Internet-based security incidents and the potential for false alarms

M.P.Evans and S.M.Furnell

Network Research Group
School of Electronic, Communication and Electrical Engineering
University of Plymouth
Plymouth
United Kingdom

email: nrg@jack.see.plym.ac.uk

Abstract

The provision of security represents an important consideration for Internet-based systems and media reports suggest that certain sites experience a significant number of remote attacks. However, the validity of such claims depends upon what type of event is considered to constitute an attack and whether an attempted security breach would be the only explanation for its occurrence.

The paper explains the background to experimental work that was conducted with the aim of measuring aspects of the WWW (specifically the average lifetime of a web link and the impact of the 'Millennium Bug'), but which inadvertently caused two perceived security breaches on remote systems. The paper explains the nature of these incidents and considers why, when over 700,000 IP addresses were randomly sampled in the experimental study, only two sites considered the activity to be an attempt to breach their security. It is concluded that, while the appropriate protection of Internet-based systems is undoubtedly of importance, the problems experienced during the experimental study suggest a lack of uniformity in what different organisations will class as a security breach.

Keywords

Internet Security, Denial of Service, Attack Detection, World Wide Web

Introduction

The Internet and the World Wide Web (WWW) represent two of the most significant technological developments of the last century. The Internet offers a previously unimaginable potential for connectivity - the possession of two IP addresses enables a seemingly direct connection between two systems, no matter where they may be physically located. With network connectivity and appropriate software utilities, it is possible to determine the existence of a remote system (and, to some extent, the services

it can provide), even if you cannot actually log into it. From a security perspective, this can represent a problem, as the mere knowledge of a system's existence serves to make it a potential target.

The paper describes the problems posed by Internet-based attacks and the resulting attitude that they demand on the part of network security administrators. The discussion then proceeds to consider the problem that this effective state of 'connection paranoia' may represent for normal Internet users wishing to conduct innocent network activities. This is supported by the example of an experimental study into web site longevity that the authors attempted, but which was complicated by inadvertent security side-effects.

Internet security issues and examples

Computer-based crime and abuse has been recognised as an unfortunate side-effect of the information technology revolution for many years, with computer hackers (crackers) and viruses representing the most widely recognised causes of the problem (Audit Commission, 1998; Furnell and Warren, 1996). However, the Internet has significantly enhanced the threat posed by attackers, giving rise to a new series of opportunities for abuse. For example, a useful tool in the hacker arsenal is the port scanner, which enables the inspection of a remote system to determine what software and services it is running. This knowledge can be useful in facilitating a more direct attack if known vulnerabilities of the discovered software can then be exploited. Probably the most well-known example of a scanner program is the Security Administrator Tool for Analyzing Networks (SATAN), the rationale for which is described by Farmer and Venema (1993). SATAN, in common with other scanner software, offers the facility to find a machine or network, find out what services are being run and then automatically test those services for known security holes. As the full name suggests, SATAN offers a useful tool to system administrators wishing to ensure the security of their network. However, the public availability serves to make it an accessible tool for hackers as well. SATAN has been followed by a wealth of other tools, such as NetInfo, PortPro, IPprober and HackTek, from both the commercial and hacker communities.

Even if an attacker cannot directly enter a system, they can often still cause problems through denial of service (DoS) attacks. Such an attack is one in which a target system is rendered inaccessible or unusable and will generally involve the consumption of a system's resources (such as memory, storage space and/or network ports), such that it is unable to provide adequate service for its legitimate users. At a minimum, the end result can be inconvenience for legitimate users attempting to use the system or wishing to gain access to it. Such incidents represent a growing problem on the Internet and account for a significant proportion of reported security problems (CERT, 2000). The attacks themselves can take many forms and are frequently system-specific, taking advantage of known vulnerabilities on a particular target platform (e.g. running a particular operating system, web server or other system software). However, some generic categories of DoS attack can also be identified and two Internet-based examples are briefly described below (Escamilla, 1998).

– **"Ping of death"**

  Relies on a flaw in some TCP/IP stack implementations. The attack relates to the handling of unusually and illegally large ping packets (which some systems, e.g. Windows NT and 95, can generate). Remote systems receiving such packets can crash as the memory allocated for storing packets overflows. The attack does not affect all systems in the same way – some systems will crash, others will remain unaffected.

– **SYN flooding**

  Exploits the fact that establishing a connection with the TCP protocol involves a 3-phase handshake between the systems, as follows:

  1. Connecting host sends a SYN packet to the receiving host
  2. Receiving host sends a SYN|ACK packet back
  3. Connecting host responds with an ACK packet

  In a SYN flood attack, an attacking host sends many SYN packets and does not respond with an ACK to the SYN|ACKs. As the receiving host is waiting for more and more ACKs, the buffer queue will fill up. Ultimately, the receiving machine can no longer accept legitimate connections.

The facility to launch such attacks can be found as a 'feature' of several cracker toolkits, such as HackTek. Given that they can be automated in this way, the mounting of such attacks does not require any skill or expertise on the part of the hacker (indeed, more dedicated hackers refer to those who rely upon such techniques as 'script kiddies'). As such, it has been conjectured that around 90% of hacking is conducted by people using such methods (Akass, 2000). This is not to underestimate the serious consequences that attacks such as denial of service may have. For example, in the context of a system used for a sensitive application, such as providing access to patient records or controlling direct care provision in a healthcare environment, any unavailability or performance degradation could have significant consequences.

The sheer range of companies and organisations now represented on the Internet and the WWW means that there are a significant number of high-profile targets for potential attackers. One such organisation is the Pentagon in the United States, which experienced a total of 5,844 recorded attacks in 1998. Although this is itself significant (averaging over 16 attacks per day), the number recorded in 1999 was significantly higher and, by November, well over 18,000 attacks had been identified (Daily Telegraph, 1999). More recently, a distributed DoS attack was experienced by a series of major Internet sites in February 2000. The attacks affected a number of notable and popular sites, including Amazon.com (books), eBay (online auctions) and CNN (news) and had a significant impact. For example, it was reported that, within a few minutes, the Amazon.com web site became 98.5% unavailable to legitimate users (McCullagh and Arent, 2000). In addition, Keynote, a US-based Internet monitoring company, reported that the average

performance of the Internet was degraded by "as much as 26.8%" (Keynote, 2000). These statistics serve to reinforce the significance of the problem that DoS attacks can represent.

As a consequence of factors such as those described above, many organisations are sensitive to the threat of Internet-based attacks and take measures to guard against them. Success here relies upon being able to accurately detect the signs of an attempted attack in progress. However, as the next section illustrates, it is not always possible to reliably differentiate between attacks and other forms of network activity.

Measuring the Web - a problematic study

The authors have first-hand experience in causing Internet security alerts - although in an entirely innocent context, through the conduct of Web-based experimental research unrelated to security. The nature of the intended experiment, and the problems that subsequently arose, are described in the sections that follow.

Experimental background and procedure

The experiment that caused the problems was involved in measuring aspects of the web. Specifically, the experiment was designed to determine both the average lifetime of a web link (that is, the average period of time before the resource pointed to was removed), and the impact of the 'Millennium Bug' on the Internet. Both of these activities were conducted as part of a wider research programme relating to web-based content migration (Evans et al. 1999).

The principal aim of the experiment was to determine the average lifetime of a web link. To do this, it was necessary to collect a large sample of links at random. Each link would then be tested periodically, and the date and time would be recorded if and when the link failed (i.e. the resource pointed to by the link could no longer be found). To ensure the randomness of the links, the authors attempted to compile a database of web servers from a large list of Internet servers chosen at random. Once the list had been compiled, the intention was to let a web crawler search through the various HTML documents on the web servers, and choose for itself a set of links at random. This would ensure no bias had crept into the link selection process. Unfortunately, however, the experiment never reached this stage.

A secondary aim of the experiment was to determine the impact of the 'Millennium Bug' on the Internet. The list of randomly selected servers was being compiled between October and November 1999. As such, it would be trivial to extend the experiment and periodically test the state of these servers after midnight on 1 January 2000. Although not related to the primary aim of the experiment, this would be interesting research for little cost.

The intended experiment comprised four different stages:

1. Compile a list of random servers.
2. From this list, compile a list of random web links.
3. Periodically determine the state of each server.
4. Periodically determine the state of each link.

However, only the first stage was ever reached, as several unanticipated side effects caused two unintentional security incidents, which forced the experiment to end prematurely. The side effects were a direct result of the design of stage 1 of the experiment interacting unpredictably with a server's firewall.

In stage 1, a list of random web servers had to be compiled. This was achieved by randomly generating an IP address and sending a simple HTTP HEAD message on port 80 to attempt to retrieve the header of the default HTML page of the server. If a response was received, the server's IP address and domain name were recorded as belonging to a web server. If no response was received, however, the machine was pinged. If a response was received, the machine's IP address and domain name were recorded as belonging to an Internet server (no attempt was made to determine which type of Internet service the machine was providing). This allowed the experiment to determine the effect of the Millennium Bug on the Internet in general and on the web in particular. If no response was received from the ping, the IP address was noted as being dead, and played no further part in the experiment.

Because of the unpredictable nature of the network at the University, the experiment was designed to compensate for any network problems that occurred. For example, the network would sometimes go down for a few seconds or several minutes before returning to normal. At other times, its speed dropped significantly, due to the network loading of the University's LAN. Because of this, some ping messages seemed to take a long time before a reply was received; equally, some HTTP requests had to be sent more than once before a reply was received. To ensure that a server really was down when no reply was received, the HTTP HEAD request was sent more than once, and both the Time To Live of each ping, and the number of pings sent, was increased to allow for delayed responses.

The experiment was designed as a multi-threaded application, with 100 threads operating at a time, each of which contacted one server at random. The design was fully tested on the University's servers before being allowed to sample the whole Internet. However, over time, it became apparent that the experimental procedure was resulting in perceived security problems and, during the course of two months, two formal complaints were received from organisations whose systems had been randomly targeted. The details of these incidents, and the remedial actions that were taken, are described in the sections that follow.

Security incident one

Figure 1 below contains the text of an email received from the administrators of one of the randomly selected systems. It should be noted that elements of the figure have been

edited to preserve the anonymity of the specific machines involved at the authors' site and of the affected remote domain. As such, organisational name references have been deleted and address elements have been replaced with 'xx' where appropriate.

```
We have detected unfriendly network activity, directed at our machines,
from 141.163.xx.xx [xx.xx.plymouth.ac.uk]. The activity, which began at
19:04 EDT (GMT-0400) on October 27, 1999 was a port scan (137) and
pinging of many addresses in our subnets (xx.xx.xx.xx, xx.xx.xx.xx).

This type of activity is not desired on the [Deleted] domain and is
monitored frequently. Please advise your system managers and users that
this activity should stop immediately.

----------------------------
[Deleted]
Computer & Network Security
```

Figure 1: Email received concerning first perceived intrusion

From the perspective of the remote system, it is clear that the activity of the program was considered comparable to that of port scanning tools such as SATAN. However, in the context of the experimental study, the use of port 137 (the NetBIOS name resolution service) was not a feature that had been explicitly included in the program. Rather, it is a feature of Windows NT Server (the operating system that was used to host the experiment) that NetBIOS is used to resolve a name, followed by the Internet DNS, in response to a call to the 'gethostbyaddr()' function. This is the default behaviour of the operating system.

Although completely innocent in this case, port scans are widely used by hackers as reconnaissance, in an attempt to determine the services the victim's servers are providing. However, in this case the firewall software would have received one NetBIOS call to port 137, and one subsequent HTTP request to port 80 (although the latter was not mentioned in their e-mail communication). Further, the "pinging of many addresses" (ping sweep) that was detected listed just two IP addresses that were pinged. Both addresses were selected completely at random, and were not part of the same sub-net, or even looked as if they were remotely related. It was an unfortunate coincidence that both were pinged at the same time, and both happened to be owned by the same organisation.

Security incident two

The experiment was remounted, with NetBIOS support disabled on the originating NT system, and the ping configuration altered such that each ping packet had a timeout value of 1 second, and at most only 6 packets would ever be sent to one machine. The experiment was restarted, and for a time all went well. In fact, 700,000 different IP addresses had been tested over a period of two weeks, when the second and final

inadvertent security incident occurred. Figure 2 presents the email message received in this case. Note that the IP addresses listed have again been altered in order to hide the identities of the machines involved.

```
------------------------
Intrusion Attempt Report
------------------------

We have noticed the following behaviour originating
from IP addresses under your control:

ICMP denial of service attempt

The activity took place at approximately:

Dec 7 03:07 GMT

We consider this/these unauthorized attempt(s) to access
our networks as malicious in nature and hereby request that
you take steps to identify the person(s) involved and
arrange for this activity to halt immediately.

Here are some samples of the activity in question:

03:07:41.035397 141.163.xx.xx   xx.xx.xx.xx: icmp: echo request
03:07:41.036032 141.163.xx.xx   xx.xx.xx.xx: icmp: echo request
03:07:41.038980 141.163.xx.xx   xx.xx.xx.xx: icmp: echo request
03:07:41.039568 141.163.xx.xx   xx.xx.xx.xx: icmp: echo request
03:07:41.042556 141.163.xx.xx   xx.xx.xx.xx: icmp: echo request
03:07:41.043138 141.163.xx.xx   xx.xx.xx.xx: icmp: echo request

338 instances in 5 seconds

We further request that you reply back to us with the
Resolution achieved in this matter.
```

Figure 2: Email received concerning second perceived intrusion

At this point, it was decided that the experiment had attracted too much adverse attention and the University's Computing Service was understandably concerned that further problems could arise if it was to continue. As such, the mutual decision was taken to discontinue this element of the study.

The unfortunate situation in this case was the fact that the incident reported differed greatly from the behaviour of the experiment under testing. 338 ping messages were allegedly received in 5 seconds, when the code was explicitly written to send only 6. Indeed, the network traffic generated by the experiment had been extensively monitored before the experiment was restarted to ensure that no more than 6 packets were sent to any one machine. The University's network was used for the test, and indeed, only 6 packets were ever sent to each machine. However, the incident reported 338 such

messages. Whether this was a fault in the experiment or in the firewall on the remote machine that reported the incident is impossible to determine, as the experiment had to be discontinued and it is doubtful that the 'victim' would be willing to share their firewall's configuration details.

Discussion

In both incidents, the receipt of the complaint was immediately followed by corrective action and a written explanation of the experimental context. This was considered satisfactory and defused the possibility of further action.

Reflecting upon the experiences, it could be argued that in the first incident, the remote firewall software used was a little overzealous. The 'incident' comprised one name resolution request to port 137, one HTTP request to port 80, and several ping messages on two separate servers with widely differing IP addresses. It is debatable whether this should be considered to represent a threat. It should also be remembered that a contributing factor in the first incident was that the experimental software performed an unexpected task (i.e. a NetBIOS call) as a consequence of asking it to perform an intended function (i.e. a name resolution). In a sense, it could therefore be argued that the program author was the victim of an inadvertent 'Trojan Horse' effect. Without this prior action occurring, it is possible that the respondent organisation may not have perceived the subsequent pings to be part of a hostile attack.

In the second incident, the duration of each ping was only 5 seconds, which intuitively would not suggest that a Denial of Service attack was intended. It could be argued that this might have been an initial assault, designed to overload the firewall system and thereby enable exploitation of some other vulnerability. However, the fact that no further activity would have been apparent from the source IP address should have provided an indication that this was not the case.

In both incidents, the ping feature had been disabled by the firewall (a relatively standard practice, which is intended to guard against attacks such as 'Ping of Death'). This meant that the experiment continued sending ping requests in an attempt to determine whether the IP address was a live server or not, while the firewall silently monitored the requests, yet did not respond. In the absence of a response, the experimental software was unintentionally entrapped into appearing as a security threat.

With hindsight, however, it can also be argued that, from a security perspective, the practical approach taken by the experimental study was ill conceived. To select an IP address at random and then attempt to determine the state of the server can be seen to have the potential for mis-interpretation by a security conscious organisation. From the organisation's perspective, such a stream of traffic would have no obviously legitimate purpose and, therefore, by default would be regarded as suspicious. However, the fact that only two organisations flagged a problem during the period of the experimental study (during which over 700,000 random addresses were targeted in this way), gives a very

strong indication that organisations are not monitoring their security to a consistent degree. If it is argued that the two complainant organisations were correct to interpret the network activity as attacks, then it could also be considered that the other organisations were failing in their network security strategy. This assertion must, of course, be offset against the fact that different organisations will be dealing with systems and data of different levels of sensitivity, and therefore, in some cases, the required level of security may legitimately be lower. Having said that, it is unlikely that in a random sample only two organisations had data that they would consider sensitive.

The question remains, however, as to how to effectively monitor the growth of the Internet without upsetting somebody's security policy. As e-commerce continues to grow on the web, security is becoming more and more of an issue, with the result that firewalls are being configured ever more tightly. As has been shown, this can have the effect of seriously derailing entirely innocent applications, and has the potential to cause serious harm to the reputation of those involved (in this case, the University was advised that unless the experiment was discontinued, its network connection would risk being terminated).

In view of the practical experiences, it is worth re-examining the Pentagon attack figures cited earlier in the paper. Closer investigation reveals that the Pentagon's definition of an attack includes activities such as port scans and pings (Wayner, 1999). As such, the level of genuine abuse may not be as significant as first suggested by the bare statistics alone, as many of the incidents recorded may have been the result of activities that were not intended to breach security. This can be regarded as a counter-argument to the commonly held belief that the majority of computer-based crime goes unreported, due to fears of adverse publicity on the part of the affected organisations (Nycum and Parker, 1990). In this case, organisations may be over-stating their vulnerability to abuse. In addition, there has been at least one legal ruling (by the Norwegian supreme court) stating that probing of systems on the Internet, using techniques such as port scans, should not be considered illegal (Jones, 1998).

It can be argued that all parties involved in the practical incidents described emerged as losers from their experiences. The authors were unable to proceed with a potentially interesting experiment, whilst the remote organisations had effectively wasted resources in responding to false alarms. As such, there appears to be the need for some form of protocol through which applications such as the experiment detailed here can safely query a server on the Internet without upsetting its security arrangement. Such a protocol exists for the web in the form of the Robot Exclusion Standard (Koster, 1994). This is a mechanism through which a web server can define the permissible behaviour of 'visiting' software agents, such as a search engine's web crawler. These agents attempt to index the contents of a web server, but have the potential to cause unwelcome side effects. For example, they may flood the server with too many requests, or attempt to index areas in which they are not welcome, either for privacy reasons, or because their presence fools the server into thinking that they are a genuine user. To prevent this, the standard lets the server owner specify clear boundaries of good behaviour which the web crawler is expected to adhere to. The boundaries are encoded in a text file called Robots.txt, which

a web crawler should parse upon arrival at the server. Although the standard cannot enforce the behaviour of a web crawler, it provides an implicit contract between server owner and web crawler designer. Breaking the contract can lead to the exclusion of all traffic from the sub network from which the web crawler originated, or even legal disputes (Pallmann, 1999). As such, the Robot Exclusion Standard could make an ideal model from which to develop similar protocols for measuring the Internet. With so many different security policies in existence, it makes sense for a server to publish its policy of acceptable behaviour, rather than expect any visiting software agents to guess what that policy might be.

Conclusions

The provision of appropriate protection for Internet-based systems is undoubtedly of importance. It is necessary for systems to have 'frontline' defences in order to detect potential abuse and reduce the possibility of successful system penetration. From this perspective, the organisations that identified and responded to the activity of the experimental software can be commended for having effective security monitoring procedures that should also enable detection of genuine attacks. The fact that, in the cases described, the monitoring software caused false alarms can be excused in a security context, as this is preferable to allowing an intruder to penetrate or disrupt the system.

Incidents such as denial of service attacks represent a significant threat to Internet-based systems and, while they do not represent a direct threat to the confidentiality or integrity of data, they may be employed as a precursor to a more direct form of attack. In addition, the lack of system availability may itself represent a significant threat to individual or organisational well-being in many scenarios. As such, it is legitimate and advisable for organisations to take appropriate steps to protect their assets from such attack. However, there is also a clear need for experimental research to be performed on the Internet. Measurements on the size of the Internet and its growth are not just of academic relevance. Web masters and network managers need to know how much network traffic to expect, while business leaders need to know the importance of the Internet to their business. The techniques used to perform such measurements, however, share certain characteristics with those used by hackers. There is, therefore, the potential to trigger security alerts for innocent reasons, as the authors have discovered. This finding calls into question whether the number of 'security incidents' logged by certain organisations actually represent a realistic indication of their vulnerability to attack. The experiment discussed previously was applied to over 700,000 Internet addresses, yet only two security incidents were flagged and followed up, demonstrating a lack of uniformity in the security policies of different organisations and what they class as an attack.

In view of the above, the security needs of the Internet must be balanced with the experimental needs of the research community (and the activities of legitimate bots performing necessary services), if only to prevent future misunderstandings that could potentially lead to more embarrassing outcomes than those discussed here. The definition of a standard enabling server owners to define their acceptable behaviour

policy could prevent such situations from occurring, and lead to a more realistic measure of security incidents, reducing public fear caused by potentially exaggerated attack statistics.

References

Akass, C. (2000), "On the straight and narrow – not", *Personal Computer World*, February 2000: 57.

Audit Commission. (1998), *Ghost in the Machine*. Audit Commission Publications, UK. February 1998. ISBN 1-86240-056-3.

CERT. (2000), "CERT® Advisory CA-2000-01 Denial-of-Service Developments", CERT Coordination Center and the Federal Computer Incident Response Capability (FedCIRC). 3 January 2000. http://www.cert.org/advisories/CA-2000-01.html.

Daily Telegraph. (1999), "Pentagon under cyber-seige", The Daily Telegraph – Connected supplement, 11 November 1999: 2.

Escamilla, T. (1998), *Intrusion Detection*. Wiley Computer Publishing. ISBN 0-471-29000-9.

Evans, M.P.; Phippen, A.D.; Mueller, G.; Furnell, S.M.; Sanders, P.W. and Reynolds, P.L. (1999), "Strategies for Content Migration on the World Wide Web", *Internet Research*, Vol. 9, No. 1: 25-34.

Farmer, D. and Venema, W. (1993), "Improving the Security of Your Site by Breaking Into it". http://www.fish.com/~zen/satan/satan-demo/admin-guide-to-cracking.html

Furnell, S.M. and Warren, M.J. (1996), "Computer abuse: Vandalising the information society", *Internet Research*, Vol. 7, No.1.:61-66.

Jones, C. (1998), "Let the Web Server Beware", Wired News, 23 December 1998. http://www.wired.com/news/politics/0,1283,17024,00.html

Keynote. 2000. "Denial of Service Attacks This Week Degraded Internet Performance Overall According to Keynote". Keynote Press Release, 12 February 2000. http://www.keynote.com.

Koster, M. (1994), "A Standard for Robot Exclusion". The Web Robots Pages. http://info.webcrawler.com/mak/projects/robots/norobots.html

McCullagh, D. and Arent, L. 2000. "A Frenzy of Hacking Attacks", Wired News, 9 February 2000. http://www.wired.com/news/print/0,1294,34234,00.html.

Nycum, S.H. and Parker, D.B. (1990), "Prosecutorial experience with state computer crime laws in the United States", in *Security and Protection in Information Systems*, A.Grissonnanche (Ed.), Elsevier Science Publishers B.V., North-Holland: 307-319.

Pallmann, D. (1999), *Programming Bots, Spiders, and Intelligent Agents in Microsoft Visual C++*. Microsoft Press. ISBN 0-7356-0565-3.

Wayner, P. (1999), "Hacker 'Attacks' on Military Networks May Be Closer to Espionage", New York Times, 8 March 1999. http://www.nytimes.com/library/tech/99/03/cyber/articles/08defense.html