DARKTRACE
Evolving threats call for evolved thinking™

Self-Learning AI for Zero Trust Environments

Securing your Zero Trust Journey

CONTENTS

Key Takeaways

- Complements, enhances, and interoperates with a zero trust model and architecture
- Illuminates and interrupts novel attacks and insider threats that operate over legitimate paths or evade policy-based defenses
- Natively integrates with zero trust technologies: IAM, Web Gateways, Microsegementation, and Firewalls
- Reinforces zero trust policies through real-time visibility, continuous monitoring, and autonomous decision-making

Darktrace interoperates with zero trust technologies via native integrations, while validating current zero trust policies and informing future micro segmentation efforts with continuous, real-time visibility across the entire organization. Crucially, this continuous monitoring is adaptive in its understanding and pervasive in its scope, delineating normal and abnormal patterns across email, cloud, and collaboration platforms, as well as remote endpoints, IoT, and the corporate network.

The Next Step in Zero Trust

Where did Zero Trust come from?

The greatest challenge IT teams have had to face in recent years has been enabling the workforce for remote work. With these changes we saw a major shift in security challenges, as our digital footprint expanded to include the homes of employees. Home Wi-Fi, employee laptops and VPNs became prevalent vulnerabilities to company security. At this point, the world really started to pay attention to the latest trend: Zero trust. The concept of 'zero trust' is more a philosophy than a technology, replacing the implicit trust of the legacy device model with a more dynamic approach that assumes breach and verifies intelligently, while restricting access and operations accordingly.

"Zero trust is a response to enterprise network trends that include remote users, bring your own device (BYOD), and cloud-based assets that are not located within an enterprise-owned network boundary. Zero trust focuses on protecting resources (assets, services, workflows, network accounts, etc.), not network segments, as the network location is no longer seen as the prime component to the security posture of the resource."[1]

In other words, before a device has access to any company data, the user must verify their identity, replacing the previous model where the device itself was trusted and had access to data by nature of being a company device.

How is Zero Trust Implemented?

In practice, this is typically implemented in the form of security policies, whether via microsegmentation, Web Gateways, or least-privilege access control. It is often associated with the Secure Access Service Edge (SASE), SD-WAN, and other security and networking services designed to accommodate the new shape of digital business. In other words, legacy security tools were designed for a castle and moat security architecture, but our digital estate has an everchanging perimeter that can no longer fit that philosophy.

Zero trust technologies enforce guardrails for organizations with rules and policies designed to reduce risk exposure by eliminating unnecessary access and privileges across critical IT systems. But zero trust isn't a status that can be achieved – rather, it's a philosophy that organizations adopt. There never will be a 'state' of zero trust and therefore there never will be a state of zero risk. People, processes, and technologies are constantly changing so risk management efforts will be constantly underway. Zero trust technology should be dynamic by nature as the risk it intends to mitigate is as well. Darktrace AI constantly and dynamically analyzes your entire infrastructure, whether it is in the cloud, on premise, or even in software applications.

> The shift to remote and hybrid work has increased the attack surface for organizations and underscores the importance of securing the identity of each user. Organizations need to assume attackers will still inevitably breach their perimeter defenses, including identity controls.
>
> / Max Heinemeyer, Chief Product Officer, Darktrace.

1. https://csrc.nist.gov/publications/detail/sp/800-207/final

Challenges of Securing a Remote Workforce

The shift to remote and hybrid working has taken many employees out of firewalled offices, and created new security risks. With employees working from offices and homes, hotels and coffee shops, companies face a myriad of cyber-threats. Legacy security systems were designed to protect static data at a centralized location, whether it was a datacenter or an office. The moment the employee or the data itself leaves this supervised environment, tools lose visibility and capacity such as the ability to respond. Meanwhile, cyber-criminals are constantly developing new techniques to exfiltrate and encrypt the information these devices hold.

On the other hand, the users themselves are becoming more of a challenge to secure. Many organizations rely on legacy systems that may not be able to adapt to zero trust principles. The use of multi-factor authentication and authentication tokens can often present an obstacle to companies on this journey due to technical difficulties or internal bureaucracy. Furthermore, many organizations are considering the use of more advanced authentication technology or zero trust tools when the users themselves are still posing a massive risk to the organization. If a company implemented an advanced zero trust architecture but the admin password is in an unencrypted .txt file on their desktop, the entire security stack is made completely vulnerable by that one file.

In other words, zero trust tries to avoid account compromises, but if employees aren't trained, or the security fundamentals aren't there, more advanced technology will be rendered ineffective from that single point of failure. There is a natural human tendency to assume zero trust technology is that 'one stop solution' that fixes our problems, but the reality is organizations still need several layers of defense. Every new piece of technology a company adopts, whether it is an email solution or a cloud application, should include a second layer of monitoring that supervises that activity as well.

The eternal balance between operation development and security has dramatically shifted due to work from home trends and modern technology adoption like the cloud. The zero trust philosophy is the answer to this shift in balance. Naturally, the more tools companies include to enable the workforce the more roles companies need to manage and supervise. Identity Access management becomes more relevant as additional identities are required, either because the toolset expands, or a company grows. Ultimately, we must all face the question of how to scale security when both the perimeter and users are constantly expanding.

Regrettably, the expectations of zero trust have not been met by the realities of its implementation. Whether it is the way organizations implement policies, or the technologies they use, the lack of dynamic flexibility has led to the same static limitation. In other words, if an organization creates a zero trust policy but only reviews it periodically, or if a security team establishes a VPN authentication tool based on a set of defined rules, the same stagnant inadequacies are present. The shift to remote work was revolutionary, but it was just one manifestation of the reality of modern user behavior. Today, users, devices, and even data are borderless and dynamic. Security solutions should not apply old technologies to new philosophies, but generate new technologies based on new philosophies. Modern security technology should be able to understand the ever-changing behavior of assets and adapt in real time to that fluctuation.

Threats Facing the Remote Workforce

The NIST Special Publication on zero trust, highlights a variety of threats associated with zero trust architectures that can be placed into three main categories:

- Credential Abuse (Subversion of the Zero Trust Architecture Decision Process + Stolen Credentials/Insider Threat)
- Network Supervision (Network Visibility + Denial of Service + Storage of System and Network Information)
- Zero Trust Technology Limitations (Reliance on Proprietary Solutions + Use of Non Person Entities).

The final threat category for the purpose of this white paper is out of context as these are mere limitations of the zero trust technology itself and doesn't reflect on the current threat landscape. However, there is extreme relevance in discussing both the threats of Credential Abuse and Network Supervision.

Credential Abuse

Credential Abuse is a prevalent threat we can find in any network architecture, zero trust or not. In the end, the keys to the castle are often the first target for any attacker. However, in a zero trust architecture it is even more relevant, as the entire premise of zero trust relies on the idea that you will not get access to any part of the network until you have identified yourself.

But what happens if the attacker gets hold of a legitimate credential? The attacker could configure the authentication policies to his liking or directly exploit the network moving laterally towards places they have legitimate access to. Legacy IAM tools will focus on set rules like location or time of login, but human behavior is much more complex than that. In fact, the adoption of Multi Factor Authentication (MFA) has curtailed a large number of threats due to the added layer of security, but it is again the human element that has revealed its weaknesses.

A simple but clear example of the weaknesses in MFA is when an attacker manages to force a second factor authorization and without checking its legitimacy, a user approves it. Multi Factor Authentication is designed to make a user think twice before approving access, if we undervalue the double check, it is equally as effective as a single factor. Simple audits of credential usage might miss comprised accounts as an admin user will not be infringing on any policies by accessing strange parts of the network, they have ultimate privileges. For this reason, it is imperative tools adopt behavioral pattern analysis so they can alert when an admin credential is acting unusually and not strictly based on the rigid policies implemented.

Network Supervision

Network Monitoring in itself seems to be a security weakness for many zero trust architectures. Currently, online services make it easy and cheap to deploy a denial of service attack. Imagine if this type of attack were launched towards your IAM tool, preventing your users from authenticating, and dramatically impacting your production.

Conversely, what if there is an infected device already present in the network? Would you be able to detect the subtle symptoms of such an attack? Is your internal network segmentation and policy of least privilege implemented properly?

Zero Trust Technology Limitations

Limited visibility and single points of failure can quickly cause a zero trust architecture to crumble if it is not properly monitored from a network perspective. Furthermore, companies are struggling to keep up with the diversification of assets: IoT, OT, Mobile, Cloud, SaaS, etc. Each of these different data sources often requires a specialist tool to monitor. All zero trust architectures should take into consideration where the data resides and ensure that the security tools implemented span the entire breadth of the digital ecosystem.

Case Study: Solar Winds

Particularly relevant to the conversation on zero trust is the recent attack on the Solarwinds platform[2]. It is claimed that APT Cozy Bear was involved in the deployment of this attack by embedding a trojan in the update of the Solarwinds Plug-in on the Orion platform. Once installed, the malware disguised itself within the network for over two weeks, pretending to be part of the Orion Improvement Program protocol. The threat-actor set the hostnames on their later-stage command and control (C2) infrastructure to match a legitimate hostname found within the victim's environment. This allowed the adversary to blend into the environment, avoid suspicion, and evade detection. They further used C2 servers in geopolitical proximity to their victims, further circumventing static geo-based trusts lists. Darktrace is unaffected by this defense evasion as it does not have implicit, pre-defined trust of any geo-locations.

Once the attacker gained access to the network with compromised credentials, they moved laterally using multiple different credentials, including that of the admin user. The credentials used for lateral movement were always different from those used for remote access. The key element here is that the attackers were able to compromise a high privilege user. In other words, they have breached past the credential infrastructure, and are now able to move within a trusted environment. The value offer of Darktrace/Zero Trust is to learn about your user behaviour in order to autonomously respond to a misuse of credentials or if you do have zero trust technology in place, be the second line of defense that leverages Self-Learning AI to enhance threat detection. The plethora of data feeds we can tap into to provide the necessary context is also a valuable asset to zero trust tools that are designed for specific focuses.

However, past the credential compromise, Darktrace can detect further symptoms exhibited down the line. The attacker used a temporary file replacement technique to remotely execute utilities: they replaced a legitimate utility with theirs, executed their payload, and then restored the legitimate original file. They similarly manipulated scheduled tasks by updating an existing legitimate task to execute their tools and then returned the scheduled task to its original configuration. These movements would trigger a variety of alerts related to uncommon service controls, DCE-RPC requests, AT Service Tasks, SMB Lateral Movements, etc.

By understanding where credentials are used and which devices talk to each other, Cyber AI has an unprecedented and dynamic understanding of business systems. This empowers it to alert security teams to enterprise changes that could indicate cyber risk in real time. The models highlighted above are not specifically designed to detect SolarWinds modifications – they are designed to detect the subtle but significant attacker activities occurring within an organization's network.

These alerts demonstrate how AI learns 'normal' for the unique digital environment surrounding it, and then alerts operators to deviations, including those that are directly relevant to the SUNBURST compromise. It further provides insights into how the attacker exploited those networks that did not have the appropriate visibility and detection capabilities.

It's clear that companies who are moving towards a zero trust philosophy, require an additional layer of analysis either from the network or device perspective, in order to compensate for the natural limitations of zero trust tools. In fact, if you look at zero trust environments that are properly executed, we can expect the trusted environment to be much less chaotic and organized, even better for a Self-Learning AI that is designed for learning normal.

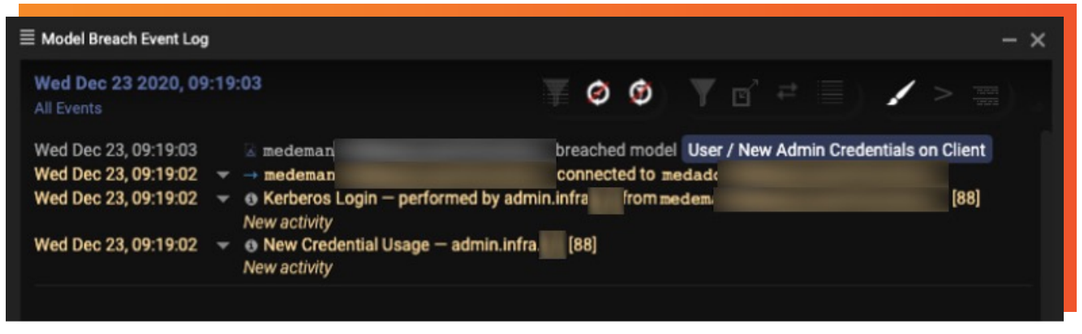

Figure 1: Example breach event log showing anomalous admin login.

2. https://www.radware.com/security/ddos-threats-attacks/threat-advisories-attack-reports/ddos-as-a-service/

How Darktrace Supports a Zero Trust Model

Darktrace delivers unified and adaptive protection across heterogenous, hybrid, and service-based microsegmented architectures, including email, cloud, and application environments as well as remote endpoints, IoT, ICS, and the corporate network. The wide coverage Darktrace provides is complemented by its depth. Darktrace delivers deep visibility into all user and machine activity down to the packet layer, enabling a full assessment of the data environment and architecture to autonomously discover resident threats or malicious activities flowing over legitimate paths. When organizations combine a robust zero trust architecture with Darktrace's Self-Learning AI, attackers have nowhere to hide.

The beauty of leveraging an AI-based tool like Darktrace is that the threat detection method is perfectly aligned with the core tenant of zero trust: assume breach. Darktrace indiscriminately inspects asset activity (data, apps, devices) for suspicious behavior without contrasting it against a list of approved activity. Traditional default access controls are incapable of seeing past the binary approach of authenticated or not. Darktrace by default never has a trusted source, its real-time monitoring analysis continuously looks for attack symptoms and suspicious events even within authenticated or authorized paths. By providing unified visibility that adapts to the business as it evolves, Darktrace can enforce zero trust policies, inform future micro segmentation efforts, and incorporate telemetry from IAM tools and Web Gateways into a broader understanding of the organization.

> Detecting anomalies and enabling response requires visibility to the entirety of your data and the totality of your network. Darktrace not only focuses on threat detection but also on providing visibility into traffic flow within your environment.
>
> / Dr. Chase Cunningham, VP & Principal Analyst, Forrester Mitigating Ransomware with Zero Trust

A holistic approach to zero trust goes beyond an analysis of suspicious network activity. Every aspect of a defense framework must incorporate zero trust, much like every component of the Darktrace Cyber AI Loop™. For example, Darktrace PREVENT™ ensures that the zero trust philosophy is incorporated in your proactive security efforts. A solid zero trust architecture will not place trust based on being an internal or external entity. Darktrace/ASM and Darktrace/End-to-End work together to harden the internal and external defenses, connecting external threats to internal targets. On the other hand, Darktrace RESPOND™ leverages the intelligence in PREVENT and DETECT in order to dynamically enforce your organizations zero trust policies wherever your data resides.

Securing zero trust architecture at every layer

Darktrace supports key zero trust tenants throughout every stage of an incident lifecycle - securing what matters most to your business.

How Darktrace Augments your Zero Trust Technologies

Darktrace/Apps, Darktrace/Cloud and Darktrace/Zero Trust integrate with enterprise software and cloud platform solutions to bring visibility and threat analysis to critical systems. Extending Darktrace's Self-Learning AI beyond the physical enterprise network, each module brings the insight of an AI-genreated natural language incident report (Cyber AI Analyst) and Darktrace's unique 'pattern of life' anomaly detection to enterprise software and cloud-based environments, ensuring that user activity is monitored whether it originates inside the network or from remote locations.

Whether your organization is just beginning to adopt a zero trust philosophy or has advanced architecture in place, Darktrace is there to secure the journey. From any device, any user will have to pass an authentication or authorization process locally or from a third-party solution, before accessing company data. The idea is that the moment that process begins, Darktrace will have visibility over the connections involved. Whether they are using authentication or access-focused technology, Darktrace can ingest the relevant data at the source and enhance the built-in security features by leveraging Self-Learning AI and Autonomous Response. The zero trust journey is not complete until the user accesses the data, but Darktrace will continue to monitor even once the user has been granted that infamous trust. A true zero trust architecture continues to distrust and assume a breach even after the user has authenticated. Darktrace is designed to focus on patterns of behavior rather than rules and signatures, which gives us the edge when it comes to detecting threats via legitimate paths.

When malicious activity occurs despite the enforcement of zero trust rules and policies, Darktrace can instantly alert and trigger a proportionate response to contain the attack. When deployed with zero trust technology, the scope of activity visible to Darktrace widens, and its AI technologies can analyze, contextualize, and act in that realm as well. Upon detecting unusual behavior indicative of a clear cyber-threat, Darktrace's Autonomous Response can directly take appropriate action via the relevant API, ranging from actions as surgical as blocking connections between two endpoints to a complete termination of all device-specific activity.

Identity and Access Management

IAM tools are uniquely positioned in protecting zero trust architectures as they collect valuable data around user behavior. Native Integrations in this area allow Darktrace to ingest telemetry data around user behavior, enriching AI detections and Cyber AI Analyst investigations across the business. Without having to stipulate was constitutes as appropriate IAM activity, Darktrace can alert and act on the anomalous behaviors of any account, including unusual and potentially unsanctioned activity. In particular, administrative users can quickly become a vulnerability for digital environments due to their elevated privileges. Unique to Darktrace, it may also detect unusual administrator activity around newly added user permissions and third-party software to allow-lists or anything that might widen the range of risk exposure. The corresponding module retrieves both administrator and authentication activity from the API via an integration key created during configuration. Darktrace can also tie its finding into IAM flows via SIEM and SOAR tools, and natively supports authentication via SAML 2.0 Single Sign On.

Web Gateway

Darktrace integrates with cloud gateways to enrich its understanding of user and device behaviour while offering the ability to share bespoke AI insights that inform zero trust policy engines.

In particular, the Darktrace Zscaler ZIA integration ingests weblogs from a ZIA device to simulate connection data. Web events produced by the Zscaler logging will be associated with a device of the same hostname. If a device of that hostname does not already exist, Darktrace will create a new device. Connection events created from Zscaler logs will be available to core Darktrace analysis and accessible in Advanced Search. Devices which have ZIA simulated connectivity associated will be automatically tagged with the ZIA tag.

In Zscaler Private Access environments, remote users initiate connections to internal resources through a ZPA App Connector located locally in an organizations private network. Darktrace observes connectivity from the ZPA App Connector to these internal resources, but is unable to resolve these connections back to a specific end user or IP address. The Darktrace Zscaler ZPA integration ingests "User Activity" logs produced by ZPA, allowing connectivity patterns seen in network traffic to be matched back to originating remote users. Devices which have connectivity mapped through ZPA ingestion will be automatically tagged with the ZPA tag.

Microsegmentation

Darktrace RESPOND can take Autonomous Response actions via microsegmentation tools, while the systems self-learning visibility provides microsegmentation validation and data flow monitoring to identify opportunities for further segmentation around application, networks and workloads.

Firewalls

Darktrace RESPOND can deliver coordinated Autonomous Response actions by interfacing with third-party firewalls. Log ingestion is also available for further enrichment.

> While Zscaler's Zero Trust Exchange reduces the attack surface and enforces cyber security policies, the integration with Darktrace AI behavioral detection and response allows customers to correlate Zscaler telemetry with data from across the enterprise to improve threat response further.
>
> / Amit Raikar, Vice President, Business Development and Technology Alliances at Zscaler.

> Darktrace's integration with Duo Security is fantastic. The logs we get from Duo are fed into Darktrace Cyber AI Analyst, providing us with detailed, automated investigations that are extremely beneficial. Darktrace's zero trust integrations mean that more data can be ingested into its Self-Learning AI, improving its understanding of normal for our business. I'm thrilled that Darktrace is expanding these integrations to improve zero trust security and help Darktrace better serve its customers.
>
> / CISO, a major financial services organization in the United States

Securing your zero-trust journey

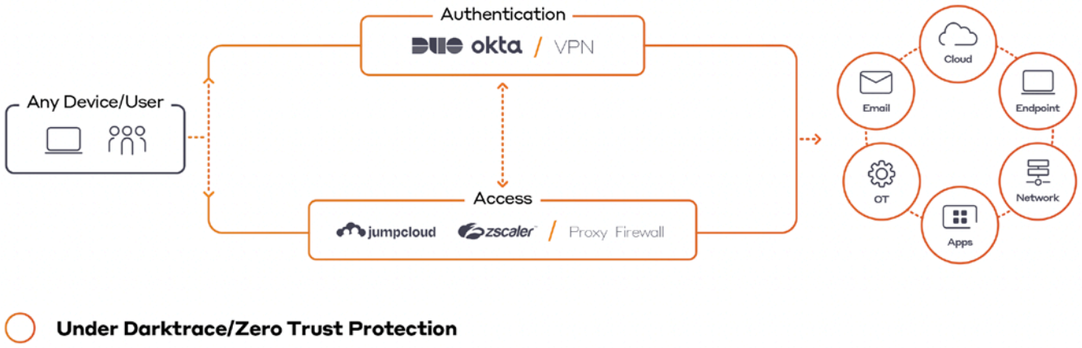

Under Darktrace/Zero Trust Protection

Customer Reviews

"Detecting anomalies and enabling response requires visibility to the entirety of your data and the totality of your network. Darktrace not only focuses on threat detection but also on providing visibility into traffic flow within your environment; this can help you understand application dependencies and identify opportunities for microsegmentation before the attack."

/ Dr. Chase Cunningham, VP & Principal Analyst, Forrester, Mitigating Ransomware with Zero Trust

"We soon realized the extra benefits we could get of the additional Darktrace products was greater visibility around user behaviors and the alerting of potential data loss. The instant reporting of certain DLP activities is now one of the most useful tools in Australian Grand Prix Corporation's cyber security toolbox."

/ Clint Waterson, Division Manager of IT at the Australian Grand Prix Corporation

"Darktrace detects things that are unusual and shouldn't be happening. So, when you think about Darktrace and anomaly detection, I think it helps with malicious insiders and insider threat, because there's not a better way to detect insider threat than anomaly detection systems

/ Dr. Robert Spangler, Associate Executive Director of Operations and Information Technology, The New Jersey State Bar Association

"The AI provides unparalleled visibility into our network traffic patterns and alerts us to behaviour that falls outside of this which could be potentially malicious. This has helped significantly in identifying suspicious events and containing them in their very earliest stages – enabling our security team to work on higher level projects like improving our baseline security and getting the basics right at all times."

/ -George Ho, Vice President of Information Technology at RioCan

Customer Deep Dive: HGG / Cardenas Markets Llc.

Heritage Grocers Group (HGG), one of the leading Hispanic and ethnic focused grocers in the country.

Moving HGG's Zero Trust Philosphy Forward:

Heritage Grocers Group (HGG), is a leading Hispanic- and ethnic-focused grocer group operating in California, Nevada, Arizona, and Illinois. Cyber security is a priority for HGG because of the impact cyber disruption could have on business operations as well as brand reputation. Unfortunately, they must deal with a wide attack surface since the business encompasses multiple brands and operates in several states.

HGG sought out a security tool that would reduce the workload of its security team. Today, HGG leverage Darktrace DETECT and RESPOND for email, network, and Microsoft 365. A crucial differentiator for HGG was how well Darktrace aligned with its zero trust endeavor.

"Early on, HGG made the security policy decision to move from implicit trust to zero trust. With the distributed nature of our stores, business unit offices, increased use of cloud business applications, and our team increasingly working remotely, zero trust was the only approach for us." Prabash Coswatte, Chief Operating Officer at Cardenas Markets LLC.

Darktrace DETECT and RESPOND complements and enhances zero trust postures with Self-Learning AI that identifies, interrupts, and investigates unpredictable cyber-threats that get through, even if they operate over legitimate paths. This includes advanced external attacks like ransomware, zero-days, and supply chain risks, as well as compromised, careless, or malicious insiders with privileged access. Ultimately, the zero trust attitude is to distrust implicitly. An AI that is designed to constantly ask the question: "Is this normal?" is the perfect partner.

"Darktrace enhances zero trust postures by leveraging its Self-Learning AI and distinguishing normal and abnormal behavior. Darktrace helps us understand what normal looks like and alerts us when we have any external threats or insider behaviors that are out of the norm." Prabash Coswatte, Chief Operating Officer at Cardenas Markets LLC.

By deploying Darktrace alongside a robust zero trust architecture, organizations benefit from a layered security strategy that combines a protective posture of 'default deny' with autonomous smart systems that adapt as the business and workforce evolve, leaving attackers with nowhere to hide.

"Traditional zero trust tools are static. Threat actors know this and have many tools at their disposal to counter them. Darktrace AI's ability to automatically detect and respond to incidents reduces triage time and frees up my cyber security team. We're pleasantly surprised by the visibility it provides us on non-malicious insider activity. Visibility to some of these non-malicious insider activity helps us prevent larger issues from occurring later on. (data loss, unusual login locations, file deletions...)" Prabash Coswatte, Chief Operating Officer at Cardenas Markets LLC.

Establishing and maintaining a robust zero trust architecture across these varied environments requires time, resources, and money. As with many other companies worldwide, HGG has had to quickly expand its digital footprint. The demands of our current fast-paced economy increase the dependence on interconnected technologies that aid company productivity, often at the cost of security. The challenge is to strike a balance between an implicit distrust and enhanced productivity.

"It's not humanly possible to keep up with the demands of a fully integrated zero trust architecture. We are not in the business of security, we're in the grocery business. IT security is simply the cost of doing business in today's environment. Darktrace's Self-Learning AI helps us keep our human capital interments in IT security low while giving us the peace of mind that the AI is constantly monitoring, adapting, and dynamically scrutinizing all network, user and device behaviors. Darktrace integrates with HGG's zero trust protective skin to detect and responds to threats in an adaptive and continuous manner." Prabash Coswatte, Chief Operating Officer at Cardenas Markets LLC.

Ultimately, DETECT provides the security team at HGG with visibility into activity in their network, email, and Microsoft 365 which then informs RESPOND's autonomous actions within each environment, making sure that the security team is augmented with the AI's insights and speed. Instead of having to manually monitor or generate rules and playbooks that attempt to curtail user behavior, as traditional network tools would require, Darktrace's Self-Learning AI can mold itself to user and device behavior to alert and act upon those anomalies without any demand of your security team's time and maintaining your zero trust posture.

About Darktrace

Darktrace (DARK.L), a global leader in cyber security artificial intelligence, delivers complete AI-powered solutions in its mission to free the world of cyber disruption. Its technology continuously learns and updates its knowledge of 'you' for an organization and applies that understanding to achieve an optimal state of cyber security. It is delivering the first ever Cyber AI Loop, fueling a continuous end-to-end security capability that can autonomously prevent, detect, and respond to novel, in-progress threats in real time. Darktrace employs over 2,200 people around the world and protects over 8,100 organizations globally from advanced cyber-threats.

Scan to LEARN MORE

DARKTRACE
Evolving threats call for evolved thinking™

North America: +1 (415) 229 9100
Europe: +44 (0) 1223 394 100
Asia-Pacific: +65 6804 5010
Latin America: +55 11 97242 2011
info@darktrace.com
darktrace.com