Detecting targeted cyber attacks in the cloud

Alexander Adamov - May 15, 2015 - Advanced Persistent Threats | APT | Cyber security

This is an article I wrote at Mirantis with Alexander Adamov as the SME to market his presentation at OpenStack Summit. If an SME is involved in the writing process as a consultant, Mirantis gives him or her the byline and doesn't credit the writer/editor, which I was in this article.

Advanced Persistent Threats (APTs) are the most modern method for committing targeted cyber attacks, running silently for long periods of time as they collect specific information on a victim's computer or network. These victims are often government and military institutions in various countries. In a recent high-profile attack, both the White House and the U.S. State Department were compromised by the CozyDuke (aka CozyBear, CozyCar, or Office Monkeys) APT.

Intruders use a variety of methods to unleash APTs in all kinds of organizations. I'll be presenting detailed information on intrusion methods and detecting targeted cyberattacks in the cloud at the OpenStack Summit in Vancouver, highlighting:

1. APTs – cyberattack mechanisms
2. Antivirus solutions and limitations
3. Practical steps to enforce cloud security with intrusion detection systems (IDS)

For now, let's look at how APTs work.

APTs – the modern cyberattack mechanisms

These days attackers do not try to penetrate a security perimeter in a straightforward manner by scanning and exploiting vulnerabilities, as these methods may attract too much attention to the attack which will be blocked in a matter of minutes. Intruders prefer more sophisticated techniques based on social engineering that allow a spy program to operate undetected for an unlimited amount of time. By not attracting attention, a program can harvest sensitive information and send it to a command and control (C&C) server.

The most popular techniques used in targeted attacks are spear phishing emails, watering hole attacks, and 0-day exploits. When finally discovered, enterprise owners and stakeholders often keep these cases secret to protect their organizations' reputations.

When the APT perpetrator spear phishes its targets, he or she sends emails containing a link to a hacked website or attaches malware, mostly exploits. Sometimes the website is a high profile, legitimate site such as "diplomacy.pl", which hosts a ZIP archive with malware, as in the case of the CozyDuke attack.

In a watering hole attack, the attacker guesses or sees the websites used by an organization's members and infects the site with a malicious inclusion, which triggers an exploit. Eventually some members become infected, making a watering hole an efficient strategy for organizations that are resistant to spear phishing.

A 0-day exploit is malicious code that takes advantage of a yet unknown vulnerability and as a result, returns a high rate of proliferation within a network. An example is a publicly disclosed vulnerability in Microsoft's print spooler service used in the Stuxnet attack, that although fixed later, briefly allowed remote code execution if an attacker sent a specially crafted print request to a vulnerable system that had a print spooler interface exposed over RPC. In this 0-day exploit, the attacker used the print spooler service to spread itself in a local network. Overall, Stuxnet exploited five vulnerabilities, four of which were 0-days.

When attackers run targeted attacks they know pretty well who the victims are and how a targeted environment looks. For example, some of the enterprise APTs are designed to infect SCADA systems such as WinCC. In one such case, the Stuxnet worm was unleashed at Iran’s nuclear facilities in 2010.

The attackers may target an unlikely person’s system, for example first infecting a low-level employee in the organization. Having installed a backdoor in the company’s security perimeter, the attacker can now use penetration testing tools, such as 0-day exploits, to scan a corporate network for other victims.

In other highly successful attacks, targets receive a phony flash video as an email attachment. A clever example is “Office Monkeys LOL Video.zip”. The executable in the zip not only plays a flash video (see video screenshot in Figure 1), but drops and runs CozyDuke APT executables.

Fig. 1 – CozyDuke (APT) flash video screenshot

Victims quickly pass the video around the office with delight while systems are infected in the background silently. Many of the CozyDuke APT’s components have phony Intel and AMD digital certificates and send checkin requests to the C&C servers, as seen in Figure 2.

Fig. 2 – Check-in requests to a C&C server

An interesting fact about the CozyDuke APT is that it is believed to belong to the family of Duke APTs, supposedly created by a group of Russian authors. This assumption is based on artifacts Kaspersky Lab researchers found in samples, including MiniDuke, which attacked NATO and European government agencies, and OnionDuke, which uses the TOR network to download the Trojan-Dropper, which in turn installs a backdoor on a victim's computer.

Antivirus solutions and limitations

So how you can detect if your private cloud has been compromised? While a host or hypervisor antivirus can be a good solution, they have several limitations:

1. Antiviruses, even those equipped with a heuristic engine in the majority of cases, do not detect unknown 0-day malware. To avoid being detected, modern cyber espionage platforms like EquationDrug which is used by Regin malware, and Epic Turla use the Windows kernel-mode rootkit driver to hide files, registry keys, and processes by hooking some of the Native API functions.
2. An antivirus solution can be disarmed by evasive APTs once detected on a victim's computer.
3. A private or public cloud in many cases is a heterogenous (hybrid) environment built on different operating systems and hypervisors that increase deployment and operational costs, also making protection inflexible and vendor dependent.

So how can you protect a cloud against cyber attacks without a big investment in an enterprise antivirus complex deployed through a private cloud?

Practical steps to enforce cloud security with intrusion detection systems (IDS)

As you’ve seen, all APTs and backdoors communicate with C&C servers to get commands from attackers and send back collected information. The network traffic is usually encrypted and sent via an http port, preventing it from being detected by Network Data Leakage Prevention (DLP) or antivirus systems. However, every APT generates specific traffic that contains unique network indicators of compromise (IoCs). Network behavior analysis of recent APTs reveals patterns of targeted attacks that can detect previously unknown cyber espionage campaigns.

For example, let us take a look at the communication protocol of the CozyDuke (Office Monkeys) APT. It sends http requests to the www.sanjosemaristas.com C&C server, which replies with details for this victim’s set of payload modules. The transmitted data is Base64 encoded. Once we take Base64 off, we see some binary content that still seems to be encrypted.

If we open the backdoor in a debugger we can figure out that the encryption algorithm used to encrypt the data is RC4. See Figure 3.

Fig. 3

Moreover, we can even find the RC4 key in memory when the backdoor exports it in a BLOB format to be stored as a header before the encrypted data. The key is 16 bytes and is generated every time a new session is created, as shown in Figure 4.

Fig. 4

Before being encoded with Base64, the network packet contains the key at the beginning, so the server can use it to decrypt the message as seen in Figure 5.

Fig. 5

The same key is used to encrypt the backdoor's local configuration file called racss.dat. The XML config file shows the C&C server addresses, as seen in Figure 6. According to the information in the <Servers> section, we see two URLs used to establish the connection with the C&C server.

These seem to be new ones as they are not mentioned in the Kaspersky Lab report.

Fig. 6

If we open the second server main page we see that this is a legal sport website "CIF Southern Section" as shown in Figure 7.

Fig. 7

As you can see, after the short analysis of the threat we found two new network IoCs that can be blacklisted in your cloud network to detect the malicious activity of "Office Monkeys":

www.sanjosemaristas.com/app/index.php

www.cifss.org/product_thumb/index.php

Improving cloud security

While cyber attacks are evolving, we're also developing better diagnostics and tools to combat the threats. We'll build on the information presented here for an in-depth discussion in improving your cloud security on Thursday, May 21 4:10 pm at the Openstack Summit in Vancouver. See you soon!