



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>

Patch Management, Getting Started

Introduction

Keeping track of the latest vulnerabilities and fixes that apply to your environment can become a complex process. By knowing how the software vendors communicate vulnerabilities and fixes related to their products and by taking advantage of the numerous mailing lists and newsgroups you can begin to simplify this process. Developing and documenting a standard process for tracking and deploying fixes can turn this into a routine function.

This paper will provide the reader with a starting point for managing service packs and security updates for Microsoft and Red Hat systems by defining the need to keep systems updated, identifying ways to stay current with security vulnerabilities and suggesting how to track relevant vulnerabilities. Additionally, the reader will be introduced to vendor, commercial and freeware options for deploying system updates.

A Little Background

It is estimated that the number of bugs in published software ranges from five to twenty bugs per 1,000 lines of code. Microsoft Windows 3.1 had approximately three million lines of code, Windows 2000 is estimated at 35 million lines of code and Red Hat Linux 7.1 has been estimated at 30 million lines of code. This means there could be anywhere from 150,000 to 600,000 bugs in the operating system you are currently using. Some of these bugs are responsible for the security vulnerabilities discovered on a daily basis while others are merely cosmetic and do not impact your system or security. Microsoft defines a security vulnerability as "...a flaw in a product that makes it infeasible – even when using the product properly – to prevent an attacker from usurping privileges on the user's system, regulating its operation, compromising data on it, or assuming ungranted trust."¹

Most software vendors, from the one person shop to the largest vendors in the world, are concerned about their software working as expected and try very hard to keep their software up to date and free of bugs. To support that fact, most vendors provide a contact person or e-mail address where the user can report new bugs. Once the vendor fixes the bug, or plugs the security hole, they will distribute a patch, a service pack or possibly even a new version of the software.

Since service packs and patches have completely different scopes it is important to understand the differences. According to Microsoft, a service pack is a

scheduled periodic update that corrects a number of problems in one version of a product. For example, there have been six service packs for Windows NT 4.0 since its original release. Some of Microsoft products use the term service release rather than service pack, but the terms mean the same thing. A patch, sometimes referred to as a “hotfix”, is an unscheduled update that occurs between service packs. Most patches are coded to correct specific security vulnerability, but may also be used to correct critical stability or performance issues.

It is important to know the software life cycle of your vendor. Microsoft generally develops service packs only for the current and next-to-current versions of a product. Patches are maintained for the current service pack. Patches for the next-to-current service pack is maintained up to twelve months after the current service pack is released. The following table shows the current time line of security patches for Windows desktop software.

Desktop Operating System	End of Security Patch Support
MS DOS x.xx	December 31, 2001
Windows 3.xx	December 31, 2001
Windows 95	December 31, 2001
Windows NT 3.5x	December 31, 2001
Windows 98/98 SE	June 30, 2003
Windows NT 4.xx	June 30, 2003
Windows Millennium Edition	December 31, 2004
Windows 2000 Professional	March 31, 2007
Windows XP Professional	December 31, 2008
Windows XP Home Edition	December 31, 2006

Red Hat will discontinue support for versions 6.2 and 7.0 on March 21, 2003 and continue supporting versions 7.1 through 8.0 until December 31, 2003. Red Hat has committed to support the Red Hat Linux Advance Server at least until May 31, 2005.

New patches usually require the most recent service pack be installed before applying the patch. If you are using a product or service pack that is no longer supported, a patch may not be available, leaving your system open to a particular vulnerability. This is another reason to keep your software up to date.

Stay current on alerts

How do you become aware of a new patch or vulnerability? There are various avenues available for keeping track of new alerts and available fixes for the operating system and applications.

The notification process you choose depends on the complexity of your network. To improve your chances for getting reliable information you should use a combination of the following options.

- **Vendor websites and mailing lists** are probably the most popular. The patches are released by the vendor and are considered reliable when downloaded from their site.
 - Microsoft's current process is to post the patch to the Microsoft Download Center or the Windows Update site, publish the Knowledge Base article on the Microsoft Web Site, and post the bulletin on the Microsoft TechNet Security web site. They also send the bulletin to over 100,000 subscribers via their free Microsoft Product Security Notification Service.
 - The main source for Red Hat alerts is the Red Hat Network (RHN). Once you register your system profile with Red Hat, an e-mail alert will be sent to your registered address containing information on the latest patch available for your system.
- **Third-Party websites** often report vulnerabilities before the vendors do. They cover a variety of products and may offer alternative solutions. These websites however may not contain all of the details for patching your system and may contain a "quick fix" or stop gap solutions rather than the vendor's approved solution.
- **Third-Party mailing lists and Newsgroups** allow interaction with other system administrators and the opportunity to take advantage of their experiences. Information learned from these sources should be validated for correctness since the chance of obtaining wrong information is greater. One of the more popular lists is Bugtraq monitored by SecurityFocus (<http://online.securityfocus.com/archive>).
- **Vulnerability scanners**, such as the freeware Nessus (www.nessus.org) or the commercial product Internet Scanner (www.iss.net), provide instant feedback by identifying vulnerabilities on a particular host or a list of hosts.
- **Vulnerability databases** such as the NIST ICAT index (<http://icat.nist.gov>) can provide a wealth of information and are usually the quickest to report new vulnerabilities.
- **Other resources include** Microsoft's Critical Update Notification application and third party tools including the Cassandra tool (<https://cassandra.cerias.purdue.edu/main/index.html>). Cassandra is a customizable search tool designed to retrieve information from the NIST ICAT index and send you e-mail alerts based on a criteria you have established.

Tracking the fix

Now that you are aware of the vulnerabilities, how do you keep track of which patches to implement and when?

Not only is patching your system important, the speed and manner of how you do it is just as important. A quick look at the Code Red worm may help encourage you to act as quickly as you can. In June of 2001, a vulnerability was discovered in Microsoft's web server, Internet Information Server (IIS). Within days Microsoft released a patch to eliminate the vulnerability. In July of 2001 the Code Red worm was released and infected more than 300,000 computers because the patch was not installed in a timely manner. Within two months the Nimda virus (admin spelled backwards) was released and took advantage of the same hosts affected by the Code Red worm. If the administrator of these systems had a systematic process for updating their systems, they would not have been hit by these two exploits.

You should start with a current inventory of systems on your network. Before you can track which systems to maintain, you need to know which systems are on your network. At the very least, your inventory should contain the operating system including the version currently installed, the most recent patch applied, critical applications running on your system, the owner and contact information for the system and a priority code. For consistency purposes, the priority code could be based on the restore order of your disaster recovery process.

Depending on the size of your network, you may want to create a tracking database or simply use a spreadsheet. In a separate table or spreadsheet, create a means for tracking the vulnerability. This table should contain the name of the vulnerability and the severity level, as issued by the vendor, the application and version affected and a description of the vulnerability or fix. It may also be helpful to store a link to the vendor's notice and fix as well as storing a copy of any files required just in case you are unable to get to them when you are ready to apply the fix.

Once you are aware of a new vulnerability or patch, the pertinent information should be entered into your tracking system and compared against your inventory. With the information recorded, you can generate a report or search for devices that may be affected. Based on the severity of the vulnerability and the priority code assigned, you now have the information available to formulate your plan for applying the fix or workaround.

Test the fix

Once you have downloaded the patch or service pack, scanned it for viruses and have read all of the accompanying documentation, it is time to head for the computer lab or a test machine. Before a patch is rolled out to the production systems, you should test the installation process and the patch itself on a similar system in a controlled environment whenever possible. This time should also be used to understand what the patch is actually doing. Reading the patch documentation will provide a good understanding of what happens when

applying the fix. Does it correct the vulnerability? Are there any software dependencies or conflicts? Furthermore, you should not assume the system is safe and reliable after the patch or service pack has been applied. It is possible the patch may have opened an old vulnerability or even created a new one. To check your system for old, read as “known”, vulnerabilities, you should test the system by running a vulnerability scanner against it. Also, did the fix reduce the reliability of the system or possibly degrade performance? You should be able to validate this in the lab and by consulting the mailing lists or newsgroups for potential issues.

Deploy the fix

So everything looks just fine. You learned about a new vulnerability from one of your many resources, you’ve researched it and identified which machines in your network are affected. You’ve researched the patch in the newsgroups and on the Internet to understand what it does and installed the patch in a test environment. You have verified the system is stable and no longer vulnerable. Now you are ready to deploy the patch.

If it is a server or a critical system that needs to be fixed, more than likely you will schedule some down time to backup the device, install the patch then test the system before bringing it back on line. Before you patch any Microsoft servers, you should review Microsoft’s service pack installation and deployment guides for Windows 2000 and Windows NT. What if the patch needs to be deployed to a number of your workstations on your network? Depending on the size of your network, this could become a full time job.

Software vendors, as well as third party companies, understand the importance of keeping your systems protected and have developed many options for deploying the fix. Additionally, It is not uncommon these days for an application to have a feature included that will connect to a predetermined site, check for software updates, notify you of a new version and ask if you want to install the update. Microsoft’s automatic update feature is one example. These automatic updates may be enough to manage a home system or a small network but will make it difficult to keep a larger network current and consistent.

Microsoft Options for Deploying the Fix

Microsoft has a number of options for automatically deploying updates to multiple servers and workstations. They include the Windows Update Catalog, the Microsoft Update Service and the previously mentioned Windows Update feature.

All versions of the Microsoft operating systems since Windows 98 include the automatic update feature called “**Windows Update**”. Using the automatic update icon in the control panel, you can configure windows update to notify you when a

new update is available. This requires you to run the windows update application in order to obtain the update. The second configuration option is to have the update automatically downloaded in the background to your system after which you are notified of the update. Windows update minimizes the impact of your online experience by only downloading the updates when the bandwidth is available. The third option is to disable windows update, which requires you to manually check for updates. Unless you can remember to check for updates at least daily, this option is not recommended.

For larger networks with centralized administration, you can use **Windows Update Catalog** to download the patches that only apply to your environment. This option is recommended for the more experienced administrator and is used for obtaining updated system files, service packs, new Windows features, and device drivers for Windows 98, Windows Millennium Edition, Windows 2000, Windows XP, and Windows .NET Server 2003-based computers. Note that Windows Update Catalog does not support Windows NT. Once a customized search of operating systems or drivers is configured, the administrator can select which updates to download by adding them to the download basket. Once patches are retrieved, the administrator can take advantage of an existing software distribution method such as Microsoft's SMS or Novell's ZenWorks.

If you do need to manage multiple machines and are not currently using a software distribution solution you can take advantage of the **Microsoft Update Services** (MUS). This service makes deploying the updates automatic, reliable and manageable. MUS is similar to the automatic update feature except the workstation receives the update from a predetermined server within your network. You configure the server to get updates from the internet then configure your workstations to get updates from your designated server. This option enables you to test the update before deploying it to all of your workstations. MUS also has the ability to track which workstations have installed the updates you made available. Unfortunately Microsoft Update Services only updates Windows 2000 operating systems and higher.

Red Hat Options for Deploying the Fix

Red Hat's equivalent to Microsoft automatic update is called the "Red Hat Network Alert Notification Tool". It is a GNOME applet that appears on the panel and alerts when software package updates are available. In order to take advantage of this applet you must register a system profile with Red Hat and activate the entitlement. Entitlements are basically subscriptions. The first subscription is free; any additional systems require a fee. If you do not want to pay the small fee of \$60 per system, you can actually manage a small number of systems by activating and deactivating entitlements. Although, the time it takes to shift the entitlements from system to system before downloading the updates clearly justifies the cost of the subscription, not to mention the priority access and other benefits you receive with the subscription. If you are not running GNOME

on your system, you can still take advantage of the up2date feature from the command line; a subscription is still required.

Third-Party Options for Deploying the Fix

The vendors of the operating systems are not the only resources available for monitoring your systems patch level or deploying updates. There are many commercial and freeware products offered.

Commercial

The cost of the commercial packages can drastically vary, depending on the number of devices you need to manage. Network Computing Magazine reviewed five products on a 1000 host network, the cost of the software ranged from \$11,000 to \$30,000.

There are basically two types of architecture used to manage the service pack level and software version for servers and workstations. The non-agent based applications work by scanning the hosts to determine which service pack and hotfixes are installed. This requires the scanning system to have domain or local administrator access on the system further restricting the types of networks these applications can support. If you have a large number of traveling or remote users, non-agent applications may not be the best option due to the number of devices that will be missed during the scan. This could potentially increase the number of scans required to check all devices.

The other option is agent based applications. These work by running a small application in the background that periodically polls the patch server for any updates. This is the better option for users who are not continuously connected to the network. Since the agent must be running on the system, some upfront work will be required to install the agent on each host.

When selecting a commercial product consider an application that allows you to group your workstations by operating system. This will help control the deployment of patches to similar systems and make it easier to troubleshoot version related issues.

Freeware

Just as there are many commercial products available for the windows environment, there are also various freeware products available for the Linux world. One of the products tested by the author is called AutoRPM by Kirk Bauer. AutoRPM was very easy to install. By default, it is configured to contact a mirrored site once a day for current updates. All new files are downloaded to a predetermined location for future installation; the configuration can be changed to

have all updates applied immediately after download, including kernel fixes. Obviously you would perform this on a test machine first or rely heavily on your backups should something go wrong. Once the new updates are downloaded, the user can list all of the downloaded files. The list will indicate whether the files are either updates to files already installed on the system or new versions of software not currently installed on the system. At this point the user can issue a command to install the updates followed by another to remove all files from the predetermined directory. AutoRPM can also be used to keep other systems on the network current.

Conclusion

As you look around your environment and observe all of the computerized systems used on a daily basis, it is not difficult to recognize the importance of maintaining those systems. As we continue to rely on the ability of software vendors to make our lives easier, it is important we are aware of the opportunities for others to exploit bugs within the software. One of the most important aspects of maintaining a reliable system is to keep the system software current by applying the recommended service packs and patches expeditiously. The first step is to develop a repeatable process to track, test and protect against new and existing vulnerabilities.

¹ <http://www.microsoft.com/technet/columns/security/essays/vulnrbl.asp>

References

Bauer, Kirk. "AutoRPM." URL: <http://www.autorpm.org> March 2002. (14 December 2002).

Culp, Scott. "The Definition of a Security Vulnerability." December 2000. URL: <http://www.microsoft.com/technet/columns/security/essays/vulnrbl.asp> (27 December 2002).

Culp, Scott. "Why Service Packs are Better Than Patches." URL: <http://www.microsoft.com/technet/columns/security/essays/srvpatch.asp> (27 December 2002).

Mell, Peter and Miles C. Tracy. Department of Commerce. Procedures for Handling Security Patches. Washington: GPO, 2002.

Microsoft Corporation. "Before Installing a Windows NT Service Pack" October 2002. URL: <http://support.microsoft.com/default.aspx?scid=kb;en-us:Q165418&sd=tech> (12 December 2002).

Microsoft Corporation. "HOW TO: Download Windows Updates and Drivers from the Windows Update Catalog" December 2002. URL: <http://support.microsoft.com/default.aspx?scid=kb;en-us:323166> (26 December 2002).

Microsoft Corporation. "Microsoft Windows 2000 Service Pack Installation and Deployment Guide." 2001. URL: <http://www.microsoft.com/technet/prodtechnol/windows2000serv/support/spdeploy.asp> (12 December 2002).

Microsoft Corporation. "Windows Desktop Product Life Cycle Support and Availability Policies for Businesses." October 15, 2002. URL: <http://www.microsoft.com/windows/lifecycle.mspx> (12 December 2002).

Mueller, Patrick. "Patchlink Helps Keep Windows Closed." Network Computing September 2002 (2002): 77-88

Red Hat, Inc. "Red Hat Network Basic: User Reference Guide 3.3." Chapter 5: Red Hat Network Alert Notification Tool. 2002. URL: <http://www.redhat.com/docs/manuals/RHNetwork/ref-guide/applet.html> (10 December 2002).

Red Hat, Inc. "Errata: Security Alerts, Bugfixes, and Enhancements. 2002. URL: <http://www.redhat.com/apps/support/errata>. (10 December 2002).