LANguard S.E.L.M.

Technical Evaluation

An NSS Group Report

First published November 2001 (V1.0)

Published by The NSS Group
Oakwood House, Wennington, Cambridgeshire, PE28 2LX, England

Tel : +44 (0)1487 773307
Fax : +44 (0)1487 773268
E-mail : info@NSS.co.uk
Internet : http://www.NSS.co.uk

TABLE OF CONTENTS

TABLE OF FIGURES

The NSS Group

The NSS Group is Europe's foremost independent security and network testing facility and consultancy organisation.

Based in Cambridgeshire, England, and with an advanced "super lab" and conference centre in the South of France, The NSS Group offers a range of specialist IT, networking and security-related services to vendors and end-user organisations throughout Europe and the United States.

The Group consists of three wholly-owned subsidiaries :

- *NSS Network Testing Laboratories*
- *Network Security Services*
- *NSS Consultancy Services*

NSS Network Testing Laboratories are available to vendors and end-users for fully independent testing of networking, communications and security hardware and software.

NSS Network Testing Laboratories also operates certification schemes for vendors and certification bodies, and currently provides certification of firewalls, VPN's, crypto products and PKI products.

Output from the labs, including detailed research reports, articles and white papers on the latest network-related technologies, are made available free of charge on our web site at http://www.nss.co.uk

The conference centre in Moux in the south of France is the ideal location for sales training, general seminars and product launches, and NSS can also provide technical writing services for sales, marketing and technical documentation.

Network Security Services provides a range of security-related services to vendors and end-users including security policy definition, firewall and VPN implementation, network security auditing and analysis, and penetration testing.

NSS Consultancy Services offers a range of network consultancy services including network design, strategy planning, directory design and Internet connectivity

INTRODUCTION

Whenever a company connects its network to the Internet, it opens up a whole can of worms regarding security. As the network grows, it will play host to numerous bugs and security loop holes of which you have never heard - but you can bet intruders have.

Many organisations are recognising the value of a good security policy to define what is and is not allowed in terms of network and Internet access Then they deploy a number of tools to enforce that security policy – usually in the form of a firewall or two.

Firewalls may be billed as commodity items, but the “shrink wrap” element certainly doesn’t extend to their configuration. A detailed knowledge of what a hacker can do and what should and shouldn’t be allowed through the firewall is required before embarking on the configuration adventure, and a slip of the mouse is all it takes to open up a hole big enough for your average hacker to drive the proverbial bus through. The problem is, a badly configured firewall can be worse than no firewall at all, since it will engender a false sense of security.

To protect an organisation completely, therefore, it is necessary to audit the network on a regular basis, and in order to achieve this, a whole new category of software has emerged in the last couple of years: ***Intrusion Detection Systems (IDS).***

When it comes to computer and network security, there are a number of analogies that can be drawn with the “real world”. Such analogies are particularly useful for answering such questions as “I already have a firewall, why do I need Intrusion Detection Systems as well?”.

Depending on how you approach the security of your home, for example, you may opt for high quality locks on your doors and windows. That will help to keep intruders out, and could be thought of as the equivalent of the firewall – perimeter defences. It’s nice to feel secure, but the determined burglar can often find ways around these measures. He can always throw a brick through your back window, for instance, and get in that way – or perhaps you simply forget to lock your door one day.

Once he is inside your home he is free to wreak havoc, perhaps making it obvious he has been there by stealing or wrecking things, or perhaps simply taking copies of any keys he finds so he can come and go later at his leisure. Whatever happens, you don’t want your first knowledge of the break-in to be when you return home to the ransacked contents.

That is why many people install a burglar alarm as well. Should the intruder gain access through the perimeter defences, the burglar alarm alerts you or your neighbours to the break in immediately, and provides an additional deterrent to the would-be thieves.

IDS, therefore, are the equivalent of the burglar alarm. To be used alongside firewalls, they are a recognition of the fact that you can never have a 100 per cent secure system. However, should someone be clever enough to breach your perimeter defences, you want to know about it as soon as possible. It would also be nice to know what they have been up to while they were inside too.

The final part of the analogy is the vulnerability scanner, which is the equivalent of your local crime prevention officer, testing your security and advising you of any potential weaknesses.

Within the IDS market place are four broad categories of product:

Vulnerability Assessment Scanners (VA)

Also known as “risk assessment products”, these take two forms.

The first is a passive scanner, which usually allows the network administrator to define a security policy for the machines on his network (perhaps a different policy for each operating system or type of server). The scanner then audits every machine automatically, producing a report that details exactly where each machines security settings differ from the defined policy and what needs to be done to fix the problem.

The second type of VA scanner takes quite a proactive stance – a sort of “hacker in a box” – providing a number of known attacks (Web server exploits, Denial of Service attacks, and so on) with which a network administrator can probe his or her network resources. By probing a network with one of these tools, the network administrator can often obtain a clear picture of potential weaknesses in his system, and even an indication as to how those weaknesses can be eliminated.

Some of these systems will make multiple passes of a network, using information gleaned on early passes to effect a more comprehensive attack in subsequent attempts. For example, a scanner may find an unprotected password file on a desktop machine in one pass. In the next pass it could actually use those passwords to attempt to gain access to protected resources as an administrator. You would be surprised how often this works!

Host IDS (HIDS)

These employ an agent that resides on each host to be monitored. The agent scrutinises event logs, critical system files and other auditable resources looking for unauthorised changes or suspicious patterns of activity. Whenever anything out of the ordinary is noticed, alerts or SNMP traps are raised automatically.

For instance, they will monitor attempted logins and take note of when an attempt is made to access an account with an incorrect password. If the attempt fails too many times within a short time span the system may conclude that someone is trying to gain access illegally and an alarm can be raised.

Another thing they can do is monitor the state of system and application files, or the Windows Registry. They do this by making an initial pass of a clean system and storing a condensed “snapshot” of how that system should look. If an intruder – or some sort of Trojan Horse - does manage to gain access to the system and make changes, the IDS will spot this (maybe not in real time) and raise an alert. There are systems on the market that specialise in this type of operation, and they tend to be referred to as File Integrity Assessment (FIA) products.

Most host-based systems tend to be reactive rather than proactive – that is they often have to wait until something has actually happened before they can raise the alarm. Some, however, attempt to be proactive, monitoring and intercepting system calls to the kernel or APIs in order to prevent attacks as well as log them. They may also monitor data streams and the environment specific to a particular application (file locations and registry settings, for example) in order to protect that application from generic attacks for which no “signature” yet exists.

These products sometimes go by the description “intrusion prevention”, since their aim is to stop intrusions dead, rather than simply report on them as or after they occur.

Network IDS (NIDS)

These monitor traffic on the wire in real time, examining packets in detail in order to spot denial of service attacks or dangerous payload before the packets reach their destination and do the damage. They do this by matching one or more packets against a database of known “attack signatures”. These databases are updated regularly by the vendors as new attacks are discovered.

When suspicious activity is noticed, a network based scanner is capable of both raising alerts and terminating the offending connection immediately (as are some host-based scanners). Some will also integrate with your firewall, automatically defining new rules to shut out the attacker in future.

Most of the network-based IDS available to date work in what is known as “promiscuous mode”. This means that they examine every packet on the local segment, whether or not those packets are destined for the IDS machine (much like a network monitor, such as Sniffer). Given that they have a lot of work to do in examining every single packet, they usually require a dedicated host on which to run due to their heavy use of system resources.

For instance, most attacks are not based on the contents of a single packet, but are made up of several, sometimes sent over a lengthy period of time. This means that the IDS has to store a number of packets in an internal buffer in order to compare groups of packets with its attack signature database. You will also need one per segment, since they are unable to see across switches or routers, and they have problems keeping up with heavily loaded Fast Ethernet segments (never mind Gigabit).

At the time of writing, one of the more recent developments in the Network IDS world is the introduction of a stateful approach to detection. As we have already noted, larger and larger databases of attack signatures, coupled with multiple minor variations of many of the attacks, means that a pure pattern-matching approach will find it increasingly difficult to keep up with heavily loaded networks.

By adopting a stateful architecture, IDS products are able to determine which attacks actually pose a genuine threat to the host system, and only those attacks are reported. This is achieved by combining protocol decoding with extensive state tables (similar to those used in stateful firewalls) in order to track active sessions on the wire.

For example, if a NIDS engine sees an attack for a Linux Web server aimed at a Microsoft Web server, it can ignore it. Nor will it simply trigger on suspicious patterns spotted on the wire unless a valid connection is first established between the source and destination machines. So if it sees the string “GET /cgi-bin/phf” on the wire (normally indicative of a PHF attack) it will not raise an alert unless there is already a valid HTTP session in place between the two machines.

Of course, this raises an interesting question regarding the role of an Intrusion Detection System. Admittedly, stateful IDS products are able to offer much higher performance than pure pattern matching (also known as “network grepping”) products, and will hopefully provide more accurate reporting with fewer false positives. But it is worth bearing in mind that by the time they raise an alert, there is a good chance that an attacker has already launched a successful attack against a host on the protected network – by definition, a stateful IDS will only alert on an attack that it regards as having a high probability of being successful.

Before a successful attack is perpetrated, however, many attackers may have tried numerous others. Would it not be better to be warned when the first IIS attack is launched against your Apache Web server, rather than waiting for the attacker to reach your Microsoft server farm?

Network Node IDS (NNIDS)

Recently we have seen a new type of “hybrid” IDS agent appear which overcomes some of the limitations of the network-based IDS.

This new agent works in a similar manner to the network-based IDS in that it takes packets from the wire and compares them against database entries. However, this new “micro agent” is only concerned with packets targeted at the network node on which it resides, thus giving rise to the term *Network Node IDS (NNIDS)*. The fact that it is no longer expected to examine every single packet on the wire means that it can be much faster and take less system resources, and this allows it to be installed on existing servers without imposing too much overhead.

It also makes it particularly suitable for heavily loaded segments, switched network environments, or VPN implementations with encrypted traffic on the wire, all areas where traditional network-based IDS have problems.

Obviously you will now need a number of these micro agents – one for every server you wish to protect – and they will all have to report back to a central console. Most systems may well opt for a combination of the two – micro-agents on individual servers in switched server farms, and traditional network-based IDS on less heavily used segments, where a single IDS can protect a large number of hosts.

For those with a reasonable security budget, we would recommend purchasing a firewall and at least one product from each of the above categories.

The firewall guards your perimeter, whilst the IDS’ monitor what is happening on your network, guarding against slip-ups by the firewall as well as internal mischief-makers.

Both host-based and network-based scanners are worth investing in, since they each have their own strengths. Network-based IDS will monitor the wire for suspect packets and are adept at spotting Denial of Service type attacks and unwelcome probes – usually from outside our network. Host-based systems, on the other hand, are watching the “crown jewels” – the actual data on the file servers, monitoring for inappropriate activities or changes to critical files from unauthorised sources.

Although network-based products seem to get most of the publicity, given that the FBI figures still point to over 70 per cent of hack attacks coming from inside a network, the host-based system can be particularly useful.

Finally, you can use the vulnerability scanner to continually test your defences and update your security policies accordingly. Only by continual vigilance and refinement will you stay one step ahead of the hackers.

LANGUARD S.E.L.M.

LANguard's Security Event Log Monitor (S.E.L.M.) is a Host-based Intrusion Detection System with a difference. Firstly, it does not rely on agent software on the hosts being monitored, and secondly, it monitors the Windows NT/2000 event logs only.

Windows NT/2000 provides the means to record all security-related events in its Security Event Logs. Logon activity, failed logons, supervisor activity, file access – they can all be logged in the Security Event Log. Unfortunately, there are a number of problems that make this data less than useful:

- **Configuration –** None of this data is recorded by default. It is the responsibility of the administrator to establish an audit policy and specify which events should be monitored. Novice administrators are rarely up to this task.
- **Reactive –** There is no real-time monitoring or alerting capability built in to the operating system. This means that critical security events can often remain unnoticed until an administrator decides to check the log files.
- **Lack of analysis –** There are no reporting or analysis tools built in to the operating system, making correlation of events extremely difficult. The only tool available for examining events is the Windows Event Viewer – a rudimentary tool offering very basic viewing facilities and little else. Likewise, there is no means to consolidate alerts from multiple machines – each log file is stand alone. Finally, there is no automatic archival capability.
- **Lack of detail –** Each event in the Security Event Log is assigned a cryptic numeric code rather than a meaningful description. Nor is there any detail on the possible cause of the alert. This makes analysis difficult for the inexperienced administrator.

LANguard S.E.L.M. provides the means to monitor multiple Security Event Logs around a corporate network in real time, providing instant notification of critical security events. Configuration and management is performed from a central console, as is log consolidation and detailed analysis.

Architecture

LANguard S.E.L.M works by retrieving on a real time or schedule basis all the events from the server and workstation event logs. It then analyses each event and determines the security level of the event, alerting the administrator when necessary (depending on how critical the event is). All events are then archived automatically, offering subsequent centralised reporting and reviewing of security events.

LANguard S.E.L.M. consists of the following modules:

- **Collector Agent -** This module retrieves all the events from the remote hosts being monitored. The Collector Agent is a high performance service that can retrieve events from many computers using an advanced scheduling algorithm based on computer security levels. It is not necessary to install this on each host to be monitored - a single, central Collector Agent is required, and this uses native Win32 APIs to collect security events from other computers on the network.

- **Alerter Agent -** This module alerts the administrator to security events. Alerts can be transmitted via email, SMS or pager (using an email-to-SMS or email-to-pager service).
- **Archiver Agent -** This module saves each and every event record which is read and processed by the LANguard S.E.L.M. collector agent to a centralised database back-end, which can either be an MS Access database or an MS SQL Server. Storage in a standard database format allows the administrator not only to use the built-in reporting capabilities, but also third-party reporting tools such as Crystal Reports.
- **Event Viewer -** The LANguard Event Viewer combines all features found in the standard Windows Event Viewer, but adds much more advanced searching, filtering and event management options, providing increased scope for detailed analysis (especially since LANguard events provide much more detail than standard Windows events).
- **LANguard S.E.L.M. Configuration –** This module allows the administrator to configure which machines are to be monitored, as well as set the operational parameters for the other LANguard S.E.L.M. components.
- **LANguard S.E.L.M. Reporter –** This module allows the administrator to create numerous reports based on the events which have been collected and processed by the Collector Agent.

Microsoft's Message Queue technology is used to maintain high performance communication between the internal components of LANguard S.E.L.M..

As with any other Host-based IDS, LANguard S.E.L.M. is not impaired in its operation by the use of high-speed switching infrastructures or the use of encryption across the network. In fact, LANguard benefits greatly from the fastest LAN possible, since the faster the connection, the faster the log retrieval.

Installation

Depending on the components already installed on your network (LANguard requires at least MMC 1.2, IE 5, MDAC 2.5 and Microsoft Message Queuing Services (MSMQ) in order to operate) the installation of LANguard S.E.L.M. can be extremely straightforward. Most users of recent Microsoft operating systems will find that MSMQ is the only item that needs to be installed, and this is covered in plenty of detail in the excellent documentation.

Having covered the preliminaries, the installation of LANguard S.E.L.M. itself is quick and simple, since it involves only a single, central console. During installation, the administrator is taken through an initialisation wizard, at the end of which the product is ready to begin collecting event log data.

Before that can happen, of course, it is necessary to configure each host with a security policy that specifies which events and objects should be logged. Auditing can be set to monitor both operating system events – logons and logoffs – as well as individual object accesses. An object in windows NT/2000 is anything from a file system object to a registry key to a printer. It is thus possible to monitor critical programs such as *cmd.exe*, *net.exe*, *tftp.exe* and *ping.exe* for use in unusual circumstances (either or both successful or failed access attempts can be monitored).

For example, when an attacker runs *cmd.exe* using the UNICODE exploit, it is actually run by the Internet Guest Account (IUSR_machinename). On the other hand, a successful buffer overflow exploit may leave an attacker running *cmd.exe* as the SYSTEM account. Since neither of these users should legitimately be running *cmd.exe*, LANguard can log such events and inform the administrator immediately.

Once again, the excellent documentation takes the administrator step-by-step through configuring an appropriate audit policy and applying it to a large number of workstations and servers across a corporate network. In an NT4 environment, this can be a painful and laborious process. Administrators of Windows 2000 networks using Active Directory, however, can simplify this process by using Group Policies.

Configuration

Everything in LANguard S.E.L.M. is controlled from the *Configuration* MMC snap-in module.

The first task is to select the hosts that will be monitored by LANguard, and this can be done by entering IP addresses directly or by browsing the network. For each host monitored, it is possible to set a normal operational time, a notional security level (high, medium or low, where a domain controller might be high, and a user workstation low), a scanning schedule, and whether the event logs on the host should be purged after they are transferred.

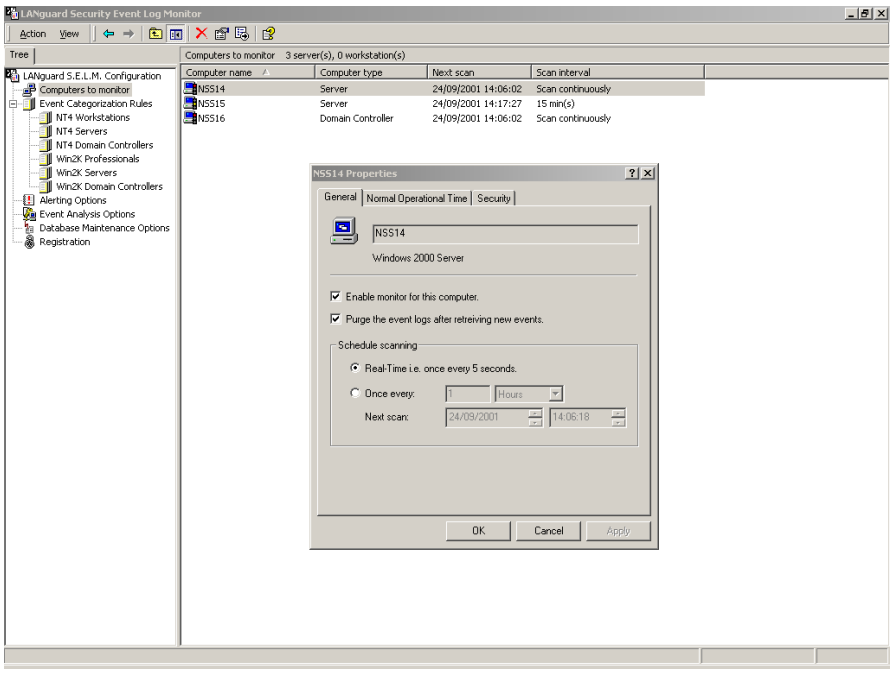

Figure 1 - Defining hosts to be monitored and scanning schedules

The scanning schedule can be “real time” (i.e. every five seconds) or at longer intervals, and every machine can have an individual schedule. This is useful, since a balance needs to be struck between operational efficiency and network resources – every machine on the network transferring log files every five seconds might well create a Denial of Service attack of its own!

With individual schedules, however, it is possible to have critical machines – such as domain controllers or key eCommerce servers located in the DMZ – scanned every few seconds, and low priority user workstations scanned once or twice a day.

The *Event Categorisation Rules* are the heart of the LANguard S.E.L.M. security policy, since they specify how a combination of time, host security level and event ID will be treated – whether the resulting alert will be allocated a *Critical*, *High*, *Medium* or *Low* status, and whether the administrator should be notified immediately or the event simply recorded for later analysis.

Different sets of rules are provided for domain controllers, servers and workstations, and for NT4 and Windows 2000. The same event can thus be interpreted differently depending on the role of the host and the operating system installed. Sensible default policies are provided out of the box, but it is a very straightforward matter to amend these to suit individual requirements.

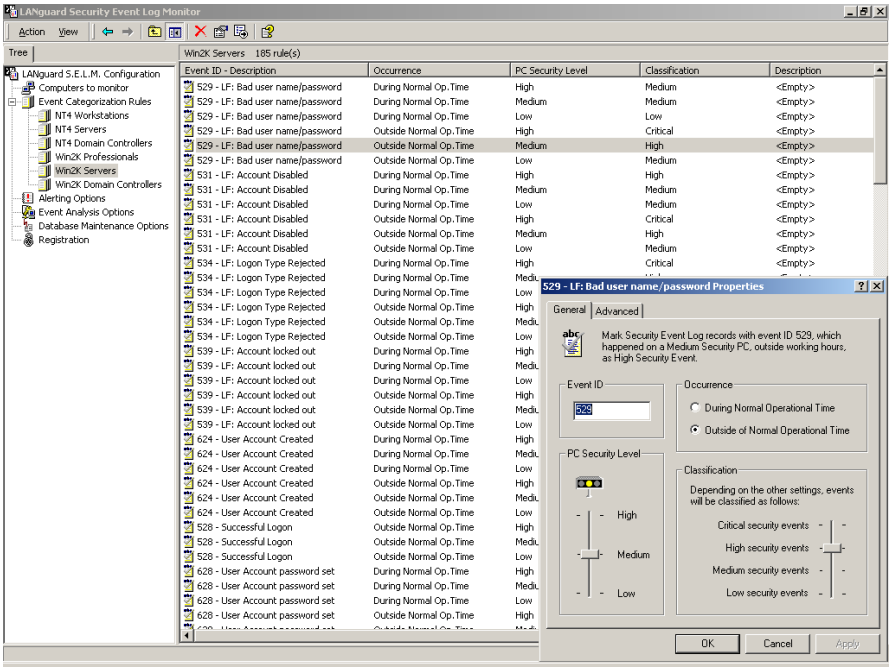

Figure 2 - Event Categorisation Rules

An example of where these distinctions are important is in the case of *network logons*. When a connection is made to a computer over the network – to access a shared folder, for example – Window 2000 logs event ID 540, but NT4 logs event ID 528 with logon type 2. Network logons to domain controllers and servers are obviously a common occurrence and shouldn't be regarded as suspicious during normal working hours.

However, in a network with centralised servers, users would not normally access resources on other workstations, and consequently the same event on a workstation could be regarded as suspicious. Thus a combination of the security level of the host, the time of day, the operating system and the event ID are brought together intelligently by LANguard S.E.L.M. in order to make in informed decision on whether a suspicious activity is in progress.

The only real means of instant alerting is via e-mail (SMS and pager alerts are provided via the appropriate e-mail gateways). It would be nice to see Winpopup message and SNMP alerting capabilities added to the product. Low priority alerts are simply recorded in the central database for later reporting and analysis purposes.

In normal operation, the Collector Agent uses the scanning schedule to determine how often it needs to retrieve the security event log entries from each machine it is monitoring. As it retrieves the event log entries, it compares each one with the Event Categorisation Rules in order to determine the severity of the alert and takes the appropriate action. Once this has been completed, it simply moves on to the next host, and continues in a loop until the Collector Agent service is stopped or until the central console host is shut down. Clearly it does not matter if the central console is unavailable for any length of time (except that critical alerts may not be raised immediately) since all events are stored in the individual security event logs until the Collector Agent operation is resumed.

A well-orchestrated attack on a poorly configured system could conceivably gain administrator authority on the computer and clear the log before LANguard's next scheduled collection. However Windows faithfully records a specific (and non-deletable) event whenever the log is cleared (even if auditing has been disabled) which is classified by S.E.L.M. as a critical event on all types of computers by default.

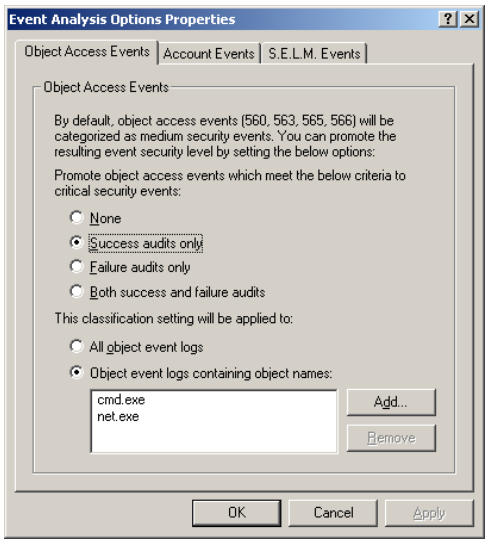

Figure 3 - Auditing object access events

It is important that the maximum size of log files and the collection intervals are set such that the possibility of log files filling (and therefore losing valuable entries) is minimised.

Note that transmissions between Collector Agent and the hosts being monitored are performed using native Win32 API calls and are not encrypted in any way. LANguard S.E.L.M. is thus not suitable for running across a public network without secure third-party encryption in place (recent Windows operating systems include VPN capabilities, of course). GFI recommend that LANguard is not run across slow WAN links anyway, preferring instead that a separate LANguard S.E.L.M. Collector Agent be installed at each site.

By deploying LANguard S.E.L.M. to monitor all workstations, member servers and domain controllers, it is relatively easy to obtain a comprehensive picture of network activity. In an enterprise-wide deployment, LANguard S.E.L.M.'s default rules numerous potentially suspicious activities including:

- *Failed logons*
- *Account lockouts*
- *After-hours account creation and group membership changes*
- *After hours logins to high security systems*
- *Entry gained to user workstations from over the network*
- *Audit policy changes*
- *Cleared security logs*
- *Successful or failed file access including specific file names*

Reporting and Analysis

The intelligent processing of security log events is only part of the LANguard S.E.L.M. story. One of the biggest problems with Windows event logs is the lack of any useful means of reporting and detailed analysis.

One major improvement with LANguard is the replacement of the Windows Event Viewer with the LANguard S.E.L.M. Event Viewer. Actually, it does not replace it, since the Windows Event Viewer is still available for use, though no one is likely to use it once the S.E.L.M. Event Viewer is installed (although S.E.L.M. can write events back to the individual local event logs should this be required).

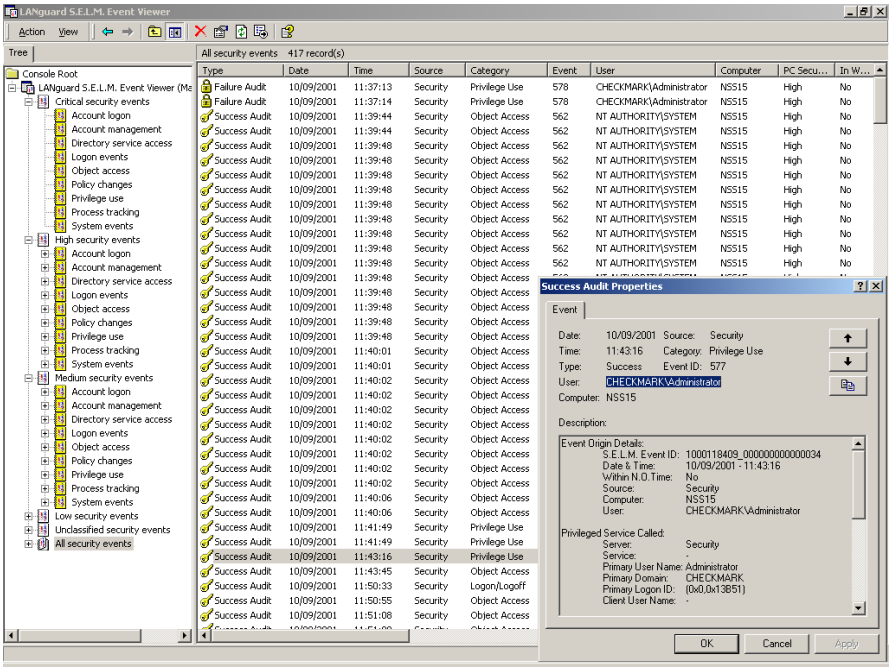

Figure 4 - The LANguard S.E.L.M. Event Viewer

Although it looks very similar to the Windows Event Viewer, the S.E.L.M. Event Viewer offers numerous improvements. The first is that it presents all of the fields comprising each event in column format, thus allowing searching, filtering and sorting on the contents of event log entries.

In addition, a hierarchical tree view of the event log entries is available providing views by severity level: *Critical*, *High*, *Medium*, *Low*, *Unclassified* and *all* levels. New nodes can easily be added to this hierarchical view to provide filtered views of the event log, filtering on any of the event data fields such as date, time, type, user, computer, security level, and so on.

The biggest difference offered by the S.E.L.M. Event Viewer, however, is the amount of additional information provided about each event. In addition to the usual date, time, computer details and purely numeric event ID, we are now offered a full description of the event (corresponding to the event ID), possible causes, and recommendations. This instantly makes the event log a much more useful tool.

Complex queries can also be applied to virtually all the information which is carried within an event record, thus increasing the forensic analysis capabilities of the administrator.

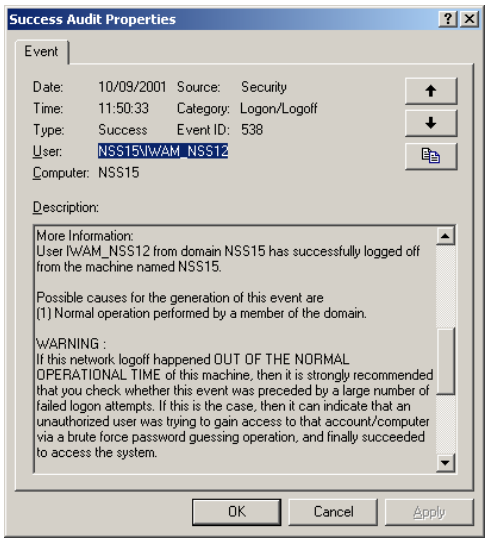

Figure 5 - Detailed information for each event

For a more detailed analysis of events, the LANguard S.E.L.M. Reporter provides a large number of pre-defined reports, including:

- ***User-based reports –*** *failed, successful or first logon events, ranking users with the highest number of generated events on top.*
- ***Machine-based reports –*** *displaying events generated on particular machines per user, such as user logons, account lockouts, object access and account activity.*

The Reporter window is divided into two panes – the left pane displays the *report types*, such as user reports in percentages, first/last day user event reports, computer event reports, and computer event reports in percentages. Report types are similar to *templates* on which individual reports are based. For example, all reports generated under the type *User Reports In Percentages* will have a similar layout. However any number of reports can be created under the same report type, each with different settings such as different periods, different events, and so on. Individual reports are listed under the report type branches in a hierarchical tree format, but initially these will be empty.

The right pane will display the contents or properties of the item selected in the left pane. As reports are run, the report output is also displayed in the right-hand pane, and the properties (report contents, reporting period, and so on) are saved for future use, allowing the same report to be easily run over and over again.

The report output is clear and easy to read, and bear in mind that since the S.E.L.M. data is stored in a SQL database it is also possible to use third-party reporting tools to create completely custom reports if required.

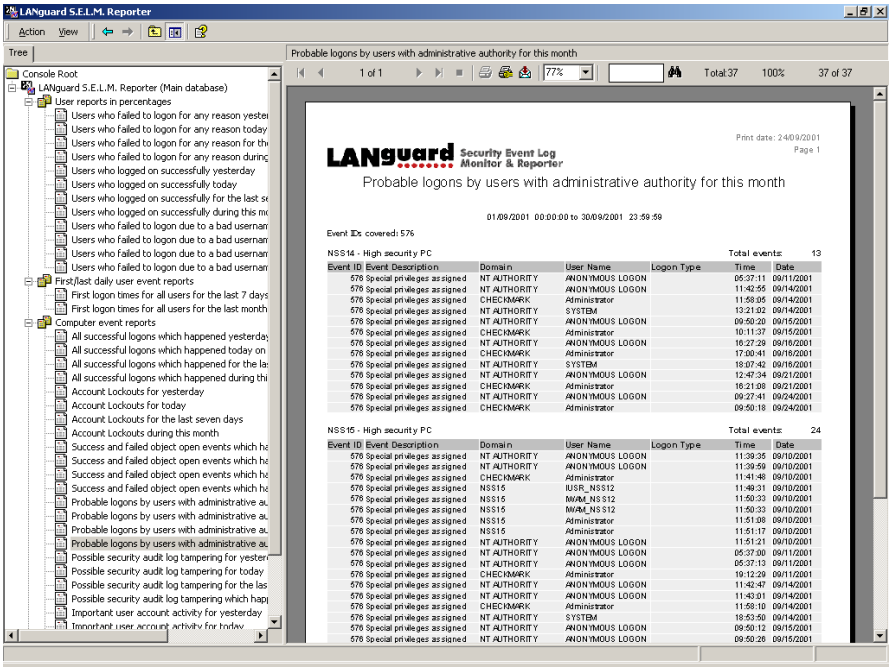

Figure 6 - Viewing reports

Verdict

Although LANguard S.E.L.M. itself is incredibly easy to install and configure, the initial configuration of the network to support the required auditing policy may be more problematical for the novice administrator, particularly where the advantages of Active Directory Group Policies can not be realised.

It is also vitally important to ensure that the S.E.L.M. installation itself is secure and that log file sizes and collection intervals are optimised to prevent excessive network traffic and the possibility of lost data due to log files filling to capacity.

However, once these issues have been dealt with, the daily operation of LANguard S.E.L.M. is virtually idiot proof.

Any number of individual event logs can be retrieved from around the network using a customised collection schedule, and the events within those logs examined and compared against a set of rules. This allows LANguard S.E.L.M. to apply an intelligence to event log processing that has been missing until now.

In addition to intelligent classification of security events and real-time alerting, however, S.E.L.M. also provides automated archival and analysis tools, with far more detail being provided to the administrator via the LANguard S.E.L.M. Event Viewer and Reporter than has ever been available via the standard Windows tools.

For the first time, the Windows administrator can perform detailed analysis of Windows security event logs, as well as gaining an instant appreciation of the severity of events thanks to much more detailed and understandable event descriptions coupled with real-time alerts.

Contact Details

Company: GFI Software Ltd.

E-mail: sales@gfi.com

Internet: http://www.gfi.com

Address:
Communications House
Mediterranean Street, SGN 07,
St Julians,
Malta.

Tel: +35 6 382418
Fax: +35 6 382419

UK office: GFI Software Ltd.

E-mail: sales@gfi.co.uk

Address:
5 Princeton Mews
167-169 London Road,
Kingston-upon-Thames,
Surrey KT2 6PT, UK.

Tel: +44 (0)20 8546 0640
Fax: +44 (0)20 8546 0741

US office: GFI USA

E-mail: sales@gfi.com

Address:
105 Towerview Ct.
Cary
NC 27513
USA

Tel: +1 (888) 2 GFIFAX
Fax: +1 (919) 388-5621