

Intelligent Security: Using Machine Learning to Help Detect Advanced Cyber Attacks

A progressive, data-driven model of cybersecurity emerges to speed up detection time and reduce risk.



Attackers aren't going to wait for your security software to catch up – industry reports show advanced cyber-attacks can go undetected for approximately 200 days. In today's threat environment, organizations need intelligent security solutions that continually evolve to keep up with the latest threats as they emerge.

Is your organization able to find the signal in the noise of data-points? Continue reading to learn how a progressive security model can help you reduce your risk.

TABLE OF CONTENTS

04	200 days before detection? Industry reports paint a murky picture
05	The \$1 billion threshold: Risk exposure is steeper than ever
07	Modus operandi: Advanced attacks at work
09	One step ahead: Moving to a proactive security model
11	Improving detection: The importance of clear signal
12	From months to minutes: Applied analytics and continuous improvement

200 days before detection?

Industry reports paint a murky picture.

When security professionals detect a breach, it's almost certain that the attacker has been active in the victim's environment for some time. But how long?

For many in the industry, "[200 days](#)" has been accepted as a standard to frame the average. But this "standard" is also problematic for a couple of reasons.

For one, let's face it, that's a long time. It's roughly six-and-a-half months that a sophisticated cyber attacker or syndicate has been at work inside your systems. Throughout this dark and exposed time, your organization's sensitive data and intellectual property have been potentially exposed, moving closer to inevitable compromise.

The fear of what goes on during those 200 days has made this statistic a yardstick for CISOs, CSOs, and even CEOs. Today, companies, security professionals and the tech industry at large are thirsty for new, more advanced security measures to drive that number down.

As a practical matter, "200 days" is just a milestone, a figure used to measure and discuss the industry's progress. CISOs and CSOs know that the number of days isn't the most important element of a breach. What keeps us all up at night is the fact that one day is too long, and by the time you find out, it's always too late. Shrinking that number to zero is the ultimate goal.

To do that, organizations need a more intelligent approach to detect threats earlier and turn the tide against sophisticated cyberattacks. This whitepaper is designed to give readers a glimpse into how advanced threats are working to compromise your sensitive information, and how the advanced computing power of the cloud, combined with data science and human experts, can help reduce the time it takes for your organization to detect an attack.

At the same time, how accurate and useful is this somewhat arbitrary number? Estimates vary even within the security community. Here's a quick snapshot of how different companies see the problem:

146 days:

[M-Trends report \(2016\)](#), Mandiant, a FireEye Company

229 days:

Lockheed Martin's current [Advanced Threat Monitoring](#) page

200 days:

Microsoft's own [Advanced Threat Analytics](#) page

On the rise again?

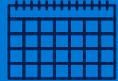
Adding to the mix, Verizon's most recent [Data Breach Investigation Report](#) declines to put a specific number on it, but shows the deficit actually grew over the past year after improving slightly in 2014.

The \$1 billion threshold: Risk exposure is steeper than ever.

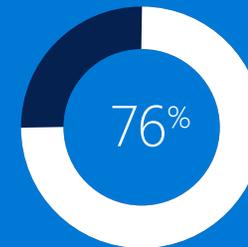
Regardless of whether it's a small business or the largest enterprise, attackers that deploy advanced exploits are a constant concern that goes well beyond the initial costs of a breach. Highly skilled, well-funded and constantly evolving, these perpetrators have motives that range from theft, to industrial espionage, to full-blown nation-state attacks.

First, there are the financial concerns. Today the many malicious actors and authors that utilize advanced attacks are looking to profit from their efforts. It's no surprise, then, that the damages keep going up.

In 2015, a new threshold was reached when a sophisticated attack ring successfully breached more than 100 banks across 30 countries, with losses estimated to [exceed \\$1 billion](#). Because of the heightened risk, cyber-insurance policies are becoming a new operating expense for many companies, with premiums for that emerging offering [set to triple by 2020](#), approaching \$7.5 billion.

200+ 

The median number of days attackers stay within a network before detection.



Compromised credentials make up to 76% of all network inclusions.

The total cost of cybercrime to the global economy could be as high as:

\$500B

There are also the less quantifiable and potentially costlier scars that successful cyberattacks leave such as damaged brands, wary customers, stagnant growth, and compromised diplomatic relations. While not directly attributable to a dollar sign, these impacts can have lasting negative effects on an organization: driving down customer loyalty, driving up public skepticism and ultimately impacting security operations staff who must be held accountable for breaches.

Other attacks are motivated not by financial incentives, but by a quest for sensitive information. Take [STRONTIUM](#), for example. STRONTIUM is a well-known activity group whose targets include government bodies, diplomatic institutions, journalists and military forces. They're not after money and don't care about size of a target. They're after the most sensitive data they can find. Similarly, the [Red October](#) attack group uncovered in 2013 was found to have been infiltrating government and diplomatic institutions for at least five years.

Although it sounds like something out of a spy novel, it's a real issue. Unseen costs of security breaches are something that even two decades ago would sound like the plot of a sci-fi story. With so much at stake, it's no wonder that budgets are increasing and companies are hungry for new solutions to address the growing problem of advanced cyber attacks.

Average cost of a data breach to a company:

\$3.5M

One in five small and medium businesses are targeted in cybercrime attacks.



The impact of lost productivity and growth caused by cybercrime is estimated to be: \$3T

-[Ponemon Institute](#)

Modus operandi: The advanced attack at work

What does an advanced attack do for those 200 days after it's gained entry to your network? Today attackers employ a mix of methods, using traditional techniques alongside new ones as they constantly explore ways to exploit both people and technologies. The longer an undetected attack lives in your system, the more intel it can glean, underscoring the importance of early detection.

Nearly [80 percent](#) begin with a good old-fashioned con job, using spear phishing attacks with compelling ruses to get users to compromise their information. But as security provider McAfee [noted recently](#), more sophisticated attacks are on the rise, including new integrity attacks that can modify internal processes and re-route data as it flows through the network. (This was the technique used in that \$1 billion bank heist.)

Attackers continue to evolve with new forms of malware that can better hide from detection or erase themselves altogether. Attack vectors are also changing: No longer content with targeting PCs and servers residing in the corporate headquarters, attackers look to compromise satellite offices, workers' home computers, and even the software inside of [cell phones](#), [wearable devices](#) and automobiles.

Understanding the Cyber Kill Chain®

Breaches generally involve six clear phases, known in the security intelligence community as the Cyber Kill Chain® (a phrase [trademarked by Lockheed Martin](#)). These phases can occur sequentially, in parallel or in a different order altogether, and each also offers an opportunity to gain intelligence to defeat attackers:



Reconnaissance

The attacker explores his target. This may involve technical procedures or simply browsing the company's web site. It often goes undetected, but the potential is there to correlate seemingly benign behaviors for advanced warning of an attack.



Exploitation

The code compromises the system. Sometimes the delivered code begins immediately to do the attacker's bidding. Other times the attack takes on multiple phases, such as when the initial package begins downloading other code, exposing itself to network alerts.



Weaponization

The attacker creates a shell to hide a malicious payload. It's not always possible to detect the attack's particular weaponization vehicle, but once discovered and reverse-engineered, it becomes a clear footprint for similar attacks later on.



Command and control (C2)

The attacker and the code work together to exploit the system. This may take the form of lateral movements designed to acquire higher-value credentials, or directly exploring the network to find the targeted data assets.



Delivery

The attacker infects the system with malicious code, or dupes a user into downloading it. This is the critical phase where the attacker gains entry and begins to do his work.



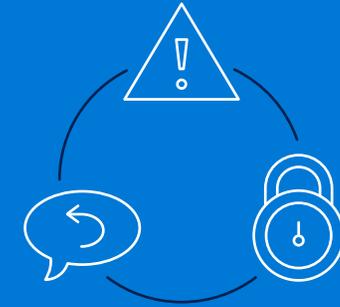
Actions on intent

Sensitive data is taken. At this point, the attack has been successful. Whether it's your customers' financial information, top-secret documents, or the blueprints for your next-generation product, it's now in the attacker's hands.

One step ahead: Moving to a proactive security model

Due to the stealth nature advanced attacks employ in carrying out their attacks, companies must shift to a more proactive security model that focuses on improving their ability to sniff out the attacker and stop him in his tracks.

Whereas the traditional model of enterprise security began with protecting the network perimeter, experts now suggest a more proactive approach that begins with detection enabled by robust security analytics. Under this model, everything else flows from there in a constantly improving cycle, as pre-breach defenses are continually improved with new intelligence from post-breach detection and response.



Detect

By focusing on analytics-based detection, companies are better poised to defend against advanced attacks and their evolving tactics.



Respond

Rather than just patching a vulnerability, the response phase becomes a valuable source of new intelligence.



Protect

The data gained in the detect and response phases is applied to help continually improve pre-breach defense technologies.

For the past few years, CISOs and CSOs have been working to make this shift by implementing security intelligence measures that use data and analytics in an effort to rapidly detect the next attack and improve defenses overall. This includes steps like:

- Investing in advanced security software and secure hardware.
- Training employees on security imperatives and risks.
- Deploying a Security Intelligence Event Management (SIEM) solution.
- Subscribing to (often multiple) threat intelligence feeds.
- Developing processes to correlate threat data, and even hiring data scientists to analyze it.

Thus far, these tools and processes have comprised the bulk of the industry's response to advanced attacks. And like many early stage efforts in the tech industry, they have had mixed results.

It's not that they aren't effective. Mandiant's [2016 M-Trends report](#) shows that when companies are successful at detection using their own systems, the time of an advanced attack's residency is cut drastically. But there are also complaints – including the expense, cumbersome integration, and the inefficient manual process of correlating threat data and feeding it into the system.

And once everything is in place with your SIEM, there's another big problem — noise. There are simply too many alerts, too much data, for even the most advanced enterprise companies to make sense of it all.

If the goal of all these efforts is to shorten those 200 days to near real time, then cutting through the noise has become a major roadblock, and part of what keeps detection a (costly) step behind.

To keep up with advanced attacks, organizations should continue investing in their SIEMs and associated process. Only the cloud can offer next-generation protection, detection and remediation at the scale needed today — including alert mechanisms integrated through platform sensors — in a way that constantly evolves to improve protections with true security intelligence.

Improving detection: The importance of clear signal

In reducing the time it takes to detect an attack, enterprise companies are struggling with a contradictory dilemma between having too much security-related data to process yet still not having enough information to separate the signal from the noise and understand an incident quickly.

The challenge here is not just sheer volume, but also separation. Many indicators of attack either seem innocent on their own, or are separated by industries, distances and timeframes. Without clear insight into the whole dataset, early detection becomes a game of chance. Even the largest enterprise companies are facing these limitations:

- Real threat intelligence requires more data than most organizations can acquire on their own.
- Finding patterns and becoming smarter in that huge data pool requires advanced techniques like machine learning along with massive computing power.
- Ultimately, applying new intelligence so that security measures and technologies constantly improve requires human experts who can understand what the data is saying, and take action.

This is where Microsoft is working to turn the tide. As a platform and services company, Microsoft's threat and activity data comes from all points in the technology chain, across every vertical industry, all over the world.

Microsoft's security products and cloud technologies are designed to work together to report malicious threat data as problems occur. This provides a "flight data recorder" that enables us to diagnose attacks, reverse engineer advanced threat techniques, and apply that intelligence across the platform.

The scope of Microsoft's threat intelligence spans literally billions of data points:

35 billion
Messages scanned monthly

600,000
Known spam email addresses tracked

250 million
Windows Defender users worldwide

600 million
Computers reporting monthly

8.5+ billion
Bing web-page scans per month

1 billion
Customers across enterprise and consumer segments

200+
Cloud services

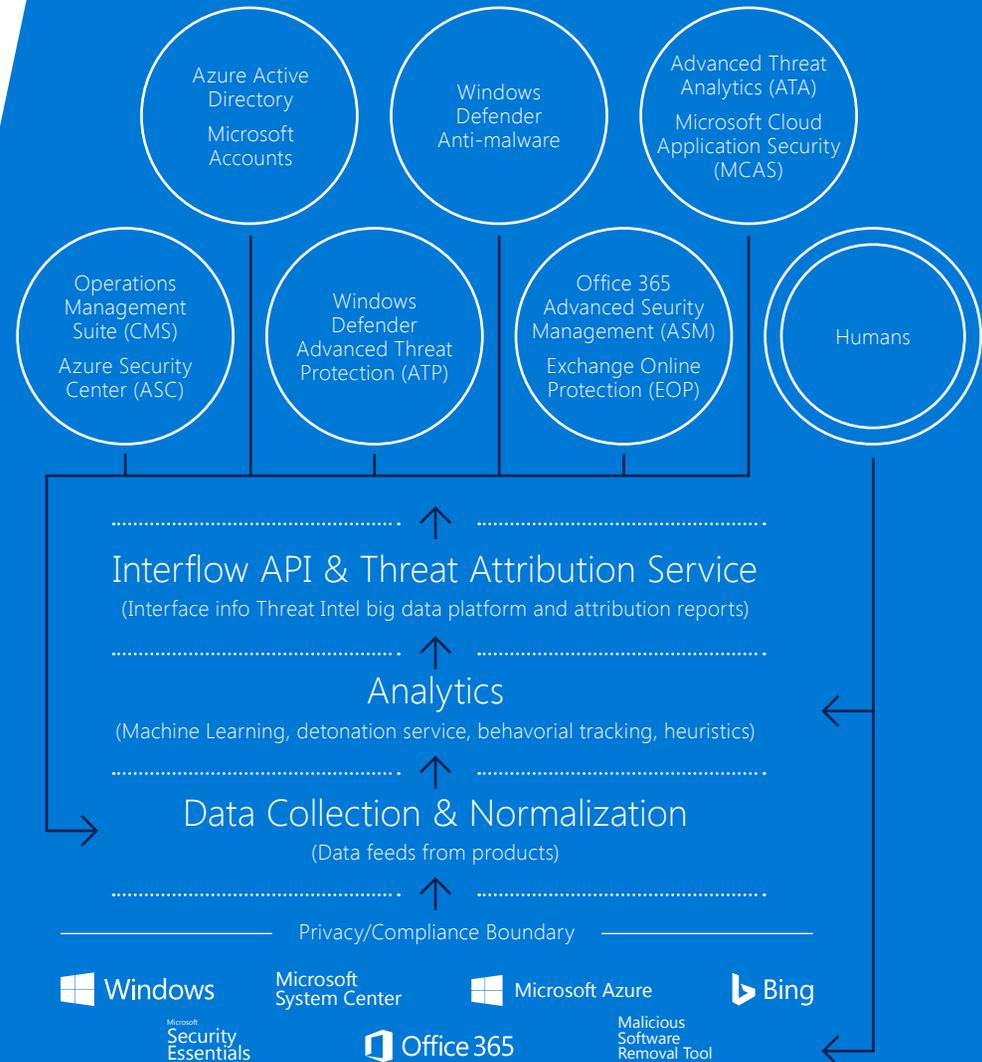
From Months to Minutes: Applied Analytics and Continuous Improvement

For nearly two decades, Microsoft has been turning threats into useful intelligence that can help fortify its platform and protect customers. Since the [Security Development Lifecycle](#) born from early worm attacks like Blaster, Code Red and Slammer, to modern security services woven into our platforms and services, the company has continually built processes, technologies and expertise to detect, protect, and respond to evolving threats.

Today, with the immense computing advantages afforded by the cloud, the company is finding new ways to use its rich analytics engines driven by threat intelligence to protect our customers. By applying a combination of automated and manual processes, machine learning and human experts, we are able to create an intelligent security graph that learns from itself and evolves in real-time, reducing our collective time to detect and respond to new incidents across our products.

Intelligent Security

Through this intelligent security graph, Microsoft is creating the most comprehensive and agile mechanism in the industry to share threat intelligence, apply analytics, and improve detection across its products and services portfolio — not in 200 days, but right now.



Learn More About Security Intelligence at Microsoft

Like the threat landscape itself, Microsoft's approach to Security Intelligence is continually evolving. Customers can learn more and stay current with new developments through several resources:

[The Microsoft Secure Blog](#)

Microsoft experts on the evolving threat landscape offer their thoughts on how both Microsoft and the industry at large are working to stay ahead.

[Azure Active Directory Identity Protection](#)

One product that is continually updated with the latest security intelligence is Azure AD Identity Protection. Customers are encouraged to access the public preview to see for themselves how the data surfaces to help ensure security protocols are up to date.

[The Microsoft Security Intelligence Report](#)

Twice a year Microsoft publishes an in-depth report of security trends and underlying data. Customers are encouraged to use the SIR as a way to help prioritize and focus security efforts.



©2016 Microsoft Corporation. All rights reserved. This document is provided "as-is." Information and views expressed in this document, including URL and other Internet Web site references, may change without notice. You bear the risk of using it. This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes.