SE Labs

INTELLIGENCE-LED TESTING

DCD

RANSOMWARE

Data Corruption Detection

NetApp

ONTAP Autonomous Ransomware Protection with AI

June 2024

SE Labs ® tested **NetApp ONTAP Autonomous Ransomware Protection with AI** against a range of ransomware attacks designed to extort victims. These attacks were realistic, using the same tactics and techniques as those used against victims in recent months.

Target systems, protected by **NetApp ONTAP Autonomous Ransomware Protection with AI**, were attacked by testers acting in the same way as we observe ransomware groups to behave.

Attacks were initiated from the start of the attack chain, using phishing email links and attachments, as just two examples. Each attack was run from the very start to its obvious conclusion, which means attempting to steal, encrypt and destroy sensitive data on the target systems.

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
Jarred Earlington
Gia Gorbold
Anila Johny
Erica Marotta
Jeremiah Morgan
Julian Owusu-Abrokwa
Joseph Pike
Georgios Sakatzidis
Dimitrios Tsarouchas
Stephen Withey

Publication and Marketing
Colin Mackleworth
Sara Claridge
Janice Sheridan

IT Support
Danny King-Smith
Chris Short

Website selabs.uk
Email info@SELabs.uk
LinkedIn www.linkedin.com/company/se-labs/
Blog blog.selabs.uk
Post SE Labs Ltd,
55A High Street, Wimbledon, SW19 5BA, UK

SE Labs is ISO/IEC 27001 : 2013 certified and
BS EN ISO 9001 : 2015 certified for The Provision
of IT Security Product Testing.

SE Labs is a member of the Microsoft Virus Initiative (MVI);
the Anti-Malware Testing Standards Organization (AMTSO);
the Association of anti Virus Asia Researchers (AVAR);
and NetSecOPEN.

Contents

Document version 1.0 Written 1st June 2024

Introduction

Ransomware on the Network

Results from the largest public ransomware test

Ransomware is the most visible, most easily understood cyber threat affecting businesses today. Paralysed computer systems mean stalled business and loss of earnings. On top of that, a ransom demand provides a clear, countable value to a threat. A demand for "one million dollars!" is easier to quantify than the possible leak of intellectual property to a competitor.

One reason why ransomware is so 'popular' is that the attackers don't have to produce their own. They outsource the production of ransomware to others, who provide Ransomware as a Service (Raas). Attackers then usually trick targets into running it, or at least into providing a route for the attackers to run it for them. Artificial intelligence systems make the creation of such social engineering attacks easier, cheaper and more effective than ever before.

Doesn't Endpoint Protection Solve Ransomware?

One approach to data protection is to secure endpoint systems and assume that all is well. The idea is that if the endpoints are 100% secure then the data is also safe. The problem with that assumption is that it relies on the business' data being only on those endpoints, and that those endpoints are completely invulnerable. Neither of those requirements are likely to be true.

Business data on the network is vulnerable to ransomware attacks, even if the endpoints that access the data are fully locked down. This is the problem that **NetApp ONTAP Autonomous Ransomware Protection with AI** seeks to solve. It monitors the integrity of the business' data on the network, aiming to detect unwanted changes made by ransomware.

For example, a file server contains the bulk of the target's files. The attacker obtains access remotely, using stolen credentials. (This is a very common type of breach.) The attacker then deploys ransomware and the files are encrypted. At no stage does the endpoint security software running on the company's laptops get a chance to stop this attack. Those endpoints might still work, but their access to the bulk of the important data is now stopped, which means business is halted.

Ransomware Attacks on the Network

In this report we have assembled a wide distribution of known ransomware malware and added variations designed to evade detection. We've listed the ransomware families used in **Threat Intelligence** on page 9. Each of these ransomware attacks targeted data that was monitored by NetApp's solution.

If it can detect the changes made by known version of each of these files, all well and good. But if it can also detect changes made by each of the ransomware's variations then we can conclude that the detection available is more proactive than simply reacting to yesterday's unlucky victims.

If you spot a detail in this report that you don't understand, or would like to discuss, please **contact us**. 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 **LinkedIn**.

Executive Summary

We tested **NetApp ONTAP Autonomous Ransomware Protection with AI** against thousands of known and unknown variations of ransomware attacks. Its job was to monitor business data and detect any unwanted changes that the attacks made to those files, such as encrypting files.

We examined its abilities to detect the behaviour of:

- Known ransomware from 15 different attacker groups
- New, similar variants of ransomware from the same groups
- Legitimate operations on the network

NetApp ONTAP Autonomous Ransomware Protection with AI performed exceptionally well, providing 99% detection coverage of the ransomware's destructive behaviour. It was completely effective at handling legitimate operations, making no mistakes.

Executive Summary

Product Tested	Detection Accuracy (%)	Legitimate Accuracy Rating (%)
NetApp ONTAP Autonomous Ransomware Protection with AI	99%	100%

The Detection rating shows how effective the product was at detecting the ransomwares attacks.

Data Corruption Detection (Ransomware) Award

The following product wins the SE Labs award:

NetApp
ONTAP Autonomous Ransomware Protection with AI

1. How we Tested

Testers can't assume that products will work a certain way, so running a realistic advanced security test means setting up real networks and hacking them in the same way that real adversaries behave.

In the diagram on the right you will see an example network that contains workstations, some basic infrastructure such as file servers.

The techniques that we choose for each test case are largely dictated by the real-world behaviour of online criminals. We observe their tactics and replicate what they do in this test. To see more details about how the specific attackers behaved, and how we copied them, see **Hackers vs. Targets** on page 9.

Test Network Example

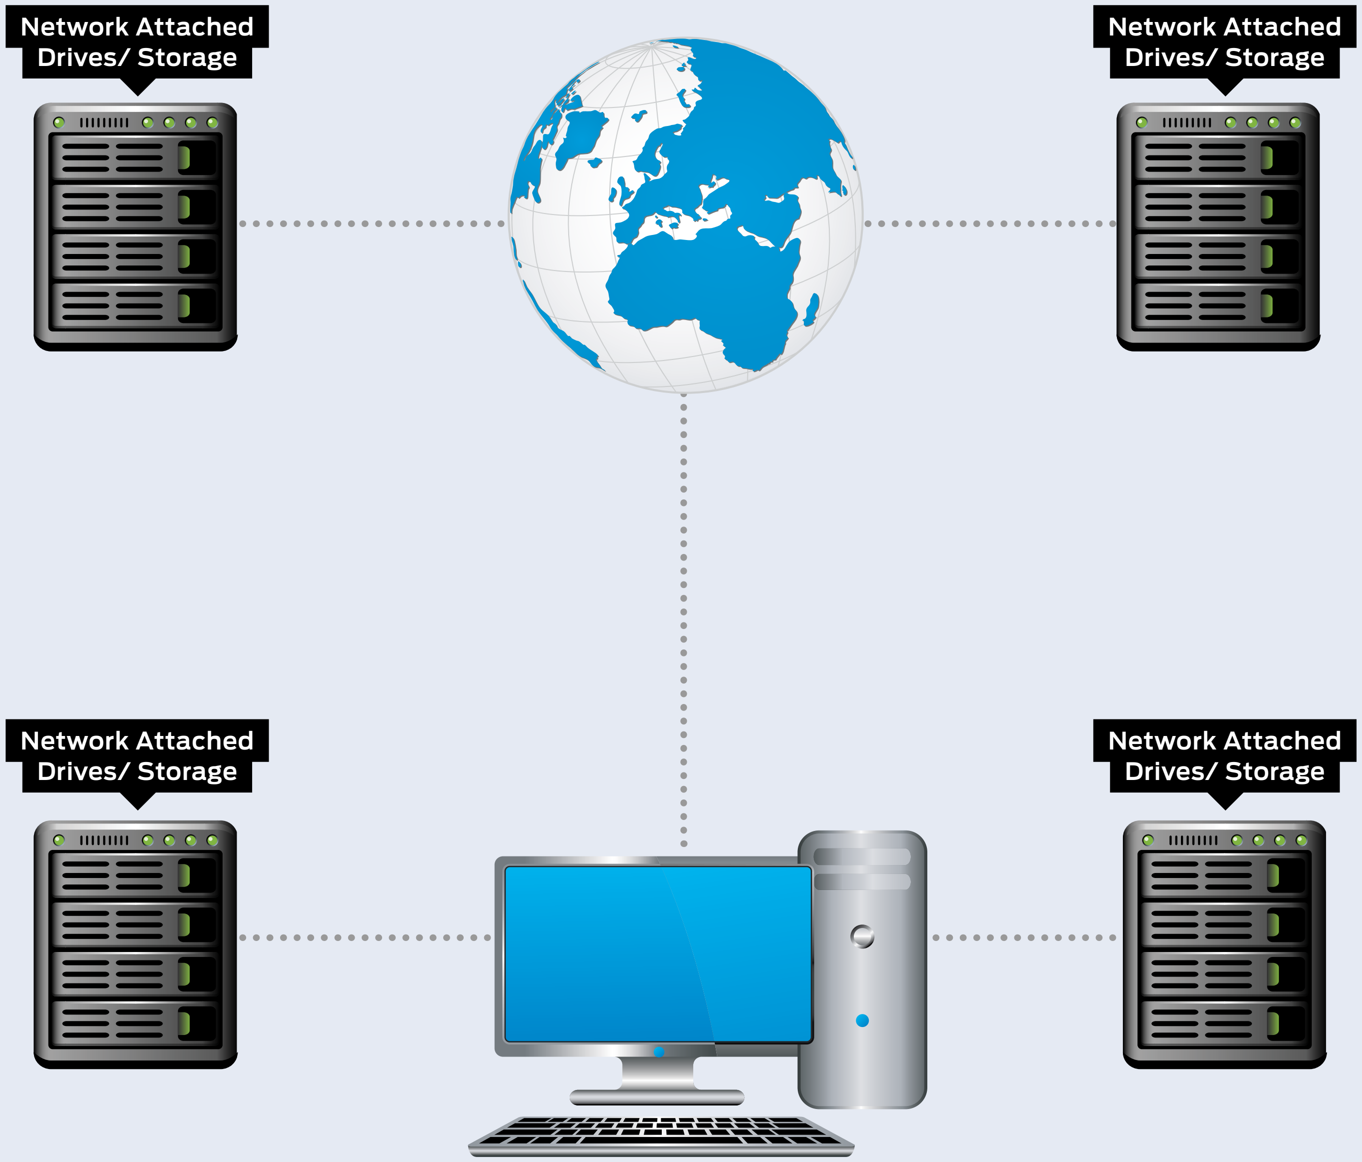
Network Attached Drives/ Storage
Network Attached Drives/ Storage
Network Attached Drives/ Storage
Network Attached Drives/ Storage

This example of a test network shows one possible topology and ways in which enterprises and criminals deploy resources.

Threat Responses

Full Attack Chain: Testing Every Layer of Detection and Detection

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

We measure how a product responds to the first stages of the attack with a detection and/ or detection 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 contains 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 (steps 5-7).

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.

Attack Chain Stages

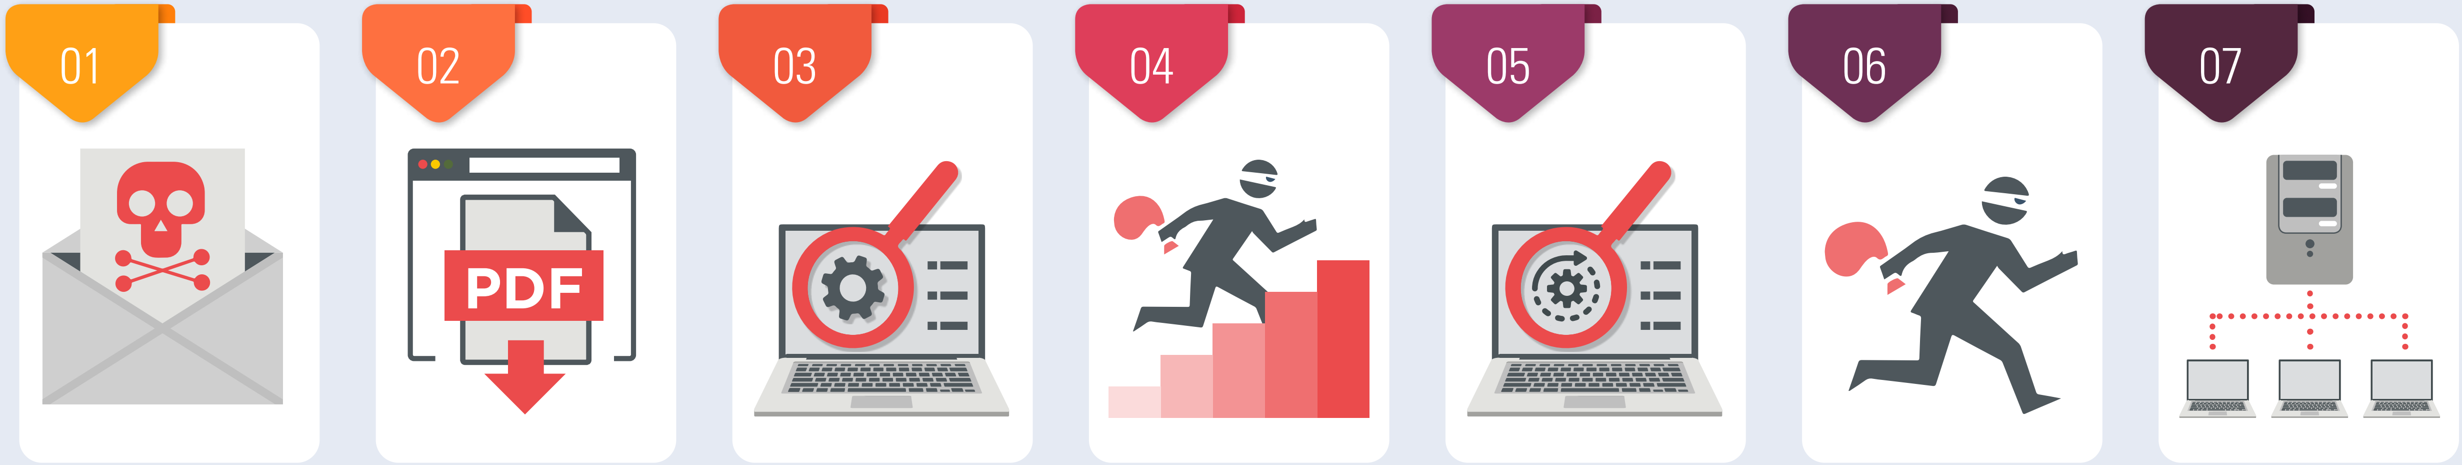

Figure 1. A typical attack starts with an initial contact and progresses through various stages, including reconnaissance, stealing data and causing damage.

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 for an attack to run in a different order with, for example, the attacker attempting to connect to other systems without needing to escalate privileges. However, it is common for password theft (see step 5) to occur before using stolen credentials to move further through the network.

It is also 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.

Attack Chain: How Hackers Progress

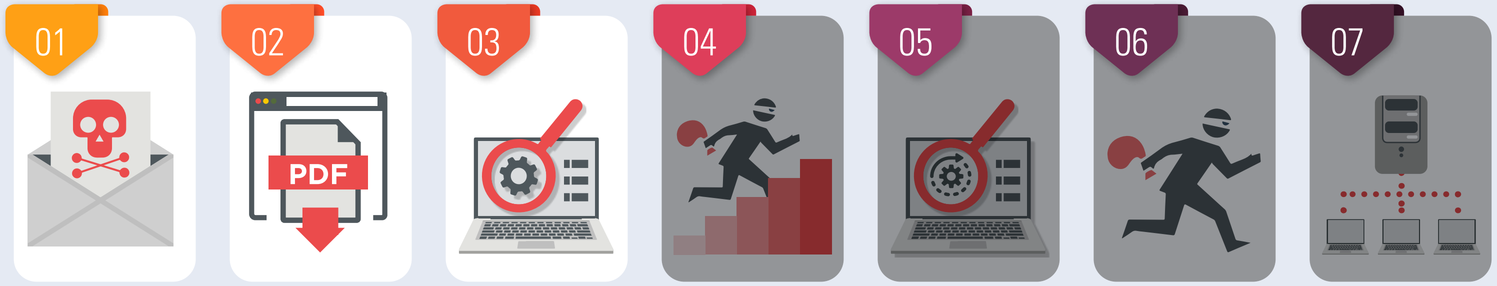

Figure 2. This attack was initially successful but only able to progress as far as the reconnaissance phase.

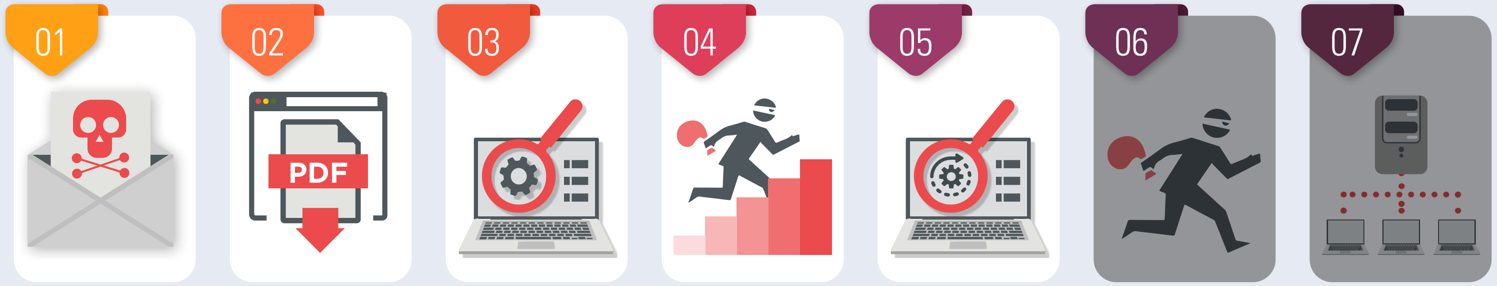

Figure 3. A more successful attack manages to steal passwords but wholesale data theft and destruction was blocked.

Threat Intelligence

When testing services against targeted attacks it is important to ensure that the attacks used are relevant. Anyone can run an attack randomly against someone else. It is the security vendor's challenge to identify common attack types and to protect against them. As testers, we need to generate threats that in some way relate to the real world.

All of the attacks used in this test are valid ways to compromise an organisation. Without any security in place, all would succeed in attacking the target. Outcomes would include systems infected with ransomware, remote access to networks and data theft.

But we didn't just sit down and brainstorm how we would attack different companies. Instead we used current threat intelligence to look at what the bad guys have been doing over the last few years and copied them quite closely. This way we can test the services' abilities to handle similar threats to those faced by global governments, financial institutions and national infrastructure.

The graphic on this page shows a summary of the attack groups that inspired the targeted attacks used in this test. If a service was able to detect these then there's a good chance they are on track to detect similar attacks in the real world. If they fail, then you might take their bold marketing claims about defeating hackers with a pinch of salt.

Hackers vs. Targets

Attacker/ APT Group	Method	Target	Details
AvosLocker	✉	$	This malware is hired out as Ransomware as a Service (RaaS) and so is used against a wide variety of targets.
BlackMatter	🕵	🌉	This threat uses LDAP and SMB file-sharing protocols to spread via Active Directory to all hosts on a network.
Cerber	🐴 ✉	⚙💀	This malware is hired out as Ransomware as a Service (RaaS) and so is used against a wide variety of targets.
DarkSide	✉ ✉	⊕ 🎓	This ransomware threat is developed and distributed as a Ransomware as a Service (RaaS) for attackers to deploy independently.
DeathRansom	📊	⚙💀	Shared through Malware as a Service operations, this ransomware avoids encrypting Windows system files to ensure the target system continues to operate. It targets both local and network resources.
GandCrab	✉ ✉	⚙💀	This malware is hired out as Ransomware as a Service (RaaS) and so is used against a wide variety of targets.
Lockbit	💵	$	An RaaS threat used across a variety of industries and continues to be prolific in 2024.
Maoloa	✉	⚙💀	This malware is hired out as Ransomware as a Service (RaaS) and so is used against a wide variety of targets.
Netwalker	✉ ✉	⚙💀	Written in PowerShell, this ransomware's payloads are executed by VBS scripts embedded in Microsoft Office documents.
Phobos	RDP ✊	🌉	Phobos operates as Ransomware as a Service (RaaS) and goes beyond simple data encryption, leading to disruption of critical systems.
PYSA	✉ ✉	🎓	Targeting US and UK educational systems, more than 12 have reported themselves as having been compromised by this threat.
Ragnar Locker	RDP ✊	⚙💀	Highly customised, this ransomware threat is operated by a group known to leak stolen data.
Ryuk	✉	⚙💀	Focussed on businesses, this group is known to leak stolen data even if the ransom is paid.
TeslaCrypt	📄	⚙💀	A so-called 'file-less' threat that gains access to targets and then deploys ransomware.
WastedLocker	ZIP	🏛	This group uses decoy files and PowerShell to spread through networks, deploying the Cobalt Strike penetration testing tool.

Key

🏛 Banking and ATMs	🎓 Education	$ Financial	⚙💀 Generic RaaS	⊕ Healthcare	🌉 Infastructure

2. Detection Scores (Ransomware Network Attacks)

This table shows the overall level of detection.

Detection Scores		
Product	**Detection Score**	**Detection Score (%)**
NetApp ONTAP Autonomous Ransomware Protection with AI	3,585	99%

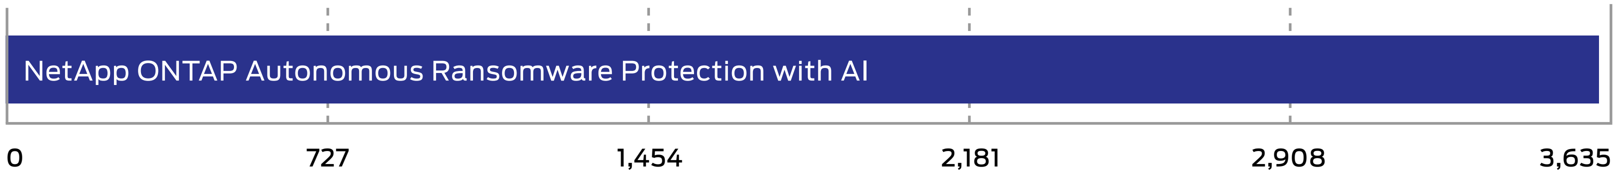

Detection Scores are a simple count of how many times a product protected the system.

3. Legitimate Software Ratings

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

We also take into account the prevalence (popularity) of the file types used in this part of the test, applying stricter penalties for when products misclassify very popular file types.

Legitimate Software Ratings		
Product	**Legitimate Accuracy Rating**	**Legitimate Accuracy (%)**
NetApp ONTAP Autonomous Ransomware Protection with AI	1,957	100%

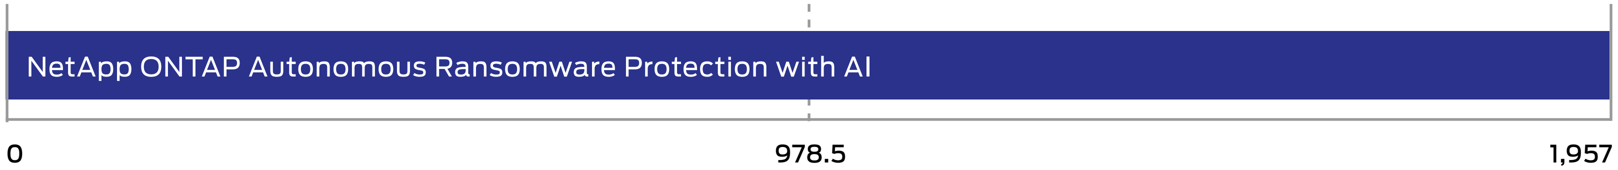

Legitimate Software Ratings can indicate how well a vendor has tuned its detection engine.

4. Conclusions

This report looks at how effectively a security product can protect against a wide range of ransomware attacks. It investigates the product's capabilities in tracking the changes made to business data by ransomware.

Ransomware Network Attacks

In this test we ran full, advanced hacking attacks against the target system, which stored a large number of business files. The ransomware attacks were given the opportunity to run in a completely realistic way, encrypting files in a wide variety of ways.

We wanted to assess how well **NetApp ONTAP Autonomous Ransomware Protection with AI** could track these changes to the files and alert administrators of the extent of the attacks.

The test cases comprised multiple variations from 15 different ransomware attack 'families', which are listed on page 9 (**Threat Intelligence**). This allows the security product to demonstrate its ability not only to detect the behaviour of known threats but similar, new ones. In turn this provides an indication of its ability to work on tomorrow's threats, rather than historical ones.

In addition to detecting threats, the product should not incorrectly claim attacks when none occur. In this test **NetApp ONTAP Autonomous Ransomware Protection with AI** was completely accurate. **NetApp ONTAP Autonomous Ransomware Protection with AI** detected the vast majority of the threats, in both original and variant form.

Overall, **NetApp ONTAP Autonomous Ransomware Protection with AI** provided 99% detection of the advanced ransomware attacks. It detected 3,585 of the 3,635 attacks, earning a AAA award.

If you spot a detail in this report that you don't understand, or would like to discuss, please **contact us**. 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 **LinkedIn**.

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 test was conducted between 3rd May to 22nd May 2024.
- The product was configured according to its vendor's recommendations.
- Targeted attacks were selected and verified by SE Labs.
- Malicious and legitimate data was provided to partner organisations once the test was complete.
- SE Labs conducted this endpoint security testing on physical PCs, not virtual machines.

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 **We are a customer considering buying or changing our endpoint detection and/ or endpoint detection and response (EDR) product. Can you help?**

A Yes, we frequently run private testing for organisations that are considering changing their security products. Please contact us at **info@selabs.uk** for more information.

Appendix C: Product Versions

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

Product Version			
Vendor	Product	Build Version (start)	Build Version (end)
NetApp	ONTAP Autonomous Ransomware Protection with AI	9.15.1*	9.15.1

*Available for technical preview

SE Labs

SE Labs Report 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.