Available online at www.sciencedirect.com
SciVerse ScienceDirect
journal homepage: www.elsevier.com/locate/cose
ELSEVIER
Computers & Security

Performance of automated network vulnerability scanning at remediating security issues

Hannes Holm*

Industrial Information and Control Systems, Royal Institute of Technology, SE 10044 Stockholm, Sweden

ARTICLE INFO

Article history:
Received 24 May 2011
Received in revised form
22 November 2011
Accepted 28 December 2011

Keywords:
Network security
Security tools
Vulnerabilities
Vulnerability remediation
Vulnerability detection

ABSTRACT

This paper evaluates how large portion of an enterprises network security holes that would be remediated if one would follow the remediation guidelines provided by seven automated network vulnerability scanners. Remediation performance was assessed for both authenticated and unauthenticated scans. The overall findings suggest that a vulnerability scanner is a usable security assessment tool, given that credentials are available for the systems in the network. However, there are issues with the method: manual effort is needed to reach complete accuracy and the remediation guidelines are oftentimes very cumbersome to study. Results also show that a scanner more accurate in terms of remediating vulnerabilities generally also is better at detecting vulnerabilities, but is in turn also more prone to false alarms. This is independent of whether the scanner is provided system credentials or not.

1. Introduction

Technical vulnerabilities such as software vulnerabilities (e.g. CVE-2008-4250) and configuration errors (e.g. weak passwords) are key contributors towards the risk in enterprise networks as they can provide the attacker means to affect the confidentiality, integrity and availability of their assets. For example, stealing or manipulating information and causing denial of service incidents. Efficient management of vulnerabilities is therefore an important activity in modern enterprises security efforts. To manually keep track of all vulnerabilities present in systems and remediate them appropriately is a strenuous task unfit for today's complex IT environments. Fortunately, there are tools aimed to provide automated support for this process available (Arkin et al., 2005).

One commonly applied solution involves the use of network vulnerability scanners (Welberg, 2008). A network vulnerability scanner is an appliance or software which is used to scan the architecture of a network, report any detected vulnerabilities and suggestions for how to remediate them. An important distinction here is the one between vulnerability detection and remediation detection; a scanner can fail to detect a vulnerability but still manage to provide remediation guidelines for it. For example, the Microsoft Security Bulletin MS09-012 contains four vulnerabilities and as such it is enough to find one of these vulnerabilities in order to be able to enable remediation of all of them. In the same sense, a vulnerability can be found but not given proper remediation suggestions for. However, this is quite rare – if a vulnerability is correctly identified there are generally several useful references that can be consulted, often for both workarounds and software updates.

The normal procedure of auditing a network with a network vulnerability scanner generally involves three parts: network scanning, vulnerability scanning and vulnerability analysis (Manzuik et al., 2007).

Network scanning involves identifying which hosts that are alive in the computer network, which operating systems that

* Tel.: +46 8 790 68 26; fax: +46 8 790 68 39.
E-mail address: hannesh@ics.kth.se.

0167-4048/$ – see front matter © 2012 Elsevier Ltd. All rights reserved.
doi:10.1016/j.cose.2011.12.014

they use, and what services they run. During the *vulnerability scan* a database of vulnerability signatures is compared to the information obtained from a network scan to produce a list of vulnerabilities that are presumably present in the network. Most tools thereafter attempt to verify the presence of these vulnerabilities through carefully constructed queries which aim to verify the vulnerability's presence without disrupting the service. Some tools also provide the possibility of actually exploiting vulnerabilities to fully verify their presence.

Network and vulnerability scanning base their assessments on signatures of operating systems used, services running, and their corresponding vulnerabilities. These signatures do not always provide the correct result, which causes issues for the security management staff. Sometimes they result in failure to identify required remediation efforts (i.e. false negatives); sometimes they result in erroneously reporting existing patches or non-vulnerable configurations to not be present/vulnerable (i.e. false positives). If scans produce remediation guidelines containing false positives and false negatives it will impede efficient mitigation – false positives will result in efforts to manage nonexistent problems and false negatives may lead to unexpected security problems.

The third part concerns *vulnerability analysis*. As it can be resource- and time consuming to remediate vulnerabilities it is important to identify the largest problems and remediate those first. While network vulnerability scanners generally are fairly immature in this respect, there are several tools which utilise information from network vulnerability scanners in order to provide sound decision support. A few such academic projects include (Homer and Ou, 2009; Ingols et al., 2009; Jajodia et al., 2005; Buschle et al., 2011) and two commercially available tools include Skybox (Security, 2011) and SecurITree (Aminaza, 2011). The currently most common scenario in practice is however that analysis is carried out using solely the remediation guidelines of a scanner.

The accuracy of network vulnerability assessment tools is of great importance, no matter if the vulnerability scanner report is interpreted by the decision maker or if it is the input to an analysis tool. While network vulnerability scanners have been around for more than a decade, there have been few thorough studies of their detection rate (Holm et al., 2011) and no thorough studies of their accuracy at providing remediation guidelines for vulnerabilities – a critical issue in need of research. This paper expands on the results regarding vulnerability detection rates described in (Holm et al., 2011) through studying remediation detection rate – how many errors that actually are eliminated if one would follow all the guidelines in a remediation report – by analysing remediation reports of seven popularly used network vulnerability scanners.

The primary focus of the study lies in *accuracy*; how many vulnerabilities that can be remediated given the guidelines of a scanning report (RQ1). This is a very important issue as remediation reports by network vulnerability scanners is one of the main means of eliminating network security vulnerabilities of today.

RQ1: How accurate are vulnerability scanners in terms of detecting remediating suggestions for vulnerabilities?

As seven network vulnerability scanners are examined, it is also possible to study the relations between other variables of importance for decision makers and vulnerability scanner developers. One such topic is that of the relation between false negatives (i.e. remediation detection rate) and false positives (i.e. false alarms) (RQ2). A strong correlation would suggest that there is a significant trade-off to be considered, and that a decision maker should choose a tool (or several tools) that complies with his or her preference in this area. Some information-critical environments might value security in front of service uptime (due to maintenance such as patching) and other environments likely favour low rates of false positives to remove any unnecessary service downtimes.

RQ2: What is the relationship between remediation detection rate and false alarms?

Another interesting topic is that of remediation and detection (RQ3). If there is a large difference between vulnerability detection rate and remediation detection rate it would suggest that using vulnerabilities found through scanners as input in analysis tools such as MulVAL (Ou et al., 2005) actually could be less useful then manually reading the remediation reports. Also, tests of vulnerability scanners (such as the present) are very resource-intensive to carry out and as such it would be beneficial if it is enough to look at one of these aspects.

RQ3: What is the relationship between remediation detection rate and vulnerability detection rate?

Finally, it is possible to provide network vulnerability scanners with authentication credentials of the studied systems and thus enable more in-depth analysis of them. However, credentials are not always readily available. Furthermore, following the argumentation regarding cost of tests – it would certainly be useful if it is enough to consider only one of these types of scans when evaluating the effectiveness of a scanner.

RQ4: What is the relationship between authenticated and unauthenticated scans in terms of remediation detection, vulnerability detection and false alarms?

The remainder of the paper is structured as follows: Section 2 discuss related works. Section 3 provides descriptions of the studied variables. Section 4 show the methodology of the study and Section 5 presents the results. Section 6 discusses the assessed results and Section 7 concludes the paper.

2. Related works

There has to the author's knowledge not been published a single study regarding vulnerability remediation rate of automated network vulnerability scanners. Furthermore, no previous academic work has been found that quantitatively evaluates other accuracy aspects of vulnerability scanners (i.e. detection rate and false positives). While there are a few