



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

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

**SANS Security Essentials
GSEC Practical Assignment
Version 1.2d**

Linda A. LeBlanc

To: certify@sans.org

**Linda_LeBlanc_gsec.doc
Linda_LeBlanc_gsec.zip**

**GSEC v1.2d
Portsmouth, NH**

© SANS Institute 2000 - 2002, Author retains full rights.

Automating Remote Windows NT Auditing: A Procedure in Three Parts

INTRODUCTION

The purpose of this paper is to discuss the availability of appropriate remote monitoring software for Windows NT environments. Two software packages, Radmin and MIT's new (beta) event-syslogger, will be examined. Native NT implementations will also be discussed. While the primary focus is not on securing an NT workstation or an environment at large, it will be addressed as it pertains to the larger issue of auditing.

In the relatively lax security of the default Windows NT environment, any seasoned network administrator will stress the need for auditing individual machines and the network as a whole to prevent compromises; whether deliberate or unintentional failures of existing security features. Unfortunately, Microsoft doesn't provide for any easily maintainable features to implement security auditing at the network level.

First, we will discuss securing the local area network, then we will look at the options available to conduct remote auditing for workstations and finally we will look at what specific messages in the auditing environment mean. We'll look at when to increase vigilance and when to take immediate action, and what types of action might be appropriate for given situations.

I. Setting up and securing the small office or department workstation:

We will be working with the basic assumption that Windows NT has just been installed on a fresh hard drive, and now necessary steps are being taken to secure the machine from intruders.

The first step is to install Microsoft's Security Patches. A general recommendation would be to install Service Packs 3, 5 and 6a and any hot fixes released after 6a. There is some discussion that once a higher level Service Pack is installed, all fixes for previous service packs are also installed. However, trouble with a specific brand of video card in a popular brand of PC led me to the discovery that "cascading" installs are sometimes required. It is certainly easier to install three service packs than to troubleshoot some arcane incompatibility without knowing what it might be. Since many of these types of files can be packaged into automatic installations, removing the necessity of sitting in front of the monitor and pressing the enter button 23 times, the onus of installing multiple service packs is minimal.

Obviously, the service pack list is subject to change as Windows is always being updated and service packs are released on a regular basis. Below is a short discussion of current hot fixes and what known vulnerabilities they attempt to address. At the end of this section, I have included links to the primary pages that contain either service packs or hot fixes and their attendant information pages.

Hot fixes are broken into two categories, Critical Updates and Recommended Updates. I have categorized some hot fixes by vulnerability. Some do not lend themselves to easy description and since each hot fix is described on Microsoft's website, will not be addressed here. Note that some available services in the Windows NT operating system were designed to default to a vulnerable state and as such have a great number of fixes to be applied. (IIS is a prime example.) Sometimes it is necessary to determine how indispensable a piece of software is in the face of security trade offs.

Privilege Elevation – malicious users gain inappropriate administrative access.

Denial of Service – malicious users prevent access to services and resources.

Malformed Requests – malicious users submit malformed requests to some services such as FPSE or TCP/IP causing denial of services.

Buffer Overflows – malicious users “overload” specific length fields in certain applications to run malicious code or take administrative control; it can also be used for denial of service.

Request Parsing – malicious users run malicious code by passing specific strings of commands between applications and the operating system.

Spoofing – malicious users fake credentials with the operating system and take control of specific services or resources. Also used to impersonate a valid user on websites with insecure cookie implementations.

Communications Issues – application interfaces that are not completely compatible with hardware constructs or operating system modules causing denial of specific communications or network services.

For WinNT 4.0 service packs and fixes:

<http://www.microsoft.com/downloads/default.asp?Search=Product&LangIDCODE=20%3Ben-us&Value=10018&OpSysID=252&Show=Alpha>

For WinNT Server service packs (some duplication):

<http://www.microsoft.com/ntserver/nts/downloads/default.asp>

Microsoft has recently added a page <http://windowsupdate.microsoft.com>, which has centralized updating materials for all Windows versions as well as updates for Internet Explorer. The only bad thing about this site is that you must be using IE4.0 or later to access it. All other browsers are redirected to the standard download site for the various OS's.

Neohapsis has also compiled a listing of MS technical support sites and elsewhere on their web page are outstanding discussions of information security at every level.

It's vitally important that a WinNT system be maintained with the latest hot fixes. A particular environment's applications and software might necessitate the use of certain hot fixes and preclude the use of others. Systems administrators must carefully consider the implications of not patching and the potential for damage versus a software incompatibility or increased level of end user complexity, when a conflict exists between existing infrastructure and new vulnerability patches. Examples are application specific hot fixes (for IIS as an example), third party applications that are incompatible with hot fixes (such as kerberos enabled applications) and particular hardware interface requirements. In particular, an enterprise that is running IIS would want to ensure that all IIS specific hot fixes are installed as soon as they become available. There are separate patches for the OS and application software so it is imperative that both sets are maintained. Some configurations of encryption software do not "play well" with aspects of NT, this can be exacerbated by new hot fixes if compatibility testing isn't done. Finally, local area networks that operate behind physical firewalls and have specific configuration issues that must be addressed when installing any type of communications software or hot fix.

Following the installation of the software and appropriate service packs and security updates, an NT workstation or server should be "locked down" against potential intrusion. This involves removing or disabling all unnecessary services, configuring user management policies and file sharing rights and permissions. This task is highly site specific and is best started with a written document authorized by management stating what delimiting factors are accepted. This policy statement can then be used as the basis for setting system policies for users, groups and machines, as well as setting file sharing permissions. Approved comprehensive policy documents protect the systems administrator when compromises occur.

Once the workstation has been installed and secured or locked down, it is important to bear in mind the continuing threat from outside. Viruses are the most frequently contracted security compromise. In their delivery package, they can bring with them the entire spectrum of other attacks, from trojan horses to worms. As hackers get more sophisticated in using email as a delivery tool, compromises of this sort will continue to increase. Antivirus software has become more user friendly in terms of installation and configuration. There are enough different vendors with credible software packages that other than to emphasize the need for virus protection, there isn't much in the way of generic information that can be provided. The most important thing to understand is that anti-virus software is the final piece to the security software puzzle.

II. Auditing Events on Workstations and the Server

Every NT installation comes with an event viewer in the administrative tools that allows the administrator to see operating system events. Types of events that are logged include: failed logins and locked out accounts, calls to networked printers, and backup termination

status if using service controlled software such as IBM's Tivoli backup and Tardis time synchronizer.

The systems log records almost every transaction between server and client. Print commands to network printers, printer installations and deletions, remote access attempts, and unauthenticated sessions between workstation and server. Some of these such as unauthenticated sessions will also have corresponding messages in the security log.

In the security log, looking at the event details (by double clicking on the given event) will provide specific information that can be used to track problems. This is where repeated attempted logons will be reported as well as account lockouts. The detail window will show which computer the attempt was originating from, which computer was being accessed and the domain involved. This provides clues to what type of activity is occurring. If it's the appropriate time for an authorized user to log in, and there are repeated failed attempts, and the domain name is incorrect. It is very likely that they aren't paying attention to what they are doing, not that someone else is trying to hack into the network. This is a simple example of the importance of understanding the nature of network traffic for system administrators.

One of the most common types of attacks are port scans. Unfortunately, WinNT does not provide any type of reporting from port scans and even more critical is the fact that a number of exploits can be run against insecure ports in the default NT environment.

Software that runs as a service will generally send log messages to the event viewer. Messages from applications such as time synchronizers and automated backup services are common. Automated updates of virus software will also report success or failure here depending on the vendor.

The event logs are found under the administrative tools in the start menu. It is necessary to configure the logs using the Log/Log Settings menu selection, to maximize their effectiveness for a given environment. Choices should be made based on the specific needs of the given workstation and network. If logs are going to be maintained on a daily basis, then small log sizes and automatic overwriting is an acceptable alternative. If logs are to be checked for long term trends, larger file sizes and manual deletion is recommended. Filters can be set under the View/Filter Events menu based on source, time, and other factors.

It is possible to view the event logs of any workstation within a domain. Anyone with access to the event viewer (not a difficult issue if you know where in Winnt/profiles to look), access to the remote workstation, and the appropriate permissions can view the event logs from their machine. Unless there are specific permissions set up in advance, basic users have the ability to view logs by default. Any type of hack that allows privilege elevation to administrator also allows access to logs remotely.

Regular monitoring of event logs over the network in this manner is tedious however. For every machine there are three logs to be viewed and each machine must be

connected to imposing an attendant time lag. As the size of the network increases this becomes a decreasingly viable option.

Remote administration packages are an alternative for auditing log traffic from a distance. Unfortunately, these have historically been plagued with security problems of their own ranging from trojan horses to poor encryption methods. Furthermore, they are not designed to facilitate auditing but simply provide a path to the log files. Some of these “remote tools” are little more than shareware packages with preset backdoors allowing malicious users to compromise your system without a trace. Others such as Back Office have so many known exploits available on the web it has a very small security confidence factor. There are also packages available which are simply glorified telnet programs that provide little or no security against network sniffing. Since the object is to create a secure environment and maintain it, unencrypted telnet is not an alternative. Poorly encrypted telnet is worse since it can engender a false sense of security.

Radmin

An article on TechRepublic's website (<http://www.techrepublic.com/article.jhtml?id=r00220001205jim03.htm>) discusses the installation and setup of Radmin and a second article (<http://www.techrepublic.com/article.jhtml?src=search&id=r00220001207jim01.htm>) discusses some of the features available in the package.

Remote Administrator is a software package from Famatech that provides remote administration tools to the Windows Operating System. Radmin relies on the video hook driver to negotiate the interface with the remote client, because of this Radmin is not compatible with other remote access packages such as PCAnywhere, LapLink or Timbuktu. Another drawback is that it is not compatible with NetMeeting 1.2 or higher, which is a default installation package on later versions of Windows. TechRepublic is careful to point out in its articles that remote administration poses a significant security risk because it exacerbates any type of network compromise.

Radmin does give you the option of granting remote access only to specific groups or users. Radmin allows the user to configure which IP addresses it will accept an Radmin connection from and defaults to a well-known port number for configuration behind firewalls. It comes with a large suite of graphical interfaces given a seamless Windows feel to it.

Based on discussions on the TechRepublic website, there appear to be several issues with Radmin as an alternative for remote monitoring. (Remember, this product was not designed with remote monitoring in mind but remote administration.) Lack of support from the vendor is a primary concern. If there are problems with the security of the software itself, and assistance is not forthcoming from the creators, the package is worse than useless. Weak encryption has been used in the software design process. Similar to telnet programs with weak or no encryption this can frequently lend a false sense of

security. Default settings in the software can cause the CPU to cycle close to maximum capacity. While this may seem to be a simple fix (resetting the delimiters) if the software package is compromised it is a simple matter to set the defaults thus causing a denial of service.

As usual, each system or network's needs must be weighed against the surrounding security environment. What is the absolute necessity for remote administration? Is there a more secure method to achieve the same ends?

MIT's beta-version remote system audit program

Until this time, automating log auditing on several Windows NT workstations has been one primarily made up of tedium in either physically visiting each workstation to collect log information or utilizing the server's ability to view workstations in the domain and viewing logs individually. Now, however, a new utility, currently in beta testing, is being developed by MIT's pismere group (<http://web.mit.edu/pismere>) that forwards log information to a host UNIX machine where it is handled by the syslogd daemon. There it can be evaluated using tools (such as grep and diff) that are more readily available in UNIX than in Windows.

The software has a gui interface (see figures 1-3) for configuring and filtering events to be reported. In its current form it comes in a Windows installer package (.msi file) that is approximately 280k. When installed it consists of five files that are primarily help and configuration files and runs as a service in the control panel. The default setting is automatic start up.

The gui handles the interface for the syslogd daemon by allowing the user to select the levels of notification priority, event types and facility. These are the standard pieces to an entry in the syslog.conf file on the host UNIX machine. The filter lets the user select what to report from which log and also allows for exclusionary filters. The parameters tab handles internal error messaging and event log clear messaging and will forward those to the remote host as well as standard log events. The final tab allows the user to deselect logs for monitoring. If a user is only interested in login attempts then the application log might be deselected to prevent filling the syslog file with irrelevant information.

From the opposite end, the syslogd must be configured appropriately to receive the incoming material from the WinNT machine. The daemon itself must be started in such a manner as to accept incoming log entries from remote machines. The default operation (to accept or not accept) is OS specific. Some systems require that the log file already exist before the daemon is started, while others will create the file with the first entry. Once logging has begun, the file can be viewed in any text editor or specific lines can be listed using grep. If a syslog file is to be kept for multiple days and saved as part of a cron job, the previous day's log can be "diffed" against the current log to show new entries.

The beta version is being tested on Solaris and AIX as well as a couple of other UNIX versions. A public release date is unknown at this time due to Institute licensing restrictions and development issues.




Figure 1: Set the IP address of the UNIX machine to which events should be forwarded. Create one filter for each type of event, or blanket coverage by sending all event types from each log.

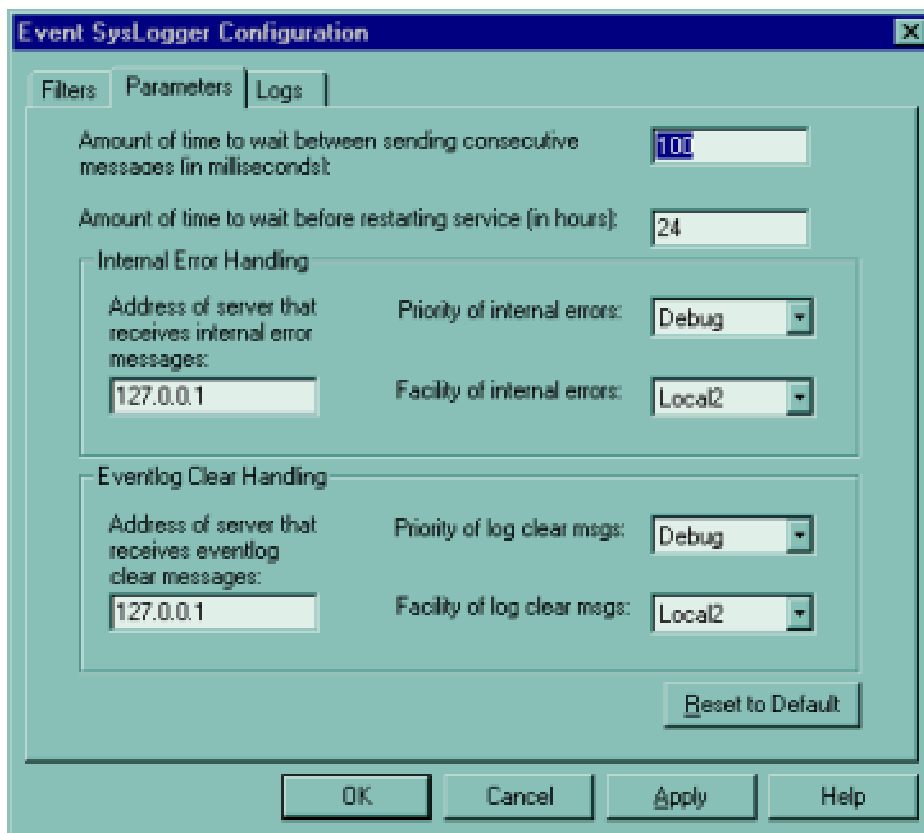


Figure 2: These are default settings for message handling. The event logger supports up to eight “local” and the other standard locations in conjunction with the syslogd protocol in UNIX. (Locations include mail, kernel, daemon, user, UUCP and others.)

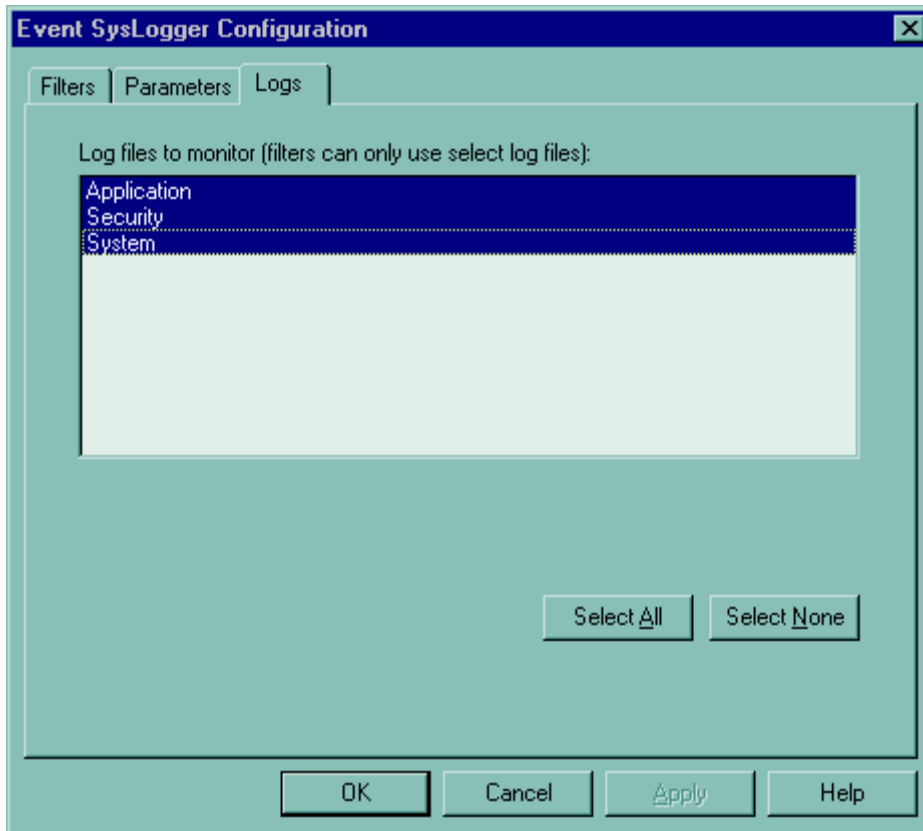


Figure 3: The standard three logs to choose from, default is all.

III. What now? Knowing what do with the information you collect.

Finally, what to do with the information collected. It isn't enough to read the logs it helps to know what to do with the information. As discussed above, some log information is more relevant than others. Knowing what to act on and what to watch are vital elements of a security program.

There are three types of event logs, application, security and system. Each of these logs has multiple levels of alert. The application log has information, warning, and error messages as does the system log. The security log has two types of messages, success audit and failure audit. In order to start advanced security log auditing it must be turned on through the user manager. In the start menu, under administrative tools, select user manager and then policies from the menu bar. Select audit and then check those items to be reported in the security log. A good starting point is to select success and failure for logon/logoff and restart/system/shutdown. This will provide a listing of who has accessed the system and who has attempted to access the system. It will also indicate when a system has been shutdown and restarted. This information can be of value when a virus has been introduced by floppy Account lockouts will also appear. Depending on how lockouts are handled by the system, if a username is compromised but not the password, repeated lockouts over a period of time are sometimes indicative of a brute

force password cracking attempt. Generally speaking, a legitimate user will call the systems administrator and ask for assistance.

Figures 4 and 5 shows an interesting anomaly about the success and failure audit reporting in Windows NT. In some instances an account lock out is considered a failure and in others it is reported as a success. The reason for this is that they are actually two different types of events. (Figure 4 is a logon/logoff category event and Figure 5 is an account management category event.) It is not that there are two different things happening: both are users attempting to log on to remote machines. One however, is an attempt to logon to the domain, and the other is an attempt to logon to a local system through the domain without proper permissions. This highlights a very important issue. System administrators need to be familiar with the type of traffic being generated on their systems in order to identify a serious threat to their network. Activity such as this needs to be identifiable in one of three basic categories: is the user confused about what he/she is trying to do? is the user trying to purposely get somewhere he/she isn't authorized to be? is someone else attempting to spoof the user's credentials to gain access to the system? Only by knowing and understanding the nature of the event logs will a systems administrator be able to make this determination.




Figure 4: A failed logon attempt with an account lock out. Shown as a failure audit event.

By consistently reviewing log activity, system administrators will recognize suspect traffic and be in a position to determine its seriousness and act. If the network is part of a twenty-four hour operation, admins will begin to recognize who should be logging on and when. If permissions have been set up to only allow logons during scheduled working hours, inappropriate logins will be easier to distinguish as well.

Repeated attempts to log in with an invalid user name (or as seen earlier a valid user id), is a good indication a network or system is being targeted by a hacker. A single

persistent attempt is most likely an attack by opportunity, where an individual has acquired an IP address or user and system name and is trying to make it work before going on to other targets. Repeated attempts over a period of time shows an interest in the targeted system or network in particular. This may be for no other reason on the part of the hacker than bragging rights, as is frequently the case at MIT. But this type of persistence needs to be stopped before a system is breached and damage is done.

Tracking an attempted (or successful) intruder with tools such as traceroute, will provide information that can be used to request assistance from the ISP sponsoring the activity. Ensuring that security measures are up to date and that legitimate users practice good security procedures helps lower the risk of a successful intrusion.

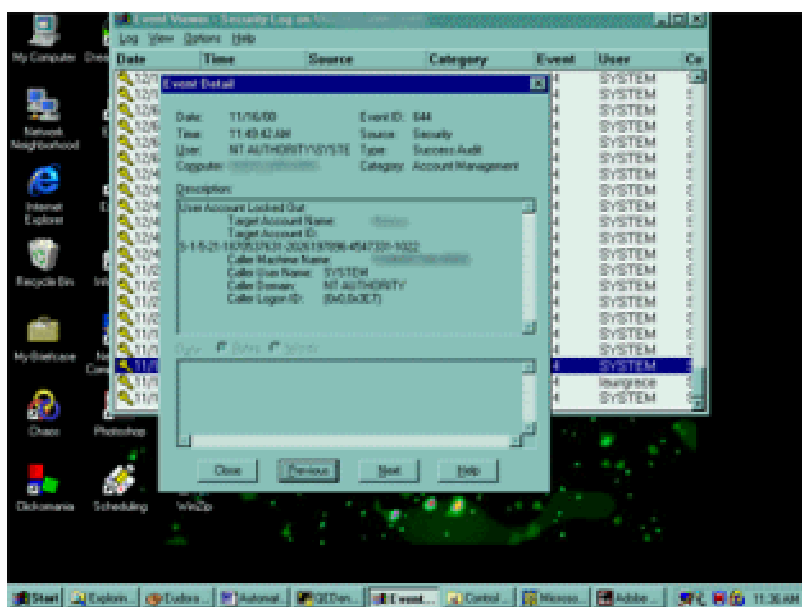


Figure 5: A failed logon attempt with an account lock out. Reported as a success audit event.

Cleaning up after a system has been compromised is always the most difficult. The validity of data and the integrity of software is questioned. Frequently, the most direct method to restoring a system is to format and reinstall. Utilizing data that has been backed up over a period of time is sometimes frustrating, but it is preferable to the distrust of compromised information. The loss of whatever data might have been changed after the most recent backup is slight compared to the possible loss of reaching conclusions based on corrupt data.

As systems administrators, cleaning up also involves examining the compromise itself and determining how best to avoid compromises of the same nature in the future. Investigating the methods used and the vulnerability exploited allows administrators to report such hacks to watchdog groups such as SANS and CIAC, so others can be aware of the possibility of attack. Policy can be drafted that will outline more secure procedures

to prevent repeated incidents. Software can be developed to close any existing holes in whatever application was utilized to gain access.

Conclusion

This is obviously a cyclic evolution. A system is installed, secured, audited and when bad things happen to it, it gets reinstalled, secured and more auditing goes on. Staying up to date with security patches for known vulnerabilities in both operating systems and application software makes the hacker's job that much harder. Maintaining permissions on a level in keeping with users' justified needs with the backing of management, helps prevent careless actions that could cause a compromise. The most important thing in this cycle however, is auditing. Habitual auditing of event logs is the only way to know what is going on a system or network. A compromised system can be discovered by the damage done or by the event logs. Event logs are simply more economical.

© SANS Institute 2000 - 2002, Author retains full rights.

References

Remote administration tools:

<http://www.famatech.com/> Radmin

<http://www.merxsoft.com/> REXX

<http://www.jpsoft.dk/products.php> Remote Home

Hacker Sites (good for seeing what you're fighting)

<http://www.wiretapped.net/>

<http://www.hoobie.net/>

<http://www.2600.com/>

Informational Sites

<http://www.slashdot.org/>

<http://www.sans.org/>

<http://www.ciac.org/ciac/>

<http://www.techrepublic.com/>

Pismere Group

<http://web.mit.edu/pismere>

WindowsNT Resource Kit, Microsoft Press, 1996