

# Security Watch

A QUARTERLY PUBLICATION

## CONTENTS

Gaining Control of Security Patch Management . . . 1

Application-layer Filtering: Moving Security up the Network Stack . . . . . 5

Introducing Microsoft Identity Integration Server 2003, Enterprise Edition . . . . . 10

Microsoft Rights Management Services Addresses Information Security Challenges . . . . . 16

Microsoft Partners . . . . . 24

## CASE STUDIES ADVERTISER-SPONSORED

Aspelle . . . . . 20

DigitalPersona . . . . . 21

N2H2 . . . . . 22

Websense . . . . . 23

## SECURITY WATCH AUGUST 2003

This special advertising section was produced by the *Windows & .NET Magazine* Custom Media Group in conjunction with Microsoft. This supplement appears as an insert in the August 2003 issue of *Windows & .NET Magazine*. For information about Custom Media opportunities, contact Don Knox (dknox@winnetmag.com).

# Gaining Control of Security Patch Management

by Randy Franklin Smith

An annual computer crime and security survey conducted by the Computer Security Institute and the U.S. Federal Bureau of Investigation tallied more than \$201 million in quantified financial losses in 2003. Ironically, the vast majority of successful attacks by individual hackers and by fast-spreading worms rely on the presence of known bugs in the target computers—bugs for which patches are already available (Table 1). Deploying patches to thousands of computers manually is not an option. In the past, however, the processes, tools, and techniques for deploying and managing patches have not always been readily available.

Software vendors and customers alike must meet the challenge of security patch management to counter the escalating cost of security breaches and to enable companies to further leverage the Internet for integrating business processes with partners and customers. Cooperation among software vendors and customers can greatly reduce the quantity of successful attacks and the staggering costs associated with them. Microsoft, for example, has made keeping your environment secure and reliable a priority. As part of Microsoft's Trustworthy Computing initiative the company leads the way in meeting the challenge of keeping your environment up-to-date and secure while addressing the equally important issues of cost and stability.

Microsoft has produced a suite of free automation tools for identifying unpatched computers and rolling out updates. These tools are described in the *Microsoft Guide to Security Patch Management*, a comprehensive guide to implementing a security patch management process for your

TABLE 1: Historical Attack Examples and Patch Availability

MSRC Bulletin	Attack Name	Attack Discovered	MSRC Date	Days patch available prior to attack
MS03-007	Trojan.Kaht	5-May-03	17-Mar-03	49
MS02-039	SQL Slammer	24-Jan-03	24-Jul-02	184
MS01-020	Klez-E	17-Jan-02	29-Mar-01	294
MS00-078	Nimda	18-Sept-01	17-Oct-00	336
MS01-033	Code Red	16-Jul-01	18-Jun-01	28

entire Windows environment. In this article, I'll introduce you to the guide and the automation tools. In addition, I'll highlight best practices that will help you make the most of these resources and thereby reduce risk and costs of unpatched systems.

### THE MICROSOFT GUIDE TO SECURITY PATCH MANAGEMENT

I can't overstate the need for a well-planned, strategic security patch management process. Today, many companies use a bottom-up approach for development and support, in which the IT department develops a security plan without executive sponsorship. To address business needs and to be successful, upper management must appreciate the importance of security patch management—and ownership of the process must be clearly defined. If you don't have upper management support or clear ownership of the process, political obstacles will sometimes defeat champions of patch management, coverage will be spotty, and the overall patch management process will fall out of alignment with your organization's larger business objectives.

The *Microsoft Guide to Security Patch Management*, available at <http://go.microsoft.com/fwlink/?LinkId=16284>, gives management and technical staff a road map for implementing a successful security patch management process. The 150-page guide leads you through planning your security patch management process, starting with non-technical management-level decisions. The guide takes you through each step, from setup activities through change initiation, change

management, and release management—and explains how to optimize the ongoing process. The guide also offers a detailed examination of where and how Microsoft products and technologies such as Software Update Services (SUS), System Management Server (SMS), and Microsoft Baseline Security Analyzer (MBSA) fit into your security patch management process and how they can automate the deployment of Microsoft product patches. SMS, although not required to implement patch management, provides added value for some users.

The guide also includes newly written scripts that help you further leverage and automate tools like MBSA. And you'll find effective methods for maintaining stability by testing and recognizing problem patches before your rollout to the whole enterprise and how to recover when patches do cause problems.

### Preparing for Patch Management

Implementing an effective security patch management process requires preparation. The guide shows you how to combine the necessary people, processes, and technology to achieve success. First, conduct an assessment of your environment to determine the assets that need patch management and the state of protection measures already in place. Assets are the computers and software products installed on those computers. Protection measures such as firewalls, content filtering, anti-virus, disabling unneeded features, and written policies often partially mitigate the risk of many exploits and may thus reduce the urgency for patching certain assets. After your initial assessment of assets and protection measures you begin

the core of the project—developing your implementation plan.

### The assessment phase


The major objective of the assessment phase is to determine what people, processes, and technologies need to be put in place for proactive security patch management to be successful. During the assessment phase you inventory all the OSs, versions, and patch status of the assets that need to be addressed. You identify all your company's standard builds and the processes currently in place to update those builds and existing systems. Microsoft's patch management guide provides detailed help for performing this assessment with time saving tools like MBSA. You'll also find guidance for integrating your security patch management process with other elements of your organization such as security administrators, written security policies, and build teams.

### Forming a team

Developing your security management process includes forming a patch management team that will own the process and ensure that each step of the process, from monitoring through testing and deployment, remains capable of meeting your company's ever-changing business and technology requirements. At the beginning of the project, as well as whenever your infrastructure changes, your team must baseline your environment. Baselining is part of the security patch management life cycle.

### The Security Patch Management Life Cycle

The end-product of your implementation plan defines not only your patch management process but also a security risk contingency plan for dealing with issues such as how to rollout patches during an attack and how to recover after rolling out defective patches. You must identify what current procedures must be changed, what new procedures must be



initiated, who will perform the procedures, and how the procedures will be performed. For instance, who will be responsible for identifying new patches relevant to your various assets and then respond? While Microsoft provides one source for notifying customers of security patches for any Microsoft product, you must also keep up with security patches for non-Microsoft products, including management software and device drivers from hardware vendors. Finally, you enter the operation phase when your security patch management process goes live. Once the ongoing process becomes active you can optimize it as necessary.

### **Baselining your environment**

Baselining includes not just maintaining an inventory of computers, software, and version numbers but also making sure that all software is maintained at a level supported by the vendor. For instance, Microsoft supports products maintained at the latest and previous service pack levels. Unsupported assets pose a security risk because patches won't be available as new exploits are discovered. The guide shows how you can use Microsoft tools like MBSA and existing third-party tools to speed up and automate the base-lining process.

A useful technique for baselining involves identifying different asset categories, subcategories, and common baselines that you can use during the security patch management process. For instance, you probably have end-user PCs, data center servers, and Internet servers. Later in the planning phase, you can use these asset categories when prioritizing resources to address the most important assets first during patch deployment. Different assets have different values. For example, a server is more valuable than a workstation not only in terms of replacement cost, but also in terms of its role in your business. You naturally tend to devote more resources to securing more valuable assets, but some exploits can affect high-value

assets through assets that are lower in value yet share the same network or other resources. For example, Nimda used infected workstations to stage attacks against servers on a network and cause network traffic increases so that the network itself experienced denial of service problems.

While the value of an asset may help determine the urgency with which it should be patched, your team's overall goal should be to keep all assets secure. Additionally, stability becomes a critical issue once you put a patch management infrastructure in place because it becomes very easy to negatively affect thousands of computers with one mistake. Moreover, a patch management infrastructure itself presents an attractive target to attackers. For example, if someone compromises your infrastructure he or she could roll out a Trojan or denial of service agent to computers throughout your network. All these factors make the integrity of your change management process extremely important.

### **Managing patches while addressing risks**

After your team baselines your environment you can use the Microsoft guide to set up a "change initiation" process that manages patches while also addressing risks. An iteration of the change initiation process can be triggered by several different activities carried on by the patch management team. Regular vulnerability scanning reports, virus or intrusion detection reports, alerts from security Web sites, or bulletins from Microsoft or other software vendors, all monitored by the team, may identify a new vulnerability that must be investigated. First, you must analyze the exploit's relevance to your environment. Which asset category does the exploit affect, what mitigating controls are already present, and how severe is the threat? For example, a SQL Server vulnerability puts only certain systems at risk. Or, if the vulnerability is specific to an unused feature of IIS that is disabled

in your standard Web server builds, you might classify the risk as low and wait for the next service pack to address the problem.

After identifying a patch that needs to be deployed you begin the release management process. The guide steps you through planning, testing, deployment, rollout monitoring, and rollback. For rollout, the guide provides detailed techniques for using SMS (if it is part of your environment), SUS, WU and other technologies such as scripting to deploy patches to thousands of computers automatically. The guide describes several alternative ways to accomplish testing before rollout, such as through pilot testing with a limited number of production assets or with a lab built to resemble your production environment. The guide also defines best practices such as rolling out the patch in phases across your network to limit exposure if problems are discovered and to save monitoring and support resources. The guide shows how to accomplish such a staged rollout using Microsoft technologies like SUS and SMS, and provides insight into how to monitor deployment using Windows logs, and by watching network utilization, scanners and service desk calls.

### **TOOLS AND TECHNOLOGIES FOR PATCH MANAGEMENT**

Understanding the need to patch systems isn't new for most administrators, but having the means to do something about it is new. When faced with hundreds or thousands of servers and workstations where do you begin? Microsoft provides a number of tools and technologies to help you automate patch management. Microsoft has adapted the baseline and update services originally offered on its Web site into tools that corporate IT departments could use internally. For help with baselining you can use Microsoft Baseline Security Analyzer (MBSA) and the Office Update Inventory Tool, both of which are free

at <http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/tools/Tools/mbsahome.asp> and <http://www.microsoft.com/office/ork/xp/journ/offutool.htm>. You can use MBSA to scan all the computers on your network for missing OS security updates, and you can use the Office Update Inventory Tool to find missing security updates for Microsoft Office. MBSA 1.1.1 is available today and includes Windows Server 2003 support. Version 1.2 will include the Office Update Inventory Tool to make scanning your entire Microsoft environment easier. You can use Software Update Services (free at <http://www.microsoft.com/windows2000/windowsupdate/sus>) to deploy OS security updates, and you can use your existing Active Directory/Group Policy infrastructure to deploy service packs and security updates to applications like Office. For maximum manageability and scalability within very large networks the guide shows you how to use Systems Management Server with the SUS Feature Pack to baseline and deploy OS and application security updates.

### Microsoft Baseline Security Analyzer

When you use MBSA and initiate a scan you can specify a single computer or a set of computers, using either a domain name or IP subnet, which lets you scope your scan to logical domain boundaries or to physical areas using subnets. MBSA uses your administrator authority to connect to each computer and scans for missing security updates to the OS (NT 4.0, Windows 2000, Windows XP, or Windows Server 2003). MBSA 1.1.1 also scans IIS 4.0, 5.0 and 6.0; SQL Server 7.0, 2000, and Microsoft Data Engine (the desktop version of SQL Server); Internet Explorer 5.01 and later; Exchange 5.5 and 2000; and Windows Media Player 6.4 and later. MBSA provides a Web page report for each computer. Before scanning the computers MBSA downloads a fresh database of updates from Microsoft.

MBSA also checks for common security misconfigurations in the above products, except for Windows Media Player. Shavlik Technologies, which developed MBSA for Microsoft, also offers HFNetChkLt (free) and HFNetChkPro, which have advanced capabilities including the ability to install missing updates and scan Microsoft Office and Commerce Server.

### Office Update Inventory Tool

The Office Update Inventory Tool is comprised of two tools. First, you run the inventory program against each client computer to produce a log. The inventory tool can scan Windows 98, ME, NT 4.0, 2000, XP, and 2003 computers to find missing security updates for Office 2000 or Office XP. Then you use the conversion program, which processes all the client logs into a comprehensive report.

### Software Update Services

Software Update Services (SUS) comprises server and client components. The server component requires a Windows 2000 or later server, which maintains an up-to-date catalog of all available updates to Windows 2000 and later OSs. You administer your SUS server via Web pages that let you select which updates you want SUS to install on your computers. Automatic Updates, the client component, regularly checks your SUS server for new updates approved for installation and then installs them. You can control how often and when AU installs updates and how it handles reboots. The AU client comes with Windows 2000 SP2 and later. For more information about SUS, see Software Update Services, Part 1 and 2 at <https://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/tools/Tools/sadSUS1.asp>. Microsoft's patch management guide provides information about how to integrate SUS into your patch management process, including how to use multiple SUS servers to accomplish testing and staged deployments.

### Systems Management Server

For larger enterprises SMS 2.0 and the Software Update Services Feature Pack make it easy to manage security updates. SMS incorporates and extends the capabilities of MBSA and the Office Update Inventory Tool to give you enterprise-wide Web-based reporting of missing Windows and Office security updates and the ability to install those updates with full support for the best practices for patch management described in the guide. One treasure from the feature pack is the Distribute Software Updates Wizard, which provides a turnkey solution for getting your enterprise up to date without having to be an SMS expert. The wizard uses SMS's inventory and software distribution features to automate the baseline and deployment process. You can use the wizard to scan your organization for missing Windows and Office security updates. Then the wizard suggests a list of updates, which you can review and authorize on an update-by-update and computer-by-computer basis. Next the wizard downloads the updates you authorized and builds the required SMS packages and advertisements, which results in automatic deployment to your computers.

Microsoft is developing enhancements to MBSA, SUS, and SMS that will go even further toward helping you automate the patch management process. MBSA 1.2 will include support for scanning Office 2000 and later for missing security updates. SUS 2.0 will let you install more than just OS updates. With SUS 2.0 you'll be able to install service packs, device drivers, Office updates, and updates to SQL Server and MSDE, Exchange Server, and other applications including Visio.

Patch management is a real problem and Microsoft provides real help with the *Microsoft Guide to Security Patch Management* and with free tools like MBSA and SUS. With these resources you can gain control of the patch management problem and protect your organization from the risk of unpatched systems. ◆

# Application-layer Filtering: Moving Security up the Network Stack

by **Randy Franklin Smith**

**Y**ou probably remember the first time you encountered “firewalls” and the rationale for implementing one before connecting your LAN to the Internet. Today, the importance of traditional firewalls that enforce policy at the transport and packet levels is indisputable. Traditional firewall developers have reduced the problem to more of a science than a black art. These devices do a good job of solving packet-level security problems with packet filtering and stateful inspection. In fact, traditional firewalls are so effective that they’ve pushed attackers further up the stack to focus on application-layer attacks that occur on a very different plane than where traditional firewalls operate.

To decide whether to allow a packet onto your intranet, traditional firewalls typically just examined the packet’s header and made the decision based on destination and source addresses and, more to the point, protocol type (UDP/TCP) and port numbers. But you can no longer rely on what a packet’s header says it contains. For instance, Back Orifice, the Trojan/backdoor hacker tool, originally communicated between agent and controller over TCP port 31337, making it easy to detect, alert, and block Back Orifice traffic. However, more recent versions of Back Orifice can use any port, including port 80, which traditional firewalls will simply assume to be http—fairly harmless Web browsing traffic.

Attackers are not the only ones creating the need for application-layer filtering; the direction of software development and the rise of Web services are overloading port 80 and creating a general need for higher firewall intelligence that can make decisions based on data inside packets instead of just packet header information. For example, rather than simply blocking or allowing RPC traffic, you might need the ability to allow RPC for certain applications like Exchange and still block any other type of RPC traffic. That decision requires an intimate understanding of how RPC works and how to decode RPC, which resides in a TCP packet’s data section.


Web services represent another scenario rapidly becoming a concern. Web services use Simple Object

Access Protocol (SOAP) and XML documents. SOAP traffic can contain high-level messages that make requests against security-critical business servers. You might already have Web-service-enabled servers and not know it. Take this case in point: SAP’s CRM software supports SOAP out of the box. SOAP messages and related XML documents ride inside http packets. Traditional firewalls only look so deep into a given packet—basically at the packet’s header, which specifies source and destination address and then the source and destination port numbers. However the bytes that make up SOAP attacks don’t reside in the packet’s header. SOAP attacks such as invalid schema, attempts to access unauthorized objects or messages, or digital signature and encryption attacks reside in what traditional firewalls consider the “payload,” or data part of a packet and thus don’t get caught.

## ISA SERVER ARCHITECTURE

For large enterprises, Microsoft Internet Security and Acceleration (ISA) Server 2000 is a complementary technology to existing traditional firewalls. You can integrate ISA Server into your existing perimeter security infrastructure to address your application-layer security needs without the massive “rip and repair” costs associated with changing firewalls. For small to medium organizations, ISA Server can fulfill the total network and application gateway requirement. ISA Server provides the features you’d expect in traditional firewalls such as packet filtering, stateful inspection and network address translation (NAT) as well as Web content caching. But the product’s focus is on application-layer filtering for protecting internal servers that are available on the Internet and for controlling internal user access to the Internet. In fact Microsoft’s prescriptive guides for enterprise and Internet data centers employ partner products as traditional firewalls and utilizes ISA Server just for its application-layer filtering capability.

Microsoft built ISA Server to provide secure access to certain “published” resources on the internal network while blocking access to everything else. Microsoft’s



philosophy for publishing a resource (a Web server, for example) securely is to put an application gateway between the Internet and the server. The gateway, serving as a reverse-proxy, poses as the published server, accepts the request and inspects it, and then makes an equivalent request to the protected server. Packets from the outside never touch the protected server.

The most common scenario for implementing Microsoft's method is where ISA Server functions as protection for a Web server. ISA Server checks each request at the application-layer before allowing the traffic to continue to the protected Web server. If the published Web server uses SSL, encryption (and optionally authentication) terminates at the ISA Server. This approach is important because ISA Server relieves the Web server from being the target of attacks. ISA Server can fully apply security checks to the request and validate the legitimacy of the traffic—a significant improvement over allowing uninspected anonymous traffic to arrive directly at the Web server where a malicious request might succeed in harming or exploiting the data that the Web server exposes. When the application uses an authentication type supported by ISA Server, the client on the Internet can authenticate to the ISA Server before a request ever hits the protected server. This scenario lets you place the actual server out of the DMZ and back into the warm and cozy environs of the internal network where it can more easily interface with other servers to complete transactions. Now ISA Server becomes the bastion host—the “sacrificial lamb” in the DMZ. If the ISA Server is compromised by an application-layer attack—which, by the way, hasn't yet occurred in the more than two years this product has been on the market—the protected server is still untouched.

ISA Server's architecture is extensible, allowing Microsoft customers

and ISVs to build Web and application filters that plug into ISA Server and endow it with the ability to perform application-layer inspection for any type of application. Web filters provide http-specific filtering, and application filters handle non-http application protocols.

## APPLICATION FILTERS INCLUDED

Besides filters to inspect Web traffic (discussed later), ISA Server comes with several other application filters already installed. The HTTP redirection filter catches Web requests from internal users and redirects them to the Web proxy component of ISA Server. The DNS intrusion detection filter intercepts incoming DNS traffic and analyzes the traffic for buffer overflows and illegal zone transfers to prevent poisoning or reconnoitering of your internal DNS servers.

The FTP filter enables FTP connections across the ISA Server for both FTP clients and FTP servers on the internal network. For FTP clients on the internal network, the filter handles the problems associated with the arcane way FTP uses 2 port connections, one initiated by the client and the second and problematic connection initiated from the server back to the client. The filter also lets you publish FTP servers on the internal network to the Internet. ISA Server presents itself as the published FTP server to clients on the Internet and then forwards requests and data to the actual FTP server on the internal network.

The H.323 filter provides ISA Server with support for IP telephony and control for incoming and outgoing calls, audio, video, and application sharing. The POP intrusion detection filter scours POP3 traffic for buffer overflows before Internet users can access your email server.

## PUBLISHING EXCHANGE RPC

ISA Server includes two RPC publishing filters: a basic filter for publishing

any RPC service and a specific filter for Exchange RPC. The RPC filter lets you publish RPC servers to the Internet without opening any more ports than needed for each individual connection. While most application protocols require just a few well-known ports, RPC is different. RPC applications do not use pre-established port numbers for communication. Instead, each application has a Universally Unique Identifier (UUID).

When any RPC server application starts up, it opens a random TCP port in the range of 1024-65535. Thus an RPC server opens a different port for each RPC service running on it. After the port opens, the RPC service registers its port number and UUID with the RPC port mapper. When an RPC client wants to connect to a corresponding server application, the client first contacts the RPC port mapper on well-known port 135/tcp and asks for the appropriate port number for the UUID in question—Exchange, in this example. The server responds on port 135, and then closes the connection. Finally, the client connects to the server on the newly learned high port. For example, an RPC server application called WidgetServer starts up and opens port 19292. Then WidgetServer registers its UUID and port 19292 with the RPC port mapper. Next a client elsewhere on the network connects to the server's RPC port mapper on port 135 and queries the port mapper with WidgetServer's UUID. The port mapper responds with port 19292. The client disconnects and then connects directly to WidgetServer over port 19292.

RPC works well enough on a trusted internal network. But when you want clients on the Internet to be able to access an RPC server behind a normal traditional firewall, you must open access to port 135 and ports 1024-65535, which creates a gaping hole in your defenses. ISA Server's RPC filter serves as a middleman between Internet client and internal RPC server and succeeds in just

opening the necessary ports for each connection.

Let's use the WidgetServer example again but this time with an ISA server between the client and server. The client connects to the ISA Server on port 135 which ISA Server directs to the RPC filter. The filter looks at the UUID and determines if a valid publishing rule exists for it. If so, the filter then connects to the specified server using the typical RPC client/port mapper/server interaction described above. At this point, the filter on the ISA Server is a client of the published RPC service on a server in the internal network. Next the filter opens a high numbered port on ISA Server for access by the associated client only and sends that port number back to the client. The client then connects to that port and initiates authentication. The RPC filter, acting as a proxy, relays messages back and forth between the client and server.

The Exchange RPC filter is designed specifically for publishing Exchange RPC over the Internet. By publishing Exchange Servers with ISA Server's RPC filter, mobile employees can seamlessly access Exchange with Outlook from anywhere on the Internet without first having to establish a VPN connection to the internal network. The RPC filter contains special functionality for publishing and protecting Exchange Servers. When you create an Exchange Server publishing rule the RPC filter specifically looks for Exchange Server's UUID when clients attempt to connect. The filter immediately drops connections specifying any other UUID thus protecting other RPC services. After getting a valid connection request to Exchange, the filter forwards the request to the protected Exchange server. The filter continues to monitor the interaction between the client and server as authentication proceeds. If the filter does not see successful authentication take place, it immediately terminates the connection, protecting the Exchange server from mal-

formed request attacks such as buffer overflows, portmapper attacks, and RPCDump reconnaissance.

### PROTECTING SMTP

ISA Server's SMTP filter is rich with features for controlling incoming SMTP email. You can limit attachments by size, file name, or extension as first line of defense against unauthorized software or malware from making it into your network via email. For instance you could block vbs files attached to email messages at the firewall. You can block email by sender or domain name. You also can implement protection against sensitive information making it outside your network by filtering for keywords in emails and attachments in outgoing emails.

The SMTP filter also lets you selectively enable or disable SMTP commands and limit their length. For example, you can stop SMTP reconnaissance attacks cold in their tracks by disabling VRFY\* which is a common way of getting SMTP servers to dump a list of all user accounts that the attacker can then use for a variety of nefarious purposes. By limiting command length, the SMTP filter prevents buffer overflow attacks from ever reaching the SMTP server.

### APPLICATION-LAYER FILTERING FOR HTTP AND HTTPS

Traditionally, Web servers reside in the DMZ. The forward firewall filters network-level attacks but the Web server is on its own to defend against application-layer attacks. Also, the rear firewall must allow the Web server to access second-tier servers on the internal network such as database servers. With application-layer filtering, ISA Server takes the place of the Web server in the DMZ and the Web server moves into the safer confines of the internal network. ISA Server protects published Web servers by recognizing application-layer attacks before they ever reach the destination

Web server.

When you publish a Web server you specify one of ISA Server's Internet IP addresses and optionally a path name as well as an internal Web server and optional path name. From a client's perspective on the Internet, the Web server is the ISA Server at the specified IP address and path. But again, acting as an application gateway, ISA server interprets the request, reformulates the request with the protected internal Web server in mind, and requests the Web page from the internal server on the user's behalf. When using ISA Server with an SSL protected Web site, the location of the Web server certificate used to authenticate the Web server to the client moves from the protected Web server to the ISA Server. Encryption terminates at the ISA Server giving it an opportunity to inspect the traffic before repackaging it and sending it on to the protected server in either cleartext HTTP or re-encrypting it to HTTPS.

You can configure ISA Server to inspect HTTP requests for a wide variety of suspicious patterns such as abnormally long URLs that might signal a denial of service attack or encoded characters that might indicate an embedded command the attacker hopes to trick the Web server into executing. ISA Server accomplishes these URL inspections using a specially adapted version of the URLScan ISAPI filter originally designed for Internet Information Services (IIS). ISA Server's URLScan filter comes with ISA Server Feature Pack 1 and checks HTTP requests before they reach the Web server. You can block URLs that contain specific verbs like PUT, POST and REPLY which attackers use to modify or deface Web site content. You can disable certain file extensions from being requested to guard against attackers trying to get the Web server to run executables and scripts. You also can disable certain character sequences that IIS or other servers

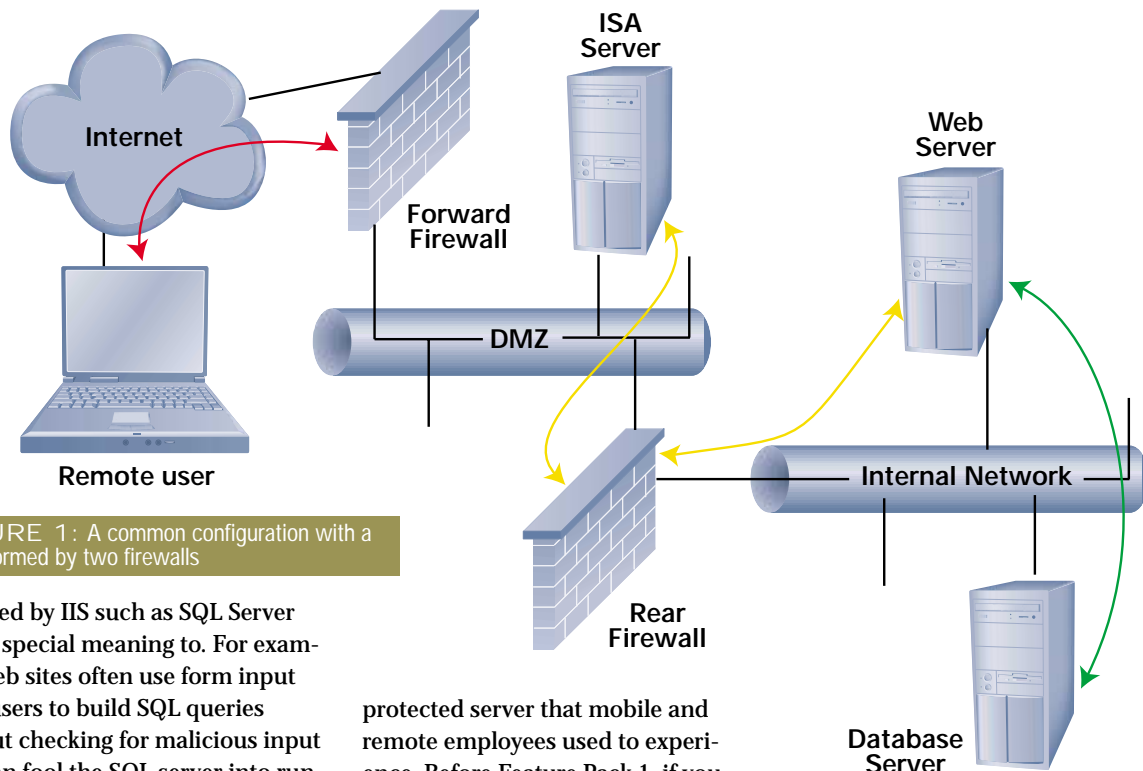


FIGURE 1: A common configuration with a DMZ formed by two firewalls

accessed by IIS such as SQL Server assign special meaning to. For example, Web sites often use form input from users to build SQL queries without checking for malicious input that can fool the SQL server into running arbitrary commands an attacker specifies. For a full discussion of URLScan's many features see <http://www.windowswebsolutions.com/Articles/Index.cfm?ArticleID=25581&pg=2>.

The link translator filter, also included with Feature Pack 1, solves the problem with publishing internal Web sites that contain hard-coded links to internal resources. On the fly, the link translator filter recognizes internal links that would typically be broken links when sent back to the client. The filter translates those broken links to links that ISA Server will be able to resolve back to the appropriate internal address when the client later requests them.

ISA Server makes sure a request is safe and valid before passing it on to the Web server. With yet another component of Feature Pack 1, you can take things a step further with delegated authentication, which enables ISA Server to even pre-authenticate the client before allowing it to access the Web server. Delegated authentication bridges the former gap between authentication to ISA Server and authentication to

protected server that mobile and remote employees used to experience. Before Feature Pack 1, if you published an authenticated Web site such as Outlook Web Access (OWA), the remote employee first had to authenticate to ISA Server and then again to the IIS server running OWA. Now with delegated authentication, the remote employee simply authenticates to ISA Server. The ISA Server then connects to the Web server with the user's authentication credentials, providing seamless single-logon authentication to HTTP and HTTPS-based applications, including OWA, SharePoint Portal Server or custom-written Web applications.

If you deploy ISA Server into your existing perimeter network, you can implement application-layer filtering. Figure 1 shows a common configuration with a DMZ formed by two firewalls. The forward firewall is connected to the Internet and the DMZ network. The rear firewall is connected to the DMZ and internal networks. You can move your Web, email, and other servers currently in the DMZ to the more protected internal network and deploy ISA Server in place of them in the DMZ. Then you configure publishing rules on the ISA Server so that it inspects and reverse-

proxies requests from the Internet to the appropriate server on the internal network. Previously, you had to open ports for Web servers in the DMZ to communicate with second-tier servers such as databases on the internal network. Now Web servers can directly communicate with those second-tier servers, and the second-tier servers are never directly touched by servers in the more vulnerable DMZ.

The ISA Server must be able to authenticate remote employees against their Active Directory accounts on the internal network for you to implement delegated authentication. You could make the ISA Server a member of the internal domain. But this option is not wise because of the problems associated with separating member computers from the domain controller by a firewall.

A better way to facilitate delegated authentication requires that you create a dedicated DMZ AD forest. However you won't need to create any accounts for remote users. Instead, create a one-way trust relationship so that the DMZ forest trusts the internal



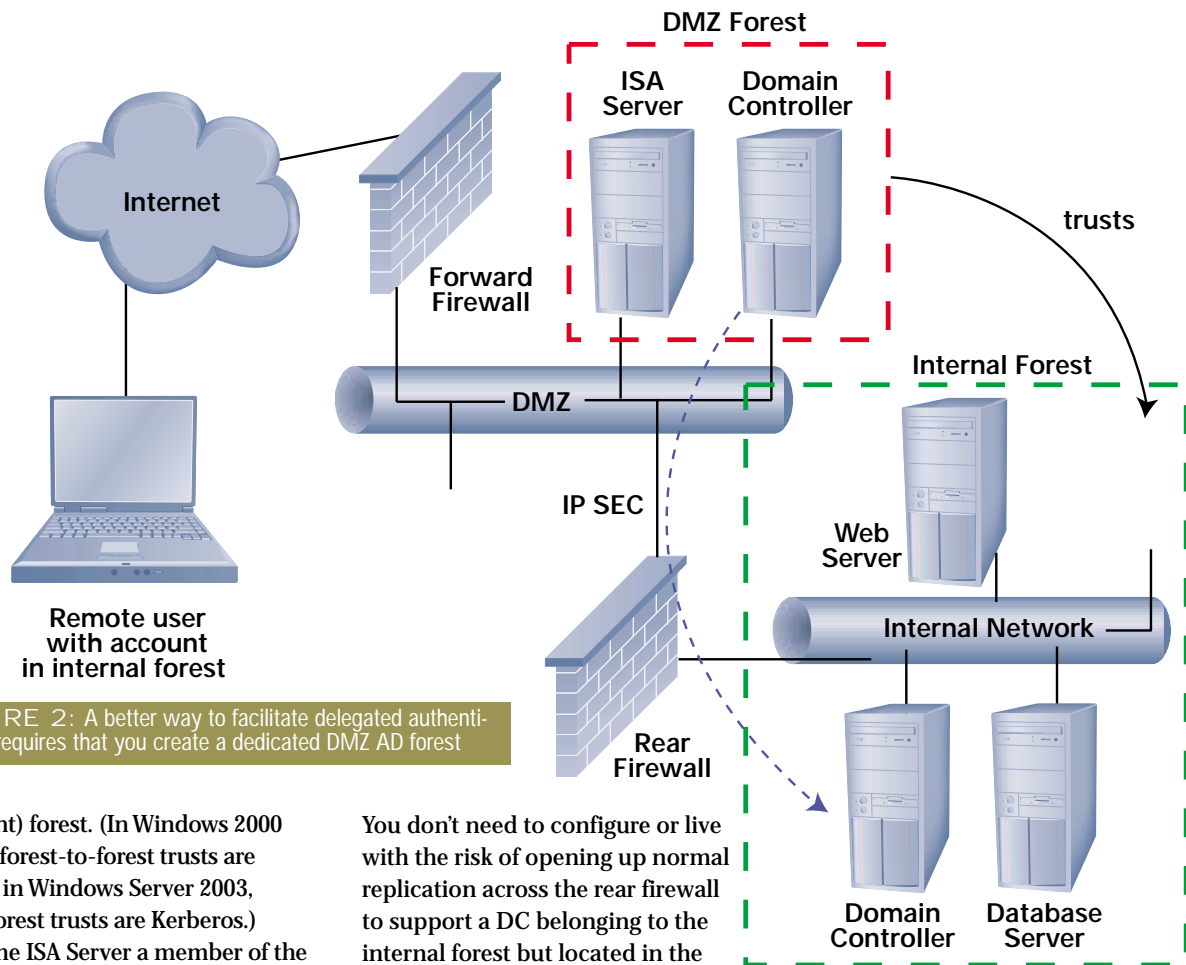


FIGURE 2: A better way to facilitate delegated authentication requires that you create a dedicated DMZ AD forest

(account) forest. (In Windows 2000 Server, forest-to-forest trusts are NTLM; in Windows Server 2003, cross-forest trusts are Kerberos.) Make the ISA Server a member of the DMZ forest. The trust relationship allows traveling employees to authenticate to the ISA Server with their normal AD account in the internal forest.

To facilitate this authentication, the DC(s) in the DMZ forest must be able to communicate with DCs in the internal forest across the rear firewall. The easiest and most secure way to allow this communication is to configure IP security policies on the DMZ and internal DCs that require IPsec for all inter-forest communication. Then, on the rear firewall, simply enable IPsec to flow between the DMZ and internal DCs. A minimal number of openings are required:

- IKE (port 500/udp)
- DNS (53/udp and 53/tcp)
- IPsec AH (IP protocol 51) or IPsec ESP (IP protocol 50)

With the configuration in Figure 2, you get the best of both worlds.

You don't need to configure or live with the risk of opening up normal replication across the rear firewall to support a DC belonging to the internal forest but located in the DMZ. The internal forest remains isolated from the Internet. But we still get to use accounts from the internal forest, so there's no need to create redundant accounts for remote employees. Note also that you could use the DMZ forest to host user accounts created for customers or business partners that need to access applications and Web sites on the DMZ.

Delegated authentication complements application-layer filtering. It provides a seamless user experience while enabling ISA Server to absorb attacks and only pass authentic, valid requests on to the published server.

### APPLICATION-LAYER FILTERING TAKES ON THE NEW BREED OF ATTACKS

Application-layer filtering is a powerful tool for combating attackers as they move up the food chain from

packet level to application level attacks. Application-layer filtering catches attacks before they ever reach the server.

ISA Server already includes a powerful array of application filters but more is on the way. By adding ISA Server Feature Pack 1, you can publish and protect Web, FTP, RPC, and Exchange servers as well as Outlook Web Access and SharePoint Portal Server sites. In future releases Microsoft is looking at building enhanced filtering to help you publish and protect all types of emerging applications and protocols. ♦

Randy Franklin Smith is a contributing editor for *Windows & .NET Magazine* and the primary instructor and course developer for MIS Training Institute's Windows NT/2000 security program. His firm, Monterey Technology Group, provides security consulting.

# Introducing Microsoft Identity Integration Server 2003 Enterprise Edition

In an increasingly complex environment, managing the lifecycle of digital identity—who users are, how they prove it, what they can access and how and when to retire that privilege—is primarily a process problem. Using technology to automate and streamline that process enables businesses to unlock the value of information stored in their IT systems and get that information into the hands of customers, employees, business partners, and contractors that need it most, when they need it, securely. The lack of comprehensive identity management and integration prevents the easy, secure flow of information among the disparate and scattered systems and applications at large companies today. A lack of identity integration and management creates security holes, holds up projects, and creates expensive extra work for human resources, management, and IT staff as they try to provision new accounts, handle access control changes, keep contact information up-to-date, and find and disable the many accounts a terminated employee gained during their tenure.

In the last issue of Security Watch, I introduced you to Microsoft's solution for identity integration, which features as keystone technologies Active Directory (AD) and Microsoft Identity Integration Server 2003 (MIIS). In this article, I take an in-depth look at MIIS and show you what it can do for your organization and how it works.

With a product like MIIS you can integrate any number of disparate identity repositories, regardless of format. Ultimately, MIIS becomes the root of identity information within your enterprise. MIIS detects changes to a user's identity information in one repository, such as your human resources application, and then synchronizes that change with all other identity repositories according to rules you specify. MIIS's integration functionality enables you to eliminate duplicated work associated with maintaining the same user's information in multiple places, reduce discrepancies that impede productivity and undermine security, and make accurate identity information widely available.

**MIIS IS INTEGRATED WITH WINDOWS**  
MIIS runs on the Windows Server 2003, Enterprise Edition platform for better scalability and dependability and enhanced capabilities. Microsoft has replaced MIIS's original data store with SQL Server for its performance, integrity, and clustered availability capabilities.

Microsoft has made MIIS available in two ways: Enterprise Edition and the Identity Integration Feature Pack for Microsoft Windows Server Active Directory. The feature pack is made available at no charge to customers who have already licensed Windows Server 2003, Enterprise Edition. The feature pack lets you solve directory synchronization problems that sometimes arise in larger, more partitioned Windows environments. For example, if your company is divided into multiple forests, identity integration across these forests becomes more difficult. If everyone resides in one forest, you can easily locate each other and find other directory published resources because an AD forest is searchable through one global catalog that encompasses all domains in the forest. But in multiple forest environments, users must be trained to search multiple global catalogs to find users or resources in another forest. Likewise, when you have multiple Exchange 2000 or Exchange 2003 environments, you no longer have one Global Address List, which impedes email communication, information sharing, and workflow activities implemented on Exchange. (For general information about MIIS, go to <http://www.microsoft.com/windowsserver2003/technologies/directory/miis/default.mspx>.)

## THE IDENTITY INTEGRATION FEATURE PACK

With the Identity Integration Feature Pack for Microsoft Windows Server Active Directory you can solve the problems related to multiple forests by integrating the identities of all your AD forests, Exchange 2000 or Exchange 2003 implementations, and ADAM installations. The feature pack synchronizes users and contacts between multiple AD



forests. When you only need to synchronize a portion of forest with another you can limit synchronization to organizational units or even specific users. You can also provision groups, user accounts, and contacts across forests. MIIS can also create an account in a company's email or other systems. Or you could set up synchronization rules in MIIS that automatically create appropriate application level accounts and access when you hire a new employee. For example, you could create a rule that specifies whenever you create a new user account in AD and assign it to the HRStaff group, MIIS also creates a user account in PeopleSoft's HR module.

MIIS Enterprise Edition supports a wide variety of other directory services and systems including products from Lotus, Netscape, Sun, IBM, SAP, PeopleSoft, Oracle and even telephone switches. MIIS supports Microsoft SQL Server and other relational databases such as Oracle so that you can synchronize identity information in internally developed and ISV applications. To round out MIIS' support for any imaginable source of identity information, MIIS can read and write to industry-standard formats such as LDAP Directory Interchange Format (LDIF), common text file formats, Directory Services Markup Language (DSML) 2.0, and XML-based files. For a full list, see Table 1. MIIS enterprise edition, as well as the feature pack, requires Windows Server 2003, Enterprise Edition. The feature pack requires Microsoft SQL 2000 Standard or Enterprise as its data store and MIIS Enterprise Edition requires Microsoft SQL 2000 Enterprise.

## THE MIIS ARCHITECTURE

An MIIS implementation typically comprises several components: connected directories (aka, connected data sources), management agents, connector spaces, and the metaverse—the core of MIIS (Figure

1). Connected directories are the directories, databases, or other identity repositories between which MIIS performs synchronization or provisioning. For a typical organization, connected directories would include obvious sources such as AD, Exchange, NT domains or Lotus Notes—as well as applications such as PeopleSoft or SAP. But you might be surprised by other technologies that could benefit. For example, with MIIS you could integrate your PBX and the access control system for your buildings and gates into your identity management infrastructure.

Management agents move data between connected directories and the metaverse. The metaverse is the integrated view of all the identity information about a given person that has been joined from all the connected directories into a single, authoritative entry in the metaverse. From the metaverse other management agents notice the change and then apply it to their corresponding connector spaces and then to their connected systems.

## Management agents

Management agents are logic modules within MIIS that link a given connected identity repository to the metaverse formed by MIIS and its connections to all other directories. Each type of connected system has a corresponding management agent that understands the structure and protocols inherent to the connected repository so that the agent can import and export data from the repository. For example, MIIS has management agents for AD and Lotus Notes and for generic file formats like comma separated value files. The management agent watches the connected system for changes to identity data and replicates those changes into MIIS.

When a change to identity data occurs in some other system connected to MIIS the management agent analyzes the change and, if appropriate, the agent replicates that change into its connected directory. For example, let's say that someone in HR changes Bob's job title from Purchasing Agent to Production

TABLE 1: Formats supported by MIIS

- Active Directory – supporting Windows 2000, Windows Server 2003, Exchange 2000 and Exchange 2003
- Active Directory Application Mode (ADAM)
- Global Address List (GAL) Synch – supporting Exchange 2000 and Exchange 2003
- Netscape/iPlanet/SunONE Directory – supporting version 4.x/5.x (includes "changelog" support)
- SQL Server – supporting SQL 7 and SQL 2000
- Oracle Databases – supporting version 8i and 9i
- Directory Services Markup Language (DSML) – supporting DSML version 2.0
- LDAP Interchange Format (LDIF)
- De-Limited Text
- Fixed-Width Text
- Attribute-Value Pair Text
- NT4 Domains
- Exchange 5.5
- Exchange 5.5 Bridgehead
- Lotus Notes – supporting versions 4.6 and 5.0
- Novell eDirectory – supporting versions 8.6.2 and 8.7
- IBM Informix, DB2, dBase, Microsoft Access & Excel and OLE DB – via SQL Data Transformation Services (DTS)

niap.nist.gov/cc-scheme

October

# SECURITY VALIDATION REPORT



National Information Assurance Partnership

**MICROSOFT WINDOWS 2000**  
Awarded Common Criteria Certification  
Evaluation Assurance Level EAL: 4+  
Report Number CCEVS-VR-02-0025

→ EA  
Re  
Com  
An  
150  
11  
150  
2



Common Criteria Testing Laboratory: Science Applications International Corporation  
Columbia, Maryland

<http://www.microsoft.com/servers/>

ber 25, 2002

PLAY IT SAFE

MSDN Online

Deepfire Content Part

**Trust the facts first.** Microsoft® Windows® 2000\* operating system has been awarded the highest level Common Criteria Certification for the broadest set of real world scenarios yet achieved by any operating system, as defined by the Common Criteria for Information Technology Security Evaluation (CCITSE). The Common Criteria (CC) is an internationally endorsed, independently tested, stringent set of standards that evaluate the security of technology products.

Windows 2000 desktop systems and server family were evaluated against CC requirements beyond the Controlled Access Protection Profile to include the following features:

- Sensitive Data Protection Device
- Enterprise Directory Service
- Virtual Private Network (VPN)
- Software Signature Creation Device
- Single Sign On
- Network Management
- Desktop Management

These results arm you with an unmatched level of confidence as you select secure products for your environment and mark an important milestone on the road to providing you the highest level of assurance in Windows and all Microsoft products.

For all the facts on Windows 2000 Common Criteria Certification and managing your security risk with Windows 2000, visit [microsoft.com/CCcertification](http://microsoft.com/CCcertification) **Software for the Agile Business.**

**Microsoft®**

Supply Management Analyst. The HR application's management agent would notice the change and replicate the change from the HR application into MIIS. Other management agents would analyze the job title change against their configuration rules and, if appropriate, apply the change to their respective systems. For instance, because AD includes a job title field, AD's management agent would execute the job title change on Bob's AD user account.

With Visual Studio you can customize the management agents that ship with MIIS by plugging in scripts that perform additional processing when certain events occur or conditions are met. MIIS lets you use whichever Visual Studio language you are comfortable with, including Visual Basic .NET, Visual C++ .NET, Visual C#, Visual Perl, Visual Python, Pascal, or COBOL. For example, you could configure MIIS's management agent for your HR application to recognize employee terminations and send emails to security guards to notifying them immediately that the former employee should not be admitted regardless of the person's badge.

### Connector spaces

Within MIIS, data such as Bob's job title change flows between connector spaces and the metaverse. Each connected directory has a staging area in MIIS, called a connector space that the management agent uses to move data in and out of the connected system. The connector space contains a replica of all the objects in the connected system. In the case of a connected AD, the connector space would have an entry for each of the user objects in AD. Entries in the connector space contain only a subset of the attributes of their corresponding objects in the connected directory. The subset of

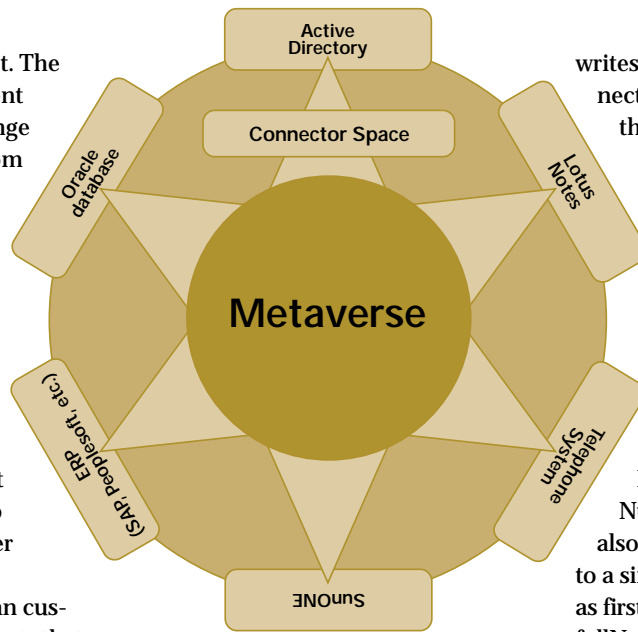


FIGURE 1: The Metaverse—the core of MIIS

attributes is defined by how you configure the corresponding management agent.

Entries in the connector space only need those attributes relevant to synchronizing the object with corresponding objects in other connected systems. The connector space stores state information about each attribute of each object, including the attributes current value in the connected system as well as the new value if changes from other directories have updated the attribute in the connector space. Having an intermediate connector space between the actual connected system and the metaverse helps MIIS solve classic problems with directory synchronization such as referential integrity, query performance and complexity and the occasional unavailability of connected systems. If a change to identity data occurs in one connected system and one of the other connected directories that should receive the change is temporarily unavailable due to network problems, this presents no problem to MIIS. MIIS simply writes the change to the connector space. When the associated directory is back online, its management agent

writes the change cached in the connector space to the actual object in the directory.

### Mapping attributes

Management agents understand how to map attributes from an object in one directory to the corresponding attributes on objects in other directories. For example, MIIS maps the Office-TelephoneNumber attribute in Notes to the LDAP telephone-Number attribute in AD. You can also map multiple source attributes to a single destination attribute such as firstName and lastName to fullName. To avoid creating mapping rules to handle the thousands of permutations that would result from trying to map every attribute in every directory to every other directory, MIIS uses the metaverse as the common junction for attribute flows. Therefore, it is only necessary to map attributes in each directory to the master attributes in the metaverse.

Changes to a given attribute such as phone number can originate in many different directories. For example, you can change a person's phone number attribute in HR, AD, Notes, the telephone system and, no doubt, in other directories. Without appropriate rules synchronization data management could quickly become chaotic with confusion and arguments arising over who changed what. To address this issue, MIIS implements the concept of data ownership and lineage. Ownership means that you can define a specific directory as being authoritative for certain attributes. For instance HR might be the owner of employee name and job title but the telephone system is the owner of the employee's telephone number. You can even specify multiple owners in order of precedence so that if the most preferable directory can't supply a value for a certain attribute, but another directory can, MIIS will use it.

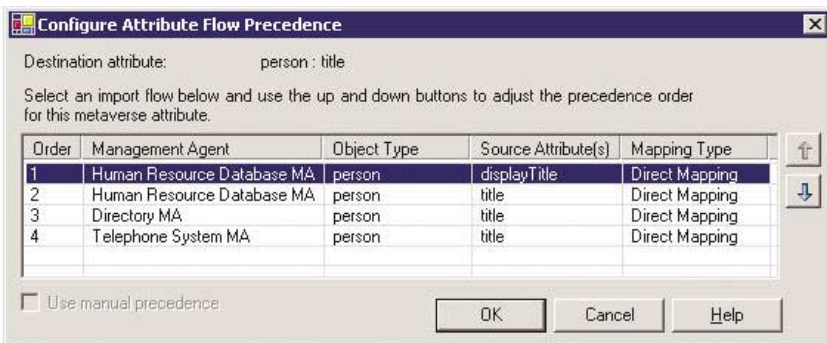


FIGURE 2: Example attribute flow precedence

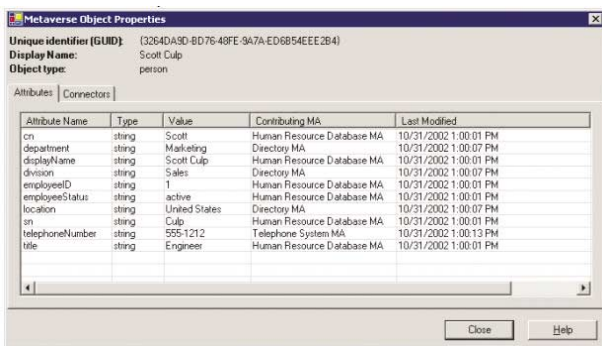


FIGURE 3: Metaverse Object Properties

Figure 2 shows the attribute flow precedence for the job title attribute. As you can see, MIIS is configured to prefer to HR's displayTitle or title attributes, but if those attributes are blank in HR for a given employee, MIIS will use the title attribute from the company's Active Directory or, if necessary, the telephone system. By implementing a concept called lineage, MIIS keeps track of where the data for each attribute comes from so that you can track synchronization through the directory. With lineage you can always explain why and how a certain attribute got to be the way it is. Figure 3 shows the metaverse entry for Scott Culp. The left side of the window lists all of Scott's attributes. The "Contributing MA" column shows which management agent that supplied each attributes value; the "Last Modified" column shows the last time each attribute was updated. You can see that while the HR system supplied attributes like job title and employee ID, the telephone system supplied Scott's phone number.

actions such as creation, deletion and moves. You can configure MIIS so that when a new employee object appears in HR, MIIS will automatically provision the employee with an AD account, Exchange mailbox, and objects in other directories as appropriate. Likewise, when you terminate an employee MIIS can deprovision the employee's accounts and other identity records throughout the enterprise. You can configure MIIS to completely delete the terminated user's objects or just make an attribute change such as disabling the employee's AD account. For exception accounts, such as service or application accounts in AD that you don't want replicated throughout the enterprise, you can specify disconnecter objects in the connector space. A disconnecter object is a placeholder in the connector space that blocks synchronization for the corresponding object in the connected directory. If you have certain directories that must be kept

synchronized with each other you can use MIIS's transactional capabilities to guarantee that a given change is successfully synchronized from the connector space to all connected directories, or to none at all. Because exceptions always occur in the real world, no matter how thorough your process is, you can configure MIIS to recognize exceptions and errors provide special handling or notification.

## TAMING IDENTITY MANAGEMENT

With MIIS you can tame the multi-headed hydra of identity integration in today's multi-identity repository environments. You can react more nimbly to business opportunities and problems, more cost effectively provision and maintain identity information, and maintain security by keeping directories and access consistent and accurate on a timely basis. When a new project, reorganization or acquisition occurs you can quickly provision new accounts and adjust access control. When users leave you can quickly close down that user's access throughout the entire enterprise.

With MIIS's architecture of management agents, connector space, and metaverse you can handle the massive complexity associated with enterprise-wide identity management and meet the other challenges of system availability, management, and audit trails—MIIS keeps running, despite errors and unscheduled system outages.

And while MIIS gives you the power to automate the drudgery of these processes, you don't give up control. With MIIS preview, ownership, attribute precedence, and lineage you can be confident about what MIIS will do to your systems before you make a configuration change—and you can always track changes back to their source, as well as the rules that governed the synchronization. ◆

# Microsoft Rights Management Services Addresses Information Security Challenges

**F**ile permissions are great for protecting information residing on a file server, but as soon as a user opens a file, the user must protect the information in that document and follow corporate policy regarding its use, modification, and distribution. Users can easily copy files to removable media, email them to unauthorized parties, or print them. Once copied, emailed, or printed, anything can happen to that information. Anyone can view it, change it, or share it with someone else. The same goes for email messages and internal Web sites. It's so easy to inadvertently send information to the wrong person, which can lead to embarrassment, legal issues, and regulatory problems. Even if you take advantage of the encryption feature of some applications to encrypt a file, you must still share the decryption information with users with whom you share the file. After they decrypt the file, they can do anything they want to with it.

A new information security technology from Microsoft, called Microsoft Windows Rights Management Services (RMS) for Windows Server 2003, addresses these information security challenges. Office 2003 will be the first Microsoft application to include RMS capabilities; however because RMS is a platform, third parties are expected to build RM-capable applications as well.

## RMS FUNDAMENTALS

Microsoft spent more than 18 months interviewing corporate customers regarding the basic problem of keeping information from falling into the wrong hands or being misused. Those interviews generated the requirements for RMS, which is designed to help enterprises exercise much more control over their information and have greater confidence that written policies are being followed.

RMS is the server side of Microsoft's Rights Management (RM) solution. RMS works with RM-capable applications such as Outlook 2003, Word 2003, Excel 2003, PowerPoint 2003, and Internet Explorer (IE) to enable users to share information in strictly controlled ways. For instance, with RMS and Word 2003 Professional Edition,

you can create a Word document and limit its usage to a specific user or group distribution list. Note that the document itself contains the usage policy, so no matter how or where the document is copied or distributed, only authorized users will be able to access it—and only in permitted ways. And RM can protect other information types besides just documents. For instance, you can protect Web pages with RM.

In addition to protecting documents, RM can protect any type of data; it is completely agnostic to application and data format. Furthermore, you can limit how authorized users can access a document and for how long. For instance, if you're a plastic parts manufacturer, you could send to Bob, a salesman for an injection-mold manufacturer, a request for bid (RFB) that contains sensitive design information about parts you're building for a new product for a customer. You could specify that only Bob can open the document and that Bob can only view the document. He can't print, save, modify, or even copy and paste from the document. About the only way Bob could compromise the information would be to memorize the information and reproduce it elsewhere. Furthermore, you could set a time limit on Bob's access so that after a week, he won't even be able to open the document for viewing purposes.

RM integrates into your existing Active Directory (AD) infrastructure so that users enjoy single sign-on based on their initial authentication to Windows. For users outside your enterprise, such as business partners, you can authorize RM-protected documents to Microsoft Passport accounts or use advanced features of RMS to collaborate with users in other organizations AD infrastructures. When the external user receives your document, he or she can open it in the native RM-capable application or the user can open it in Internet Explorer, after installing the Rights Management Add-on for IE, available from Microsoft for download at no charge.

This protection works through licensing, certificate, and authorization services provided by the deployment of RMS, a Windows RM client, a unique RM Lockbox DLL created



for each computer, and RM-capable applications. RM-capable applications include Word 2003, Excel 2003, PowerPoint 2003, Outlook 2003, applications from participating ISVs, and even your own applications. RM also uses two Microsoft Web services to help establish a collaborative trust ecosystem to help protect users from the vulnerabilities inherent in software and provide assurance. These Web services don't store any of your information, and the keys they generate aren't the keys directly involved in the subsequent licensing of RM-protected information to users. Therefore, Microsoft can't decrypt information protected by an organization's RMS servers. Together, the RM components use encryption to enforce control over your information according to the trusted entities you define as well as usage rights and conditions.

An RM system issues certificates and licenses for protecting documents and other information created with RM-capable applications. There are two kinds of certificates. RM uses account certificates to identify users in response to license requests and to encrypt information with a certificate's public keys so that only the certificate's owner can decrypt the information using the corresponding private key. RM also uses client computer certificates when handling license requests and to protect the account certificates and licenses stored on the local computer.

There are two types of licenses. Each RM-protected document contains a publishing license, which holds the symmetric key used to encrypt the document's information and rules that specify who can use the information and how. I explain how the publishing license is protected below. The other type of license is a use license, which is created when an authorized person tries to use an RM-protected document. The RM-capable application must obtain a use license, which lets the application decrypt and render the document according to the

usage policies defined within the use license. Figure 1 shows the components of an RM system.

### SETTING UP RMS

To set up an RM system in your existing AD infrastructure, you first install RMS on one or more servers. One of the root technologies of an RM system is public key infrastructure (PKI). Each RMS server is, in part, a Certification Authority (CA). Just like a PKI, your RM system can consist of just one RMS server or a hierarchy of servers. RM systems with multiple RMS servers have one root RMS server and one or more subordinate RMS licensing servers. To support availability and scalability, the root RMS server can be a server cluster. Small, one-server installations can use Microsoft SQL Server Desktop Engine (MSDE) Service Pack 3 (SP3) as RMS's data store. SQL Server 2000 SP3 is recommended for multiple and/or clustered RMS-server installations. RMS, fundamentally, is an ASP.NET Web Service that is the front-end to the aforementioned SQL or MSDE database. (For general

information about RMS, go to <http://www.microsoft.com/windowsserver2003/technologies/rightsmgmt/default.aspx>.)

The root RMS server issues licensing certificates to subordinate licensing servers, client computer certificates to each computer that participates in the RM system, and account certificates to each user in the RM system. You must obtain a Windows RMS Licensor Certificate for the root RMS server from the Microsoft Enrollment Service over the Internet. Microsoft's involvement ensures that the RMS ecosystem can be trusted by all parties that play a part by authenticating the organizations. First, you install a VeriSign Class 3 X.509 certificate on the root server. Then, the root server must obtain a Windows RMS Licensor Certificate from the Enrollment Service, which uses the information that the root server passes it solely to issue your certificate. Next, you set up licensing servers as necessary; then you set up your client computers.

You must install the RM client component on each computer that you will use to publish or use RM-

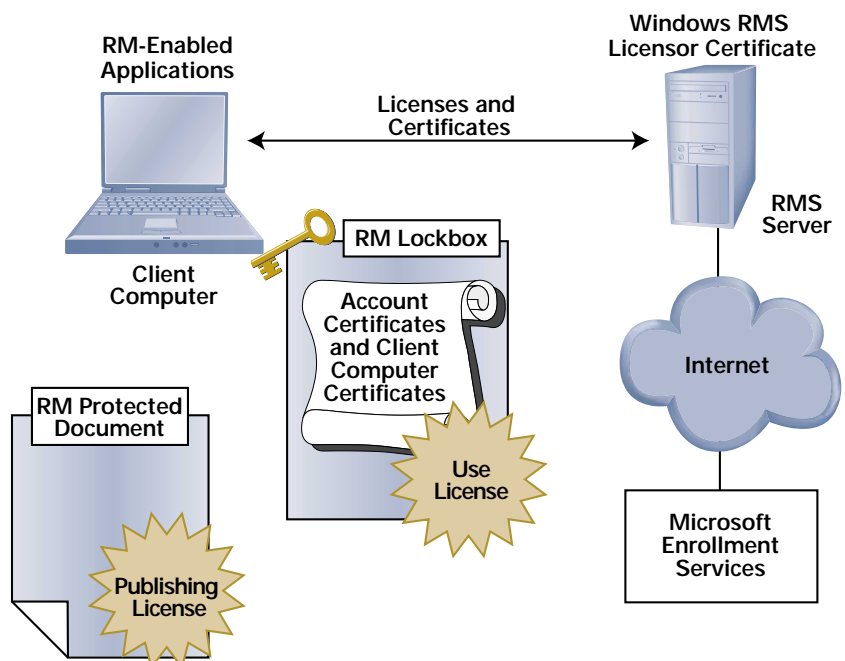


FIGURE 1: Specifying a template's name and users

protected information. To install the RM client, you can use software deployment tools such as Windows Update, Group Policy's Software Installation feature, or Microsoft Systems Management Server (SMS). Then, you activate each computer so that it becomes a trusted entity. When you activate a client computer, it generates a unique hardware identifier (HID). Next, the client creates a one-way hash of the HID and sends it to the RMS server. The RM Activation Service uses the information to create the client computer's certificate and public/private key pair and embed them in a heavily obfuscated unique RM Lockbox DLL. The RM Activation Service then sends these items back to the RMS server, which in turn sends them back to the client. Microsoft doesn't store the HID hash or other customer information. After your client computers are activated and have RM-enabled applications installed, you can use the computers to publish or access RM-protected information.

Because RM is a software-based rights management solution, you might wonder about how easily an attacker could break RM by inserting arbitrary code into or otherwise tampering with RM client files or RM-enabled applications. All significant RM encryption and decryption takes place inside each client's unique RM Lockbox, which contains the computer's private key. Microsoft put significant effort into designing the RM Lockbox so that it's useless if stolen and used on another computer. Microsoft also implemented a number of safeguards that make it extremely difficult, if not impossible, for an attacker to dissect the RM Lockbox and insert his or her own arbitrary code at crucial execution points. To guard against attackers tampering with RM-capable applications, Microsoft requires ISVs to sign a set of application manifests with a key that Microsoft adds to the root RM key hierarchy, making the application a

certified RM application. When you load an RM-capable application, the RM client on the computer checks the bytes of the executable you just loaded against the signed manifests to make sure that the executable in memory is the same as what the ISV entered into the signed manifests when they created the application. Thereafter, the RM client regularly scans areas of the computer's memory relevant to RM to detect any malicious attacks against that application.

Finally, users must be certified for RM on the computers they'll be using. The first time a user tries to publish or access RM-protected information, the user's workstation uses Windows authentication and the user's AD user account to authenticate to the RMS server. The RMS server creates an account certificate based on the user's authenticated credentials and a public/private key pair. The RMS server then encrypts the user's private key using the client computer's key. The client computer stores the account certificate locally so that the user can publish or use RM-protected information from that computer. You can certify a user on as many computers as necessary. Each time, the RMS server creates a new account certificate that ties the user to the specific computer but that has the same public/private key pair as the user's other account certificates. After a user is certified on a client machine, he or she can publish or use RM-protected information from that computer.

### RMS IN ACTION

How do all these pieces work together? Here's an example of Office Professional Edition 2003's implementation. Fred creates a Word document and enables RM protection for the document with usage policies that limit the document's use to read access by Tom and Sarah and that end that access on December 1, 2003. Using the RM client APIs, Word 2003 professional edition creates a publishing license that states the usage

policy that Fred defined. Word then generates a symmetric key that will serve as the core encryption key for the document, encrypts the document and specific parts of the publishing license with the key. The key is then encrypted by the user's Client Licensor Certificate (CLC) public key, a certificate issued to users by the RMS server for the purposes of publishing information, and embeds the publishing license in the document. Fred sends the document to Tom and Sarah.

When Sarah tries to open the document on her computer, Word 2003 realizes that it's RM protected and takes the following actions. The RM client checks Sarah's authentication using her Windows logon session credentials and her account certificate before sending the request. The client then determines the licensing server from the publishing license embedded in the document, and sends a request for a use license to the server. The request includes Sarah's account certificate and the publishing license. The licensing server checks the publishing certificate to make sure she is a named user of the document. Then the server creates a use license.

To build the use license, the server first takes the symmetric key from the publishing license and decrypts it using the CLC private key. Then the server re-encrypts the symmetric key using the public key of Sarah's account certificate and puts the symmetric key in the use license along with any conditions, such as that Sarah's read access expires on December 1, 2003. When the server sends the use license back to Word 2003 at Sarah's workstation, Word decrypts the symmetric key using Sarah's private key, then uses the symmetric key to decrypt the document and displays it for her to read. (I've said that Word performed the encryption, certificate, and license operations, but the computer's unique RM Lockbox DLL actually performs these tasks on the application's behalf because the lockbox.dll must run in-

FIGURE 2: Specifying a template's name and users

## Template properties

Specify template name and description.

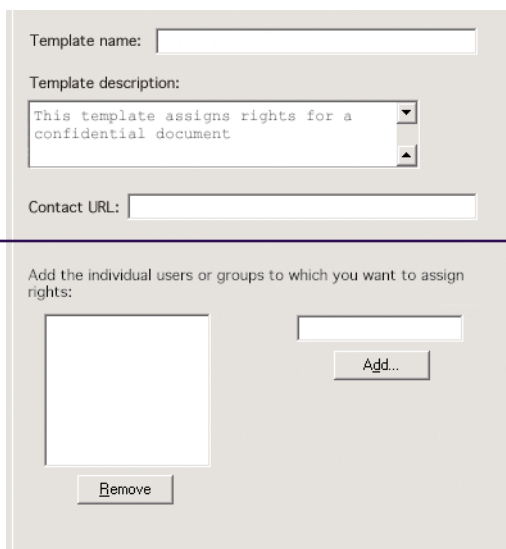
Specify a URL (can be mailto:) that users can navigate to, to request more permissions than those granted by template.

## Users

Add or remove individual users or groups that you want to assign rights for. Use the format `userid@company.com`.

In addition you can use the special role `Anyone` by simply typing it directly into the user field.

Select an individual user or group to specify rights for that user or group.



## ADDITIONAL RM FEATURES

As you can see, RM provides protection that travels with information wherever it goes. RM strikes a profitable balance between security and usability that lets you use RM-enabled applications to protect any kind of information. RM includes features for connecting your RM system with the RM systems of your partners as well as letting users publish RM-protected information even when they're offline and can't communicate with your RMS server. RM includes security management features for revoking compromised certificates, logging certificate operations, and logging licensing operations.

Attackers are always discovering new vulnerabilities in software that the attackers then exploit, so RM lets you create exclusion policies that block compromised software. As a failsafe against possible exploits in the RM Lockbox, you can exclude computers that have an earlier Lockbox version. To handle the more likely event of security holes in RM-enabled applications, you can exclude specific applications or application versions. To help users easily select the right level of protection, the appropriate group of users, and other usage policies, you can create policy templates that map to corporate, departmental, or project-specific classification levels such as Internal, Confidential or Private, or Project 23 Eyes Only. Figures 2 and 3 show the dialog boxes for creating a policy template.

RM is flexible enough to handle offline users, trusted business partners, and Passport-authenticated users on the Internet but strong enough to stand up to compromised software and user mistakes or intentional misuse. With RM, companies can stay in control of their information, even when it leaves their doors. ◆

FIGURE 3: Specifying a template's policies

## Template policies

Specify expiration and global template policies.

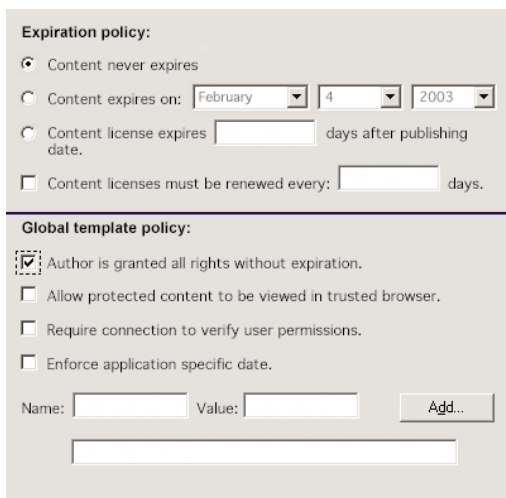
Specify the date on which content expires, and users can no longer view the content.

Specify the period in days after publishing that the license for the content expires. Content will no longer be accessible after this period expires.

Specify the period of time before an use license created from a template must be renewed.

Specify whether the author of the document has full perpetual rights.

Specify whether to allow licensed content to be viewed in the Microsoft Trusted viewer.



process with all RM-capable application while they are performing RM operations. The application and RM Lockbox DLL communicate via the RM APIs. Thus, account certificates and licenses are protected from theft by malicious code on the computer.)

On Tom's end, his laptop is stolen after he receives the file from Fred. However, the document is safe because it's encrypted with the symmetric key. The symmetric key is stored with the document in the publishing license, but the key is encrypted with the licensing server's

public key. Only the licensing server can decrypt the key and let it be used by someone else. If an attacker tries to add himself as an authorized user to the document's publishing license, the licensing server will detect the tampering because the publishing license is signed with a key belonging to the server. Attackers can't forge their own account certificate because account certificates are signed by the root RMS server for the organization and are checked by the RMS licensing server whenever the user requests a use license.

# TANDBERG Television Deploys *Aspelle Everywhere*<sup>™</sup> to Facilitate Remote Sales Demonstrations

TANDBERG Television (Oslo, Norway) is a market leader in providing open solutions for the digital broadcasting of audio, data, and video across various networks. The company, an innovator in digital broadcasting, has operations in Asia, Australia, Europe, and the United States. TANDBERG Television's customers include major broadcasters, network operators, and convergence players around the world.

## THE CHALLENGE

Having a global sales force in place is not enough; a successful company must have tools that enable the sales force to dazzle prospective clients with the breadth of its technology offerings.

Initially, TANDBERG Television gained this competitive advantage for its sales force by deploying a dial-up virtual private network (VPN) to facilitate access to its corporate resources. However, the company soon realized that a traditional VPN solution was too complex, expensive, and troublesome for most of its sales teams. The VPN was resource intensive and was unable to deliver the level of security, functionality, and flexibility the company needed. Employees experienced problems with the delivery of email and other critical business functions while using the VPN. And while it provided secure access to the corporate network, dial-up restrictions such as speed and latency were problematic for a company building its reputation as a pioneer in digital broadcasting.

TANDBERG Television needed a fast, secure, and easy-to-use solution that would let its mobile sales force easily gain access to corporate resources and, more importantly, serve as a mechanism for demonstrating broadcast products in real time at a customer site or in the field. The company also wanted to eliminate some of the multiple points of failure associated with its VPN solution.

Because some of TANDBERG Television's business partners require a VPN to meet their security guidelines, the company could not eliminate it as a resource. Instead, it selected *Aspelle Everywhere* to augment and support its VPN, making it more flexible and better able to meet the needs of its sales force.

## THE SOLUTION

The first product on the market to provide secure, managed, clientless access to all corporate applications without the limitations found in a VPN, thin client, or remote access product, *Aspelle Everywhere* requires only a Web browser and an Internet connection. In less than one day, TANDBERG Television installed, tested, and deployed *Aspelle Everywhere*, delivering a more comprehensive remote access solution to its sales employees.

"Once we saw a demonstration of *Aspelle Everywhere* the decision process was simple," said Lindsay Morgan, IS Service Delivery Manager for TANDBERG Television. "Using *Aspelle's* product to access live demonstration systems has enabled our sales force to show potential customers complex broadcast offerings that were previously unavailable or logistically difficult to display. As a result, we have a competitive edge in the field, which has helped us secure sales—which in turn has already paid for our investment in the system."

*Aspelle Everywhere* delivers a comprehensive, secure enterprise access solution. Designed to meet the demands of TANDBERG Television's global sales force, *Aspelle Everywhere* maximizes the company's previous technology investments by seamlessly integrating with its existing VPN and network.

"*Aspelle Everywhere* was the most cost-effective and easily deployed remote access product we reviewed," commented Morgan. "Additionally, we were very pleased with *Aspelle's* customer support during the implementation process. With the support of their dedicated sales and engineering team, we were able to test and implement the product in one working day."

## KEY DECISION FACTORS

*Aspelle Everywhere's* enterprise-class security features were particularly appealing to TANDBERG Television. With *Aspelle Everywhere*, the company can maintain all corporate applications inside the enterprise on existing network ports without exposing additional ports to the Internet-facing firewall. Data is 128-bit encrypted and authentication, authorization, and access are controlled separately. Moreover, *Aspelle's* innovative single sign on technology lets TANDBERG Television employees connect from public resources (e.g., Internet café, client computer) without risking key data.

TANDBERG Television also benefits from *Aspelle Everywhere's* easy-to-use management capabilities. It offers a secure, centralized point of management and the uniform delivery of corporate applications, eliminating many time-consuming and costly administrative and maintenance tasks. And *Aspelle Everywhere* was designed to expand to new locations in minutes, merge with new businesses overnight and add new users and applications in seconds.



# DigitalPersona's Fingerprint Recognition System Adds Needed Level of Security for Telecommunications Company

Telmex is a leading telecommunications company in Mexico and currently holds the distinction of being the largest private company in Latin America. With technology investments of more than \$27 billion, Telmex prides itself on a history of successful technology innovation. Supporting more than 100,000 users, the Telmex IT infrastructure is complex and diverse, and it includes IBM mainframes; IBM AS/400; Unix servers; Microsoft NT servers; desktop PCs, and laptops running Microsoft Windows 95, 98, 2000, and XP.

## SECURITY CHALLENGE

Given the challenges of a diversified infrastructure, Telmex embarked on a project to reinforce its security infrastructure and simplify login for employees. The security group was concerned about potential risks associated with single-password access. As a result, they concluded that they would benefit from a stronger authentication solution to ensure the latest, best-of-breed security technology at desktops and servers throughout the organization. The need for a complementary and stronger method of access auditability was also identified as a desirable outcome.

Telmex also identified a number of shortcomings to a token-based approach they were using. These shortcomings included the high rate of token loss, the costs of replacing the lost tokens, and the disruption of managing through the replacement process. In addition, tokens can be easily shared and therefore failed to meet the absolute identity requirement.

## THE SOLUTION: U.ARE.U PRO FOR AD

Incorporating DigitalPersona's U.are.U Pro for Active Directory into the security infrastructure satisfied both objectives. In both cases, DigitalPersona's fingerprint authentication system strengthened single sign-on authentication to meet the security criteria established by Telmex, addressed the need for stonger auditability, and created a more convenient solution for end-users and administrators.

DigitalPersona's U.are.U Pro was first installed for senior managers, providing them with secure, convenient, and non-shareable access to sensitive data. Although a solution built on the existing Microsoft Windows NT domain was initially planned, Telmex and Biometria Aplicada determined that a combination of U.are.U Pro for Active Directory, Active Directory, and Windows 2000 for the desktop, would better meet Telmex's stringent security and performance requirements.

In addition to the benefits previously noted, replacing tokens provided significant additional value to Telmex. Specifically, it eliminated the recurring costs of the tokens and the token server, as well as the replacement costs for lost tokens.

## THE RESULTS

Implementing U.are.U Pro for Active Directory and Active Directory has enabled Telmex to further increase the security and auditability of access to their corporate network and applications, enhance the security of its single sign-on system, and reduce administrative support costs. Telmex also has increased the security of senior executives, protecting access to their systems and to laptops used the Telmex network.

"U.are.U Pro for Active Directory has enabled us to strengthen the Telmex security schema," said Javier Barajas Gonzalez, Computer Information Security Officer, TelMex. "We are now more comfortable that the person logging in to our network has been unambiguously identified."

Overall, already strict security policies have been further strengthened while end users have a more streamlined user experience defined by "single touch; single sign-on."

## ABOUT DIGITALPERSONA™

DigitalPersona, developer of the innovative U.are.U® fingerprint recognition systems, brings both heightened security and convenience to business and government. As proven by independent tests and millions of customers, U.are.U delivers time-tested, leading-edge biometrics technology.

## ABOUT BIOMETRIA APLICADA

Biometria Aplicada, is a premium DigitalPersona partner in Mexico City dedicated to providing biometric authentication systems.



digitalPersona

DigitalPersona

650-261-6070

sales@digitalpersona.com

www.digitalpersona.com

# PaeTec Communications Extends Corporate Policies to the Internet with N2H2's Sentian

Based in Fairport, New York, PaeTec Communications, Inc. is an integrated communications provider offering local, domestic, and international long-distance services, high-speed Internet access, advanced data services, and communications management services to medium- and large-sized businesses, colleges, universities, hospitals, hotels, governmental organizations, and affinity groups in more than 27 markets nationwide. Recently ranked the Top Private Company in the greater Rochester, New York area, PaeTec has experienced 134,000 percent growth since its inception in May 1998.

## THE CHALLENGE

The PaeTec company Web site features a statement that describes "PaeTec Culture." The statement says, "PaeTec's culture is one in which employees look forward to excelling at their jobs and our customers look forward to interacting with our employees. Every employee of PaeTec is a part owner, thus creating an operating environment where teamwork and consideration for customer concerns naturally occur."

When PaeTec provided Internet access to its employees, company management became concerned about the potential display of offensive and inappropriate material, because this could undermine the environment of mutual respect and teamwork the company had worked to build. Other companies have experienced nightmarish problems related to unfiltered Internet use. In 1999, the Dow Chemical Co. fired 50 employees and disciplined 200 others after an e-mail investigation turned up hard-core pornography and violent subject matter.

## THE SOLUTION

PaeTec Director of Enterprise Technology, Jim Raub, began evaluating filtering solutions in December 2001. Raub was seeking a solution that allowed him to further extend company policies to the Internet by setting different levels of Internet access for different departments across PaeTec. In order for such a solution to not become an administrative headache, the solution would need to have a simple interface. Additionally, the solution would need to have a highly accurate and comprehensive database of blocked sites, so that IT resources were not wasted by constantly adding and deleting sites to the database.

Raub found that N2H2's Sentian for the Microsoft Internet Security and Acceleration (ISA) Server 2000 more than met these requirements. "Sentian has wonderful administrative abilities, with an interface that is great for applying rules and levels of access to different groups of users," said Raub. "The quality of the database is very good, and the few times we have needed to customize our block list, it was a fast, easy fix. I would recommend Sentian without hesitation."

***"Sentian has wonderful administrative abilities with an interface that is great...I would recommend Sentian without hesitation."***

**Jim Raub**

Director of Enterprise Technology  
PaeTec Communications, Inc.

For PaeTec, managing Internet access makes sense. PaeTec management can feel secure in knowing that their company mission statement is being carried out, in part, by using filtering software, which helps to keep the PaeTec workplace a safe, productive, team environment.



**Sentian**<sup>™</sup>

**N2H2**<sup>™</sup>

INTERNET CONTENT FILTERING

N2H2

877-336-2999

[www.n2h2.com](http://www.n2h2.com)

# Websense Enterprise Protects Carnival Cruise Line From Rising Tide of Internet Threats in the Workplace

As the Internet has become more prevalent in the workplace, it has also created new distractions for employees that represent new risks for companies. These risks extend beyond non-work-related Web sites and include peer-to-peer (P2P) file sharing and instant messaging (IM) protocols and applications. In addition, the problem of illegal content in corporate networks is quickly becoming a significant legal issue for corporations.

## PROACTIVE RISK MANAGEMENT

Carnival Cruise Lines decided it was necessary to proactively manage its employees' computing resources both on land and at sea. After an extensive product search and evaluation, Carnival selected Websense software as its employee Internet management (EIM) solution of choice.

Carnival uses Websense Enterprise, integrated with Microsoft Internet Security and Acceleration (ISA) Server 2000, to manage Web use for more than 3,000 Internet-enabled employees at the company's corporate headquarters in Miami. In the future, the company plans to roll out Websense software in other office locations and on its fleet of ships.

"We chose Websense because it was by far the best Internet filtering product on the market. With its ease of management, comprehensive and flexible features and ability to integrate with a host of products, it was the obvious choice to make the Net more productive for our employees," said Rodney Orange, lead systems engineer, Carnival Cruise Lines. "Websense won easily over the competition with its flexible management tools and ability to integrate with any number of existing Internet infrastructure products, such as Microsoft ISA."

Although a company may initially recognize the bandwidth drain of employees' non-work-related Internet use, it is the security risks associated with the Internet, such as P2P file sharing and IM, which are far more troublesome. These threats can enter a company in three ways: via Web access, by tunneling across various network protocols and by launching on individual employee desktops. As a result, it's important that IT administrators install EIM software that offers all three levels of protection to manage employee-computing security.

## MANAGE P2P FILE SHARING

In addition to draining corporate bandwidth, P2P applications carry security risks because they communicate directly with other users' computers, bypassing a compa-

ny's firewall and often entering the corporate environment without being scanned for viruses.

According to IDC, P2P is primarily used for swapping copyrighted material, so that the problem of illegal content in corporate networks is quickly becoming a significant legal issue for corporations. As a result, the Recording Industry of America (RIAA), which estimates that more than 2.6 million files are copied illegally every month, has warned companies that they could be held liable for violating copyright laws.

With Websense Enterprise, companies can manage P2P file sharing by blocking employee access to P2P Web sites, network-level protocols or desktop applications.

## STANDARDIZE IM

Because IM is a form of P2P file sharing, companies face security vulnerabilities. Up to 84 percent of all organizations use some sort of IM application, according to a report issued by Osterman Research. Unfortunately, employees at less than one-third of the companies surveyed use approved IM software.

With Websense Enterprise, companies can standardize IM by blocking employee access to all network-based IM protocols except the one designated by corporate management. In addition, companies can easily enforce a policy in real-time that minimizes bandwidth drain by blocking IM, except by those departments for which it's an essential business function.

As employees' computing environments have evolved, so too have the threats to infiltrate companies. Because the threats can enter a company in three ways, they require a three-tiered EIM solution. Without all three levels of Internet security protection, a company is as vulnerable to threat from employee Internet use as a cruise ship passenger is to seasickness in the midst of a storm.



Websense, Inc.

800-723-1166

sales@websense.com

www.websense.com

# Microsoft Partners

## TECHNOLOGY PARTNERS

### 8e6 Technologies

Orange, CA  
www.8e6technologies.com/isaserver

### ACP

Birmingham, AL  
www.acp-inc.com/isaserver

### Aelita Software

Powell, OH  
www.aelita.com/isaserver

### AEP

Boston, MA  
www.aep.ie/isaserver

### Akonix

San Deigo, CA  
www.akonix.com/isaserver

### Aladdin Knowledge Systems

Arlington Heights, IL  
www.ealaddin.com/isaserver

### Aspelle

Boston, MA  
www.aspelle.com/isaserver

### Authenex

Oakland, CA  
www.authenex.com/isaserver

### Bindview

Houston, TX  
www.bindview.com

### Burst Technology

Bonita Springs, FL  
www.burstek.com/isaserver

### Castify Networks

Alexandria, VA  
www.castify.net/isaserver

### Cereton

Waltham, MA  
www.cereton.com/isaserver

### Chutney Technologies

Atlanta, GA  
www.chutneytech.com/ISAserver/

### Cobion

Kassel, Germany  
www.cobion.com/isaserver

### CornerPost Software

Duffield, VA  
www.cornerpostsw.com/isaserver

### F5 Networks

Seattle, WA  
www.f5.com/isaserver

### Finjan Software

Los Gatos, CA  
www.finjan.com/isaserver

### GFI Software

Cary, NC  
www.gfi.com/isaserver

### Intellitactics

Bethesda, MD  
www.intellitactics.com/isaserver

### FutureSoft

Houston, TX  
www.futuresoft.com/isaserver

### ITWorx

Burlington, MA  
www.fileway.com/ISAServer.htm

### Internet Security Systems

Atlanta, GA  
www.iss.net/isaserver

### N2H2

Seattle, WA  
www.n2h2.com/isaserver

### nCipher

Woburn, MA  
www.ncipher.com/isaserver

### NetIQ

San Jose, CA  
www.netiq.com

### Network Associates

Santa Clara, CA  
www.nai.com

### Nexus Technology

United Kingdom  
www.webconsent.net/isaserver

### Oblix

Cupertino, CA  
www.oblix.com

### OpenNetwork

Clearwater, FL  
www.opennetwork.com

### Panda Software

Buenos Aires, Spain  
www.pandasoftware.com/isaserver

### PatchLink Corporation

Scottsdale, AZ  
www.patchlink.com/isaserver

### Radware

Mahwah, NJ  
www.radware.com/ISAServer

### Rainbow

Irvine, WA  
www.rainbow.com/isaserver

### Rainfinity

San Jose, CA  
www.rainfinity.com/isaserver

### RSA Security

Bedford, MA  
www.rsasecurity.com/isaserver

### Sane Solutions

North Kingstown, RI  
www.sane.com/isaserver

### Secure Computing Corporation

San Jose, CA  
www.securecomputing.com/isaserver

### Stonesoft

Atlanta, GA  
www.stonesoft.com/isaserver

### SurfControl

Scotts Valley, CA  
www.surfcontrol.com/isaserver

### Symantec

Cupertino, CA  
www.symantec.com/isaserver

### Trendmicro

Cupertino, CA  
www.trendmicro.com

### Venation

United Kingdom  
www.venation.com/isaserver

### Wavecrest Computing

Melbourne, FL  
sales@wavecrest.net  
www.wavecrest.net/isaserver

### Websense

San Diego, CA  
www.websense.com/isaserver

### WebSpy

Kirkland, WA  
sales@webspy.com  
www.webspy.com/isaserver

## GLOBAL DELIVERY PARTNERS

### Accenture

Chicago, IL  
www.accenture.com

### Avanade

Seattle, WA  
www.avanade.com

### Ernst and Young

Boston, MA  
http://ey.com/security

### HP

Boston, MA  
www.hp.com/hps/tech/security/

### IBM

San Jose, CA  
www-1.ibm.com/services/its/us/  
/isa\_server.html

### PricewaterhouseCoopers LLP

New York, NY  
www.pwcglobal.com/security

### Schlumberger SEMA

Houston, TX  
www.slb.com

### Unisys

Boston, MA  
www.unisys.com/security

## REGIONAL DELIVERY PARTNERS

### Canada

**CMS Consulting Inc.**  
Toronto  
www.cms.ca/isaserver/

### Codefusion Communications Inc.

Toronto  
www.codefusion.com/isaserver/

### CyberSecure, Inc. Rothsay

www.cybersecure.ca/isaserver/

### LegendCorp

Toronto  
www.legendcorp.com/isaserver/

## U.S.

### ACP

Birmingham, AL  
www.acp-inc.com/isaserver/

### Convergent Computing

Oakland, CA  
www.cco.com/isaserver.htm

### EYT

Chantilly, VA  
www.eyt.com/isaserver

### Netivity Solutions

Waltham, MA  
www.netivitysolutions.com/  
/isaserver/

### Quilogy

St. Charles, MO  
www.quilogy.com/isaserver

### RDA

Atlanta, GA  
www.rdacorp.com/ISAServer/

### Corbett Technologies

Alexandria, VA  
www.corbett-tech.com/isaserver/

### Extreme Logic

Atlanta, GA  
www.extremelogic.com/isaserver

### Guardent, Inc.

Providence, RI  
www.guardent.com/isaserver

### InDepth Technology

Dublin, OH  
www.indepthtech.com/isaserver/

## SECURITY RESELLERS

### CDW

Vernon Hills, IL  
www.cdw.com/isaserver

### Insight

Tempe, AZ  
www.insight.com/isaserver

### Zones, Inc.

Renton, WA  
www.zones.com/isaserver