TRAINING NETWORK MANAGERS TO RECOGNISE INTRUSION ATTACKS

Colin Pattinson, Kemal Hajdarevic
School of Computing, Leeds Metropolitan University, Leeds, UK

Keywords: Training network managers, denial of service attack, simulation

Abstract: One of the major challenges facing the e-Business community, and the broader telecommunications network world, is the threat of electronic attack. Of the sub-categories of such attacks, the denial of service attack, in which the intruder's objective is to prevent legitimate users from accessing some or all of an organisation's computing resource, regularly creates headlines in the popular press. Whilst significant research effort is being expended on the development of automated tools to recognise such attacks, for many businesses (particularly the small business sector) network management (including security and intrusion detection) is the responsibility of an individual employee (the "network manager"), among whose responsibilities is the observation and monitoring of network behaviour, and who will be expected to monitor data, detect the signs of intrusion, and take action, ideally before the attack has taken effect. Traditionally, this skill has developed through a hands-on process, learning "normal" behaviour, using this knowledge to detect anomalies, undertaking further investigation to determine more details of the cause. This will involve interaction with the "live" network, and the first experience of an attack will be when it actually occurs. This is counter to good training practice, in which a trainee will have had experience of "problem situations" in a controlled environment, and will have the opportunity to develop their responses, review actions and repeat the activity, so that when the situation occurs "for real", responses are semi-automatic. This paper describes a simulation-based training tool in which student network managers experience the symptoms and effects of a denial of service attack and practice their responses in a controlled environment, with the aim of preparing them more effectively for the time they meet such an attack in reality.

1 INTRODUCTION

Reports of attacks on computer networks, aimed at disrupting the service provided to users of those networks, and the consequential effects on business and consumer confidence are commonplace in the popular press. These attacks vary in their method, purpose and effectiveness, but one of the most frequently launched attacks – probably because it is also one of the most effective and easy to bring about – is denial of service (DOS).

DOS has many flavours and variants, but the common factor is that one or more nodes within a network are subjected to some form of traffic which requires the use of resources to handle it, for example by directing an excessive volume of a particular request message, to which the host is obliged to respond. Resources (processing time and buffer space) allocated to this task are resources which cannot be allocated elsewhere, and, *in extremis*, there are insufficient resources left to meet the service requirements of legitimate users. The result is that those users are denied service, so the attack has succeeded.

The (human) network manager is the member of staff with responsibility for maintaining the service levels of the network, and it therefore appropriate to consider how such an individual might be assisted in identifying that an attack is underway, and the ways in which the effects can be mitigated, if not prevented. Consider first the identification phase, a problem of this nature will be noticed because it presents variations from the "normal" behaviour of the system. In the best case, this will be early enough to allow defences to be deployed against the attacker. In the worst case, the anomaly will be some form of failure. Clearly therefore, early detection of anomalies is an important factor. Whilst the development of automated anomaly detectors is a fertile research area (see below), in many smaller

Pattinson C. and Hajdarevic K. (2004).
TRAINING NETWORK MANAGERS TO RECOGNISE INTRUSION ATTACKS.
In *Proceedings of the First International Conference on E-Business and Telecommunication Networks*, pages 269-274
DOI: 10.5220/0001386202690274

organisations the network manager will carry out this task manually, using the current generation of network monitoring tools. In essence, the network manager will be expected to learn through observation the "normal" signs of operation, and also to spot any variation from "normality" in sufficient time to react effectively. The most popular network monitoring tools make use of the Simple Network Management Protocol (SNMP) in one of its variants, and monitoring in this situation is a case of observing trends in the various counters and related data, with anomalies being revealed as changes in those trends.

This is a skill which develops over time and with experience, but unfortunately, the most common experience is that the onset of a real attack is the first practical experience many network managers have, they "learn through doing" in a very direct fashion. We argue that it is desirable to have a more structured approach to learning, and we borrow from the "flight simulator" style of training to propose a means by which novice network managers can practice their skills in a controlled environment. We have developed such a training tool, described in detail elsewhere, this paper describes how we have incorporated the effects of a DOS attack into that tool, to give a student network manager the opportunity to recognise these symptoms and to practice their response to the attack.

Our network management simulator tool uses a combination of synthetically varied SNMP Management Information Base (MIB-2) objects and external "messages from other systems and users" (themselves produced according to a pre-set script), to create the effect of different network scenarios. The details of the simulation process has been described elsewhere (Pattinson, 2000), where we also make it clear that we use an existing SNMP-capable network management platform to collect and display data from our simulated SNMP data set. In the case described in this paper, we show how these input stimuli are manipulated to create the impression of a DOS attack.
The structure of this paper is as follows: we state the need to detect and respond to DOS attacks, including a review of current work in automated detection systems; we then present the results of our analysis of the symptoms of such attacks, with particular reference to the effect on observed SNMP data; we then describe how these symptoms are represented in our simulation system, together with the other related symptoms; we conclude with an analysis of the effectiveness of our approach and a description of further work.

2 EARLY ATTACK DETECTION

Early detection of attacks or other anomalies which could harm a network in some way is the subject of much current research.

An automated system which is able to recognise and prevent attacks or other anomalies, and work with already standardized mechanisms for network management, is preferable to those systems and projects whose goal is to design new hardware equipment such as D-WARD (UCLA, 2002). Some research groups, such as MINDS, (University of Minnesota, 2003) use more than one mechanism to gather data about network behaviour such as sniffers, syslog, tcpdump or net-flow (capturing packet headers) while this offers better understanding of what is really going on in each packet, it makes the overall system more complex – data has to be gathered and processed from different sources. Dokas et al (2001) report work on the detection of anomalies by setting threshold values and treating an overstepping of this threshold as an event worthy of further investigation. The major concern with these automated systems is that of reliability – in this context meaning identifying (and not responding to) false alarms, while remaining sufficiently alert to detect genuine problems early enough to mount a defence. Typically, this leads to the MINDS approach of gathering more, and different, data, but for our work we are interested in the situation where the (human) Network Manager relies on manual interpretation of data, and on SNMP MIB-2 data alone.

Our research has used only MIB-2 instances to discover anomalies in the behaviour of network protocols and resources. The reason for this is that MIB-2 is a standardised database already installed on many pieces of network equipment. If all network devices (routers, switches, PCs etc.) are MIB-2 capable it might be possible to detect an attack or other anomaly in its early development stages, allowing enough time to prevent any serious damage to the attacked systems. Many attacks come from outside a secured network and are delivered through network gateways.

Clearly, the collection of counters in MIB-2 (located at various points across the network) will show the impact of any attempt to overload a device within the network. Although this impact will be most obvious at the node being attacked (as we show below), there will also be (probably lesser) impacts on the near neighbour nodes, and particularly on the gateway router devices. In view of the possible

failure of the "attacked" node, we must also pay attention to the way in which neighbouring nodes represent the developing attack, and include those in our simulation.

We wished to develop a simulation in which changes in behaviour of specific MIB-2 instances could allow a network manager to practice their interpretation skills and detect the onset of an attack. We begin this process by presenting an analysis of the effects of a DOS attack on a small network of the type common to the environment in which our trainees are likely to work. We then present a structured simulation and show, through a number of screen shots, what a trainee user will experience when presented with this simulation.

3 SIMULATING AN ATTACK THE TCP/SYN EXAMPLE

The TCP/SYN attack (CERT 1999) is a well-known variant of DOS, in which an attacker launches a series of TCP/SYN packets (connection establishment requests) at a target. These packets typically each have a different (spoofed) source IP address, which is unreachable from the target. The "rules" of TCP connection establishment require the target to issue a SYN-ACK packet in response, and to allocate resources (buffer and table space etc.) for the "new" connection. Whilst this is sufficient to tie up resource at the target, and is therefore a DOS attack in its own right, a consequence of attempting to send a packet to an unreachable IP address is that the next hop (router) will return an ICMP "Destination unreachable" message to the sender, so further consuming buffer space. Any existing (legitimate) connections to or from the target will experience a gradual reduction in their throughput, and may, as the attack develops, timeout. New legitimate attempts to connect are likely to fail via timeout before completion.
Therefore, in order to represent the effect of such an attack on node 1 of our simulated network (presented at figure 1), we need to vary the MIB-2 objects representing the behaviour of node 1 in the manner shown in table 1. Figure 2 shows how this behaviour might be traced during operation.

4 THE EFFECTS ON OTHER HOSTS WITHIN THE NETWORK

As noted above, other users and systems will also be affected by a DOS attack, in this case, the major impact will be on those users who are, or would wish to be, connected to the attacked node. We represent this by reducing their "traffic flow" during the onset of the attack, and by terminating some "connections" as the effects of the attack worsen.

The simulation also allows us to generate "pop-up" messages on the trainee's console reporting (or complaining of) connection loss / slowdown. Other affected activities are those devices which share a connection (bandwidth resource) with the attacked node. In this case, this is particularly relevant to other nodes sharing the router node X. They too should experience a slow-down in connection speed, with the occasional connection loss due to time out. This gives rise to simulated "complaints" (via pop-ups) about the service being offered.

Finally, the router which detects that the SYN/ACK frames generated by the target cannot be forwarded will inform the target (via ICMP) that the destination is unreachable. In our simulation, this means that this router should be seen to produce a large number of these ICMP messages.

By arranging that these symptoms develop over time, we create the representation of a developing attack, and the trainee's task is then to establish, as soon as possible, the nature and impact of the attack, and to specify remedial action to recover from it.

From this starting point, a number of possible activities are possible: we can use the system as it stands, allowing trainee users to carry out a set of training activities (limited only by our creation of different scenarios); we can record the users' activities and review those records to determine situations where other actions might have been more appropriate, the ability to replay the same scenario means that this can then be easily done; we can use data mining techniques to determine whether patterns of use emerge which might give us insights into the way in which users react to these situations, offering the opportunity to develop new network management tools to assist by automating the beneficial reactions. Work on the latter is already underway, as described by Donelan *et al* (2004).

Table 1: Variation to MIB objects representing a simulated DOS attack

MIB-2 object	Change	Comment
tcpInSegs	Increase at "attack rate" +	Each attack is a tcp segment, others may still also be transmitting data
tcpOutSegs	Do.	Each SYN should generate a SYN/ACK response
tcpPassiveOpens	Increase at "attack rate"	Each attack is a connect request event
tcpCurrEstab	Do.	Number of "established" connections counter
tcpConnTab entries	increasing number of connections in state SYN_RCVD with non-recognised remote IP address	Connection table shows state of current connections
tcpAttemptFails	Increase at "attack rate" after onset of attack	After some time the "attack" connections will time out. (Stevens 1997)
icmpInDestUnreachs	Do.	The router to which the SYN/ACKs are sent will be unable to find an outgoing route for these frames, and will inform the sender.

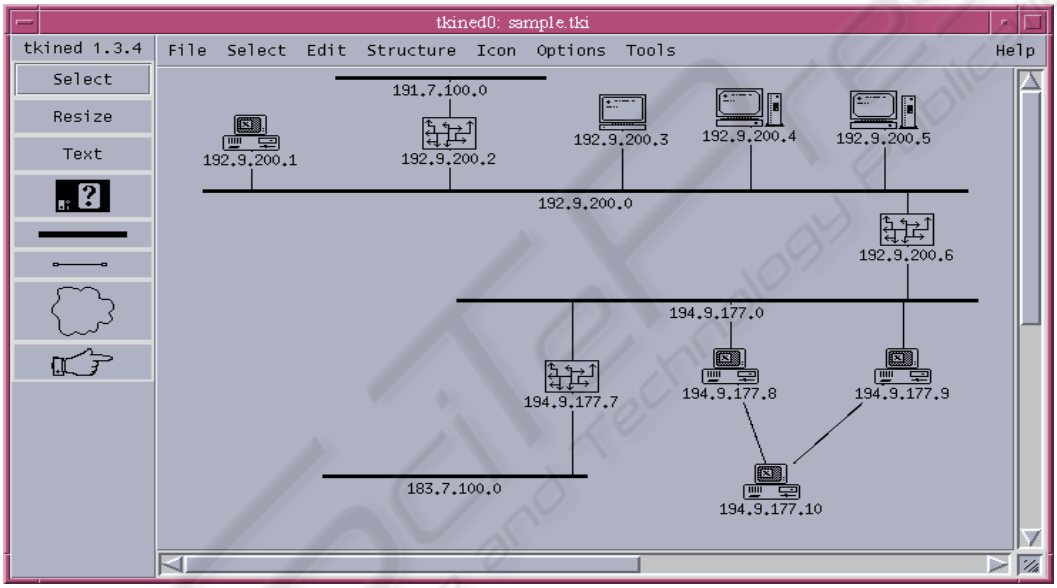

Figure 1: The simulated DOS attack will be mounted on node 192.9.200.1 (upper left), and will arrive from the internet through router 192.9.200.2.*

* Note we use the tkined network editor (tkined, 2000) to present this information, but with MIB data taken from a simulation rather than a "real" network – the work discussed here is our simulation technique.

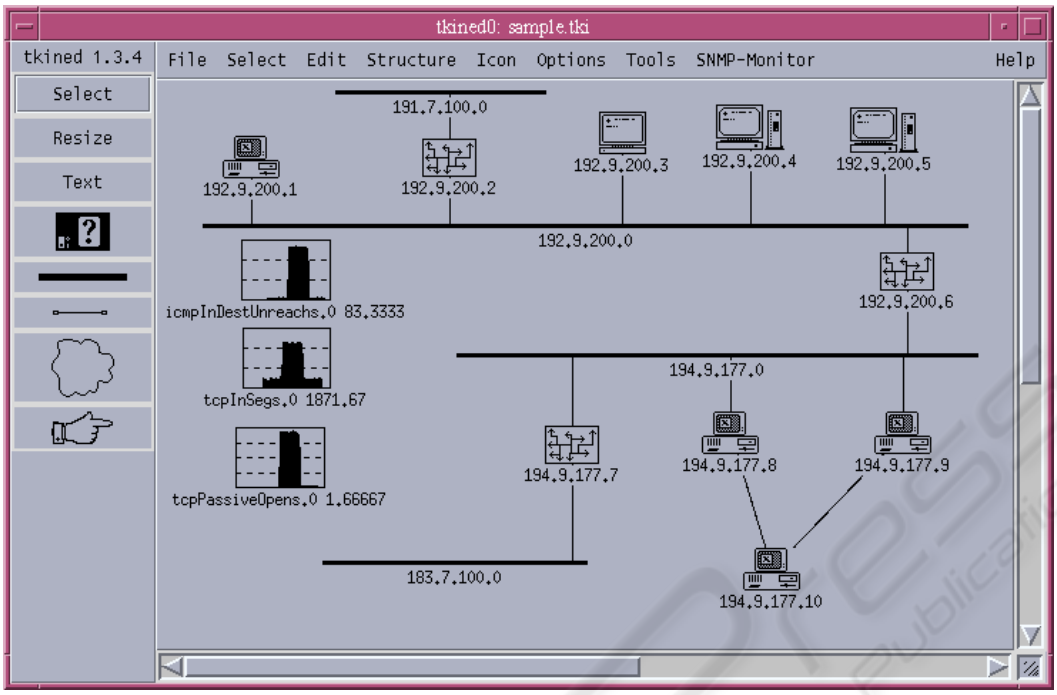

Figure 2: Graphs monitoring the delta values of different MIB counters on node 192.9.200.1 show the surge of a DOS attack against lower “normal” background traffic.

5 RESULTS AND CONCLUSIONS

We are aware that our model represents a form of attack which is now somewhat “dated”, and that newer forms of attack present different symptoms. However, we believe that the basic premise of our work does indicate the need to develop tools of this form in order to meet the needs of the large number of network administrators who may not have access to large, complex and automated network detection systems, but who are still expected to take responsibility for the integrity of their network provision.

Whilst our system does not address the problem of *prevention* of DOS attacks, we believe that we are able to present our trainees with a valid experience of the kinds of network behaviour they may experience in managing a network which is subjected to such an attack. We also believe that a system such as this is a valuable addition to the learning experience of novice network managers. We wish to continue development to present further simulated network attacks, and to expand and increase the size and variety of network types involved. One particular area of interest is in the incorporation of simulated mobile devices with the simulated network. We are also recording the users’ responses, we then plan to study these responses to gain greater insight into how network managers react to fault situations, which may in turn allow us to influence the development of tools to assist managers in their task.

REFERENCES

P. Dokas, L. Ertoz, V. Kumar, A. Lazarevic, J. Srivastava, Pang-Ning Tan, (2001) Data Mining for Network Intrusion Detection. *Decision Sciences Journal*, 32, Number 4 Fall 2001, Decision Sciences Institute http://www.decisionsciences.org/index.html

H. Donelan, C. Pattinson, D. Palmer-Brown (2004), The analysis of user behaviour of a network management training tool using a neural network. International Conference on Education and Information Systems: Technologies and Applications (EISTA 2004), Orlando, USA, July 21-25 2004.

C. Pattinson (2000) A simulated network management information base Journal of Network and ComputerApplications 23 April 2000 pp. 93 – 107

W.R. Stevens (1997) TCP/IP Illustrated Volume 1. Addison Wesley Longman, Reading, MA.

The Carnegie Mellon University CERT Coordination Center, (1999) 1999– UDP DoS. http://www.cert.org/incident_notes/IN-99-07.html [Accessed 12 October 2003]

Tkined (2000) is the network management interface provided by scotty
http://wwwhome.cs.utwente.nl/~schoenw/scotty [Accessed 12 February 2004]
UCLA, 2002 D-WARD Project, UCLA, Computer Science Department. http://lasr.cs.ucla.edu/dward/ [Accessed 12 October 2003]
University of Minnesota Dept. of Computer Science & Engineering, 2003 Minnesota Intrusion Detection System
http://www.cs.umn.edu/research/minds/MINDS.htm [Accessed 12 February 2004]