White Paper

McAfee®
An Intel Company

Overcome the Challenges of Database Patching in Production Environments

Achieve compliance to reduce downtime, testing, and resource requirements

Table of Contents

The Realities of Database Patching

- Given the complexity of database platforms and the varied ways we use them, and with new exploits being discovered frequently, vulnerabilities are a reality that must be carefully managed
- Cybercriminals—whether they are outside hackers or malicious insiders—are attracted to databases, as they typically hold an organization's most valued and sensitive information
- Critical patch updates from database vendors are issued regularly to address known vulnerabilities, but for a variety of reasons, organizations are often unable to install them in a timely manner, if at all
- The most critical period for vulnerability extends from the time that database management system (DBMS) vendors issue the security patch until it is applied, and it is during this window that hackers now know that a weakness is probably still exploitable
- Automation tools used by hackers and sharing of vulnerabilities over social networks increase the frequency and sophistication of database attacks
- Compliance standards and governance policies often dictate a maximum time until patches must be implemented, but many organizations struggle to meet these strict mandates

What are the underlying threats posed to sensitive data by not patching regularly? We propose a practical and reasoned method to protecting databases when vendor patches cannot be installed promptly after they are released.

Protecting Databases From Attacks—The Stakes Have Never Been Higher

No one would argue against the proposition that data protection and privacy remain major concerns for businesses and government institutions worldwide. If anything, the stakes have grown even higher. The hacker community has matured from the early days of the brilliant loner whose main motivation was to provide proof of high-level code-cracking ability for the purpose of ego satisfaction and the thrill of wreaking havoc. Within the modern security landscape, hacking has become a serious and sophisticated business, often practiced by organized crime syndicates and sometimes with assistance from privileged insiders at targeted organizations.

The negative consequences for businesses can be very serious in terms of damaged reputation, fines for noncompliance, and the potential loss of existing customers whose confidence has been badly shaken. And the exodus of existing customers could be compounded by difficulty in signing up new customers who have heard or read about the security breach. Certainly, corporate security administrators have frequent and painful reminders about the cost of breaches.

Although there is a scarcity of security breach data that is tied specifically to organizations failing to patch, two things remain crystal clear:

- The cost of a breach is far too painful for user sites to be willing to absorb
- Without the protection of applying either real or virtual patches, databases are exposed

Leaving databases unpatched just isn't worth the risk. Although it may be difficult to tie a security breach directly to the absence of a patch, it's not difficult to understand the price that user sites pay when they are breached.

For example, early in 2010, the Connecticut attorney general's office filed a lawsuit against Health Net of Connecticut, alleging that Health Net failed to secure patient medical records and financial information prior to a security breach. In this case, 446,000 customers may have been affected by the breach.

And the stakes have also grown higher because of a tightened regulatory and compliance structure. For example, public companies must adhere to data security controls outlined in Sarbanes-Oxley (SOX), organizations handling private medical records are subject to Health Insurance Portability and Accountability Act (HIPAA), and anyone accepting credit cards or debit cards must follow the provisions of the Payment Card Industry Data Security Standard (PCI DSS).

The risks of not complying with these and many other regulations include fines and potential business disruption, as well as the direct costs of notification in the case of a breach. Beyond external factors, many organizations also have their own IT governance policies that mandate the timely application of security patches to all software, including database systems.

"The ability to apply Critical Patch Updates to all database systems also plays a key role in preventing database breaches and insuring compliance with corporate security policies regarding patching and configuration. It ensures that a consistent policy is enforced company-wide—in a uniform fashion. However, less than one-third (29 percent) will apply the patches within the time frame they are released (within one to three months). Almost one out of 10 (9 percent) say the patches are either applied after a year or not at all."

Databases Are More at Risk than EVER: 2011 IOUG Data Security Survey

Vendor Security Patches—Improved Protection for Databases

While businesses across nearly every industry have realized the crucial nature of data security and privacy, database software vendors have also responded to the intensifying threat posed by hackers and malicious insiders. Oracle, Microsoft, IBM, and others regularly issue security patches to plug vulnerabilities that have been discovered either by their own testing, the work of independent ethical researchers who report their findings, or when one of their customer systems is breached.

Oracle uses its Critical Patch Updates (CPUs) as the primary method of releasing security fixes for its products. They are released quarterly—in January, April, July and October—and a significant number of serious vulnerabilities in the core database are typically addressed in every CPU. Similarly, Microsoft utilizes a monthly process called "Patch Tuesday" and, although it less frequently includes fixes for Microsoft SQL Server, on those occasions when it does, users are encouraged to update database servers as soon as it is feasible.

Despite the fact that vendors regularly issue security updates, many organizations often delay installing them, or simply do not install them at all. While this hesitancy to install patches is understandable and stems from a variety of reasons, it creates a dangerous situation for the organizations because it introduces a "window of vulnerability"—the period of time between the issuance of a vendor security patch and the installation of the patch.

One of the real ironies of the vendor patch cycle is that the release of a patch—obviously aimed at providing added security—actually increases the likelihood that more attacks will take place. Patch releases are well publicized, giving hackers a signal that the time to strike is at hand. Within the window—from announcement until the patch has been applied to each system—not only are databases not fully protected and therefore susceptible to attack, but, in fact, the hacker community is now fully aware of the vulnerability and can determine precisely how to exploit it.

Within days (or hours, in some cases) of a vendor patch release, the social network of black hats is rife with details on the nature of the underlying problem and how to leverage it to penetrate database servers that have not yet been patched. Of course, a few sophisticated hackers also discover their own zero-day vulnerabilities, but this is much more difficult than simply taking advantage of published vulnerabilities that many user sites have not yet patched.

Surprisingly, the reluctance to promptly install security updates is a more common practice than one might suspect. In September of 2010, The Independent Oracle Users Group released data gathered from a survey of 430 Oracle database administrators, consultants, and developers. The results were both eye-opening and disturbing.

The data indicated that only 37 percent of this group installed Oracle CPUs within three months after they were released. While 16 percent were unsure of their patching frequency (therefore it is likely not a high priority for them), more than 40 percent of the respondents indicated they were not applying patches for at least three months, relying entirely on their network security to prevent a breach of their critical databases. By ignoring these important security updates, organizations are simply rolling the dice and hoping that their databases will not be the target of an attack.

Some of the industry's top analysts who focus on database security do not mince words when pointing out the importance of installing patches.

If Patches Are Available, Why Aren't They Installed?

As we've seen, although vendor-issued security patches are crucially important in maintaining a safety net covering an organization's databases, those same systems are remaining unpatched and, therefore, highly vulnerable to attack. To say the least, this reluctance is counterintuitive, but there are often valid reasons for this hesitation.

The gap from patch release to patch application leaves a long way to fall.

- Patching is an update to the DBMS kernel and requires database downtime. This is of extreme concern in business-critical environments where systems must operate with 24/7 availability.
- To ensure that the update does not negatively impact other software, most organizations will first perform complete regression testing of all applications running on top of the database prior to patching
- Many application vendors will only certify their applications to run on top of specific releases of DBMSs. Users must wait for this certification or risk potential problems when requesting application technical support.
- There are older versions of databases that are still widely used (for example, Oracle 8i and 9i) but are no longer officially supported by the vendor, and therefore patches are not made available for these versions. In many cases, upgrading applications to a newer database release is not possible, as the source code may not be available or the development team has more critical priorities.

So, while it is understandable that organizations want to avoid downtime, minimize regression testing, and continue using older database versions, the fact remains that unpatched systems are trouble waiting to happen. Would-be attackers are waiting in the weeds to make their move and will use techniques such as SQL injection, buffer overflows, and privilege escalation to gain access through a vulnerability that allows security to be bypassed.

The net result is that a compromised database could allow an attacker to deny availability to legitimate users, inject malicious content, or access sensitive information. In some cases, changes can be made to the database structure itself, even to the extent of dropping tables entirely.

Closing the Window of Risk with Virtual Patching

At this point, it's clear that the threat landscape for databases remains daunting. Troubling reports of new and costly breaches continue to make their way into the public consciousness, leading to more and more unease on the part of credit card companies, health insurance providers, and all of their customers. When vulnerabilities are left unpatched, auditors of public companies subject to Sarbanes-Oxley and the Gramm-Leach-Bliley Act (GLBA) are understandably concerned with certifying that a company has adequate control over financial records.

But for all of that, organizations are still reluctant to endure the downtime and the other complications that go along with installing vendor patches in a timely and consistent manner. Fortunately, there is a way to bridge the gap between uninstalled DBMS vendor patches and the ever-present threat posed by hackers. It comes in the form of virtual patching.

Virtual patching is a way to protect the database against exploits without modifying the DBMS binaries. This creates a security layer around the database that, unlike vendor patching, does not require downtime or application testing and can also protect older, unsupported DBMS releases.

By monitoring all activity occurring in the database memory and detecting "fingerprints" of known exploits and vulnerabilities, virtual patching identifies attempted attacks. When a match occurs, an alert is issued in real time, the suspicious session can be terminated, and the originating user can be quarantined for a specified period until the nature of the suspected attack is investigated.

Virtual patching has proven to be an effective way to quickly and easily meet compliance patching requirements by reducing the risk created by long intervals between vendor patch releases and their deployment. And it is the only way to protect older DBMS versions that are no longer supported by the vendors—Oracle 8i/9i, and Microsoft SQL Server 2000.

McAfee® Virtual Patching for Databases software protects databases against many known vulnerabilities as soon as they are discovered—and also from many common threat vectors used in zero-day attacks.

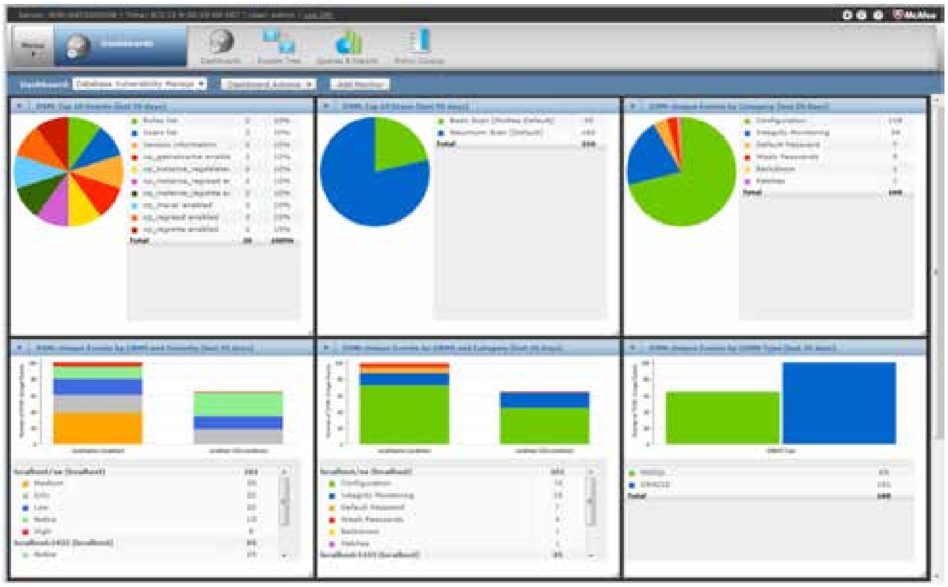

Figure 1. The McAfee® ePolicy Orchestrator® dashboard makes crucial database vulnerability information easy to find.

The hybrid approach of implementing vendor patches as soon as they can be reasonably deployed, in combination with virtual patching in the interim, provides an effective balance between the challenges of patch deployment and the need to implement adequate security for compliance purposes.

In addition to protecting from attacks that exploit many specific known vulnerabilities, McAfee Virtual Patching for Databases protection also includes many “generic” rules to identify and protect from “zero-day” vulnerabilities. By looking for certain patterns of misuse (for example, a system package attempting to escalate privileges), the common threat vectors likely to be used by future exploits can be proactively closed.

McAfee Virtual Patching for Databases is provided on an annual subscription basis and offers numerous features:

- Host-based software that uses a small, non-intrusive sensor on each database server to protect the database with a set of protections in the form of virtual patches that detect and prevent attempted exploits of DBMS vulnerabilities
- As new vulnerabilities are discovered, shared by other researchers who cooperate with our research team, disclosed publicly by third parties, or included in a vendor patch, updates with new virtual patch rules are released by McAfee and can be automatically distributed to all covered databases through a centralized management console
- McAfee Virtual Patching for Databases installation is non-intrusive, as the sensor is read-only, installed as a user process, makes no changes to the DBMS software itself, and doesn’t require any database downtime. It does not require the use of native DBMS auditing and utilizes minimal resources on each database server.

McAfee Virtual Patching for Databases features allow the virtual patching process to move quickly and not interfere with the work of production databases, requiring no downtime for application of new virtual patch updates. What's more, McAfee Virtual Patching for Databases directly monitors the database memory cache and has full visibility into all database activity, regardless of the source of an attack, whether coming from an application over the network, a user connected directly to the database server, or from within the database itself, for example, utilizing a stored procedure or trigger.

As an added benefit, sites running McAfee Virtual Patching for Databases also have increased visibility into what are ultimately failed efforts to exploit vulnerabilities. Even on systems where a physical patch has already been installed, McAfee Virtual Patching for Databases alerts indicate attempts to utilize these weaknesses, acting as a "canary in the coal mine." While the attack was destined to fail as the database was already patched, the valuable information that internal or external hackers are targeting the organization's systems can be used as an early warning system to proactively restrict access from the origination point of the threat.

Current and Future Challenges—The Impact of Cloud Computing

When it comes to cloud computing environments, security also becomes a major concern because computing resources are not under the direct control of the "owner" of the database. Whether you are deploying a private cloud in your own data center using virtualization technologies or implementing a database on an infrastructure-as-a-service (IaaS) in the public cloud, such as Amazon's EC2, the same issues of database security apply. McAfee Database Activity Monitoring and the McAfee Virtual Patching for Databases solution can be of particular benefit to these environments because their software-only architecture is easy to deploy, and they function well, even in these highly distributed and dynamic implementations.

A Case in Point—A Regional Bank Achieves Compliance with Virtual Patching

The use of McAfee Virtual Patching for Databases simplifies and accelerates the process of complying with standards and regulations like SOX, PCI DSS, Statement on Auditing Standards No. 70 (SAS 70), HIPAA, and GLBA, among others.

For example, one McAfee Virtual Patching for Databases customer, a West Coast regional bank faced an upcoming audit mandated by GLBA, which governs the financial services industry. For various reasons, the organization had not been able to consistently patch all database servers—and a GLBA audit was due in a couple of weeks. With hundreds of databases, spanning several releases (some no longer supported by the vendor) and dozens of applications, there simply wasn't enough time to get the job done without significant disruption to business operations.

McAfee Virtual Patching for Databases was deployed on about 200 databases in the first week, with zero disruption to ongoing database operations. Tests were conducted to verify that the McAfee Virtual Patching for Databases solution did not affect database performance or introduce any application instability. The use of McAfee Virtual Patching for Databases technology was then expanded to other critical databases and the company began the GLBA audit process. The bank's new "hybrid" patching strategy—under which databases with the most sensitive data would be patched with a vendor patch at regular intervals and McAfee Virtual Patching for Databases would be used to protect during the interim—enabled them to meet compliance requirements.

The installation time of two weeks—rather than the many months and the resources it would have taken to install vendor-issued physical patches—was a huge boon to the financial services company, not only saving them time, but also helping them avoid disruption, since they didn't have to take down their systems to install physical patches.

For the bank, McAfee Virtual Patching for Databases technology provided much-needed flexibility, allowing them to create an attainable plan for securing their databases with a combination of physical patching on a reasonable schedule and virtual patches to protect in the interim.

Risks Are Here to Stay, But Virtual Patching Can Help

The frequency and sophistication of attacks shows no signs of declining, and today's economic climate does not allow for more resources to be applied to the challenging task of testing and applying physical patches promptly upon their release. Combined with the compliance framework of SOX, HIPAA, PCI DSS, and the like, the need to demonstrate adherence to a strict patching policy will only become more demanding, and the use of McAfee Virtual Patching for Databases makes perfect sense for closing the window of risk, while saving dollars and minimizing business disruption.

For more information, visit www.mcafee.com/dbsecurity, or contact your local McAfee representative or reseller near you.

About McAfee Endpoint Security

Next-generation McAfee endpoint security solutions provide security across all of your devices, the data that runs through them, and applications that run on them. These comprehensive and tailored solutions reduce complexity to achieve multilayer endpoint defense—without impacting productivity. It's the perfect blend of traditional smart malware scanning, dynamic whitelisting, behavioral zero-day intrusion prevention, unified management, and integrated threat intelligence. Find out more at www.mcafee.com/endpoint.

About McAfee

McAfee, a wholly owned subsidiary of Intel Corporation (NASDAQ:INTC), is the world's largest dedicated security technology company. McAfee delivers proactive and proven solutions and services that help secure systems, networks, and mobile devices around the world, allowing users to safely connect to the Internet, browse, and shop the web more securely. Backed by its unrivaled global threat intelligence, McAfee creates innovative products that empower home users, businesses, the public sector, and service providers by enabling them to prove compliance with regulations, protect data, prevent disruptions, identify vulnerabilities, and continuously monitor and improve their security. McAfee is relentlessly focused on constantly finding new ways to keep our customers safe. http://www.mcafee.com

McAfee
An Intel Company

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