor are they designed to stop such threats before damage is done.

Firewalls provide protection only at the network level—examining packets and determining whether an incoming request should be given access to a given port. They do not understand database vulnerabilities or protocols (such as SQL) that may be used by attackers.



Though necessary, perimeter security solutions like a firewall do not recognize database security measures, they simply decide if a given user can use a given port. They have no basis to determine if that user executes a buffer overflow, has guessed a default password, or is injecting dangerous SQL commands.

As a result, most firms are creating a layered defense that includes direct protection of the database.

For example, a firewall examining the flow of packets to and from specific ports has no way of detecting if a user is trying to execute a buffer overflow, has guessed a default password or is injecting dangerous SQL commands into the database. Nor can it identify, for example, a “dictionary” attack that tries thousands of common passwords to gain illegitimate access to a database, or to issue an alarm when a user upgrades their own access privileges (one sign a hacker may have given himself the power to corrupt or steal from a database). Firewalls also do not understand the specific weaknesses that can affect individual database platforms, such as Oracle, Microsoft’s SQL Server and IBM’s DB2. Oracle, for example, ships with 15 default accounts and passwords, each of which can provide an avenue for attack if not properly secured.

Firewalls are updated by their vendors to protect against new network-level threats, but not against new security vulnerabilities discovered by database vendors. They are also typically located on the edge of the network, where they are ideally situated to watch for attacks from outside the enterprise, but not from insiders—which some experts estimate account for three quarters or more of all hacking attacks. Even when watching for outside threats, firewalls cannot inspect increasingly common types of traffic such as SSL (Secure Socket Layer) encrypted packets, or the XML (Extensible Markup Language) traffic used in Web services.

Finally, in a modern enterprise, firewalls simply have to let too much traffic through to provide foolproof application protection. In a world of virtual organizations and electronic commerce, an enterprise cannot afford to completely lock out customers, suppliers, distributors, remote employees or contractors.

Similarly, though many enterprises have deployed intrusion detection systems (IDS) to improve network security, these too do little to protect core databases and applications. Such systems scan the network, comparing traffic and usage patterns to either historic trends or against the “signatures” of known network attacks. However, most IDSs are passive, scanning for suspicious traffic and alerting the network administrator, but not taking any action to stop the attack. They are also designed as forensic tools, gathering evidence to analyze an attack after the fact rather than stopping it in real time.

Firewalls and IDS each have a place in a multi-layered security system. But they are not enough to protect organizations from internal and external threats while allowing appropriate access

to applications and databases. The modern enterprise needs application-intelligent equivalents of its existing network and host-based security platforms, which can discover, assess and dynamically protect applications and databases against rapidly changing security threats.

REQUIREMENTS FOR ENTERPRISE-CLASS APPLICATION SECURITY

What capabilities, then, are required to provide true security for the application layer? For a proven framework, look no further than the methodology organizations have already successfully applied at the network and host operating system levels. Just as at the host and the network perimeters, application-aware security solutions must provide vulnerability assessment, real-time intrusion protection and audit, and encryption. To achieve these goals, such application-level tools must provide:

Audit/Proactive Hardening: The system must audit the status and configuration of all application components and perform security tests and proactive hardening of such components while producing detailed security audit reports before and after application deployment. It must also ensure all current patches have been installed; default passwords have been changed; and recommended security configurations (such as changing the default ports on which applications run) have been implemented. As with the network and host OS, assessing the vulnerability of application components is the bedrock upon which any security strategy is built. Without it, an enterprise cannot either proactively minimize risk or gauge ongoing compliance with its security policies.

Real-Time Protection: The ability to detect and block attacks as they happen. Not only are more hackers creating more attacks than ever, but the malware they create is spreading more rapidly than ever. In addition, “the time between the disclosure and widespread exploitation of a vulnerability continues to shrink. Between the announcement of a new vulnerability and the development and deployment of a patch, companies are open to attack.²” Given today’s rapidly propagating threats and the time needed to deploy patches, organizations require real-time protection to complement the proactive hardening provided by ongoing vulnerability assessments.

Attacks can begin at any time. One growing threat is from “zero-day” attacks that target vulnerabilities before their existence is published and before patches are available for them. “If such an outbreak occurs, widespread damage could occur before users are able to effectively patch their systems.³” This

points up the need for behavioral-based intrusion prevention systems that can detect, and block, application-level attacks for which there is no known signature to scan for, nor any patch to apply.

Not all security threats are created equal. Some will pose more severe threats than others; and some threats will be of greater danger to some types of organizations than others. For this reason, administrators must be able to tune their response to the danger posed by the security threat for their specific enterprise.

Encryption: The ability to encrypt the most sensitive data as a “last line of defense” even if the database itself is compromised without incurring the overhead or complexity of encrypting the entire production database. Selective encryption also prevents unauthorized access to data by legitimate users. For example, a database administrator needs administrative access to the application in order to grant, revoke or change users’ access rights, but should not be able to see, change or copy the actual information in the database, such as customers’ credit card numbers. Any such encryption solution must be transparent to the application components it protects, meaning that the encryption will still function even as needed changes are made to individual components.

Internal and External Protection: The ability to detect and protect against application or database attacks from inside as well as outside the firewall. Many organizations focus their secu-

rity attention on attacks from outside the organization, and believe that a secure perimeter (such as firewalls) will eliminate most threats. But Gartner, Inc. estimates that 70 percent of security incidents that cause loss (rather than mere annoyance) to organizations involve insiders. Since an insider has trusted access to corporate systems, he or she is (by definition) inside the firewall—meaning that perimeter-based defenses will never see their attacks.

Multi-Tier Protection: Hackers increasingly are creating “blended” attacks that might use a port scan to find a way into a Web front-end, a password dictionary attack to gain illegal access to an application and a SQL injection attack against the database itself. A multi-tier protection approach is necessary to protect against attacks against any tier of the IT infrastructure, including the Web front-end, the application and middleware, and the back-end database. Application-level security must work to protect every tier of the IT infrastructure.

Enterprise-Class Infrastructure: As organizations move towards flexible, service-based IT architectures, applications may run on any number of tiers (or platforms) throughout the enterprise. The number, and nature of tiers on which an application depends may change unpredictably as business or technical needs change. Organizations cannot afford to pay skilled personnel to monitor multiple security scanning tools, nor can their networks afford the bandwidth it takes for those scanners to look for threats and report their results. Just as with network

Vulnerability Differences Report for Application	
Company Name:	Print Date: 9/10/2004
Application Name: Oracle8i Database (oracle) on 172.16.0.91, port 1521	Test Date 1:
Report Description:	Test Date 2:
AppDetective can be used to run individual tests on the applications on your network. You can also run later test on an application to determine what has changed for the vulnerabilities.	
The report provides a list of three categories:	
Fixed vulnerabilities - this lists the vulnerabilities that existed when you ran the first test but were fixed and therefore no longer exist.	
New vulnerabilities - this lists the vulnerabilities that did not exist when the first test was run but do exist in the second test. This reflects an application that has a new vulnerability.	
Outstanding vulnerabilities - this lists the vulnerabilities that existed for both tests. This indicates the vulnerabilities that have not been fixed.	
Fixed Vulnerabilities (exists in Test #1 but NOT in Test #2)	
▶ High	<input type="checkbox"/> <input type="checkbox"/>
Vulnerabilities still outstanding (exists in both Test #1 and Test #2)	
▶ High	<input type="checkbox"/> <input type="checkbox"/>
▶ Medium	<input type="checkbox"/> <input type="checkbox"/>
▶ Low	<input type="checkbox"/> <input type="checkbox"/>

Reports need to contain all of the vulnerabilities discovered from each Audit and Pen Test performed against an application.

and host-level security tools, organizations need scalable, enterprise-class application security tools that can grow to meet their future needs. A unified scanning infrastructure that works in a common fashion and provides the same capabilities within each tier of the application environment is required.

Distributed Management/Centralized Reporting: Distributed management and centralized reporting is the ability to delegate the responsibility for, and the work involved in, monitoring and managing application and database security across geographies or business units, while providing for centralized reporting of audit results. Modern businesses outsource more work than ever to consultants, contractors, or business partners such as distributors or contract manufacturers. An application-level security system must be flexible enough to delegate responsibility to such outside entities for keeping their portion of shared information systems secure. Even within a single organization, multiple business units, divisions or geographies must cooperate in keeping data secure, and take responsibility for securing that data. At the same time, however, the security infrastructure must provide a single, centralized security audit to provide for centralized accountability and enforcement of security processes.

SUMMARY

Applications and databases form the core of an organization's information technology infrastructure. Without the business processes they support (such as sales, marketing, finance, manufacturing, distribution and accounting) and the data they hold (such as customer names, production status, credit card data, and account histories) the business cannot function. Yet applications and databases have been disturbingly neglected within the enterprise compared to the security provided for networks and servers. Organizations that understand the importance of their applications and databases recognize the need for proactive, dynamic tools that can find and stop attacks on applications and databases before they cripple the enterprise. Fortunately, hard-earned experience securing the network provides a ready-made blueprint for an effective approach to securing enterprise applications: vulnerability assessments, real-time intrusion protection and audit, and encryption at the application layer.

¹M. Nicolett, J. Pescatore, *Gartner Group*, "Security Demands Drive Shift to Vulnerability Management," November 2003

²*Symantec*, "Internet Security Threat Report," March 2004.

³*Ibid*

ABOUT APPLICATION SECURITY, INC. (APPSECINC)

AppSecInc is the leading provider of database security solutions for the enterprise. AppSecInc products proactively secure enterprise applications at more than 250 organizations around the world by discovering, assessing, and protecting the database against rapidly changing security threats. By securing data at its source, we enable organizations to more confidently extend their business with customers, partners and suppliers. Our security experts, combined with our strong support team, deliver up-to-date application safeguards that minimize risk and eliminate its impact on business. Please contact us at 1-866-927-7732 to learn more, or visit us on the web at www.appsecinc.com.

**APPLICATION
SECURITY, INC.**

www.appsecinc.com