

Blue Lane PatchPoint® System
Securing Unpatchable Applications



Introduction

Not long ago, most organizations had a very different strategy for dealing with updates, hotfixes and patches from software vendors. Unless the patch addressed an area of concern specific to the organization, the update was considered optional and was frequently postponed, sometimes indefinitely.

This was for good reason. In the IT organization, business continuity is typically top priority. Software updates introduce change, and change introduces risk. So every patch from a software vendor would undergo some form of risk/reward analysis and if the benefit of the patch outweighed the risk, it would be installed.

Fast-forward to the present. In the current reality of the remotely exploitable vulnerability, a constant stream of patches is released from a wide variety of software vendors. Because of the proliferation of worms and viruses, business continuity is now threatened from both sides: the patch might cause a system outage and failure to patch increases the odds of a security compromise.

Where security patches are concerned, the risk/reward equation has been altered dramatically: the luxury of postponing installation no longer exists and, worse, the only reward is the elimination of a vulnerability. Adding yet another driver, the increase in security risk has resulted in regulatory pressure as well as directives from software vendors to patch systems quickly. This new reality is now accepted by most organizations; the burden of constant patching as part of software's total cost of ownership. But this new patching paradigm has left behind two large segments of the IT infrastructure: the hard-to-patch server and the unpatchable server.

Hard-to-patch and unpatchable servers exist in a variety of vertical industries as well as horizontal applications. An order processing system in the retail world is a relevant example because of the strictly regimented maintenance window procedure within that vertical. Voice over IP (VoIP) is a prime example of a horizontal application that is hard-to-patch because of restrictions by the VoIP software vendors around patching the underlying operating system. This paper seeks to provide a comprehensive list of these servers and, more importantly, outline current strategies for dealing with the problem.

Hard-to-Patch and Unpatchable, Defined

At the easy end of the "patching-difficulty spectrum," an organization might patch and reboot a server quickly and easily, perhaps even automating part of the process, and rarely encounter any disruption to business continuity. Patch management software is frequently employed to patch these systems with even greater efficiency.

Hard-to-patch servers, by comparison, require significant amounts of hands-on care during patching cycles and take significantly longer to patch:

- The server/application is part of a larger, more complex system where any change to one component could dramatically alter the functionality of the whole system
- The application itself is complex and special skills are required in order to deploy a patch

- A third-party application vendor must first test, validate, and then approve any patching of underlying operating systems
- The change management process is extremely rigid or the maintenance windows of the organization are extremely tight

Alternatively, “unpatchable” servers occupy the extreme end of the patching-difficulty spectrum, and simply cannot be patched at all without causing a significant outage or imposing a significant cost on the organization:

- The system is required to run 24x7 and cannot be rebooted
- The application is part of an embedded system
- The software vendor has declared “end of life” for the product and will release no more security updates

Hard-to-Patch/Unpatchable Horizontal Applications

- Oracle® is a prime example of a sophisticated application that requires a significant level of skill to patch properly. Oracle releases patches on a quarterly cycle and can create a significant amount of challenge for the DBAs tasked with deploying the updates. Compounding the problem is the fact that Oracle is often deployed within an organization to run mission-critical applications. The critical nature of these applications further complicates the effort to patch and results in more restrictive change management processes.
- Windows® NT is an application that, despite its end-of-life status from Microsoft®, is still widely deployed and heavily relied upon by many organizations. The primary obstacle to patching NT is the fact that Microsoft no longer provides security patches under standard support programs. In the absence of negotiating special support, organizations must either upgrade to a supported version or simply assume the risk of running unpatched systems. For many organizations, upgrading takes a long time and leaves them exposed until an upgrade can be planned and executed.
- VoIP systems have increasingly begun to replace traditional phone systems in corporations both large and small, primarily due to the advantages offered in terms of cost and ongoing maintenance. The leading vendors in VoIP share a few common characteristics that make patching these systems problematic. VoIP software relies on an underlying operating system to function properly and, additionally, is frequently integrated with the corporate email system. Because the telephone is widely considered a mission-critical tool, without which many businesses would cease to exist, VoIP vendors have stringent requirements with regard to patching the underlying OS and mail servers. The vendors themselves must approve the application of any patch and will even go as far as to void the service contract should the organization apply a patch without approval.
- Custom applications often present problems during patching cycles. While patching a server running commercial software may result in unintended consequences, the results of patching the underlying OS of a custom application are entirely unpredictable. Complicating the issue further is the fact that many custom applications have poor or missing documentation and, because the application is unique, there is no peer group to reference when problems are encountered. Custom applications that perform mission-critical functions are often considered unpatchable.

Hard-to-Patch/Unpatchable Systems in Vertical Industries

- Retail is an example brimming with hard-to-patch and unpatchable servers. In the United States, it is common for a retailer to close a majority of its annual business during the months leading up to and immediately following the Christmas holiday. Consequently, retailers have some of the most restrictive and lengthy periods between maintenance windows. Many retailers lock down critical systems (e.g. inventory management, order processing, etc.) from September until January and simply refuse to allow any patching for fear of disrupting business continuity. Also unique to the retail industry is the recent introduction of an initiative called Payment Card Industry (PCI) Data Security Standard, driven by Visa and MasterCard, which mandates rapid patch installation and can impose significant fines for failure to comply.
- Process manufacturing provides examples of systems that are required to run continuously without any possibility of disruption. Taking down a system from the process manufacturing floor means taking down the entire manufacturing process, which can result in a significant loss of revenue and an increase in costs. Patching in this type of environment is typically not an option.
- Healthcare is another example of an industry being squeezed from multiple directions. Health Insurance Portability and Accountability Act (HIPAA) regulations are placing increased pressure and scrutiny on the IT organization. Servers in the healthcare industry store some of the most sensitive, private data among all industries. Additionally, many healthcare organizations run proprietary applications on legacy systems that can be extremely difficult to patch.
- Call centers exist in a variety of vertical industries, including insurance, banking, travel and hospitality, retail and many others. In an effort to provide ever higher levels of customer service, many of these call centers operate on a 24/7 basis and field calls from customers all over the world. Much like process manufacturing, to disrupt these operations with a patch-related outage or to suffer a security event due to running an unpatched server would result in a significant problem for the organization. Additionally, many of these call centers utilize software that stores extremely sensitive customer information, increasing the risk of data compromise and placing these organizations under greater scrutiny from the public as well as regulators.

Strategies for Coping

The "Patch & Pray"

The concept of patch and pray is executed just as it sounds. The perceived security risk of the server remaining unpatched is high enough, or the pressure to patch the server (from auditors, software vendors, or internal mandates) is high enough, that the installation of the patch is performed regardless of the consequences. This is rarely an option, primarily because most of the systems described in this paper are so mission-critical to the organization. There are cases however where the consequences of not patching are considered too costly and the only recourse is to move forward with the patch and hope for the best outcome.

The “Don’t Patch, & Pray”

Firewalls are the tool of choice for establishing a network perimeter. The typical firewall deployment aims to prevent all traffic except for that which is known to be good. Firewall rules can establish what protocols and services are allowed to access the “inside” of the network from the “outside,” over which ports. Unfortunately, firewalls do little to prevent unauthorized access, intrusions and service interruptions.

Because a firewall is ill-equipped to deal with many of these threats, the intrusion prevention system (IPS) has become a much more important component of a layered security approach. Contrary to standard mode of operation for a firewall, an IPS is typically deployed to pass all traffic except for that which is known to be bad. In an effort to minimize disruption of legitimate traffic into and out of the perimeter, the IPS vendors must develop extremely precise signatures to identify specific exploits. The loose-grained rules of the firewall contrast sharply with the fine-grained signatures of the IPS. Together, these two solutions are effective at preventing the majority of Internet-borne threats. However, Internet-borne attacks represent only one of the many threats to IT infrastructure and the whole notion of a network perimeter has begun to erode.

Patch Proxy

Blue Lane™ Technologies developed PatchPoint®, the world’s first inline patch proxy, to solve the specific problem of server patching. PatchPoint sits in the network as close as possible to the servers. It is an inline device but not a perimeter security solution. PatchPoint provides instant patch protection for servers without requiring any patches to be physically installed on the servers hosting business applications.

PatchPoint relies on sophisticated application proxies to properly interpret server-bound traffic. Each application supported by PatchPoint (e.g. Microsoft Exchange®, Oracle DBMS, etc.) must be accompanied by an appropriate proxy. PatchPoint uses these proxies to transparently “listen” to client-server transactions and to establish when the conditions exist for a particular patch is to be applied.

In the same way that Blue Lane develops an application proxy for each application, an inline patch is developed for every applicable vendor patch on that platform. The inline patch, when applied to the network traffic, performs the same corrective action inline that the software patch would do on the server. The result is that PatchPoint is able to deterministically apply a fix to server-bound traffic without having to heuristically determine whether or not traffic contains malicious content.

The PatchPoint inline patch proxy is entirely architected around the functionality of the original software patch and how that patch would operate within the context of the overall software application. By porting that logic to the network however, PatchPoint is able to secure servers without incurring any of the inherent risk involved in the physical deployment of the patch. And because the inline patch proxy deals with software vulnerabilities at the root cause, rather than focusing on specific exploit traffic patterns or behavior, PatchPoint can fix the vulnerability instantly and permanently, no matter how many exploit variants may emerge in the future.

Summary

With all of the problems involved with patching Oracle servers, it's amazing that enterprises continue to be able to stay current with Oracle patch levels. Of all of the critical problems that IT staffs deal with during normal operations, rapid patch testing and deployment can now be removed as top priority. Considering the recent emergence of patch emulation technologies, these problems can be measured and controlled to release precious IT resources from ad hoc, rapid patch deployments. Users of the PatchPoint System leverage the ActiveUpdate service to maintain up-to-date ActiveFix[®] patch information, the PatchPoint Enterprise Manager for control and maintenance of the PatchPoint System and the PatchPoint Gateway appliance, which implements the inline patch emulation. Today with the Blue Lane Technologies PatchPoint System, IT staffs can fix vulnerabilities inline efficiently and feel confident that using the PatchPoint System has returned control and determinism to their critical server patch processes.

About Blue Lane Technologies

An inline patch proxy is entirely architected around the functionality of the original software patch and the way in which that patch would operate within the context of the overall software application. By implementing that application logic via the network, however, the patch proxy is able to secure servers without incurring any of the inherent risk involved in the physical deployment of the patch. And because the patch proxy deals with software vulnerabilities at the root cause, rather than focusing on specific exploit traffic patterns or behavior, it can fix the vulnerability instantly and permanently, no matter how many variants of the exploit may emerge in the future. The primary benefit with regard to the hard-to-patch and the unpatchable is that there is zero footprint on the server, zero downtime, and no rebooting required.

Due to the widespread proliferation of worms and other network-borne exploits, many organizations no longer have the luxury to delay patching. Security patches must be installed quickly, regardless of how "patchable" the target server is. With the ability to quickly eliminate vulnerabilities while preserving application availability, the inline patch proxy is the first viable solution for patching the unpatchable.

For more information about Blue Lane Technologies and for a complete list of coverage, please visit www.bluelane.com, or call 1-866-488-PROTECT.

Contact Information

www.bluelane.com
info@bluelane.com
10450 Bubb Road
Cupertino, CA 95014
1-866-488-PROTECT