

PROACTIVE INTRUSION DETECTION AND SNMP-BASED SECURITY MANAGEMENT: NEW EXPERIMENTS AND VALIDATION

J.B.D. Cabrera¹, L. Lewis², X. Qin³, C. Gutiérrez¹, W. Lee³ and R.K. Mehra¹
*Scientific Systems Company*¹:*University of New Hampshire*²:*Georgia Institute of Technology*³

Abstract: In our earlier work we have proposed and developed a methodology for the early detection of Distributed Denial of Service (DDoS) attacks. In this paper, we examine the applicability of Proactive Intrusion Detection on a considerably more complex set-up, with hosts associated with three clusters, connected by routers. Background TCP, UDP and ICMP traffic following Interrupted Poisson Processes are superimposed on the attack traffic. We have examined six types of DDoS attacks. In four of the attacks we have obtained valid MIB-based precursors with no false alarms in all experiments. In the remaining two attacks precursors were obtained, but false alarms were observed. Procedures for eliminating these false alarms are discussed.

Keywords: Security Management; Data Warehousing and Statistical Methods in Management; Network and Systems Monitoring; Information Modeling

1. INTRODUCTION

In [2] we developed a methodology for utilizing traffic MIB variables for the early detection of Distributed Denial of Service (DDoS) attacks: Proactive Intrusion Detection, as opposed to the passive detection enabled by current IDSs.

In the present paper we evaluate the Proactive Intrusion methodology in a more realistic scenario. We performed the experiments on a network formed by three sub-networks connected by routers. The DDoS master, the DDoS slave and the target are placed in different clusters, allowing us to investigate the effect of routing in our ability to extract the precursors. We produced traffic for TCP, ICMP and UDP following Interrupted Poisson Processes and superimposed it on the attack traffic. Finally, we experimented with six types of DDoS attacks. We refer the reader to [2][3] for a detailed discussion of the Proactive Intrusion Detection Methodology.

The original version of this chapter was revised: The copyright line was incorrect. This has been corrected. The Erratum to this chapter is available at DOI: [10.1007/978-0-387-35674-7_66](https://doi.org/10.1007/978-0-387-35674-7_66)

2. THE NEW TEST BED AND EXPERIMENTS

Topology: R1 - R3 are routers; H1 through H6 are normal hosts.

Data collection: 64 MIB variables from the `ip`, `udp`, `tcp` and `icmp` groups were collected for 2 hours, at a sample rate of 3 seconds.

Attack Runs and Normal Runs: 12 attack runs corresponding to 6 types of attacks were obtained.

The attacks are: TFN2K SYN Flood, TFN2K Ping Flood, TFN2K Targa3 Flood, TFN2K Mix Flood, TFN2K UDP Flood and Trin00 [6]. Two runs are available for each attack type. 12 normal runs are also available, in which no attacks are present.

Placement of Malicious Agents: The master agent is placed in H2 during all attacks. Slave agents are placed in H1 and H4. The target is H6.

Background and Malicious Traffic: During normal runs, H1 to H6 generate TCP, ICMP and UDP traffic according to Interrupted Poisson Processes [5]. During attack runs, attack traffic was produced by malicious agents at H1, H2 and H4 and superimposed on the normal traffic. Live TCP connections among the various nodes of the network allowed us to observe the phenomenon of connection “time-out” caused by congestion.




Figure 1 The NCSU Test Bed

3. EXPERIMENTAL RESULTS

3.1 Step 1: Determining key variables at the target

Key variables were selected by comparing the time series of normal and attack runs. Since the attacks can be classified as flood attacks, MIB variables related to packet transmission are the best candidates for attack detection. `ipInReceives`, belongs to this category and is the most promising as it appears in all attacks.

3.2 Step 2: Determining key variables at the attackers

We search for possible causal relationships connecting MIB variables that could be considered key variables for each attack at the attacker hosts (H1 through H5) with the MIB variable selected in step 1 (`ipInReceives`) at the target host (H6). The value for the causality index between these variables was obtained using the

Granger Causality Test (GCT) [2]. The MIB variables found are then compared to the ground truth variables, selected from domain knowledge about the attacks [6].

Two types of precursors are defined. The first, called *T2 precursors*, correspond to the communication between master and slave. Second are *T3 precursors*, which correspond to the type of flood depending on the DoS attack.

3.3 Steps 3 and 4: Determining key events at the attacker

We break finding the key events at the potential attacker hosts that can be considered as precursors of an attack into two steps: step 3, *training*, where we determine the precursors at the potential attacker host and step 4, *testing*, where we test these precursors. In step 3, we define a precursor as an abnormal change on a precursor variable. To find this abnormal change, a Normal Profile for each MIB variable is constructed from the normal runs. Any variable whose change in each attack run was larger than the one in the Normal Profile was considered a precursor event for step 3 for that run.

Table 1. TFN2K Ping Flood Precursor Events - Run 1

Training Attack	Precursor Variable	Test Attack Data					Normal Data									
		Detections					False Alarms									
		H1	H2	H3	H4	H5	H1	H2	H3	H4	H5	H1	H2	H3	H4	H5
PING_1_1 through PING_1_6	ip.ipOutRequests	x														
	tcp.tcpRetransSegs							x								
	tcp.tcpInErrs	x			x											
	udp.udplnErrors	x			x											

In step 4, the precursor events of one attack run were tested on the other run. Table 1 shows the results for the first run of TFN2K Ping Flood. *Precursor variables* are the MIB variables that were selected following step 2 and step 3. Under *Test Attack Data*, *Detections* corresponds to the variables that belong to ground truth, and *False Alarms* are variables detected that do not belong to ground truth. Under *Normal Data*, the variables that are marked are those whose jump during a the normal test run was greater than that in the Normal Profile, in other words, a false positive.

Table 1 shows that there are T2 and T3 precursor events in H1 and H4, and tcpRetransSegs appears as a false alarm in H3. H1 and H4 can be detected as the attack hosts during this TFN2K Ping Flood.

4. SUMMARY AND CONCLUSIONS

The overall results demonstrated the applicability of Proactive Intrusion Detection using more realistic background traffic and sub-networking. In three of the attacks (TFN2K Syn Flood, TFN2K Mix Flood and TFN2K Targa3) both T2 and

T3 precursors were obtained, with no false alarms. For TFN2K UDP Flood, no false alarms were recorded, but T3 events were not obtained. For TFN2K Ping Flood and Trin00 UDP Flood false alarms were observed, related to the variable `tcpRetransSegs`. During an attack the network becomes congested and the open TCP connections start to time-out and TCP packets are retransmitted. `tcpRetransSegs` is detected as a precursor at H3, which is neither an attacker nor a target. Since the TCP time-out effect and retransmission occur after the attack starts a second pass may eliminate these false alarms. We are currently exploring an application of precursors which is more directly related to “conventional” (not proactive) Intrusion Detection. It has been noted [1][4] that the high rate of false alarms in current IDSs represents perhaps the main challenge for the deployment of IDSs. Precursors could be useful to corroborate the “hits” from a conventional IDS.

Finally, using more general traffic models instead of Interrupted Poisson Processes for background traffic would not have changed the results for this particular set of experiments and topology. We noted that key precursors MIBs are not related with traffic counting processes, but with counting anomalies in the protocol stack.

ACKNOWLEDGEMENTS

This work was supported by the Air Force Research Laboratory (Rome, NY) under contract F30602-01-C-057 to Scientific Systems Company and by Aprisma’s University Fellowship Program 1999/2000. SSCI acknowledges the support from the Defense Information Warfare Branch at the Air Force Research Laboratory.

REFERENCES

1. S. Axelsson. *The base-rate fallacy and the difficulty in intrusion detection*, ACM Transactions on Information and Systems Security, vol. 3, no. 3, 2000.
2. J.B.D. Cabrera, L. Lewis, X. Qin, W. Lee and R.K. Mehra. *Proactive Intrusion Detection and Distributed Denial of Service Attacks – A Case Study in Security Management*, Journal of Network and Systems Management, vol. 10, num. 2, pp. 225-254, June 2002.
3. J.B.D. Cabrera and R.K. Mehra. *Extracting Precursor Rules from Time Series – A Classical Statistical Viewpoint*, Proceedings of the Second SIAM International Conference on Data Mining, Arlington, VA, pages 213-228, April 2002.
4. S. Northcutt. *Intrusion Detection: An Analyst’s Handbook*, New Riders, 1999.
5. H.G. Perros. *An Introduction to ATM Networks*, John Wiley and Sons, 2001.
6. P.J. Criscuolo. *Distributed Denial of Service – Trin00, Tribe Flood Network, Tribe Flood Network 2000, and Stacheldraht*. Technical Report CIAC-2319, Department of Energy (CIAC – Computer Incident Advisory Capability), February 2000.