

Independent Tests of Anti-Virus Software



Advanced Threat Protection - Consumer Enhanced Real-World Test - Targeted Attacks

TEST PERIOD: SEPTEMBER-OCTOBER 2022

LAST REVISION: 10TH NOVEMBER 2022

WWW.AV-COMPARATIVES.ORG

Content

INTRODUCTION	3
TEST PROCEDURE	5
TESTED PRODUCTS	6
TEST RESULTS	7
AWARD LEVELS REACHED	9
TEST CASES EMPLOYED	10
ABOUT THIS TEST	11
COPYRIGHT AND DISCLAIMER	13

Introduction

"Advanced persistent threat" is a term commonly used to describe a targeted cyber-attack that employs a complex set of methods and techniques to penetrate information system(s). Different aims of such attacks could be stealing/substituting/damaging confidential information, or establishing sabotage capabilities, the last of which could lead to financial and reputational damage of the targeted organisations. Such attacks are very purposeful, and usually involve highly specialised tools. The tools employed include heavily obfuscated malicious code, the malicious use of benign system tools, and non-file-based malicious code.

In our Advanced Threat Protection Test (Enhanced Real-World Test), we use hacking and penetration techniques that allow attackers to access internal computer systems. These attacks can be broken down into Lockheed Martin's Cybersecurity Kill Chain, and seven distinct phases - each with unique IOCs (Indicators of Compromise) for the victims. All our tests use a subset of the TTP (Tactics, Techniques, Procedures) listed in the MITRE ATT&CK[®] framework¹. A false alarm test is also included in the report. The tests use a range of techniques and resources, mimicking malware used in the real world. Some examples of these are given here. We make use of system programs, in an attempt to bypass signature-based detection. Popular scripting languages (JavaScript, batch files, PowerShell, Visual Basic scripts, etc.) are used. The tests involve both staged and non-staged malware samples, and deploy obfuscation and/or encryption of malicious code before execution (Base64, AES). Different C2 channels are used to connect to the attacker (HTTP, HTTPS, TCP). Use is made of known exploit frameworks (Metasploit Framework, PowerShell Empire, commercial frameworks, etc.).

To represent the targeted PCs, we use fully patched 64-bit Windows 10 systems, each with a different AV product installed. In the enterprise test, the target user has a standard user account. In the consumer test, an admin account is targeted, although every POC is executed using only a standard-user account, with medium integrity. Windows User Account Control is enabled and set to the default level in both tests. With regard to vendors whose products were tested in both the Consumer and Enterprise ATP Tests, please note that the products and their settings may differ. Hence, the results of the Consumer Test should not be compared with those of the Enterprise Test.

Once the payload is executed by the victim, a Command and Control Channel (C2) to the attacker's system is opened. For this to happen, a listener has to be running on the attacker's side. For example, this could be a Metasploit Listener on a Kali Linux system. Using the C2 channel, the attacker has full access to the compromised system. The functionality and stability of this established access is verified in each test-case. If a stable C2 connection is made, the system is considered to be compromised.

The test consists of 15 different attacks. It focuses on protection, not on detection, and is carried out entirely manually. Whilst the testing procedure is necessarily complex, we have used a fairly simple description of it in this report.

AV Consumer Main-Test-Series vendors were given the opportunity to opt-out of this test before the test started, which is why not all vendors are included in this test. Some vendors are continuing to perfect their products before joining this advanced test. We congratulate all those vendors who took part in the test, even those whose products did not get the best scores, as they are striving to make their software better.

¹ <https://attack.mitre.org/matrices/enterprise/windows/>

Scope of the test

The Advanced Threat Protection (ATP) Test looks at how well the tested products protect against very specific targeted attack methods. It does not consider the overall security provided by each program, or how well it protects the system against malware downloaded from the Internet or introduced via USB devices and shared network drives.

It should be considered as an addition to the Real-World Protection Test and Malware Protection Test, not a replacement for either of these. Consequently, readers should also consider the results of other tests in our Main-Test Series when evaluating the overall protection provided by any individual product. This test focuses on whether the security products protect against specific attack/exploitation techniques used in advanced persistent threats. Readers who are concerned about such attacks should consider the products participating in this test, whose vendors were confident of their ability to protect against these threats in the test.

In the ATP test, we focus on testing different kinds of POC C2 malware, based on different adversary tactics and techniques. We use a variety of delivery scenarios to include the possible adversary strategies. The goal of the ATP Test is to demonstrate the prevention capabilities of the respective products. To accomplish this, we use different POCs, all of which try to open a stable C2 channel after execution, thus simulating a successful initial compromise. In cases where a POC was not prevented and the attacker was able to open a stable C2 session, the target PC was considered to be compromised. The test does not check across different stages of an attack (which is done in our EPR test).

Test procedure

Scripts such as VBS, JS or MS Office macros can execute and install a file-less backdoor on victims' systems and create a control channel (C2) to the attacker, who is usually in a different physical location, and maybe even in a different country. Apart from these well-known scenarios, it is possible to deliver malware using exploits, remote calls (PSexec, wmic), task scheduler, registry entries, Arduino hardware (USB RubberDucky) and WMI calls. This can be done with built-in Windows tools like PowerShell. These methods load the actual malware directly from the Internet into the target system's memory, and continue to expand further into the local area network with native OS tools. They may even become persistent on machines in this way.

Fileless attacks

In the field of malware there are many (possibly overlapping) classification categories, and amongst other things a distinction can be made between file-based and fileless malware. Since 2017, a significant increase in fileless threats has been recorded. One reason for this is the fact that such attacks have proved very successful from the attackers' point of view. One factor in their effectiveness is the fact that fileless threats operate only in the memory of the compromised system, making it harder for security solutions to recognise them. It is important that fileless threats are recognised by consumer security programs as well as by business products, for the reasons given below.

Attack vectors and targets

In penetration tests, we see that certain attack vectors may not yet be well covered by security programs, and many popular AV products still provide insufficient protection. Some business security products are now making improvements in this area, and providing better protection in some scenarios. As mentioned above, we believe that consumer products also need to improve their protection against such malicious attacks; non-business users can be, and are, attacked in the same way². Anyone can be targeted, for a variety of reasons, including "doxing" (publishing confidential personal information) as an act of revenge. Attacking the home computers of businesspeople is also an obvious route into accessing their company data.

Attack methods

In the Advanced Threat Protection Test, we also include several different command-line stacks, CMD/PS commands, which can download malware from the network directly into RAM (staged) or base64 encoded calls. These methods completely avoid disk access, which is (usually) well guarded by security products. We sometimes use simple concealment measures, or change the method of the stager call as well. Once the malware has loaded its second stage, an http/https connection to the attacker will be established. This inside-out mechanism has the advantage of establishing a C2 channel to the attacker that is beyond the protection measures of the majority of NAT and firewall products. Once the C2 tunnel has been established, the attacker can use all known control mechanisms of the common C2 products (Meterpreter, PowerShell Empire, etc.). These can include e.g. file uploads/downloads, screenshots, keylogging, Windows shell (GUI), and webcam snapshots. We expect attacks to be blocked regardless of where/how they are hosted and where from/how they are executed. If an attack is detected only under very specific circumstances, we would say the product does not provide effective protection.

² <https://www.bleepingcomputer.com/news/security/google-youtubers-accounts-hijacked-with-cookie-stealing-malware/>

False Positive (False Alarm) Test

A security product that blocks 100% of malicious attacks, but also blocks legitimate (non-malicious) actions, can be hugely disruptive. Consequently, we conduct a false-positives test as part of the Advanced Threat Protection Test, to check whether the tested products are able to distinguish malicious from non-malicious actions. Otherwise a security product could easily block 100% of malicious attacks that e.g. use email attachments, scripts and macros, simply by blocking such functions. For many users, this could make it impossible to carry out their normal daily tasks. Consequently, false-positive scores are taken into account in the product's test score.

We also note that warning the user against e.g. opening harmless email attachments can lead to a "boy who cried wolf" scenario. Users who encounter a number of unnecessary warnings will sooner or later assume that all warnings are false alarms, and thus ignore a genuine warning when it comes along.

Tested Products

The following vendors participated in the Advanced Threat Protection Test. These are the vendors who were confident enough in the protection capabilities of their products³ against targeted attacks to take part in this public test. All other vendors in the Consumer Main-Test Series opted out of the test.



Vendor	Product	Version
Avast	Free Antivirus	22.9 – 22.10
AVG	Free Antivirus	22.9 – 22.10
Avira	Prime	1.1
Bitdefender	Internet Security	26.0
ESET	Internet Security	15.2
G Data	Total Security	25.5
Kaspersky	Internet Security	21.3
McAfee	Total Protection	26.4
Microsoft	Defender	4.18

All consumer products were tested with default settings.

³ Information about additional third-party engines/signatures used inside the products: **G Data** and **VIPRE** use the Bitdefender engine. **AVG** is a rebranded version of Avast.

Test Results

Below are the results for the 15 attacks used in this test:

Test scenarios																	
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	FPs	Score
Avast	✓	✓	✗	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✓	N	10
AVG	✓	✓	✗	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✓	N	10
Avira	✓	✓	✗	✓	✗	✗	✗	✗	✓	✓	✗	✓	✓	✗	🛡️	Y	N/A
Bitdefender	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	N	13
ESET	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✓	✓	✓	N	14
G Data	✓	✓	✓	✓	✓	✓	✗	✓	✓	✓	✓	✓	✓	✗	✗	N	12
Kaspersky	✓	✓	✓	✓	✗	✓	✗	✓	✓	✗	✓	✓	✓	✓	✓	N	12
McAfee	✗	✗	✗	✓	✗	✓	✗	✗	✓	✗	✗	✗	✗	✗	✗	N	3
Microsoft	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✓	✗	✗	N	12

Key

✓	Threat blocked, no C2 session, system protected	1 point
🛡️	No alert shown, but no C2 session established, system protected	1 point
✗	Threat not blocked, C2 session established	0 points
✓	Protection result invalid, as also non-malicious scripts/functions were blocked	N/A

In our opinion, the goal of every AV/EPP/EDR system should be to detect and prevent attacks or other malware as soon as possible. In other words, if the attack is detected before, at or soon after execution, thus preventing e.g. the opening of a Command and Control Channel, there is no need to prevent post-exploitation activities. A good burglar alarm should go off when somebody breaks into your house, not wait until they start stealing things.

A product that blocked certain legitimate functions (e.g. email attachments or scripts) would be categorised only as "Tested".

Observations on consumer products

In this section, we report some additional information which could be of interest to readers.

Detection/Blocking stages

Pre-execution (PRE): when the threat has not been run, and is inactive on the system (static).

On-execution (ON): immediately after the threat has been run (dynamic).

Post-execution (POST): after the threat has been run, and its actions have been recognised (in-memory).

Test scenarios															
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Avast	POST	PRE	-	PRE	ON	ON	ON	ON	ON	ON	-	-	-	-	ON
AVG	POST	PRE	-	PRE	ON	ON	ON	ON	ON	ON	-	-	-	-	ON
Avira	ON*	ON*	-	PRE*	-	-	-	-	ON*	ON*	-	ON*	PRE*	-	??*
Bitdefender	PRE	PRE	ON	PRE	ON	ON	ON	PRE	PRE	PRE	ON	PRE	PRE	-	-
ESET	POST	ON	PRE	PRE	ON	PRE	ON	POST	PRE	ON	PRE	-	ON	ON	ON
G Data	PRE	PRE	ON	PRE	POST	ON	-	PRE	PRE	ON	ON	PRE	ON	-	-
Kaspersky	PRE	ON	ON	ON	-	ON	-	POST	PRE	-	PRE	ON	ON	ON	ON
McAfee	-	-	-	ON	-	ON	-	-	ON	-	-	-	-	-	-
Microsoft	PRE	PRE	PRE	PRE	ON	ON	PRE	POST	ON	POST	PRE	-	PRE	-	-

Avast, AVG: Most detections occurred on-execution.

Avira: Detections occurred mostly on-execution by Avira. In one case, Avira hindered the opening of the Word document, which stalled and did not open even after 20 minutes' wait. This is very strange behavior, especially as no detection/alert was provided by the product. *) In at least one case, Avira did not distinguish between malicious and non-malicious scripts/actions, so the results do not prove how well it can block attacks without also blocking legitimate scripts/actions.

Bitdefender: Detections occurred either pre- or on-execution.

ESET: Detections occurred mostly pre- or on-execution, with two post-execution.

G Data: Detections occurred mostly pre- or on-execution, with one post-execution.

Kaspersky: Detections occurred mostly pre- or on-execution, with one post-execution.

McAfee: Detections occurred on-execution.





Microsoft: Detections occurred mostly pre- or on-execution, with two post-execution.

All the tested vendors continuously implement improvements in the product, so it is to be expected that many of the missed attacks used in the test are now covered.

Award levels reached

From our experience, we know that many consumer AV programs do not provide effective protection against the types of threat used in this test. For this reason, a consumer security app that detects even 5 out of 15 threats is worthy of an award for "Advanced Threat Protection" (ATP). Precise criteria for awards in this test are given below:

	Blocked Threats (out of 15)			
	0-4	5-8	9-12	13-15
<i>No false alarms/functionality blocking</i>	TESTED	STANDARD	ADVANCED	ADVANCED+
<i>False alarms/functionality blocking seen</i>	TESTED	TESTED	TESTED	TESTED

AWARDS (based on protection rates and false alarms)	PRODUCTS
	<ul style="list-style-type: none"> ✓ ESET ✓ Bitdefender
	<ul style="list-style-type: none"> ✓ G Data ✓ Kaspersky ✓ Microsoft ✓ Avast ✓ AVG
	-
	<ul style="list-style-type: none"> ✓ Avira* ✓ McAfee

*) Products that blocked certain legitimate functions (e.g. email attachments or scripts) were downgraded to "Tested".

Test cases employed

We used five different [Initial Access Phases](#), distributed among the 15 test cases (e.g. 3 testcases came via email/spear-phishing attachment).

- a) [Trusted Relationship](#): “Adversaries may breach or otherwise leverage organizations who have access to intended victims. Access through trusted third-party relationship exploits an existing connection that may not be protected or receives less scrutiny than standard mechanisms of gaining access to a network.”
- b) [Valid accounts](#): “Adversaries may steal the credentials of a specific user or service account using Credential Access techniques or capture credentials earlier in their reconnaissance process through social engineering [...]”
- c) [Replication Through Removable Media](#): “Adversaries may move onto systems [...] by copying malware to removable media [...] and renaming it to look like a legitimate file to trick users into executing it on a separate system. [...]”
- d) [Phishing: Spearphishing Attachment](#): “Spearphishing attachment is [...] employs the use of malware attached to an email. [...]”
- e) [Phishing: Spearphishing Link](#): “Spearphishing with a link [...] employs the use of links to download malware contained in email [...]”

The 15 test scenarios used in this test are very briefly described below:

- 1) This threat is introduced via Spearphishing Link. A malicious binary executes x86 shellcode to open a meterpreter C2 channel via http.
- 2) This threat is introduced via Spearphishing Link. A malicious JavaScript file executes x64 shellcode to open a meterpreter C2 channel via http.
- 3) This threat is introduced via Spearphishing Link. A malicious obfuscated JavaScript injects x64 shellcode into an Office process to open a meterpreter C2 channel via http.
- 4) This threat is introduced via Valid Accounts. A malicious HTA file opens a meterpreter C2 channel via http.
- 5) This threat is introduced via Valid Accounts. A malicious PowerShell command with some defense evasion capabilities opens a meterpreter C2 channel via http.
- 6) This threat is introduced via Valid Accounts. A malicious Batch file opens an Empire C2 channel via http using a non-standard port.
- 7) This threat is introduced via Trusted Relationship. A malicious, obfuscated binary with some defense evasion capabilities and file extension spoofing, opens a PowerShell Empire C2 channel via http using a non-standard port.
- 8) This threat is introduced via Trusted Relationship. A malicious CPL file executes a PowerShell payload, which opens an Empire C2 channel via http using a non-standard port.
- 9) This threat is introduced via Trusted Relationship. A malicious XSL file is executed via WMI, which opens an obfuscated Empire C2 channel via http using a non-standard port.

10) This threat is introduced via Spearphishing Attachment. A malicious binary with a spoofed file extension executes an Empire payload to open an Empire C2 channel via http using a non-standard port.

11) This threat is introduced via Spearphishing Attachment. A malicious JavaScript file with some defense evasion capabilities enables to execute malicious code via the control panel application, and opens a C2 channel to a commercial C2 framework via https.

12) This threat is introduced via Spearphishing Attachment. A malicious binary with some defense evasion capabilities opens a C2 channel to a commercial C2 framework via https.

13) This threat is introduced via Removable Media. A malicious DLL opens a C2 channel via https to a commercial C2 framework.

14) This threat is introduced via Removable Media. A malicious binary with advanced defense evasion capabilities opens a C2 channel via https to a commercial C2 framework.

15) This threat is introduced via Removable Media. A malicious office document injects into another user-space process and opens a C2 channel to a commercial C2 framework via https.

False Alarm Test: Various false-alarm scenarios were used in order to see if any product is over-blocking certain actions (e.g. by blocking by policy email attachments, communication, scripts, etc.). None of the tested products showed over-blocking behaviour in the false-alarm test scenarios used. If during the course of the test, we were to observe products adapting their protection to our test environment, we would use countermeasures to evade these adaptations, to ensure that each product can genuinely detect the attack, as opposed to the test situation.

About this test

The Advanced Threat Protection Test for consumer products is an optional part of the Public Main-Test Series⁴. We congratulate those vendors who chose to take part. They have obviously worked hard on their products, and are using these public third-party tests as independent verification that their products do what they claim. The complex nature of the test means that automation is not possible, and it has to be performed entirely manually, making it cost-intensive to run. However, vendors in the Consumer Main-Test Series had the opportunity to participate in the Public Advanced Threat Protection Test of 2022 at no additional cost to themselves.

As some of the attack methods used in the test make use of legitimate system programs and techniques, it would be fairly easy for a vendor to stop such attacks e.g. simply by blocking the use of these legitimate processes. However, this would result in the product concerned being marked down for false positives, in the same way that a security program would be marked down for e.g. blocking all unknown executable program files. Likewise, in this test, preventing an attack e.g. by simply blacklisting used servers, files or emails originating from a particular domain name would not be allowed as a means of preventing a targeted attack. Similarly, we do not accept an approach which does not distinguish between malicious and non-malicious processes, but requires e.g. an admin to whitelist ones that should be allowed.

⁴ <https://www.av-comparatives.org/consumer/>

In our Consumer Main-Test Series, products are tested with their default settings. In the Business Main-Test Series, vendors are allowed to configure the products as they see fit – as is common practice with business security products in the real world. However, precisely the same product and configuration is used for all the tests in the series. If we did not insist on this, a vendor could turn up protection settings or activate features in order to score highly in the Real-World and Malware Protection Tests, but turn them down/deactivate them for the Performance and False Positive Tests, in order to appear faster and less error-prone. In real life, users can only have one setting at once, so they should be able to see if high protection scores mean slower system performance, or lower false-positive scores mean reduced protection.

We had requests from vendors regarding the attack methods to be used in the test. We did not divulge specific details of the attack methods, but after the test, we provided each participating vendor with sufficient data to demonstrate the missed test cases.

The test is very challenging, but at the same time it also reflects realistic scenarios. Penetration testers see the real capabilities of products in their tests every day. Our comparison test tries to create a level playing-field that allows us to fairly compare the protection capabilities of the different products against such attacks. This lets users see how well they are protected, and allows vendors, where necessary, to improve their products in the future.

Whilst this test is for consumer products, the attack techniques used are the same as for our Enterprise ATP test. Hackers may be highly motivated to attack the home computers of specific, high-profile individuals, e.g. politicians or the very wealthy. We also note that targeted attacks on enterprises may begin by taking control of the home computer of the CEO, for example.

As regards the Windows User Accounts used in the test, none of the test scenarios required administrator permissions on the targeted system. Hence, from an attacker's perspective it made no difference whether the user was logged on with an Administrator Account (as used for the Consumer Test) or Standard User Account (as used for the Enterprise Test).

In some of the attacks, as noted in the testcase descriptions, Initial Access vectors such as Trusted Relationships and Valid Accounts were used. That is to say, the attacker already had the necessary user credentials needed to proceed with the advanced attack. Various studies show that such scenarios (e.g. using stolen credentials for Initial Access) are very common nowadays.

In some cases, the test made use of redirected drives. As the ATT&CK framework does not have a specific category for such cases, we feel that they fall best into the category "removable media", which we have noted in the descriptions. Nevertheless, how the malware was introduced into the system made no technical difference in practice.

In general, we have received positive feedback on the thorough and realistic methodology of this test from security vendors. We are happy to say that a number of vendors not in this year's test are improving their products' protection against real-life targeted attacks, and aim to join the test next year. We have also considered some suggestions for improvement from some vendors, and will endeavour to include these in next year's test where appropriate.

Copyright and Disclaimer

This publication is Copyright © 2022 by AV-Comparatives®. Any use of the results, etc. in whole or in part, is ONLY permitted after the explicit written agreement of the management board of AV-Comparatives prior to any publication. AV-Comparatives and its testers cannot be held liable for any damage or loss, which might occur as result of, or in connection with, the use of the information provided in this paper. We take every possible care to ensure the correctness of the basic data, but a liability for the correctness of the test results cannot be taken by any representative of AV-Comparatives. We do not give any guarantee of the correctness, completeness, or suitability for a specific purpose of any of the information/content provided at any given time. No one else involved in creating, producing or delivering test results shall be liable for any indirect, special or consequential damage, or loss of profits, arising out of, or related to, the use or inability to use, the services provided by the website, test documents or any related data.

For more information about AV-Comparatives and the testing methodologies, please visit our website.

AV-Comparatives
(November 2022)