The Dwell Time Issue in Cybersecurity

Whitepaper

Crytica Security
7655 Town Square Way, Suite 212
Reno, NV 89523

Email: info@cryticasecurity.com

The Dwell Time Issue

Contents

1. INTRODUCTION 3
2. THE DWELL TIME PROBLEM 3
3. THE CONSEQUENCES OF HIGH DWELL TIMES 4
4. WHY ARE THE DWELL TIMES SO HIGH? 6
5. HOW CAN DWELL TIMES BE REDUCED? 10
6. NEXT STEPS 11

1. INTRODUCTION

After years in the shadows, Cybersecurity is now front and center. The on-going and disastrous epidemic of ransomware attacks, combined with a multitude of marquee data breaches, have commanded the attention of the mainstream media and the general public. Again and again, with almost monotonous repetition, large institutions, as well as small businesses and private individuals are exploited by seemingly unstoppable cybercriminals. There is an unpalatable truth that cybersecurity professionals have known for years, one that the public is just beginning to understand. It is that: the question is not "if" one will get attacked, but "when"? And yet, with all the cybersecurity options available and myriad tools available, why is cybercrime continuing to grow at such an exponential rate?

A basic tenet of all security states that "If you can't detect an attack, if you can't detect an intrusion, you can't protect against it." If malware can be injected into a system and remain undetected for months at a time, and often is detected only when the attack itself is launched, then the cybersecurity protections in place have failed to perform a critical role: timely and actionable malware detection.

Unfortunately, all the independent studies of "dwell time", that is, "the time between malware injection and malware detection", demonstrate convincingly that the current state-of-the-art malware detection measures are not effective. When average dwell times are measured in months rather than minutes, the inescapable answer to the question of why cybercrime is continuing to grow exponentially, is that: the current detection systems cannot detect, and therefore they cannot protect.

To stem the tide of this onslaught, Crytica Security has developed a new malware detection system, one with "Zero-Day Detection" capability, that can detect even zero-day infections (malware never previously identified) in minutes rather than months. In doing so, Crytica is reversing the trend of exponentially increasing cybercrime, and arming the cybersecurity professionals with the ability to detect, so that now they can truly protect.

2. THE DWELL TIME PROBLEM

Dwell Time is a metric that measures the elapsed time between malware infection (injection) and malware detection. Obviously, the longer the dwell time, the longer the malware is afforded the time it requires to create the insidious foundations for its ultimate attack. During this time, the malware can, among things:

- Collect intelligence about: the environment invaded, its defenses, its topology, its network traffic, and its user community.
- Proliferate and replicate itself throughout the entire infrastructure to be attacked.
- Morph itself so that it remains invisible to all the deployed malware detection tools.

Once the malware is ready to launch, it will, by then, have already spread itself to all the devices it intends to attack, and it will have determined the most effect attack vectors for that attack.

In the case of a ransomware attack, it will be able to lock up all the intended targets, virtually simultaneously. In the case of a data exfiltration attack, it will be able to start exfiltrating data from all the intended targets, again virtually simultaneously. And in the case of human-engineering and/or spear phishing attacks, these too can be launched with a timing and an efficacy that is optimized for the environment.

Confirmation that dwell times are unconscionably high[1] is detailed in the following graph (Figure 1) from the IBM Cost of Data Breach Study 2024 (in which MTTI is "Mean Time To Identify" and MTTC is "Mean Time To Correct").

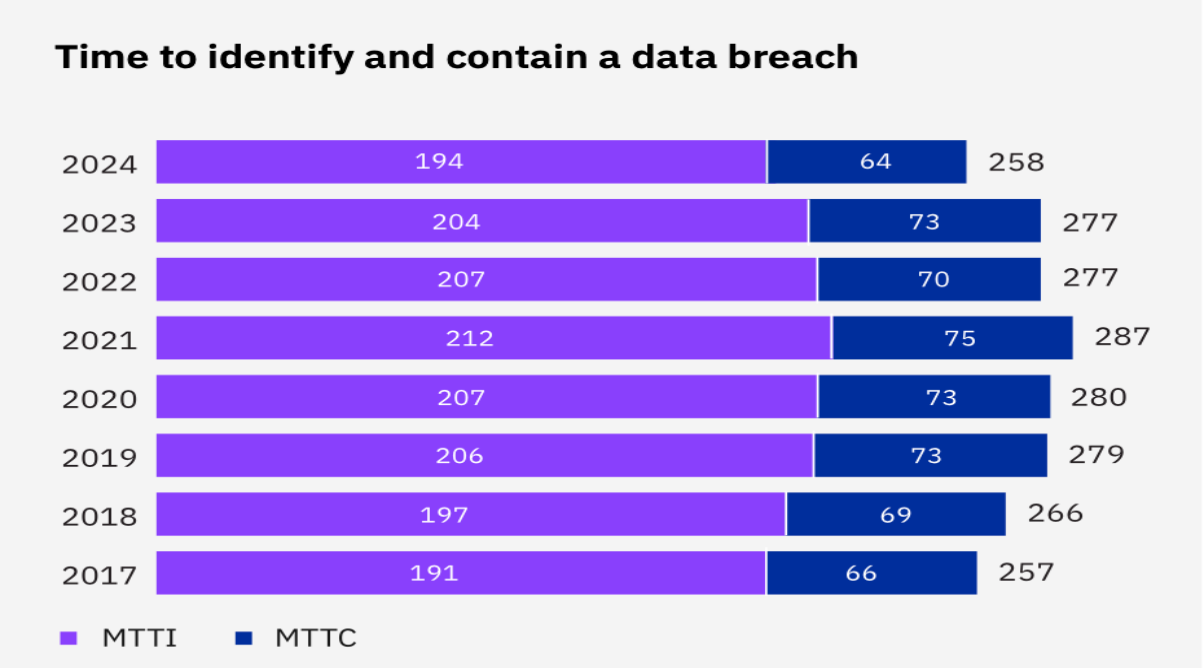

Figure 1 – Unconscionably High Average Dwell Times[2]

3. THE CONSEQUENCES OF HIGH DWELL TIMES

The results of this inability to detect malware infections are staggering. According to Cybersecurity Ventures, this year a new ransomware attack will occur every 11 seconds. The overall annual cost of ransomware attacks is expected to reach $20 billion.[3] Along with the growing number of cyber breaches and adversely impacted businesses, cybersecurity technology and services costs continue to rise annually.

[1] It will be noted that the times have dropped a bit in 2004, even though they are still unconscionably high. This is not necessarily due to improved detection. Any malware that resides undetected for many months is almost always discovered when it launches, and not by the by cybersecurity systems in place. A shorter Dwell Time here indicates malware that is more efficient in its ability to infect and proliferate. It does so faster, and therefore launches sooner.
[2] IBM Corporation. (2024). (publication). Cost of a Data Breach 2024. Armonk, New York.
[3] Morgan, S. (2021, April 27). Cybercrime to cost the world $10.5 trillion annually by 2025. Cybercrime Magazine.

However, even fully deployed, managed security solutions with the most state-of-the-art technology cannot keep pace with the new, ever-evolving threats. That is the inescapable conclusion from the graph in Figure 1 above.

According to IBM, the cost of a data breach is broken down into three components of monetary loss: fines, lost business, and response costs. Since many cybercrimes target the customers' information, the loss of customer confidence can be devastating. Consumers lose trust in businesses that have experienced breaches, and consequently, often take their business elsewhere. Some companies, such as the telemarketing firm The Heritage Company, actually have been forced to shut down after a successful ransomware attack.

Cybercriminals are equal-opportunity attackers. Every industry is at risk for data compromises. Among the top targets, however, are healthcare, financial institutions, pharmaceutical companies, and technology companies.

One very public attack against a financial institution took place in 2017. The credit reporting company Equifax experienced a highly impactful hack. The intruders dwelt for 76 days in Equifax's systems before they were discovered. This dwell time allowed the hackers to move from the initial compromised server to many other components of Equifax's network. According to the Federal Trade Commission, Equifax stated that approximately 150 million people's data were stolen, including social security numbers and credit card information. The US Federal Trade Commission ultimately brought charges against Equifax. Those charges were settled for a fine of approximately $425 million. This was followed by an additional $275 million in fines from various other governing bodies.

Many other big name financial institutions have suffered breaches. These include Visa, Capital One, CNA Financial, and JP Morgan Chase. All these institutions have the financial and technological resources to deploy the best state-of-the-art cyber defenses. And yet, still they were breached.

As the financial industry is no stranger to malware attacks, neither is the retail market. Fifty-six million credit card numbers were stolen from Home Depot. Many of the other high-profile retailers have been the successful targets of cybercriminals, e.g., Target and Nieman Marcus. The Home Depot breach resulted in costs of $13 million in customer compensation, millions more on credit monitoring, increased staff within the help call center, legal battles, and professional services; concomitant with the cost of diminished brand equity and reputation among customers, shareholders, suppliers, and partners.

Healthcare is also a prime target for cybercriminals. It "boasts" the highest cost per breach. According to the above cited IBM report, the average cost of a healthcare breach is $9.23 million. For healthcare companies, the highly confidential nature of private stolen data affords the hackers additional leverage for demanding ransoms. Loss of personal data is a HIPAA violation. Yet according to the HIPAA journal, over forty-four million healthcare records were exposed or stolen in 2021.

There is a clear correlation between dwell time and the damage a malware infection causes. Figure 2, below, graphically illustrates this relationship, one that is also intuitively obvious. The longer the dwell time, the greater the damage and the greater the costs incurred. The important questions are:

- Why are dwell times so high?
- How can dwell times be reduced?

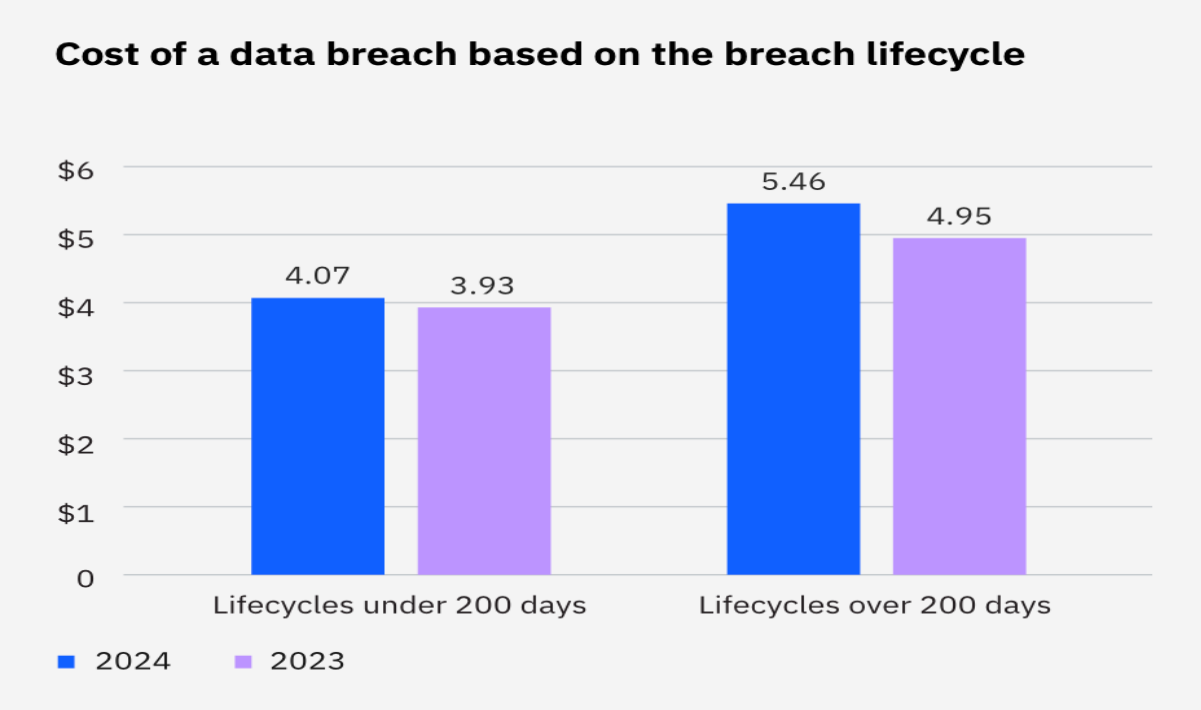

Figure 2 – Cost (in $millions) as a function of Dwell Time[4]

4. WHY ARE THE DWELL TIMES SO HIGH?

There are many reasons why the average dwell times are so high. All can be traced to shortcomings in how a malware infection is detected. Consider some of the primary approaches. There are some underlying common themes. Most either rely upon history, that is, activities/objects previously identified, or they rely upon the detection of malicious behavior, that is, the attack has already launched. They fall loosely into three general classes:

- Virus/Malware Detection
- Attack Detection
- Cyber Threat Intelligence

[4] IBM Corporation. (2024). (publication). Cost of a Data Breach 2024. Armonk, New York.

Virus/Malware Detection

Virus "Signature" matching

A Virus "Signature" is the computer code that constitutes the executable malware. Virus Signature matching reads through every file in an environment and, for each one, searches to see if any of the virus signatures in its database of virus signatures are present in the file.

- Pros: When a match is found, it is an accurate match. There are no "false positives."
- Cons:
  - With new malware constantly being created, by some estimates every 4 seconds, the number of viruses already identified numbers in the billions. Therefore, to search through every file and compare it to billions of known viruses is a long, time-consuming, and resource-hampering process.
  - The number of viruses increases significantly daily. Thus, keeping the database of known viruses up to date is a Sisyphean task. It is impossible to keep current the databases of known viruses.
  - Most sophisticated viruses self-morph, that is they change their own virus signatures over time and/or by location. A morphed virus, although previously identified and associated with a known signature is no longer be identifiable after it morphs.
  - **Virus Signature matching can only detect malware previously seen and identified. It can never detect any new, previously unidentified, malware**.

Virus Signature Pattern Matching

Pattern Matching was developed as a direct response to self-morphing malware. Instead of merely looking for an exact match to a virus signature, Artificial Intelligence (AI) pattern matching algorithms are used to identify computer code that is sufficiently similar to the original known malware as to be possibly a morphed version of it. If Virus Signature matching is a dictionary lookup, then Pattern Matching is a thesaurus lookup.

- Pros: Pattern Matching is able to detect some morphed malware that direct virus signature matching would miss.
- Cons:
  - Pattern Matching is subject to a plague of false positives; so much so that many providers of Pattern Matching detection also employ banks of human experts to filter out as many of the false positives as possible before passing detection alerts along.

- **Virus Signature Pattern Matching can only detect malware previously seen and identified. It can never detect any truly new, previously unidentified, malware.**

Sandboxing

Sandboxing is the practice of putting all new files into a "sandbox" (i.e., a controlled isolated environment) before permitting them access to a live environment. In the controlled environment, each new file is monitored to see if it exhibits any malicious behavior. If it does, the file is flagged and dealt with as malware. If the file does not exhibit malicious behavior, it is permitted to enter the live environment.

- Pros:
  - Sandboxing detects fast-acting malware and prevents it from entering the live environment. The malware is stopped before it can do any damage.
  - Sandboxing can detect even Zero-Day malware, i.e., malware never previously identified and/or documented.
- Cons:
  - Because files cannot be stored indefinitely in a sandbox, slower-acting, time-delayed ("time bombs") cannot be detected using a sandbox.
  - Detection of "malicious" behavior is not an exact science. Thus, false positive alerts and missed true positives can be a problem.

Attack Detection

Behavioral Anomaly Detection

Behavioral Anomaly Detection systems also use AI. By observing an operational environment, they learn the norms of an environment's operational behavior, and especially the human behavior. After the norms have been baselined, the AI system can then detect any statistically significant variations from the baselined normal behavior. However, as all functioning environments are dynamic and evolve over time, the Behavioral Anomaly systems must continue "to learn", to monitor the environment and update the baseline as the behavior of the monitored environments evolves.

- Pros: Behavioral Anomaly Detectors do not rely upon the detection of injected malware to detect incidents. Rather they rely upon norms of behavior. As such, they can detect attacks by even previously unknown, previously unseen (i.e., Zero-Day) malware.
- Cons: Behavioral Anomaly Detectors cannot detect sophisticated malware that also uses AI. For example, they cannot detect attacks that use AI to re-educate the malware detection system. In this situation, the malware's AI can slowly, incrementally, modify the environmental behavior, staying below the threshold

of alert-triggering changes, and thereby "teaching" the AI learning algorithms that certain malicious behavior is part of the environmental baselined norm.

Various other Attack Detection Techniques

There are many attack detection techniques that rely upon monitoring various aspects of system operations. For example, Rootkit detection is used to check and observe the actions that a program is attempting to execute. Based on those actions, the Rootkit detection system determines whether the program is malicious.

- Pros: By detecting malicious behavior as early as possible after the malware launches, the damage the malware causes can be minimized.
- Cons:
  - By waiting until the malicious behavior has actually begun, some of the damage will have already been done.
  - By waiting until the malicious behavior has actually begun, affords the malware time to collect intelligence and proliferate prior to launching a massive attack.
  - Determining, at a system level, what constitutes "malicious behavior" is not clearly and precisely defined. Often heuristics are involved. As with all non-deterministic determinations, the risk of false positives and missed true positives is significant.

Cyber Threat Intelligence

Cyber Threat Intelligence has been defined as "Intelligence analysis on threats in the cyber domain."[5] As with almost all intelligence, Cyber Threat Intelligence comes from many sources. These include on-going analyses of recent cyber-attacks, evaluations of the socio-political-economic environment from whence and for which attacks originate, and the monitoring of "chatter" on the Dark Web, the messages being exchanged and posted by the cybercriminal community.

- Pros: Global perspectives and multi-dimensional analyses can reveal information about impending attacks. It can often do so with a long enough lead-time to prepare adequate defenses.
- Cons:
  - Intelligence gathering and analysis of any type are very non-quantitative, very qualitative endeavors. The results are usually open to various interpretations. It is by no means an "exact science". False positives can abound, as well as missed true positives.
  - Information garnered from the Dark Web and the cybercriminal community will never include the most serious and most sophisticated

[5] Carnegie Mellon University: "Cyber Intelligence Tradecraft Report" 2019

threats. Those are almost always kept secret and away from the "chatter" until well after the attacks are launched.

General Trends

If one considers the overall general trends of malware detection, despite some very sophisticated and effective tools, the vulnerabilities are manifest:

- Virus Signatures: These tools rely upon history, e.g., malware dictionaries and/or malware patterns. However, history is incapable of detecting new (Zero-Day) malware attacks. As the financial investment community is quick to point out: **"Past performance is no guarantee of future results."**
- Artificial Intelligence: Although AI is a very power tool, it is also a double-edged sword. The cyber criminals have access to the same AI tools as the cybersecurity community; and often the criminals are better funded. Superior AI can always be used to defeat inferior AI. The use of AI for cyber defense, unless aided by other technologies, is ultimately a futile dead-end. Its use becomes an eternal arms-race, a race to ensure that the cybersecurity's AI will always be superior to the AI of the cyber criminals.
- Attack Detection: There can be no doubt that detecting an attack as soon as it is launched is extremely important. Attack detection must be intrinsic to all cybersecurity arsenals. However, attack detection does not prevent nor mitigate long dwell times, nor does it prevent the malware from collecting valuable intelligence and proliferating throughout an environment.
- Cyber Threat Intelligence: Threat Intelligence is always important for any serious and effective security system. However, intelligence can never be standalone. It must always be part of an entire security infrastructure. It requires a foundation of other tools in order to be of value.

5. HOW CAN DWELL TIMES BE REDUCED?

There is only one way dwell times can be reduced, and that is to have rapid, timely, and actionable detection. This mean having the ability to detect:

- Even Zero-Day malware infections.
- Malware infections anywhere they occur, even in peripheral and IoT devices, using code that is small enough and efficient enough to run in the smallest of devices.
- On an almost continuous basis, and not just on scheduled or on-demand intervals, using code that is efficient enough to run using minimal resources and thus having a negligible impact upon its host device.
- With the resiliency to continue detecting even when significant components of the detection system have been themselves attacked and compromised.

- With the speed necessary to keep a malware infection from spreading and proliferating throughout a device and/or to other devices.

Although these abilities have been absent in the past, they now exist. Crytica Security Inc. has produced the world's first such malware detection system. It is a system designed to, and is capable of, integrating into existing cybersecurity products and "security stacks". It is the missing link for any effective security platform. By reducing the dwell time from months to minutes, Crytica strengthens all the other Cybersecurity tools. With Crytica "inside", cybersecurity systems can now detect, so that they can now truly protect.

6. NEXT STEPS

To find out more about dwell time reduction, Crytica's Rapid Detection & Alert (RD) system, including Persistent Detection™, and about how Crytica has solved this critical problem, contact Crytica Security, Inc.

Website: CryticaSecurity.com
Email: info@cryticasecurity.com