Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

Interested in learning more?

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

GIAC Security Essentials Certification (GSEC) Practical Assignment
Version 1.4b, Option 1

Accommodating Visitors with Untrusted Devices
Reconciling Hospitality and Security in a Business Environment

Submitted by: Jefferson Scher
Date: January 1, 2004

Abstract
Every day, companies are asked to allow visitors to connect to the internet from their offices. In an ideal world, untrusted devices would be completely isolated from the company network. In many cases, however, the connections serve a legitimate business purpose and, as a practical matter, some connections are inevitable. Such connections raise a number of troubling security issues, including the possibility of exploiting vulnerabilities inside the LAN, and undesirable traffic to and from the internet.

A company's security policy should address these threats using a layered approach that includes rules and educational measures for employees, and acknowledgements signed by visitors. A range of physical and technical security measures are available, and each company will need to balance their monetary and administrative cost against their benefits in determining how best to manage these risks.

Introduction
News stories about threats to network security routinely feature the usual suspects: hackers, crackers, corporate spies, and script kiddies. But there is another threat that deserves the attention of security analysts. Every day, companies open their doors to a steady stream of trusted visitors, including clients, vendors, contractors, friends, and relatives. The longer the visit and the greater the distance from her or his home base, the more likely that the guest will wish to access the internet from the company's offices. Everyone can relate to at least one of these scenarios:

- A client of the company needs to retrieve a confidential document from her firm's computer network, without the risk of disclosure that could result from having it faxed or e-mailed to your offices;

- A vendor needs to connect to its public web servers to demonstrate a hosted solution your company is evaluating, and there is not enough time to copy data files and presentation materials onto one of your computers before show time;

- A contractor needs to download patches from Microsoft's web site to burn to CD for someone about to leave the office, and no one is authorized to provide the contractor a network logon to use one of your computers;

- The CEO's son needs to submit his term paper electronically, and did not leave sufficient time to drive to Starbucks to use a public hotspot.

Eventually, with everyone having the best of intentions, an untrusted laptop computer or other device will be attached to the company's network — regardless of a company's official security policy. The question is not "if" or "when," but "how often?"

Connecting an untrusted device to either a wired or wireless network raises a number of troubling issues, including the risks of data loss, disclosure of confidential information, unavailability of services, and potential legal exposure. In an ideal world, such devices would be completely isolated from the company's network, but in some environments the resulting inconvenience might be unacceptable. In many cases, the connections serve a legitimate business purpose and, as a practical matter, some connections are inevitable. Therefore, a company's security policy should take a layered approach, addressing the desirability of minimizing such connections, and providing suitable technology for managing the risks associated therewith.

Modeling Threat and Response
A well-meaning visitor can be trusted not to actively attack the network, but the same cannot be said for his portable computer: the software running on the device may well answer to others with less innocent intentions. The host offering her visitor an Ethernet port cannot be certain, without intrusive inspections, that the device is free from potentially damaging software. It is as though the host network is connecting with all the other networks the visitor has used, protected only by the visitor's and his former hosts' exercise of "safe computing practices." One would be wise to exercise caution.

1

As part of GIAC practical repository.

In this discussion, I assume the visitor will not be granted credentials (such as a username and password) to access internal servers that require authentication. Such credentials seldom should be required by casual visitors and, if they are, the visitor can be expected to tolerate more intrusive measures to verify that his machine is "clean." Unfortunately, a great deal of exposure remains.

Untrustworthy Software

Microsoft Windows is the most commonly used operating system on desktop computers and servers in a typical business environment. Over the years, numerous vulnerabilities have been found in components of Windows, Internet Explorer, and related applications installed on millions of computers. Many of these vulnerabilities can be exploited only through a user's volitional act, such as opening an e-mail message or executing a file, and in many cases the user must sign on to the system before taking such acts. But a number of others can be exploited, without any special privileges, simply by sending a packet to a port on which an unpatched service or application is listening. Recent examples of the exploitation of such vulnerabilities include the following:

- The "Blaster" worm exploited a buffer overflow vulnerability in Microsoft's Distributed Component Object Model (DCOM) interface within the Remote Procedure Call (RPC) service of Windows NT, 2000, XP and 2003.[1] An RPC interface provides a standard protocol by which a process on one computer can request services from a process on another computer. One of the features made available by RPC on these versions of Windows was DCOM, a Microsoft technology for allowing software components to communicate with one another over a network. By transmitting a crafted packet to one of the ports listening for RPC requests, the Blaster worm was able to take control of the machine and use it to spread itself to other machines in the same manner. The worm also contained code to launch a denial-of-service attack on servers at windowsupdate.com. Adding inconvenience to injury, Windows XP and 2003 shut down within sixty seconds when the worm crashes the RPC service, making it extremely difficult to remain connected to a network under attack.[2] New worms continue to be written targeting the same vulnerability, often combined with attacks on other vulnerable components.[3]

---

[1] The vulnerability, dubbed CAN-2003-0352 in the Common Vulnerabilities and Exposures (CVE) database, is described in CERT's Vulnerability Note VU#568148. Finlay, Ian A. and Damon G. Morda. "Vulnerability Note VU#568148: Microsoft Windows RPC vulnerable to buffer overflow." Aug. 14, 2003. URL: http://www.kb.cert.org/vuls/id/568148 (Jan. 1, 2004).

[2] The operation of the Blaster worm and the available patches are described in more detail in the online articles cited in the References section.

[3] For example, the AGOBOT family of worms attempts to gain access to a machine using three different ports. Te, Darwin. "WORM_AGOBOT.AZ" Trend Micro Virus Encyclopedia. Dec. 1, 2003. URL: http://www.trendmicro.com/vinfo/virusencyclo/default5.asp?VName=WORM_AGOBOT.AZ&VSect=T (Jan. 1, 2004).

2

- The "SQL Slammer" worm exploited a buffer overflow vulnerability in Microsoft's SQL Server 2000 software, including the widely deployed embedded version, Microsoft Desktop Engine (MSDE).[4] Because SQL Server 2000 could host multiple instances of the database engine, the SQL Server Resolution Service was added to direct requests to the proper port for the desired database. By transmitting a crafted packet to the port listening for such requests, the Slammer worm was able to take control of the machine and use it to spread itself to other machines in the same manner.[5] Some users of other Microsoft and third party applications which utilize MSDE were caught by surprise, as they were unaware of the use of SQL Server technology. Posting or circulating lists of such applications became, at least for a time, a popular pursuit.[6] The current "Data Access Downloads" page on Microsoft's MSDN site continues to offer the unpatched version of MSDE (Service Pack 2), with a recommendation to apply a separate patch, rather than a fully patched up version.[7]

The Blaster and SQL Slammer worms spread rapidly around the world, and it is reasonable to speculate that their ability to move from one machine to another without any user action or authentication was a key ingredient in their success. An untrusted device infected with either of these worms, or others like them, would be very unwelcome inside a company's network.

Access to More Trusted Zones

Most business networks are divided into segments or "zones." Electronic communications travel relatively freely *within* a segment, but may be prevented from moving *between* zones by devices implementing a company's security policy. A company's employee workstations, file and intranet servers, and other computing assets intended for routine internal use typically form one or more segments referred to

---

[4]   The vulnerability, dubbed CAN-2002-0649 in the CVE database, is described in CERT's Vulnerability Note VU#484891. Lanza, Jeffrey P. "Vulnerability Note VU#484891: Microsoft SQL Server 2000 contains stack buffer overflow in SQL Server Resolution Service." Mar. 26, 2003. URL: http://www.kb.cert.org/vuls/id/484891 (Jan. 1, 2004).

[5]   The operation of the Slammer worm and the available patches are described in more detail in these sources: Microsoft Corporation. "Finding and Fixing Slammer Vulnerabilities." Feb. 21, 2003. URL: http://www.microsoft.com/security/slammer.asp (Jan. 1, 2004).; Danyliw, Roman. "CERT® Advisory CA-2003-04 MS-SQL Server Worm." Jan. 27, 2003. URL: http://www.cert.org/advisories/CA-2003-04.html (Jan. 1, 2004).

[6]   For a compilation of lists and diagnostic directions, see the February 3, 2003 issue of the e-mail newsletter Woody's Office Watch. Leonhard, Woody. "TOP TIP: DO YOU HAVE MSDE?" Woody's Office Watch. Vol. 8, no. 4 (2003). URL: http://www.woodyswatch.com/office/archtemplate.asp?v8-n04 (Jan. 1, 2004).

[7]   See Microsoft Corporation. "Data Access Downloads." Latest file dated Oct. 15, 2003. URL: http://msdn.microsoft.com/downloads/list/dataaccess.asp (Jan. 1, 2004).

3

as the local area network (LAN). The devices in these "more trusted" zones typically are connected by switches that do not enforce any access rules. A firewall prevents devices outside the more trusted zones from accessing devices inside the more trusted zones unless they comply with the company's access rules. Devices in less trusted zones, by contrast, tend to be more accessible from outside those zones. Such devices usually include public web servers and mail servers, and such a segment often is called a demilitarized zone, or DMZ. A firewall separates each less trusted zone from other zones, and from the wholly untrusted public internet. Defenses that filter traffic entering or exiting a zone often are termed "perimeter" security.

Companies often simplify the management and reduce the cost of network security by assuming that they have complete control over the software running on the devices connected to the more trusted zones of the network. For example, all employee workstations may be protected by a centrally managed antivirus software package which is forcibly installed and updated when the employee provides her credentials to the logon server.[8] Because the antivirus software substantially eliminates the possibility of malware[9] running on workstations in the more trusted zones, the company may forego providing firewall protection *between* those workstations, and instead rely on backups to ameliorate the risk of data loss caused by any unforeseen threat. In general, the devices in the LAN segment are trusted not to pose a security risk to one another.

Potential Threats from Visitor Connections. The assumption that all devices are trustworthy breaks down, however, if a visitor is allowed to connect to the LAN through an Ethernet port or wireless access point. The untrusted device potentially can introduce any known exploit — including "old" ones such as Blaster and SQL Slammer — into this previously clean and safe environment. While most such exploits could be thwarted by keeping all employee workstations and servers current with operating system and application patches, many companies defer this type of maintenance, either to allow time for testing, or due to other priorities.[10] For such companies, which typically rely on firewall rules at the perimeter to protect vulnerable ports and unprotected shares

---

8   One such product is Trend Micro's OfficeScan Corporate Edition, whose central server pushes virus pattern and software updates to workstations while receiving incident reports and quarantined files in return, all using the HTTP protocol. Further product details are available at http://www.trendmicro.com/en/products/desktop/osce/evaluate/overview.htm.

9   Malware is an umbrella term that covers viruses, worms, trojans, and other forms of malicious code.

10   The difficulty of staying on top of vendor patches is frequent fodder in publications directed toward corporate IT personnel. Recent examples include: Fontana, John. "How to handle patch management." Network World. Dec. 1, 2003. URL: http://www.nwfusion.com/research/2003/1201howtopatch.html (Jan. 1, 2004).; eWEEK Labs. "Labs Answers Patch Management Questions." eWEEK. Sept. 8, 2003. URL: http://www.eweek.com/article2/0,4149,1257572,00.asp (Jan. 1, 2004).

4

inside the LAN, the untrusted device could significantly disrupt operations, compromise confidentiality of information accessible to employee workstations, and possibly even press those workstations into service as zombies participating in attacks against computers inside or outside the more trusted zone.

Furthermore, the risk of exploitation potentially extends beyond workstations to the servers in the more trusted zone. An unpatched Windows server might permit an unauthenticated "null session" to gain access to user names.[11] A spider could crawl any internal intranet that permits anonymous access, harvesting what might be extremely valuable proprietary information. The visitor need not intend or even be aware of these activities; he only needs to have failed to practice safe computing habits at some point in the past.

If the LAN has been configured and maintained in the manner described above, a company has three choices: (1) ensure that no untrusted device ever connects to the more trusted zones, (2) strengthen security in those zones, or (3) a combination of both approaches.

Preventing Access. A decree that "none shall connect" — in itself — is unlikely to provide sufficient security. Such a rule should be accompanied by (i) a convenient alternative, and (ii) user education about the importance of observing the policy. Traditionally, the most common alternative was to wire publicly accessible areas, such as conference and telephone rooms, to a separate network segment with little or no access to the network's more trusted zones. Recently, vendors have begun to offer companies products to create public wireless access points, or "hotspots," as a substitute for or supplement to wired access.

Both of these approaches can be convenient for visitors, but both have their drawbacks for the company itself. In the case of isolating conference rooms, company users working in those areas will need to use virtual private networking (VPN) or other authenticated access technology to connect to the LAN to open files, print, and perhaps even to check their e-mail. Such technologies require additional configuration and training, can make access more fragile, and can reduce performance. These are reasonable trade-offs to make in most cases, but users might attempt to circumvent them through the use of wireless access to the more trusted zones of the network.

A visitor hotspot requires careful attention to security design so that network access can truly be limited to its intended users. The company probably would want to restrict access to certain Media Access Control (MAC) addresses, and avoid broadcasting the wireless segment's Service Set Identifier (SSID). Without these minimal measures, the access point essentially would be available to anyone within its signal radius. On the other hand, it may be inconvenient to require that the access point be apprised of the

---

11   One example of such a vulnerability in Windows NT is described in entry CVE-2000-1200 in the Common Vulnerabilities and Exposures (CVE) database (http://www.cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2000-1200).

visitor's MAC address, and that the visitor's software be configured to use the proper SSID. While the configuration details are beyond the scope of this paper, one possible approach would be to "lend" visitors a wireless NIC to ease this process.

Another alternative would be to implement a solution specifically designed for enterprise visitor access. Proxim Corporation's ORINOCO AP-2500 product, for example, combines a feature-rich gateway with a wireless access point.[12] The product documentation indicates that the visitor need only open a browser in order to be connected to the access point for authentication, regardless of network settings, and provides the host the ability to enable and disable various services. Authentication can be by MAC address, or a username and password, that the company hotspot operator adds to an authorized user list (either through a web interface to the device or on a RADIUS server). Once access is granted, the user can be connected to the internet either through a network switch (preferably on a port segregated into a separate virtual LAN (VLAN)), or through a firewall or router. A user who is not authenticated will have no access to the wireless network. While the AP-2500 supports only WEP encryption, it is compatible with certain IPSec-based VPNs, allowing the visitor to make a more secure connection to his home network.[13]

Whichever options are implemented, there remains a reasonable chance that visitors will be permitted to connect to a more trusted zone of the network, either through inadvertence, or because they are working in an area of the building where the wired or wireless "visitor segment" is not readily available. Technologies could be adopted to *literally* prevent such devices from obtaining access to the network, by eliminating the network's implicit trust of devices connected to particular ports on a network switch. In this scenario, every device must supply certain credentials for accessing the network. For example, the network switches could be configured to permit access only to devices with specified MAC addresses (perhaps by creating a dynamic VLAN[14]), or to devices able to supply a unique code.[15] Erecting such barriers involves a considerable amount

---

12   References to products are for illustration and should not be construed as an endorsement. Any representation or warranty that said products would be suitable for use by a particular reader are hereby disclaimed.

13   Further product information is available on Proxim's web site at http://www.proxim.com/products/wifi/ap/ap2500/.

14   A VLAN traditionally is defined by specifying particular ports on one or more switches that belong to that VLAN. A dynamic VLAN is defined using the MAC addresses of the connecting devices: a particular MAC address will belong to a specific VLAN regardless of the port on which it connects. Conversely, a MAC address that is not defined as part of any VLAN will be denied access to the network.

15   Strong authentication is available for users, but it is difficult to extend this protection to devices. SecurID, for example, is a system consisting essentially of two components: a token (implemented either as hardware or software) assigned to a user and a server to authenticate the user's input. Both the token and server generate a new key every sixty seconds. The code generated by the user's

© SANS Institute 2004,   As part of GIAC practical repository.   Author retains full rights.

of administrative time and cost, and providing for exceptions, similarly, is burdensome. Each company must make its own assessment as to whether the frequency of anticipated connections justifies taking these extra steps.

Strengthening Security.   Servers behind the firewall can be protected against anonymous attacks following the well-documented techniques for "hardening" computers exposed directly to the public internet.[16]   Steps such as turning off unneeded services, closing unused ports, and keeping the patch level up-to-date (subject to reliability concerns) are equally applicable to internal as to more publicly exposed servers. Because servers tend to house a company's most important information resources, the additional investment of time required to secure them against potential threats from within their own segment (or in the event that the firewall fails or is bypassed[17]) are well justified.

Employee workstations can benefit from hardening as well, but the case for improved security must be balanced against significantly greater administrative costs, both in time required to apply ever-more-frequent software updates, and in potentially decreased utility of tools used for central administration.[18]   Software firewalls might be appropriate for some devices, particularly laptop computers, but due to the likelihood of increased support requests to respond to alerts or reduced functionality, many companies will not want to deploy them widely.

An intermediate step that may be appropriate for many companies is to segment the LAN into multiple VLANs. By limiting the number of devices on each segment, and tightly controlling access between segments, the company may well be able to limit the damage from untrusted software running amuck to a tolerable number of machines.

---

token, combined with the user's PIN, allows the SecurID server to reliably identify the user. Because SecurID authenticates users, rather than devices, users would need to be trained never to allow a visitor to use her token and pin.

[16]   A Google search for "hardening Windows 2000 server" returns over 80,000 references, including one from the source: Microsoft Corporation. "Microsoft Windows 2000 Security Hardening Guide." Apr. 11, 2003. URL: http://www.microsoft.com/technet/security/prodtech/win2000/win2khg/default.asp (Jan. 1, 2004).

[17]   Peer-to-peer (P2P) file sharing software, if connected to a folder on the server, may advertise the availability of files on the server to other members of the network. Due to flaws in various P2P programs, it may be possible for third parties to navigate outside of the folder the user intended to share and obtain access to critical files.

[18]   For example, while null sessions on employee workstations can be blocked with a change to the registry, this may interfere with certain management tools that connect anonymously with workstations.

While VLANs require care in configuration,[19] and are not airtight against a determined attacker,[20] the separation they afford may well be sufficient for trusted visitors.

Summary. Following the principle of defense-in-depth, a company should take a layered approach to minimizing the threat of untrusted device connections. Steps to consider include:

- A written policy against such connections, combined with user education;

- A convenient alternative for providing internet access for untrusted devices;

- To the greatest extent possible, hardening servers and keeping them patched up to date;

- To the greatest extent possible, hardening employee workstations and keeping them patched up to date;

- If feasible, segmenting the LAN into multiple VLANs to limit the damage resulting from the connection of an untrusted device to any one segment; and

- If additional control is required, eliminating the implicit trust of devices connected to the LAN and requiring additional credentials to access the network.

Other security threats, such as the risks posed by untrusted *persons* potentially gaining access to the wired or wireless LAN, might better justify the greater expense and administrative burden of the latter steps on this list. However, most businesses do not find such incidents sufficiently probable to take those steps. Adding the everyday scenario of accommodating trusted visitors to the assessment may tip the scales in favor of transitioning to those network designs.

Protection of the Public
After isolating the network's most critical assets from the trusted visitor's untrusted device, the company also should consider how its connection to the public internet is being used by that device. Malware installed on the visitor's computer might launch an attack against third party computers, or send unsolicited e-mail, or "share" copyrighted materials, or download pornography, or take any number of actions that — if associated with the company — could be highly embarrassing at best, and grounds for legal liability at worst.[21] Furthermore, a visitor using a wireless NIC may unknowingly create a conduit for third parties, through the company's wired or wireless visitor segment, to the

---

19   See, for example, Ou, George. "Implementing VLAN trunking." June 2003. URL: http://www.lanarchitect.net/Articles/VLANTrunking/Implementation/ (Jan. 1, 2004).

20   See, for example, Taylor, David. "Are there Vulnerabilites [sic] in VLAN Implementations?" July 12, 2000. URL: http://www.sans.org/resources/idfaq/vlan.php (Jan. 1, 2004).

21   The author of this paper, an attorney, expresses no opinion on this issue.

8

internet. These third parties may well not share the good intentions of your trusted visitor.

Potential Threats from Visitor Connections. Most companies have only a limited number of "public" IP addresses, and therefore use a range of private addresses on their networks and Network Address Translation (NAT) to enable communications between devices on the LAN and the outside world. At the gateway between the company's network and the internet, then, a firewall or router replaces the private IP address in an outgoing packet with its own address, one of the public IP addresses assigned to the company. In a less trusted segment, a company might use public IP addresses to simplify access to mail and web servers, and any other devices made available to the public. Whether visitor access is configured using NAT or via a distinct public IP address, the implication is clear: an address associated with the company will appear in every packet your trusted visitor sends. By consulting various logs, third parties will be able to trace such packets back to the company.

If the company permits its visitors unfiltered access to the internet, then the consequences for others could be serious. In addition to such specialized malware as Blaster and SQL Slammer, the visitor might be or — after accessing e-mail or the web — become infected by a mass mailing virus that proliferates messages across the internet. In addition to its interest in good internet citizenship, the company should be concerned that it might bear some legal responsibility for the consequences of providing the transport for such attacks.

Peer-to-peer (P2P) file sharing software will, at a minimum, burden a company's internet connection with substantial search traffic, and it may present grave security issues due both to technical flaws in the software and the frequent distribution of malware in the guise of desirable content.[22] Companies also should remember that the Recording Industry Association of America (RIAA) is monitoring file-sharing networks and contacting institutions that it believes may be hosting unauthorized copies of copyrighted music. Even if a company hosting trusted visitors ultimately were vindicated, substantial legal fees and potential negative publicity would follow any serious allegation of copyright infringement.

To limit the range of potential uses of the internet, companies should admonish visitors that access is being provided to facilitate their business with the company, and require a signed acknowledgement that they will not use the connection for other purposes, such as posting on public bulletin boards, viewing online entertainment, and so forth. But visitors are not always aware of, or able to control, the software running on their computers. Therefore, additional technical measures will be necessary to limit visitors' use of the company pipe.

---

[22] Windows Peer to Peer File Sharing ranks ninth on SANS' list of the two ten Windows vulnerabilities. SANS Institute, The. "The Twenty Most Critical internet Security Vulnerabilities (Updated) ~ The Experts Consensus." Version 4.0, Oct. 8, 2003. URL: http://www.sans.org/top20/#w9 (Jan. 1, 2004).

Filtering Outbound Traffic. The first step in filtering traffic is to close every unneeded port, that is, every port not required to provide essential services. Determining which services are essential may be a matter of some debate. Is AOL Instance Messenger or RealAudio a business-critical service? Each company will need to find an appropriate balance between maximum security and a manageable number of special requests. Presumably the balance will be similar for employees and for trusted visitors, although there may well be practical differences in their needs.

A second possibility is signature-based filtering. Because many programs can communicate over port 80, or other ports that are likely to kept open, a company may have to inspect outbound packets to determine whether they match the "signatures" of prohibited applications. Such signatures, and instructions for using them, have been posted on the internet,[23] and purpose-built commercial products now target both instant messaging applications and the P2P networks.[24]

One issue that requires special care is the handling of e-mail access. Some visitors may desire RPC access to a Microsoft Exchange Server, while others might want to "borrow" the company's SMTP server. To the extent possible, both of these suggestions should be resisted. Most visitors should be able to obtain access to Exchange mailboxes using a web browser or, in the latest version, RPC-over-HTTP. This is far safer than opening up the segment for general RPC traffic. Most visitors also should be able to send mail using authenticated access to their own mail servers, or using a generic SMTP host provided by the company's ISP. Creating a general "open relay" to accommodate trusted visitors is an invitation to future problems, particularly on a wireless network.

Filtering Outbound and Inbound Content. In addition to the relatively lightweight content scanning involved in checking for application signatures, the company might want to prevent the transmission and retrieval of pornography or other materials considered objectionable in the workplace. Because it is unlikely that a visitor could be asked to install such filters on his machine, the company would have to filter web (and possibly other) traffic at the gateway. A number of popular firewall products integrate with WebSense Enterprise, a filtering application that allows companies to block access to sites in the supplied database — which is divided into numerous categories — or sites

---

[23]   See, for example, Ballard, Josh. "File Sharing Technologies." Nov. 11, 2003.
        URL: http://www.oofle.com/filesharing.php (Jan. 1, 2004). Laura Chappell
        explains how to set up filter for signatures in Novell's Borderguard product in:
        Chappell, Laura. "Security Alert: Just Say Gno!" Novell Connection. Sept. 2001.
        URL: http://www.novell.com/connectionmagazine/2001/09/gnutella91.pdf (Jan. 1,
        2004).

[24]   For example, SurfControl plc offers SurfControl Instant Message Filter, described
        in further details at http://www.surfcontrol.com/products/im/.

whose URLs contain selected keywords.[25] WebSense does not actually attempt to interpret or modify the contents of a web page, and therefore is not a foolproof solution to offensive content. Conversely, some content filtering solutions have been known to overreach and block access to "legitimate" pages. The company must balance effectiveness against the potential offense taken by visitors to the "this page violates our policy" message.

Finally, the company should not rely on others upstream to perform its filtering. In the case of the Blaster and SQL Slammer worms, or any other software with a recognizable signature, the company's internet access provider may have put in place egress filters to prevent such packets from reaching their destination. This is no substitute for an internal barrier: the company's agreement with its access provider may well provide for suspension or termination of service in the event that attacks are issuing from the company's network. Accordingly, to avoid a breach of contract and the potential downtime and serious headaches that would follow, the company should take care to control such traffic and not rely on provider filtering to address these threats.

Summary. Following the principle of defense-in-depth, a company should take a layered approach to minimizing the risks of visitor access to the internet. Steps to consider include:

- A written policy, acknowledged by visitors prior to connecting, that access is for specific business purposes only;

- Limiting outbound traffic to necessary ports;

- A strict rule against allowing visitors to use the company's SMTP server to send mail;

- If feasible, signature-based packet filtering that limits use of peer-to-peer and instant messaging networks; and

- If feasible, content filtering designed to prevent the download and display of pornographic and other potentially offensive materials.

Because most of these steps could apply equally to employees and visitors alike, a company might already have much of the required infrastructure in place. If a company is considering these issues for the first time, it should strongly consider taking the same measures in the more trusted zones as well as in the visitor segment.

Security and Privacy of the Visitor
With all of the potential problems that a visitor's untrusted device can create for a company, it may be galling to consider the possibility that the company may have a degree of responsibility to protect its visitor from similar attacks. If a client's computer is

---

[25]   Further product information is available on the company's web site at
         http://www.websense.com/products/about/Enterprise/.

damaged or disabled by attacks from unfiltered employee workstations, or from the internet, or if her expectations of privacy are violated by covert monitoring of her internet usage, the company may do more harm than good to its relationship with its client by providing a network connection.

A good starting point is to provide protections similar to those given employees. In the case of the enterprise firewall, this may be trivial or even automatic, but other layers of defense might be impossible to replicate. For example, the license agreement for an enterprise antivirus system might not permit the company to install it on visitors' computers. Must the company provide such protection in another way? Without expressing a legal opinion on the matter, I find it doubtful that the company has such an obligation. Most business users should be familiar with the need to install antivirus software on their computers and to keep it up-to-date. Similarly, the company might offer its employees protection against "spyware" that it cannot easily extend to visitors. If the visitor restricts her activities to the business purpose for which access was sought, the risk of installing spyware appears quite low.

Regardless of the extent to which the company extends protections to the visitor, it would be well advised to obtain the visitor's acknowledgement of the risks of using the internet and approval of any e-mail or other monitoring. This can be combined with the policy documentation described under the previous heading.

Conclusion
Companies are under more pressure than ever to allow visitors to connect to their network, whether for collaboration or to try to keep a customer happy. A sensible network design, clear policies, and attention to education can be combined to prevent the well-intentioned visitor from introducing serious risks into the network environment.

References

In addition to these more formal sources, the SANS course materials and various mailing lists aided in gaining an understanding of the practical impact of various incidents and technologies.

Blaster Worm and other RPC/DCOM Attacks

Dougherty, Chad, and Jeffrey Havrilla, Shawn Hernan, and Marty Lindner. "CERT® Advisory CA-2003-20 W32/Blaster worm." Aug. 14, 2003. URL: http://www.cert.org/advisories/CA-2003-20.html (Jan. 1, 2004). This article describes the DCOM vulnerability and mitigation options in more detail.

Finlay, Ian A. and Damon G. Morda. "Vulnerability Note VU#568148: Microsoft Windows RPC vulnerable to buffer overflow." Aug. 14, 2003. URL: http://www.kb.cert.org/vuls/id/568148 (Jan. 1, 2004).

Microsoft Corporation. "What You Should Know About the Blaster Worm and Its Variants." Aug. 22, 2003. URL: http://www.microsoft.com/security/incident/blast.asp (Jan. 1, 2004).

Microsoft Corporation. "PSS Security Response Team Alert - New Worm: W32.Blaster.worm." Sept. 10, 2003. URL: http://www.microsoft.com/technet/security/alerts/msblaster.asp (Jan. 1, 2004).

Te, Darwin. "WORM_AGOBOT.AZ" Trend Micro Virus Encyclopedia. Dec. 1, 2003. URL: http://www.trendmicro.com/vinfo/virusencyclo/default5.asp?VName=WORM_AGOBOT. AZ&VSect=T (Jan. 1, 2004).

Trend Micro Incorporated. "Solution 15898." Aug. 12, 2003. URL: http://kb.trendmicro.com/solutions/solutionDetail.asp?solutionID=15898 (Jan. 1, 2004). The Knowledgebase article describes and illustrates the shutdown problem in more detail.

SQL Slammer Worm and Related Vulnerabilities

Danyliw, Roman. "CERT® Advisory CA-2003-04 MS-SQL Server Worm." Jan. 27, 2003. URL: http://www.cert.org/advisories/CA-2003-04.html (Jan. 1, 2004).

Lanza, Jeffrey P. "Vulnerability Note VU#484891: Microsoft SQL Server 2000 contains stack buffer overflow in SQL Server Resolution Service." Mar. 26, 2003. URL: http://www.kb.cert.org/vuls/id/484891 (Jan. 1, 2004).

13

Leonhard, Woody. "TOP TIP: DO YOU HAVE MSDE?" Woody's Office Watch. Vol. 8, no. 4 (2003). URL: http://www.woodyswatch.com/office/archtemplate.asp?v8-n04 (Jan. 1, 2004).

Microsoft Corporation. "Data Access Downloads." Latest file dated Oct. 15, 2003. URL: http://msdn.microsoft.com/downloads/list/dataaccess.asp (Jan. 1, 2004).

Microsoft Corporation. "Finding and Fixing Slammer Vulnerabilities." Feb. 21, 2003. URL: http://www.microsoft.com/security/slammer.asp (Jan. 1, 2004).

Network Management and Configuration Issues

Ballard, Josh. "File Sharing Technologies." Nov. 11, 2003. URL: http://www.oofle.com/filesharing.php (Jan. 1, 2004).

Chappell, Laura. "Security Alert: Just Say Gno!" Novell Connection. Sept. 2001. URL: http://www.novell.com/connectionmagazine/2001/09/gnutella91.pdf (Jan. 1, 2004).

eWEEK Labs.   "Labs Answers Patch Management Questions." eWEEK. Sept. 8, 2003. URL: http://www.eweek.com/article2/0,4149,1257572,00.asp (Jan. 1, 2004).

Fontana, John. "How to handle patch management." Network World. Dec. 1, 2003. URL: http://www.nwfusion.com/research/2003/1201howtopatch.html (Jan. 1, 2004).

Lim, Keng Hoe. "Security Guidelines for Wireless LAN Implementation." Aug. 27, 2003. URL: http://www.sans.org/rr/papers/index.php?id=1233 (Jan. 1, 2004).

Ou, George. "Implementing VLAN trunking." June 2003. URL: http://www.lanarchitect.net/Articles/VLANTrunking/Implementation/ (Jan. 1, 2004).

SANS Institute, The. "The Twenty Most Critical internet Security Vulnerabilities (Updated) ~ The Experts Consensus." Version 4.0, Oct. 8, 2003. URL: http://www.sans.org/top20/#w9 (Jan. 1, 2004).

Taylor, David. "Are there Vulnerabilites [sic] in VLAN Implementations?" July 12, 2000. URL: http://www.sans.org/resources/idfaq/vlan.php (Jan. 1, 2004).

Other

The author also has drawn on his experience as a moderator in "Woody's Lounge," an online forum in which security issues and products are regularly discussed in a variety of contexts. The Lounge can be found at http://www.wopr.com/cgi-bin/w3t/wwwthreads.pl.