DEPARTMENT OF ELECTRICAL AND INFORMATION TECHNOLOGY | LUND UNIVERSITY | PRESENTATION 2016-06-01

MASTER'S THESIS Bypassing modern sandbox technologies
An experiment on sandbox evasion techniques
STUDENT Gustav Lundsgård and Victor Nedström
ADVISOR Paul Stankovski (LTH)
EXAMINER Thomas Johansson(LTH)

Bypassing modern sandbox technologies

POPULAR SCIENCE PAPER **Gustav Lundsgård and Victor Nedström**

Cybercrime is an increasing problem with both organisations and governments developing malicious software. Sandboxes are a hyped solution to this, but how good are they really? This experiment investigates techniques to bypass them.

Malicious software ("malware") is a continuously growing problem which affects both companies and individuals. Malware grows both in numbers and complexity, and in 2014 alone, over 320 million new malware were seen. The people developing malware can make a fortune on exploiting valuable computer systems or sabotaging them. What can be done to protect us from malware?

Traditional anti-malware products are well prepared to catch known malware which has been observed before. However, they fall short when facing new, previously unseen malware. As a solution to this, something called *sandboxes* were introduced. A sandbox is, just like the ones kids play in, an isolated "playground" where you can test things without having it affect and potentially break something outside of it. In the context of malware, a sandbox is a defense system which runs suspicious files before they reach the computers of regular users. When a file is run in a sandbox, an internal program observes what the file does and tries to determine whether or not it is malicious. Malicious files are quarantined, while benign ones are delivered to the end users.

Recently, malware has been observed which contain *sandbox detection* techniques. Simply put, some malware try to detect if they are running inside a sandbox or on the computer of an unsuspecting user. If a sandbox is detected, the malware hides it malicious behavior, otherwise it executes as intended and tries to infect the user. It is therefore very important that sandboxes resemble regular computers as much as possible, so that malware cannot distinguish between the two. In this Master's Thesis, such techniques for sandbox detection were developed to investigate different approaches and how good top-of-the-line sandboxes are at mitigating them.

The experiment showed that an average of roughly 43% of the sandbox detection techniques successfully managed to detect sandboxes. Clear patterns of weaknesses were found among the sandboxes, and some techniques proved successful on all sandboxes tested. Besides, some of the highly effective techniques were also trivial to implement.

On the bright side, the authors believe that some of the successful sandbox detection techniques can be mitigated quite easily by configuring the sandboxes a bit more thoroughly. On the other hand, there are also techniques which are much more difficult to prevent. The results of the experiment have been shared with the sandbox vendors, who have been happy to cooperate with the authors to improve their products even more.