



ANSA

**Autonomic Network System Administration
Whitepaper**

Automating Patch Management

Information in this document is subject to change without notice. This document may be distributed freely, but must remain whole and no changes are allowed without written consent of Autonomic Software, Inc.

All trade, corporate names and products may be trademarks or registered trademarks or registered trademarks, and are used only for identification, without intent to infringe.

Please see <http://www.Autonomic-Software.com> for more information about Autonomic Software, Inc.

The Problem

Patch management for years has been a low visibility, labor-intensive job. However, in today's climate of mission critical computer systems and security breaches, patch management has come to the forefront. A global survey of 4,900 Information Technology professionals across 30 nations, conducted by InformationWeek Research and fielded by PricewaterhouseCoopers LLP, estimates that some 50,000 firms in the U.S. are sufficiently large enough to be impacted by and are able to accurately compute the cost of a software virus. In total, the bill to these U.S. firms this year for viruses and computer hacking will amount to \$266 billion, or more than 2.5% of the nation's Gross Domestic Product (GDP). The price tag worldwide soars to \$1.6 trillion.

In a symposium on Terrorism and Business at DePaul University, Ambassador Michael Sheehan spoke prophetically of these changing threats. "...We are going to need to look at these changing threats because I also think we are going to be vulnerable on a commercial front. We talk about this threat because when enemies cannot attack American military abroad, they may attack us at home in our critical infrastructure or private sector, as well as abroad. As it becomes more difficult to attack an embassy or military base, terrorists may go increasingly after commercial interests..." (Michael Sheehan, "International Terrorism: Trends and Responses", (2000)).

Network Exposure and Patch Management

It is an overwhelming task for the System administrators to keep systems up to date and ensure that applications and operating systems are current. In order to stay current, the System administrator must deploy patches to systems that are vulnerable. These patches fall into two broad categories: major revisions, and updates.

The number and complexity of patches challenge administrators. Is it a homogenous or heterogeneous environment (i.e. more than a single operating system to support within the network)? Within that network, are there different groups of computers requiring different patch levels? Is there conflict between patches? Can patches be deployed on top of one another? Numerous questions need answering; meanwhile, the administrator still needs to test and program the patch deployments. In the interim, networked systems are vulnerable to security threats. Systems need to be patched, but patch management is a full-time job requiring specialized knowledge.

The Increasing Number of Patches and Service Packs

A software patch often repairs a known problem (bug) in a program, whether in an operating system or an application. Some of these bugs cause little or no concern (i.e., software does not “hang” and systems are not vulnerable as a result). Other bugs cause catastrophic failures, including loss of mission critical data. Additionally, even coding styles may cause unintended consequences; e.g. buffer overruns can often be exploited, enabling a rogue process to take full control over a machine.

Once a bug is detected and a solution is created, it can often be downloaded from the software manufacturer’s web site. The problem may be as simple as a correction to a typo, or much more serious such as a security vulnerability or system instability. In the case of a serious performance issue, the patch may replace core system files, with possible interactions or conflicts with other files. The patch may also alter the system’s registry or intrinsic files, with other potential ramifications.

Major releases (Microsoft “service packs”), which are released much less frequently than updates or interim patches (Microsoft “hotfixes” or “updates”), include improvements and enhancements to operation or performance of a program, in addition to including most updates issued since the last service pack or major release. These changes are nearly always tested by the vendor to ensure that the patch will not cause more problems than it solves. Service packs are essentially new OS deployments and require much more thorough planning. On the other hand, hotfixes are typically unplanned deployments which happen frequently and are not regression tested. It is not safe to assume that a company can ignore periodic updates, and only implement patch procedures when major releases to come out. The reason is that many interim updates are released to repair a newly discovered security hole. If ignored, serious consequences may result.

The Patch

A major problem with patches is keeping up to date on available patches, sources of new patches, and retrieval / distribution of said patches. To keep current, administrators must periodically check for updates at the software manufacturer's web site. This, by itself, is a time-consuming activity. From the web site, an administrator is supposed to be able to determine whether a particular patch or update is required to be deployed onto machines within their network. Unfortunately, that theory is seldom the case.

The problem oftentimes is that there are too many patches to keep track of. With thousands of software patches it is impossible for IT and security professionals to be knowledgeable enough to ensure quality deployments. Installing patches that haven't been properly tested can cause more problems than they correct, yet not installing patches is too risky. This is the conundrum network administrator's face.

Add to this the fact that not all patches are required to be installed on all machines, which forces administrators to keep track of different configurations and determine either: is a patch necessary, or, if deployed, will it cause problems.

Patch Management and the Effect on your Business

Administrators investigate what patches are available, determine which machines are affected, test the patch, deploy the patch, and ensure successful deployment. Many times, updates are frequent and are not announced prior to their release, which makes planning for them next to impossible. When an urgent patch is made available, the administrator must drop everything to manage the situation. Obviously, this has a negative effect on the administrator's other job duties and time management.

Accompanying the normal risks associated with updating systems, an urgent patch compounds the normal risks associated with patch deployment (disruption, potential downtime associated with reboots, etc). All patches are capable of causing a conflict or making the patched product unstable. Accordingly, no matter how urgent a patch may be, it still needs to be tested prior to deployment.

Analyzing, testing and deploying patches on numerous systems over an enterprise is a very time consuming process. There is initial time for considering available patches, decision making in what to deploy to the network, actual deployment time and effort, and follow up to confirm that the patch is serving its purpose properly, while causing no ill effects. Time management is critical in this process, for the more time it takes, the larger the vulnerability window.

Network Vulnerabilities

Simply using anti-virus software or firewalls will **not** protect your network. It is estimated that 98 percent of successful attacks are a result of software bugs or other vulnerabilities that could have been prevented by applying available patches. Modern worms such as “Code Red” and “SQL Slammer” do not rely on any of the methods of transmission guarded by most virus protection systems. These new strands of viruses are designed to attack the computer system directory by exploiting faults in the software used by the computer to perform its operations. The viruses are therefore able to crack corporate networks and replicate without the intervention of anti-virus software.

A valuable lesson was learned recently in Boston, Massachusetts. In an article by Paul Roberts IDG News Service, 08/14/03 Titled “At a Boston hospital, Lessons learned from Slammer”, Roberts reports: "I'm proud to say we don't have a single copy (of Blaster) in the hospital," said John Halamka, CIO of Beth Israel Deaconess Medical Center (BIDMC) in Boston, a Harvard University research institution. On the server side, the IT staff held what Halamka called an "all nightmare-athon" patching session in late July, applying the relevant Microsoft patches to the hospital's 130 Windows servers. It was also a marked contrast to the scene at BIDMC 7 months earlier after the SQL Slammer worm crippled the hospital's computer systems for about six hours, forcing medical staff to resort to paper-based records to track patients.

In an article written by Joel Snyder Network World, 09/29/03 titled “When will we ever learn?” Joel compares the two-year old Code Red virus and two new worms. Joel states: “Two worms that hit this summer, W32/Blaster (also known as W32/Welchia, W32/Nachi and Lovsan) and SoBig (also known as SoBig.F) spread exactly the same way. Microsoft published bulletins, but people ignored them. Patches were issued, but no one applied them. The worms came in through firewalls that shouldn't have let them in. Infected systems continued spreading the worms because we didn't have adequate tools to contain them. Two years after Code Red, there are still fundamental problems in the way we manage and secure systems that make us vulnerable to this kind of attack.”

As IT professionals learn from previous mistakes, they are looking for ways to automate Patch Management so they can proactively manage their networks, instead of constantly being in reactive mode.

Flaws in current Patch Management

There are many reasons why Patch Management isn't working at a majority of companies today.

- **If it isn't broken, don't fix it.** While this attitude works with some of life's difficulties, it is a recipe for disaster when it comes to network security. By the time the network becomes broken, it has already had a major impact on your business.
- **People are already spread so thin with their daily work activities that they don't have time to keep abreast on available patches.** It is a daunting task to try to stay on top of all patches available for your enterprise, yet the danger of not doing so is immense.
- **There is no need to patch the non-critical systems.** Many viruses such as the Slammer worm now target non-critical systems and are intended to cause wide-scale disruptions rather than attacking a company's high profile systems.

Patch Management Solutions

Traditionally, patch management falls into three categories: manual solutions, in-house/legacy solutions, and third party solutions.

Manual Solutions

A manual solution involves employees researching what patches are needed and patching servers and workstations one by one (“Sneakernet”). This solution is very slow and prone to errors. Many patches and / or machines can be overlooked. The manual solution is just that; manually deploying patches by traveling from machine to machine is no way to handle a corporate network.

In-House/Legacy Solutions

Many companies that started with a manual solution soon created an in-house solution. An in-house solution normally requires an “in-house expert” who has programming skills and extensive knowledge of the solution. In-house solutions are usually rigid and take extensive programming to maintain and enhance. The unintended absence of the aforementioned “in-house expert” can cause disastrous consequences.

Third-Party Solutions

Third-Party solutions allow companies to automate patch management, which will save money and reduce the risks involved. For mid to large sized companies, automated tools from third parties may reduce both risks and costs associated with patch management.

Manual and in-house/legacy solutions are costly, error prone, and time consuming ways to perform patch management. Patch management should be easy to administer,

responsive, flexible, reliable and cost effective. Automating patch management will allow a company to become more proactive in their network security.

Making the Decision

According to Gartner Research (April 24, 2002 research report), the following questions should be asked when deciding whether to adopt a third-party solution:

- What is the security risk involved if patches are not applied?
- What will it cost to stay abreast of the latest patches and download them to each system?
- Would it be less expensive to deploy patches from a central server rather than manually installing them on each machine?
- Can manual installation be easily validated to determine that the patches have been installed and are properly functioning?

When evaluating which third-party patch management tool to employ, companies should address broad issues relating to flexibility and reliability as well as feature sets. There are some broad issues to keep in mind when assessing alternatives: (1) whether to use an agent versus a non-agent based solution, (2) the scope and reliability of the external management database that supports the product, (3) cost of the product, (4) ease of deployment and administration.

Supporting Disconnected Networks

Since patch management solutions require regular updates from external sources for critical patch data, it is impractical to support systems from a network with no connections to the outside world. ANSA offers an exclusive feature that allows administrators to replace being connected to the outside world with portable “packaged updates” functionality. This means that the administrator can manage machines orphaned from the Internet by using centralized controllers, a web interface, and a database hosted by ANSA that contains all patch materials pertinent to the customer.

The Patch Management Database Issue

Most patching tools use a limited set of patches and patch information supplied from public sources (such as MSSecure.XML). Free patch detection tools, such as HFNetChk available from the Microsoft web site, and most third-party patch management solutions, use a common database of patches. This database is supplied to the public and is limited in scope, as not all critical patches are included, nor are all applications supported. Another shortfall is that public databases do not include the ability to add proprietary patches. In addition, testing is limited and does not include analysis of patch sequence, prerequisites, co-requisites or conflict issues related to coexisting patches shared by more than one application.

The patch management databases of third-party vendors fall into two groups:

- Those that use a public source database and test just one patch at a time.
- Those that have a third-party patch database and do comprehensive, independent testing on a combination of patches and deployment scenarios.

When assessing patch management tools, companies will want to know what type of patch database (public or private) supports the tools as well as the mechanism through which patches are tested before being deployed.

Benefits of ANSA™ Patch Management

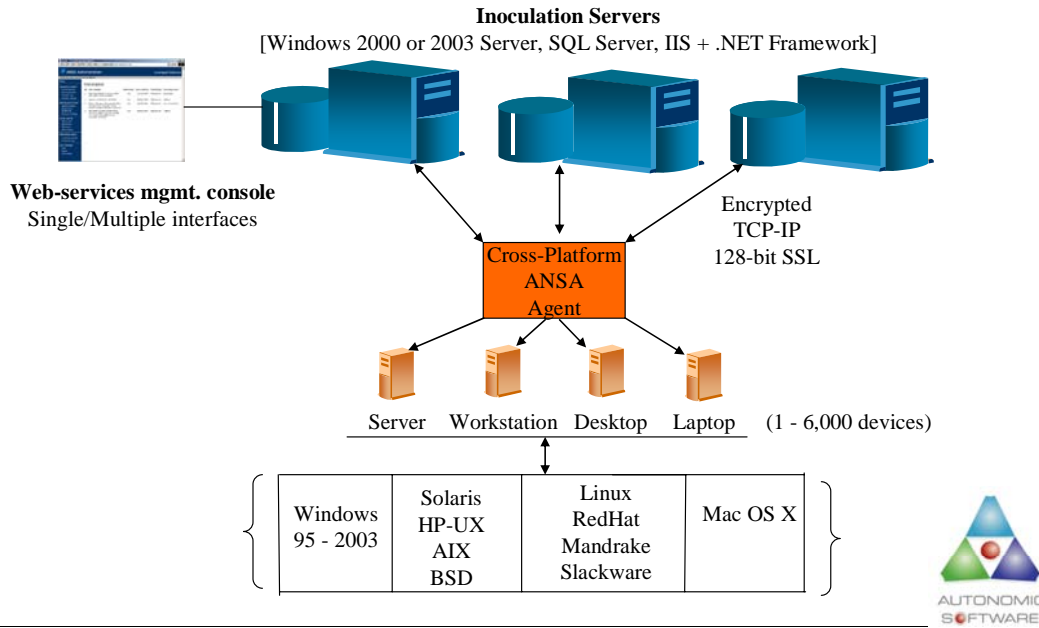
ANSA™ Patch Management automates the patch management process across multiple platforms, yet still allows the network administrator the flexibility to pick and choose which patches he would like to deploy. ANSA's™ Global Update Repository (GUR) is a hosted database containing up-to-date patches for multiple operating systems and applications, allowing the administrator to easily manage the patches available for the entire enterprise. Through a single interface, the administrator can control the entire network, even remotely if needed.

By using ANSA™ to automatically inoculate systems before viruses are able to take advantage of their weaknesses, corporations can prevent many of the modern invasive agents from entering their network, and therefore, reduce losses incurred from such entry. Furthermore, it is a well known fact that a sufficient amount of network and system administrator time is utilized on keeping track of security fixes, downloading these patches, and applying these fixes across the corporation. Corporations that utilize the ANSA software will receive the benefit of recouping this time, allowing their employees to spend it on other important network and system administration duties, thus realizing a direct increase in IT productivity and a corresponding reduction in expenses.

Flexibility of ANSA™ Patch Management

ANSA™ Patch Management is an extremely robust solution in that it is installed enterprise-wide (to Microsoft based clients as well as non-Microsoft based clients) across ALL platforms. It can be deployed either automatically or manually, creates configurable variations of audit reports, and requires few system resources, thereby lowering the load it places on networked computers. All scheduled events are under the control of the system administrator through a friendly user interface. Administrators have control over the software patches that will be installed at the lowest granularity – applications and operating systems running on individual Inoculation Clients (workstations). Reports are configurable, as are the sources from which patches and updates will be retrieved. Grouping of computers is also configurable – ANSA allows the system administrator control over how patches are deployed to functionally different groups of clients, defined by the system administrator. Policies at the client level can be established; for instance, a specific client can have rules that will only allow patch installs to happen during off-hours or while specific applications are not active. As the system administrator creates these various policies and groupings, ANSA automatically deploys updates in accordance with these policies.

ANSA's Distributed Architecture



How ANSA™ Patch Management Works

ANSA is composed of four basic components: The Global Update Repository Server, the Inoculation Server, the Inoculation Client(s), and a Web Management Interface.

ANSA™ Patch Management works by using a Spider to scan the various reporting services and application manufacturer's websites for recently created security upgrades, "hot fixes", and service packs. The Spider then retrieves these patches into a Global Update Repository (GUR). A key feature of the GUR Spider is that it is not limited to Microsoft patches and updates; this Spider mines, retrieves, and archives external updates from a multitude of other sources. This feature differentiates ANSA from most other offerings in this space. The GUR is a hosted solution maintained by the ANSA GUR support team. Additionally, ANSA Professional Services can be engaged to replicate a user environment; i.e., the ANSA GUR team has the ability to provide mirrored environments in order to test patches which are to be subsequently deployed to licensed ANSA users. When a patch is made available to ANSA by a vendor, turnaround times under normal circumstances will not be greater than 12 hours (unless additional testing is dictated by the licensed user in question).

The Inoculation Server (IS) then compares the company's information to the GUR information (via XML) to determine whether the patch is applicable to any machines on the network which have the ANSA agent running on them. The Inoculation Server runs on a Windows 2000 or 2003 server with .NET Framework 1.1 installed, either SQL Server or MSDE, and Internet Information Services (IIS). The IS components are comprised of:

- An internal inventory control engine, which is used to scan for application updates within the GUR, comparing them with the configured clients through the client status report.
- A distribution engine is notified by the inventory control engine, which schedules external package and patch installations, recording the status of all client updates upon receipt of status information from the individual clients.

The Administrator then manages and schedules the patches they want to apply to their enterprise through a web-based user interface. For performance reasons, patches are cached on the server and deployed to clients from that server, not from the source of the patch (typically a vendor site). This eliminates issues with high traffic constraints as well as bandwidth considerations when downloading patches to client machines.

The Inoculation Client (IC) then queries the IS to detect any patches scheduled for download. The IC communicates back to the IS with the status update. The Inoculation Client is a very small (around ½ megabyte) application that is installed on any number of servers or workstations across an organization's network. The IC is an agent designed to communicate to the IS its own operating system type and version information, as well as all pertinent installed applications which would be affected by available OS and application patches. The client is installed as a service for Windows environments and as an application for all other environments.

The IC communicates to the IS in the following scenarios:

To initiate connection with the IS to provide initial, as well as ongoing, OS and application software information. The communication mechanism used in this instance (and all other instances within ANSA Patch Management) is XML.

To query the job queues on the IS to see if any available external updates are ready to be installed.

When the IS directs the IC to apply an update or patch, the IC communicates back to the IS the status of the update (pass/fail).

In the event of power failure or other network disturbance, the IC(s) will be able to restart an installation process where necessary. If/when an IC is unavailable (for instance, a laptop is removed for the evening or a desktop computer is powered down), the IS will

restart the patch installation where it left off when the IC again becomes available. This feature differentiates ANSA from scanner-type solutions; the benefit being that ANSA will ultimately install any required patches to a known client even though the client was unavailable to the network at the time the patch was initially distributed. All it takes is for that IC to be successfully connected to the network; from that point on, ANSA Patch Management restarts itself and automatically installs any available patches. Non-agent based solutions (scanners) will not do this.

Conclusion

Manual patch management is very time consuming and prone to errors, while In-House/Legacy Solutions are somewhat more efficient, they still require time effort and full-cycle programming from an internal IT group. Using ANSA™ Patch Management removes all the headaches from patch management and allows employees to concentrate on other job activities. ANSA™ Patch Management greatly reduces costs involved with patch management by reducing the amount of manpower required to manage and deploy patches. Patch compliance has become critical to the success of organizations.

ANSA™ Patch Management works across multiple platforms (Microsoft Windows 9.x, NT, 2000, XP, Linux, Solaris, Mac OS X), this is extremely important for any company that runs or is planning to run multiple platform applications. It allows the Administrator the ability to manage the complete enterprise through one tool.

ANSA™ Patch Management provides a comprehensive web based user interface (single point) for Administrators. The interface allows the ANSA Administrator to control exactly which updates will be allowed on their defined clients, and establish policies for patch management at three levels: server, groups of computers within servers, and discrete client levels. These policies enable control over how and when patches will be deployed.

ANSA provides you with a complete record of your network devices and computer hardware and software, in XML format, importable into SQL data structures. This enables you to make decisions based on fact, not assumptions, using reports that are customized for your needs.

ANSA™ Patch Management will reduce costs, save time, improve security and reduce patch management errors.