

SELabs

INTELLIGENCE-LED TESTING



www.SELabs.uk



info@SELabs.uk



[@SELabsUK](https://twitter.com/SELabsUK)




www.facebook.com/selabsuk



blog.selabs.uk

PREDICTIVE MALWARE RESPONSE TEST

MARCH 2018





SE Labs tested **CylancePROTECT** in an offline environment against major threats that subsequently appeared in the wild. The test explores the product's ability to prevent new threats from attacking endpoint systems successfully.

CylancePROTECT contains technology designed to identify and block malware using what it claims to be an "artificial intelligence" (AI) model. This model can be updated over time. However, in this test we used the model created in May 2015 and did not permit further updates so that the software was unable to receive new models or edit the existing one.

The test exposed systems protected by this older version of **CylancePROTECT** to very impactful threats discovered and reported widely after May 2015. In this way the test shows to what extent the product was able to predict how future threats would appear. This "Predictive Advantage" (PA), the advantage that users of the product have against future adversaries, is presented in this report.



Contents

Introduction	04
Executive Summary	05
1. Predictive Advantage by Threat Family	06
2. Predictive Advantage by Individual Campaign	07
3. Legitimate Software Handling	09
4. Conclusions	10
Appendix A: FAQs	11
Appendix B: Sample Validation	12
Appendix C: Product Versions	12

Document version 1.0. Written 28th March 2018



Simon Edwards

Director

WEBSITE www.SELabs.uk

TWITTER @SELabsUK

EMAIL info@SELabs.uk

FACEBOOK www.facebook.com/selabsuk

BLOG blog.selabs.uk

PHONE 0203 875 5000

POST ONE Croydon, London, CR0 0XT

MANAGEMENT

Operations Director Marc Briggs

Office Manager Magdalena Jurenko

Technical Lead Stefan Dumitrascu

TESTING TEAM

Thomas Bean

Dimitar Dobrev

Liam Fisher

Gia Gorbald

Pooja Jain

Ivan Merazchiev

Jon Thompson

Jake Warren

Stephen Withey

IT SUPPORT

Danny King-Smith

Chris Short

PUBLICATION

Steve Haines

Colin Mackleworth

SE Labs is BS EN ISO 9001 : 2015 certified for The Provision of IT Security Product Testing.

SE Labs Ltd is a member of the Anti-Malware Testing Standards Organization (AMTSO)

Introduction

A common criticism of computer security products is that they can only protect against known threats. When new attacks are detected and analysed security companies produce updates based on this new knowledge, which can then be applied to endpoint, network and cloud security software and services.

But in the time between detection of the attack and application of the corresponding updates, systems are vulnerable to compromise. Almost by definition at least one victim, the so-called 'patient zero', has to experience the threat before new protection systems can be deployed. While the rest of us benefit from patient zero's misfortune, patient zero has potentially suffered catastrophic damage to its operations.

MINORITY REPORT

Security companies have, for some years, developed advanced detection systems, often labelled as using 'AI', 'machine learning' or some other technical-sounding term. The basic idea is that past threats are analysed in deep ways to identify what future threats might look like. Ideally the result will be a product that can detect potentially bad files or behaviour before the attack is successful.

It is possible to test claims of this type of predictive capability by taking an old version of a product, denying it the ability to update or query cloud services, and then exposing it to threats that were created, detected and analysed months or even years after its own creation. It's the equivalent of sending an old product forward in time and seeing how well it works with future threats.

This is exactly what we did in this test. Using **CylancePROTECT's** AI model from May 2015 we collected serious threats dating from February 2016 all the way through to November 2017.

Such threats included WannaCry, a mid-2017 ransomware-based attack that was spread using the NSA's EternalBlue exploit; Petya, a ransomware attack from early 2016; and GhostAdmin, malware from 2017 capable of taking remote control of victim systems and exfiltrating data.

These results demonstrate that **CylancePROTECT** users would have been safe from the zero-day attack types used in the test even if they had not updated their software for two years and nine months.

Executive Summary

The product is scored according to how far into the future its protection is seen to reach. For example, if it protected against a threat that was created one year after the product was built, then it would have a predictive advantage of 12 months.

Malware campaigns can run over a period of time, with those in control making changes to the malware to add features or evade detection. For this reason we used different variants for each 'family' of attack. For example, we used five different versions of the Cerber ransomware attack, with samples dating from December 2016 through to February 2018.

CylancePROTECT's Predictive Advantage (PA) varied, depending on the threat. It ranged from

11 months up to 33 months, with an average PA of 25 months. In other words, in some cases it was able to recognise and protect against threats that would not appear in real life for up to two years and nine months into the future. Generally speaking it was effective, without updates, against threats just over two years into the future.

While it is good practice to keep security products fully updated, in many cases keeping endpoint security products continuously up to date is challenging. The purpose of this test is to examine how effective past AI models could be against newer threats. For this reason a version of **CylancePROTECT** from early 2015 was used against threats from 2016, 2017 and 2018.

1. Predictive Advantage by Threat Family

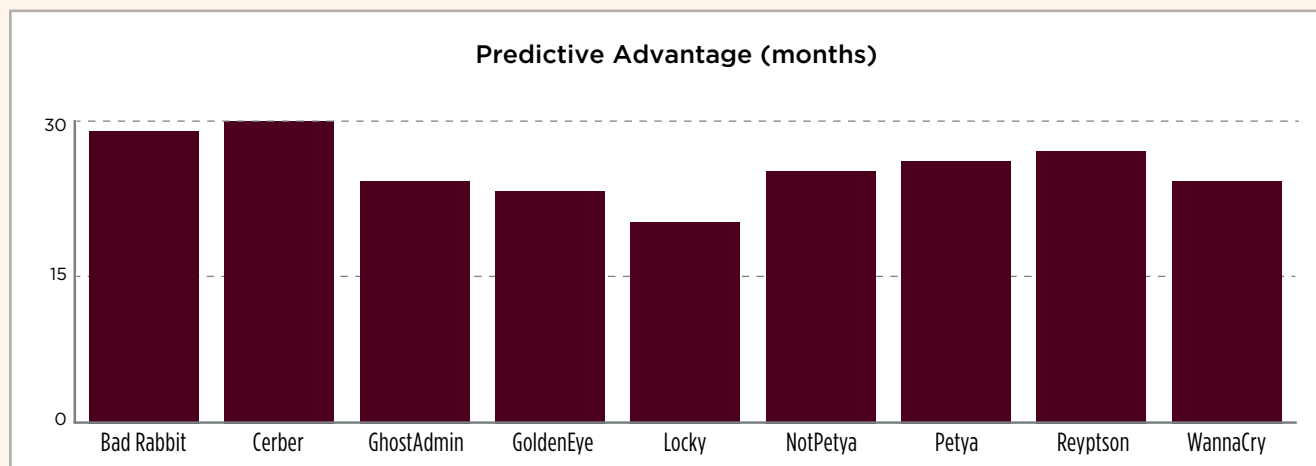
Predictive Advantage (PA) is the time difference between the creation of the model and the first time a threat is seen by victims and security companies protecting those victims.

The model represented in this test was created in May 2015. This is the same model as that deployed in the real world with **CylancePROTECT's** agent, version 1300.

We exposed the model to a range of threats. These comprised nine different 'families' that featured in well-publicised campaigns. Each family set contains five variants as found in the wild.

The graph below shows the average PA value for each threat family. The higher the number, the greater the distance in time from the model's creation date to the first known detection of that specific set of files. Higher PA values are more impressive, as they show the model's ability to predict threats further into the future.

Predictive Advantage by Threat Family	
Threat Family	Predictive Advantage (months)
Bad Rabbit	29
Cerber	30
GhostAdmin	24
GoldenEye	23
Locky	20
NotPetya	25
Petya	26
Reyptson	27
WannaCry	24




2. Predictive Advantage by Individual Campaign

Malware campaigns can run over a period of time, with those in control making changes to the malware to add features or evade detection. For this reason we used different variants for each 'family' of attack. Variants within one family group may appear in the real world at different times as a campaign develops. For example, the GoldenEye

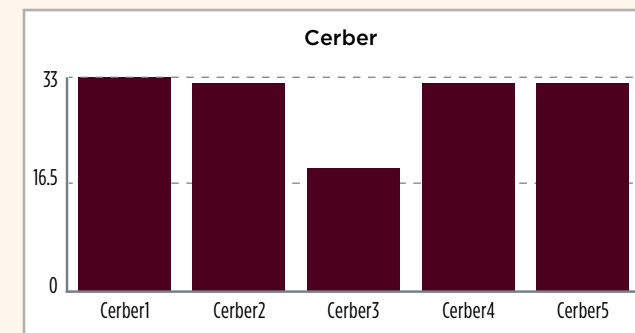
samples range from December 2016 through May 2017 until July 2017.

The graphs below shows the different PA values for the individual threats, which are grouped into their own families.

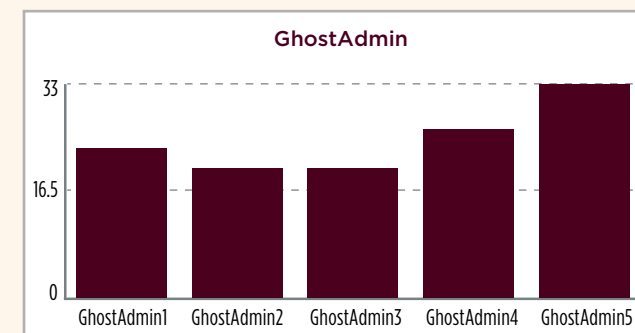
CAMPAIGN: Bad Rabbit	
Threat Variant	Predictive Advantage (months)
Bad Rabbit1	29
Bad Rabbit2	29
Bad Rabbit3	29
Bad Rabbit4	29
Bad Rabbit5	29



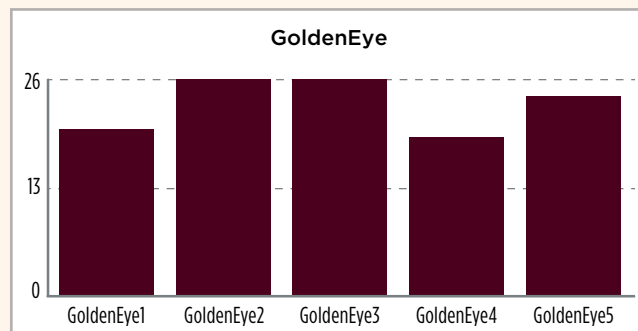
CAMPAIGN: Cerber	
Threat Variant	Predictive Advantage (months)
Cerber1	33
Cerber2	32
Cerber3	19
Cerber4	32
Cerber5	32



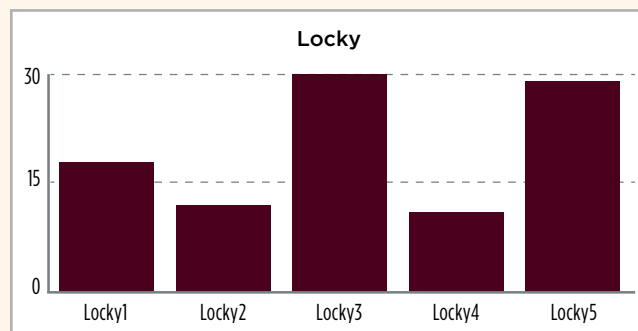
CAMPAIGN: GhostAdmin	
Threat Variant	Predictive Advantage (months)
GhostAdmin1	23
GhostAdmin2	20
GhostAdmin3	20
GhostAdmin4	26
GhostAdmin5	33



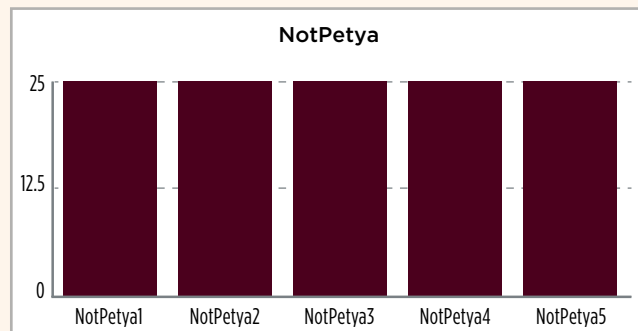
CAMPAIGN: GoldenEye	
Threat Variant	Predictive Advantage (months)
GoldenEye1	20
GoldenEye2	26
GoldenEye3	26
GoldenEye4	19
GoldenEye5	24



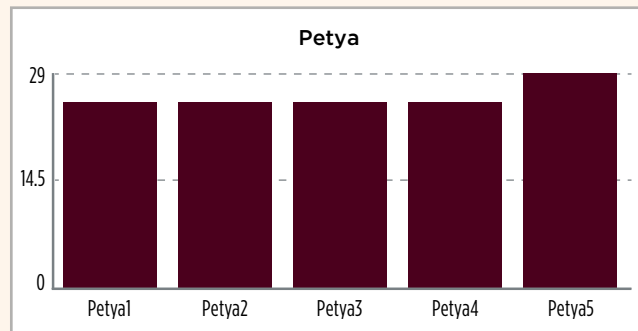
CAMPAIGN: Locky	
Threat Variant	Predictive Advantage (months)
Locky1	18
Locky2	12
Locky3	30
Locky4	11
Locky5	29



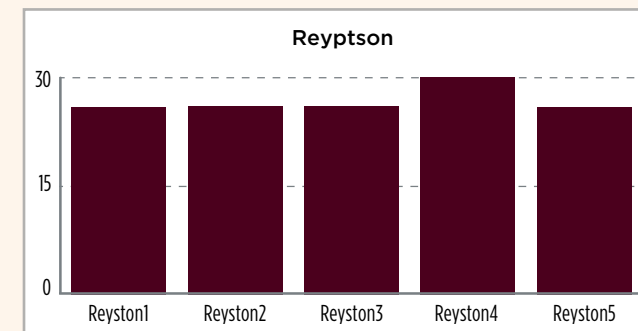
CAMPAIGN: NotPetya	
Threat Variant	Predictive Advantage (months)
NotPetya1	25
NotPetya2	25
NotPetya3	25
NotPetya4	25
NotPetya5	25



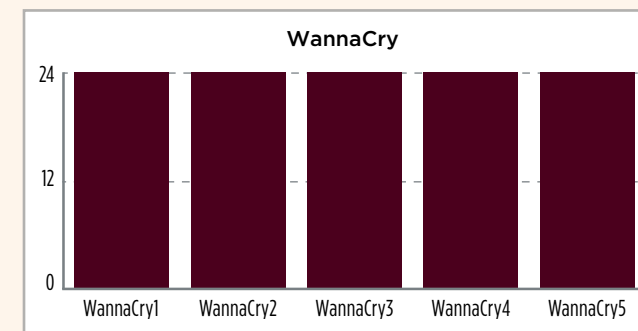
CAMPAIGN: Petya	
Threat Variant	Predictive Advantage (months)
Petya1	25
petya2	25
Petya3	25
Petya4	25
Petya5	29



CAMPAIGN: Reypson	
Threat Variant	Predictive Advantage (months)
Reypson1	26
Reypson2	26
Reypson3	26
Reypson4	30
Reypson5	26



CAMPAIGN: WannaCry	
Threat Variant	Predictive Advantage (months)
WannaCry1	24
WannaCry2	24
WannaCry3	24
WannaCry4	24
WannaCry5	24



3. Legitimate Software Handling

It is necessary, when testing a security product's ability to handle threats, to also measure how it handles legitimate code. Failure to do so means that a product that blocks both good and bad effectively will win a test but cause extreme disruption in the real world.

In this test we measured any incorrect classifications of files already present on the system and of products and websites downloaded during the time of testing. We downloaded 50

popular business and consumer applications, and visited 50 highly popular websites, according to Alexa's index, and found only one sub-optimum case. In this case a spreadsheet viewing utility was quarantined. Subsequently we discovered that this utility was bundled with potentially unwanted code so users may well be advised not to install it.

As such, there were no 'false positives' and no other types of sub-optimum handling of legitimate files.

4. Conclusions

This test was designed to examine **Cylance's** claim that the Artificial Intelligence (AI) technology at the heart of its endpoint protection product is self-contained, in terms of being effective without relying on regular updates or cloud queries. It was also intended to determine whether or not an AI model created some months and even years in the past could identify and handle threats that subsequently attacked systems on the internet.

Predictive Advantage (PA) is the time difference between the creation of the model and the first time a threat is seen by victims and security companies protecting those victims.

Out of 45 threats, 43 were detected and prevented from compromising the system with an average PA of 25 months. The threats used in the test were

discovered in the wild at dates ranging from 11 months to two years and nine months (33 months) after the creation of the AI model.

Not only does the data demonstrate that **CylancePROTECT** (agent v1300, model May 2015) was capable of preventing threats that did not exist at the time the AI model was 'trained', but it provides an insight into how far ahead in time it could be effective without new knowledge. In practical terms, this indicates that regular updates to the product are not always needed, although we would expect **Cylance** to develop and deploy newly-trained models over time, simply because product development is an ongoing process and machine learning continues to take into account new threats to predict future ones.

Appendix A: FAQs

A full [methodology](#) for this test is available from our [website](#).

- The test was sponsored by **Cylance**.
- The artificial intelligence models used in his test were chosen and provided by **Cylance**.
- The test was conducted between 28th January 2018 and 24th March 2018.
- The test was conducted without internet or other access to back-end systems.
- Threats and legitimate applications were independently located and verified by SE Labs.
- Malicious and legitimate data was provided to **Cylance** once the full test was completed.
- SE Labs conducted this test using virtual machines.

Q Did SE Labs provide Cylance with access to the threats used before the test started?

A No, threats were identified, collected and verified before testing, and only then made available to **Cylance**.

Q Is the data shown in this report the full data from the test, or a selection of

A The full set of data collected during testing of the featured model is represented in this report. We have not excluded any sub-optimum results.

Q Did Cylance have an opportunity to dispute any sub-optimum results?

A Yes, after every test we give partner vendors an opportunity to discuss and dispute results.

Q So did Cylance persuade you to drop any sub-optimum results?

A No, all results gathered during testing are represented in this report. We did not drop any sub-optimum results for any reason.

Q Isn't this kind of offline testing unrealistic, given that most systems are connected to the internet in the real world?

A This report's conclusions show that the model under test would have been able to protect against theoretical threats that subsequently became real. It demonstrates the model's potential against unknown threats. Disconnecting from the internet was necessary to show that **Cylance** was not enhancing its product's abilities using online updates.

Appendix B: Sample Validation

To ensure the integrity of the results we validated that the threats were real, matched the family descriptions used, were fully functional and first known to be active at the dates given.

The validation process involved exposing unprotected systems to the threats to observe successful attacks; cross-referencing code samples with threats as identified by online malware repositories and third-party anti-malware scanners; while checking the 'first-seen' dates stored in various third-party malware analysis services including, but not limited to VirusTotal.

Samples of the threats were collected by SE Labs independently of any anti-malware vendor and were not downloaded from VirusTotal.

The legitimate applications used to verify that the product would not simply block everything it encountered were validated to be malware-free, and identified as being highly popular downloads.

In addition, any false positive cases occurring as a result of the product interacting with the operating system's file were noted.

Finally, the product's AI model was also validated by checking its compilation date. The models were provided by **Cylance** and verified by SE Labs' review of their digital signature's timestamps.

The model for which we present results in this report was built in May 2015, while the threats used appeared in the wild throughout 2016 to 2018. We also verified that the ability to detect these same threats did not disappear in versions of the model that had been trained later. We did this by re-testing all attacks against models dated at October 2015, June 2016 and April 2017, obtaining similar results.

Appendix C: Product Versions

CylancePROTECT has been released in different agent versions over the last few years, and each agent may have a different AI model included. This is the summary of the agents we tested.

Product Versions		
Agent Version	Agent Date	Model Date
1300	May 2015	May 2015
1320	October 2015	October 2015
1380	June 2016	June 2016
1460	February 2018	April 2017

Disclaimer

1. The information contained in this report is subject to change and revision by SE Labs without notice.
2. SE Labs is under no obligation to update this report at any time.
3. SE Labs believes that the information contained within this report is accurate and reliable at the time of its publication, which can be found at the bottom of the contents page, but SE Labs does not guarantee this in any way.
4. All use of and any reliance on this report, or any information contained within this report, is solely at your own risk. SE Labs shall not be liable or responsible for any loss of profit (whether incurred directly or indirectly), any loss of goodwill or business reputation, any loss of data suffered, pure economic loss, cost of procurement of substitute goods or services, or other intangible loss, or any indirect, incidental, special or consequential loss, costs, damages, charges or expenses or exemplary damages arising his report in any way whatsoever.
5. The contents of this report does not constitute a recommendation, guarantee, endorsement or otherwise of any of the products listed, mentioned or tested.
6. The testing and subsequent results do not guarantee that there are no errors in the products, or that you will achieve the same or similar results. SE Labs does not guarantee in any way that the products will meet your expectations, requirements, specifications or needs.
7. Any trade marks, trade names, logos or images used in this report are the trade marks, trade names, logos or images of their respective owners.
8. The contents of this report are provided on an "AS IS" basis and accordingly SE Labs does not make any express or implied warranty or representation concerning its accuracy or completeness.