



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

A Practical Methodology for Implementing a Patch management Process

Executive Summary

The time between the discovery of an operating system or application vulnerability and the emergence of an exploit is getting shorter, sometimes only a matter of hours. This imposes pressures on IT managers to rapidly patch production systems which directly conflicts with configuration management best practices of quality assurance testing. Many organizations are struggling to keep current with the constant release of new patches and updates. At the same time, they are under pressure to provide near 100% availability of key business systems. IT organizations must develop a process to ensure the availability of resources, install required security patches and not break existing systems in the process. This paper presents one methodology for identifying, evaluating and applying security patches in a real world environment along with descriptions of some useful tools that can be used to automate the process.

Understand the Risk of Patching vs. Not Patching

While it is essential to protect company IT assets from attack, patching vulnerabilities is only one part of the risk equation. A responsible system administrator must also look at the potential threat along with the vulnerability to determine the risk of having an unpatched system.

"Patch management is a subset of the overall configuration management process" (Colville, p.1). This means that an organization should have in place a strategy for establishing, documenting, maintaining and changing the configuration of all servers and workstations according to their function. Configuration management underlies the management of all other management functions: security, performance, accounting and fault. Fault is the management of device failures. Establishing a patch management plan can be considered a dress rehearsal for developing a configuration management strategy. Developing a risk management strategy goes hand in hand with creating a patch management plan. A risk assessment should be performed on all servers on the network. This assessment should include the criticality of the data on the server, the impact of server downtime on enterprise operations and the vulnerability of the server to internal and external attack. Risk management also affects the decision to apply patches and fixes. Rather than blindly applying every patch and hotfix that is released by vendors, a process should be developed to evaluate the criticality and applicability to the software patch. This is where configuration management, risk management and patch management merge. If a server's configuration is well documented, a decision as to whether a patch

needs to be applied becomes easier to make. The risk assessment as to whether to apply the patch should include the risks of not patching the reported vulnerability, extended downtime, impaired functionality and lost data. Anyone responsible for government IT security must follow the Federal Information Security Management Act (FISMA) of 2002. This act spells out the information security responsibilities of all agencies of the federal government. Section 301, subchapter III, paragraph 3544, subparagraph (b) spells out the responsibilities of federal agencies to develop, document and implement an agency-wide information security program. This is the section that addresses patch management through the following guidelines: implemented policies must be based on risk assessments, cost effectively reduce information security risk to an acceptable level, and ensure that information security is addressed throughout the lifecycle of the system. A patch management process that includes risk analysis and mitigation strategies, implementation of automated tools, and puts in place a repeatable process to maintain the patch level of all enterprise computing platforms will address all of these guidelines.

A good patch management plan consists of several phases. The plan outlined below consists of seven. The actual number and order of the phases may vary between organizations due to organizational size, structure or established procedures but the basic process is the same. Where appropriate, tools are identified to help automate some of the tasks.

Phase 1 – Baseline and Harden

Gather and consolidate inventory data on every server, switch, router, printer, laptop and desktop in the enterprise. Although this information can be collected manually, ideally an automated tool linked to a database should be used. This would enable data collection dynamically and help ensure that the data is always current as opposed to static information collected manually. Data to be collected should include hostname, location, IP address, MAC address, operating system and current revision level. For servers collect their function and services actively running.

Many inexperienced administrators accept the default options when installing operating systems. If documentation on what services were installed as part of the operating system installation is unavailable, consider running vulnerability scans against the server to uncover unnecessary services that should be disabled or removed. Use caution when securing workstations. Some organizations are overzealous in locking down desktops and only make distributing updates more difficult. For example, deploying the SMS client requires server service to be running, file and print sharing enabled and remote registry access enabled and running. These same items are often turned off or disabled while hardening desktops. There are numerous tools available that will scan systems for vulnerabilities and a few of them are described below. Many of these are free of charge while others are expensive.

Microsoft provides for free the Security Configuration and Analysis (SCA) tool as part of Windows 2000 and above. It can be launched from the Microsoft

management Console (MMC). The SCA can be used to compare the host configuration against a predetermined template. Microsoft includes several templates with Windows 2000 and XP or additional templates can be obtained from other sources such as the Center for Internet Security (CISecurity.org). The SCA will fix security holes however caution should be exercised before doing so. Make sure that the consequences of making changes to a system are fully understood and in any event, only implement changes one or two at a time and only on test systems first.

The Center for Internet Security (CIS) also has available free benchmarking and scoring tools for Windows 2000 and NT, Linux, Solaris and HP-UX. These are host-based tools and are designed not to impact the systems or applications of the host that they are running on. The tools will compare the security configuration of the test system against a CIS Benchmark for that operating system. The results are displayed in an easy to understand scoring report and detailed explanations of the meaning of the scores is provided. These tools are useful for identifying configuration weaknesses and getting all machines to a common baseline.

A free tool is available from Nessus (Nessus.Org) that will scan for security vulnerabilities on multiple flavors of Linux and Unix as well as Windows. Nessus is very powerful and easy to use. It will not make any assumptions about the server configuration. It will scan all ports for running services and attempt to exploit those it discovers. It is highly configurable through the use of plug-ins that are targeted towards specific vulnerabilities such as FTP, remote file access and DOS. Each plug-in has an even more targeted selection of specific vulnerabilities to choose from such as whether anonymous FTP is enabled or whether Solaris FTPd is configured to tell whether a user exists. Nessus takes about 2 hours for a competent administrator to get up and running. In a comparison test of seven vulnerability scanners in Network Computing magazine, January 2001, Nessus was the top scorer against several commercial scanners. Nessus found 15 of 17 vulnerabilities in the tests. Another key strength of the Nessus scanner was the fact that if it made assumptions a service that may not be entirely accurate, it warned of the assumption so that the administrator could investigate more thoroughly. Other scanners reported false positives requiring additional analysis by the operator. Through the use of the Nessus Attack Scripting Language (NASL) administrators can script custom probes and even attacks. The one noted weakness in the Nessus product is its weak reporting capability. However, this was conducted in 2001 and the current version should be significantly improved.

A different approach to vulnerability scanning is the QualysGuard Intranet Scanner. This product is an appliance which can be plugged in and configured in 15 minutes. Compared to Nessus, QualysGuard didn't report as many vulnerabilities in comparison testing done by Federal Computer Week, its reporting capabilities were considered superior. QualysGuard cost \$2,995 for the appliance as well as a licensing fee for the hosts.

Each server should also have an indication of its criticality to the enterprise mission. The higher the rating, the more mission critical the system. Factors to

consider when determining the mission critical status of a system would include: system role in the enterprise mission, impact on the mission of system down time and time and effort required for disaster recovery. The mission critical status translates into a risk level to the enterprise of the system being unavailable. This risk factor becomes important when making the decision of if, when and how to apply a patch. The servers in an enterprise can be divided into three environments:

- Mission critical – an environment in which even one hour of downtime will have a significant impact on the business service, and availability is required at almost any price. Examples would be e-commerce sites where downtime can translate into significant lost revenue and consumer confidence.
- Business critical – an environment in which business services require continuous availability, but breaks in service for short periods of time are not catastrophic. Examples would be payroll processing servers, E-mail servers.
- Business operational – an environment in which breaks in service are not catastrophic. Examples include print servers, file servers, E-mail gateways. (Radhakrishnan, p. 5)

Many organizations have situations where the responsibility for maintaining the server hardware and operating system falls on one group but the maintenance of the applications running on the server are the responsibility of another group. In this situation it is vital that proper change management procedures be implemented and adhered to. These servers should have standard hardware configurations as far as that is possible with the constant advancements in technology. For each server, develop a change control document. This document should contain the function of the server, the primary and backup point of contact including after hours contact information, any special procedures required prior to making a configuration change, and detailed disaster recovery procedures.

Patch managers should be aware of security precautions in place in their environment. If they do not personally manage the company firewall they should obtain configuration information from the firewall administrator. Ensure that there is available documentation as to what traffic is being allowed through to the internal network. This will help in the evaluation of threats posed by known vulnerabilities and assign a risk factor to them.

Once the data is gathered it should be documented and distributed to all system owners. Put in place a process to keep the data current.

Phase 2 – Develop a Test Environment

Once the environment is baselined, build a test environment that mirrors the

production environment. At a minimum, the test environment should have test servers representing all mission critical applications. Ideally, every type of platform in the enterprise should be represented in the test environment. In many cases, if applications are developed in house there should already be servers that can be used for testing security patches.

It may not be possible to maintain a test environment that mirrors the production environment, especially for small organizations with tight IT budgets. In this situation, patches should be deployed to the least critical, easily recoverable servers first. These would be servers without a lot of data or applications that need to be restored. An example would be print servers. These can be rebuilt quickly from registry backups. Ideally, the organization should have multiple print servers with the queues divided between them in such a way that if one fails, a user could find a print queue on another server that is in the same physical location as the one on the failed server. When installing patches on E-mail servers, update the gateway before the database server.

Personnel designated to evaluate patch stability should have expertise in mission critical systems and be capable of verifying stability of systems after patch installation.

One cost effective means of establishing a test lab is to use VMWare to create a "Lab in a box". While this method won't account for hardware variables in patch testing, it is a good way to test patch compatibility with the OS as well as any applications that are running on production servers. VMWare supports Windows as well as Linux operating systems. A replica of the production environment can exist on a single piece of hardware allowing the patch testers to evaluate multiple configurations of operating systems and applications and their interaction with each other before and after patch installation.

Phase 3 – Develop Backout Plan

Before any patch is installed, a full backup of all data and server configuration information must be made. Best practices for disaster recovery recommend periodic testing of the restore process to ensure the integrity of the backed up data. Create Emergency Repair disks for all servers after updating. This way, it won't have to be done before the next update.

When updating workstations, establish a group of test users who are the first to obtain the new updates. After successful deployment to the test group, expand to the rest of the enterprise. Users should be storing their critical data on network shares and have minimal desktop customization to facilitate rapid restoration from a standard image.

Phase 4 – Patch Evaluation and Collection

Keeping current with hotfixes and updates can be a daunting task. It is important to be able to quickly evaluate which updates are critical, which ones are merely useful and which ones are unnecessary. An automated tool makes this job a little easier by either maintaining a database of monitored systems and their patch status or scanning them on demand. These results are then compared to a database of the ideal configuration and systems needing to be updated are identified. Gartner Group has identified nine functional requirements that should be considered by enterprises that are considering automated solutions for patch management:

1. The solution should be able to create and maintain an inventory of server and desktop systems. It should be able to discover new systems without requiring the distribution of an agent.
2. The automated solution should be able to provide information about installed service packs and patches for the operating system as well as each major installed component.
3. It should be able to evaluate patch prerequisites. This will reduce the labor requirements of patch management.
4. The automated solution should maintain a current, dynamically refreshed inventory of patches and information about them. This will help the enterprise prioritize patch installations based on the criticality from a security perspective.
5. The automated solution should be able to report the patches that are needed by each individual server and workstation.
6. The automated solution should support role-based administration and system grouping. This allows the workload to be distributed among groups of system owners.
7. This may be obvious but automated “patch management tools should provide patch distribution and installation functions, including the ability to automate the installation of patches that require intervention”. (Nicolett, p.3)
8. Since Microsoft still dominates the desktop environment, most patch management solutions have greater Microsoft support. That is beginning to change and as will be described later, some are beginning to add Unix, Linux and even Novell support.
9. There are two types of automated solutions. Agentless architectures rely on scans of target machines to determine their update status. This type is easier to set up and configure but consumes more network bandwidth to push out patches. Agent based systems are more efficient users of network bandwidth and provide more functionality but they also have higher deployment and maintenance costs. However, effective patch management, especially “with respect to mobile users, is likely to require the functionality of an agent based approach” (Nicolett, p.3). Organizations should leverage as much as possible any established software distribution agents for patch management.

This is where all of the preliminary work will pay off. The next three phases can be broken down into 5 steps: receiving information on latest software updates and vulnerabilities; auditing the enterprise for applicable software updates; assessing and authorizing available software updates; deploying authorized software updates within the enterprise in a timely, accurate, and efficient manner; tracking update deployment across the enterprise. (Systems Management Server Version 2.0, Enterprise Software Update management Using Systems Management Server 2.0 Software Update Services Feature Pack, White Paper, p. 5). Tools are available for analyzing the current patched status of systems, downloading available patches from a central database and managing the installation of the patches. Some of these tools are Solaris Patch Manager Tool for SUN Solaris, Ximian Red Carpet Enterprise for Linux, Microsoft Systems Update Services (SUS) and the SUS Feature Pack for Microsoft Systems management Server (SMS) for Windows 2000 and up. These products all maintain a database of systems and installed patches, analyze patch dependencies, deploy approved patches to clients and track patch installation status. Some of them also provide a rollback feature to return to the previous version of the software in case of problems.

Microsoft SUS is fairly easy to get up and running in a Windows 2000 environment. The configuration usually consists of two SUS servers. One is used for downloading the patches from the Microsoft web site and deploying them to the test workstations. Once the patch stability is verified, they are copied to the production server and advertised to the clients. Windows 2000 machines with service pack 3 or greater and Automatic Update configured will then download and install the updates according to the settings configured by the administrator. In Active Directory enable domains, the client settings can be deployed through group policy. In non Active Directory environments, the client settings can be configured through registry changes deployed via the login script, Windows NT-4 style system policy or SMS if it is available. The limitation of SUS is that it will only distribute patches and updates available from the Windows Update site. These consist of security hotfixes and patches and service packs for the Windows operating system and related components such as Internet Explorer.

Enterprises using SMS have the option of employing the SUS Feature Pack for SMS. The SUS Feature pack has a few advantages over straight SUS. It can be used to distribute service packs and updates for Office applications as well as OS updates. It also gives the administrator more control over the distribution schedule as well as tracking the status of the client installations. The SUS Feature Pack uses the HFNETCHK scan agent, developed by Shavlik, to inventory current patch status of client machines.

Shavlik sells a GUI version HFNetChk called HFNetChkPro. HFNetChkPro differs from most patch management products in that it doesn't use an agent which makes it easier to install and manage. Like SUS and the SUS Feature pack, HFNetChkPro supports only Windows.

PatchLinkUpdate from Patchlink is a cross platform patch management solution. It supports Windows 95 through 2003, Novell NetWare, Unix including Linux, Solaris, AIX and HP-UX. PatchlinkUpdate is an agent based solution and in tests done by eWeek, was the most consistent in deploying patches across the enterprise. Patchlink maintains a database of patches released by OS and application vendors. They conduct additional tests of the patches in their labs before they make the patches available for download. For an additional fee, they will test patches against an image supplied by the customer. For heterogeneous environments, PatchlinkUpdate may be the perfect solution for managing updates.

Ximian makes Red Carpet Enterprise which supports only Unix based machines including the Red Hat, Mandrake, SuSE and Debian flavors of Linux as well as Solaris 8 and 9. Red Carpet users subscribe to “channels” to keep track of available updates. This allows users to monitor specific projects or collections of files beyond the standard security updates and bug fixes for essential packages. (Hall Linux Planet)

Remote workstations are the bane of most administrators. Keeping them current with anti-virus software is enough of a challenge without adding security updates to the mix. Most users are still using slow dial up connections to access the company intranet and they have little tolerance for delays while waiting for software downloads. Some IT organizations spend an inordinate amount of time trying to develop strategies for deploying updates to home based workstations. Other organizations are taking a different approach. They are using products like Citrix Metaframe for their remote users. Citrix is a client server solution where no data is transferred between the client and the server. Only keystrokes and video refresh data is sent over the network and all processing occurs on the server side. While this solution doesn't protect the remote clients, it does prevent any potential vulnerabilities present on the client machine from spreading through the network.

Phase 5 – Configuration management

After the patch has been tested and is ready to be deployed, the proposed changes to systems and the results of the testing should be documented and approved by system owners. The system owners should be prepared to standby in case disaster recovery steps are required. The helpdesk should be aware of the planned updates, any possible side effects and remediation instructions if users are affected. If automated systems monitoring is active, the appropriate personnel should be notified if any monitored systems will be going offline and triggering alerts. If any adverse events do occur during the deployment, the details of what occurred and on what systems should be documented and incorporated into future testing. And finally, capable personnel should be available to test systems after patch deployment.

Phase 6 – Patch Rollout

Once the patch has passed internal testing and configuration management review, it is time to deploy it. If one of the previously described tools is being used to monitor patch status and gather patches from vendors, it can also be used to distribute the patches to clients. Most of the tools have the ability to schedule patch distribution and don't require user intervention so that deployments can be done during off peak hours but even better, no one has to stay late to monitor them. Patching of mission or business critical servers should be done manually during off hours in case disaster recovery plans need to be implemented. If the patch is not an emergency fix, it can be applied during a regularly scheduled maintenance window. Make sure that the maintenance window allows for the recovery process if required. Patching of business operational servers can be accomplished through the use of the same tools as the workstations. Enterprises that don't have access to these tools will have to rely on alternate methods of patch distribution. They can utilize login scripts to deploy patches and free utilities such as HFNETCHK to report on the status. Or they can post the patches to an intranet site and provide users with instructions for installing them.

Phase 7 – Maintenance Phase – Procedures and Policies.

Maintaining the enterprise resources at current level is a function of establishing and following documented policies and procedures. Documented can't be emphasized enough because the policies and procedures must be able to survive staff turnover. Below are some guidelines to establishing patch management policies.

1. Designate patch management lead person or team. Ensure that they have support from top management and authority to get the job done.
2. Establish policies for patch updates. Non-critical updates on non-critical systems will be performed on regular scheduled maintenance windows. Emergency updates will be performed as soon as possible after ensuring patch stability. These updates should only be applied if they fix an existing problem that the server is experiencing. Critical updates should be applied during off hours as soon as possible after ensuring patch stability.
3. Establish procedures for checking for the existence of available patches, assessing the applicability of the patches and testing the patches. The more thoroughly the process is documented, the less vulnerable it is to staff turnover and loss of institutional knowledge. Ensure that the testing team contains members who are familiar with every application used in the enterprise.
4. Constantly update the workstation images for new PCs with the latest updates. Make sure that all workstations utilize a standard security configuration and don't prevent authorized access to install updates.

5. Provide regular reports for management. IT personnel can often enjoy more personal freedom if their management knows that they are on top of important issues.

© SANS Institute 2003, Author retains full rights.

References

- 1)) "Automated Vulnerability Remediation – The Cure for Security's Common Cold An Executive White Paper", The Aberdeen Group inc., November 2002
URL: <http://www.aberdeen.com/ab%5Fabstracts/2002/12/12023072.htm>
- 2) Benchmark and Scoring Tools for Windows 2000 Professional, Windows 2000 Server, Linux, Solaris and HP-UX.
URL: <http://www.cisecurity.org/>
- 3) Bishop, Vincil, Greer, Earl "Vulnerability Scanning: It's All About Control", June 9, 2003, Federal Computer Week
URL: <http://www.fcw.com/fcw/articles/2003/0609/tec-scan-06-09-03.asp>
- 4) Colville, R., Wagner, R., Nicolett, M., "Patch management Benefits, Challenges and Prerequisites", DF-18-0680, 4 November 2002, Gartner Group.
- 5) Harangsri, Banchong "Introduction to Nessus, a Vulnerability Scanner", 6/07/2002, LinuxSecurity.Com
URL: http://www.linuxsecurity.com/feature_stories/nesusintro-printer.html
- 6) Information Technology Security, Practices & Checklists/Implementation Guides.
URL: <http://www.csrc.nist.gov/pcig/cig.html>
- 7) Nicolett, M., Colville R., "Robust Patch Management Requires Specific Capabilities", Technology, T-19-4570, 18 March 2003, Gartner Group
- 8) "Patch Management Using Microsoft Systems Management Server - Operations Guide"
URL: <http://www.microsoft.com/technet/technet/itsolutions/msm/swdist/pmsms/pmsmsog.asp>
- 9) PatchLink Update 4.0 White Paper, Cross Platform Security Patch Management, 2002
URL: <http://www.patchlink.com/support/documents/PUW4.html>
- 10) Radhakrishnan, Ramesh "Patch Management Strategy for the Solaris Operating Environment" January 2003 Sun Documents,
URL: <http://www.sun.com/solutions/blueprints/0103/817-1115.pdf>
- 11) Sturdevant, Cameron "Patch Work Gets Harder", June 2, 2003, eWeek
URL: <http://www.eweek.com/article2/0,3959,1111839,00.asp>

12) Systems Management Server Version 2.0, Enterprise Software Update management Using Systems Management Server 2.0 Software Update Services Feature Pack, White Paper

URL: <http://www.microsoft.com/technet/prodtechnol/sms/deploy/confeat/smsfpdep.asp>

© SANS Institute 2003, Author retains full rights.