SE Labs

INTELLIGENCE-LED TESTING

Endpoint Security

Enterprise

Jul - Sep 2022

EPS PROTECTION

SE Labs tested a variety of anti-malware (aka 'anti-virus'; aka 'endpoint security') products from a range of well-known vendors in an effort to judge which were the most effective.

Each product was exposed to the same threats, which were a mixture of targeted attacks using well-established techniques and public email and web-based threats that were found to be live on the internet at the time of the test.

The results indicate how effectively the products were at detecting and/or protecting against those threats in real time.

MANAGEMENT
Chief Executive Officer Simon Edwards
Chief Operations Officer Marc Briggs
Chief Human Resources Officer Magdalena Jurenko
Chief Technical Officer Stefan Dumitrascu

TESTING TEAM
Nikki Albesa
Thomas Bean
Solandra Brewster
Gia Gorbold
Anila Johny
Erica Marotta
Luca Menegazzo
Jeremiah Morgan
Julian Owusu-Abrokwa
Joseph Pike
Georgios Sakatzidis
Dimitrios Tsarouchas
Stephen Withey

IT SUPPORT
Danny King-Smith
Chris Short

PUBLICATION
Sara Claridge
Colin Mackleworth

Contents

Anti-Virus Testing
How hard should a security test be?

Thank you for opening this report. I hope you'll be able to use it to get a better idea about which anti-malware products you might want to buy (or get rid of!)

The report starts off with a list of products, each of which win impressive-looking awards. But have you considered what those awards mean? How come there aren't any massive losers in the list? How hard is this security test anyway?

Baseline Testing
There are lots of ways you can test products. You could prod a teddy bear and say, "well, that looks good enough," or you could take it to pieces and analyse every component forensically for build and functional quality. "This toy looks safe, its parts are large, soft and non-toxic, and we can't burn it easily. Plus, it's got big, cute eyes." This could be a baseline for cuddly toys: SAFE, with cuteness as an extra bonus.

For anti-malware products we have to consider a few different things, including the following:

- Is it really an anti-malware product? Is it at least basically functional?
- Can it determine a good quantity of common malware, without blocking lots of useful software?
- Can it stop the malware, as well as simply detecting it?

How hard do you want your security testing to be? We could take a product, scan a real virus (or the harmless EICAR test file) and record that it detected a threat. Is that good enough? It's good enough to answer the first question above, and to reassure you that you've installed the software correctly. But you can't tell if it's better than other anti-malware products, because most will all react the same.

You also can't tell if the product has extra functionality capable of detecting and stopping other types of threats, of which there are many. Scanning files is a basic way to test anti-malware. In most cases, it's too basic.

Attacks For Everyone
Let's turn up the dial and throw a wider range of attacks at the products. We can use malware that bad guys use to attack everyone, every day. We call these threats 'commodity attacks' because they are all over the place, indiscriminately damaging computers all over the world. Our hope is that all main anti-malware products will detect each of these threats. To us, that's the baseline. If a product can't do that, there's something very wrong going on.

We go further than that, though. By using forensics, we can tell not only how many threats a product can recognise, but how well it can protect against them too. It might say, "I see a virus and I've blocked it!" but we don't trust that claim. We check that it's really done what it said it did.

Attacks Just For You
We could go harder, and use targeted attacks designed to evade anti-malware. So we do. That's the 'targeted attacks' part of the test. You can see how the products handled them in **Appendix D**, buried at the end of this report. For me, that's the most interesting part!

You don't generally see products that perform less well in our tests because we think you probably don't even want to consider them. Don't waste your time. But sometimes they won't let us test them. If your favourite anti-malware vendor isn't in our report, please tell them to play!

Any A, AA or AAA awards that products achieve in our reports show that they go well beyond basic functionality. It shows that they can handle both common and customised threats, without blocking the software you need to run on your computer. For even more in-depth testing check out the Enterprise Advanced Security test reports on our **website**.

If you spot a detail in this report that you don't understand, or would like to discuss, please contact us via our **Twitter** or **LinkedIn** accounts. SE Labs uses current threat intelligence to make our tests as realistic as possible. To learn more about how we test, how we define 'threat intelligence' and how we use it to improve our tests please visit our **website** and follow us on **Twitter**.

Executive Summary

Product Names

It is good practice to stay up to date with the latest version of your chosen endpoint security product. We made best efforts to ensure that each product tested was the very latest version running with the most recent updates to give the best possible outcome.

For specific build numbers, see **Appendix C: Product Versions** on page 19.

Executive Summary			
Products Tested	Protection Accuracy Rating (%)	Legitimate Accuracy Rating (%)	Total Accuracy Rating (%)
Kaspersky Endpoint Security	100%	100%	100%
Microsoft Defender Antivirus (enterprise)	100%	100%	100%
Sophos Intercept X	100%	100%	100%
ESET Endpoint Security	99%	100%	100%
Broadcom Endpoint Security Enterprise Edition	100%	98%	99%
CrowdStrike Falcon	97%	100%	99%
Fortinet FortiEDR	98%	99%	99%
Trellix Endpoint Security	97%	100%	99%
VIPRE Endpoint Security	98%	98%	98%

Products highlighted in green were the most accurate, scoring 85 per cent or more for Total Accuracy. Those in yellow scored less than 85 but 75 or more. Products shown in red scored less than 75 per cent.

For exact percentages, see **1. Total Accuracy Ratings** on page 6.

● **The endpoints were generally effective at handling general threats from cyber criminals...**
Most products were very capable of handling public email- and web-based threats such as those used by criminals to attack Windows PCs, tricking users into running malicious files or running scripts that download and run malicious files. However, three were not completely effective.

● **... but targeted attacks caused problems for a couple of the products.**
Seven of the nine products provided complete protection against the targeted attacks used in this test. It only takes one targeted attack to breach an organisation, so this is a concerning result.

● **False positives were not an issue for the products.**
Most of the products were perfectly good at correctly classifying legitimate applications and websites. Three products delegated the decision to allow or block to users (advising to block) in one case each.

● **Which products were the most effective?**
Products from **Microsoft**, **Sophos** and **Kaspersky** produced extremely good results due to a combination of their ability to block malicious URLs, handle exploits and correctly classify legitimate applications and websites. All products performed well enough to achieve AAA awards.

1. Total Accuracy Ratings

Judging the effectiveness of an endpoint security product is a subtle art, and many factors are at play when assessing how well it performs. To make things easier we've combined all the different results from this report into one easy-to-understand graph.
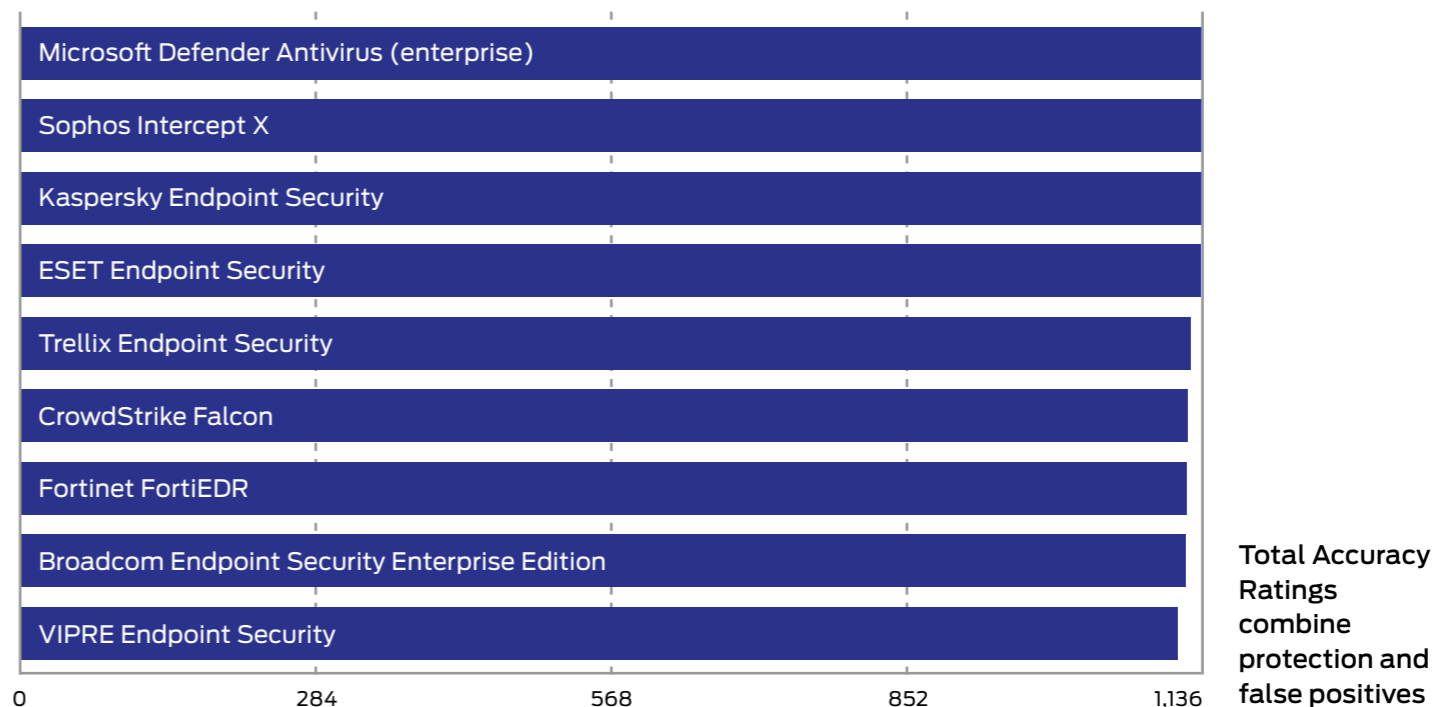
The graph below takes into account not only each product's ability to detect and protect against threats, but also its handling of non-malicious objects such as web addresses (URLs) and applications.

Not all protections, or detections for that matter, are equal. A product might completely block a URL, which stops the threat before it can even start its intended series of malicious events. Alternatively, the product might allow a web-based exploit to execute but prevent it from downloading any further code to the target. In another case malware might run on the target for a short while before its behaviour is detected and its code is deleted or moved to a safe 'quarantine' area for future analysis. We take these outcomes into account when attributing points that form final ratings.

For example, a product that completely blocks a threat is rated more highly than one that allows a threat to run for a while before eventually evicting it. Products that allow all malware infections, or that block popular legitimate applications, are penalised heavily.

Categorising how a product handles legitimate objects is complex, and you can find out how we do it in **6. Legitimate Software Ratings** on page 14.

Total Accuracy Ratings			
Product	Total Accuracy Rating	Total Accuracy (%)	Award
Microsoft Defender Antivirus (enterprise)	1,135	100%	AAA
Sophos Intercept X	1,135	100%	AAA
Kaspersky Endpoint Security	1,134	100%	AAA
ESET Endpoint Security	1,133	100%	AAA
Trellix Endpoint Security	1,125	99%	AAA
CrowdStrike Falcon	1,124	99%	AAA
Fortinet FortiEDR	1,120	99%	AAA
Broadcom Endpoint Security Enterprise Edition	1,119	99%	AAA
VIPRE Endpoint Security	1,111	98%	AAA

Microsoft Defender Antivirus (enterprise)
Sophos Intercept X
Kaspersky Endpoint Security
ESET Endpoint Security
Trellix Endpoint Security
CrowdStrike Falcon
Fortinet FortiEDR
Broadcom Endpoint Security Enterprise Edition
VIPRE Endpoint Security

0   284   568   852   1,136

Total Accuracy Ratings combine protection and false positives

Enterprise Endpoint Security Awards

The following products win SE Labs awards:

- **Microsoft** Defender Antivirus (enterprise)

- **Sophos** Intercept X

- **Kaspersky** Endpoint Security

- **ESET** Endpoint Security

- **Trellix** Endpoint Security

- **CrowdStrike** Falcon

- **Fortinet** FortiEDR

- **Broadcom** Endpoint Security Enterprise Edition

- **VIPRE** Endpoint Security

SE Labs
AAA
Jul - Sep 2022
Enterprise Endpoint Security
EPS Protection

SE Labs Monthly Newsletter

Don't miss our security articles and reports

- Test reports announced
- Blog posts reviewed
- Security testing analysed
- **NEW:** Podcast episodes

SE Labs
INTELLIGENCE-LED TESTING

News from the security lab

RANSOMWARE

FREE

SUBSCRIBE NOW!

2. Threat Responses

Full Attack Chain: Testing every layer of detection and protection

Attackers start from a certain point and don't stop until they have either achieved their goal or have reached the end of their resources (which could be a deadline or the limit of their abilities).

This means, in a test, the tester needs to begin the attack from a realistic first position, such as sending a phishing email or setting up an infected website, and moving through many of the likely steps leading to actually stealing data or causing some other form of damage to the network.

If the test starts too far into the attack chain, such as executing malware on an endpoint, then many products will be denied opportunities to use the full extent of their protection and detection abilities.

If the test concludes before any 'useful' damage or theft has been achieved, then similarly the product may be denied a chance to demonstrate its abilities in behavioural detection and so on.

Attack Stages
The illustration below shows some typical stages of an attack. In a test each of these should be

Attack Chain: How Hackers Progress

Figure 1. A typical attack starts with an initial contact and progresses through various stages, including reconnaissance, stealing data and causing damage.

Figure 2. This attack was initially successful but only able to progress as far as the reconnaissance phase.

Figure 3. A more successful attack manages to steal passwords but wholesale data theft and destruction was blocked.

attempted to determine the security solution's effectiveness. This test's results record detection and protection for each of these stages.

We measure how a product responds to the first stages of the attack with a detection and/ or protection rating. Sometimes products allow threats to run but detect them. Other times they might allow the threat to run briefly before neutralising it. Ideally they detect and block the threat before it has a chance to run. Products may delete threats or automatically contain them in a 'quarantine' or other safe holding mechanism for later analysis.

Should the initial attack phase succeed we then measure post-exploitation stages, which are represented by steps two through to seven

below. We broadly categorise these stages as: Access (step 2); Action (step 3); Escalation (step 4); and Post-escalation (step 5).

In figure 1. you can see a typical attack running from start to end, through various 'hacking' activities. This can be classified as a fully successful breach.
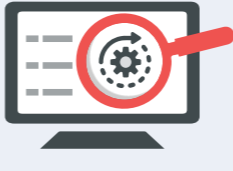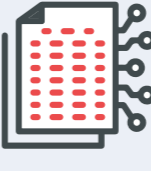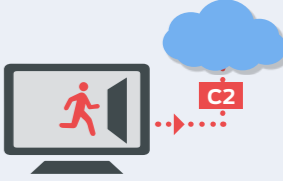
In figure 2. a product or service has interfered with the attack, allowing it to succeed only as far as stage 3, after which it was detected and neutralised. The attacker was unable to progress through stages 4 and onwards.

It is possible that attackers will not cause noticeable damage during an attack. It may be that their goal is persistent presence on the

systems to monitor for activities, slowly steal information and other more subtle missions.

In figure 3. the attacker has managed to progress as far as stage five. This means that the system has been seriously compromised. The attacker has a high level of access and has stolen passwords. However, attempts to exfiltrate data from the target were blocked, as were attempts to damage the system.

The table below shows how a typical way in which security testers illustrate attackers' behaviour. It is largely the same as our images above, but more detailed.

MITRE Example Attack Chain Details

Initial Access	Execution	Privilege Escalation	Credential Access	Discovery	Collection	Command and Control	Exfiltration
Spearphishing via Service	Command-Line Interface	Bypass UAC	Input Capture	File and Directory Discovery	Input Capture	Data Encoding	Exfiltration Over C2 Channel
Spearphishing Link	PowerShell		OS Credential Dumping	Process Discovery	Data from Local System	Data Obfuscation	
	Scripting			System Information Discovery			
	User Execution						
Spearphishing Link	Scripting	Bypass UAC	OS Credential Dumping	Process Discovery	Data from Local System	Data Obfuscation	Exfiltration Over C2 Channel

3. Protection Ratings

The results below indicate how effectively the products dealt with threats. Points are earned for detecting the threat and for either blocking or neutralising it.

■ **Detected (+1)**

If the product detects the threat with any degree of useful information, we award it one point.

■ **Blocked (+2)**

Threats that are disallowed from even starting their malicious activities are blocked. Blocking products score two points.

■ **Complete Remediation (+1)**

If, in addition to neutralising a threat, the product removes all significant traces of the attack, it gains an additional one point.

■ **Neutralised (+1)**

Products that kill all running malicious processes 'neutralise' the threat and win one point.

■ **Persistent Neutralisation (-2)**

This result occurs when a product continually blocks a persistent threat from achieving its aim, while not removing it from the system.

■ **Compromised (-5)**

If the threat compromises the system, the product loses five points. This loss may be reduced to four points if it manages to detect the threat (see Detected, above), as this at least

alerts the user, who may now take steps to secure the system.

Rating Calculations

We calculate the protection ratings using the following formula:

Protection Rating =
(1x number of Detected) +
(2x number of Blocked) +
(1x number of Neutralised) +
(1x number of Complete remediation) +
(-5x number of Compromised)

The 'Complete remediation' number relates to cases of neutralisation in which all significant traces of the attack were removed from the target.

These ratings are based on our opinion of how important these different outcomes are. You may have a different view on how seriously you treat a 'Compromise' or 'Neutralisation without complete remediation'. If you want to create your own rating system, you can use the raw data from **5. Protection Details** on page 13 to roll your own set of personalised ratings.

Targeted Attack Scoring

The following scores apply only to targeted attacks and are cumulative, ranging from -1 to -5.

■ **Access (-1)**

If any command that yields information about the

target system is successful this score is applied. Examples of successful commands include listing current running processes, exploring the file system and so on. If the first command is attempted and the session is terminated by the product without the command being successful the score of Neutralised (see above) will be applied.

■ **Action (-1)**

If the attacker is able to exfiltrate a document from the target's Desktop of the currently logged in user then an 'action' has been successfully taken.

■ **Escalation (-2)**

The attacker attempts to escalate privileges to NT Authority/System. If successful, an additional two points are deducted.

■ **Post-Escalation Action (-1)**

After escalation the attacker attempts actions that rely on escalated privileges. These include attempting to steal credentials, modifying the file system and recording keystrokes. If any of these actions are successful then a further penalty of one point deduction is applied.

Protection Accuracy

Product	Protection Accuracy	Protection Accuracy (%)
Broadcom Endpoint Security Enterprise Edition	399	100%
Microsoft Defender Antivirus (enterprise)	399	100%
Sophos Intercept X	399	100%
Kaspersky Endpoint Security	398	100%
ESET Endpoint Security	397	99%
Fortinet FortiEDR	393	98%
VIPRE Endpoint Security	391	98%
Trellix Endpoint Security	389	97%
CrowdStrike Falcon	388	97%

Average 99%

Broadcom Endpoint Security Enterprise Edition

Microsoft Defender Antivirus (enterprise)

Sophos Intercept X

Kaspersky Endpoint Security

ESET Endpoint Security

Fortinet FortiEDR

VIPRE Endpoint Security

Trellix Endpoint Security

CrowdStrike Falcon

0   100   200   300   400
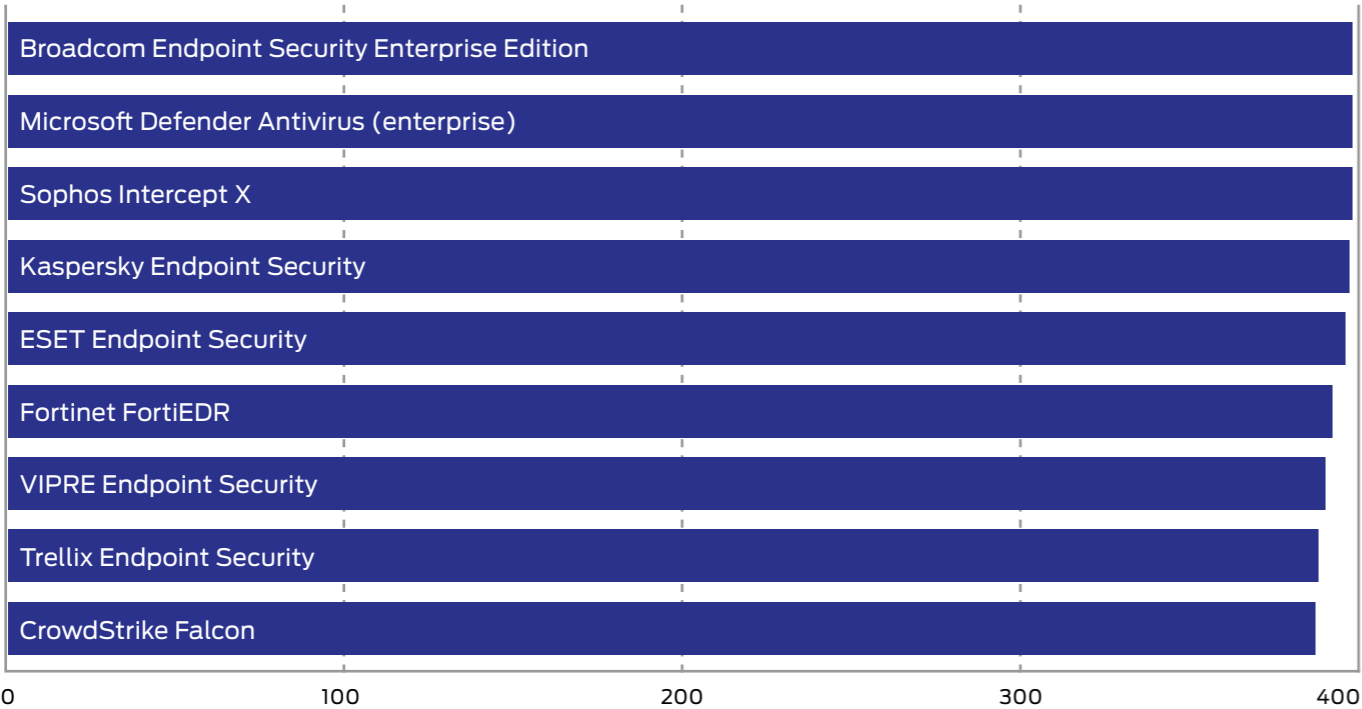
Protection Ratings are weighted to show that how products handle threats can be subtler than just 'win' or 'lose'.
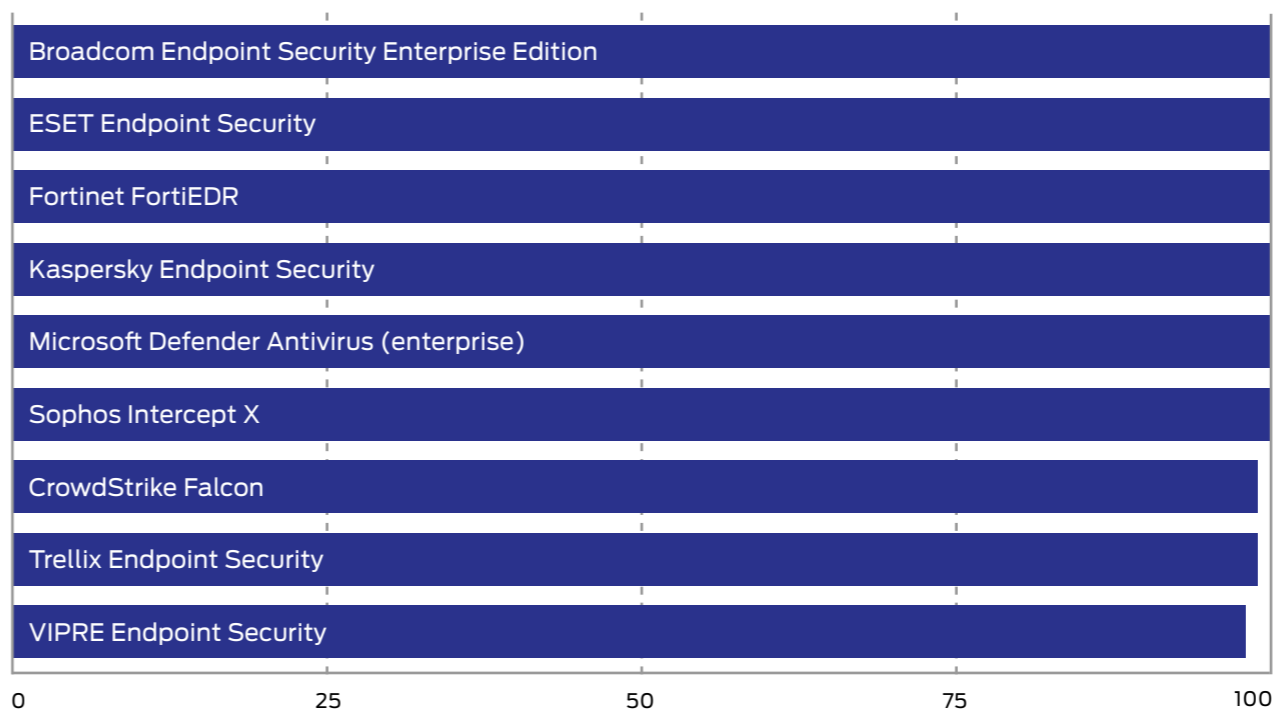
4. Protection Scores

This graph shows the overall level of protection, making no distinction between neutralised and blocked incidents.

For each product we add Blocked and Neutralised cases together to make one simple tally.

Protection Scores	
Product	Protection Score
Broadcom Endpoint Security Enterprise Edition	100
ESET Endpoint Security	100
Fortinet FortiEDR	100
Kaspersky Endpoint Security	100
Microsoft Defender Antivirus (enterprise)	100
Sophos Intercept X	100
CrowdStrike Falcon	99
Trellix Endpoint Security	99
VIPRE Endpoint Security	98

Broadcom Endpoint Security Enterprise Edition

ESET Endpoint Security

Fortinet FortiEDR

Kaspersky Endpoint Security

Microsoft Defender Antivirus (enterprise)

Sophos Intercept X

CrowdStrike Falcon

Trellix Endpoint Security

VIPRE Endpoint Security

0   25   50   75   100

Protection Scores are a simple count of how many times a product protected the system.

5. Protection Details

These results break down how each product handled threats into some detail. You can see how many detected a threat and the levels of protection provided.

Products sometimes detect more threats than they protect against. This can happen when they recognise an element of the threat but aren't equipped to stop it. Products can also provide protection even if they don't detect certain threats. Some threats abort on detecting specific endpoint protection software.

Protection Details					
Product	Detected	Blocked	Neutralised	Compromised	Protected
ESET Endpoint Security	100	100	0	0	100
Sophos Intercept X	100	100	0	0	100
Broadcom Endpoint Security Enterprise Edition	100	99	1	0	100
Kaspersky Endpoint Security	100	99	1	0	100
Microsoft Defender Antivirus (enterprise)	100	99	1	0	100
Fortinet FortiEDR	100	98	2	0	100
Trellix Endpoint Security	99	99	0	1	99
CrowdStrike Falcon	100	96	3	1	99
VIPRE Endpoint Security	100	98	0	2	98

This data shows in detail how each product handled the threats used.
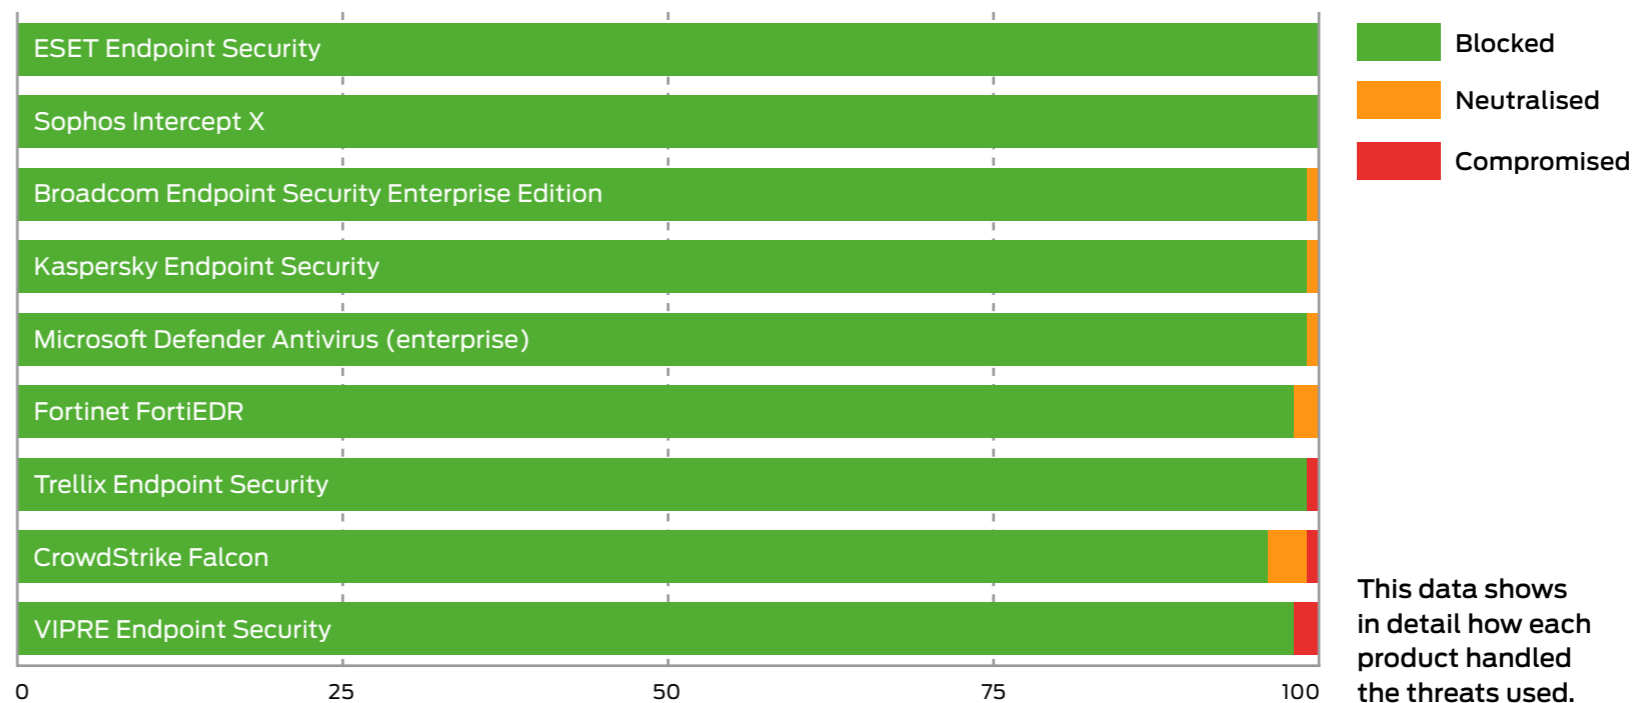
6. Legitimate Software Ratings

These ratings indicate how accurately the products classify legitimate applications and URLs, while also taking into account the interactions that each product has with the user. Ideally a product will either not classify a legitimate object or will classify it as safe. In neither case should it bother the user.

We also take into account the prevalence (popularity) of the applications and websites used in this part of the test, applying stricter penalties for when products misclassify very popular software and sites.

To understand how we calculate these ratings, see **6.3 Accuracy Ratings** on page 16.

Legitimate Software Ratings		
Product	Legitimate Accuracy Rating	Legitimate Accuracy (%)
CrowdStrike Falcon	736	100%
ESET Endpoint Security	736	100%
Kaspersky Endpoint Security	736	100%
Microsoft Defender Antivirus (enterprise)	736	100%
Sophos Intercept X	736	100%
Trellix Endpoint Security	736	100%
Fortinet FortiEDR	727	99%
Broadcom Endpoint Security Enterprise Edition	720	98%
VIPRE Endpoint Security	720	98%

CrowdStrike Falcon

ESET Endpoint Security

Kaspersky Endpoint Security

Microsoft Defender Antivirus (enterprise)

Sophos Intercept X

Trellix Endpoint Security

Fortinet FortiEDR

Broadcom Endpoint Security Enterprise Edition

VIPRE Endpoint Security

0   184   368   552   736

Legitimate Software Ratings can indicate how well a vendor has tuned its detection engine.

6.1 Interaction Ratings

It's crucial that anti-malware endpoint products not only stop – or at least detect – threats, but that they allow legitimate applications to install and run without misclassifying them as malware. Such an error is known as a 'false positive' (FP).

In reality, genuine FPs are quite rare in testing. In our experience it is unusual for a legitimate application to be classified as 'malware'. More often it will be classified as 'unknown', 'suspicious' or 'unwanted' (or terms that mean much the same thing).

We use a subtle system of rating an endpoint's approach to legitimate objects, which takes into account how it classifies the application and how it presents that information to the user. Sometimes the endpoint software will pass the buck and demand that the user decide if the application is safe or not. In such cases the product may make a recommendation to allow or block. In other cases, the product will make no recommendation, which is possibly even less helpful.

If a product allows an application to install and run with no user interaction, or with simply a brief notification that the application is likely to be safe, it has achieved an optimum result. Anything else is a Non-Optimal Classification/Action (NOCA). We think that measuring NOCAs is more useful than counting the rarer FPs.

	None (allowed)	Click to Allow (default allow)	Click to Allow/Block (no recommendation)	Click to Block (default block)	None (blocked)	
Object is Safe	2	1.5	1			A
Object is Unknown	2	1	0.5	0	-0.5	B
Object is not Classified	2	0.5	0	-0.5	-1	C
Object is Suspicious	0.5	0	-0.5	-1	-1.5	D
Object is Unwanted	0	-0.5	-1	-1.5	-2	E
Object is Malicious				-2	-2	F
	1	2	3	4	5	

Interaction Ratings		
Product	None (allowed)	Click to block (default block)
CrowdStrike Falcon	100	0
ESET Endpoint Security	100	0
Kaspersky Endpoint Security	100	0
Microsoft Defender Antivirus (enterprise)	100	0
Sophos Intercept X	100	0
Trellix Endpoint Security	100	0
Broadcom Endpoint Security Enterprise Edition	99	1
Fortinet FortiEDR	99	1
VIPRE Endpoint Security	99	1

Products that do not bother users and classify most applications correctly earn more points than those that ask questions and condemn legitimate applications.

6.2 Prevalence Ratings

There is a significant difference between an endpoint product blocking a popular application such as the latest version of Microsoft Word and condemning a rare Iranian dating toolbar for Internet Explorer 6. One is very popular all over the world and its detection as malware (or something less serious but still suspicious) is a big deal. Conversely, the outdated toolbar won't have had a comparably large user base even when it was new. Detecting this application as malware may be wrong, but it is less impactful in the overall scheme of things.

With this in mind, we collected applications of varying popularity and sorted them into five separate categories, as follows:

1. Very High Impact
2. High Impact
3. Medium Impact
4. Low Impact
5. Very Low Impact

Incorrectly handling any legitimate application will invoke penalties, but classifying Microsoft Word as malware and blocking it without any way for the user to override this will bring far greater penalties than doing the same for an ancient niche toolbar. In order to calculate these relative penalties, we assigned each impact category with a rating modifier, as shown in the table above.

Legitimate Software Prevalence Rating Modifiers	
Impact Category	Rating Modifier
Very High Impact	5
High Impact	4
Medium Impact	3
Low Impact	2
Very Low Impact	1

Applications were downloaded and installed during the test, but third-party download sites were avoided and original developers' URLs were used where possible. Download sites will sometimes bundle additional components into applications' install files, which may correctly cause anti-malware products to flag adware. We remove adware from the test set because it is often unclear how desirable this type of code is.

The prevalence for each application and URL is estimated using metrics such as third-party download sites and the data from Alexa.com's global traffic ranking system.

6.3 Accuracy Ratings

We calculate legitimate software accuracy ratings by multiplying together the interaction and prevalence ratings for each download and installation:

Accuracy rating = Interaction rating x Prevalence rating

If a product allowed one legitimate, Medium impact application to install with zero interaction with the user, then its Accuracy rating would be calculated like this:

Accuracy rating = 2 x 3 = 6

This same calculation is made for each legitimate application/site in the test and the results are summed and used to populate the graph and table shown under **6. Legitimate Software Ratings** on page 14.

6.4 Distribution of Impact Categories

Endpoint products that were most accurate in handling legitimate objects achieved the highest ratings. If all objects were of the highest prevalence, the maximum possible rating would be 1,000 (100 incidents x (2 interaction rating x 5 prevalence rating)).

In this test there was a range of applications with different levels of prevalence. The table below shows the frequency:

Legitimate Software Category Frequency	
Prevalence Rating	Frequency
Very High Impact	32
High Impact	32
Medium Impact	16
Low Impact	12
Very Low Impact	8

7. Conclusions

Attacks in this test included threats that affect the wider public and more closely targeted individuals and organisations. You could say that we tested the products with 'public' malware and full-on hacking attacks.

We introduced the threats in a realistic way such that threats seen in the wild on websites were downloaded from those same websites, while threats caught spreading through email were delivered to our target systems as emails.

All of the products tested are well-known and should do well in this test. While we do 'create' threats by using publicly available free hacking tools, we do not write unique malware so there is no technical reason why any vendor being tested should do poorly.

The results were generally strong but there were differences in the way the products handled threats. For example, CrowdStrike protected against all of the targeted attacks, but missed one of the more common threats. On the other hand, VIPRE stopped all of the general attacks but missed two targeted threats.

Six products stopped all of the attacks. These belonged to Broadcom, ESET, Fortinet, Kaspersky, Microsoft and Sophos. Trellix and CrowdStrike came in strongly behind, each missing only one attack.

Although the performances from CrowdStrike and VIPRE were weaker, they were still strong enough to make it into the AAA rating zone.

Targeted attacks are usually more of a challenge because they are not as widely known, so it's interesting that Trellix missed a general threat but stopped all of the more evasive targeted attacks.

Most of the products handled the legitimate applications correctly, with no mistakes. Three asked the user to make the decision in one case each, recommending the block option. Despite that, this is a particularly strong performance from all of the products.

Unusually all of the products in this test win AAA awards. The strongest, from Kaspersky, Microsoft and Sophos stopped all threats and allowed all legitimate applications. Broadcom, CrowdStrike, Fortinet and Trellix received total ratings of 99%, which is also outstanding.

Appendices

Appendix A: Terms Used

Term	Meaning
Compromised	The attack succeeded, resulting in malware running unhindered on the target. In the case of a targeted attack, the attacker was able to take remote control of the system and carry out a variety of tasks without hindrance.
Blocked	The attack was prevented from making any changes to the target.
False positive	When a security product misclassifies a legitimate application or website as being malicious, it generates a 'false positive'.
Neutralised	The exploit or malware payload ran on the target but was subsequently removed.
Complete Remediation	If a security product removes all significant traces of an attack, it has achieved complete remediation.
Target	The test system that is protected by a security product.
Threat	A program or sequence of interactions with the target that is designed to take some level of unauthorised control of that target.
Update	Security vendors provide information to their products in an effort to keep abreast of the latest threats. These updates may be downloaded in bulk as one or more files, or requested individually and live over the internet.

Appendix B: FAQs

A **full methodology** for this test is available from our website.

- The products chosen for this test were selected by SE Labs.
- The test was unsponsored.
- The test was conducted between 4th July and 26th August 2022.
- All products were configured according to each vendor's recommendations, when such recommendations were provided.
- Malicious URLs and legitimate applications and URLs were independently located and verified by SE Labs.
- Targeted attacks were selected and verified by SE Labs.
- Malicious and legitimate data was provided to partner organisations once the test was complete.
- The web browser used in this test was Google Chrome. When testing Microsoft products Chrome was equipped with the Windows Defender Browser Protection browser extension (**https://browserprotection.microsoft.com**). We allow other browser extensions when a tested product requests a user install one or more.

Q **What is a partner organisation? Can I become one to gain access to the threat data used in your tests?**

A Partner organisations benefit from our consultancy services after a test has been run. Partners may gain access to low-level data that can be useful in product improvement initiatives and have permission to use award logos, where appropriate, for marketing purposes. We do not share data on one partner with other partners. We do not partner with organisations that do not engage in our testing.

Q **I am a security vendor and you tested my product without permission. May I access the threat data to verify that your results are accurate?**

A We are willing to share a certain level of test data with non-partner participants for free. The intention is to provide sufficient data to demonstrate that the results are accurate. For more in-depth data suitable for product improvement purposes we recommend becoming a partner.

Appendix C: Product Versions

The table below shows the service's name as it was being marketed at the time of the test.

Product Versions			
Vendor	Product	Build Version (start)	Build Version (end)
Broadcom	Endpoint Security Enterprise Edition	14 (14.3.RU5) Build: 8268 (14.3.8268.5000)	Version 14 (14.3 RU5) Build: 8268 (14.3.8268.5000)
CrowdStrike	Falcon	6.39.15316	6.44.15803.0
ESET	Endpoint Security	9.0.2046.0	9.0.2046.0
Fortinet	FortiEDR	5.2.0.2068	5.2.0.2068
Kaspersky	Endpoint Security	11.8.0.384	11.9.0.351
Microsoft	Defender Antivirus (enterprise)	Antimalware Client Version: 4.18.2205.7 Engine Version: 1.1.19500.2 Antivirus Version: 1.373.999.0 Antispyware Version: 1.373.999.0	Antimalware Client Version: 4.18.2205.7 Engine Version: 1.1.19500.2 Antivirus Version: 1.373.999.0 Antispyware Version: 1.373.999.0
Sophos	Intercept X	Core Agent: 2022.1.1.3	Core Agent: 2022.2.1.9
Trellix	Endpoint Security	Endpoint Security v. 10.7 Agent: v.5.7.6.251 Data Exchange Layer: v.6.0.3.646	Endpoint Security: v.10.7 Agent: v.5.7.6.251 Data Exchange Layer: v.6.0.3.646
VIPRE	Endpoint Security	12.0.7874 Definition version: 102320-7.92208	12.0.7874 Definition version: 103768-7.92690

Appendix D: Attack Types

The table below shows how each product protected against the different types of attacks used in the test.

Attack Types			
Product	General Attack	Targeted Attack	Protected (%)
Broadcom Endpoint Security Enterprise Edition	75	25	100
ESET Endpoint Security	75	25	100
Fortinet FortiEDR	75	25	100
Kaspersky Endpoint Security	75	25	100
Microsoft Defender Antivirus (enterprise)	75	25	100
Sophos Intercept X	75	25	100
Trellix Endpoint Security	74	25	99.5
Crowdstrike Falcon	75	24	98
VIPRE Endpoint Security	75	23	96