

An Analysis of Network Scanning Traffic as it relates to Scan-Detection in Network Intrusion Detection Systems

Richard J Barnett¹ and Barry Irwin²
Security and Networks Research Group
Department of Computer Science
Rhodes University
Grahamstown, South Africa
E-Mail: ¹barnettrj@acm.org ²b.irwin@ru.ac.za

Abstract—Network Intrusion Detection is, in a modern network, a useful tool to detect a wide variety of malicious traffic. The ever present prevalence of scanning activity on the Internet is fair justification to warrant scan detection as a component of network intrusion detection. Whilst current systems are able to perform scan-detection, the methods they use are often flawed and exhibit an inability to detect scans in an efficient and scalable manner.

Existing research by van Riel and Irwin has illustrated a number of flaws present in the open source systems Snort and Bro. This paper builds on this by describing current research at Rhodes University in which these flaws are being addressed. In particular, this research will address the flaws in the scan-detection engines in Snort and Bro by developing new plug-ins for these systems which take into consideration the improvements which are identified over the course of the research.

Index Terms—Network Security, Intrusion Detection, Port scanning, Snort, Bro.

I. INTRODUCTION

THE continued prevalence of scanning activity on the Internet is of interest to network administrators as scanning is often used by malicious persons as a prelude to attack. This is done to determine specific vulnerabilities in networked hosts. The detection of scanning activity is, in the simple case, a well defined and understood exercise and current generation Network Intrusion Detection Systems (NIDS) are capable of performing simple scan detection.

Existing research [1], [2] in the Department of Computer Science at Rhodes University has identified limitations in how the open source NIDS Snort [3] and Bro [4] handle scan detection. This paper describes current research which addresses these flaws.

The remainder of this paper is structured as follows: Section II presents some related work in this field. Section III will present the planned outcomes of this research. Following this, Section IV will present our research strategy and lastly,

The authors would like to acknowledge the financial support of Telkom SA, Business Connexion, Comverse SA, Stortech, Tellabs, Amatole, Mars Technologies, openVOICE and THRIIP through the Telkom Centre of Excellence in the Department of Computer Science at Rhodes University. They would also like to acknowledge the support of the National Research Foundation.

Section V will present those conclusions which can be drawn from our research at this stage.

II. RELATED WORK

A wide variety of literature is available on the detection of malicious traffic in a network and on scanning activities in general. We present a very small subset of this research as it relates to our work.

The identification of scanning on the network is core to this research and Allman *et al.* [5] describe methods for doing so through the classification of connections and hosts. Their work allows for the assignment of connections into three categories and hosts into two. Connections can be classified either as *good*, *bad* or in some cases as *unknown*. Hosts are either *scanners* or *not scanners*.

Good connections are those which are successful, whilst those connection attempts which do not lead to established connections can be classified as *bad*. In cases where it is unclear if the connection was successful, the connection is marked as *unknown*. Hosts which are labelled as *scanners* are those which make a majority of *bad* connections, hosts which are *not scanners* are those that make a majority of *good* connections.

van Riel and Irwin [6] present a visual tool for the graphical analysis of network traffic. This tool, known as InetVis, is designed to permit the rapid analysis of significant quantities of network traffic. In addition to InetVis, van Riel and Irwin present some preliminary findings [2] from their analysis of network telescope traffic. These findings are the basis for our own work.

III. RESEARCH OUTCOMES

The importance of performing scan detection in a network has been seen. The problems associated with the identification of scans in current NIDS does, however, partially invalidate the usefulness of these features. This is because false negatives - or, in this context, missed scans - may lead network administrators to believe that their network is more secure than it really is and give them a false sense of security.

This research aims to achieve a number of distinct outcomes. Firstly, it aims to validate the prior research performed by van

Riel and Irwin [2] which isolates a number of scans present in Network Telescope traffic but remain undetected by Snort and Bro. According to the Snort module *sfPortscan*'s documentation [7], it is designed to detect the different types of scan which can be generated by the common portscanning tool Nmap [8]. However when tested, we found that Snort does not detect all possible scans generated by Nmap.

Having validated the outcomes of that research, we intend on using those outcomes to produce modules for the open source NIDS Snort and Bro with the intention of improving their ability to perform scan detection. This process has a number of outcomes of its own, some of which will be achieved through the validation of van Riel and Irwin's research. Most notably, this process requires the identification and isolation of different scan types. These can then be used in conjunction with the existing Snort and Bro modules to produce algorithms which can be applied to traffic to determine if it constitutes a scan.

IV. RESEARCH STRATEGY

The process of developing algorithms for use in NIDS involves a number of approaches. Initially, several years of network telescope traffic is analysed with InetVis [2], [6] and scans are identified, classified and extracted. Our preliminary work in this regard has shown that there are a number of categories of scans and each category has a number of scan types.

Thus far, it would appear as though the categories of port sweeps and pseudo random phenomena are the most common, but port scans and hybrid scans are also seen.

A port sweep is defined as a scan which scans the same port across a number of hosts. Our work has shown that different agents perform sweeps in different ways, which may prove to aid in the evasion of NIDS. Pseudo random phenomena are scans which appear to be random network traffic, but over time show a defined pattern. This category of scan is more prevalent in the higher, less frequently used port ranges, whilst more traditional sweeps are usually seen in the reserved port range.

Port scans have not, at this stage, been seen extensively but are scans which target a large number of ports on a single host. Hybrid scans are scans which encompass different aspects of the previous categories.

On completion of this process, the authors intend on applying the extracted scans to Snort and Bro to determine if their current scan detection systems are capable of alerting on these scans. Snort and Bro will also be tested with scans crafted using current scanning tools such as Nmap. The process of applying both real-world and theoretical data to these NIDS will permit us to evaluate exactly how effective they are in dealing with scans.

The outcome of this process will permit us to take the existing scan detection systems and improve upon them by developing our own modules for both Snort and Bro which make use of more sophisticated techniques. Whilst we have yet to perform work in this area, it is envisaged that this will include a statistical analysis of packet data and the adaptation of current algorithms to determine if a scan is in progress or

not. Our analysis will also consider the scalability of such scans, especially with regard to memory utilisation during scans which occur in a non-trivial temporal reference frame.

V. CONCLUSIONS

At this stage, our research is in its infancy, however, this paper has highlighted some significant aspects of it.

Firstly, the preliminary research has shown that there are flaws in the scan detection algorithms in the NIDS Snort and Bro, and that they are unable to identify all scans produced by tools such as Nmap. Secondly, a number of categories of scans have been identified and the relative frequency of these scans has been presented.

These aspects will be utilised in our continuing research and will be used in the development of sophisticated detection systems for the open source NIDS Snort and Bro.

REFERENCES

- [1] J.-P. van Riel and B. Irwin, "Identifying and investigating intrusive scanning patterns by visualizing network telescope traffic in a 3-d scatter-plot," in *Proceedings of 6th Annual Information Security South Africa (ISSA) 2006*, H. Venter, J. Eloff, L. Labuschagne, and M. Eloff, Eds. Information Security South Africa, 2006.
- [2] B. Irwin and J.-P. van Riel, "Inetvis: a graphical aid for the detection and visualisation of network scans," in *Conference on Visualization Security (VizSec2007)*, 2007.
- [3] "Snort - the de facto standard for intrusion detection/prevention," Online: <http://www.snort.org/>, Accessed: 28/01/2008.
- [4] "Bro intrusion detection system - bro overview," Online: <http://www.bro-ids.org/>, Accessed: 28/01/2008.
- [5] M. Allman, V. Paxson, and J. Terrell, "A brief history of scanning," in *IMC '07: Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*. New York, NY, USA: ACM, 2007, pp. 77–82.
- [6] J.-P. van Riel and B. Irwin, "Inetvis, a visual tool for network telescope traffic analysis," in *Afrigraph '06: Proceedings of the 4th international conference on Computer graphics, virtual reality, visualisation and interaction in Africa*. New York, NY, USA: ACM, 2006, pp. 85–89.
- [7] D. Roelker, M. Norton, and J. Hewlett, *sfPortscan*.
- [8] "Nmap - free security scanner for network exploration & security audits." Online: <http://nmap.org/>, Accessed: 29/05/2008.

Richard Barnett received his BSc (Hons) degree in Computer Science from Rhodes University, Grahamstown. He is currently reading towards his MSc degree in the Security and Networks Research Group in the Department of Computer Science at Rhodes University. His research interests include Computer Security and Network Management.

Barry Irwin received his MSc degree in Computer Science from Rhodes University, and is a Senior Lecturer in the Department of Computer Science. His research interests include Computer and Network Security with a focus on Data mining of Network traffic through visual methods.