PRODUCT Reviews Guide 2007

TESTING & ANALYSIS TO HELP YOU MAKE PURCHASING DECISIONS

Information Security **magazine's 2007 Product Review Guide is a compilation of the single and comparative reviews published in 2007, an indispensable guide for information security managers tasked with evaluating and purchasing security hardware and software in 2008.**

contents

INFORMATION SECURITY

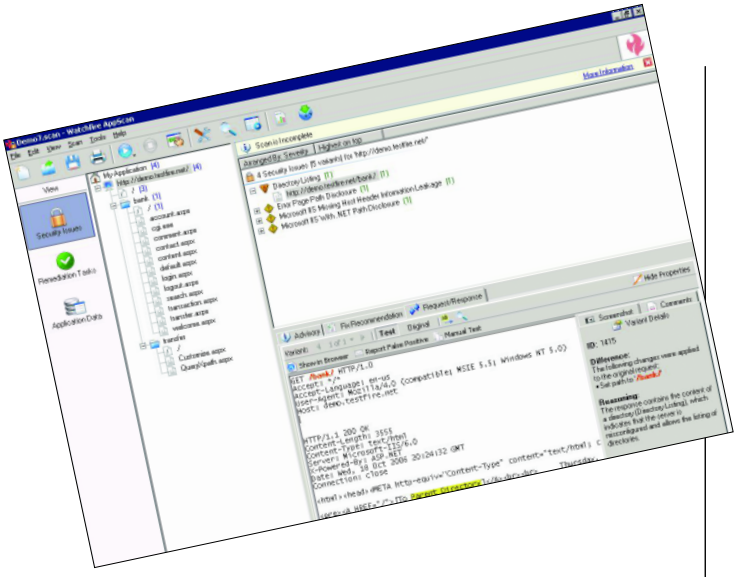

APPLICATION SECURITY

AppScan 7.0

REVIEWED BY PHORAM MEHTA

Watchfire

www.watchfire.com

Price: **Starts at $14,400; Reporting Console (including AppScan 7.0) starts at $35,000**

The failure to incorporate sound security practices into software development has left business-critical Web applications open to attack, but that's changing as corporations adopt secure coding requirements. To that end, Watchfire's AppScan 7.0 provides sound application security testing for developers, quality assurance teams and penetration testers.

Installation/Usage B+

The wizard-driven installation took five minutes; AppScan runs on Windows XP, Vista or 2003 Server.

To initiate a scan, a wizard walks you through the information required, from assessment type (Web application or Web service), starting URL, login parameters, test policy (default, app only, infrastructure, invasive) and scan options (full scan or explore/crawl). There are plenty of advanced settings and customization options, like two-factor recorded login and privilege escalation.

There are more than 75,000 individual security checks distributed across various policy files; advanced users can create custom tests in a few steps.

Advanced Features B

AppScan has tried to create a one-stop solution for Web application and services assessment by incorporating multiple advanced techniques. Tools like HTTP Request Editor, Encode/Decode and Regex Tester come in handy for vulnerability assessment and other QA tests. You can add external tools by linking to the executable.

Above all, AppScan gives you a single interface to open all the tools and techniques required to test your Web apps. Users have lots of options, from customizing existing policies to recording two-factor login information. Unfortunately, the login information is not stored in an encrypted format.

Testing methodology: AppScan 7.0 was run multiple times using default and custom settings against two production Web applications based on .NET, PHP, Apache Tomcat, Oracle and others.

Performance B

The AppScan dashboard gives users multiple real-time views of the structure, results summary and details of vulnerabilities discovered. The number and severity levels of vulnerabilities are displayed in the bottom taskbar.

We ran the tool against two production Web applications, both of which handle sensitive data and use different application and infrastructure technologies. AppScan discovered common issues, and a few subtle flaws.

We weren't blown away with the scanning speed, but were impressed with the adaptive scanning technique: Once the tool determines that a particular technology, say IIS, is not used, it removes all the corresponding tests from the queue.

If you elect to report a false positive to Watchfire, AppScan generates an unencrypted email to the tech support team, so be sure to scrub any sensitive data from the files before sending the email.

Reporting A

AppScan's reporting capabilities are as good as we've seen in any tool. Report categories include security, industry standard, regulatory compliance and delta analysis. Each of these categories has multiple templates and options to customize reports. Reports can be exported in numerous formats.

AppScan Reporting Console (sold separately) enables users to consolidate vulnerability data into one centralized location to better control who has access to sensitive data. Because it is Web-based, you can create dashboards and for multiple users, such as QA and development.

Verdict

Consultants and in-house app security testers will appreciate AppScan's accuracy and efficiency. The reporting options alone are enough to wow management.

DATABASE SECURITY

Guardium SQL Guard 6.0

REVIEWED BY JAMES C. FOSTER

Guardium
www.guardium.com
Price: **Starts at $50,000**

In an industry flush with products for securing the network perimeter, Guardium's SQL Guard 6.0 serves as an important addition for monitoring and managing connections to and from a wide variety of enterprise database products.

SQL Guard continues to address one of the most typical database audit failure points. Most auditors will not issue a "pass" if you leverage a database's native logging features because they are owned and controlled by the groups you are trying to monitor (for example, DBAs should not be responsible for configuring and monitoring DBAs). SQL Guard ensures a system of checks and balances between the security and database engineering teams.

The solution consists of local database agents, network-based appliances to passively gather traffic or to actively work as a firewall, and aggregation servers that collect and analyze data.

Installation/Configuration B

The preconfigured Linux-based 1U Dell appliances can be plugged directly into the span port on a switch that controls traffic to the databases. The administrative account is created during installation, along with a series of default user roles—common users, administrators, DBAs, security, application developers, auditors, network engineering—that can be used to create other accounts.

Passively collecting network traffic is as easy as running a sniffer; installing agents will require admin credentials and console-level access.

The classification feature helps you identify potentially sensitive information on a live database. You can create rules based on SQL Guard's Perl Compatible Regular Expression (PCRE) engine to search for data, specific permissions, or even conduct a catalog search. The results can be categorized and assigned additional rules for protection.

You can create any number of levels of classification depending on the complexity of your environment or business (low, medium, high, or severe, critical, sensitive, compliance, etc.).

Reporting A

Guardium has all of the bases covered here. Reports are grouped and labeled under three tabs for templates, custom reports and alerts. Templates include high-level or technical information on database activities, sensitive object usage, data markup language exceptions, overall performance and permanent schema changes.

The strong custom reporting is built atop a SQL querying engine.

The new incident management dashboard provides a clear-cut summary on policy violations and incidents. It permits you to quickly dig deep into the incident, via a click, to identify the timestamp, source/destination IP, user, full SQL string, technical incident specifics and more. The breadth of information is impressive.

Management/Monitoring B+

In addition to monitoring database connections, 6.0 has added application layer monitoring, providing JD Edwards, Oracle, PeopleSoft, SAP and Siebel filters.

Alerts are triggered in one of two ways: statistical or real time. Both save the same type and amount of data; however, one is merely logged into the back-end Guardium database and the other is logged and then passed to one of four notification mechanisms.

Organizations looking to monitor databases in real time will be best served leveraging SQL Guard's integration capabilities as opposed to its console. SQL Guard can easily integrate with SIEM or aggregation platforms via SMTP, SNMP, syslog, or a custom Web-based Java class.

Verdict

SQL Guard has evolved from an impressive technology to an enterprise-class data security product that should be on every organization's radar.

Testing methodology: We tested a Guardium G2000 appliance testing a lab that contained DB2 8 and Oracle 9i and 10g on Linux 2.6, Informix 7 on AIX 5.3, SQL Server 2000 and 2005 on Windows Server 2003, and Sybase 15 on Sun Solaris 9.

DATABASE SECURITY

Lumigent Audit DB 6.0

REVIEWED BY JAMES C. FOSTER

Lumigent
www.lumigent.com
Price: **Starts at $10,500**

Lumigent Audit DB 6.0 helps organizations, particularly those with significant compliance-related issues and failed audits, protect and audit production databases.

The overarching problem remains that database administrators and engineers do not have the security background required to lock down their databases. Audit DB 6.0 is designed to audit, monitor and protect databases.

Audit DB has two main auditing components. The first captures network traffic to and from the database. Matching patterns or out-of-policy actions are identified, triggering alerts. The second mode uses agents to pull database log files, activity information and general database configuration information.

Installation/Configuration B

Installing and configuring Audit DB is not an afternoon project. Plan for key stakeholders within the database, network and security teams to provide input into the Lumigent product. You will need system and database administrative credentials for the target databases as well as admin-level access for the system that is going to house the Audit DB reporting engine. Network administrators will help identify placement of the Audit DB NetWatch sniffer agents, which can reside on the target databases or on nearby systems. Your design will depend on the number of databases you must audit; for multiple databases, it's probably best to monitor them centrally as opposed to leveraging individual agents. Each component's policy can be tailored to the rules that matter to the organization or geared toward a regulation.

Lumigent supports all major database platforms including Microsoft SQL Server, IBM DB2, Sybase and Oracle, all with good documentation.

Testing methodology: We tested Lumigent Audit DB 6.0 running on an Intel-Red Hat Enterprise Linux machine auditing an Oracle 10g R2 database on Red Hat Linux and SQL 2000 on Windows 2003 Server.

Reporting B+

Audit DB 6.0 is built atop a role-based reporting engine that allows you to create and schedule reports based upon organizational components, specific servers or technologies, and audit requirements. Its strength is its compliance/audit reporting capability, which supports SOX, GLBA, SAS 70, HIPAA, PCI, SB 1386 and Basel II. It also includes frameworks such as COBIT, COSO and ITIL. These preconfigured reports are easily customized.

Executive and/or managerial dashboard-level views allow you to drill down to the audited systems.

Of particular interest is Audit DB's user reporting capability. Most regulations place heavy emphasis on user provisioning and decommissioning. Audit DB has strong features meeting compliance standards for user management and can be a valuable tool for identifying obsolete or dormant users, validating password policies, and identifying privileged users and entitlements.

Management/Monitoring A

The most beneficial feature of Audit DB 6.0 is its API, which comes with a complete 98-page reference, flush with SOAP interface details and code samples. This integration allows database administrators and developers to leverage Audit DB's functionality natively within its environments. Through the API, database and application developers have the ability to write events, logs and other data to the Lumigent repositories, while administrators can automate common maintenance.

The SOAP interface is efficient and clean. If you've ever implemented an RSS or XML feed you won't have an issue integrating this feature.

The tabs in the Web-based management interface allow you to access all data collected by Audit DB. Drill-down reports take you into the details of an event with timestamps, user information, data sources and activity.

Verdict

Audit DB is a strong tool for organizations that are mandated to achieve and report compliance on their database servers.

DATABASE SECURITY

Oracle Database Vault

REVIEWED BY JAMES C. FOSTER

Oracle

www.oracle.com

Price: **Starts at $20,000 per CPU or $400 per user**

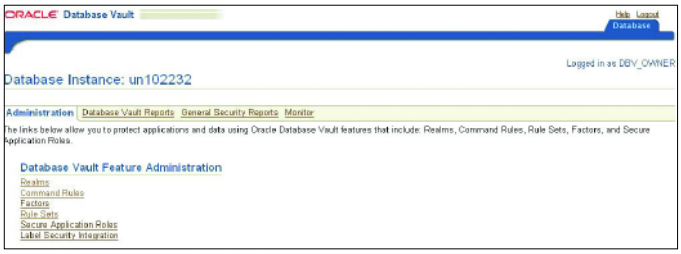

Oracle Database Vault enables advanced separation of duty to help organizations meet compliance and data security business challenges.

While database administrators and engineers may be responsible for securing, managing, backing up and performance tuning, they shouldn't need access to data. Vault allows admins and application owners to manage databases and applications without accessing credit card numbers, customer information, company secrets, etc.

Installation/Configuration B+

Set aside one morning to complete the installation; you'll be installing it on your current Oracle server (Oracle 9i R2, 10g and 11 are supported), and will need both system and database admin accounts—strong passwords are required. The configuration agent helps automatically configure the key components—adapter configurations, DNS name, host name and host file updates.

Security Features A

User Realms ensure data protections are implemented properly. A Realm is similar to a database software firewall. They can be put around an entire application or a particular table within an application.

Vault uses two technology concepts to control application access, Command Rules and Factors. Factors are properties or elements—users, IP addresses, network ranges and specific databases—that can be included within Command Rules.

Vault implements Command Rules to control the execution of SQL commands and can control Data Definition Language and Data Manipulation Language SQL commands. This level of protection can be useful in locking down permissions and accessibility for application service accounts and internal users alike. For example, a rule could be created to disallow any application user from executing a CREATE DATABASE LINK command on a particular database, a command that is typically reserved when creating applications. Or, you might prohibit any application user or service account from leveraging SQL INJECT commands to thwart injection attacks.

Auditing/Reporting B

Each created Realm can include auditing, or in this case, event logging. If enabled, auditing comes in two flavors, audit on failure and audit on success or failure. The audit on failure option enables you to see who is attempting to break the rules/Realms, while the more robust audit on success or failure option will give you a picture of everyone who successfully or unsuccessfully attempts to conduct an operation that is protected by Vault.

While the reporting options are straightforward and somewhat effective, there is room for significant improvement if you intend to use these for daily operations. For instance, it would be beneficial to run operational reports within specific windows of time, or to correlate events across all databases throughout the enterprise.

Global reporting allows you to analyze results from the entire database. The auditor and executive reports include high-level statistics such as number of successes and failures as well as user permission reports. User reports and statistics can help you identify users, their corresponding roles and access levels.

Verdict

Surprisingly mature for a first release, Oracle Database Vault may prove valuable for large environments that have made heavy investments in Oracle.

Testing methodology: We tested Oracle Database Vault on Oracle 10g with Red Hat Linux. All components of the application were tested to include user administration and application development.

INFORMATION LEAKAGE

Content Inspection Appliance 1500

REVIEWED BY MIKE CHAPPLE

Code Green Networks

www.codegreennetworks.com

Price: **Starts at $25,000 for networks with up to 250 users**

As organizations increasingly turn security focus from outside attackers to the threat from within, they are beginning to consider information leakage tools. Code Green's Content Inspection Appliance 1500 (CI-1500) is among the still-maturing handful of products designed to detect sensitive information leaving the enterprise.

Policy Control A-

Code Green's primary detection engine uses proprietary technology to develop many short "fingerprints" of each piece of submitted content.

Code Green also supports the use of regular expression matching rules to protect against users extracting content from databases and other structured data sources.

The challenge here is determining and managing rules for what constitutes sensitive data. This may be a straightforward task for companies with clearly defined document-classification policies, but it could pose a considerable challenge for less mature organizations.

You create policy rules based on content match, traffic source and destination, protocol, and desired action.

A document triggers an alert if it contains at least one fingerprint from the RedList of restricted content (the GreenList holds permitted content). Actions include notifying the administrator, retaining a copy of the offending document, sending a syslog report and blocking/rerouting email messages.

Testing methodology: The CI-1500 was tested in a network of Windows and Macintosh clients using a variety of common file formats, including Microsoft Office files and Adobe Acrobat documents.

We were impressed with the range of file types covered—390 according to Code Green—including all standard productivity software, such as Microsoft Office, Acrobat and WordPerfect, as well as other formats such as AutoCAD, Flash and .exe/.dll binary files.

Configuration/Management A

The appliance ships with default policies and settings. We simply booted it up, attached it to a network tap and were monitoring traffic.

We then explored the content registration and policy creation process. The CI-1500 allows you to register individual content files through a Web upload, scan Windows and NFS file shares, and integrates with Stellent and EMC Documentum content management systems.

We evaluated both the Web upload and Windows file share registration options and had no difficulty adding new content. We also created policies based on customized regular expressions.

Effectiveness B

We ran the CI-1500 through a variety of tests designed to thoroughly evaluate its content-inspection capabilities. It successfully detected Microsoft Office, Adobe Acrobat and text documents that we attempted to send via Web upload and email. We tried cutting and pasting document sections into other formats, and the appliance detected our circumvention attempt. It also scanned and detected protected content in .zip files.

The appliance's Achilles' heel, common to content-inspection products, is its inability to inspect encrypted traffic. An insider aware the organization performs content filtering can simply encrypt outbound traffic (for example, using Gmail) to bypass the appliance's scrutiny. The ideal solution would be to integrate a Web proxy server to allow for SSL decryption and inspection.

Reporting B+

The CI-1500 provides a fairly robust reporting mechanism. Administrators may use one of nine predefined reports (such as top-matched policies, policy violations by protocol and top-source emails of policy matches) or create their own based on specified criteria (such as time, policy violated, source/destination IP, content).

Verdict

Code Green's CI-1500 could prove valuable to organizations with well-defined classification schemes.

DOCUMENT CONTROL

Workshare Professional 5

REVIEWED BY MIKE CHAPPLE

Workshare

www.workshare.com

Price: **Starts at $349 per workstation**

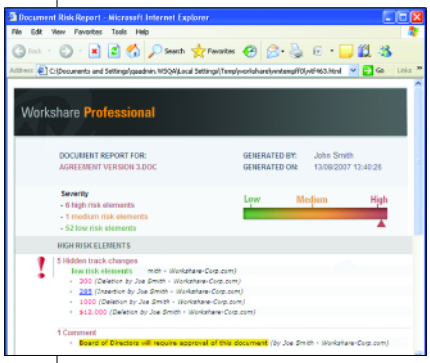

Workshare Professional allows organizations to maintain control over information throughout the document lifecycle. It provides the ability to send documents to internal and external reviewers, manage changes with a strong audit trail, and protect documents from data leakage.

Workshare has strong integration with Microsoft Office 2007 but is significantly limited in effectiveness by its ties to specific email platforms.

Configuration/Management A-

Installation of Workshare Professional on the client uses an intuitive wizard-driven process. Workshare automatically installs Office integration features, allowing the user to quickly begin operating in a Workshare environment. Office 2007 users will find a Workshare pane in the Office Ribbon that offers all Workshare functionality in a familiar format. Administrators create and manage policy through the Workshare Policy Designer in the desktop version of the product. Enterprise installations may distribute policies through a centralized server.

Policy Control A

Workshare allows for an impressive degree of granular policy control over user actions, allowing administrators to specify broad or narrow criteria and choose from a large set of policy violation responses. For example, we created a policy to intercept all outbound email including Microsoft Office attachments containing reviewer comments, and scrub those comments from the document before transmitting it to the end user. In addition to cleaning documents, administrators may create policies that block transmission, alert administrators to the violation, send the document for review prior to transmission, convert the document to a secure PDF or compress the file in a secure ZIP. Administrators may create policies based upon document contents (keywords, phrases or regular expression matches) or attributes (classification, hidden metadata, file name, file type, file size).

Effectiveness C

Workshare met all of its stated document control objectives and correctly enforced our defined policies when we attempted to send protected content via Outlook email. The biggest drawback to this product is its sole reliance upon three supported email environments: Microsoft Outlook, Lotus Notes and Novell GroupWise. This limitation makes the content distribution control features of Workshare useful only to prevent accidental data leakage or deter a novice attacker. In fact, we were able to defeat the content controls by simply uploading the file via a Web-based email connection. Workshare does offer a complementary network appliance that filters all SMTP and HTTP connections, but this device does not support encrypted connections, making it straightforward for a determined attacker to undermine the system's controls.

Reporting B+

Workshare offers three types of reports: an audit report that itemizes all changes made to a document by any reviewer, a history report that shows participation in the document review cycle, and a risk report that shows hidden metadata that has the potential for inadvertent information disclosure. Workshare will generate reports in either HTML or XML format.

Verdict

Workshare Professional is a good document control solution for organizations seeking to prevent inadvertent disclosure of sensitive information. Organizations should supplement it with strict controls on outbound network traffic to filter all content leaving an enterprise.

Testing methodology: We tested Workshare Professional 5 in a Windows XP environment using Microsoft Office 2007.

ENDPOINT SECURITY

Bit9 Parity 3.5

REVIEWED BY GREG BALAZE

Bit9

www.bit9.com

Price: **$35 per desktop**

Bit9's Parity 3.5 is designed to give you control over what users can do on company computers, and prevent executables from unauthorized or malicious apps from running on your desktops.

Configuration/Management B

Bit9 Parity Server was easy to install and didn't have much to configure. A step-by-step screen walks you through setting everything from IP addresses and ports to selecting the creation of a self-signed or previously generated certificate. It automatically installs SQL Server 2005 and Apache Web Server, which is used for remote administration.

Small client agents for Windows XP/2000 (Vista is coming) are generated or updated automatically when a policy is created or modified for a group. The agents can also be downloaded off the Web, or distributed by application deployment software such as SMS. The agent and server communicate via a SSL tunnel.

Policy Control B

Policies are applied based on groups set up within Parity Server that specify the file types it will block. Security condition levels, set by group, determine what happens when a file violates policy—various combinations of allowing or prohibiting file execution with or without notification. For policy enforcement, you can identify executables by name, or hash the file. Although malware can use an altered name to pose as a legitimate app, Parity will report on renamed programs. We recommend using hashes, though this means additional administrative overhead before deploying software.

Programs can be authorized to run from trusted individuals, trusted directories or trusted deployment applications, eliminating the need to manually add to the policy for each software deployment.

Recognizing the problems posed by mobile workforces, Bit9 allows for different security conditions when attached to the local network, and when disconnected.

Effectiveness A

Parity is effective at stopping programs from executing, as the agent goes through a lengthy process of inventorying the host workstation and reporting executable files to the Parity Server. This can take a while, especially in large enterprises with many clients.

Parity Server uses a combination of blacklisted applications and Bit9's signature database of known malware. The latter prevents the rapid spread of viruses and spyware from host to host by identifying the offending program and preventing its subsequent execution on other protected systems.

The Parity agent allowed executables to run according to policy, and quickly caught changes we made to a file. For example, we renamed Kazaa, a prohibited app, but still couldn't run it.

Reporting C

Bit9 has some work to do to beef up reporting capabilities. While several canned reports give quick access to important information, the sparse main reports page gives limited statistics on important file activities. We were disappointed that there's no way to graph the statistics, which would be especially useful for trending reports. There's no syslog support, nor can reports be exported to another format.

Verdict

Bit9 Parity 3.5 does a good job of preventing unwanted programs from running, although we didn't see any new methodologies or technologies that make it stand out from established competitors.

Testing methodology: We installed Parity Server on a Windows 2003 SP1 machine to manage several fully patched XP and Windows 2000 VMware clients. We used a variety of applications, such as Skype, Kazaa and µTorrent, to test executable blocking.

HIPS

CA Host-Based Intrusion Prevention System

REVIEWED BY BRAD CAUSEY

CA

www.ca.com

Price: **$40 per client with enterprise-level maintenance**

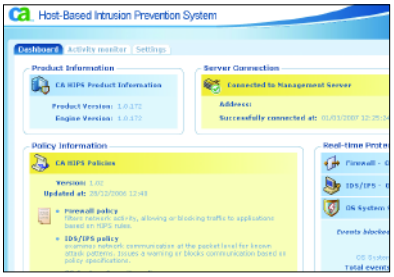

CA Host-Based Intrusion Prevention System (CA HIPS) combines standalone firewall and intrusion detection and prevention technologies to provide security, access control, policy enforcement, intrusion prevention management and deployment from a central console.

Configuration/Management A

Management server setup was simple and fast. The server will automatically enumerate all LDAP users and groups, ultimately providing the ability to control policies and rule sets by Active Directory group membership. The management server is configured by default to check for LDAP changes every 180 seconds.

The real beauty of CA HIPS is its Learning Mode, which allows you to monitor a group of systems to determine what constitutes acceptable behavior. This will help avoid the problem of a policy deployment that is either too restrictive, causing application/network outages, or too liberal, limiting the effectiveness of the product.

Client deployment is a cinch for desktops and servers. Log in to the intuitive, Web-based console and configure the client with the options you desire. Click "Build" and you are presented with an installation package.

Testing methodology: We assigned policies to users and computer groups in a single Active Directory domain. Attacks were launched on the clients in the form of viruses, worms, remote exploits (some with no known patches) and spyware.

Policy Control A

The intuitive policy creation and deployment workflow simplifies what could be a complicated process.

You begin by creating "Common Objects"—basically, the targets of security policies, such as USB drives, registry entries or network protocols. There are thousands of default objects, which can be easily customized. Next, you define the list of rules associated to the objects. Rules are broken down into categories such as application, firewall, operating system and IDS/IPS.

Customizing or creating rules is simple, and they can be grouped into manageable collections such as high-security laptops and DMZ Web servers.

Policies allow you to apply groups of rules to subnets, hosts, users, groups and a number of other criteria. A simple deployment wizard pushes the policy to clients.

Effectiveness B

CA HIPS provides effective defense against threats, known and unknown. Clients are heavily protected, but in a way that will not adversely affect the functionality of the system. By sandboxing an application or questionable device, it protects against unknown threats.

We executed a number of viruses, Trojans and other malware, many of which didn't have signatures. Software that wasn't approved was sandboxed based on the policy, and escalated for approval/restriction. Known malware was restricted from accessing the OS and network.

You'll still need an antivirus product to remove viruses and Trojans. CA's antivirus product works with CA HIPS to remove malicious code, but the products aren't completely integrated yet.

Reporting D

Reporting seems immature. There are around 100 prebuilt reports that offer a plethora of information, but there's no method to modify them or create custom reports from inside the Web interface. The only criteria that can be supplied to filter the results are based on time frames, such as the last week or month.

CA does provide an API to create custom reports, but that shouldn't be necessary to create a report from data that is stored on the management server.

Verdict

CA HIPS offers comprehensive threat defense in a flexible and easy-to-use product. Its reporting weaknesses may rule it out for some organizations.

DEVICE MANAGEMENT

HotPick INFORMATION SECURITY®

Pointsec Protector

REVIEWED BY SANDRA KAY MILLER

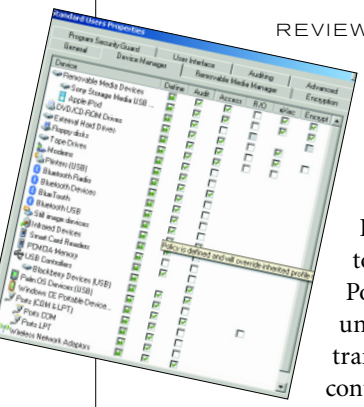

Check Point Software Technologies
www.checkpoint.com
Price: **Starts at $45 per seat**

Pointsec Protector (formerly Device Protector prior to Check Point's acquisition of Pointsec) addresses the growing problem of unsecured ports and endpoint devices while transparently delivering encryption, filtering content, enforcing policies and maintaining an audit trail, even when mobile devices are disconnected from the network.

Configuration/Management A

Getting the product running was effortless, thanks to well-designed wizards and a straightforward installation process. Protector ties in with Microsoft Active Directory and Novell eDirectory for user and group synchronization when assigning device access rights, encryption and policies.

The administration console is intuitive, and multiple tiers of administrative access can be assigned for distributed management. We were able to easily manage users, groups and devices, policies, alerts and encryption, and create and view audits, logs and reports.

Policy Control A

We began by editing Protector's default profile through a series of tabs to choose what types of devices and removable media to permit/deny access, define encryption, create email alerts, and assign stringent policies for groups that fell under compliance regulations (e.g., finance) and less stringent for others.

Policies are layered, so the default policy is applied to every group to which it is assigned. When another policy profile is created, it can inherit from the default policy or become a new profile. For example, in the default profile we globally banned iPods and enabled encryption on all USB storage devices. The next policy, while it inherits the default profile, may define access to approved devices, such as portable hard drives, on which encryption from the default policy will be enforced.

Policies can be assigned on a user, group or device basis. Administrators can restrict the types of files that may be transferred or the launching of unauthorized applications from removable media.

Protector uses combinations of whitelists and blacklists to block access to devices and files without any legitimate business purpose, while still allowing users access to critical tools, applications and data defined by brand, model and file type.

Reporting A

Protector excels in logging and reporting. With detailed auditing, administrators can determine what devices are being used and in what way. Alerts are easily set up to be sent via email; we assigned each AD group a different notification recipient simulating department managers being alerted to their employees doing such things as downloading music at work or copying sensitive files to portable media.

Logs can be customized, filtered by column heading and exported to CSV. Reports are equally flexible and can be exported in HTML.

Effectiveness B

Protector enforces all policies and offers a high level of control and auditing over offline devices. Even with local admin rights, we were prevented from disabling or uninstalling the client software from our test laptop thanks to anti-tampering controls.

The encryption feature works transparently when the user is logged on to the network. For offline machines and mobile devices, users simply drag and drop files on or off of the encrypted device through a password or challenge/response.

Protector lacks centralized control for Linux and Mac, and doesn't have data shadowing, meaning administrators could record all information sent to a particular device or port.

Verdict

Pointsec Protector is an affordable and scalable solution that will work well in both SMB and enterprise environments.›

Testing methodology: The testing environment included Windows clients, AD and SQL Server. We tested the use of portable storage devices, including USB flash drives, FireWire external hard drives, CD-RW drives and floppy disk drives.

ENDPOINT SECURITY

Blink Professional 3.0

REVIEWED BY STEVEN WEIL

eEye Digital Security

www.eeye.com

Price: **$59 per computer per year**

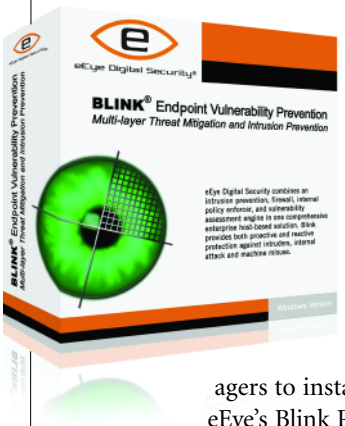

Antivirus software is no longer enough to protect your company's computers. Prolific laptops, desktops and critical servers are facing threats from many fronts: malware; insecure protocols and applications; lost, stolen or misused portable storage devices; and network traffic. Host IPS, antivirus and storage device control programs can mitigate certain threats, but force security managers to install and manage multiple applications.

eEye's Blink Professional 3.0 is among the increasing number of host-based endpoint security products that use a layered, consolidated approach to defend Windows computers against different attacks.

Configuration/Management B

Following eEye's well-written documentation, we were able to quickly and easily install Blink. Blink's interface is intuitive and easy to use; we were able to effectively navigate among the many local settings.

We liked the well-designed wizard programs that are used to create rules and signatures. We also liked being able to add references, such as CVE and Bugtraq IDs, to IPS rules. Blink can be configured to automatically check for software and signature updates.

Policy Control B

Blink deploys a single agent and common management of its multiple security capabilities: a host firewall, monitoring inbound and outbound traffic; an application firewall that controls the network activity of installed applications; signature- and protocol analysis-based host intrusion prevention; antimalware protection against worms, viruses and spyware; antiphishing capabilities; system protection against buffer overflows; and controls over which applications may access the registry and/or be launched.

Blink can also block the use of storage devices, such as USB flash drives, and conduct local vulnerability assessment scans.

Security managers can configure Blink locally, or configure it to regularly check and download a centralized policy. Blink can also be integrated with eEye's REM Security Management Console for creation and management of dynamic policies. It also centralizes logging.

We were able to create numerous granular firewall rules, IPS signatures and system protection rules, which defined the actions to be taken (allow, log, block, alert).

Effectiveness A

Blink did an excellent job of protecting host computers. We ran numerous manual and automated attacks against our test computers; our attacks included sending malicious data and executing unexpected protocol actions. Blink always took the correct action, such as blocking or logging attacks. Permitted traffic was correctly allowed. Blink also correctly blocked the use of prohibited storage devices and detected malware we installed on the test computers.

Reporting B

Integration with REM enhances Blink's reporting abilities. In standalone mode, Blink locally logs system, firewall and IPS events, and can send SNMP traps. Individual log events are easy to understand, but the logs can only be exported as a .csv file.

Via Blink's local event log interface, an administrator can select an individual log event and, as appropriate, block an IP address, go to the rule that logged the event or create a new rule in response to an event, such as allowing traffic that was blocked. Administrators can configure Blink to pop up a user alert when a specific event occurs, such as an RDP connection to a server.

Blink also generates useful reports after an antimalware and/or vulnerability assessment scan is run, but they cannot be exported.

Verdict

Blink is well-designed, and its multilayered approach makes it a good choice for protecting Windows computers throughout an organization.

Testing methodology: Our test network included a Windows 2003 laptop, an unmanaged switch and three Windows 2003 servers.

ANTIVIRUS

NOD32 Antivirus 2.7

REVIEWED BY MIKE CHAPPLE

ESET

www.eset.com

Price: **Starts at $39 per workstation for a one-year license**

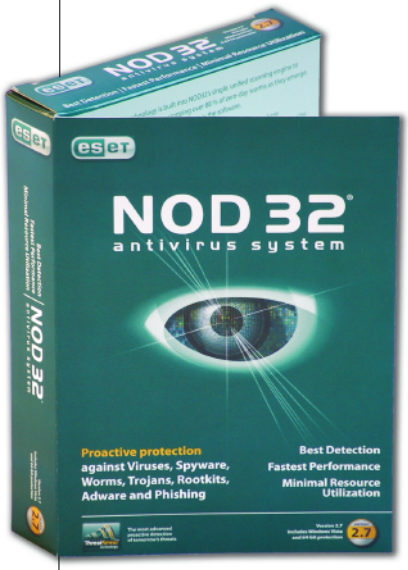

NOD32 Antivirus provides effective virus and spyware protection, albeit without many of the bells and whistles found in competing products.

Our review supports ESET's claim that NOD32's lightweight design allows it to scan systems more efficiently than the competition. While it might appeal to administrators seeking to scratch out every last cycle of CPU performance, it's not a product for those seeking a highly integrated enterprise solution.

Configuration/Management B

Client installation uses standard deployment techniques—push the package onto client systems, create a login script or distribute installation packages through traditional means. The central administration capabilities require installing the remote administration server, and the remote administration console on administrator systems. We used the wizard-guided installation processes to install all these packages without difficulty.

Policy Control C

NOD32 allows administrators to apply policy to individual systems or groups of systems through the remote administrator feature. Policy control is quite granular: you can specify scan types and frequency, excluded files, etc. You can apply different policies on a group basis.

However, the policy control GUI is very awkward. To create or modify a policy, you must open a separate policy editor and then save the resulting policy as an XML file. You can then return to a separate window to apply the XML file to systems/groups.

NOD32's Remote Administrator package provides limited integration with Active Directory. You are able to synchronize NOD32's internal groups with an Active Directory server, but you can't manage NOD32 within the native AD environment. Additionally, you can only import machine information, not user information. This highlights one of the package's significant shortcomings: NOD32 does not allow you to create individual administrator accounts, so admins must use a shared server password—hardly a best security practice.

Effectiveness A-

NOD32 delivers on its promise. It successfully detected the test viruses we placed on the evaluation system and alerted the administrator to their presence.

NOD32 met its efficiency claims. We ran it against a leading product in a head-to-head test on identical environments. Both packages detected our test viruses, but NOD32 completed its scans 16 percent faster.

An active monitor watches for file system changes, while the on-demand scanner performs full system scans. The Internet monitor scans HTTP and POP3 traffic. While NOD32 does not support scanning of IMAP connections, an add-on supports Microsoft Exchange. Another module scans Microsoft Office documents.

NOD32 doesn't include a centrally managed host firewall service or a host-based IPS. These shortcomings may limit the product's viability for some organizations.

Reporting A

NOD32's smorgasbord of reports will please even the pickiest data hound. Predefined templates include reports on top viruses, top clients with most alerts and alerts by module, among others. Administrators can filter reports based upon specific clients, servers or viruses and the HTML-based reports can be exported to a .csv file for use in other analysis packages.

Verdict

NOD32 is an efficient, effective antivirus solution with decent central administration tools. However, it lacks strong enterprise-level features.

Testing methodology: We tested NOD32 in a Windows XP virtual machine environment with separate machines serving as the client and administrator. We performed the competitor comparison on an identically configured virtual machine.

ENDPOINT SECURITY

F-Secure Client Security 7.0

REVIEWED BY BRENT HUSTON

F-Secure

www.f-secure.com

Price: **$29.75 per user for 100 licenses**

In the face of criminal zero-day exploits, targeted attacks and mobile workforces, companies are increasingly seeking comprehensive endpoint security tools. F-Secure Client Security, a business-grade, centrally managed suite, addresses emerging threats, as Version 7 adds rootkit scanning and host-based intrusion prevention (HIPS).

Installation & Features B+

Both the Policy Manager and client software install in a snap. By default, the antivirus scanner does not look inside archives, or force a complete system scan upon installation, saving system resources and time. It removes other antivirus or firewall products, although it left Comodo Firewall on one of our test systems. Client Security suite consumed about 80 MB of memory on our test systems—a little high, but not excessive.

In addition to new rootkit scanning and HIPS capabilities, Client Security suite features an antivirus and antispyware scanner, mail and Web proxies, and client firewall. The proxies act as a middleman, inspecting email and data passed via Web browsers. This allows Client Security to remove malicious software before it has completely reached the system. F-Secure's excellent antivirus engine is known for its speed and efficiency; it displays the same competence in this suite.

Testing methodology: The management server was run on Windows Server 2003, the Client Security Suite on Windows XP SP2 clients. We tested against common viruses and malware, as well as newer malicious software.

Management B

The Java-based Policy Manager console interface is clean, with tabbed windows allowing easy navigation. The console allows you to install F-Secure remotely, as well as manage configuration, updating and monitoring. The policy management tool auto-detects Windows machines, facilitating installation.

The console allows admins to centrally control all installed systems. Pressing one button allows you to scan every system for rootkits with the new BlackLight feature (rootkit and spyware detection), and remotely update all clients with the latest virus definitions.

Effectiveness A

Client Security is unobtrusive during normal desktop operations. It caught everything we threw at it: common viruses and rootkits, as well as hacking tools planted on test systems. We browsed Web sites that install malicious software to test the HTTP streaming scanner, which stopped all threats before they could be downloaded. Client Security also performed well in cleaning up already infected systems.

The new HIPS component, DeepGuard, which is designed to protect against new threats, blocked malware from installing when the antivirus scanner was disabled. It also stopped software downloads most organizations block by policy, such as some browser toolbars. DeepGuard can also be configured to prompt anytime something tries to change the registry, instead of relying on its AI to detect if the change is malicious or not.

Reporting B-

Client Security reports are simple HTML pages. The Policy Manager console reports are a bit more extensive, allowing current status, trend and detailed list reports. Policy Manager stores data in a lightweight SQL database, but it too can only generate reports in HTML. However, the reports are well designed, and included an easily printable version. The bookmarking feature generates a new report with the latest data for the report type you specified every time you select a bookmark.

Verdict

Client Security is fast, efficient and reasonably priced. The Policy Manager is free, a definite value-add.

MALWARE ANALYSIS

Norman SandBox Analyzer Pro

REVIEWED BY TOM LISTON

Norman

www.norman.com

Price: **Starts at $5,000 for 100 users**

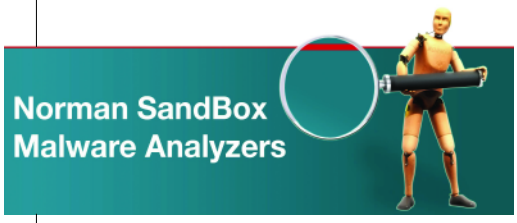

Relying solely on antivirus to protect you from malware is no longer an option. Antivirus software is reactive; vendors only release signatures for malware they've seen. With the growing prevalence of more targeted viruses, the bigger your company, the more likely you are to be hit by something that no one, not even an antivirus vendor, has seen before. In response, many companies are developing in-house malware analysis capabilities.

Norman SandBox Analyzer Pro is a unique malware analysis tool that allows potentially malicious code to execute within a simulated environment that effectively mimics a generic Windows installation. All actions taken by the code under analysis are monitored. Any permanent changes that the test code attempts to make are trapped by the sandbox (files don't get written to the file system, keys don't get changed in the registry) but everything appears normal from the point of view of the code under test.

Analysis Tools A-

Analyzer Pro provides analysts with an almost overwhelming amount of information about the inner workings of the code under test. From the files it attempts to create, to the registry entries it adds or changes, to the network connections it attempts to make, Analyzer Pro sees and logs all.

One incredibly useful feature is the ability to allow mediated access to the Internet using powerful filtering tools. Access to the Internet can be controlled in many ways—remote connections can be "faked" by Analyzer Pro, access to the real Internet can be allowed, or the analyst can alter packets being sent or received from the Internet on-the-fly.

Recent malware often has a networking component that can only be fully investigated using this feature. For example, the behavioral aspects of a bot program can be fully understood if it is allowed to contact its command and control server.

Usability B-

Analyzer Pro is a powerful tool for combining code-level analysis with extensive behavioral monitoring and logging, but it has a steep learning curve. The main analysis tool is a specialized debugger that allows the analyst full control over the execution of the program at a granular level.

This is not a tool for neophytes. Even with years of experience using debuggers and code analysis tools, we found Analyzer Pro to be very confusing at times. We had to analyze several dozen pieces of code before we felt reasonably comfortable with the tool's quirks.

If your organization is looking to start analyzing malicious code, we would suggest staying away from Analyzer Pro until you hire experienced malware analysts or develop internal expertise.

Documentation D

Perhaps the greatest problem is documentation. Analyzer Pro was obviously originally developed by Norman as an internal analysis tool, and that heritage is evident in its documentation. It is poorly written, confusing and assumes a level of expertise that makes Analyzer Pro unsuitable for anyone but a seasoned malware analyst.

Verdict

Although it lacks polish in its user interface and its documentation, SandBox Analyzer Pro's powerful and flexible feature set makes it a desirable tool for seasoned malware analysts. Beginners will find it frustrating and confusing, but mature code analysts will find it a welcome addition to their toolkit.

Testing methodology: Analyzer Pro was tested on a Windows XP Professional machine with a 1.8 GHz processor and 1MB of RAM. Testing was done by analyzing a variety of sample code (from the reviewer's malware "zoo") using the tools provided. Tests were performed using known benign code and previously analyzed malware samples.

ENDPOINT SECURITY

ZENworks Endpoint Security Management 3.5

REVIWED BY SANDRA KAY MILLER

Novell

www.novell.com

Price: **$69 per device**

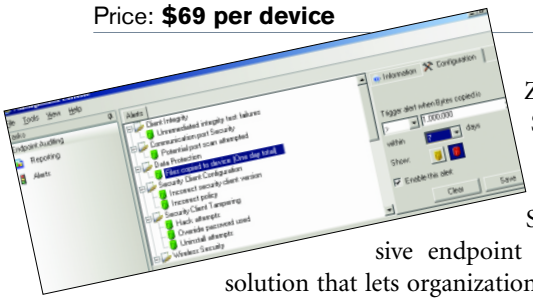

ZENworks Endpoint Security Management 3.5 (formerly Senforce Endpoint Security Suite) is a comprehensive endpoint security management solution that lets organizations control applications, protocols and removable storage devices. It delivers encryption to files and folders, and network access control to ensure protection levels are current.

Configuration/Management B-

We installed ESM's three components—Policy Distribution Service, Management Server and Management Console—on a server. Larger deployments require that the Management Server and Policy Distribution Service be installed on separate servers.

Installations were straightforward; the client required us to choose between obtaining policy updates through ESM or files. The Policy Distribution Service checks what is sent out against the Management Server, which interfaces with directory services. Password protection for the client prevents removal and tampering.

Setting up the server required extensive networking, security and SQL knowledge.

Multiple installs on secured machines connected to the server are possible, but a Web-based interface would make configuration and management easier.

The console allows navigation through the taskbar and expandable submenus, but we'd prefer to see items like reporting and alerts accessible through a click.

Testing methodology: The single-server installation was deployed on a Windows-based network behind a firewall. Clients were installed on a variety of endpoint systems located within and outside of the firewall. Policies were enabled for a variety of scenarios, including remote and mobile endpoints.

Review how we grade at searchsecurity.com/grading_criteria.

Policy Control A

ESM earned top marks for the granularity and scope of security and control policies.

There are extensive policy options for wireless and wired networks, communications hardware, firewall settings for multiple locations, antivirus/spyware and Microsoft patches. You can use advanced scripting rules for customized rule sets and set features and alerts specific to regulatory actions.

Policies are distributed via SSL through a Web services application, pulling users and groups from directory services. Policies are easily edited and instantly updated.

Reporting B

Alerting and reporting are tough to locate, hidden in the Management Console's Tasks menu.

Alert thresholds are adjustable. For example, we enabled an alert if data in excess of 5 GB is copied to removable storage media or device.

ESM offers 10 reporting categories—adherence, alert drill-down, endpoint activity, encryption solution, client self-defense, integrity enforcement, outbound content compliance, administrative overrides, endpoint updates and wireless enforcement. If you want to create custom reports, however, you'll have to use an ODBC-compliant app such as Crystal Reports.

Effectiveness B

While ESM provides a multitude of security and control features in a single suite, there are a number of features available in similar products we would like to have seen. For instance, we could not assign storage device control policies when encryption for the particular device was required. Also, an additional USB Drive Scanner Tool had to be installed separately from the initial installation to be able to scan and identify devices attached to USB ports.

Despite those shortcomings, policies were automatically distributed to clients. ESM blocked noncompliant clients that were assigned specific requirements (such as up-to-date antivirus signatures).

Verdict

ZENworks Endpoint Security Management is a comprehensive solution for managing and enforcing security policies on networked devices, regardless of their location and connectivity.

ENDPOINT SECURITY

Desktop Authority 7.5

REVIEWED BY HARRIS WEISMAN

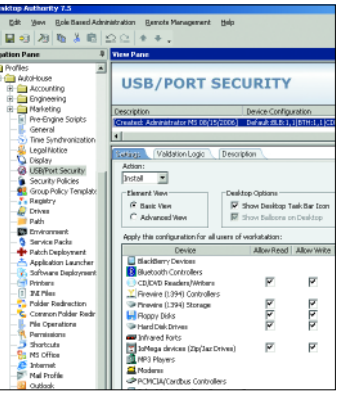

ScriptLogic
www.scriptlogic.com
Price: **$30-$40 per seat**

As the name suggests, ScriptLogic's Desktop Authority 7.5 is a centralized desktop management solution. What makes this product unique is the optional add-ons that address patch management, spyware removal and USB/port security. The total package makes this a useful tool for day-to-day desktop and security management.

Configuration/Management B+

Server installation was straightforward: The software recognized that the base Windows 2000 software was missing several core components and added them.

When a user logs on to the network, Desktop Authority executes a script, SLogic, that installs the client software and configures the registry. This script is assigned to users through User Manager for Domains or Microsoft Active Directory, or from within Desktop Authority itself.

Through a dedicated central console, admins can manage the configuration of the user's environment, maintain a complete inventory of the enterprise's hardware and software, and support users remotely.

Granular profiles allow for group and individual user management for various environments. As the user logs on to the network—either on the network or by connecting remotely—the script activates and applies the group's rules to the workstation. This allows admins to maintain consistent and updated policy controls.

Desktop Authority bundles several optional security components, including patch management via Shavlik Technologies and antispyware from Aluria Software (antispyware deployment relies on the installation of the Shavlik module). ScriptLogic also offers its USB/port control tool, which secures data against theft or accidental loss.

Testing methodology: Our test platforms included Windows 2000 Server, Windows XP Home and Windows XP Professional.

Policy Control B+

Desktop Authority makes it easy to administer several key security functions. Patches can be tailored to each individual group's needs and deployed accordingly. This also works for the antispyware component. One nice feature is the ability to exempt specific "spyware" applications, which is how many security tools are classified.

The USB/port control component can manage use of Bluetooth, CD/DVD drives, FireWire devices, floppy disks, hard disks, infrared ports, Zip disks, MP3 players, modems, PCMIA controllers, PDAs, USB devices, parallel/serial ports and Wi-Fi devices.

Administrators have granular control of what can be accessed; for example, you can grant read-only access to CD-ROM burners, allowing users to access portable disk content but not copy sensitive data.

Effectiveness B+

In our testing, we focused on the security modules, developing a series of groups that denied and granted access to specific devices and media. The system prevented us from violating the assigned access control lists. Using these same groups, we also deployed patches to Windows XP desktops and laptops. We deployed the antispyware component using the patch management tool.

Reporting A

The reporting tool is very effective, allowing the administrator to use either packaged or custom reports. All three security components use Desktop Authority's reporting tool, which allowed us to generate canned reports on antispyware, USB/port and patch activity, as well as administrative desktop management. Each of these areas features several detailed options; for example, under USB/port management, you can generate reports by a specific type of device or workstation. Each canned report can be modified to suit the enterprise's needs.

Verdict

Desktop Authority is an effective desktop management/security tool. While it does not work with non-Windows systems, it would be a cost-effective solution in a Windows shop.

VIRTUALIZATION

BufferZone Enterprise

REVIEWED BY GREG BALAZE

TrustWare

www.trustware.com

Price: **Starts at $2,899 per 100 licenses**

With the information security field seemingly saturated with every possible appliance and software, it would seem there's little room for an innovative approach. TrustWare's BufferZone belies that notion.

TrustWare's BufferZone works by quarantining suspect or restricted applications, creating a protected environment for each Web- or network-based application, such as Web browsers, IM, email and P2P applications, preventing viruses or malware from entering and affecting the rest of the workstation

Configuration/Management B

Instead of designing a management console, BufferZone relies on Microsoft's native Group Policy Objects to manage and deploy BufferZone and its installation file. This allows easy integration with Active Directory and reduces the learning curve, through a group policy template that uses the familiar management console (MMC). Simply copy the administration file to the c:\windows\inf\ directory and add it to the GPO administration templates. Deployment was just as easy, using GPO and software installation packages by copying the .MSI file to a shared drive from which it's deployed to workstations as users log in.

This should be fine for most organizations, though some may prefer products with more robust proprietary consoles.

Testing methodology: BufferZone was run on two workstations, Windows XP SP2 and Windows 2000 SP4, that were in a standalone AD domain.

Policy Control B

Depending on the policy set, you can easily prevent P2P applications, such as Kazaa, from being launched by a user, or allow ActiveX applications to be run only in the BufferZone-protected area. You can add applications or files by simply typing the name of the executable or DLL in the applicable dialog box. Any application or attachment that is launched by a Web navigator, P2P application, IM or mailer is quarantined by default.

You can choose from four policy settings under which files can be run: BufferZone, so files run only in the protected area (this prevents a file from affecting other areas of the workstation hard drive or memory space, or removable media such as a CD, flash drive or MP3 player); Forbidden, in which users have no access to the files; Confidential, which means any file or path matching the policy is invisible to applications run in BufferZone; or Trusted.

In testing, we found that using both file path and wild cards was best for policy enforcement (for example, *\MY DOCUMENTS*.doc and *torrent.exe). BufferZone includes a switch to allow digitally signed executables to run outside of the controls set for a certain media or file.

Effectiveness B

Although BufferZone was excellent at stopping potentially malicious executables and preventing CDs or USB devices from being accessed, it was obvious that it was only for files already known. A new policy would have to be created each time an unknown application or executable was installed. Fortunately, in an enterprise environment where typically standardized applications are installed, this shouldn't be too much of an issue.

The lack of any reporting capabilities may give pause to some enterprises.

Verdict

BufferZone does a remarkable job controlling which files can be run or downloaded on a workstation. The typical entry points, such as Web browsers, can be locked down, preventing unwanted access.

ENDPOINT SECURITY

Yoggie Gatekeeper Pro

REVIEWED BY PETER GIANNACOPOULOS

Yoggie

www.yoggie.com

Price: **$220**

Yoggie Gatekeeper Pro is an interesting new take on the perennial problem of using personal firewalls to secure individual PCs. It is a Linux-based USB device that serves as a full-blown firewall, VPN gateway and integrated antivirus/Web filter gateway. All of this capability is packed into a form factor that is approximately the size of a pager.

We tested the standalone version to look at the core capabilities of this unique product, but an enterprise version is available that provides centralized management, policy control and reporting.

Policy Control B

The policy control is rather basic, but effective. The administrator can open any necessary firewall ports, configure IP addresses and perform other tasks. Some rather pleasant surprises included a fairly effective Web filter control that can block access to malicious Web sites or those that violate corporate policies, such as adult entertainment and gambling sites. It also offers integrated antivirus and automatic updating of the device firmware and virus definitions.

These features are excellent capabilities bundled into the basic product, but we would have liked more fine-grained control.

Testing methodology: We installed Yoggie on a Windows XP SP2 test machine. We used the EICAR test suite, along with samples of known viruses, to validate the antivirus capabilities, and used Nmap to verify open and closed ports on the host.

Effectiveness A

We ran several Nessus and Nmap scans as well as some custom penetration testing tools against the Yoggie-protected test system to see if we could either take out the device, or tunnel a way through it. The Gatekeeper Pro passed with flying colors and shrugged off the attacks, while allowing the test system continued access to the Internet.

We enabled Web filtering and attempted to hit some dubious Web sites. We were denied access, and all the attempts were logged. Gatekeeper Pro's antivirus likewise caught the EICAR test suite—as expected—as well as the sample viruses we had in the lab. Gatekeeper Pro worked as advertised.

Configuration/Management B+

Installation is gloriously simple: Plug in the device to the USB port of your PC or laptop, connect your Ethernet port to the clearly labeled PC port on the device using the bundled cable, connect your broadband connection to the Net port, and you're ready to surf the Web. If you plan to use the Gatekeeper Pro with a WLAN connection, you'll need to install a driver.

Configuring the device is accomplished through a built-in Web GUI that will immediately look familiar to anyone who has set up any type of home networking equipment from D-Link or Linksys.

Though effective, the Web GUI is somewhat Spartan. For example, the user is allowed to choose from low, medium or high security levels, but there is no clear explanation of the specific differences between the levels, either in the online help or in the product documentation. We set the security level to high and were able to perform common tasks such as instant messaging, Web browsing and email without any difficulties.

Verdict

We see a definite sweet spot in the market for Yoggie's offerings, especially the enterprise version in large environments that want to protect disparate endpoints without getting bogged down with configuring individual hosts. ›

COLOR ME COMPLEX

REVIEWED BY ED SKOUDIS & MATT CARPENTER

Image by JOHN KUCZALA

TODAY'S DESKTOP AND LAPTOP COMPUTERS ARE a complex melange of software: office applications, specialized business programs, homegrown client apps, and ever more complex browsers and operating systems. As an enterprise security pro, you're lucky if you can get at least two agent-based security tools included in this zoo of your standard builds to protect all of that software. Some of us get only one. Yet desktop security technology is rapidly advancing, with host-based intrusion prevention systems (HIPS), personal firewalls and other defenses augmenting traditional antivirus and antispyware tools. Because of the severe constraints on the number of host security products our enterprises can deploy and manage, major security vendors have responded with integrated endpoint security suites, rolling a bunch of desktop defenses into a single package *(see "End-to-End Endpoint Security," p. 24)*. These endpoint security products have introduced a new dynamism into our industry, as antivirus vendors augment their wares with fresh features to compete against each other and hungry challengers. To help sort out all of this, *Information Security* evaluated seven enterprise endpoint security solutions. We graded each on its management capabilities, reporting, ability to detect and block malware, detecting and thwarting exploit attempts, and integration of the various desktop security capabilities in one package.

Funneling integrated endpoint security features into one product makes for a murky mix of complexity and immaturity.

Specifically, we tested CA Threat Manager 8.1 and Host-Based Intrusion Prevention System 8; eEye Digital Security Blink Enterprise Edition; IBM ISS Proventia Desktop Endpoint Security 9.0; McAfee Total Protection for Enterprise; Sophos Endpoint Security and Control 7.0; Symantec Endpoint Protection 11.0; and Trend Micro OfficeScan 8.0.

Bearing witness to the rapidly evolving nature of the endpoint security space, the three giants of the information security industry—McAfee, Symantec and Trend Micro—responded with beta versions of their suites that were nearly finalized for shipping. (We requested every product we analyzed be available for general purchase by our publication date.)

Many of the problems we encountered with testing and, in some cases, retesting updated versions of these products reflected the difficulties in dealing with beta builds of highly complex packages. But, further, our testing suggests this class of integrated endpoint security products is, for the most part, far from mature.

MANAGEMENT

The immense complexity of these tools can be overwhelming, with more features than almost any distributed system in today's enterprise. If a given product provides really good security, but cannot be managed across an enterprise in a coherent fashion, it just isn't useful.

The typical endpoint security suite must parse hundreds of types of files for antivirus and antispyware functionality, analyze numerous packet types for firewall and some IPS activities, police hundreds of operating system settings and running processes to detect malicious action, and much more. All of these features are controlled with thousands of settings applied via default policy templates or custom policies for specific enterprise needs. Admins roll out these policies to a multitude of managed workstations and servers across the network, all to be managed in real time.

THE IMMENSE COMPLEXITY OF THESE TOOLS CAN BE OVERWHELMING, WITH MORE FEATURES THAN ALMOST ANY DISTRIBUTED SYSTEM IN TODAY'S ENTERPRISE.

We looked at the ease with which each product could be used to configure systems, quickly determine their security state, and update settings based on attack activity.

We were particularly impressed with Symantec's management capabilities, but McAfee's completely new ePolicy Orchestrator (ePO) is a major disappointment.

Symantec is top-notch from a large enterprise perspective, with intuitive GUIs for policy configuration and status checking. Its overall dashboard clearly identifies potential problems associated with infection, out-of-date signatures or disabled functionality on managed hosts, offering advice to an administrator on how to fix each issue. The management GUI comes in two flavors: a full-blown Java-based interface for all aspects of the administration console, and a scaled-down Web-based interface that can be used for status checks and reporting, but not policy management.

Sophos also provided very solid management capabilities, organized, the vendor told us, around the KISS principle, which we assume stands for "Keep It Simple, Sysadmin." Sophos' GUI is designed to reduce the time and effort needed to configure and deploy the product. Sure, you don't have access to a lot of the fine-grained policy settings, but the overall options available for configuration

CA Threat Manager 8.1, Host-Based Intrusion Prevention System 8

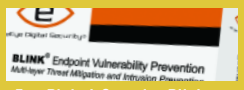

eEye Digital Security Blink Enterprise Edition

IBM ISS Proventia Desktop Endpoint Security 9.0

McAfee Total Protection for Enterprise

Sophos Endpoint Security and Control 7.0


Symantec Endpoint Protection 11.0

Trend Micro OfficeScan 8.0

ABOUT THIS REVIEW

We tested solutions that provided a minimum of signature-based antivirus and antispyware, personal firewall, host-based IPS, and central management and reporting capabilities. We selected a mix of leading traditional antivirus vendors and relative newcomers in the market, testing comprehensive endpoint security products from CA, eEye Digital Security, IBM ISS, McAfee, Sophos, Symantec and Trend Micro. (Because of space and resource limits, we opted not to include products from four other companies that responded positively to our invitation to apply for consideration for testing: Fortinet, F-Secure, Grisoft and Kaspersky Lab. In addition, ESET, Norman Data Defense Systems and Third Brigade declined our invitation to be considered.)

Our analysis test bed consisted of a Windows 2000 Server with Service Pack 4 running the enterprise management software, managing and protecting four Windows XP client systems. Each XP client had no service pack or patches. As a control, for each test, we utilized one target Windows XP machine with exactly the same configuration but lacking the endpoint security protection.

—ED SKOUDIS & MATT CARPENTER

are excellent. Checking the status of managed workstations is snappy, and alerts about systems that deviate from policy are easy to understand.

The Trend Micro management interface worked quite well in configuring and analyzing managed systems, especially for antivirus and antispyware. The new product architecture enables Trend Micro to package new endpoint product building blocks into plug-ins for rapid deployment, a design decision that will benefit Trend and its customers.

However, we were concerned that we couldn't discern client signature updates for its Intrusion Defense Firewall (the component that implements the firewall and HIPS functionality). Such information is vital in signature-based IPS products such as this one, which applies network-based IPS signatures to traffic going into the protected host. Trend Micro licensed this functionality from Third Brigade to create the first plug-in for its new architecture.

eEye Digital Security features a well-organized, intuitive management interface. However, the client GUI is clearly more mature than the enterprise management console itself, offering finer-grained insight into the configuration and alerts generated by the tool.

CA's management console has improved significantly since we last looked at it in our antispyware analysis in May 2006. Its latest version is much faster and more interactive than previous versions. Still, checking the status of different workstations required moving between different screens, and policy configuration of this purely Web-based GUI was more difficult than with other products.

We found the IBM ISS product quite difficult to manage. Determining the current status of clients from the management console was cumbersome, and managing all of the separate features was complicated and confusing. Also, at several points we encountered cryptic error messages that didn't explain the problems we encountered in installing and configuring the product. Finally, the IBM ISS endpoint product is exclusively for Windows clients; it cannot be used to manage servers, Windows or otherwise. Server security is available only as a completely separate product.

McAfee's new enterprise management server, ePolicy Orchestrator (ePO) 4.0, was a great disappointment. The more you loved the previous versions of McAfee's ePO, the more frustrated you will likely be with the new version.

McAfee has completely rewritten its flagship management product with a Web-based GUI, letting admins manage it from any browser in their enterprise.

ENDPOINTS | Management

THE GOOD NEWS Symantec is top-notch from a large enterprise perspective, offering easy-to-understand GUIs for policy configuration and status checking.

THE BAD NEWS The more you loved the previous versions of McAfee's ePO, the more frustrated you will likely become with the new version.

The well-laid-out and quick GUI of earlier ePO versions has been replaced with a complex and bewildering Web-based interface. The easy drag-and-drop features of the thick-client ePO have been replaced with countless Web-based drop-down menus in screens that make it difficult to find what you need.

The difficult-to-use management GUI and default policy are of significant concern. While testing, we accidentally applied a baseline medium security policy to the McAfee management server itself. In the complex ePO 4.0 GUI, such mistakes are frustratingly easy to make—we did it while being guided step-by-step on the phone by McAfee

DETECTION

ANTIMALWARE SCANNING RESULTS

Chart shows the percentage (lower percentages are better) of 8,114 recent malware specimens that survived on protected systems after (1) real-time scans of malware as we attempted to copy them to the target system; (2) on-demand scans of the malware surviving the real-time scan; and (3) on-demand scan of all specimens with real-time scanning disabled.

	CA	eEye Digital Security	IBM ISS	McAfee	Sophos	Symantec	Trend Micro
1 Real-time scan	8.7%	9.3%	100%***	22.3%	100%**	17.6%	8.0%
2 On-demand scan following real-time scan	8.7%	9.3%	58.6%	22.3%*	36.7%	8.3%	7.8%
3 On-demand scan of all malware specimens	8.7%	9.3%	34.7%****	100%*	36.7%	8.3%	7.9%

*McAfee's default action is to alert only during an on-demand scan. Changing this default, 11.6% of malware remained after both real-time/on-demand and pure on-demand scans.
**Sophos real-time scans block only read-actions by default, not write-actions.
***IBM ISS does not block network-writes across Windows shares, even when configured to scan for all write-actions.
****IBM ISS only scored this number after repeated runs from scans that ended prematurely with a messages saying "Successfully Completed"

support. By simply applying the default security policy to the management server, ePO killed itself. We were unable to get access to any of the management capabilities, and had to reinstall everything from scratch to resume testing.

ANTIMALWARE SCANNING

To gauge each vendor's ability to detect and block malware found in the wild, we ran three tests using 8,114 recent malware specimens from a private collection graciously provided by antispam researcher Bill Stearns. Our zoo included a large variety of worms, bots, backdoors and viruses. For each test, we recorded the percentage of specimens not eradicated in each round of testing *(See "Antimalware Scanning Results," p. 22).*

Our first test was designed to evaluate each product's real-time signature-based defenses by copying the malware from a hardened machine to a shared directory on the protected target system. We then recorded the percentage of malware specimens that made it into the target's file system, escaping detection by the product's real-time scanning capabilities.

We then performed an on-demand scan of all malware that survived our first test, to assess the combined real-time and on-demand scan capabilities for identifying and eradicating malware.

Finally, we conducted on-demand scanning independently by disabling real-time scanning, copying all malware to the target file system, and then executing a scan of the entire zoo.

Trend Micro, CA and eEye all did very well, generally detecting and blocking or removing all but about 8 to 9 percent of the malware we threw at them in all tests.

Symantec was close behind, missing 17.6 percent of specimens on the real-time scan, but performing on a par with Trend, CA and eEye in the on-demand scans.

McAfee was next, with 22.3 percent of our specimens eluding the real-time scan. The follow-up on-demand scan, however, produced some surprising results: Another 10.7 percent of specimens were detected, but none of those were deleted or quarantined. Likewise, in the pure-play on-demand scan, all of the 8,000-plus malware specimens survived, despite an avalanche of alerts. That's because this new McAfee product has a default action of alert-only for on-demand scans, in contrast to the competition and a departure from most earlier McAfee products.

With the help of McAfee support, we used the McAfee client to conduct an on-demand scan with a delete action, a process that requires several rather nonintuitive steps. After that scan, 11.6 percent of our initial specimens remained for both the on-demand scan following real-time scanning and the pure-play on-demand scan. Notably, McAfee blocks all .exe files from a network copy, even benign test files, due to another default setting. Such a feature is likely to cause problems in environments attempting to distribute programs via network file shares, and is certain to be disabled in some enterprises.

When we tested Sophos, every one of our specimens survived the initial copy because, by default, Sophos' real-time defenses only look at "read" actions, not "write" actions. Such an approach, possibly done to improve file system performance, prevents the malware from executing, but does not stop infiltration of malware into a file system. Sophos does offer an option for changing this default behavior.

In the end, both the real-time/on-demand combo test and the pure on-demand test left 36.7 percent of the specimens on the target machine.

Sophos' default behavior is to perform "in-place" quarantine, preventing future access of the file but leaving it in its current location. All the other products move malware to a separate quarantine directory or delete it. Sophos says its approach makes restoration of files misidentified as malware easier. If your antivirus tool makes false-positive matches on legitimate files, restoring access in their normal locations is a lot easier than scraping them out of a quarantine directory and finding their homes again. The Sophos tool can be configured to perform traditional quarantine or deletion.

IBM ISS was rated lowest in this series of tests, crashing several times and scoring so poorly as to cause us to double-check that protection was enabled. IBM ISS leaves signature-based antivirus turned off by default, another indication that this product is typically used to augment another vendor's antivirus solution. IBM ISS has licensed BitDefender's antivirus and antispyware functionality in its endpoint suite, which we activated before starting our test regimen. The initial real-time test completed without the tool blocking a single file. According to IBM ISS support personnel, file copies across Windows network shares are not scanned, even with the on-write scanning option enabled. This stance mystifies us, considering that users could copy infected files on a file server back to their clients without any real-time protection.

The on-demand scanning was hardly better. The follow-up on-demand scan started off as expected, but halfway through the scan (according to the progress bar) scanning stopped and we were greeted with the message "Successfully Completed." However, the same GUI listed "Number of Files Remaining: 4,430" and we still counted 58.6 percent of our malware in the target machine's file system. This stop and start repeated several times during the scan. We re-ran this test several times, but 34.7 percent was the best IBM ISS managed in repeated on-demand scans.

ENDPOINTS | Antimalware Scanning

THE GOOD NEWS Trend Micro, CA and eEye all did very well, generally detecting and blocking or removing all but about 8 to 9 percent of the malware thrown at them.

THE BAD NEWS IBM ISS crashed several times, scoring so poorly as to cause us to double-check that the protection was enabled.

EXPLOIT PROTECTION

Every vendor in our analysis claims to protect systems against exploitation using some form of HIPS technology. Different vendors use this term for a variety of disparate technical defenses *(see "HIPS Hydra," p. 26)*. Regardless of approach, we wanted to see how each vendor would fare against exploitation attempts in a series of three tests. We disabled each product's firewall component to focus the test exclusively on HIPS functionality.

First, we attempted to exploit client-side software running on the protected hosts, trying to attack Internet Explorer via the IE CreateObject vulnerability (MS06-014) and VML flaw (MS06-055). We also tried to exploit the Firefox browser using the Mozilla_CompareTo vulnerability.

Our second test measured how well each product defended listening services on the protected system, particularly services associated with Windows networking. We attempted to exploit both the MSRPC DCOM buffer overflow flaw (MS03-026) and the LSASS buffer overflow issue (MS04-011). To add some variety to this network service testing, we attempted each exploit with both a standard command shell payload and the Metasploit Meterpreter shell, a more sophisticated and often harder-to-detect attack that provides specialized remote shell access running from within an exploited process. Our client-side and server-side testing relied on Metasploit Framework version 3.0, using all default settings except for the HTTP port for browser exploits, which we changed from 80 and 8080 to another number to simulate an attacker who tricks a client into clicking on a link with a port number in it.

Our third test was designed to look at how each vendor could defend against zero-day exploits of third-party applications. We created our own network-listening program with a buffer overflow flaw, and wrote some code to exploit it to give remote command-shell access on the target machine.

Overall, eEye performed best in detecting exploits. It was the clear leader in identifying client-side attacks, alerting on all of our tests, but by default did not block; it simply displayed the alert "Application Protection: Suspicious System Call." This default behavior could be altered to block such exploits, as a global switch applied to all such suspicious system-call detection. While blocking is the goal, concern over false-positive blocks makes eEye's default setting reasonable.

eEye successfully stopped all service exploits. However, in the process of blocking the MSRPC DCOM exploit, it killed the svchost.exe process, which made our Windows machines reboot themselves within 60 seconds. It is generally considered better to kill an exploited process rather than run the attacker's code, but re-booting could result in loss of valuable data.

eEye detected and alerted on our zero-day exploit; when

ADDITIONAL FEATURES

END-TO-END ENDPOINT SECURITY

Some vendors provide additional capabilities that could be very helpful in an enterprise environment. Some organizations will certainly desire these features, while others will not. For that reason, we did not include them in our overall product ranking, but still detail them here for organizations that will take them into account in their buying decisions. (See "Building Blocks," opposite page.)

Web site reputation analysis helps combat phishing attacks, spyware downloads and related nastiness. The technique assesses the URLs accessed by a browser to determine whether they are associated with a malicious or suspicious Web site. If the endpoint product detects such a URL, it prevents the browser from surfing there, either displaying an alert message or redirecting it to a safe site. McAfee and Trend Micro deploy an impressive and vast infrastructure of monitoring software to discover malicious Web sites and update their extensive blacklists.

Application execution control enables enterprises to prohibit the execution of unwanted game and P2P applications by using a blacklist. Similarly, application control blacklists can help limit the outbreak of a fast-spreading worm or browser exploit by blocking execution of the malware on the end system. Using whitelist functionality, an enterprise can severely lock down a machine, allowing it to run only required applications. Most endpoint security solutions can look at the name, file system location, and/or MD5 hash of a given executable to determine whether the application should be allowed to run on the protected machine.

All of the vendor products we tested, except for Trend Micro, supported some form of application execution control. eEye and Symantec offer the most flexibility, with custom-designed whitelists and blacklists based on executable

BUILDING BLOCKS

	CA	eEye Digital Security	IBM ISS	McAfee	Sophos	Symantec	Trend Micro
Antivirus	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Antispyware	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Personal firewall	Yes*	Yes	Yes	Yes	Yes	Yes	Yes
HIPS/buffer overflow protection	Yes*	Yes	Yes	Yes	Yes	Yes	Yes
Web site reputation analysis	No	No	No	Yes	No	No	Yes
Application execution control	Yes*	Yes	Yes	Yes	Yes	Yes	No
USB control	Yes*	Yes	No	No	No	Yes	No
Vulnerability scanner	No	Yes	No	No	No	No	Yes
Supported operating systems	AV/AS: Windows NT, 2000, XP, 2003, Vista. HIPS: XP only.	Windows NT, 2000, XP, 2003.	Windows 2000 Pro, XP. No server support in endpoint product, so no Windows 2003**	Windows NT, 2000, XP, 2003, Linux	Windows 95/98/Me, NT, 2000, XP, 2003, Linux, Mac OS X (PowerPC and Intel)	Windows 2000, XP, 2003, Vista. Separate, unmanaged clients for Linux and Mac OS X.	Windows 2000, XP, 2003, Vista. Mac OS X support by end of year.

* CA's endpoint solution requires the purchase of two separate products: its anitvirus/antispyware solution and its HIPS solution. All items marked with a * are available only with the purchase of the HIPS solution.
**Server protection is a separate product.

path, hash, or regular expression matching. CA's blacklist and whitelist capabilities were also impressive, augmented by a graylist function that allows admins to define rules for execution of specific programs only by certain users and at certain times of day. Sophos has the most limited application control, supporting only blacklists for those applications for which it generates a specific signature.

USB control addresses the concern of the many enterprises that have been stung by employees who bring malware into the organization on USB thumb drives or spirit gigabytes of sensitive data out on them. CA, eEye and Symantec feature device control capabilities.

eEye's USB control is an all-or-nothing feature—all USB tokens are allowed for all users of a given policy group, or all users are forbidden from using the devices.

CA and Symantec's solutions are much more fine-grained, allowing policies to be defined for various Windows devices such as USB, infrared, FireWire, SCSI, and serial or parallel ports. Their device control policies can specify certain kinds of files (e.g., Word documents, executables, etc.) that can be accessed or denied on specific devices.

Vulnerability scanners assess the local host on which they are installed to determine its security stance, looking for unpatched operating system components or weak security settings. In particular, eEye includes a copy of its comprehensive Retina security scanner on every host; it's a powerful suite that can discover several hundred potential security flaws and propose fixes. Trend Micro includes a far more limited scanning tool that can search the system for applications that the IPS functionality has signatures to protect.

—ED SKOUDIS & MATT CARPENTER

THE GOOD NEWS Overall, eEye performed best in detecting exploits.

THE BAD NEWS CA fared poorly in detecting and blocking client and services exploits.

we tweaked the configuration to an action of "Terminate Process," it blocked as well.

IBM ISS came next. On the client side, it detected and blocked the VML exploit. However, the alert messages for the IE CreateObject and Firefox attacks didn't indicate that the product had detected the exploit action, only that it identified a Microsoft Windows shell banner passing across the network. An attacker could launch such an exploit without creating a banner, thereby dodging this form of detection.

IBM ISS identified and blocked all services-based attacks, with an alert that cited the specific exploit we used, the ideal behavior for the product under these tests.

It allowed our zero-day attack, again merely alerting to the presence of a Windows shell banner.

Sophos delivered reasonable performance in our client-side testing, alerting on two exploits as "Buffer Overflow" behavior, but missing the CreateObject exploit. The default action is to alert, but Sophos can be configured to block the attacks.

All of our services attacks were detected, but by default they were allowed through, giving the attacker control of the system. Sophos neither detected nor blocked our zero-day exploit.

McAfee detected and blocked our VML and Firefox exploits, but failed to detect our CreateObject exploit. McAfee detected and blocked all of our service exploits. For zero-day defenses, McAfee requires administrators to configure specific applications to be protected on a machine. By default, nothing other than specific Windows components is protected, so our zero-day attack went undetected. As an experiment, we configured McAfee to add zero-day protection to our custom vulnerable application. Unfortunately, our exploit still went undetected.

Trend Micro and Symantec came next in our exploit testing. Neither identified nor blocked a single client exploit. Trend Micro support personnel indicated that the HIPS protection it licensed from Third Brigade (as well as the protections offered by other vendors) is often configured by default to look for browser exploits only on TCP ports 80 and 8080. Again, independent of our scoring, we tweaked our test to verify this claim, and Trend Micro did detect our attacks on those ports. Administrators can add lists of additional ports for browser and other HTTP-related defenses. Ideally, an admin would configure the endpoint security suite so it monitored for HTTP and HTTPS attacks on all ports allowed out through the enterprise's network firewall. In many organizations, unfortunately, the number of ports allowed outbound are rather high and change on a regular basis, making this synchronization of network firewall and endpoint security tool difficult.

SEVEN-HEADED PROTECTION

HIPS HYDRA

Host-based intrusion prevention system (HIPS) functionality means many different things to the vendors that include such capabilities in their endpoint security suites. The goal, of course, is to prevent the end system from being compromised by an attacker, but the technological approach of the vendors implementing HIPS varies widely. We interviewed each vendor, asking them to describe their technical approach to blocking exploitation attempts. We wanted to focus specifically on defenses against buffer overflow and related code execution exploits. Based on our interviews, we identified seven essential forms of such exploit detection and prevention:

- System call backtracing analyzes various system API calls to ensure the calling address exists in a known code segment.
- Spawn blocking limits which programs can run new programs (for example, blocking a browser from running a new command shell process).
- Behavior checking monitors system calls for combinations that historically have indicated that an attack is under way.
- DLL loading checking looks for unusual or unexpected DLLs to be loaded into running applications on the machine.
- Call verification ensures the return address for the current function is immediately preceded by a call instruction.
- SEH validation protects against exploits that overwrite exception handlers by validating the Structured Exception Handler chain.
- Network-based IPS monitors network traffic for known vulnerabilities and exploits.

CA implements spawn blocking, DLL loading checks and network-based IPS. eEye relies on system call backtracing, call verification and network-based IPS. IBM ISS uses system call backtracing and network-based IPS. McAfee has created a patented "generic buffer overflow protection," although it declined to share details with us before press time, as well as network-based IPS. Sophos uses system call backtracing. Symantec implements behavior checking and network-based IPS. Trend Micro focuses exclusively on network-based IPS. ›

—ED SKOUDIS & MATT CARPENTER

Both Trend Micro and Symantec detected and blocked all of our services exploits, but neither detected our zero-day attack.

CA fared worst of the seven products in this series of tests, failing on most. It didn't detect or block any of the client exploits with its default security policy. Although not part of the scoring, we experimented with its "Restrictive Policy," which did block all of the exploits, but also prevented Firefox from accessing the network.

The next set of results were, if anything, poorer, as it did not alert or block our services exploits, even when we applied Restrictive Policy.

The one success was that CA detected and blocked our zero-day exploit under default policy.

REPORTING

We evaluated each product's reporting functionality, used to pull information such as long-term attack and infection trends, policy compliance information, and lists of the most problematic groups of machines. In particular, we looked at comprehensiveness, flexibility and ease of use

Though McAfee's management GUI was disappointing, ePO's reporting features are excellent, including more than 70 different reports that break down all aspects of the enterprise. The point-and-click custom report creation tool is stellar, making it easy for people who are not database experts to massage the information into highly useful reports.

Symantec is also solid, offering more than 70 reports, with impressive performance. Symantec's custom reporting capabilities are focused on defining filters for its existing reports to create useful subsets, a valuable capability but somewhat less flexible than McAfee.

The IBM ISS reporting tool provided good coverage, addressing long-term trends and top attacked and infected machines. However, getting at the report files is a little obscure. Admins have to remember where they were generated in the file system to open the report from within the management GUI. Further, to open a report, you have to right-click on it and go to "Properties," a bizarre GUI twist that takes some getting used to.

Trend Micro's reporting is handled by a separate product, Trend Micro Control Manager, which is not tightly bundled into the existing management GUI, making a little more work for installation and use. On the positive side, this separate reporting tool applies to all Trend Micro enterprise products, including gateway security appliances, antispam products, etc. It's included in the purchase of the endpoint suite, and provides a full complement of well-laid-out reports.

ENDPOINTS | Reporting

THE GOOD NEWS McAfee ePO's reporting features are excellent, including more than 70 different reports that break down all aspects of the enterprise.

THE BAD NEWS Sophos' reporting capabilities are quite skimpy. Only about a dozen reports are available.

eEye's built-in reporting features are decent and offer some features for creating custom queries in its published database schema. However, building custom or tweaked report queries is a complicated process, even using the built-in templates.

CA's reporting for antivirus and antispyware is stellar, with more than 70 reports available. Unfortunately, CA's HIPS and firewall features offer very little reporting, with only about a dozen high-level reports providing much less visibility into these important aspects.

Sophos' reporting capabilities are quite skimpy. Only about a dozen reports are available. They don't include Top 10 style reports of most infected systems, users or groups. The look and feel of the reporting engine makes the product appear better suited for small and medium businesses, rather than large enterprises. However, Sophos publishes its database schema for customers to use with third-party reporting tools, such as Crystal Reports.

INTEGRATION OF COMPONENTS

Endpoint security suites should integrate disparate components into a coherent, manageable whole. Most of the vendors have worked hard to integrate various aspects of their solution, with high marks going to eEye, McAfee, Sophos and Symantec.

The IBM ISS product has good integration, but often looked like it was packaged around complementing another vendor's antivirus and antispyware, rather than providing a whole solution. This approach makes sense given that IBM ISS HIPS capability has been an added defensive layer to complement traditional antivirus/antispyware products. IBM ISS has licensed BitDefender's antivirus/antispyware technology, but the management GUI still appears as though it is merely grafted in.

ENDPOINTS | Integration

THE GOOD NEWS Most of the vendors have worked hard to integrate various aspects of their solution, with high marks going to eEye, McAfee, Sophos and Symantec.

THE BAD NEWS CA was never one for deep integration of components in its antimalware solutions.

Trend Micro's antivirus/antispyware integration is decent, but integrating the personal firewall and HIPS, licensed from Third Brigade, needs some work. These components are a plug-in inside the management GUI, with a separate set of configuration screens that don't have the same look and feel of the configuration of antivirus and antispyware. Further, on the client side, antivirus and antispyware is a completely separate program from the HIPS software.

CA was never one for deep integration of components in its antimalware solutions. Its endpoint security product continues to separate management of antivirus and antispyware on different screens, but at least they both are available in one GUI application. CA's HIPS, on the other hand, is a separately purchased product. It is installed and man-

MAKING THE GRADE

	CA Threat Manager 8.1, Host-Based Intrusion Prevention System 8 www.ca.com	eEye Digital Security Blink Enterprise Edition (Blink Professional 3.1, REM Security Management Console 3.5) www.eeye.com	IBM Internet Security Systems Proventia Desktop Endpoint Security 9.0 www.iss.net	McAfee* Total Protection for Enterprise www.mcafee.com	Sophos Endpoint Security and Control 7.0 www.sophos.com	Symantec* Endpoint Protection 11.0 www.symantec.com	Trend Micro* OfficeScan version 8.0 www.trendmicro.com
Enterprise management 30%	B	B	C	D	A-	A	B+
Antimalware scanning 20%	A-	A-	D	C	C-	B+	A-
Exploit protection 20%	D	A-	B	B-	B-	C	C
Reporting 20%	B-	B	B+	A	C	A-	B+
Integration 10%	C	A	A-	A	A	A	B
VERDICT	**B-** Lacks integration of separate products. Pros: Good antimalware scanning Cons: Weak exploit protection and integration	**B+** A worthy new entry with a comprehensive solution. Pros: Excellent protection with decent management Cons: Client GUI better than Enterprise Management GUI	**C+** Product much more geared around augmenting an antivirus/antispyware solution, not replacing it. Pros: Decent exploit protection Cons: Malware scanning default settings weak and prone to failures	**C+** A disappointing new management GUI could make transition difficult. Pros: Excellent reporting Cons: New management interface complex and difficult to use	**B-** A moderate performer in all categories. Pros: Very straightforward management capabilities and integration Cons: Skimpy reports that can be augmented with third-party tool	**B+** Solid product for the enterprise. Pros: Excellent enterprise management Cons: Exploit protection relatively weak	**B** Decent product, but HIPS integration needs work. Pros: Good antimalware scanning and new plug-in architecture for expanded functionality Cons: Client-side exploit protection tied to port number

*Beta. Symantec Endpoint Protection 11.0 was made available Sept. 27; all tested components of McAfee Total Protection for Enterprise have been available since Oct. 11; Trend Micro OfficeScan 8.0 will be available Nov. 15.

aged using its own GUI and a separate client package is installed on protected workstations.

In the End, Take It Slow

Symantec's new offering looks very solid, and eEye is a worthy new competitor in the endpoint security space. Trend Micro has a decent solution and a promising plug-in architecture for future expansion. CA and Sophos did reasonably well, but neither shined consistently. Finally, we were very disappointed with the numerous glitches, unfortunate design decisions and poor performance of McAfee and IBM ISS.

Regardless of which vendor you choose, keep in mind that the endpoint security market is relatively immature—witness our beta testing of three major vendors—and the complexity of any of these products warrants a carefully planned deployment strategy. We urge you to experiment with the products on your own laboratory test systems with images from your production environment to make sure they don't have any adverse consequences on your particular application mix.

Double check default policy settings to make sure they offer reasonable protection, and if not, adjust them for your environment and risk profile. And, finally, have your support staff become familiar with the various quirks of these management GUIs before production roll-out.

Ed Skoudis is a co-founder of Intelguardians, an information security research and consulting company. He is also a fellow with the SANS Institute, where he teaches the Hacker Techniques, Exploits and Incident Handling course.

Matt Carpenter is a senior security analyst for Intelguardians with expertise in hacker attacks, defenses and security vulnerability research. He has released several open source security research tools. Send comments on this article to feedback@infosecuritymag.com.

IDENTITY MANAGEMENT

IMAG 500

RREVIEWED BY BRAD CAUSEY

Apere

www.apere.com

Price: **$15,000 for a single device**

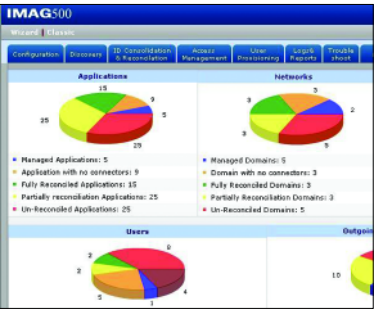

Apere's IMAG 500 appliance aims to simplify the complex maze of identity management through the automatic discovery of distributed identity stores, consolidation, reconciliation and provisioning of all user accounts, and access to practically all network resources through a single control point.

It's an attractive proposition, but we found the implementation rough going.

Configuration/Management D

Setup was difficult, ultimately requiring the company's technical support to remotely access the device to assist with the configuration process. Although Apere claims that the device can learn the location and type of applications, each application required manual configuration.

Basic device configuration is like any other appliance: You set up DNS servers, email servers, log servers and VLANs.

Things got sticky when we tried to add authentication resources and user stores, particularly for Web applications. Because user accounts don't generally live on the Web server itself, IMAG doesn't have a way to tie users to the resource they need to access. In the absence of an API, you have to enumerate users from each identity store and reassign the resource that IMAG associated with the users. You then have to get the user information, reformat it and redirect the allowed resources for those users to point to the Web server.

Rather than use an Active Directory domain, for example, the provisioning process requires an "authoritative list," in a specifically formatted comma delimited text file.

Further, IMAG wouldn't recognize user formats used by key applications such as SQL Server and Project Server for AD accounts. As a result, these IDs must be manually matched to their account on the user store.

Testing methodology: Our lab included a single Active Directory domain and a single LDAP tree. User accounts were enumerated from various sources such as MySQL, SQL Server, Web applications and various client-server applications. User roles such as administrators, power users and end users were set up to test access controls.

Policy Control D

Access control entries are based on user ID, VLAN and IP address, but IMAG lacks a proper grouping mechanism, so it can assign only one or all users, for example, to a resource. Anything in between will require extensive work to set up.

Policies provide the ability to control which users or groups of users can access particular applications, but are only usable when the device has been deployed in an in-band configuration. When deployed out of band, IMAG can still serve as a central ID consolidation and reconciliation point, but any other benefits are lost.

Effectiveness D

IMAG has some major challenges to overcome. An ID management solution needs to be able to effectively link and manage identities from stores of user accounts. Tasks such as importing users from identity stores and consolidating them proved extremely difficult. Some of the most basic applications, such as SQL Server and LDAP, were a major challenge.

Proper user provisioning requires that you manually create an account in each application and assign that user to each resource via the user interface.

Reporting B

Reporting is handled through an easy-to-use Web interface. Each report is customizable, and canned reports are available for applications, authentication stores and user provisioning.

You can research orphaned users, unmanaged resources and user stores that have not been reconciled. Each report can be filtered based on criteria such as specific users or resources.

Verdict

Apere says many of the issues we encountered are addressed in its next release, but mid-enterprise businesses may not have the tolerance for a product with so many features missing or unfinished.

PRIVILEGE MANAGEMENT

HotPick INFORMATION SECURITY®

BeyondTrust Privilege Manager 3.0

REVIEWED BY BRAD CAUSEY

BeyondTrust
www.beyondtrust.com
Price: **$30 per seat**

The least privilege security model is the de facto standard for reducing the risks of elevated user privileges. This can be a challenge in Windows environments. You don't want your end users to have general admin rights, but they may need them to run the applications required to do their jobs. There's no easy way to manage this, so companies wind up letting users have excessive privileges, leaving their desktops, user accounts and software vulnerable to attack.

BeyondTrust's Privilege Manager 3.0 solves this dilemma through a Group Policy extension that allows organizations to control permissions for selected processes and applications. BeyondTrust has also introduced a new technology, called ShatterProof process isolation, that prevents shatter attacks, a complex privilege escalation technique.

Configuration/Management A

Privilege Manager adds GPO extensions that integrate with Internet Explorer and Microsoft's Group Policy Management Console, so admins can work directly through a customized Active Directory interface. The installation was very easy and fast, consisting of an MSI with few requirements, chiefly the .NET framework and AD's Group Policy Management Console. (These can be downloaded free from Microsoft's Web site.)

Once installed, the Privilege Manager settings are available by simply opening the Group Policy Object Editor. It gives you a single interface to manage the custom add-ons as well as the default GPO settings, simplifying management and reducing administrative overhead.

Each managed computer requires client software to capture and manage permissions for processes and programs; it can be installed through standard software deployment or via Group Policy. The client reads the custom GPO settings and modifies the security token on programs as they launch, giving the user elevated privileges as defined by Privilege Manager rules.

Policy Control A

Creating policies for application privileges is simple and intuitive. Each rule allows you to identify a target process or executable name. This can be done by a number of different methods, including MSI GUID, hash, path, folder or ActiveX rules, giving you tremendous flexibility.

For each rule, you define what action will be taken, including modifying privileges and permissions for target applications. Defined privileges dictate what components of the system will be accessible when the program or process is initiated and for the duration of its run time. These rules can be configured with filters that restrict what settings apply to what group. For example, you can disable the policy for a specified application based on a wide range of criteria, such as subnet, computer name, user, security group or organizational unit. In addition, you can modify Internet Explorer behavior and ActiveX security through a custom administrative template.

Effectiveness A

Privilege Manager provides an extremely effective framework for implementing least privilege policies. The overall concept of least privilege in an enterprise environment is plagued with difficulties. Often, developers have to get involved, code has to be changed, and massive amounts of time will be spent during implementation and dealing with unknowns. Because Privilege Manager integrates with Group Policy, it will significantly simplify the management of application privileges and permissions.

Verdict

Privilege Manager will prove invaluable for implementing and managing a least privileges program. Although long-term management of each application will be complex, it helps cut the job down to size.

Testing methodology: Clients in our AD domain consisted of several Windows 2000 and Windows XP computers with various service packs. A variety of applications were tested including Web sites with ActiveX requirements, DOS-based applications, network-based applications and locally installed programs.

PASSWORD SECURITY

Enterprise Password Vault 4.0

REVIEWED BY TOM BOWERS

Cyber-Ark Software
www.cyber-ark.com
Price: **EPV server, $25,000; user pricing starting at $220 per user**

Privileged users hold the keys to your kingdom: passwords that control administrative access to devices and applications across your enterprise. Cyber-Ark's Enterprise Password Vault (EPV) is among a handful of specialized products designed to securely manage these sensitive passwords, controlling privileged accounts across a wide range of client/server and mainframe OSes, switches, databases, etc.

It provides the privileged account controls mandated by regulations, and its central repository makes it an ideal addition to identity/access management projects.

Installation/Configuration C+

Overall, this was a tedious installation/configuration process. EPV is in serious need of an installation wizard and graphics-filled documentation to help users understand the purpose of each of its components and where it sits in the architecture. The documentation, while voluminous, is disjointed and difficult to follow.

The expectation is that the four components be distributed on at least two Windows 2003 servers, and we sorely missed an overall diagram to reference the separate installations.

We were somewhat vexed, for example, when we installed the last component, Password Vault Web Access. We belatedly realized that you need IIS installed on the second server—something the documentation didn't mention until then.

Testing methodology: EPV was tested on multiple fully patched and hardened Windows 2003 servers and Windows XP workstations. We used a sample database of users and passwords, and scanned the system for weaknesses using standard penetration testing tools and forensic analysis software.

Effectiveness A

The EPV experience is superb once the system is installed. Operationally, the end user password management system is an intuitive, wizard-driven interface, requiring little to no training.

The system is organized around the vault, which contains multiple safes. Each safe is independent and may be connected to one user or group, or many of both. A person in one group or safe cannot see the existence of other safes nor access them without explicit permission. Each safe also has an owner or owners that control access. Via the safe, passwords are synchronized with the end products, such as routers, switches and servers; changing the password in the safe also changes it on them.

Essentially, the EPV takes control of the admin logon function. For example, an admin logs on to the EPV Web interface to access the password object associated with a switch they wish to manage. This object gives them the new password, they log on to the switch and conduct their maintenance. Passwords can be generated based on internal policies and/or regulations such as FFIEC or the Family Educational Rights and Privacy Act.

The architecture is very secure. That's obviously a critical point, but we don't see it often enough in enterprise security products. We encountered no way for password information to leak, either through the vault or the browser-based interfaces. A firewall on the PrivateArk server protects the host, opening a single port that allows only Cyber-Ark's proprietary protocol.

Reporting B

Reporting is very well executed, but lacks a cohesive export mechanism.

Reports are clear and concise. A nice dashboard presents reports and graphs that provide good auditing capabilities to help meet regulatory requirements.

The exporting mechanism is smooth yet somewhat disappointing. Reports can be exported only to Microsoft Access and Excel, or via CSV format.

Verdict

EPV is a valuable tool and a maturing product that performs its privileged password management function very well.

BIOMETRIC AUTHENTICATION

DigitalPersona Workstation Pro and Server 4.0

REVIEWED BY BRENT HUSTON

DigitalPersona

www.digitalpersona.com

Price: **Server, $1,499, plus $50 authentication license per user; Workstation without reader, $60, with DigitalPersona U.are.U 4000B reader, $149**

Biometric authentication has met considerable market resistance, mostly because of integration issues, accuracy and cost. With improved technology and the introduction of laptops equipped with fingerprint readers, biometrics may be starting to move into the mainstream. DigitalPersona Pro is a robust single sign-on (SSO) software suite that allows an enterprise to replace passwords with biometric fingerprint readers or provide dual-factor authentication.

Installation and Setup B+

There are two pieces to the suite: DigitalPersona Pro Workstation software for individual systems and the server component, which integrates with Active Directory on your domain controller. While the workstation software can function by itself, the server provides domain-wide SSO.

Installation is straightforward. The server installation requires a few more steps to integrate with Active Directory, but it's all detailed in the manual. After installation, the workstation software starts a wizard, which records your fingerprint. After a few repetitions, we were able to register a fingerprint in less than 10 seconds. The workstation software automatically detects any DigitalPersona Pro servers on the local network.

Both the server and workstation software can be purchased with or without DigitalPersona's fingerprint reader. The latest version of DigitalPersona Pro offers wide support for third-party readers, such as those becoming popular in new business-class laptops. The DigitalPersona optical reader is quite good; we found it to be accurate, with few false negatives and no false positives.

Testing methodology: DigitalPersona Pro Workstation was tested as a standalone product on Windows XP desktops, and in an AD environment with the server component on Windows Server 2003.

Workstation (Single User) B-

The workstation software, in standalone mode, is rather simple. It integrates with Windows logon and also provides an SSO function that seems to be geared toward home users. The SSO feature provides an automatic wizard that will detect the login fields in many applications. Unfortunately, we found there are some apps it does not support (such as terminals, like Putty). It also supports only Internet Explorer, a problem considering the growing popularity of Firefox. However, it is very easy to use, fast and accurate with applications it supports.

Server (Centralized Environment) B+

The server software is much more robust. The SSO wizard allows manual creation of login templates to support applications that the automatic wizard can't detect; these templates are pushed out to desktops via GPOs.

Creating a template is fairly easy. You need to make sure the window title is accurately reflected within the SSO administration tool. You then enter the actions required for login—e.g., entering keystrokes into a field, time delays, or x-y coordinates of a window. Templates can also be created for password-change forms, which can be used to automatically generate passwords. The created templates can be either pushed out to workstations via GPO or copied manually.

The server centrally manages fingerprint data for all users, with tight Active Directory integration. It also provides event logs for fingerprint logins to help with regulatory compliance, but lacks strong reporting capabilities. It also provides a handy query tool to easily discover who has registered fingerprints.

Verdict

Enterprises looking for a biometric single sign-on solution won't be disappointed with what DigitalPersona Pro offers. The software is easy to use, and can function with single- and two-factor authentication.

ACCESS CONTROL

eGuardPost

REVIEWED BY STEVEN WEIL

e-DMZ Security

www.e-dmzsecurity.com

Price: **Starts at $12,500 for five concurrent sessions**

Secure remote vendor and system administrator access to information systems is a critical business requirement for many organizations, but it can be a challenge to manage and audit. While VPNs are fine for most users, they can require client software and don't offer the level of audit and forensic capabilities demanded by regulatory requirements and high-security environments. eGuardPost is a hardened appliance that can be used to secure, manage and audit these sensitive connections.

Policy Control B+

eGuardPost allows security managers to apply granular access controls to remote connections. The appliance comes bundled with Security's Password Auto Repository (PAR), e-DMZ's flagship product, which securely stores and manages administrative passwords.

We were able to successfully create multiple users and enforce a variety of access controls on them.

Once users log in via HTTPS and are authenticated via RSA Security's SecurID, Secure Computing's SafeWord or LDAP (or against user accounts created and stored on eGuardPost), eGuardPost determines what type of remote access they are allowed and which systems they can connect to. Security managers can assign specific roles (e.g., requester, approver, auditor and administrator) to remote users.

eGuardPost can be configured to automatically log in specific users; it retrieves the necessary password from the local or a remote PAR. The password is never shown to, or known by, the remote user.

Testing methodology: Our test network included a Windows XP laptop, an unmanaged switch and three Windows 2003 Web, FTP and domain controller servers.

Security managers can also require that certain remote connection requests be approved by one or more designated persons. Connection requests and approvals can be sent to a ticketing system.

Configuration/Management B

Configuration is straightforward and easy thanks to excellent documentation. The appliance is managed via HTTPS. The management interface is well designed and mostly easy to navigate.

Systems to be managed are defined, users are created, and the security manager determines which users have what type of remote access to which systems. You can even limit access to specified time periods, which will be very useful for vendors and contractors, as well as admins assigned to particular tasks. Systems and users can be placed into and managed as groups.

Users do not need to install any software; eGuardPost proxies all remote connections. It can establish connections to systems via Telnet, Windows Terminal Server, SSH, VNC and X5250.

Reporting B

eGuardPost's forensics capabilities are unique, offering VCR-like recording and playback of every mouse, keyboard and screen action during a remote session. We conducted multiple remote sessions via eGuardPost then watched their recordings; each was flawlessly presented. eGuardPost can automatically move recorded sessions to designated archives.

eGuardPost can produce detailed reports of user rights and activities, security alerts, firewall events, database events and Web server events. Reports can only be exported to Excel and some of them are a bit cryptic. The appliance supports SNMP and syslog.

Effectiveness B+

We found eGuardPost to be a very effective product, correctly and efficiently managing and auditing all of the many remote connections we sent through it. eGuardPost is carefully hardened, with an embedded firewall and hard drive encrypted with 256-bit AES. Our security scans of the appliance found no vulnerabilities.

Verdict

eGuardPost is a well-designed and highly capable product that meets an important need. It has strong security and great forensics capabilities.

IDENTITY MANAGEMENT

Identity Engines Ignition Server

REVIWED BY SANDRA KAY MILLER

Identity Engines
www.idengines.com
Price: **Starts at $33,500**

Identity Engines' Ignition Server manages access controls across disparate directory services platforms (Active Directory, LDAP, eDirectory) by consolidating them into a single user store. Deployed as an alternative to RADIUS, the appliance includes a comprehensive policy engine to use with multiple access control devices (wireless access points, switches, firewalls, VPNs) throughout a heterogeneous enterprise.

Configuration/Management B+

Because of well-written documentation, we completed basic network installation in minutes. But that's where simplicity ends. Users must have an extensive knowledge of authentication protocols, directory structures, virtual provisioning and certificate management to take full advantage of the Ignition Server's features.

There are three major aspects of the Ignition Server: networked devices (authenticators), user stores (directory services) and policies.

Authenticators—devices attached to the network—can be bundled by subnet to facilitate large installations. They can be managed according to several attributes, including service categories—groups of authenticators to which policies are applied. Adding authenticators was the same as with RADIUS: Provide a name, IP and shared secret. Service category, device type (wired, wireless, VPN) and vendor are added the same way.

Ignition Server automatically connected to AD once we entered the domain name, service account name and password, and to LDAP using the service account domain name, password, IP address and port number. We could create fall-through rules across multiple directory services for a variety of situations (for example, check AD first to authenticate a VPN user, then LDAP).

Testing methodology: Ignition Server was deployed in place of the RADIUS server in our simulated enterprise network. It provided AAA services for our wired and wireless network access, as well as for a VPN.

Policy Control A

The Ignition Server is really a policy engine that speaks RADIUS. It does everything a RADIUS server would do, but it's the policy engine that sets it apart. We liked how multiple authenticators are tied together into a single service category to which three different policies—authentication, identity routing and authorization—can be easily configured and applied.

Authentication policy determines the tunnel protocols, credentials and ciphers for communication between the supplicant, Ignition Server and directory services.

An identity routing policy traverses directory services during authentication, determining which user store to apply based on the user's network domain or what device is making the authentication request.

The authorization policy controls access according to the user account.

Effectiveness A

We authenticated users to specific devices, such as wireless access points, and assigned a common policy using credentials from two directory services (AD, LDAP).

Ignition Server supports strong authentication, such as RSA SecurID and Secure Computing's SafeWord.

Security is solid. Built on a 64-bit hardened appliance running a stripped-down version of BSD, security features include onboard IDS, 256-bit AES encrypted file system, and protection against physical tampering.

Reporting C

This is Ignition Server's biggest shortcoming. While real-time statistics and logging are available, the logs could only be exported hourly, daily or weekly—nothing customized or on-demand. We'd welcome the ability to export the statistics displayed in the individual tabs.

Verdict

Organizations that need a unified policy engine to control network access using multiple authentication systems will be able to justify Ignition Server's price tag.

IDENTITY MANAGEMENT

HotPick INFORMATION SECURITY®

Sun Java System Identity Manager 7.0

REVIEWED BY BRAD CAUSEY

Sun Microsystems

www.sun.com

Price: **$50 per user**

Sun Microsytems' latest ID management solution unifies its user provisioning and auditing products, providing an impressive level of integration and functionality in a single package. Sun Java System Identity Manager 7.0 is a complete solution that allows an enterprise to use a single console for a multitude of ID management tasks, including role delegation, password synchronization, automated provisioning and compliance auditing.

Configuration/Management B

Setup was somewhat lengthy, although not difficult; a minimum installation required a dedicated server, JRE, JDK, Tomcat and MySQL. Large enterprises will need robust hardware and software components (all major databases and application servers are supported). The documentation is thorough and well written.

User data sources are added via agentless connectors. Among the supported sources, which Sun calls resources, are RSA products, BlackBerry Enterprise Server, Remedy, PeopleSoft, Siebel and all database servers. Supported resources can be added in a few simple steps, and others can be accessed through generic connectors, or custom built through the API. Sun has integrated SPML to allow for nearly any type of integration, including Web applications, which generally present a huge challenge because of their distributed nature.

Most of the common primary identity stores, such as Active Directory, require that at least one Sun Identity Manager Gateway be installed. The Gateways make Identity Manager very scalable; you add as many Gateway servers as you need.

Testing methodology: Our lab included two Active Directory domains and one OpenLDAP tree. User accounts were enumerated from various sources, including MySQL and SQL Server, Web applications, and various client-server applications. User roles such as administrators, power users and end users were created to test access controls.

Policy Control A

Policy and audit is where Identity Manager really shines. By integrating fully functional auditing capabilities into the standard interface, it allows you to provision a new user for Active Directory, RACF and Oracle, and compare the access given to current policies. If there are any violations, provisioning is automatically escalated for approval based on a process you define. You can even periodically audit existing identities for policy violations.

Delegation of duties reduces cumbersome management overhead.

Effectiveness A

Sun has done an impressive job in furnishing a comprehensive ID management solution for the large enterprise, providing fast and effective linking of users to identities. In addition to the great administration features, it handles user interaction very well. Users can easily log in to Identity Manager to handle password resets and requests for resource access. Automatic resource discovery allows a simpler approach to adding and configuring identity stores, while ID consolidation helps link various user accounts throughout the enterprise. Information from ID stores can be reconciled, eliminating inconsistencies and reducing errors.

Reporting A

Identity Manager handles all major reporting functions—getting the data, formatting and moving it—remarkably well. Clicking on the reports tab in the management interface provides access to canned reports, and you can also easily create very flexible custom reports.

Reports can be scheduled, cloned, downloaded or emailed in PDF or CSV format, or viewed in real time in a custom-built dashboard.

Verdict

Sun Java System Identity Manager excels with agentless connectors, scalability and amazing auditing.›

ACCESS CONTROL

HotPick INFORMATION SECURITY®

Symark PowerBroker

REVIEWED BY SANDRA KAY MILLER

Symark
PowerBroker 5.0
www.symark.com
Price: **Starts at $1,000 per server**

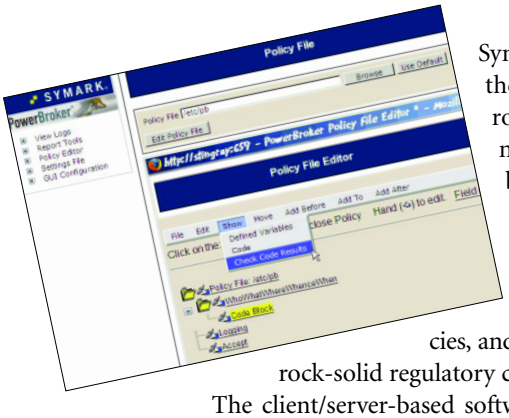

Symark PowerBroker solves the dilemma of providing root access privileges to multiple users on Unix-based systems without compromising security. It delivers comprehensive security controls through granular policies, and exhaustive auditing for rock-solid regulatory compliance.

The client/server-based software resides at the shell level, making no changes to the kernel. PowerBroker supports 30 different types of encryption—AES 256 is the default—to secure network traffic, logs and configuration files.

Configuration/Management A

Installation requires moderate expertise in Unix environments and an understanding of basic shell scripting. We used a simple batch file to disseminate the necessary files to client systems.

PowerBroker works with HP-UX, Linux, Solaris, SCO and AIX and integrates well with existing infrastructure such as routers and firewalls.

PowerBroker can be configured and managed by command line or its well-designed Web GUI, which can easily be used by someone with minimum knowledge of Unix. We used the GUI to quickly set up privileges, create and assign policies, create alerts, manage encryption, and generate and view audits, logs and reports.

Testing methodology: Symark PowerBroker was deployed in a Linux-based environment with a variety of servers requiring root privileges, including a Web server and mail server.

Policy Control A

PowerBroker's policy control is extremely granular, based on a programmable scripting language.

By assigning root-level privileges based upon on role, the actual root password is never revealed. Policies can also be assigned based upon user authentication through centralized repositories such as LDAP and SSO systems.

The new access control lists allow those unfamiliar with programming or shell scripts to write policies that control privileges through global categories such as user, system, command, time of day and day of week.

Reporting A+

PowerBroker's greatest capability is logging and reporting. Ad hoc and custom reports are easily set up and run from the Web-based report utility, drawing from massive amounts of information in the encrypted log files.

The Entitlement Report will satisfy auditors, presenting a quick overview of who can run what, and under what circumstances.

The I/O logging option records all screens and keystrokes, storing them in an encrypted file that can be used for forensic analysis or to meet rigorous regulatory requirements. It can also be used for real-time monitoring.

Data is logged in syslog format, so it can be ported to SIM/SEM products, or exported in CSV and text formats.

Effectiveness A

Everything the shell touches can be controlled through PowerBroker. Instead of logging in through bin/bash or csh, PowerBroker offers two transparent secured Korn- and Bourne-based shells. When we logged in through the PowerBroker shell, we did not have to type pbrun in front of every request we wanted to run as root.

We were impressed by the control that can be assigned to users based on role and circumstance. For example, we elevated privileges of users so they could access a particular system, such as a Web server, as root, while denying similar root privileges to a mail server. Security features include blocking predefined keystrokes, automatic termination of idle root sessions, and checksum comparisons to identify potential malicious code.

Verdict

PowerBroker is a scalable solution that effectively delegates root privileges securely and provides excellent audit trails for regulatory compliance.

INTRUSION PREVENTION

ISG 2000 with IDP

REVIEWED BY PHORAM MEHTA

Juniper Networks

www.juniper.net

Price: **Starts at $42,500**

The marriage of firewalls and intrusion prevention makes good sense, as IPS technology matures and gets serious enterprise interest. Juniper Networks' ISG 2000 appliance combines firewall, VPN and its latest intrusion detection and prevention software in an effective, high-performance package.

Installation/Configuration B+

The ISG 2000 is a multigigabit integrated firewall/VPN system with a modular architecture, enabling high scalability and flexibility. To add IDP, the organization has to get an advanced license, possibly buy extra memory and purchase up to three security modules, depending on their usage and throughput requirement.

ISG with IDP tightly integrates the software available on standalone IDP products with ScreenOS 5.4.0r2, a security-specific operating system with the capacity to handle high-speed, high-volume traffic inspection.

Although the appliance offers a console for configuration, the best way is to use the Netscreen Security Manager (NSM), a dedicated Red Hat Linux or Solaris console for managing Juniper security products. The user interface or the management client is the final component that is installed on an administrator's machine (Windows or Linux) to configure the ISG and any other ScreenOS-based devices in the network.

The user interface is designed well but still complex because of the number of settings and features available. When the device is added, NSM automatically detects the OS and the installed license, and enables/disables appropriate features accordingly. Adding IDP rules is easy and similar to adding firewall/VPN rules. Juniper provides a rich database of checks that can be used to match and drop, or just log the attack traffic between specified sources and destinations.

Testing methodology: We set up a lab with Windows and Linux PCs sending legitimate as well as malicious traffic back and forth through ISG 2000.

Effectiveness/Performance A

Juniper Networks' Multi-Method Detection (MMD) technology uses up to eight different intrusion detection methods, including stateful signature, protocol and traffic anomaly detection, and backdoor detection.

We tried—without success—to dupe the ISG 2000 using a variety of detection-evasion techniques such as splicing and fragmentation, while executing DoS and OS exploit attacks. We were amazed to see how little all those attacks affected the performance of this beast, which leverages a fourth-generation security ASIC, the GigaScreen3, along with high-speed processors.

NSM lets you view the code of the current checks and create your own checks within the IDP database.

Administration B

Like any access control system, it is imperative that the IDP rules be verified and updated on regular intervals on the basis of the normal traffic flow. It's easy to set up daily updates and many other tasks, such as importing updated configurations and rebooting devices. The management interface can be used to specify actions like SNMP, syslog or email alerts when specified criteria are matched. Because NSM stores all the information required on the server, you can take care of device and log backups like any other system.

Reporting B+

NSM's reporting module is a powerful and intuitive tool, with multiple predefined reports grouped by type of data, including firewall/VPN, IDP and administration. Each grouping includes many report templates for top attacks, attackers and targets, giving comprehensive information with graphs. You can also create custom report queries and run them automatically. Reports can be exported only in HTML format.

Verdict

ISG 2000 with IDP is an excellent appliance that offers a powerful combination of effectiveness and performance, flexibility and manageability, and low cost of ownership.

INTRUSION PREVENTION

SGI-2000S IPS

REVIEWED BY PHORAM MEHTA

Stonesoft
www.stonesoft.com
Price: **SGI-2000S, $31,900; SGI-200ANZ, $8,950**

The implementation of intrusion detection/prevention systems has increased considerably, in part because of improved effectiveness and the need to comply with federal and industry regulations. Stonesoft offers a strong new entry into this crowded market with its StoneGate IPS products.

Installation/Configuration B+

StoneGate's security platform is highly flexible and scalable, featuring a three-tier architecture—user interface, management and IPS (and firewall if you own it as well). Organizations can deploy clusters with up to 16 nodes.

We tested the SGI-2000S IPS sensor appliance and SGI-200ANZ analyzer device (for event correlation).

Since the appliances come with IPS engines installed, we only needed to install the three management components, which we put on a single Windows server. (Linux and Solaris OS versions are available.)

Wizard-driven installation and configuration of the management server, which can manage all Stonesoft products, is fairly simple.

There's no auto-update capability, so we recommend you download and install the latest IPS signature updates (released about once per week) as regularly as possible.

Effectiveness/Performance B+

Designed for gigabit networks, these beasts are equipped with an IPS engine that uses various techniques for misuse and anomaly detection, for effective intrusion detection, while minimizing the number of false positives.

Attacks are detected by using context-sensitive fingerprinting defined with regular expressions. For example, this expression matches any of the following patterns in the traffic: "/bin/{ash|bash|csh|ksh| sh|tcsh}." You can make it context-sensitive by defining situation elements that generate an alert only when the above expression is detected for traffic originating from untrusted sources.

StoneGate detects zero-day attacks through protocol and statistical anomaly detection techniques.

We tried multiple vulnerability scanners (Nessus, WebInspect, AppScan) and penetration-testing techniques—denial of service, gain remote shell, overflows, etc.—accompanied with evasion techniques, such as time delay and fragmentation. Most were detected without affecting the normal traffic when run in inline mode.

In sniffer mode, the sensors can respond to selected threats by sending TCP resets directly to the communicating parties or by giving a blacklisting command to a StoneGate firewall, if one exists on the network.

Administration/Management B+

The use of regular expressions makes creating new rules and custom checks easy. For example, if installing a certain product is prohibited, administrators can create a custom situation element and add an HTTP request URI with a regular expression containing the address the forbidden software uses to send or receive data.

The UI offers various situation context elements to help users write intelligent context-sensitive regular expressions to detect malicious traffic. Users can create custom situations, category tags and even workflow for certain events or sets of events

Reporting B

The management server provides extensive reporting tools for generating reports on the logged firewall and IPS events. You can create reports on log, alert and audit entries as well as statistical monitoring information. A variety of report designs are ready for use in the software and new reports can be designed and customized as needed. Reports can be exported into PDF or text.

Verdict

Stonesoft has delivered another strong product *(see review of StoneGate SG-4000 firewall appliance, February 2006)* that thwarts attacks and monitors traffic on internal networks without noticeable degradation of bandwidth.

Testing methodology: In a typical lab setup with multiple Windows and Linux machines, we sent legitimate as well as malicious traffic back and forth between the machines and the Internet through the SGI-2000S IPS.

LOG MANAGEMENT

LogLogic's LX

REVIEWED BY PHORAM MEHTA

LogLogic
www.loglogic.com
Price: **Starts at $50,000**

Although device logs can contain a wide variety of information, little attention was given to the review and management of these logs until regulations like SOX, HIPAA and PCI made it mandatory. Now companies are finding there is far more to gain from reviewing and auditing logs than just compliance.

LogLogic offers enterprise-class appliances for analyzing and archiving log data that enable organizations to achieve compliance, while offering decision support and improved availability. We reviewed the LX 2010, one of the LX family of appliances for real-time log data collection and analysis. (The ST series interfaces with NAS and WORM devices for mass storage.)

Installation/Setup B+

LX 2010 is a beast of an appliance with a 2 TB RAW (1 TB in RAID 10) storage capacity, dual 2.4 AMD Opteron processors and 4 GB of memory. The setup is as simple as it gets, supported by a hardened Linux kernel, MySQL database, Apache Web server and Java.

The only thing left for the user is mounting the hard drive in the slots in allotted order and changing the default IP address on the Ethernet interface used to access the Web-based management console, which is easily done through the GUI or CLI.

An appliance can be used to manage multiple LX and/or ST deployments. The access control feature allows you to restrict network access based on source IP address and destination port, similar to access lists used by routers or firewalls. The GUI provides controls to access different parts of the system, and menu items display according to the level of privileges granted to a specific user.

Testing methodology: Logs were obtained from Windows and UNIX servers, Cisco routers, Check Point firewalls and other networking devices generating logs in syslog format.

Configuration B+

LogLogic supports most of the widely deployed devices in the industry. At a sustainable rate of 4,000 messages per second, the LX 2010 can become the syslog and/or SNMP server for all servers and devices in the network. Logs can also be imported via HTTP, HTTPS, SCP, FTP or SFTP. Multiple log formats covering virtually all types of devices are supported—but not all log types. For instance, for firewall/VPN products with proprietary log formats, only Check Point Software Technologies, Cisco Systems, Juniper Networks and Nortel are supported. Email (Exchange) and database (Oracle) server support is also limited.

Configuring log sources is straightforward. Adding devices requires configuration changes on the source devices as well. The documentation provides step-by-step instructions for setting up the log transfer rules and frequency. We configured a few syslog devices, Windows servers using LogLogic's own open-source Lasso tool, a couple of Cisco routers and a Check Point firewall. Since most of the configuration happens on the log sources themselves, adding and setting up devices on LX 2010 usually takes less than a minute.

Reporting B+

Reporting is the most important component of this product. Two excellent status dashboard screens show the current mps rate, alerts, system performance and total message counters. Another screen shows all added devices and their message counters. The Real-Time Viewer tab shows log messages as they are received.

LogLogic offers many built-in real-time reports for access control, connectivity, database event logs, IBM i5/OS, IDS, email and Web activity.

Administrators can create keyword or regular expression searches to produce custom reports to monitor network security and health. The ability to replay old log data should prove very useful for incident response.

Verdict

LogLogic's LX 2010 offers much-needed help to companies in the areas of log review, analysis and archiving. It can help organizations not only with compliance but also with detection and prevention of dangerous events.

UNIFIED THREAT MANAGEMENT

FortiGate 3600A

REVIEWED BY SANDRA KAY MILLER

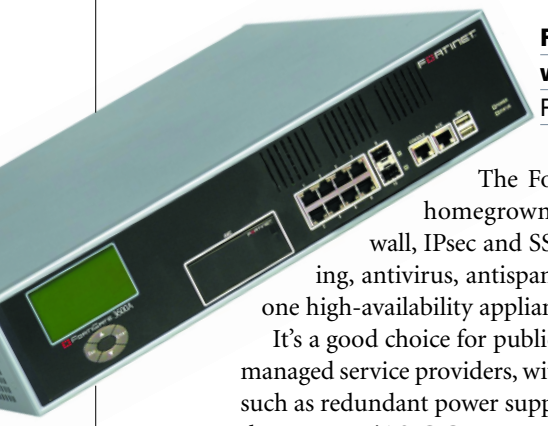

Fortinet
www.fortinet.com
Price: **Starts at $32,995**

The FortiGate 3600A rolls eight homegrown security services—firewall, IPsec and SSL VPNs, IPS, traffic shaping, antivirus, antispam and Web filtering—into one high-availability appliance built for speed.

It's a good choice for publicly facing data centers and managed service providers, with enterprise-class features such as redundant power supplies, dual-core processing, the new FortiASIC Content Processor-6, AMC network adapter expansion slots, two accelerated gigabit SFP ports and eight 10/100/1000 copper gigabit ports.

Configuration/Management B-

The 3600A can be deployed as a gateway between the Internet and private network (NAT/route mode), or on a single subnet invisible to the rest of the network (transparent mode). We chose NAT mode in order to include multiple subnets.

Using the quick-start guide, we planned our network configuration and connected to the Web-based manager in minutes. The interface is pleasantly clean and easy to navigate. Thanks to the expandable menu tree, moving through the initial setup was fairly intuitive.

For example, the VPN option expands to provide instant access to IPsec, PPTP, SSL and certificate administration. We created firewall rules, applied policies for content filtering and set up VPN tunnel associations.

Fortigate supports RADIUS, LDAP and Active Directory authentication.

Our only significant frustration was with the client software, which provides endpoint security and IPsec VPN connectivity. It was extremely slow to install and created instabilities in several instances.

Testing methodology: The FortiGate 3600A was deployed in NAT/route mode between the Internet and a simulated enterprise network. Threats specific to individual components were executed, along with behaviors and traffic denied by policy.

Policy Control B

Creating customized firewall rules, IPS signatures and adding URLs to the Web filter was straightforward.

Working primarily through the firewall, we quickly assigned numerous policies relating to network settings, logging, traffic shaping and restricting client network access based on policy compliance, such as up-to-date antivirus and IPS signatures.

However, given the extensive hardware support for high throughput (the pair of SFP connectors for optical networks), we were dismayed there was little standard policy control for VoIP. Also, there are only four IM services listed in the IM/P2P policy tab (MSN, Yahoo!, AIM and ICQ); we would have liked to seen more choices, given the explosive growth of IM clients.

Effectiveness A

We were impressed with the quality of security services on a single appliance, as well as flexibility for deployment and ease of administration. For instance, the IPS is signature- and anomaly-based, and multiple VPN technologies are included. Automatic updates and system backup and restore for multiple security services simplify life for admins and reduce the chance of human error.

Each layer of security functioned effectively when faced with common threats such as syn floods, malware, port scans and spam. Prohibited Skype traffic and potentially hazardous URLs and sites containing blacklisted keywords were blocked.

Reporting B-

Logging is outstanding. The 3600A provides three avenues for logging: local, syslog and through the FortiAnalyzer, an additional dedicated appliance for data collection and analysis from multiple FortiGate devices. The exhaustive logging was easily parsed using single-click column filtering.

Using check boxes, we set up custom email alerts for more than a dozen different events, to be sent at defined intervals. The event log is also highly customizable.

Unfortunately, there are few onboard reporting features unless the data is sent to a FortiAnalyzer, which was not included in our testing.

Verdict

Considering the costs and IT resources for managing individual products, the FortiGate 3600A offers an affordable and manageable enterprise solution.

UNIFIED THREAT MANAGEMENT

Sidewinder 2150 v7

REVIEWED BY DAVID STROM

Secure Computing
www.securecomputing.com
Price: **Varies; as tested, $35,900**

In its June issue, *Information Security* tested six Unified Threat Management (UTM) boxes; this month we review Secure Computing Sidewinder UTM, specifically the large-enterprise 2150 appliance. The new Sidewinder release was too late for the comparative evaluation, but would stack up in the middle of the pack.

Enterprise Management/Control B-

Sidewinder has a collection of different software management tools. Unlike most other UTM products, it does not have a built-in Web server but uses a Windows-based tool that doesn't run on Vista yet. One big drawback is that the product doesn't include a DHCP server for the local network; you'll need to supply your own. It took about an hour to set up.

Sidewinder doesn't allow multiple concurrent administrators to save configurations, although they can view configuration and monitor operations. It also comes with dual power supplies, which is handy if one fails. It also has two available add-in slots in the model we tested.

Daily Operations B+

We followed the same testing procedure as for the comparative review and tested how these products would work on a daily basis. While Sidewinder's IDS/IPS is wired to live inside its firewall module, it has a very flexible IPS coverage and can scan for attack signatures and behaviors. It can also explicitly detect outbound attack signatures. Sidewinder has a very useful front-page dashboard that shows alerts, CPU and memory usage, and other summary statistics in one convenient place. It is also easy to set up and change security policies.

Testing methodology: We connected the Sidewinder box on a test network with Windows XP, Vista and Apple Macintosh clients and a Windows 2003 Enterprise Server, and ran tests using Skype, AOL and Google Talk IM clients, and various security penetration techniques.

Sidewinder At-a-Glance	
Ethernet Ports	8*
Max attachment size for antivirus scans	User selected < 1 GB
AV supplier	Sophos
Content filter supplier	Own
IDS: Patterns or behavior	Both
IDS in or out of firewall?	Inside only
VPN types	IPsec
Other ports scanned	IM, P2P, VOIP, SQL
Web App firewall	Extensive
Authentication options	Radius, LDAP, AD
Multiple concurrent admins?	No

*Includes two gigabit Ethernet ports that came installed on our test unit

Authentication & Security B

Sidewinder sets up most of it security policies for each network interface, but has separate controls for content filtering, antivirus and antispam modules that are applied across these interfaces. Sidewinder offers connections to a variety of authentication servers, including Radius, LDAP, Active Directory and iPlanet servers. It includes an IPsec VPN only—no SSL.

Feature Module Integration B

Sidewinder uses Sophos antivirus scanning but also has its own SmartFilter content filtering engine. However, SmartFilter requires a separate Windows-based administration and configuration tool and its own obscure setup with nested sub-menus. This is because Secure Computing sells this as a separate product that can be run on other vendors' firewalls. We'd like to see it completely integrated into the main console. One nice feature is the ability to run several antivirus scanners in parallel on the same box to balance the processing load. A maximum 1 GB file attachment can be scanned.

Although Sidewinder was able to easily block Skype with its default settings, it doesn't have explicit protection rules for other IM/P2P protocols.

It does extensive port scanning, including ports used for VOIP, IM, P2P, SQL server and Citrix applications. It also protects against common Web server attacks, such as SQL injection and cross-site scripting.

Verdict

Sidewinder offers solid security features and is easy to set up and manage. Its strengths are extensive IDS/IPS and antivirus scanning features; its biggest weakness is its separate content filtering module.

FIREWALL

SonicWALL TZ 180W

REVIEWED BY JOEL SNYDER

SonicWALL
www.sonicwall.com
Price: **$700 for 10-user version with standard O/S**

SonicWALL has been a major player in the SME firewall market for as long as there has been a market, helping to define what SME firewalls should look like, cost and do. The SonicWALL TZ 180W UTM continues SonicWALL's tradition of products designed and sized for the small business.

Ease of Use A-

Small-office firewalls are often difficult to use, as developers jam more and more features into poorly designed user interfaces. The TZ 180W, from opening the box to final deployment, has a sophisticated and refined feel. Although we ran into a bug in the initial deployment wizard and some misdirected online help, every other part of the system was easy to use.

The TZ 180W with SonicWALL's SonicOS Standard is stripped down from an enterprise firewall, which means that some features, such as NAT, come as "one size fits all." However, SonicWALL has chosen an excellent subset of features—more than most network managers will need—for the small office version.

Security Features B

SonicWALL has set a very attractive price for its Comprehensive Gateway Security Suite, a UTM add-on service for the TZ 180W that includes software support along with content filtering, antivirus, antispyware and intrusion prevention subscription services. Consider this service a must-have for any TZ 180W, as it unlocks the full potential of the product.

We found content filtering, antivirus and antispyware effective. Intrusion prevention was less effective. SonicWALL has loaded what looks like a subset of Snort's signatures into the IPS, but not all the intelligence. For example, the TZ 180W alerted on a possible port scan, which was actually the box communicating with SonicWALL's own Web servers. It missed all but one of our outbound attacks.

Having all those IPS signatures configured to protect against inbound attacks against servers on a device primarily designed to protect end users doesn't make a lot of sense. A better approach is signatures focused on end-users, such as malware protection or browser-focused overflow attacks.

Performance A

The TZ 180W provides significantly greater performance than the earlier TZ series. Because we had a 10-user unit, we were unable to push the firewall to its advertised limit of 90 Mbps using typical Internet traffic mix. Our system ran out of CPU at about 27 Mbps. When we turned on all the TZ 180W's security services, goodput was about 9.7 Mbps, close to SonicWALL's advertised speed of 10 Mbps.

It should keep up with DSL and cable modem connections, but don't be tempted to run LAN backups or file sharing through it with security services turned on.

Wireless B

The TZ 180W has a built-in, dual-antenna 802.11b/g wireless access point. Although this only adds $95 to the price, it was disappointing that we couldn't get 802.11a or 802.11n on a new system. Wireless security is limited, but we were impressed that we were able to set up WPA2 with RADIUS authentication in seconds.

However, you can't easily set up the wireless so that insiders and guests can use it, so it's best to pick one set of users. With a cool set of features aimed specifically at guest users, it will probably fit best to give guests secure temporary Internet access without allowing them access to your internal network.

Verdict

The TZ 180W is an outstanding small-office and home-office UTM firewall, offering good value and a broad suite of gateway security services.

Testing methodology: We installed the SonicWALL TZ 180W on a production network and used both Windows and Mac laptops to validate security services. To test performance, we used Spirent's WebAvalanche and WebReflector testing devices to deliver a typical mix of Internet traffic sizes and pages.

UNIFIED THREAT MANAGEMENT

Firebox X 1250e

REVIEWED BY BRENT HUSTON

WatchGuard Technologies

www.watchguard.com

Price: **Ranges from $2,290 (plus $4,420 for UTM bundle for the Firebox X550e) to $3,790 (plus $7,400 for the UTM bundle for the 1250e)**

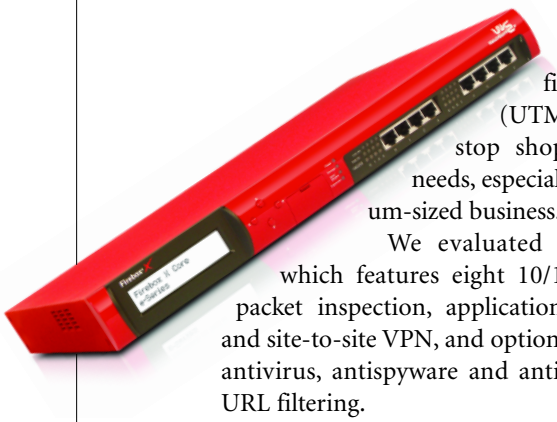

WatchGuard's unified threat management (UTM) appliances are a one-stop shop for border security needs, especially for a small- to medium-sized business.

We evaluated the Firebox X1250e, which features eight 10/100 interfaces, stateful packet inspection, application proxies, remote-user and site-to-site VPN, and optional modules for gateway antivirus, antispyware and antispam protection, plus URL filtering.

Configuration/Management A

Setup is straightforward. We followed the included quick-start guide to get the device working in less than an hour.

The management interface is one of the best we've seen. The rules setup is logical and does not require knowing any cryptic languages. The proxies and other features are well integrated, and can be configured and enabled/disabled easily for each rule.

Effectiveness B+

The firewall immediately stood out on its own, thanks to the ease of setting up rules. Rules are granular, and you don't have to worry about putting them in the correct order—Firebox takes care of that for you.

Application proxies for HTTP, FTP, SMTP and DNS, and a generic TCP proxy allow the firewall to inspect traffic and deny or allow the request based on your policy. For example, we set up a rule in the FTP proxy to deny "get" requests. The rule worked as intended and wouldn't allow any file downloads. The controls are granular; you can, for example, block the download of certain extensions, and block or allow HTTP requests or content types in the HTTP proxy.

Firebox's IPS capabilities are strong. By default, it will block anyone trying to port-scan or send suspicious packets through the device; our port scans got us quickly blacklisted. We set up a Web site behind the Firebox and attacked it using Metasploit, but all our attacks were stopped.

The antivirus module is based on open-source ClamAV, which we've found to be a competent antivirus. One issue here is that you can only use the antivirus through the HTTP and SMTP proxies, so, for instance, there is no way to scan files going through the FTP proxy.

The VPN uses IPSec and PPTP, supporting remote user and branch connections. Back-end authentication can be implemented through Firebox itself, RADIUS, Active Directory, LDAP or RSA Security's SecurID.

The VPN client only works with Windows—a restriction for some shops, which can use the less secure PPTP option.

The antispam filtering, provided by Commtouch, picked up spam that even our tuned SpamAssassin filter missed.

While Firebox's URL filtering module features many categories and blacklisted sites, it was possible to get around some by using the IP address.

Reporting B+

Reporting capabilities are good, but you can only export the results in HTML and NetIQ formats (but it derives the reports from XML data, so importing it elsewhere is not out of the question).

However, the reporting gives you an excellent breakdown of device statistics, traffic stats, and IPS alerts, and a report of hits on any rules you have in place (such as users trying to visit blocked Web sites.

There are also extensive real-time monitoring capabilities including traffic and bandwidth monitors, device statistics (memory usage, processes running) and a list of authenticated users.

Verdict

Despite some minor flaws, the Firebox X series is an excellent UTM deal, with its low entry price, terrific firewall and routing capabilities, and top-notch filtering services.

Testing methodology: We tested the Firebox X 1250e protecting two internal networks and a DMZ that included a Web server, FTP server, SMPT and POP server.

SSL VPN

SonicWALL SSL-VPN 4000

REVIEWED BY SANDRA K. MILLER

SonicWALL

www.sonicwall.com

Price: **Starts at $6,995**

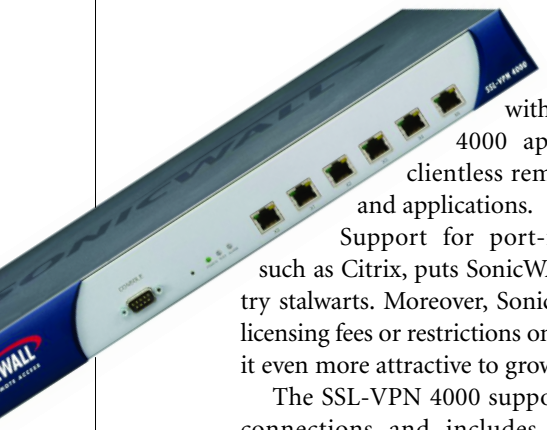

SonicWALL steps up to the enterprise market with the affordable SSL-VPN 4000 appliance, offering secure clientless remote access to files, shares and applications.

Support for port-forwarding applications, such as Citrix, puts SonicWALL in league with industry stalwarts. Moreover, SonicWALL has no per-tunnel licensing fees or restrictions on concurrent users, making it even more attractive to growing organizations.

The SSL-VPN 4000 supports up to 200 concurrent connections and includes support for two-factor authentication, such as RSA Security tokens.

Configuration/Management A

Using the administrator's guide, we were able to log on to the appliance within minutes. All major browsers and OSes are supported.

SonicWALL's familiar easy-to-use Web-based console gave us instant access to major features, each offering a subset of functionalities.

For example, the network tab includes access to the interfaces, DNS, network paths, host resolution and network objects—all straightforward. After basic network settings, we quickly set up services to which we would provide secure remote access: HTTP, HTTPS, terminal services (Java and ActiveX), VNC, FTP, Telnet, SSH (versions 1 and 2), file shares and Citrix Portal. Objects can be defined by a solitary IP address or a network segment.

Setting up individual users and groups was equally effortless. The 4000 supports user authentication via LDAP, Active Directory, NT and RADIUS.

Testing methodology: We tested SonicWALL SSL-VPN 4000 on a simulated Windows-based enterprise network behind a third-party firewall. Remote access was tested from a variety of laptops and remote machines, running an assortment of operating systems and Web browsers.

Policy Control A

We were impressed with the granular policy control, which let us assign access privileges at the user, group and global levels.

We were able to delineate authentication to our AD server, so that domain members were automatically assigned the policies and access privileges of their associated group.

Policies are granular and highly accessible. A single window enabled us to assign general settings, such as enabling single sign-on using SSL VPN credentials, creating individual policies for network objects, IP addresses and ranges, and server paths, such as for Citrix. In addition, we were able to set up detailed login policies, such as one-time passwords sent via email and logins from specific IP addresses or defined browsers.

Effectiveness A

We were extremely satisfied with SonicWALL's interoperability, including the product's Web access to email, files and Web-based applications.

Additionally, the NetExtender thin client can be automatically downloaded and installed to provide access to email using client software installed on remote machines and non-Web-based applications, such as CRM systems and proprietary software.

We simulated a variety of scenarios that tested the granularity of policy features, such as allowing global access to email while limiting access to specific file shares and applications.

Reporting C+

The VPN lacks a comprehensive view. Interface statuses are under the system tab, active user sessions are under the users tab, and viewing events requires going to the log tab.

Logging is very basic, although it supports syslog and can email logs and alerts to a single address.

Verdict

SonicWALL SSL-VPN 4000 is an affordable and capable appliance for mid-sized enterprises.

INTERNET SECURITY

HotPick INFORMATION SECURITY®

FaceTime Internet Security Edition

REVIEWED BY SANDRA KAY MILLER

FaceTime Communications
www.facetime.com
Price: **Starts at $7,125**

FaceTime's Internet Security Edition tackles the greynet challenge of sifting Web traffic to differentiate between legitimate and unauthorized use of real-time communications applications such as instant messaging, Web browsing and VoIP.

The combination of RTGuardian (RTG), a hardened Linux rack-mounted inline appliance, and Greynet Enterprise Manager (GEM), a Windows-based server, delivers security at the perimeter and endpoint by identifying malware, spyware, adware and unauthorized traffic.

Configuration/Management B

Despite well-written documentation, setup took extensive preparation and knowledge of Active Directory and domain credentials.

RTG enforces policies and ties into GEM, which provides centralized management and reporting through a secure Web interface.

GEM automatically discovers endpoints by querying the primary domain controller. Administrators can also specify a range of IP addresses and discover endpoints through ping and Windows Management Instrumentation. In both cases, GEM failed to detect several Windows desktops and all our non-Microsoft machines.

Effectiveness A

RTG controls traffic at the gateway, performing URL filtering and managing greynet applications. What really impressed us is FaceTime's approach to protecting desktops against malware, spyware and adware.

When RTG identifies malicious behavior, it feeds the data to GEM, which deploys a temporary client to clean the machine and scan for additional infection. It inoculates the machine, using ActiveX kill bits and Windows software restriction policies, which prevent the code from executing again. This feature stopped spyware cold, despite our repeated attempts to reinfect the machines.

Policy Control A

We were able to set global and granular policies for Web browsing, Web mail, IM, file sharing, streaming media and VoIP.

For example, we allowed the use of certain public IM clients while prohibiting others. No one on our network was permitted to use P2P applications, and Skype was accessible only to the sales group. We could schedule automatic scans, spyware removal and inoculation. Policies can be assigned according to multiple criteria such as IP address, host, user, domain and operating system.

Comprehensive URL filtering categories let us turn off access to generally prohibited and productivity-draining sites (porn, gambling, shopping, news, travel). Custom policies can be set by users, groups, location, file extension and content.

Reporting A

GEM provides centralized reporting and logging; the real-time reporting dashboard includes activity blocked by RTG as well as GEM activity tracking infected computers.

In addition to providing statistical analysis for everything from infections to policy violations, FaceTime offers a variety of executive and auditing reports. Administrators can quickly see the rate of spyware infections and spot trends as to what users/systems were most vulnerable and often infected, while auditors have access to detailed information about data transferred via various Web-based channels.

IM reports can be split into events and usage, providing detailed, critical information, such as transferred files. Reports can be sent via email or exported to an FTP server for automated distribution.

Verdict

FaceTime offers an effective and affordable solution to manage, control, secure and provide policy compliance around Web-based applications.

Testing methodology: The RTGuardian appliance was deployed on the span port of a DMZ switch; GEM Server was installed on a Windows 2003 Server. Numerous applications were tested, using malware including spyware and adware.

From USB drives to MP3 players to DVDs, portable storage media are an end user's dream and a security manager's nightmare.

Gone in

We evaluated six device control tools that can help you rest easier.

BY SANDRA KAY MILLER

Did your business just walk out the door? Our mobile workforce can steal or lose sensitive data quickly and without detection, from a software developer sneaking out gigabytes of valuable source code on his iPod to an executive's wireless-enabled laptop being sniffed at the local coffee shop. Think about all the ways we move and store data on mobile devices: USB ports, which support a multitude of portable storage devices, including flash drives, portable hard drives, printers, and music and video players; FireWire, PCMCIA, serial and parallel ports, CDs/DVDs, tape drives and even the lowly floppy drive. Add unprotected WiFi, Bluetooth and Infrared (IrDA) connections, and you have a real security nightmare on your hands.

Photographs by JOHN KUCZALA


a Flash

It wasn't long ago that security administrators controlled access to USB ports with epoxy or caulk and physically disabled onboard wireless. Now, however, instead of trying to ban use of portable storage devices and wireless connections, organizations can select from a fairly new but effective group of products that give them granular policy-based control over their use. Device control products can help balance productivity with security by allowing administrators to centrally authorize and monitor endpoint devices.

In a head-to-head review, *Information Security* examined six device control products, all of which provide centrally managed granular control over ports, interfaces and storage devices: DeviceLock 6.0 from SmartLine, Sanctuary Device Control 4.0 from SecureWave, Endpoint Access Manager 3.0 from ControlGuard, DeviceWall 4.5 from Centennial Software, Safend Protector 3.1 from Safend and Protect Mobile from Workshare.

Each product was graded based on its ease of installation and configuration, policy, tampering resistance, port and device control, encryption support, performance, and monitoring, alerting and reporting. Overall, we found all the products performed as advertised, but there are enough differences to consider when choosing a portable endpoint data control solution *(see "Making the Grade," p. 51).*

About this review

Centennial Software's DeviceWall 4.5


ControlGuard's Endpoint Access Manager 3.0

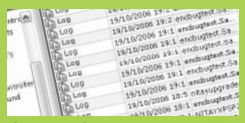
Safend's Safend Protector 3.1

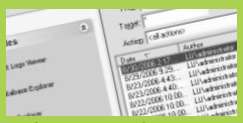
SecureWave's Sanctuary Device Control 4.0

SmartLine's DeviceLock 6.0

Workshare's Protect Mobile

Information Security deployed six portable storage device control products in our test lab.

All the products were tested in a Microsoft Windows environment with Active Directory, although some supported Novell. All the products utilized either an embedded or external version of a SQL database. Our testing environment included wired and wireless network connectivity with both desktops and laptops, supporting an array of portable storage devices including USB flash drives, FireWire external hard drives, CD-RW drives and floppy disk drives. Our testing also included PDAs and serial docking stations, smart phones with Bluetooth connections, PCMCIA wireless adapters and multifunction printer/scanner/fax/copier machines on both USB and parallel ports.

Concentrating on real-world scenarios, we blocked devices such as portable music players and storage devices (flash drives, FireWire drives) while allowing legitimate peripherals including keyboards, mice, printers, faxes and scanners. Drilling down into the granular policies, we set CD/DVD drives to read-only and disabled Bluetooth and IrDA connections, while allowing WiFi use.

Multiple attempts to introduce devices contrary to policy were performed using a variety of devices and connections, including portable storage devices infected with known malware including worms, Trojans and keyloggers.

—SANDRA KAY MILLER

Getting Started

All the products we tested have similar architectures—server, console and client/agent—based on a Windows platform, although each included support for Novell's directory services. We deployed each product on an identical simulated enterprise network *(see "About This Review," above)* using numerous desktops and laptops, supporting multiple ports and removable storage devices.

Centennial's DeviceWall was the easiest product to install, since it requires only two components—the Control Center and the Client Service. With a half-dozen ways to roll out the client, DeviceWall got our top vote for installation and configuration.

ControlGuard and DeviceLock have similar setups, consisting of a server, client and multiple Windows-based ways to administer the product (Active Directory, MMS, SMS, GPO). With two different client agents—active and passive—ControlGuard gave us more to consider during setup. We deployed both types of agents and concluded that this aspect of ControlGuard should be simplified with a single agent that could perform in either or both modes.

DeviceWall and DeviceLock had easy install wizards that walked us through setting up initial permissions and policies. DeviceLock's wizard allowed us to set permissions for ports and devices, getting us running quickly.

Installing and configuring Workshare Protect Mobile took the most effort, because it is part of an enterprise suite of three components. It delivered comprehensive endpoint protection, but didn't provide the depth of granularity or functionality as the other products.

The installation of SecureWave was the most difficult, because of its four compo-

nents—a database server, an application server with two subcomponents, the management console and the client. We also encountered several client deployment issues that required extra time reconfiguring our firewall.

Policy Configuration/Enforcement

Ultimately, everything boils down to policy and enforcement and performance. Policy granularity is a driving factor in each of these six products. For portable storage devices, our testing revealed nearly identical features, including monitoring and control over reading, writing and blocking.

Policies were determined by device types and classes, ports, connections, machines and users. With all the products, we could set up who could use what device/port/connection and when.

The policy options available are so plentiful, it's easy to get overwhelmed and confused. We found it was easier to start with our global policies and work to more detailed policies, such as those for individual users. We were also able to set different policies for the same user/computer determined by online/offline status. That means when a mobile user returns to the office and logs in to the domain, wireless interfaces can be turned off, and corporate asset protection, such as file filtering, engaged.

All the products allowed very fine-grained policy, mainly through whitelists—the more granular the policies a product supports, the better the controls. DeviceLock provided the most detailed assignment of authorized devices. For example, we were able to allow a single FireWire portable hard drive based on its serial number. The exceptions can also work in reverse; for example, you can shut down access for terminated employees or limit devices to read-only.

We liked how SecureWave's Sanctuary Device Control comes out of the box with a default deny-all policy. No data was allowed to be transferred to external storage devices until we set up authorization. Allowing only what you authorize—instead of trying to blacklist what you don't—is sound security policy.

SecureWave has a number of ways to keep tabs on traffic, including data transfer throttling and file type filtering. For example, we set policies that limited file types to Microsoft Office files no larger than 5 MB. Regardless of how we tried to save CAD files—both less than and in excess of our size limit—to flash drives, portable hard drives or write to CD, we were unable to do so.

ControlGuard earns kudos for recognizing that many mobile workers also connect directly to the corporate network. We easily set up two distinctly different policies, offline and online. We simulated a common problem that occurs when mobile workers connect their WiFi-enabled laptops directly to the corporate network—they still have a live wireless connection. For our testing purposes, when laptop users logged on to the domain, their WiFi adapters were disabled.

ControlGuard addresses another real-world scenario, exercising control over multiple users logging on to a single machine or a single user with access to multiple machines. This is where a firm understanding of policy hierarchy is required. For example, a user having rights to a USB port on one machine doesn't necessarily mean he has the same rights on another.

For organizations that want to further enforce policy through enterprise-class management systems, ControlGuard's Endpoint Access Manager is designed to integrate with third-party products like HP OpenView and CA Unicenter.

Safend offers similar policy control through role-based access and prohibiting simultaneously enabling multiple networking protocols. One feature that really caught our attention was the ability to easily print a summary of our entire policy anytime. This means corporate policy can be posted or viewed by management, which doesn't have access to

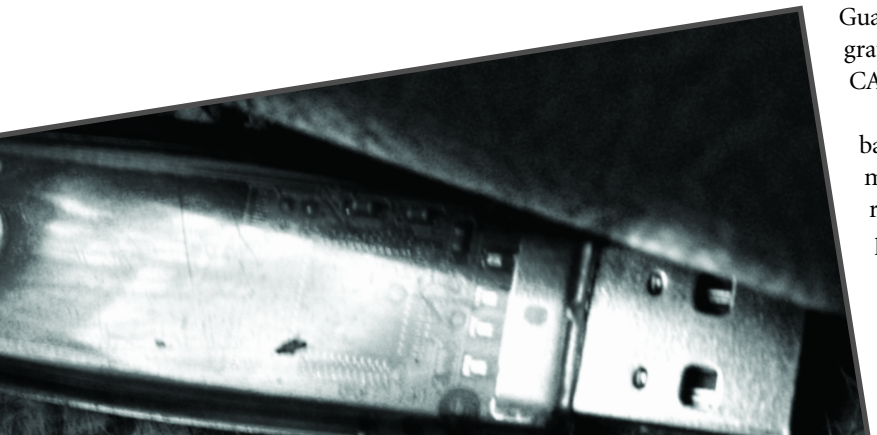

Encryption Gains ground

Data encryption has long been a strong security technology, but its use has been generally limited because of the complexity of implementing and maintaining it, as enterprises wrestle with thorny issues such as a key management and security.

That's all changing because of regulatory requirements and the exposure of data through Internet-facing applications. Nowhere has this become more evident than with the ubiquitous use of high-capacity portable storage media. Five out of six solutions *Information Security* tested for this review have integrated automatic forced encryption capabilities into their products.

Although the use of automated encryption for portable storage media is solving numerous security issues, there are still significant challenges to address.

Nate Lawson, senior researcher at Cryptography Research (www.cryptography.com), a security consulting and technology licensing firm, points out that there is plenty of room for improvement with the widespread use of encryption for storing information.

"How do I make sure I never lose or destroy that key, because if I do, it's like losing the entire set of data," Lawson says. "I won't be able to decrypt it again."

While there are lots of standards for encryption (AES, DES, 3DES, etc.) and protocols (SSL), there is little standardization for key backup.

Lawson sees this as a potential pain point, especially in M&A scenarios and because of the speed at which technology becomes obsolete.

Before organizations begin routinely encrypting portable storage media through solutions such as the ones we tested, they need to examine the lifecycle of the data being stored and ensure access to decryption tools, such as keys and software.

—SANDRA KAY MILLER

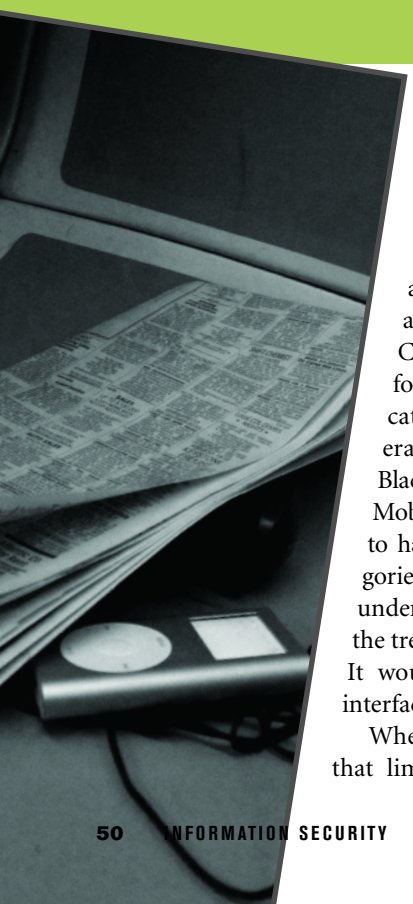

the console, but needs access to security policies.

Overall, DeviceWall's policy configuration was the least intuitive of the products tested, although the Master Policy tree accessed through the Control Center provided a clean interface for configuring 16 different device categories, including digital cameras, scanners, smart phones, and BlackBerry, Palm OS and Windows Mobile devices. We would have liked to have seen all the individual categories for mobile handheld devices under a single high-level branch on the tree, instead of each given its own. It would make the Control Center interface much less cluttered.

When building complex policies that limit or deny the availability of computer resources, there are bound to be exceptions to the rules. DeviceWall was our favorite product for bending the rules. It let us assign temporary access for up to three devices either for the current Windows session or by start time and duration. Even when we weren't connected to the network or Internet to push out a change in policy, DeviceWall gave us the option to generate a key that could be verbally exchanged or sent via text message over a mobile phone to provide temporary access to the restricted port or device.

DeviceLock's exception to policy functions similarly to DeviceWall's temporary access, but lacked the granularity to assign any length of time, giving only the option to use the restricted resource during that particular Windows Session.

Safend allows for the temporary suspension of the client, even when the computer is offline.

Tamper-proofing

We attempted to circumvent our installed clients through a variety of methods. Often, installed components can be sidestepped by local users who have administrative rights to their machine. Even with local admin rights, we were unable to modify or remove any of the installed clients.

Tens of millions of USB flash drives are sold every year, and you can bet some are going to be lost or stolen, sometimes with sensitive data. DeviceWall was our pick for the lost flash drive scenario. When we inserted a USB flash drive into a bare-bones laptop running no device control client, we received the message that our drive was not formatted and asked if we would like to format it. Had the drive contained confidential information, the cost associated with losing the data to the wrong entities could be devastating, but thanks to DeviceWall, less than a minute after plugging in the uncontrolled drive, it was wiped clean.

We also addressed the issue of theft, loss and tampering of removable storage devices and media through the products' use of encryption. DeviceLock, which was generally outstanding in other areas, was the only product in our testing that did not support any type of encryption, which brought down its overall grade.

SecureWave set the bar with two different types of encryption—centralized, which allows administrators to set the requirements, and decentralized, meaning an authorized user can decide when to encrypt. Additionally, you can export keys to a file or to the portable device for

Making the grade

Vendor	ControlGuard Endpoint Access Manager 3.0 Starts at $10 per seat www.controlguard.com	SmartLine DeviceLock 6.0 Starts at $35 per seat www.devicelock.com	Centennial Software DeviceWall 4.5 Starts at $25 per seat www.devicewall.com	Safend Protector 3.1 Starts at $32 per seat www.safend.com	SecureWave Sanctuary Device Control 4.0* Starts at $45 per seat www.securewave.com	WorkShare Protect Mobile Starts at $10 per seat www.workshare.com
Installation & Configuration How easy was the product to install and set up? 10%	B	B	A	B	C	B
Policy Configuration Level of granularity and ease of deployment. 20%	A	A	A	A	A	A
Loss, theft, tampering Client and portable storage security. 10%	A	A	A	A	A	B
Port & Device Control What does the solution cover? 10%	A	A	A	A	A	B
Encryption How well is it supported? 10%	B	N/A	C	B	A+	C
Performance How did the product stack up against testing? 20%	B	B	B	B	B	B
Monitoring, Alerting & Reporting 20%	A	A	B	A	A	B
The Verdict	**A-** Comprehensive control through client-based agents; provides complete control over all I/O devices including PCI, ISA and optical devices	**B+** Powerful shadowing feature lets companies know exactly how their critical information is moving and being stored. Would like to have seen support for encryption.	**B+** Easy to use, great reporting; includes encryption but doesn't provide the policy granularity or alerting features of competing products.	**A-** Delivers comprehensive visibility and control over endpoints, devices and network interfaces.	**A-** Offers complete I/O control over all devices on Microsoft Windows and Novell networks.	**B** Focuses on mobile devices such as laptops and PDAs.

*SecureWave's Sanctuary Device Control 4.0 is part of the Sanctuary Suite.

access to encrypted media offline, although we felt that this compromised the security of the portable storage device. SecureWave offers the strongest encryption, with AES 256.

DeviceWall offers two different ciphers—AES and Blowfish—in both global and individual user key models. For instance, a company might require its HR employees to automatically encrypt all data transmitted via WiFi or saved to portable media. However, encryption is only available for use with USB flash drives. On the plus side, DeviceWall allows you to easily back up the Global Key, so data can be retrieved if the key is lost.

ControlGuard also provides encryption for secured USB drives. We liked its "self-destruct" feature, which limits the lifecycle of the data accessible on the drive.

Workshare Protect Mobile provides the most flexible client-side encryption through PGP based upon content. Once files have been identified as requiring additional security, they are automatically encrypted.

Safend's encryption is the most transparent to users. We were able to use the same encrypted USB drive on all the machines on our network with the Safend client installed without ever realizing the device had been encrypted. Of

course, when we attempted to use the drive in a non-Safend computer, we were unable to access the drive.

One big worry with encrypted files on portable media is the decryption software won't be available when needed. Safend had the forethought for just such a scenario and includes a Home Decryption Utility that allows authorized users to access information on encrypted devices when the Protector Client is not present.

Wireless Control

Wireless covers a lot of territory on today's mobile devices. All the products we looked at included comprehensive control over WiFi, Bluetooth and IrDA interfaces.

Since its introduction, there has been a lot of hand-wringing over WiFi connections. Administrators disable onboard wireless, but still have to worry about an employee using their own inexpensive PCMCIA wireless adapter so they can hook up at home or a hotspot.

Administrators are just catching up to smart phones and PDAs, which are increasingly taking advantage of Bluetooth technology for file transfer and synchronization with laptops. An inexpensive USB Bluetooth adapter can quickly connect a PC to a Pocket PC.

And let's not forget about IrDA. Not as powerful or popular as WiFi or Bluetooth, infrared personal area network connectivity still presents a vulnerability.

Safend clocked in with the best control for WiFi, based upon MAC addresses, SSID and network security levels.

The remaining products didn't provide as much control as Safend, but they all provided basic permit/deny wireless interface blocking functionalities that identified all wireless interfaces regardless of type. For instance, we set policy to deny all WiFi with a laptop containing an onboard wireless adapter. As we added PCMCIA and USB wireless adapters, they too were disabled despite those ports not having any deny policy assigned to them.

Auditing and Reporting

While the majority of our testing was devoted to the verification of security features, in today's regulatory environment, a robust auditing feature can be just as critical as security.

The most comprehensive monitoring feature for this purpose is shadowing, which is the ability to record all data transferred to and/or from a device or port. DeviceLock and SecureWave both support shadowing.

During our testing, shadowing allowed us to capture all data sent to specific devices, including our printer/scanner/copier/fax machine. How many companies actually monitor the information sent over a fax or documents that have been scanned? Low-tech crimes are often overlooked.

The only drawback we could see with data shadowing was, ultimately, data storage. A large enterprise could generate an enormous amount of data.

Safend uses file logging; while not as robust as shadowing, it lets administrators track what files are being accessed, moved, deleted, created and modified.

Safend took the honors for the most useful logs, with excellent information for forensic investigations.

We also liked DeviceWall's detailed Policy Change Logs, which record all the policy changes made and provide comprehensive connection reporting. On the other hand, we found the graphical Audit Log Reports and Acceptable Risk meter of little use to a security professional.

Spotlight on the Endpoint

Mobility and portability make data protection a far more complicated problem than it once was. They've given this product market traction it wouldn't have seen just a couple of years ago. As the products mature, they feature improved reporting and central management capabilities, and, in most cases, combine encryption with device control for stronger endpoint security.

We're starting to see similar capabilities in more comprehensive security and data protection products. Expect to see these products extend their feature sets or get folded into broader products, as larger companies continue their pattern of acquiring key new technologies.

Time and experience have taught the security community that going to the root of a problem will often save time and money. Priced on a per-user, annual basis, for larger enterprises, these products can become pricey on top of existing security solutions licensed on an annual basis, such as antivirus/antispyware. Nevertheless, endpoint security has clearly moved center stage, and many corporations are going to do what it takes to protect their data as it moves out into the world.

Technical editor Sandra Kay Miller is a frequent contributor to Information Security. *Send comments on this article to feedback@infosecuritymag.com.*

UNIVERSAL

IS YOUR BUSINESS READY TO ROLL NETWORK SECURITY INTO A SINGLE PLATFORM?

WE EVALUATE SIX LEADING UTM APPLIANCES TO HELP YOU PUSH THE RIGHT BUTTONS.

REVIEWED BY DAVID STROM

Photograph by JOHN KUCZALA

CONTROL

UNIFIED THREAT MANAGEMENT IS A GROWING, competitive field, with more than a dozen vendors. The idea is to consolidate your security appliances into a single box and manage an integrated protection profile for your corporate network. While especially appealing for companies with several branch offices without any resident IT or security staff, the implementation isn't perfect for a corporate-wide deployment, primarily because of limits on how you are allowed to administer the component security applications.

We asked vendors to deliver a product that could act as a firewall and virtual private network gateway, and protect our test network against attack with a minimum of four defense mechanisms—antivirus, Web content filtering, intrusion prevention and antispam.

ABOUT THIS REVIEW


Astaro Security Gateway 320

Check Point Software's UTM-1 2050


Fortinet's FortiGate-1000A

IBM Internet Security Systems' Proventia NMS MX5010

Juniper Networks' SSG 550

SonicWALL Pro 5060c

We connected each UTM box on a test network with Windows XP, Vista and Apple Macintosh clients and a Windows 2003 Enterprise Server running Microsoft's IIS Web server.

Each UTM box was configured with two interfaces—a local network with a DHCP server enabled, and an external network connecting to our DSL modem. We set up firewall and intrusion rule sets, ran Outlook Express POP email clients, and used Skype, GoogleTalk and AIM messaging sessions.

We also connected to a WebDAV server to share files over the Internet. We connected to each product's built-in Web management server using both Firefox v2 and Internet Explorer 6 and 7. We also used SSH to perform command-line configuration tasks when necessary.

—DAVID STROM

We reviewed six UTM appliances in a head-to-head evaluation: Astaro Internet Security's Astaro Security Gateway 320; Check Point Software's UTM-1 2050; Fortinet's FortiGate-1000A; IBM Internet Security Systems' Proventia Network Multifunction Security MX5010; Juniper Networks' SSG 550 and SonicWALL's SonicWALL Pro 5060c.

We examined log files and configuration reports to determine how each appliance stacked up in enterprise management and control, daily operation, authentication and policies, and feature integration.

All of the products sell for between $12,000 and $18,500. But getting specific price configurations isn't easy, as each product has a complex range of user and feature licenses. Further confounding the pricing issue is that you will need to match the capacity of the product with the expected network traffic it will protect. We tried to compare appliances that had a similar number of network ports and capacity for a 1 Gbps external network connection.

We asked vendors to send us the boxes with the highest throughput possible and geared toward the largest networks. When we did our tests, we turned on all of the security modules—in the real world, this will severely limit their overall performance and is something to consider when deploying these products. However, we did not test performance. This is because testing performance is fraught with all sorts of issues. Either you test with synthetic clients to generate phony traffic so you can compare how different products respond on the "same" artificial lab network, or you do your tests on a live network and hope that the insights gained with your actual conditions are worth the loss of having the comparable traffic data. As a potential purchaser, you should match throughput specs with what you ultimately need on your network.

1 ENTERPRISE MANAGEMENT AND CONTROL

We examined how each product is managed, typically with a Web browser, and the various administrative roles that can be performed concurrently. We also looked at how particular functions are licensed, and how threat signatures are updated. Because these products handle a variety of security tasks, ease of setup is important, and being able to delegate and divide administrative roles is also critical.

CONFIGURATION. All of the products, except Check Point, are primarily configured by connecting to their built-in Web servers. Check Point actually has three configuration interfaces—command line, a Web-based initial configuration tool for basic tasks, and its SmartView Monitor Windows-based administration tool *(See Check Point screen shot, below)*. Unfortunately, you'll need to be familiar with all three. For example, you have to go to the command line interface to set up a DHCP server on an internal network.

Some of the Web interfaces are more logically designed

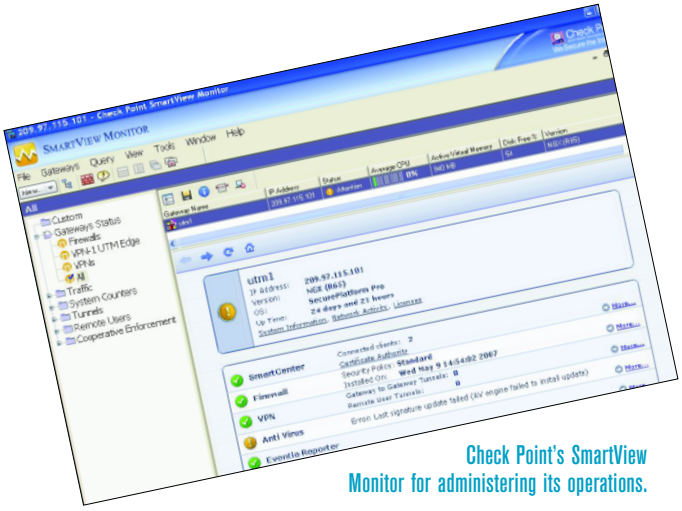

Check Point's SmartView Monitor for administering its operations.

than others. For example, Astaro, IBM ISS and Fortinet separate the functional modules—separate menu trees for antivirus and IDS, for example—and logically lay them out. Juniper has the poorest interface of the six, because its commands and controls are buried several levels down or require operators to visit multiple pages to set up even the simplest procedures, such as changing one of the antivirus settings. SonicWALL's interface is just a little better than Juniper, hiding many of its UTM features under a single "security services" menu tab.

Setting up the IBM ISS box took about an hour, and Check Point took several hours. The others were somewhere in between. While this may not be terribly important if you're installing a single box, it will add up for large deployments.

IBM ISS stood out from the pack with superior defaults, such as setting up internal network routes and activating features at the click of a button. This default-driven approach could be a bit problematic if your tastes run to doing something more sophisticated. For example, most of the other UTM appliances could handle connections to a WebDAV server for sharing files; with IBM ISS, we needed to set up a special firewall policy to allow this traffic. Nevertheless, this was a minor inconvenience—not enough to keep IBM ISS from getting the clear top grade in overall enterprise management.

LICENSING, UPDATING. Each product has intricate licensing and signature file update issues, mainly because customers will purchase varying configurations, feature sets and user counts. None of the products did a particularly good job troubleshooting licensing errors; Check Point and Juniper had the most complex and unintuitive licensing procedures. In fact, we had trouble with our Check Point licenses even after its engineer spent several hours on site setting up our box that turned up a bug. The other products make installing and upgrading licenses, and updating threat signatures, far easier.

IBM ISS makes this process a snap; it consolidates all of its updates for antivirus, IDS and firmware in a single screen. You can set it to check for updates automatically on a schedule. The others are more complex; you will have to visit multiple screens or do more than just push a single button to update everything.

ADMINISTRATION. Consolidated security administration is a key value proposition for UTM. However, getting to this consolidation won't be easy. Because these products cover a wide range of protection methods, they need to have the flexibility to be operated by multiple administrators.

Fortinet, Juniper and Astaro can handle multiple concurrent administrators and immediately post any configuration changes to their boxes in a "last one wins" scenario: This means that any intermediate changes will be ignored, which isn't ideal and means one person needs to have ultimate authority over all UTM appliances. Check Point, SonicWALL and IBM ISS only allow for a single administrator to be connected at any one time to avoid conflicts.

Check Point has the most complex and useful approach, providing great flexibility across a large deployment. Multiple administrators can run its SmartDashboard in read-only mode to view, but not change, the configuration. And it has other tools, such as the separately priced Provider-1, which can segregate roles between, say, a desktop department to handle antivirus configuration and a network group to manage the firewall setup. Juniper has something similar with its separately priced NetScreen Security Manager for managing role-based administration. (SonicWALL is coming out with a new version of its management software that will allow multiple concurrent admin users, but this wasn't available for our tests.)

2 DAILY OPERATION

A critical value UTM products offer is the ability to quickly determine if your network has been breached or if you need to adjust the various protective mechanisms, since you have access to firewall, IDS and VPNs all in the same place. This means that if you mistakenly open a firewall port for the VPN, you can receive alerts to fix it without having to compare logs from two different places.

We used a typical scenario in which we ran the box for several days, examined the reports based on an initial firewall and protection rule set, and then adjusted our rules based on two situations—places where we wanted to eliminate false positives, and places where we needed to tighten down the box to prevent typical security weaknesses. Part of this exercise was to examine how reports would be created and examined and how threats will be evaluated and acted upon by the device.

Overall, Fortinet has the best set of tools to handle the day-to-day life of a security administrator, and Juniper scored lowest with its quirky main menu that scatters controls in almost random order *(See Juniper screen shot, below)*. Juniper also requires that you visit several places to examine reports and other screens to change its protection rules.

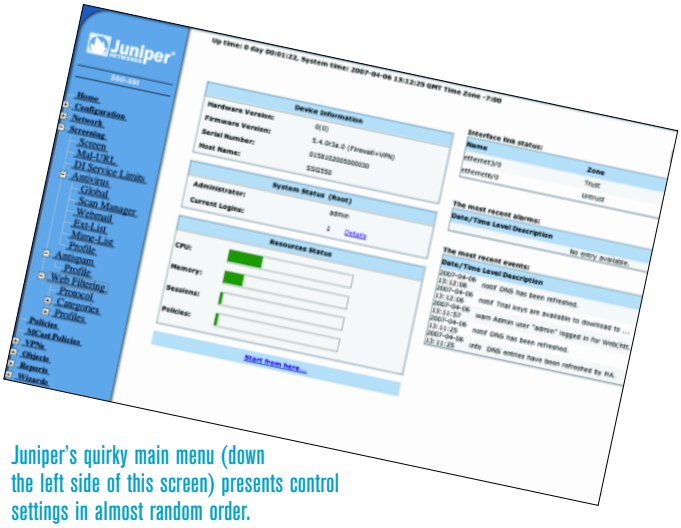

Juniper's quirky main menu (down the left side of this screen) presents control settings in almost random order.

UTM FEATURES

	Astaro Internet Security ASG 320	Check Point Software UTM-1 2050	Fortinet FortiGate-1000A	IBM Internet Security Systems Proventia MX5010	Juniper Networks SSG 550	SonicWALL Pro 5060c
Ethernet ports	8	8	10	10	4*	6
Maximum attachment size for antivirus scans	User selected	User selected	User selected < 139 MB	User selected < 1 GB	User selected < 24 MB	Unlimited
Antivirus supplier	Authentium, Clam AntiVirus, hardware-based	CA	Own	Sophos, own	Kaspersky	Own
Content filter supplier	SurfControl**	SurfControl	Own	Own	SurfControl	Own
IDS: Patterns or behavior	Both	Both	Both	Pattern recognition	Both	Behavior
IDS in or out of firewall	Either	Either	Either	Inside only	Inside only	Inside only
VPN types	SSL, IPSec	SSL, IPSec	SSL, IPsec	IPsec	IPsec	IPsec
Other ports scanned	IM, P2P, VOIP	IM, P2P, VOIP and many more	IM, P2P, VOIP	IM, P2P	IM, P2P, VOIP	VOIP, IM
Web application firewall	Minimal	Extensive	Minimal	Minimal	Extensive	Extensive
Authentication options	eDirectory, AD, RADIUS, LDAP	RADIUS	RADIUS, LDAP, AD	RADIUS	RADIUS, LDAP, SecurID	RADIUS, LDAP
Multiple concurrent admins	Yes	No***	Yes	No	Yes	No****

* The Juniper SSG box that we tested has room for six add-in cards and came with two additional gigabit Ethernet NICs installed.

** Websense recently announced it will acquire SurfControl.

*** Available with separate Provider-1 product

****Planned for version 4

The other products are capable and about equal in this area.

Fortinet's front page gives you just enough details to monitor its overall operations. You can quickly find attack summaries in its menus, and the policy definitions are easy to set, and more importantly, easy to change when you have done something wrong.

FIREWALL-IDS. Part of the usefulness of a UTM appliance is how its firewall and IDS work together, and flexibility in terms of where it can be used across different configurations of an enterprise network. In other words, some products can position the IDS module outside of the firewall to repel attacks and reject this traffic before it is processed any further, or to work with an existing firewall infrastructure at a headquarters network.

Fortinet and Astaro can also examine incoming encrypted packet streams and act on this analysis before passing these streams through other modules, thereby saving on processing power.

Check Point, Juniper, Fortinet and Astaro IDSes scan for both attack signatures and attack behaviors. SonicWALL only analyzes behaviors and IBM ISS only signatures. The IDS modules of both IBM ISS and SonicWALL UTMs can also explicitly detect outbound attack signatures.

The SonicWALL, IBM ISS and Juniper IDSes are hardwired to "live inside" the firewall, meaning that all network packets from the outside world go first to the firewall and then to the IDS for inspection. The advantage is that packets are filtered out by the firewall, reducing the inspection burden on the IDS. However, you do lose some insights because having the IDS outside the firewall can help you identify attack vectors early. This may be fine for organizations that manage both with the same administrative group, but problematic if the administrative roles are split.

REPORTING. The products have varying methods for producing reports, with different levels of details. All of the vendors except Astaro sell separate reporting tools (not evaluated for this review) that work across their larger security product lines. This assumes that you have more than just UTM boxes from these vendors and want to consolidate reports so that all firewall information is in one place, all IDS alerts are in another, and so forth. This may not work for all usage scenarios, and could be cumbersome if you have multiple vendors' products in your data center. Having to purchase add-on reporting tools somewhat undercuts the purpose of having an integrated appliance.

Astaro includes reports as part of the Web administrative interface and produces an "executive report," which doesn't do much more than show some nice graphs of traffic flows.

LIVE MONITORING. We examined several critical pieces of information available from the Web interface: real-time CPU and memory load, current alerts of potential network attacks, antivirus-related messages, and system health messages that required immediate attention.

This is helpful to see if your UTM box is overloaded or

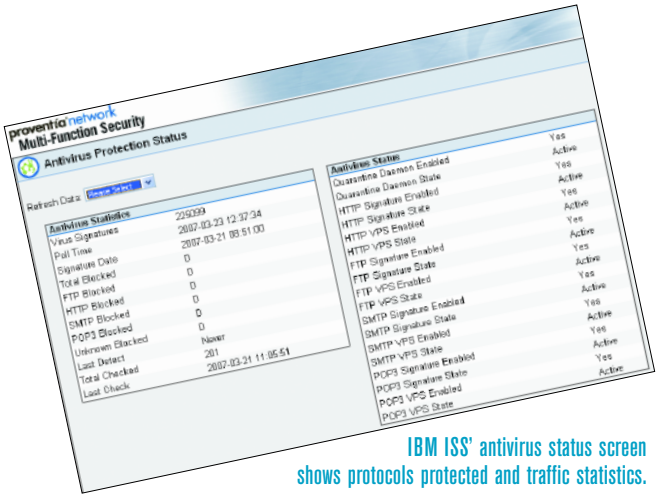

IBM ISS' antivirus status screen shows protocols protected and traffic statistics.

mismatched with the particular network traffic and inspection loads.

All of the UTM products except Check Point and IBM show the current CPU load and, in some cases, memory consumption on the home page of their Web interface, so it is easy to find and easier still to track. IBM ISS buries its status screen, while you have to visit Check Point's SmartView Monitor (a separate piece of software that comes as part of the UTM package) to get this information.

The three most useful front pages were from Astaro, SonicWALL and Fortinet, which offer all sorts of helpful summary information in one convenient place. Fortinet also includes a secure command-line console window within its Web interface, while the others require an SSH client to connect to their box if you need access to the command line. SonicWALL also tells you if you have set up the box with a known security weakness, such as allowing management from the WAN interface.

Check Point uses Windows software for its management, which means an admin must carry around a laptop with the software installed, rather than simply logging in through a browser. IBM ISS and Astaro can't be managed through Macintosh-based Firefox browsers, and we found some bugs when we administered SonicWALL with Firefox on a Mac.

Antivirus statistics are very important, since few things light up the help desk lines like email problems. IBM ISS has a simple-to-understand antivirus status screen *(See IBM ISS screen shot, above)*, showing messages blocked, signatures, and which ports are being blocked or scanned. Astaro also has a good summary display of its email traffic, but tweaking the protection results requires visiting several different sub-menus. Check Point and Fortinet put this information on summary screens; Juniper and SonicWALL have separate screens that summarize the virus penetrations.

3 AUTHENTICATION AND POLICIES

Setting up and tuning security policies for the various modules is at the core of these products. Ideally, you would want an appliance that makes it easy to figure out how to keep your network protected, but still allows users room to get actual work done, all the while providing feedback when you have too strong or too weak a policy.

SonicWALL and Fortinet clearly lead the pack in this regard with the others scoring equally behind. Even if you don't activate all of the security modules, both vendors' approach is easy to understand and provides just enough feedback so as to not overwhelm an administrator.

There are two basic approaches to how security policies are created:

• **Integrated policy** that applies to particular users or network interfaces. This has its advantages if your UTM box sits on several different network segments and you want to deploy different policies by segment or by user group (for example, one with servers on it, or one with engineering users). With this method, an administrator sets one policy that cuts across all of the individual security modules, with specifics for antivirus, IDS and so forth. Call this the traditional firewall approach, and each policy can enable different security modules for particular situations.

Fortinet and Check Point use this approach; Fortinet does a better job, setting up a series of four default protection policies that gives you a great starting point and examples that make it easy to modify them for your specific needs *(See Fortinet screen shot, below)*.

• **Separate policies** that are module-specific. This means there will be one policy for antivirus, another for general firewall tasks, and more for IDS actions. IBM ISS uses this approach; while it also has chosen lots of defaults to get you started, making modifications isn't as easy as with Fortinet, because you must make them in several places. Juniper also sets up security policies by module.

The appropriateness for your company depends largely on how you have structured your support staff. If you have an antivirus person on staff, and you have a box that requires adjusting antivirus policies in several different places, you have a lot more maintenance work than with a box where you can set these policies in a single place. However, your security staff may wear a lot of different hats and thus this might not be as much of an issue. It is really a matter of taste and organizational structure.

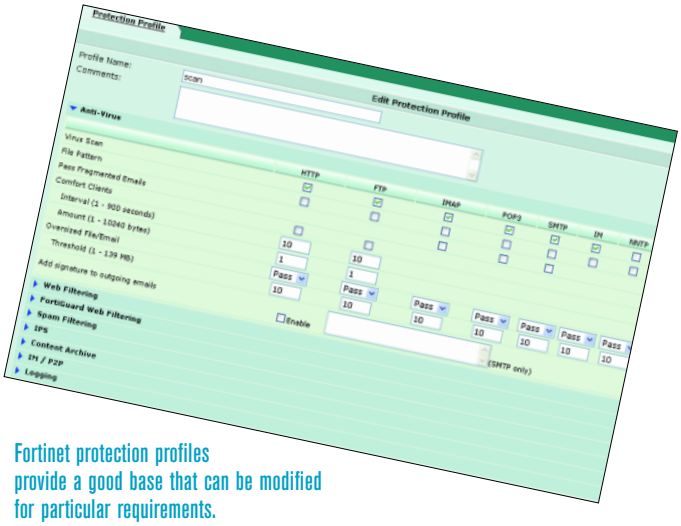

Fortinet protection profiles provide a good base that can be modified for particular requirements.

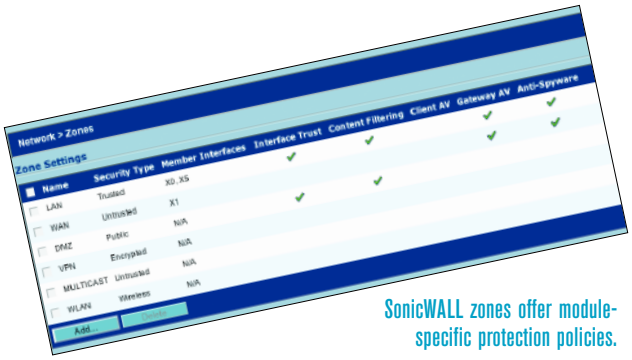

SonicWALL zones offer module-specific protection policies.

SonicWALL and Astaro mix both approaches. Astaro has policies that are based on application-layer protocols (Web, email, IM and so forth) and has separate policies for network layer events. This means that to make changes in the UTM operations, you need to touch screens in both the protocol section and the network interfaces. If you forget one or the other, you will have configuration problems or, worse yet, think you are protected when you aren't.

SonicWALL policies are module-specific, and are applied to particular network routes. That has a lot of appeal, and is why we give it top marks here. All of its protection rules are organized in a single section, and it is easy to apply them to the appropriate interface *(See SonicWALL screen shot, above)*.

Authentication capabilities are relevant if they are used for remote VPN connections. For most site-to-site VPNs, this isn't important unless you want to do some rudimentary endpoint protection or create policies based on particular user groups or roles.

All of the products support RADIUS authentication; Astaro and Fortinet can connect directly with Active Directory user store; Astaro also supports authenticating to Novell's eDirectory. Juniper can integrate with RSA's SecurID tokens directly.

All of the products offer IPsec VPNs, and Astaro, Check Point and Fortinet support SSL VPN terminations. None of the SSL modules has anywhere close to the level of features that a standalone SSL VPN box would provide.

4 FEATURE INTEGRATION

Our final series of tests looked at how the various functional modules work together. We also determined the third-party suppliers for these modules and what noteworthy features one product has that the others do not.

We gave SonicWALL the top grade because of its superior antivirus features, protection rule flexibility and implementation of IM protection across all of its security modules. Juniper and IBM ISS scored lowest because of the difficulty in making changes to their protection rules. For example, in order to implement protection or blocking of a specific protocol, you have to hunt down the rules that apply to that protocol and make adjustments in several places in the user interface. The other products fall somewhere in between in terms of complexity.

Each product uses different combinations of home-grown and third-party security services to round out its UTM coverage. Astaro, Check Point and Juniper use SurfControl for Web content filtering, while the others have developed their own content-filtering capabilities. Astaro uses Snort, while the others have their own IDS engines.

Astaro supplies three virus scanners—a proprietary one using the Authentium antivirus engine, another based on open-source Clam AntiVirus, and a PCI hardware-based antivirus capability from Sensory Networks. Juniper uses Kaspersky, and Check Point uses CA. IBM ISS uses Sophos, along with a second scanning algorithm that examines network behavior. SonicWALL and Fortinet have their own antivirus scanners.

The six products differ on how big a file attachment they will scan through their antivirus engines. SonicWALL claims an unlimited file size because it scans while streaming the packets, while the others are more limiting because they have to cache the files first. If performance bogs down, an administrator can automatically block files beyond a certain size. IBM ISS hides this setting in its advanced settings, while the others make it easier to adjust the maximum limit.

All of the products can at least monitor IM traffic *(See Astaro screen shot, below, for example)*, and some have rudimentary mechanisms to (sometimes) block particular IM protocols. SonicWALL was the only vendor that can completely block Google Talk and Skype conversations. Fortinet's IM protection is somewhat obscure. You have to go to two different places, one to handle policies for individual users and one to monitor or block the specific IM protocols.

Any solid defense against IM use will require combining Web filters to block access to particular sites as well as using the IM modules' features.

Check Point also does some very extensive port scanning, including ports that are used for VOIP, IM and P2P applications.

Web application scanning is absolutely essential if your company's Web servers are in remote locations or if you

Astaro presents a wide variety of choices to allow, monitor or block IM sessions.

MAKING THE GRADE

Company, product	**Astaro Internet Security** Astaro Security Gateway 320 www.astaro.com	**Check Point Software** UTM-1 2050 www.checkpoint.com	**Fortinet** FortiGate-1000A www.fortinet.com	**IBM Internet Security Systems** Proventia Network Multifunction Security MX5010 www.iss.net	**Juniper Networks** SSG 550 www.juniper.net	**SonicWALL** SonicWALL Pro 5060c www.sonicwall.com
Price as configured	$12,465	$15,500	$14,995	$14,890	$18,375	$10,995
Enterprise management, control	B	B-	B	A	C	B-
Daily operations, reports	B	B	A	B	B-	B
Authentication and security	B	B	A	B	B	A
Feature integration	B	B+	B	C	C	A
Verdict	B	B	A-	B	C+	B+
	Feature-rich, nice Web interface, relatively simple to set up. **Pros:** Numerous antispam and antivirus features; authenticates to both AD and eDirectory **Cons:** Confusing method of entering security policies	Licensing is its big weakness **Pros:** Flexible and extensive firewall and IDS **Cons:** Too many individual pieces of management software	Flexible and capable product **Pros:** Solid reporting and security features **Cons:** IM protection obscure	Easiest to set up and manage; nice arrangement of commands and policies **Pros:** Default settings are a big help **Cons:** Feature integration needs some work	A complex product made more so with a bad interface **Pros:** Feature-rich **Cons:** Miserable management interface and messy menus will confound configuration	Solid feature integration, reports and security features **Pros:** Web application firewall features, unlimited file attachment size for scans **Cons:** Complex configuration menus and a tired user interface

plan to set up a new Web server on an unprotected network such as at a branch office. Check Point, SonicWALL and Juniper offer protective mechanisms for preventing common Web application attacks such as SQL injection and cross-site scripting. We didn't find policy setting particularly straightforward for any of them.

The others just give lip service here, or require you to spend your days writing firewall rule sets.

For additional features, we liked Check Point's safe upgrade, requiring an administrator to complete a successful login within a specified (and user-selected) period of time; otherwise the box will roll back to a previous version. SonicWALL allows management of its wireless access points from its UTM device.

Not All Things to All People

UTM is one of those concepts that sounds great in theory, but is messy in practice. The six products tested all had their quirks, and we would have found show-stopping issues on all of the boxes if we didn't have a lot of support from each vendor.

While Fortinet and SonicWALL clearly have the best collection of features and Juniper the weakest, the others all had their good points, and the differences among each of the products is more a matter of taste and judgment than anything else.

Weigh the ability for multiple people to manage these boxes with how you organize your security staff. If you have separate groups managing firewalls and antivirus, for example, you might be better off choosing the products that separate their security policies.

You will also want to examine how a UTM deployment for your branch offices—which makes a lot of sense and can reduce your overall support burden—will be balanced with the products that you use or will use on your headquarters network. While Check Point and Juniper have solid solutions for the headquarters, they have less satisfying and less mature UTM product lines. Think carefully about what functions and modules you want to consolidate, and how you will go about managing the appliances before you invest heavily in any solution. ›

David Strom is a freelance writer, speaker and former editor-in-chief of Tom's Hardware *and* Network Computing *magazines. Send your comments on this article to feedback@infosecuritymag.com.*

FORENSICS

P2 Enterprise Shuttle

REVIEWED BY BRENT HUSTON

Paraben
www.paraben.com
Price: **$6,995**

Paraben's P2 Enterprise Shuttle is a remote digital forensic suite, allowing you to remotely conduct undetected forensic tests on Windows machines in your network without taking the machines offline.

This can be useful to acquire the data without raising suspicion of the target. It may also be used to monitor infected systems in real time.

Installation C+

Installation was pretty straightforward—a CD guide walks you through it—but we did encounter some issues after the installation. We were unable to get the proxy functioning due to a misconfiguration.

The installation automatically filled in the IP address to be used by the proxy and server with the hostname, which did not seem to work. The proxy would not start and did not really give a reason. We corrected the issue by editing the config files and changing the hostname to be the actual IP address.

The client agents can be installed directly or through the Captain, which controls agents and acquires and analyzes data from systems.

The latter allows you to place the agent without alerting the user or install agents on multiple machines.

Testing methodology: Server, Proxy and Captain were all installed on the same system. Agents were installed on a variety of Windows XP SP2 and Windows 2000 machines.

Management and Features B+

There are four major parts to the enterprise suite: the Agent, Captain, Proxy and Server. These modules interact with each other over a 128-bit encrypted channel.

The Paraben Agent is invisible to the user, although a savvy user may suspect something by the increased CPU load and network activity during acquisition. We were also able to see it with a rootkit detector.

The GUI-based Captain has a tabbed and framed design. Navigation is smooth, and buttons are easy to figure out with contextual help.

The Paraben Proxy, naturally, acts as an encrypted proxy between all of the components. It's installed on a system with an Internet connection

The Server is the main module, performing all authentication and acting as the central repository for acquired data. It verifies access permission for any actions initiated by the Captain and Agent to provide increased security. The Server should be installed on an isolated and secured system with no direct Internet connection.

You will spend most of your time with the Captain, which has quite a few tools to analyze clients. You can do a forensic dump of data, copying over each file or directory, or perform deep system inspections while the system is running. You can view running processes, what files those processes are accessing, and which registry keys they have open. Other capabilities include capturing screenshots, viewing the registry, processes, drivers and network sessions, as well as viewing the files on the system. You can create a full snapshot and save it to the database.

Reporting B

Reporting functions are fairly simple and to the point. Reports can be generated for module access such as server/proxy connecting, login and logout of the server, and agent connections. Each event is assigned a priority (fatal, error, warning and information); you can filter based on the event and the priority. Reports are generated in a table inside of Captain GUI. There are no charts or fancy graphics, but they're definitely not needed here. Reports can be saved to text, HTML or XML formats.

Verdict

Paraben's P2 Enterprise Shuttle is a good offering if you are looking for a remote forensics tool to use in a Windows environment. It provides all the tools necessary for a complete forensic analysis of a system, as well as the security to ensure the integrity of the acquired data.

SECURITY TESTING

HotPick INFORMATION SECURITY®

Core Impact 6.0

REVIEWED BY MIKE POOR

Core Security Technologies

www.coresecurity.com

Price: **$25,000 for annual subscription**

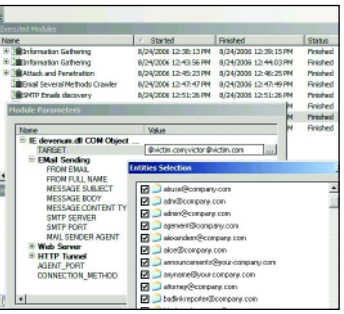

Prior to Core Impact, the vast majority of security penetration testers would use "off-the-Web" exploit code, after scouring an application code for backdoors and covert channels. Core Impact changed the security landscape by providing stable, tested and trustworthy exploits for ethical hacking.

The latest release of this automated, commercial-grade penetration-testing software platform is an invaluable tool for professional penetration testers and corporate security engineers.

Configuration/Management A

Installing Core Impact 6.0 was a breeze—download, double-click and enter the long string to decrypt the installation executable.

There are two main workflows: The rapid penetration test guides the user through the phases of reconnaissance, exploitation and reporting via a series of menu-driven wizards. You can choose the type of exploits to test, as well as the levels of risk to take (e.g., whether or not to run exploits that might crash or DoS the service).

The second workflow is conducted via modules. In this more granular mode, you choose from this version's plethora of available exploits.

Effectiveness A

We were able to run multiple exploits, test and compromise machines in minutes, giving the attacker complete command and control over our target systems and arming us with detailed information.

We first ran a rapid penetration test in which Core Impact walks you through a simple set of questions to identify the target systems. It scanned the network and identified live hosts, listening ports and OS versions, allowing us to choose the exploit modules most likely to compromise the target.

We chose remote exploits first, attacking the Microsoft Windows Plug and Play services umpnpmgr.dll vulnerability (Microsoft bulletin MS05-039). One click brings up a quick description of the exploit and the vulnerability, including links to patches and remediation information.

We then tested client-side exploits, switching to the "Modules View" to select individual attacks. We ran the IE IFRAME buffer-overflow exploit—which automatically sets up a Web server on the attack system, with a Web page serving up the exploit—and browsed the attack site to compromise our target system. Compromising a target using remote and client-side exploits demonstrates the need to patch the vulnerable software. One of the biggest benefits of running exploits against real systems is gaining insight into how credible the threat posed by vulnerabilities is in our environment.

At the end, Core Impact removes all the agents from the target system, a step often neglected by inexperienced penetration testers. Test activity is logged for review.

Reporting B

Core Impact comes with a number of report generation options, from a simple executive summary to a detailed vulnerability report. The reports are simple and straightforward. The vulnerability report includes the number of systems compromised, along with detailed information regarding the vulnerabilities exploited. Administrators could use this report to remediate the exposures by following the remediation information and links to patches.

Core Impact uses Crystal Reports, with an XML-based generation system that can be altered and customized to meet your requirements. To take advantage of this feature, however, you have to get under the hood and change the XML templates.

Verdict

Core Impact 6.0 is an amazing tool to validate your security posture. We highly recommend it to security engineers to verify the vulnerability of their networks, or confirm test results from third-party consultants.

Testing methodology: We used VMware virtualization software to install a fully patched Windows XP Pro system to host Core Impact; and a Windows 2000 Advanced Server system with a few service packs missing to play the victim.

PENETRATION TESTING

Metasploit Framework 3.0

REVIEWED BY PETER GIANNACOPOULOS

Metasploit LLC

www.metasploit.com

Price: **Free**

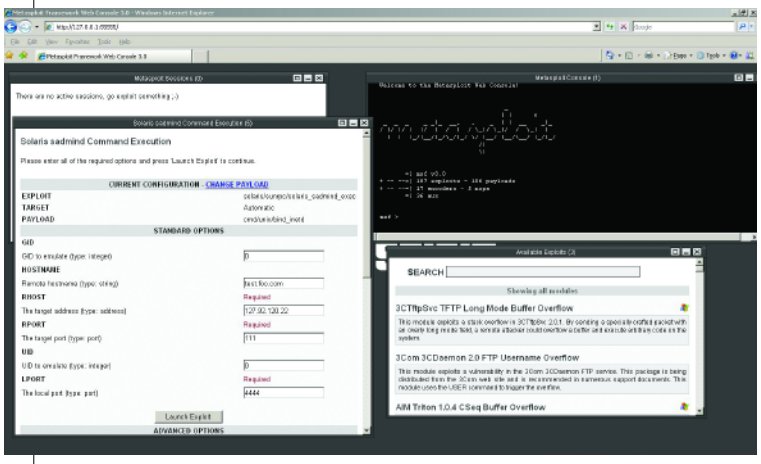

The Metasploit Framework is a platform for developing, testing and executing exploit code for all popular Unix, Linux and Windows platforms; it's an essential tool for the serious penetration tester or security professional. The latest release only serves to further cement its formidable capabilities.

Installation is a breeze: download the appropriate package (Windows or Linux) and execute. The Metasploit Framework can either be used via the familiar and capable console interface or the much improved Web interface. For a free product, it contains surprisingly good documentation for users and developers, so it's fairly easy to get productive quickly.

Metasploit has a searchable database of more than 180 exploits, targeting multiple processor architectures and operating systems, with more than 70 payloads that can be delivered to exploitable systems. Using the product is ridiculously simple: select the exploit via the Web or console GUI, specify target, payload and options, and run the exploit. It really is as easy as "point, click, own."

The payloads range from simply binding a reverse shell to injecting DLLs (like a VNC server) into the target's memory space to uploading and executing scripts or apps on the target. As if this isn't enough, there are also tools for building your own exploits, such as developing a NOOP sled to exploit a buffer overflow. Building new exploits is essentially writing code, so you'll need to have Ruby development skills (some C experience wouldn't hurt either). This shouldn't be a problem, since almost all of Metasploit's target audience will have some ability in this area or work with someone who does.

Exploits can be delivered either directly to the target host, or via a chain of proxies, which are nice for obfuscating attacks. Additionally, various browser hijacking routines will let you load malicious ActiveX controls (either your own or some that are bundled with Metasploit) to vulnerable Internet Explorer versions. One way or another, you will be able to gain a foothold in a vulnerable system and leverage it for greater access. Determining whether or not an exploit succeeds depends on the payload chosen. For example, if you elect to bind a shell, Metasploit will open a console session and connect back to the host via the specified port number.

Metasploit can continually update itself with the latest exploits and payloads developed by its sizable user community. Even if you don't possess the deep programming knowledge to make full use of its exploit development capabilities, you'll benefit from the work of others and stay current as new exploits come online and old ones are addressed by patches.

Metasploit isn't a shrinkwrap port scan or vulnerability assessment tool for the casual user. It's best to think of the product as a development environment akin to Visual Studio, but with a laser focus on developing usable exploit code. It is a serious pen tester's delight, but it's also the sort of tool that gives security officers nightmares, reinforcing the need for aggressive patching, layered defense and encryption of data at rest.

Testing methodology: We installed the Metasploit Framework console on a Windows XP SP2 and SUSE Linux 9.3 hosts with no hitches and used both platforms to successfully exploit vulnerable versions of Windows, Red Hat, SUSE and Fedora hosts.

Verdict

Metasploit Framework is a mandatory tool for every security professional. This brief overview offers a glimpse of its capabilities.

RISK/POLICY MANAGEMENT

Elemental Security Platform

REVIEWED BY BRENT HUSTON

Elemental Security

www.elementalsecurity.com

Price: **Management server, $35,000; desktop agent, $60; server agent, $600**

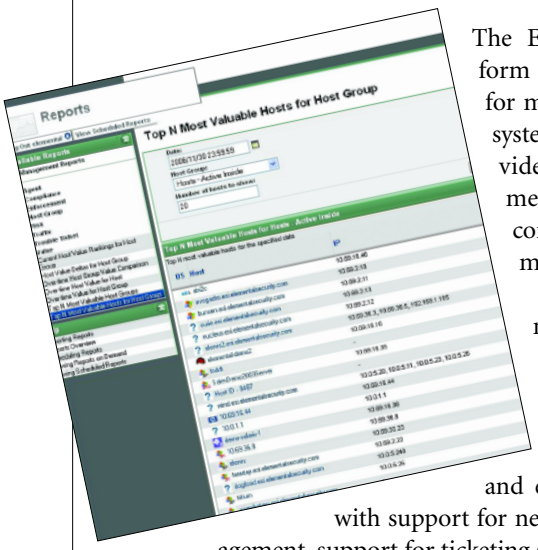

The Elemental Security Platform (ESP) is a powerful tool for monitoring and enforcing system compliance, and provides effective asset management, asset-centric access controls and risk management.

Since we reviewed Elemental Security's version 1.1, then called Elemental Compliance System (August 2005), the product has matured and extended its capabilities, with support for new client OSes, risk management, support for ticketing systems and better LDAP integration.

Setup B+

Our ESP server was preconfigured, but Elemental typically sends an engineer on-site to install the device and provide a rundown on features and usage. We were impressed with the ease of client agent installation, and getting the clients/servers running.

Agent installation simply requires giving it the address of the ESP server and answering one or two other questions, depending on the platform.

The client connects securely to the server, reports gathered information and downloads relevant policies. The server automatically gathers data about open ports and services to categorize hosts, and places them in groups that can be defined manually or imported from LDAP.

As a key to risk assessment, the ESP server assigns a value to the system, depending on what services it's running. These values can be overwritten.

Testing methodology: The system we received was preconfigured for our environment. In our tests we used a variety of OSes, including Windows, Mac OS X and Red Hat Linux.

Effectiveness B+

The user interface is clean and functional. Pages are uniform, with all dropdown menus on the left side, navigational buttons for selecting your page at the top, and relevant page information, such as reports, or policies you are creating, in the middle.

But it can still use a bit more tweaking. For instance, during policy creation, if you click on rules for a closer look, they open on the same page, so there's no facility to backtrack to where you were. So we had to hit the backspace button, which erased any rules we had already configured. You can right-click on the link and open a new window to bypass that inconvenience.

ESP can be used as a basic asset inventory tool or a granular asset-centric access control solution, depending on policy. Policies can contain a variety of rules, from packet filters, to whether the user can install a piece of software, to rules that check for compliance with baselines (such as CIS, or HIPAA security requirements).

We defined some simple policies, such as denying access to secured hosts by unsecured hosts (hosts not running the agent), by naming the policy and adding rules. Some rules require additional configuration, such as ports for the network filters.

Reporting B+

Reports can be created for any aspect of ESP for managers, and viewed on-demand or scheduled. You can view reports for each policy, as well as specific host groups under a policy.

The reports are easy to read and feature a variety of graphs and charts to effectively represent the information. Data can be exported to a variety of formats, including CSV and PDF.

Verdict

We are as impressed with the latest release of Elemental Security's tool for monitoring and continuously assessing the security posture of large, heterogeneous enterprises as we were with its early version.

POLICY MANAGEMENT

nCircle Configuration Compliance Manager

REVIEWED BY BRENT HUSTON

nCircle www.ncircle.com

Price: **Management server, $20,000, plus $199 per monitored server, $199 per network device and $49 per other assets**

Riding herd on the integrity of your security infrastructure is not an option—it's a necessity. But keeping your IT assets in compliance in a large, complex environment is daunting without automated tools. nCircle Configuration Compliance Manager brings policy compliance and security management together into one centralized suite, helping cut the job down to size.

Configuration/Installation B+

Installation is straightforward, either on a single server or multiple systems. nCircle's agentless technology simplifies the process.

Vulnerability scanning is handled by third-party scanners, such as Nessus, IBM ISS Internet Scanner and QualysGuard, allowing you to integrate existing tools. nCircle pulls vulnerability data directly into its database (we used MSDE, but larger organizations will want to use SQL Server). Integrating Nessus in our lab was a snap.

Features and Interface A

nCircle is a deep product with a lot of features. It finds assets on your network, either actively or using its passive detection technology, identifying hosts and determining basic information, such as OS and open ports.

The efficient management console presents three primary tabs to a paned interface. The changes pane contains the aggregate of the latest alerts. The inventory section lists all network assets, broken down by subnet range. The compliance view is similar to inventory, but adds columns for compliance with applicable policies, both pass/fail and by percentage.

Testing methodology: We tested nCircle Configuration Compliance Manager in our lab, including Windows (XP, Server 2003), Linux and Solaris systems, and Cisco networking devices.

Effectiveness B+

nCircle Configuration Compliance Manager assesses systems against predefined policy, vulnerability assessments and several other options, such as spider-like scans of Web servers, in response to events or at predefined times. It monitors files for activity, and tracks file attributes, runs MD5 checks for integrity and issues alerts when changes are detected.

We were impressed with nCircle's asset inventory capability, identifying and collecting detailed information on each system. For example, on Windows machines, nCircle reported every piece of installed software, users, groups, services running, shares available and updates installed. For instance, it can check for the latest AV version or unauthorized apps.

We configured nCircle Configuration Compliance Manager to reactively scan a host with Nessus, and issue a policy compliance check whenever target files were modified. Many other events can trigger tasks, such as finding new assets, or starting a task when a previous one has finished. You can also schedule scans and set tasks by single host or by group.

Compliance starts with predefined policies for various OSes. Creating policies from scratch can be daunting, but nCircle can automatically create policies from "gold standard" machine configurations.

nCircle has included a PCI compliance policy, and said it was planning to include HIPAA and SOX policies in upcoming versions.

Reporting B+

Reporting is critical for compliance tools, and nCircle's is thorough and easy to generate. Reports can include asset and file changes, vulnerability reports, installed software, risk trends, etc. They are easy to read and have colorful charts as well as technical breakdowns. Reports can be exported in .pdf, Excel, Crystal Reports, Word and RTF formats. nCircle also integrates with the Remedy ticketing system.

Verdict

nCircle Configuration Compliance Manager is a compelling package that rolls policy compliance and vulnerability detection into one usable package.

VULNERABILITY SCANNING

PatchLink Update 6.3

REVIEWED BY TOM BOWERS

PatchLink

www.patchlink.com

Price: **$1,495 per update server, plus annual subscription fee starting at $18 per client**

Keeping systems patched can be a nightmare for enterprises. Medium- and large-sized enterprises will find this automated patching tool from PatchLink an excellent, cost-effective solution.

Configuration/Management B

PatchLink's wizard-based process walks you through installation, including MSDE if it (or SQL Server) isn't already present.

PatchLink Update provides a Web-based manual installation process for one-off installs. It's also possible to automate installation by using a script or Group Policy Object with the standalone installer. Agents can be automatically deployed through the console using Active Directory/LDAP lookup or IP/DNS scanning.

One of the product's major shortcomings is weak integration with AD. Although you can deploy the agent through AD, it's not possible to manage devices through AD organizational units (OUs) or import OU membership information into PatchLink's group-based management structure.

Policy Control B+

PatchLink Update grants administrators a great deal of policy control, including the ability to schedule scans and deployments.

PatchLink will enforce minimum baselines automatically according to standards you create for admin-defined or platform-based groups of devices. When a patch becomes available, you may create a rollout schedule based on your specific needs.

As noted above, tighter AD integration would allow enterprises to leverage the time they've already invested in creating an OU structure to apply policies to different parts of the organization.

We were impressed with PatchLink's ability to create custom patch packages and deploy them to the enterprise or specific groups on a scheduled basis. These packages can also be used to change configuration settings, install software and run automated scripts.

We also like the flexibility to grant end users varying degrees of control over the PatchLink agent. For example, administrators may choose to allow users to delay patch deployments and system reboots.

Effectiveness A-

PatchLink Update is a robust and flexible automated patch tool that will go a long way toward taking some of the pain out of Patch Tuesday.

The Windows-centric product provides a baseline level of support for other OSes, including Mac OS X and several Unix/Linux variants. It also supports many Windows applications, but only a handful of Mac apps and none for *nix.

Determining which systems are unpatched and verifying successful deployments are major pain points for enterprises. PatchLink uses digital signatures for each patch and scans the host system to determine patch level. If the initial patch fails, it attempts to redeploy it up to three times.

Reporting B

PatchLink has a wide range of reports on the status of agents, enterprise-wide package compliance, patch deployment status for systems/groups and the current mandatory baseline.

Customization and filtering are limited. For example, you can filter your report based on devices or device groups, but not on complex criteria such as creating a report listing devices missing patches for a certain period of time.

Verdict

PatchLink Update 6.3 is a solid solution to the enterprise patch management problem and demonstrates its true power in a Windows environment.

Testing methodology: PatchLink Update was tested in a Windows Server 2003 and Windows XP environment within VMware Workstation. We ran the PatchLink Web server on IIS 6.0.

RISK MANAGEMENT

RedSeal Security Risk Manager

REVIEWED BY ADAM HOSTETLER

RedSeal Systems
www.redseal.net
Price: **Starts at $25,000**

Your network produces a flood of information that could tell you where your business is at greatest risk. But how do you sort through it all and determine exactly how your critical assets are threatened?

RedSeal's Security Risk Manager (SRM) enables security administrators to model and manage threats to those corporate assets and network infrastructure. The appliance transforms network device configurations, vulnerability data and system value ratings into a graphical view that shows how systems can be compromised.

Setup and Configuration B

Setup is fairly easy. We hooked up the serial cable, ran a few configuration commands and installed the SRM Java administrative application through our browser (only Windows is supported).

SRM generates risk and threat maps based on imported device configurations and vulnerability data. SRM supports popular network devices out of the box, including Cisco IOS, Cisco PIX5/6/7, Juniper ScreenOS and Check Point Firewall-1/VPN-1 NGX, as well as vulnerability sources such as Nessus and QualysGuard. Other devices can be imported with the help of RedSeal, or by creating an XML schema. Device and vulnerability data can be imported manually, or SRM can retrieve it directly from the devices or a central repository through a variety of means (FTP, SSH, HTTP/S, Telnet, CVS).

Effectiveness A

SRM's clean tabbed interface nicely displays available information and makes it easy to import data or edit device and system values. This clean interface carries over into the network map, which appears quite haphazard at first, with systems appearing out of order and all over the map. This is quickly alleviated by the auto-arrangers, which are a great feature for larger networks. The map shows connectivity between devices and networks, accounting for traffic flow restricted and allowed by ACLs.

You can assign values (from 1-100) to systems to help determine where your company is at greatest risk. SRM uses this data to generate risk and threat maps.

The threat map is similar to the inventory map, but includes threat calculations based on exposure and business value, modeling how an attacker might get to a system, and through which vulnerabilities. You can pick any point in your network to see which systems can talk to this system, or what systems your selected system can see. A "heat box" style risk map shows which systems are at greatest risks and establishes mitigation priorities.

The threat map showed us systems at risk, such as firewalls allowing improper traffic, or systems that had severe vulnerabilities. After correcting the issues and reloading the data, we could regenerate the maps and see that the issues were mitigated. For instance, SRM showed that a high-value internal database server could be attacked from an FTP vulnerability on an external server. After the issue with the FTP server was mitigated, SRM showed that the database server was no longer threatened.

Reporting B

Reports can be generated for a number of different categories, including inventory, network device configuration errors, exposure to vulnerabilities and performance data of the appliance. The network configuration checks are quite useful, as they compare your network devices against a built-in rule set to check for common configuration errors. Some reports contain colorful heat bars, or expanding bubbles to show threats, others just text.

Verdict

RedSeal SRM is a very good tool to provide an overall view of network threats and risks, and will help you prioritize mitigation measures.

Testing methodology: We tested the RedSeal SRM appliance using RedSeal-provided data that modeled a network containing a mixture of network devices, and vulnerability data, in addition to data generated in our lab.