DARKTRACE

WHITE PAPER

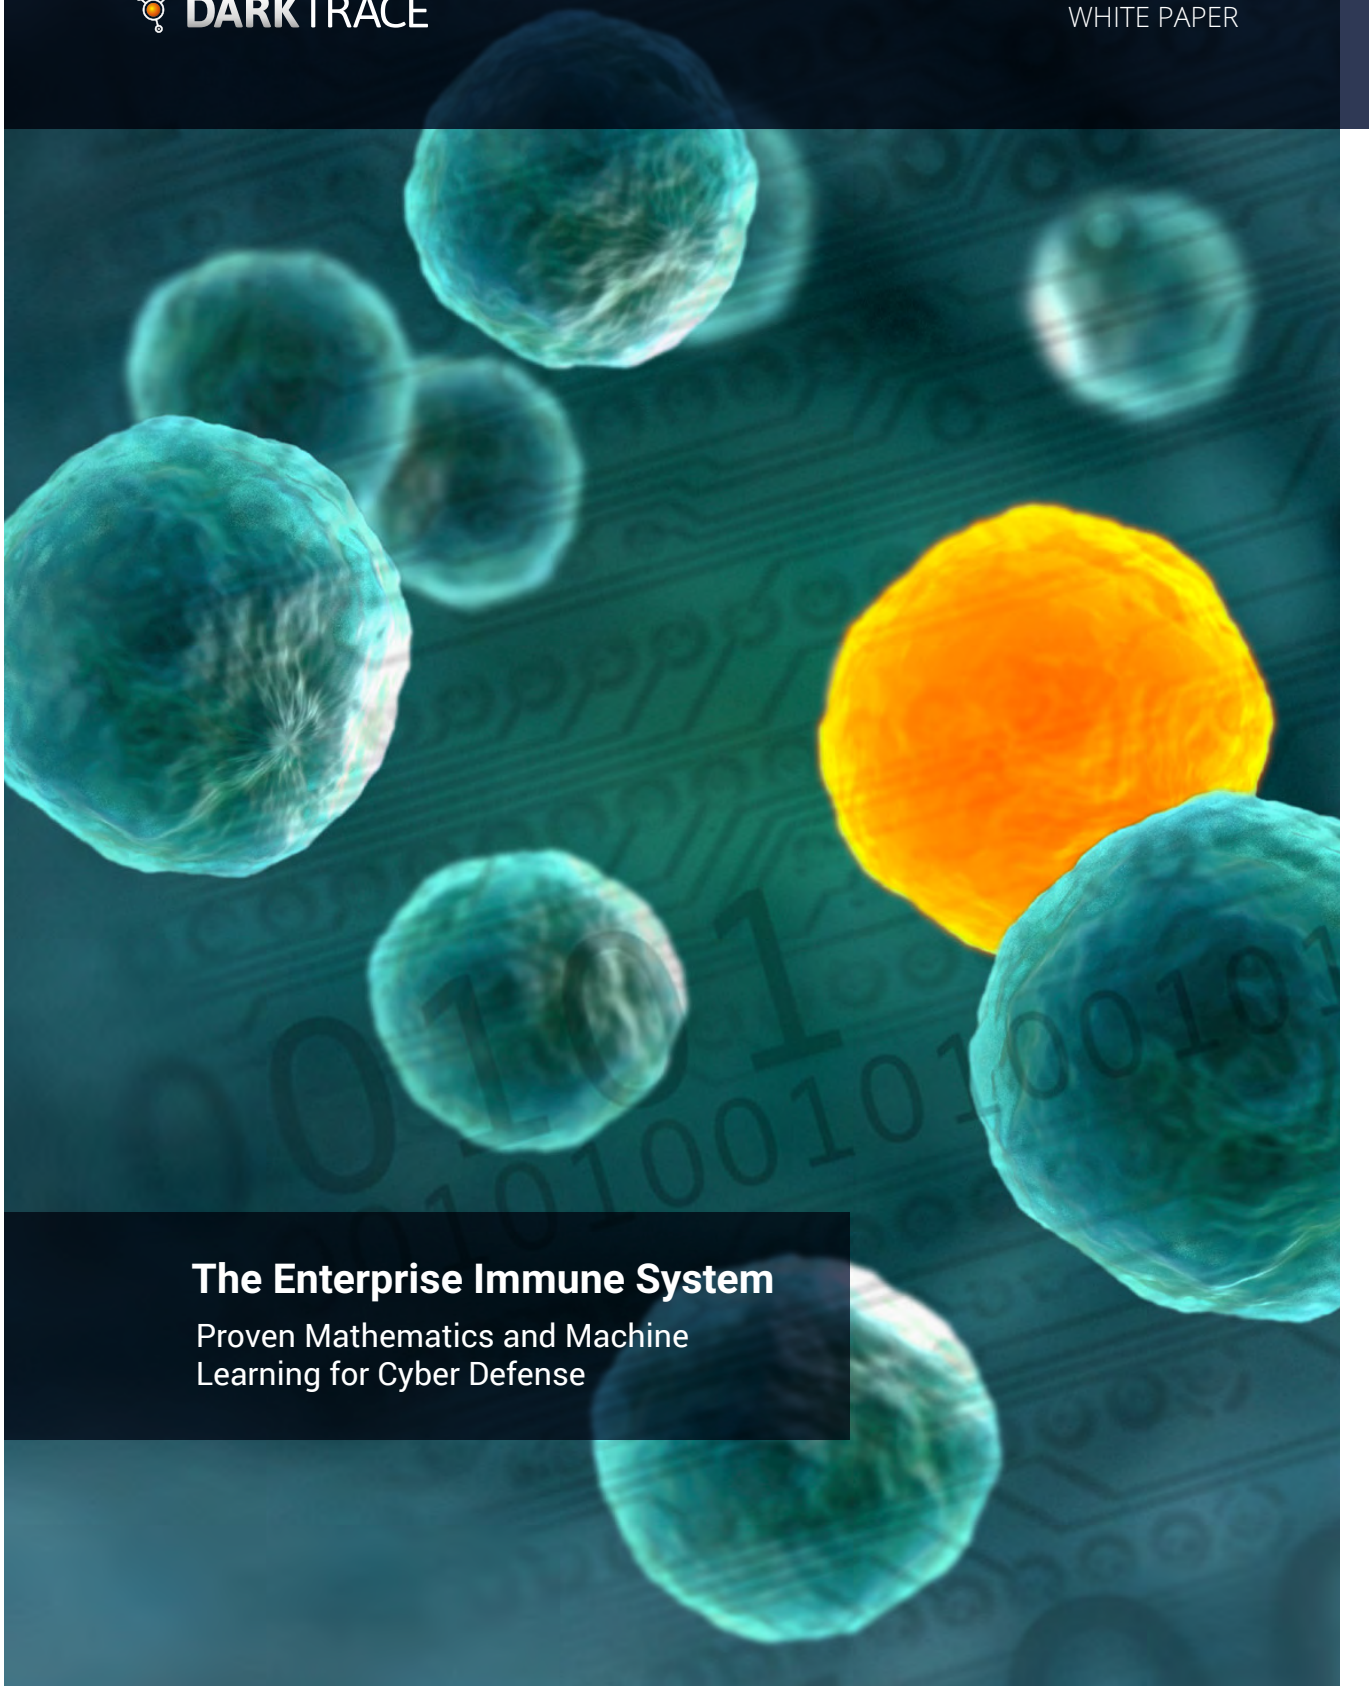

The Enterprise Immune System

Proven Mathematics and Machine Learning for Cyber Defense

Executive Summary

By leveraging the power of advanced machine learning and mathematics, Darktrace takes a fundamentally new approach to the challenge of defending against sophisticated cyber-attacks.

This approach, known as the Enterprise Immune System, represents a new category of cyber defense technology. It deploys cutting-edge technology modeled after one of the most powerful biological systems — the human immune system. By understanding what 'normal' looks like, our immune systems can distinguish 'us' from 'not us' to quickly identify potential threats. This self-learning capability gives it the means to adapt and evolve in an intelligent manner, defending against new threats and changing environments.

The self-learning mechanisms of the human immune system inspired Darktrace's flagship technology: the Enterprise Immune System. Powered by advances in unsupervised machine learning developed by specialists from the University of Cambridge, the Enterprise Immune System intelligently detects emerging threats that other security tools miss. The system can also automatically fight back with 'digital antibodies' — targeted responses that complete the end-to-end functionality of the Enterprise Immune System.

To devise this new approach, mathematicians from the University of Cambridge developed a breakthrough in Bayesian theory. This new branch of probabilistic mathematics, called Recursive Bayesian Estimation, helps to make sense of huge data sets, deal with uncertainty, and, ultimately, identify cyber-attacks amid the noise of a network.

These groundbreaking advancements generated the first-ever immune system for the enterprise. With thousands of deployments around the world, the technology has equipped organizations of all sizes to defend themselves in a world of constant threat, where attacks move at machine-speed and strike without warning.

> "Today's threats are far too unpredictable for traditional security measures to be effective – and manifestly they are failing repeatedly.
> The good news for defenders is that new technologies are now available that evolve and adapt in line with constantly-changing networks and constantly-changing threats, allowing them to spot early indicators of suspicious activity, amid the noise of the network."
>
> **Alan Wade, former CIO, Central Intelligence Agency**

The Threat Landscape Has Changed

In recent years, cyber-attacks have grown exponentially more sophisticated. It is not just a question of data theft and defaced websites. Now, organizations have to contend with threats that are far more subtle and stealthy, and may persist inside the network for many months.

Today's highly disruptive attackers often want to destroy the very integrity of data, rather than just take it. Through targeted data manipulation, cyber-attackers can undermine confidence in entire systems and companies, all while remaining undetected. Blending into the noise of the network, malicious code can activate for only a few milliseconds a year — more than enough time to change data, but not enough time for the security team to know what, if anything, has changed. For instance, a threat-actor could tweak bank account numbers to wreak financial havoc, or adjust data from an oilfield sensor to trick a company into mining a depleted area.

Moreover, modern attacks are trending away from human-operated threats and toward automated attacks and artificial intelligence. Ransomware attacks have been endemic in recent years. This new scourge is an automated attack that moves at machine-speed, encrypting an entire network in a matter of minutes. The attacker then demands a fee for the organization to regain control of the network.

The rise of ransomware signals a sustained move toward AI attacks that self-learn and blend into network traffic, without needing a human controller.

Finally, the rise of nation-state hacking represents a profound shift in the modern threat landscape. Foreign nations can now employ devastating cyber-attacks to gain an economic or political edge. As seen in the high-profile hack on the Democratic National Committee, foreign powers are now willing to use cyber-attacks to influence politics and erode public trust. With vast resources at their command, nation states can afford to deploy the most advanced tactics in these so-called 'trust attacks', which will increasingly target private sector organizations.

Given the sophisticated state of the modern threat landscape, the risk facing modern companies is both inherent and omnipresent. Organizations now have to contend with near-ubiquitous threat, and the cyber security solution for the future has to accept this new reality.

The Threat Is Already Inside

The traditional approach to cyber security relies on a distinction between inside and outside. However, boundaries are virtually impossible to define in modern infrastructures. Today's networks are global, complex, and porous. They have to be, in order to remain competitive. But while this has increased productivity, it has also introduced new vulnerabilities. Given the constant risk that organizations face, the only sensible way to approach cyber security is to assume that the threat is already inside.

Malicious code can lie dormant for years before initiating a kill-switch. Or an attacker can hide in plain sight, blending into network traffic and altering sensitive data. And the growing Internet of Things has given threat-actors an unprecedented number of entry points. Using traditional tools, it is impossible to know whether an organization has been infiltrated.

The threat posed by insiders is frequently underestimated — and remains an extremely difficult problem to solve. Edward Snowden proved that even the best-defended and most security-conscious organizations are vulnerable to lone attackers who move silently within their systems and have the means to undermine their entire operation.

While the most damaging internal attacks come from measured, intentional action, insider threats span the entire spectrum of severity and motivation. Indeed, many insider threats aren't malicious at all. Everyday actions and forgotten security protocols — like taking work home via the cloud — pose a serious security risk. Moreover, employees can fall victim to a customized phishing attack, or may misuse access privileges for the sake of convenience.

> "Good cyber security is not just about a really strong wall on the outside, but some kind of an immune system within."
>
> **Lord Evans of Weardale, former Director General, MI5**

And employees are just the tip of the iceberg when it comes to insider threats. Third-party vendors, customers, and anyone with network or physical access represents a potential threat. Without company loyalty, they have even more motivation to take shortcuts or exploit their position for financial gain. With such a broad range of potential threats, it is nearly impossible to identify high-risk users in advance.

This core principle lies at the heart of the Enterprise Immune System. Networks are inherently at risk from insider threats and external attackers. The Enterprise Immune System operates on the basis that the only way to secure a complex, fast-moving information system is to assume that it's already been compromised.

> "The reality of cyber security today is that border defenses are not enough to keep fast-moving attacks out. Using machine learning, Darktrace's unique Enterprise Immune System detects zero-day threats and suspicious insider behaviors, without having to define the activity in advance."
>
> **Michael Sherwood, CIO, City of Las Vegas**

The Legacy Approach

Network perimeters still represent the crucial first line of defense. But as the industry now recognizes, perimeters on their own are ill-equipped to handle today's sophisticated cyber-attacks.

There are three key reasons why legacy systems have failed:

1. You can't keep threat out

The traditional approach assumes that you can keep attackers out by strengthening the perimeter. Organizations have invested large amounts of time and money into perimeter controls to insulate their information systems from attack.

Unfortunately, threat-actors have proven to be more than capable of overcoming perimeter controls, and insiders can bypass the perimeter altogether. Companies have to work on the assumption that they are constantly at risk, and that many threats — and certainly the most insidious — can infiltrate their organization with relative ease.

2. You can't define the threat

Definitions and signatures lie at the heart of traditional security. This core IT principle, widely used for all manner of automated applications, was duly transmuted to cyber security. Indeed, many solutions still create definitions of what 'bad' looks like, and protect the network against that type of 'bad', if and when it is encountered again.

These systems need to be pre-programmed with signatures of past attacks. And yet, there is no guarantee that future attacks will look anything like what came before. In fact, considering the constantly evolving threat landscape, it is rather unlikely.

While this approach may protect against unsophisticated attackers who repeatedly use the same tactics and toolkit, serious attackers constantly change their strategies and use custom malware to conquer a specific target.

3. You can't assume the threat is purely technical

In a world of botnets, Trojans, and Remote Access Tools, it can be easy to forget the human behind every cyber-attack. The most serious cyber-threats are directed by skilled agents who move deftly through the network. The traditional approach is incapable of dealing with the complexity and subtlety that such attackers bring to their missions.

Both internal and external parties usually exhibit distinct behaviors before engaging in malicious acts. A contractor logging on at an unusual time, groups of files being aggregated, or an unusual volume of email traffic — these are all signs that are often meaningless to legacy tools, but form a compelling picture when correlated.

> "A machine learning approach is critical to cyber defense. The self-learning technology only focuses on the most important threats and finds abnormalities without any prior assumptions."
>
> **Vari Bindra, Head of Cyber Defense Center, Blackhawk Network**

The Enterprise Immune System

The Enterprise Immune System represents a fresh approach that has successfully transformed the cyber security landscape. Hundreds of businesses now benefit from this innovative solution, which is built around the premise that organizations face a constant level of threat. This cutting-edge technology is capable of learning 'self' on an adaptive, real-time basis — thereby understanding when and where abnormal behavior first manifests.

Like viral DNA, modern cyber-attacks constantly evolve and mutate. These sophisticated attacks avoid detection by subtly adjusting their behavior. Fortunately for us, the human immune system is just as clever. It is continually learning to understand precisely what constitutes a threat. It's not a perfect system — we still catch the occasional cold — but it plays a critical role by protecting us from threats which, if left unchecked, would be life-threatening. Its adaptiveness enables us to interact with each other and expose ourselves to risk on a daily basis.

The Enterprise Immune System works on the same premise. Built on a foundation of Bayesian mathematics and unsupervised machine learning, the system analyzes complex network environments to learn a 'pattern of life' for every network, device, and user.

Advanced machine learning techniques then correlate patterns in network traffic to detect previously unknown threats and automatically defend networks with digital 'antibodies'. The model doesn't rely on rules or signatures. Instead, it intelligently draws patterns from large sets of data to discover deviations from 'normal' that indicate live, in-progress threats.

Mathematics and Machine Learning 'Done Right'

When machines replaced manual labor in the eighteenth century, it was dubbed the Industrial Revolution. When computers began performing repetitive, rote tasks en masse, the Digital Revolution had dawned. Now, we are in the midst of the third revolution in automation — the Machine Learning Revolution.

In this new era of automation, machines have the ability to exercise precise judgement and carry out advanced, thoughtful tasks that were once reserved for human specialists. Through machine learning, computers are no longer restricted by rules and definitions. They can understand uncertainty — indeed, they embrace it.

The Enterprise Immune System would not have been possible without this revolution in machine learning. However, a more subtle evolution has taken place in recent years, which is at the heart of Darktrace's core technology. Traditional machine learning has been 'supervised', whereby a system is trained using a data set built from pre-classified behaviors. In cyber security, a program would flag unknown behavior as malicious or benign depending on how closely it resembles the known behavior.

Supervised machine learning has utility, but when applied to cyber defense it fails to identify the all-important 'unknown unknowns'. Moreover, the method requires significant human input, and it depends entirely on pre-programmed rules. Such rules belie the subtlety of modern attacks, and instead fall back on a rigid black-and-white framework.

Unsupervised machine learning, on the other hand, allows for shades of gray. Modern attacks exist on a scale of type and severity. Unsupervised machine learning captures the full spectrum of threat, as it doesn't require data sets, pre-defined labels, or any human input whatsoever. Importantly, this lets the system go beyond what its programmer knows, to discover previously unknown threats.

The Enterprise Immune System employs unsupervised machine learning to full effect, as well as a new field of probabilistic mathematics known as Recursive Bayesian Estimation. These advanced machine learning algorithms are designed to analyze network data at scale and intelligently handle the unexpected. The system does not rely on knowledge of past attacks. Instead, it discovers previously unknown threats by detecting deviations from normal behavior.

To learn 'normal' for a network, the Enterprise Immune System identifies naturally occurring groups of devices and behaviors — a task that would be impossible to do manually. Darktrace then employs advanced clustering methods to analyze network behavior in terms of similar devices on the same network. This generates a picture of 'normal' without reference to external data and without human interference.

While traditional systems adopted a binary approach, Darktrace accepts the inevitable ambiguity of such data. The Enterprise Immune System recognizes that behavior isn't merely 'malicious' or 'benign'. By correlating a broad range of factors, like server access, timing, and data volumes, Darktrace intelligently ranks threat. This simultaneously allows organizations to prioritize the most serious threats, and eliminates the problem of false positives.

Equally important is the task of learning the unique topology of intricate network structures. To achieve this, the Enterprise Immune System utilizes iterative matrix methods that reveal relationships between network features. In conjunction, Darktrace uses an innovative application of models from statistical physics to map a network's 'energy landscape' and reveal potentially anomalous substructures.

A further problem lies in how to handle the huge number of variables involved in modeling the high-dimensional structure of complex network environments. In the observation of packet traffic and host activity within an enterprise LAN or WAN, where both input and output can contain millions of inter-related features, learning a sparse and consistent predictive function is challenged by a lack of normal distribution.

In this context, the Enterprise Immune System is the most advanced, large-scale computational approach to learning sparse structure I/O models. It achieves this by extending the L1-regularized regression model – also known as the lasso method – to a family of sparse 'structured' regression models. This allows for the discovery of true associations between linked malware, C2 events (inputs), and data egress (outputs), efficiently solving convex optimization problems to yield parsimonious models.

In combination with the advanced probabilistic mathematics of Recursive Bayesian Estimation, these models generate a comprehensive picture of an enterprise network, granting full visibility of the network structure in order to spot emerging threats.

The Enterprise Immune System takes this one step further to create a truly self-learning and adaptive technology that can even fight back with 'digital antibodies'.

Darktrace's cutting-edge application of unsupervised machine learning was a pivotal moment in the cyber security industry. For the first time in history, a defense system could learn the precise structure of complex network environments to create a picture of 'normal' behavior, iteratively adapting itself to detect subtle deviations and discover previously unknown cyber-threats, all in real time and without the need for human involvement.

> "Darktrace applies mathematical models to create statistically significant views of user, device and network behaviors – an approach that makes it adept at detecting attacks that are already within the enterprise."
>
> **Eric Ogren, 451 Research**

Conclusion

In an era of pervasive threat, Darktrace's novel approach to cyber security has equipped businesses to intelligently monitor their networks and automatically fight back against the most serious cyber-attacks. The Enterprise Immune System is a cutting-edge defense system capable of learning 'normal', evolving with a network, and detecting early-stage cyber-threats.

The Enterprise Immune System allows organizations to understand threat holistically. The system is continually learning, meaning it can handle unpredictable and sophisticated threats. Crucially, its self-learning mechanisms allow the system to move in step with both the organization and the evolving threat landscape.

Organizations that have implemented an Enterprise Immune System benefit from the world's leading advances in machine learning and mathematics to protect against cyber-threats, all while maintaining the flexibility and connectivity that modern businesses thrive on. Darktrace's technology sits on the cutting-edge of cyber security, with a proprietary technology designed around groundbreaking mathematics and machine learning, and purpose-built to provide complete network visibility and detect emerging threats in real time, which would otherwise go unnoticed.

About Darktrace

Darktrace is a world-leading cyber-threat defense company. Its multi-award-winning Enterprise Immune System technology automatically detects and responds to emerging threats, powered by machine learning and mathematics developed by specialists from the University of Cambridge. Without using rules or signatures, Darktrace models the 'pattern of life' of every device, user and network within an organization, identifying and mitigating cyber-threats before damage is done. Darktrace's self-learning technology has been deployed globally and across all sectors, including energy, retail, telecommunications, manufacturing, financial services and healthcare. The company is headquartered in San Francisco and Cambridge, UK, with over 20 global offices including London, New York, Milan, Mumbai, Paris, Singapore, Sydney, Tokyo and Toronto.

www.darktrace.com

Certified Partner:

CyVent

PLAN. PROTECT. PRE-EMPT

1395 Brickell Avenue, #800
Miami, FL 33131
+1 (305) 299 1188
www.cyvent.com