White Paper

A Practical Guide to Database Security

By

Sean Roth
Manager, Database Security Product Marketing, McAfee

Table of Contents

Why is database security so important? For a company that has suffered a serious data breach, it boils down to monetary damage in its many forms: business disruption, bad publicity, stiff fines for noncompliance, and undermined customer confidence. But most damaging of all is the trouble that it creates when it comes to signing up new customers. A tarnished reputation is a big objection for sales and business development to overcome. That's why data security in general and database security in particular are a crucial part of any company's overall corporate health.

The Pain and the Price

Are database and data security really that important? Just ask the folks at Sony.

In April of 2011, a series of sophisticated attacks penetrated systems at Sony PlayStation Network, Sony Online Entertainment, and Sony Pictures, forcing the company to shutter its popular online gaming and media sites until an extended forensic investigation was complete. More than 100 million user records were found to have been compromised, including user names, passwords, birth dates and some credit card information.

The incident drew the attention of authorities worldwide, setting off widespread speculation on the ultimate cost of operational, financial and regulatory impacts. Many estimates ran into the billions, depending on the scope of ensuing identity theft and credit card fraud.

While the ultimate costs may not be known for years, Sony issued a preliminary accounting of short-term response costs, which as of late May 2011 already amounted to $171 million in direct spending and anticipated costs related to:

- Estimated costs related to identity theft protection
- Welcome back programs
- Customer support
- Network security improvements
- Legal and consulting costs
- Future lost revenue

In June, a company official described the April intrusions as a bump in the road in a long-term relationship between Sony and its gaming customers, 90 percent of whom had returned to a PlayStation Network (PSN) service that was up, operational, and secure.

Then in October, Sony's online services came under attack again, resulting in 60,000 compromised accounts at PlayStation Network and 33,000 at Sony Online Entertainment. The company was forced to lock these accounts, reset passwords, and promise refunds of any unauthorized charges made against users' account wallets.

Clearly, the potential harm to consumers, merchants, banks and the possibility of multiple class action lawsuits resulting from the breach create the profile of a painful, time-consuming, and costly experience.

The combination of damaging and highly publicized data breaches and stricter regulatory compliance demands continue to push database security to the foreground. Even at this late date, database security still has an aura of mystery about it, almost as though it were a "black art." Part of the mystery is driven by two simple facts: many database professionals are not familiar with the security aspects of database management, while a large number of security professionals have a grasp of desktop security but not database security. This is beginning to change as the crucial importance of securing databases becomes more and more apparent.

Understanding the Threat

Databases are subject to unique types of threats that cannot be handled by firewalls, intrusion detection and prevention systems, and other perimeter defenses. The threat landscape is constantly evolving and becoming more sophisticated and specialized (for example attacks through memory backdoors inside databases).

Who are the intruders?

The profile of the typical hacker has changed dramatically over the last few years. The stereotyped image of the brilliant young loner who hacks into a secure environment driven by a desire for mischief and mayhem has given way to a different and more insidious threat—members of organized crime syndicates who prefer long-term stealth to short-term drama. In the end, this new type of cyberfraud professional is much more damaging. The changing profile and purpose of the intruder has reshaped the nature of intrusion attempts from ones that try to penetrate, then perhaps deface or wreak havoc, to ones that strive to be stealthy and leave no tracks with the aim of stealing data for financial gain.

Insider threat, privileged users

Concurrently with the change in the nature of the external threat, there is increasing attention being given to insider threats. This umbrella term refers to damage caused by individuals within the organization, either maliciously or accidentally.

Are insider threats serious? They certainly are. In the 2012 study, *Aftermath of a Data Breach*, conducted by the Ponemon Institute, investigators found that insiders and third parties are the most common causes of data breaches. Of the incidents that were successfully traced to a root cause, 34 percent were attributed to negligent insiders. An additional 19 percent were traced to third-party data outsourcers and 16 percent to malicious insiders. While all insiders are certainly not suspects, it is self-evident that insiders bent on stealing data have a greater chance of succeeding than those engaged in outside intrusion attempts.

The criminalization of the general threat landscape also has an impact on "crimes of opportunity" committed by insiders. It is often easier and quicker to offer a bribe to an insider who already enjoys access privileges rather than attempt to hack from the outside. And it isn't just enterprise data that is at risk. The Defense Advanced Research Projects Agency (DARPA) has launched a project for detecting and responding to insider threats on the US Department of Defense Networks.

Still smarting from the WikiLeaks disclosures that caused such embarrassment for the US Department of Defense, DARPA will use the Cyber INsiDER (CINDER) Threat Program to look for ways to improve the speed and accuracy of insider threat detection. DARPA will focus on looking for telltale signs and network activities that should be monitored to detect malicious activity.

Many regulatory compliance requirements focus on privileged insiders as well, with special attention given to those whose actions used to go unmonitored in the past.

Vulnerabilities

As database management systems have grown in complexity, they have become more vulnerable to attacks. These vulnerabilities range from the relatively benign to weaknesses that allow unauthorized users to own the database through privilege escalation—the exploitation of a software or security vulnerability by a user to gain unauthorized access to resources.

Much has been said and written about how database management system (DBMS) vendors should cope with vulnerabilities and how quickly they should patch them. The reality over the past few years shows that the number of reported vulnerabilities is constantly rising, despite the fact that vendors are doubling their efforts to patch them.

Additionally, it usually takes a vendor several months or more to distribute a patch and it takes an additional several months for customers to install the patches. The patch install process usually requires testing and database downtime. For these and other reasons, many customers do not apply the patches at all and their databases remain vulnerable to severe attacks.

Existing Solutions Are Inadequate

Perimeter security

When it comes to databases, the traditional perimeter defenses such as firewalls and intrusion detection systems/intrusion prevention systems (IDS/IPS) are grossly inadequate due to the sophisticated nature of the threats posed to databases and the specific nature of their vulnerabilities.

While many intrusion prevention systems claim to thwart, for example, SQL injection attacks, their capabilities in this area are very weak and are usually based on signatures, which hackers can easily evade. Intrusion prevention systems are certainly incapable of detecting sophisticated attacks that exploit vulnerabilities specific to database software from a particular vendor and to a particular version of the database and platforms that are supported by the database software.

Native DBMS auditing

Virtually all database management systems have the ability to write audit logs for some or all transactions. However, they are seldom used extensively due to the detrimental effect they have on database performance and the amount of storage they need for full auditing.

From a security standpoint, the usage of open logs is inadequate anyway since the logs can easily be manipulated after the fact and privileged users can turn the audit function on and off as they please.

Compliance and Security Requirements for Databases

There are many different compliance laws and regulations nowadays. The compliance landscape has evolved over the recent past and changed the way many IT systems, applications, and data are controlled.

Here are brief descriptions of the key regulations and their effect on database management and security practices.

Sarbanes-Oxley

The Sarbanes-Oxley Act of 2002 (SOX) forces publicly traded companies to be more transparent about their financial data. This federal law does not specify technical measures or tools, but requires companies to have "effective controls" in place.

- Section 302 of SOX requires executives to certify that its financial reporting is accurate and complete and that effective controls are in place to ensure the integrity of the reported data
- Section 404 of SOX further requires executives and auditors to report on the scope, adequacy, and effectiveness of internal control structures and procedures for financial reporting

Since most financial data reside in databases, SOX auditors are increasingly looking for ways to view database activity in a way that can be easily interpreted and acted on.

SOX also enforces serious criminal penalties for noncompliance, including fines of up to $1 million and prison terms of up to 10 years for corporate officers who knowingly certify a noncompliant statement. Those who do the same thing willfully can face a maximum $5 million fine or 20-year prison sentence or both.

PCI DSS

The Payment Card Industry Data Security Standard (PCI DSS) is a joint development of the major credit card companies, which aligned their individual security standards and released version 1.0 in December 2004. Version 2.0 was released in October 2010. The standard sets out technical and procedural requirements for merchants and companies that process and store credit card data. Unlike SOX, PCI gets specific about what measures need to be put in place for protecting card data. Stiff penalties can be assessed for noncompliance, with monthly fines ranging from $5,000 to $100,000, increased audit requirements, and the potential withdrawal of card processing rights.

Of particular interest is the concept of compensating controls (section 3.4), which recognizes that cardholder data encryption is sometimes impractical, and provides for the use of other methods, including real-time monitoring of user activity.

State and federal personal data protection laws

Most US states and territories have laws requiring public notification of all data breaches involving personally identifiable information. Most of these are modeled on California's regulation, Cal. Civ. Code 1798.82 and 1798.29 (originally California SB1386), which was passed in 2002 and became effective on July 1, 2003. An effort to standardize this regulation under federal jurisdiction failed in the 111th Congress (HR 2221), but several similar measures are currently under consideration in the 112th session (HR 1707, 1841, and 2577).

HIPAA

The Health Information Portability and Accountability Act (HIPAA) is a federal law that was put in place to ensure that the freedom of patients to choose healthcare insurers and providers does not come at the expense of the privacy of their medical records.

These articles are relevant to securing database:

- § 164.312(a)(1): Access Controls, which require organizations to "Implement technical policies and procedures for electronic information systems that maintain electronic protected health information to allow access only to those persons or software programs that have been granted access rights"
- § 164.312(b): Audit Controls, which require organizations to "Implement hardware, software, and/or procedural mechanisms that record and examine activity in information systems that contain or use electronic protected health information"
- § 164.312(c)(2): Integrity, which requires organizations to "Implement electronic mechanisms to corroborate that electronic protected health information has not been altered or destroyed in an unauthorized manner"

Although HIPAA had been notorious for weak enforcement, recent legislative action has changed the picture. When the American Recovery and Reinvestment Act (ARRA) of 2009 was signed into law, it included a strong healthcare component. The Health Information Technology for Economic and Clinical Health Act (HITECH Act), a key part of ARRA, has also become law and will turn up the heat on both penalties and enforcement for HIPAA compliance.

A key element of the new framework is that enforcement authority for HIPAA security rules has been transferred from the Centers for Medicare and Medicaid to the Office of Civil Rights, which has more resources to bring to the fight. According to published reports, the Office of Civil Rights—which has 275 investigators and a $40 million budget—is in a much better position to crack the enforcement whip.

SAS 70 and SSAE 16

SAS 70 (Statement on Auditing Standards Number 70) was an important accounting industry standard developed by the American Institute of Certified Public Accountants (AICPA). It provided guidelines for auditors engaged in assessing and validating the internal controls that service organizations had implemented to meet the requirements of data security regulations such as SOX, HIPPA, and GLBA. SAS 70 was introduced in 1992 and became the basis for many similar international standards.

In 2010, AICPA introduced SSAE 16, the Statement on Standards for Attestation Engagements Number 16, as an update and replacement for SAS 70. SSAE 16 reflects the global move toward internationally accepted accounting standards. Like both SOX and SAS 70, it does not specify what measures need to be in place, but it does provide tools and procedures that facilitate the job of auditors charged with assessing the effectiveness of a service organization's internal controls.

International data privacy and security regulations

Almost every county has its own set of laws and industry standards to regulate business handling of financial and personal data—and all of them are subject to change over time. In January 2012, for instance, the European Commission released a proposed new regulation to replace the existing Data Protection Directive 95/46/EC, which has governed data privacy in the EU since 1995. If adopted in its current form, the new regulation will introduce new rights for employees, consumers, and users across Europe, creating a variety of new challenges for businesses that operate in the Eurozone. A draft of the proposed regulations is available at: http://epic.org/privacy/intl/EU-Privacy-Regulation-29-11-2011.pdf.

Regulatory Compliance and Security

Compliance does not automatically equate to security. A company may be compliant with a host of regulatory requirements, yet its databases may remain exposed and vulnerable. There are several reasons why this happens:

- Regulations are limited in scope to whatever they are intended to regulate. PCI DSS, for example, is concerned with credit card details and credit card details only. If someone stole your employee data, for instance, this would be of no concern to PCI auditors.
- Compliance is often audit focused and audits are, by nature, activities that take place after the fact. So, while you may discover a breach four months earlier, it won't do you much good in reversing that data theft.
- The requirements posed by some regulations represent a minimum baseline, not a best practice. They are deliberately created this way to allow many organizations to comply.
- Regulations are outdated as soon as they are published. The threat landscape changes faster than lawmakers write laws and regulators issue regulations.

Reconciling Compliance and Security Requirements

Regulatory compliance can be seen as a burden, or it can be seen as a chance to streamline business processes and significantly upgrade security. To make the best of this opportunity, enterprises should align their compliance projects with security requirements to ensure that the required measures are implemented at the same time.

Even if your current initiative does not involve all your organization's major databases, why not apply the relevant security improvements to any that are subject to security or compliance audits?

Consider a publicly traded healthcare insurance company that is subject to Sarbanes-Oxley, HIPAA, state regulations like Cal. Civ. Code 1798.82, and PCI DSS. Many of its databases will be subject to more than one mandate, and their compliance requirements can be mapped out as shown below.

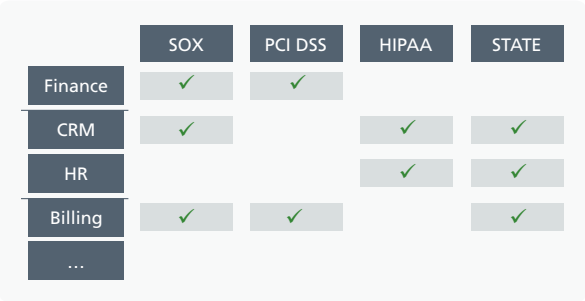

Figure 1. Compliance regulations by database type.

Whenever one database is modified for compliance with one mandate, it makes sense to upgrade all other systems with the same requirements. The additional effort is marginal and provides a more efficient and comprehensive view.

Overlapping requirements

Virtually all compliance requirements adhere to principles that are useful for improving security as a whole, and those can be leveraged for upgrading database security with no additional effort.

- *Controlling access to sensitive data*—Requires identifying sensitive data and using user authentication and monitoring to enforce access policy to that data
- *Separation of duties*—Requires that the people in charge of auditing or monitoring the database are not the same people whose actions are being monitored
- *Monitoring privileged users*—Due to the level of access given to privileged users, such as database administrators, it is important to monitor their activity for any unauthorized transactions

Five Principles for Protecting the Database

1. Think about security in everything you do

Integrate security into everything you do. Start with application development and move down to everyday tasks like user and data management. Teach your users to think the same way. Most security gaps can be traced to ignorance or inattention, so do your best to banish both.

2. Apply the least privilege principle

The least privilege principle gives users and applications only those privileges necessary to fulfill their functions. For database users, initial access privileges should be carefully restricted and regularly reviewed. If you must grant deep privileges to consultants and contract developers, be sure to rescind those privileges when the work is done.

Note that some vulnerabilities allow escalation of view privileges to provide more privileged access. Carefully consider need and risk before granting any privileges to any user.

3. Minimize the attack surface

It is more difficult to secure a large house with many windows than a small house with few windows. Databases are the same—the more complex the system, the larger the attack surface. Keep the system simple and the surface small. For new systems, install only the components your application requires. For existing systems, eliminate anything that's not in use.

4. Encrypt, but not as a panacea

Encryption is often the first thing that comes to mind when it comes to securing data and is certainly recommended for sensitive data. However, it can be expensive, difficult to use, and difficult to manage securely. Encrypt only sensitive data that requires it, be careful how you manage the encryption/decryption keys, and change them on a regular basis.

It is important to combine encryption with other means and procedures, such as activity monitoring, auditing, periodic vulnerability assessments, and user authentication.

5. Development, testing, and staging environments

Many organizations invest efforts in securing their production databases but neglect to do so in development, testing, and staging environments. As the staging environment is often copied into production when it is ready, it should be as secure as the production version. Beyond that, it is often the case that real production data is used in non-production environments without any masking. This poses a serious security risk. It is recommended that non-production environments are treated with the same tools and procedures one as production environments.

Five Practical First Steps

The following are steps that you can perform on your databases without getting into lengthy projects, vulnerability assessments, or the use of expensive tools and third-party services. They require only a minimal investment of time and attention. These steps will not make your database 100 percent secure, but they will bring you a long way towards a more secure environment.

1. Usernames and passwords

Nothing makes it easier to penetrate a database than weak passwords, default passwords, or password sharing. Oracle and other databases come with hundreds of default usernames and passwords created to simplify setup. These should always be erased after installation. There are free tools available on http://www.petefinnigan.com.

Passwords that are weak, too short, or based on easy-to-guess guessed words, names, and dates are easily cracked using readily available tools. It is imperative to use strong passwords that are not based on dictionary words, names, or dates and that use a combination of letters, numbers, and symbols.

2. Remove unnecessary components

Today's enterprise DBMSs are behemoth applications that ship with many options most people won't use. Certain vulnerabilities exploit the large attack surface created by add-ons and extensions such as Oracle's APEX. Initial database configurations should be as lightweight as possible, and should be reviewed periodically for unused components that can be removed without loss of functionality.

3. Apply security patches

New database vulnerabilities are uncovered constantly and many are patched by the database vendors themselves, who issue updates and patches to the DBMS. It is not always easy to apply those patches because they require testing and database downtime. But even deciding on a schedule where patches are applied twice a year is better than not applying them at all.

Ironically, it is precisely when patches are issued by the DBMS vendor that unpatched systems are even more vulnerable to attack. Why? Because the public announcement of the availability of such patches also alerts potential intruders to the existence of vulnerabilities in specific modules.

An alternative and complementary approach is to use virtual patching tools, such as the ones available from McAfee. These tools create an external layer of defense on top of the database that specifically addresses vulnerabilities and issues alerts or takes action to stop attempts to exploit them.

4. Follow secure coding practices

Many database vulnerabilities are exposed due to the way applications are coded and their interaction with the database. Lack of accountability and lack of secure coding practices may open the way to breaches and attacks.

For example, SQL injections in web applications can be thwarted entirely by binding variables in SQL statements. Unfortunately, many developers still do not use the bind variables method when developing applications, leaving the database exposed to SQL injections.

Architecting and designing for security, validating input, and sanitizing data sent to other systems are some of the recommended methods used in secure coding. You can visit the Computer Emergency Readiness Team (CERT) website at http://www.cert.org/secure-coding/ for additional information on coding standards and practices.

5. Monitor and audit, then monitor and audit again

You don't know what you can't see. That seems obvious, yet most DBAs and security professionals have no idea who is doing what in the database.

Auditing is an offline endeavor that looks back at the database activity over a period. Full native DBMS auditing is impractical as it significantly slows performance, but selective fine-grained auditing can be used. While certain compliance requirements may force you to audit everything and keep an infinite audit trail, it is also impractical—the more benign actions you record, the less likely you are to notice the conspicuous ones. Try to work with auditors to define the types of activities that are crucial (accessing sensitive data, privileged user access, and others), and record those.

Monitoring occurs in real time and is therefore more actionable and useful in security terms. There are tools, most of them expensive, that are available for enterprise deployments. Some of these tools have prevention capabilities. But now there is a breakthrough product that can monitor multiple databases, terminate attacks by using policy violations as a tool, provide access to virtual patches that protect systems from many known and zero-day vulnerabilities, and deliver real-time notification through email or security information and event management/security event management (SIEM/SEM) integration.

McAfee Database Activity Monitoring

Real-time activity monitoring and breach prevention

McAfee® Database Activity Monitoring is a database security solution that monitors all database activity in real time based on defined rules and policies. It issues alerts on suspicious activity and, if necessary, intervenes in real time.

McAfee Database Activity Monitoring makes use of small-footprint software agents that are installed on the database host server itself and that monitor all activity. The design is non-intrusive, easy to install, and consumes only small amounts of CPU resources (less than 5 percent of a single CPU, even on multi-CPU machines). The sensors communicate with the McAfee Database Activity server, which generates alerts in accordance with its defined rules.

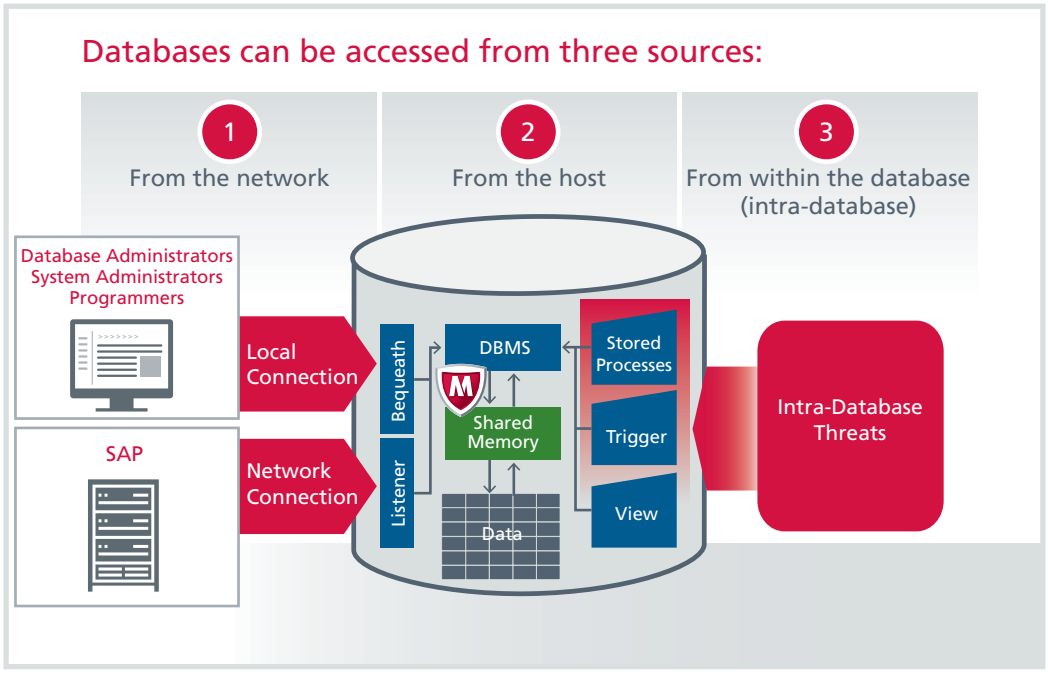

Figure 2. Database access vectors.

Policy rules apply to types of SQL statements, database objects accessed, time of day or day of the month, specific user profiles, IP addresses, and the applications used—among other parameters. The action taken when the conditions of a rule are met can be as simple as logging an event, sending an alert to a SIEM/SEM system via email or SMS, or terminating a user session to prevent malicious activity. Users can also be quarantined to prevent subsequent attempts at breaching the database. The system comes with predefined rules that prevent known attacks that exploit database vulnerabilities.

A single McAfee Database Activity Monitoring server can manage and communicate with numerous sensors on different databases, and an enterprise installation can scale to support hundreds of databases. The server also easily integrates with your IT infrastructure to facilitate central IT management and security event management.

Since the McAfee Database Activity Monitoring sensor is installed on the database machine, it is impossible to bypass. It possesses self-defense mechanisms that send out an alert if any tampering attempts are undertaken. The structure of the system ensures separation of duties, with definable roles and access rights to different users.

Key advantages of McAfee Database Activity Monitoring

- Maximizes visibility and protection from all sources of attacks
- Monitors external threats, privileged insiders, and sophisticated threats from within the database
- Minimizes risk and liability by stopping attacks before they cause damage
- Saves time and money with faster deployment and a more efficient architecture
- Gives you the flexibility to easily deploy on the IT infrastructure of your choice
- Integrates with core McAfee products, including the McAfee® ePolicy Orchestrator® (McAfee ePO™) management platform and McAfee Vulnerability Manager for Databases

Unique advantages of McAfee Database Activity Monitoring

- The only database monitoring solution that monitors all database activities and provides protection against insiders with privileged access
- Granular monitoring of database transactions, queries, objects, and stored procedures, with real-time alerts and prevention of breaches
- Flexible rules that allow enforcement of corporate security policy with minimal "false positive" alerts
- Virtual patching of newly discovered database vulnerabilities, providing immediate protection with no DBMS downtime
- Flexible audit and reporting capabilities suitable for PCI DSS, SOX, and HIPAA
- An easy-to-deploy and scalable software solution
- Multiple user rules that facilitate separation of duties

Download your trial version of McAfee Database Activity Monitoring.

2821 Mission College Boulevard
Santa Clara, CA 95054
888 847 8766
www.mcafee.com

43108wp_db-security_0312_fnl_ETMG