Acronis

WHITE PAPER

Using AI to detect malicious documents and scripts

Achieving strong multilayered protection with the help of machine intelligence

Malware attacks continue to dominate the headlines, and cybercriminals are using social engineering to craft malicious emails and exploit modern channels of communication. As a primary attack vector, nefarious emails often contain malevolent attachments which, when opened, trigger the execution of a malware and exploit a vulnerability to perform adversarial objectives. A malicious document also can be downloaded by the link a user receives over communication channels.

Acronis experts observed a 15% increase in the number of files and URLs per scanned email in 2023, and cybercriminals continue to weaponize malicious documents such as Microsoft Office and Adobe PDF files in attacks. Other industry analysts have seen a similar picture. Malware is typically embedded in a file, although the chain of its execution can vary. For example, a vulnerability is initially exploited, privileges are raised, and the malware is then downloaded and executed. Office and PDF files have been used for years in cyberattacks as they can contain embedded macros, shellcodes, JavaScript and even complete files within them. However, the problem has worsened in recent years, as hackers now use search engine optimization (SEO) to rank malicious files — especially PDFs — higher in search engine results.

Another, similar problem involves malicious scripts executed by legitimate tools like PowerShell. The number of attacks using malicious scripts, as with malicious documents, has steadily grown from year to year — Acronis experts have seen almost double the growth in such attacks over the past two years. The goal with malicious scripts is the same: infiltrate the system, raise privileges, download and then execute the payload.

For cybersecurity companies, both issues detailed above are identical, because they require you to distinguish between good and bad documents or script. The good news is that the problem can be solved much more effectively with the usage of machine learning (ML) and artificial intelligence (AI), or, as we say at Acronis, with machine intelligence (MI).

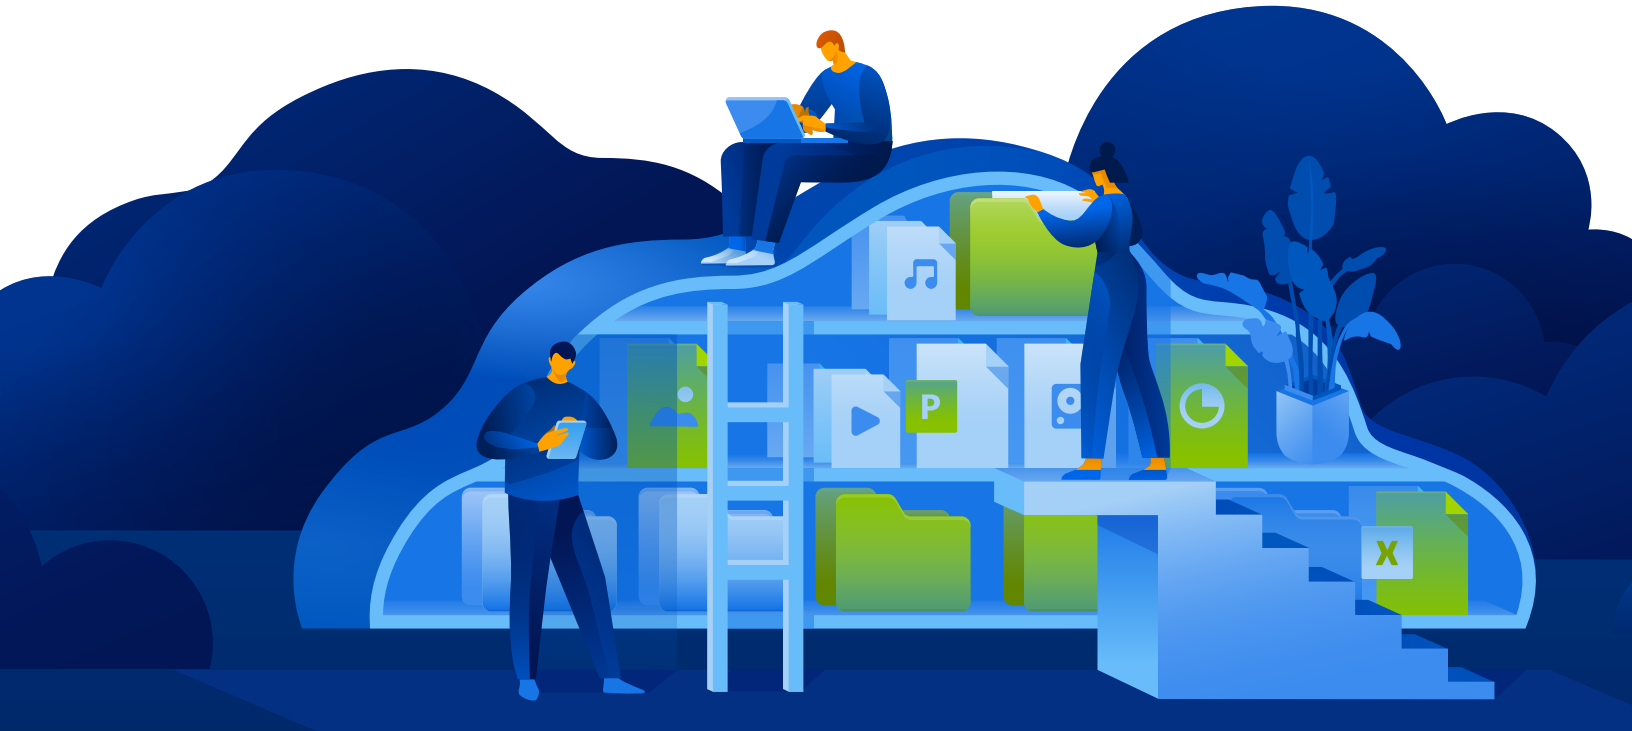

Using machine intelligence for threat detection

Pioneering machine learning, Acronis leveraged the technology as early as 2017 when we introduced Active Protection anti-ransomware. Soon after, we created a comprehensive AI-based static detection engine that is currently used in the Acronis flagship product, Acronis Cyber Protect.

This engine is constantly updated, and with every new machine learning model introduced, it improves in terms of performance and detection rate. For example, in the beginning, it was used to analyze stack traces of executed processes. Later, it started to analyze complete files and libraries to catch malicious ones. Now, it can analyze strings extracted from executables and process images. The frequency of a few hundred selected words in strings extracted as additional features has led to a more than 3% improvement in detection rate.

A similar approach can be used to detect malicious documents and scripts.

Let's start with Microsoft Office documents. Although Microsoft began disabling macros by default in internet-borne documents in 2023, this, unfortunately, can't solve all the problems. The mark of the web (MOTW) attribute, a feature that identifies files that originate from the internet, is added by Windows to files from an untrusted location, such as the internet or a restricted zone — for example, browser downloads or email attachments. The attribute applies only to files saved on an NTFS file system, not files saved to FAT32 formatted devices. What if the file comes from a legitimate office email? This will not necessarily work in an environment where a job is heavily reliant on macros — accounting, for example. Nevertheless, such a policy-enabled environment is made more secure, and some typical cases of phishing attacks, like those used in Emotet campaigns, can be successfully prevented.

But as we've seen, this policy will raise other complications for many businesses. That's why Acronis AI specialists have enhanced our detection engine to be able to spot malicious documents, thanks to the power of machine learning.

Scripts that perform one or more of the below may be perceived as malicious:

- **Creates processes.**
- **Executes scripts in PowerShell, VBA, etc.**
- **Downloads files from remote servers.**
- **Embeds itself in other Office files or Office template files.**

However, real detection is not that simple, as you need to factor in a lot of other parameters. For instance, the Acronis machine learning model analyzes the following attributes of DOCX files to arrive at a verdict:

- **Various text and VBA function features.**
- **Ratio features like comments, code, etc.**
- **The entropy of a macro itself, its code and comments.**
- **Any obfuscation in place and what is obfuscated.**
- **Known indicators of compromise (IoC) parameters like URLs, executables, etc.**

This is not a full list, but it helps demonstrate that a number of attributes are analyzed on a large dataset that is constantly being revised and updated. As a result, we are achieving an excellent detection rate with a model size of less than 1 MB compressed. Achieving such results without AI and ML is close to impossible. Keep in mind that this is just one part of robust multilayered protection that will be triggered only if the threat is not detected by other technologies beforehand, such as email security scanning engines, sandboxes or URL filtering.

A similar approach was recently initiated by Acronis experts to detect malicious AutoIt scripts, which are often used in a service provider environment. With a tiny model of around 0.6 MB, we are already able to provide a 92% detection rate, and the same with DOCX, where the model constantly improved.

Eliminating the weaponized PDF threat

Apart from Microsoft Office Word documents, the Adobe portable document format, or PDF, is a popular tool cybercriminals use to compromise a system or plant malware on a user's machine. Based on PostScript language, a PDF can contain a lot of information, including text, hyperlinks, multimedia, images, attachments, metadata, etc. PDF format has an "actions" feature that allows the opening of a web link or file, running JavaScript code, and many other operations that can be performed with malicious intent.

PDF documents can be viewed with browsers and a variety of reading software — all of which can and sometimes do have vulnerabilities threat actors can exploit. These include arbitrary code execution, buffer overflow, memory corruption, out-of-bounds read and many others. Currently, there are hundreds of CVEs for PDF readers — more than 900 known vulnerabilities for Adobe Acrobat Reader alone in November 2023. Security researchers and threat actors are finding new PDF-related exploits practically every day.

Here is an example of a CVE-2023-44372 vulnerability seen exploited in the wild: Acrobat Reader DC versions 23.0006.20360 (and earlier) and 20.005.30524 (and earlier) are affected by a Use-After-Free vulnerability. An unauthenticated attacker could leverage this vulnerability to achieve arbitrary code execution in the context of the current user. The exploitation of this issue requires user interaction — a victim must open a malicious file.

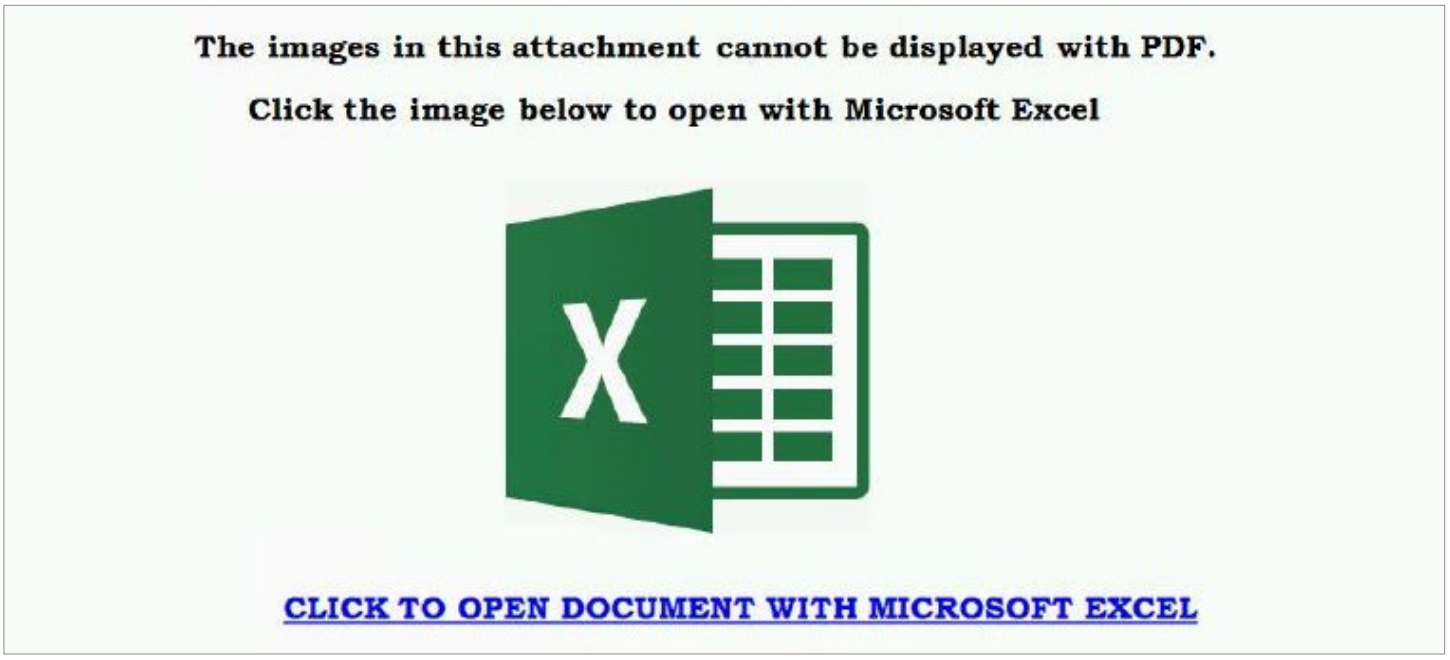

Example of phishing with a malicious link planted in a PDF file.

As with other file types described above, Acronis' machine learning-based malicious PDF detection model checks a variety of parameters to arrive at the correct verdict:

- **Entropy.**
- **Total character count.**
- **Special keyword counts.**
- **Number of lines, special assignment lines.**

As a result, an identical and high detection rate, as described above, is achieved.

Human cybersecurity awareness training can help significantly

Integrating this type of detection model into a cyber protection solution is essential. People, however, are also critical to defense. In most cases, engaging with a malicious document or script requires user input. Without it, there is no threat. Overall company and individual security postures can benefit greatly if users are aware of threats and the way they present themselves, and then react properly.

Awareness training should teach users to adhere to these well-known security community rules:

- Always check where an email or link originated from. Do you know this person? Were you expecting this file? Inspect the sender's email address — not just the sender's name or email alias — to reveal the true sender.
- If you have clicked on a file and you receive an incompatibility alert, are asked to enable something, are presented with links, a captcha, or something similar, you should either close the file or delete it, and then inform the security team immediately.
- If you have clicked on a file but it does not open, this could be a sign that that file is malicious. If file opens, but it is not what you expected, you may already have fallen victim. Inform the security team immediately.

The best rule is to stop and think before opening links and attachments. Double-check everything. If unsure, talk to the security team. Stay vigilant and stay safe.