BEYOND THE PATCH: REDUCING THE RISK OF DATABASE AND APPLICATION VULNERABILITIES

October 2016

Author: Derek E. Brink, CISSP
Vice President and Research Fellow, Information Security and IT GRC

Report Highlights

p2

Organizations should think less about a specific, tactical checklist of technical security controls (e.g., *patch management*), and more about higher-level enterprise capabilities (e.g., *reducing the risks of vulnerabilities* to an acceptable level). The most appropriate choice of controls can vary, based on your own specific context.

p5

In several common scenarios, the strategy of *virtual patching* makes the most sense for the business, both operationally and financially: when vendor patches are *lagging*, *not possible*, or *not practical*; to avoid the *time*, *cost*, and *disruption* of vendor patching; and to further reduce risk with *enhanced visibility* and *alerting*.

p7

For a mid-size enterprise with just 100 database instances, Aberdeen estimates between 220 and 660 vendor patches per year, with a most likely value of about 410. Using a traditional, vendor patching strategy requires a median of about 910 hours per year of disruption to enterprise databases and applications.

p8

Including the negative impact on revenue and user productivity in addition to the cost of administrative staff, Aberdeen's analysis estimates the total business impact from a traditional, vendor patching approach for enterprise databases and applications to be between 1% and 8% of annual revenue, with a median value of 4%.

About one third (31%) of all respondents in Aberdeen Group's benchmark research have turned to virtual patching, particularly for the scenario of enterprise databases and associated applications. Under a virtual patching strategy, the window of vulnerability is substantially shorter than with traditional vendor patching, substantially reducing the *likelihood* that enterprise databases and applications may be compromised. In addition, with virtual patching the two biggest contributors to the total annual *business impact* — i.e., lost revenue and lost user productivity during the time that enterprise databases and applications are disrupted for traditional vendor patching — are substantially eliminated. In other words: in comparison to traditional vendor patching, virtual patching is an effective strategy for addressing both aspects of security-related risk.

ABERDEEN GROUP

Aberdeen Group has repeatedly recommended that organizations should be thinking less about a specific, tactical checklist of technical security controls (e.g., *patch management*), and more about an essential set of higher-level enterprise capabilities (e.g., *reducing the risks of vulnerabilities* to an acceptable level). The most appropriate choice of controls can vary, based on your own specific context.

Leverage the Power of Community, But Still Think for Yourself

As Aberdeen Group noted in *Flash Forward: Putting "Critical Security Controls" in Perspective* (January 2015), the core concept behind the so-called Critical Security Controls movement is to leverage the power of the information security community to identify a small number of security controls that have been demonstrated to have a high payoff in terms of preventing known attacks. The attraction of such a "CSC Top 20," checklist-style approach is amplified by the reality that security teams in many organizations have found it increasingly difficult to sort through and evaluate the rich and complex stream of security technologies that solution providers continue to make available — a situation that has been aptly referred to as "The Fog of More."

At the same time, Aberdeen has consistently advocated that such checklists are not a recipe to be strictly followed, but rather a quicker path to considering the successful choices that other companies have made. Organizations should then adapt these approaches in the way that represents the best fit for their own specific **context**, which includes their *systems*; their *applications* and *data*; their *users*; their *industry* and *regulatory requirements*; their *mission* and *strategy*; and their *appetite for risk*.

In particular, Aberdeen has repeatedly recommended that organizations should be thinking less about a tactical checklist of specific technical security **controls**, and more about an essential set of higher-level enterprise **capabilities**, as in the following examples:

- ➔ **Understand what systems and applications are in your environment** (CSC 1, 2)
- ➔ **Keep your systems, applications, and networks securely configured, up to date, and protected against vulnerabilities** (CSC 3, 11, 9, 4, 8)

3

The Goal is to Reduce the Risks of Vulnerabilities to an Acceptable Level — But That Isn't Synonymous with Traditional Patching!

It's worth noting that anyone can fall victim to the checklist mentality, even the curators of the Top 20 Critical Security Controls. As a case in point, consider the following description of **CSC 4 (Continuous Vulnerability Assessment and Remediation)**, from their Practical Guidance for Implementing the Critical Security Controls (V6) (formatting added for emphasis):

- ➔ The goal of this Control is to understand the technical software weaknesses that exist in an organization's information systems, and to *remove or remediate those weaknesses*. [See sidebar at left.]
- ➔ Successful organizations implement *patch management systems* [which could include *virtual patching*] that cover both operating system and third-party application vulnerabilities. This allows for the *automatic, ongoing, and proactive* installation of updates to address software vulnerabilities.
- ➔ In addition to patch management systems, organizations must implement a commercial *vulnerability management system* to give themselves the ability to detect where exploitable software weaknesses currently exist so they can be remediated.

The point being made here is not just semantics. We can start from the premise that managing the risks of vulnerabilities to an acceptable level is a foundational capability for successful enterprise information security initiatives — but we need to understand that *how* managing those risks is achieved is not as

The Responses to Risk

This may be anathema to many old-school security professionals, but contrary to the guidance about the Critical Security Controls, not all risks need to be addressed!

Risks may be *accepted*; *ignored* (which is an inferior form of acceptance); *transferred* to other parties; or *managed to an acceptable level* through an investment in an appropriate mix of technical, administrative, and physical controls.

A great many security professionals firmly believe in their hearts — mistakenly — that the fundamental objective of information security is to counter all threats, remediate all vulnerabilities, and mitigate all risks.

But the goal is not to "remove or remediate all weaknesses" — on the contrary, the goal is to take steps to **reduce security-related risks to an acceptable level**, which also means balancing the factors of cost and context for their particular environment in the selection of a specific mix of controls.

cut-and-dried as the above guidance would imply. There's an important difference between a **capability** (e.g., an efficient and cost-effective mechanism for *managing the risk of software vulnerabilities*), and a specific type of technical **control** (e.g., traditional *vendor patching*). In other words: we need to think more broadly about how best to achieve the objective, given a particular operational context.

Scenarios When Virtual Patching Makes a Lot of Sense

Virtual patching — sometimes known as *external patching* or *vulnerability shielding* — refers to establishing a **policy enforcement point** that is external to the resource being protected, to identify and intercept exploits of vulnerabilities before they reach their target. In this way, direct modifications to the resource being protected are not required, and updates can be automatic and ongoing.

A high-level summary of common scenarios where the strategy of virtual patching makes the most sense for the business, both operationally and financially, is provided in Table 1:

- ➔ It provides an effective compensating control when vendor patches are not available, or when vendor patching is not possible or not practical
- ➔ It provides protection against both known and unknown vulnerabilities, including "zero-day" exploits
- ➔ It reduces the need for emergency patches or workarounds
- ➔ It requires fewer policy enforcement points (i.e., at selected points in the network, as opposed to applying vendor patches on every system)

Definitions

A **Policy Decision Point** is where access policies are evaluated and combined to yield a yes / no value for use by a *Policy Enforcement Point.*

A **Policy Enforcement Point** is where a yes / no policy decision from a *Policy Decision Point* is used to grant or deny access to a protected resource. Policy Enforcement Points typically reside throughout the organization, e.g., within applications, databases, file systems, network devices, and endpoint systems.

A **Policy Administration Point** is where access policies are defined and managed.

A **compensating control** refers to countermeasures or safeguards put in place to mitigate specific risks, in lieu of the nominally recommended controls, as a result of legitimate technical or business reasons.

5

➔ Related Research: *Virtual Patching and Database Security: An Effective Compensating Control*

➔ Related Research: *The Virtues of Virtual Patching*

- ➔ It gives enterprises the flexibility to implement vendor patches on the schedule of their choice — or mitigates the need to implement vendor patches at all
- ➔ It helps to mitigate the high opportunity cost of downtime (planned, or unplanned) for critical enterprise databases and applications
- ➔ It reduces the likelihood of disrupting critical enterprise applications and databases, since libraries and support code files are unchanged
- ➔ It can provide visibility and alerts on attempted exploits, suspicious activity, and unwanted database activity

Table 1: Scenarios When Virtual Patching Makes a Lot of Sense

Scenario	Examples
Vendor patches may not be available	▪ There is often a significant lag time between the public disclosure of a vulnerability, and the availability of a patch from the vendor — and some vulnerabilities are never patched ▪ Third-party application vendors need to test and certify database security updates with their own applications before enterprises can deploy updates, adding additional lag time
Vendor patching may not be possible or practical	▪ Even if new vulnerabilities are discovered, vendors no longer provide patches for older, out of support systems ▪ Vendor patches may not be available for OEM systems (e.g., where license agreements may prohibit modifications to the underlying platform), or for outsourced code

Vendor patching is costly, time-consuming, and inconvenient	▪ The patching process itself — i.e., assessing, prioritizing, testing, and implementing — is costly and time-consuming, particularly when database restarts are required ▪ The opportunity cost of downtime or system outages — e.g., lost end-user productivity, deferred or lost revenue, and the cost of administrators — provides an incentive to defer vendor patching
Vendor patching does not support up-to-date visibility into what's happening in your environment	▪ Attempts to exploit a patched vulnerability will fail, but no alert will be raised — allowing attackers the ongoing opportunity to make alternative attempts ▪ Similarly, no alerts are raised for other suspicious or unwanted database activity (e.g., access of default accounts, use of evasion techniques, or access of file systems from databases)

Source: Aberdeen Group, October 2016

Traditional Vendor Patching is a Complex, Never-Ending Process

Let's pause here, and bring a taste of quantification to our discussion regarding the cost and complexity of a traditional, vendor patching approach for enterprise databases and applications. Consider a simple example, based on a mid-size enterprise with the following handful of assumptions:

- ➔ 100 database instances
- ➔ Each database instance serves between 1 and 3 enterprise applications
- ➔ 1 to 4 database security updates per instance, per year
- ➔ 1 to 5 hours to implement each database security instance

At this point, let's keep things simple by ignoring factors such as the time it takes to validate database security updates in a test environment, to confirm they don't break anything in the production environment. How many database security updates in this scenario need to be implemented over a 12-month period?

Definitions

A database **instance** is typically used to describe a complete database environment, including the RDBMS software, table structure, stored procedures, and all other associated functionality. For example: *development*, *test*, and *production* might represent three instances of the same database environment.

7

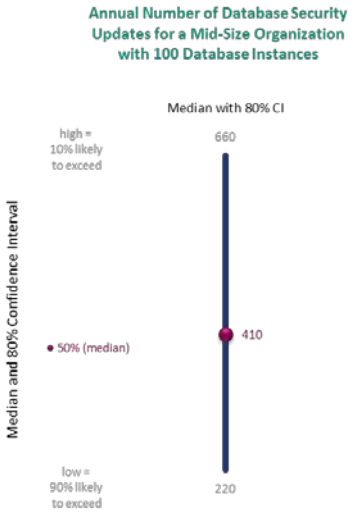

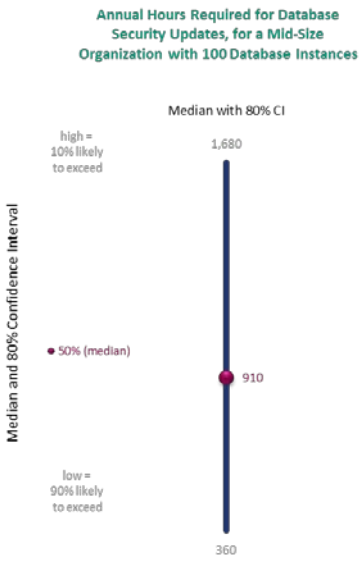

A simple Monte Carlo model provides some somewhat surprising insights. With a 90% likelihood, there are **more than *220*** vendor patches for this simple environment – and with 10% likelihood, there are ***more than 660***. The median value is ***about 410,*** which means that vendor patches most likely need to be applied to the databases and applications that support this mid-size organization's revenue and users more than 400 times per year.

In terms of *time*, the annual hours required for these database security updates is estimated to be **between 360 and 1,680 hours** (the *80% confidence interval*), with a most likely value of **about 910 hours**. This means there are more than 900 hours per year that the databases and applications that support this mid-size organization's revenue and users may be disrupted, with the corresponding loss of revenue and user productivity, along with the associated cost of administrative staff.

By itself, the fully-loaded *cost of administrative staff* for a traditional, vendor patching approach in this scenario is not too overwhelming — in Aberdeen's simple model, the median value is estimated at **about $50K per year**. It turns out that the median impact on *lost revenue* (**about $100K per year**) and *lost user productivity* (**about $3.8M per year**) during the time that enterprise databases and applications are disrupted are much larger contributors to the total business impact of a traditional, vendor patching strategy. These estimates are based on the same mid-size enterprise, with $100M in annual revenue and 1,000 users, and the simplifying assumptions that disruptions to enterprise databases and applications affect both revenue and user productivity uniformly throughout the year.

Including the negative impact on revenue and user productivity in addition to the cost of administrative staff, Aberdeen's analysis estimates the total business impact from a traditional, vendor patching approach for enterprise databases and applications to be between 1% and 8% of annual revenue, with a median of about 4%.

In total, Aberdeen's model estimates the total business impact from a traditional, vendor patching approach for enterprise databases and applications to be **between $1M and $8M per year** in this scenario, with a **median value of about $4M**. Expressed another way, that's the equivalent of **1% to 8% of annual revenue**, with a **most likely value of about 4%**.

In addition, keep in mind that vendor patches may not be immediately available, and that as a routine practice the enterprise may purposely defer implementation of vendor patches for reasons of time, cost, or (in)convenience — leaving the window of vulnerability open, and increasing the risk that the organization's databases and applications may be compromised. (The business impact of those risks is *not* incorporated in the current analysis.)

Given that managing the risks of vulnerabilities in enterprise databases and applications to an acceptable level is the overarching objective for any of this activity in the first place, this is a good illustration of why traditional vendor patch management might not always be the optimal choice of security control.

And this is for just 100 database instances — imagine the cost, time, and complexity of a traditional, vendor patching approach for a much larger, more complex IT environment. This raises some legitimate questions with respect to reducing the risk of database and application vulnerabilities: Can the organization keep up with this level of volume, complexity, and cost? Are the security implications fully considered, and consistent with the organization's appetite for risk?

A high-level summary of how a virtual patching strategy reduces the risk of database and application vulnerabilities, in comparison to a traditional vendor patching approach, is provided in Table 2.

Table 2: In Comparison to Traditional Vendor Patching, How a Virtual Patching Strategy Reduces the Risk of Database and Application Vulnerabilities

Selected Factors of Annual Business Impact	Estimated Median Annual Business Impact	
	Vendor Patching	Virtual Patching
Fully-loaded cost of administrative staff	$50K	Comparable, or lower
Lost revenue during the time that enterprise databases and applications are disrupted	$100K	Substantially eliminated
Lost user productivity during the time that enterprise databases and applications are disrupted	$3.8M	Substantially eliminated
Risk that enterprise databases and applications may be compromised, based on the window of vulnerability from public disclosure to mitigation	High: vendor patching is not available, not possible, not practical, or deferred to avoid cost and inconvenience	Substantially reduced

Source: Monte Carlo analysis based on a mid-size enterprise with 100 database instances, $100M in annual revenue, and 1K users; Aberdeen Group, October 2016

For the same illustrative scenario, when using a virtual patching approach:

- ➔ The fully-loaded *cost of administrative staff* is relatively small, and for simplicity is listed here as comparable (or lower) to that of traditional vendor patching.
- ➔ The two biggest contributors to the total annual business impact when using traditional vendor patching — i.e., *lost revenue* and *lost user productivity* during the time that

enterprise databases and applications are disrupted for vendor patching — are substantially eliminated.

➔ The window of vulnerability from public disclosure to mitigation is substantially shorter than with traditional vendor patching, substantially reducing the likelihood that enterprise databases and applications may be compromised.

Summary and Key Takeaways: How a Virtual Patching Strategy Reduces the Risk of Database and Application Vulnerabilities

The issues raised above are among the many reasons that about one third (31%) of all respondents in Aberdeen's recent benchmark study have turned to **virtual patching** — a strategy that is particularly well-suited for the scenario of enterprise databases and associated applications.

A virtual patching strategy addresses both aspects of security-related **risk**, as risk is properly defined:

➔ It substantially reduces the *likelihood* aspect of risk, by reducing the window of vulnerability that enterprise databases and applications could be compromised when vendor patching is not available, not possible, not practical, or deferred to avoid cost and inconvenience.

➔ It substantially reduces the *business impact* aspect of risk, by substantially reducing both lost user productivity and lost revenue during the time that enterprise databases and applications are disrupted by traditional vendor patching.

11

For more information on this or other research topics, please visit www.aberdeen.com.

Related Research

Quantifying the Value of Firewall Management; October 2016

Bad Bots, Good Bots, and Humans: Quantifying the Risk of Bad Bots; September 2016

Virtual Patching and Database Security: An Effective Compensating Control; April 2013

The Virtues of Virtual Patching; April 2012

Author: Derek E. Brink, CISSP, Vice President and
Research Fellow, Information Security and IT GRC

About Aberdeen Group

Since 1988, Aberdeen Group has published research that helps businesses worldwide to improve their performance. Our analysts derive fact-based, vendor-neutral insights from a proprietary analytical framework, which identifies Best-in-Class organizations from primary research conducted with industry practitioners. The resulting research content is used by hundreds of thousands of business professionals to drive smarter decision-making and improve business strategies. Aberdeen Group is headquartered in Waltham, Massachusetts, USA.

This document is the result of primary research performed by Aberdeen Group and represents the best analysis available at the time of publication. Unless otherwise noted, the entire contents of this publication are copyrighted by Aberdeen Group and may not be reproduced, distributed, archived, or transmitted in any form or by any means without prior written consent by Aberdeen Group.