An Analysis of Rogue AV Campaigns*

Marco Cova[1], Corrado Leita[2], Olivier Thonnard[3], Angelos D. Keromytis[4], and Marc Dacier[2]

[1] University of California Santa Barbara, Santa Barbara, USA
marco@cs.ucsb.edu
[2] Symantec Research Labs, Sophia Antipolis, France
{corrado_leita,marc_dacier}@symantec.com
[3] Royal Military Academy, Brussels, Belgium
olivier.thonnard@rma.ac.be
[4] Columbia University, New York, USA
angelos@cs.columbia.edu

Abstract. Rogue antivirus software has recently received extensive attention, justified by the diffusion and efficacy of its propagation. We present a longitudinal analysis of the rogue antivirus threat ecosystem, focusing on the structure and dynamics of this threat and its economics. To that end, we compiled and mined a large dataset of characteristics of rogue antivirus domains and of the servers that host them.

The contributions of this paper are threefold. Firstly, we offer the first, to our knowledge, broad analysis of the infrastructure underpinning the distribution of rogue security software by tracking 6,500 malicious domains. Secondly, we show how to apply attack attribution methodologies to correlate campaigns likely to be associated to the same individuals or groups. By using these techniques, we identify 127 rogue security software campaigns comprising 4,549 domains. Finally, we contextualize our findings by comparing them to a different threat ecosystem, that of browser exploits. We underline the profound difference in the structure of the two threats, and we investigate the root causes of this difference by analyzing the economic balance of the rogue antivirus ecosystem. We track 372,096 victims over a period of 2 months and we take advantage of this information to retrieve monetization insights. While applied to a specific threat type, the methodology and the lessons learned from this work are of general applicability to develop a better understanding of the threat economies.

* This work has been partially supported by the European Commission through project FP7-ICT-216026-WOMBAT funded by the 7th framework program. The opinions expressed in this paper are those of the authors and do not necessarily reflect the views of the European Commission. This work was also partly supported by ONR through Grant N00014-07-1-0907 and the NSF through Grant CNS-09-14845. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the ONR or the NSF. The work of Marco Cova was supported by a fellowship made possible by Symantec Research Labs.

S. Jha, R. Sommer, and C. Kreibich (Eds.): RAID 2010, LNCS 6307, pp. 442–463, 2010.

1 Introduction

A rogue security software program is a type of misleading application that pretends to be legitimate security software, such as an anti-virus scanner, but which actually provides the user with little or no protection. In some cases, rogue security software (heretofore referred to as "rogue AV") actually facilitates the installation of the very malicious code that it purports to protect against.

Rogue AVs typically find their way into victim machines in two ways. First, social engineering techniques can be used to convince inexperienced users that a rogue tool is legitimate and that its use is necessary to remediate often nonexistent or exaggerated threats found on the victim's computer. A second, stealthier technique consists of attracting victims to malicious web sites that exploit vulnerabilities in the client software (typically, the browser or one of its plugins) to download and install the rogue programs without any user intervention (*e.g.,* through *drive-by* downloads). After a rogue AV is installed on a victim's machine, it uses a number of techniques to convince (or force) a user to pay for additional tools or services, such as a "full version" of the program or the subscription to an update service. The cost of these additional programs or services ranges from $30–$100 [8].

In the last few years, rogue AVs have become a major security threat, both in terms of their pervasiveness and their financial impact. For example, over a 1-year period, Symantec's sensors detected 43 million installation attempts, covering over 250 distinct families of rogue AV software [8]. In addition, an investigation by Krebs revealed that affiliate programs alone can generate upward of $300,000 a month for the individuals that distribute rogue AVs [14].

As a consequence, different companies in the computer security industry have recently focused their attention on this threat [1,4,8]. Most of the existing works have considered individual facets of the rogue AV problem, for example, the malware code (*e.g.,* the installation techniques it employs), the sites involved in its distribution (*e.g.,* their number and geolocation), and the victims that it affects. However, little has been done to understand the rogue AV phenomenon as a whole, that is, relating how these individual pieces become combined in rogue AV campaigns.

We seek to fill this gap by providing a better understanding of the organization and dynamics of rogue AV campaigns. In particular, we focus on characterizing the infrastructure used in a campaign (*e.g.,* its web servers, DNS servers, and web sites) and the strategies used to create and manage it. We also investigate the uniqueness of our findings to this very specific threat type, and we investigate the motivations underneath these differences by exploring its economics.

The key of our approach is a method that, given a list of individual AV-hosting sites, allows us to group them into campaigns, each characterized by coherent features. More precisely, we use an extensive dataset including network and domain registration information, as well as network-observable temporal characteristics of a large number of domains that are associated with rogue AV advertising and distribution. To that dataset we apply a multi-criteria fusion algorithm to group together nodes based on a certain number of common elements likely due to the