DARKTRACE

Cyber Resilience for OT Systems:

Darktrace and NIST SP800-160 Volume 2

Introduction

Contents

"Fortifying the network perimeter is simply not enough, and you will get caught out. They are already inside. It's time to lift up the stone and uncover the creepy-crawlies blithely circulating beneath."

Dave Palmer, Chief Product Officer, Darktrace

The NIST publication SP800-160 Volume 2 provides guidance for achieving cyber resilience in OT systems. A recent revision endorses the following approach:

"[The publication] turns the traditional perimeter defense strategy on its head and moves organizations toward a cyber resiliency strategy that facilitates defending systems from the inside out instead of from the outside in. This guidance helps organizations anticipate, withstand, recover from, and adapt to adverse conditions, stresses, or compromises on systems – including hostile and increasingly destructive cyber attacks from nation states, criminal gangs, and disgruntled individuals."

— NIST SP800-160 Volume 2

Darktrace's Industrial Immune System has been putting this philosophy into effective practice since 2015. The formalization of resilience by NIST is a welcome development in government guidance for OT cyber security. This whitepaper examines the core ideas developed in the publication and relates them to the observed threat landscape for OT, as well as Darktrace's technical philosophy and solutions. It then shows Darktrace's practical benefits with real-world detections and mitigation.

The Cyber Resilience Perspective on Cyber Security

Recent high-profile attacks against OT businesses have fallen into two broad categories. First, rapidly evolving and rapidly moving ransomware with a very short attack lifecycle. Second, and representing even greater risks, Advanced Persistent Threats (APTs) targeting OT organizations — especially those which gain entry to the core OT networks within. This latter category has a history of publicly known campaigns such as Triton and Industroyer, likely with more attacks that have never been revealed or remain undetected.

“ADVERSARY PERSISTENCE AND LONG-TERM PRESENCE

Numerous reports of cyber incidents and cyber breaches indicate that extended periods of time transpired between the time an adversary initially established a presence in an organizational system by exploiting a vulnerability and when that presence was revealed or detected. In certain instances, the time periods before detection can be as long as months or years. In the worst case, the adversary’s presence may never be detected.”

— NIST SP800-160 Volume

A cyber resilience approach begins with the understanding that not all threats can be prevented from entering an organization, and, moreover, that attacks might not be easily recognisable as such. Highly resourced threat actors have proven more than capable of evading border defences and traditional internal detection methods.

Cyber resiliency as considered by NIST primarily addresses this slower-moving APT type threat scenario. However, it does also mention the main practical requirement to extend protections to faster moving threats as well as outright prevention. Notably, these fast-moving threats can be neutralized in their early stages at machine speeds with autonomous response technology, containing them at point of impact.

“THREAT DETECTION AND CYBER RESILIENCY

Cyber resiliency is based on the recognition that adversaries can establish and maintain a covert presence in systems. Therefore, many of the cyber resiliency techniques and approaches are not predicated on the assumption of successfully detecting adversity including cyber attacks.

Other techniques and approaches can provide automatic response—or can support cyber defender responses—to detected indicators of possible or suspected adversity, or to warn of potential forthcoming adverse conditions (including predictions of increased system load or announcements of planned outages of supporting services).

Two cyber resiliency techniques directly involve the detection of adversity or its effects.”

— NIST SP800-160 Volume 2

Darktrace's Antigena technology, which actions autonomous response to emerging threats, is a direct implementation of one of the techniques described above, as are Darktrace's Immune System detection methods for identifying novel and unknown threats. For these purposes, the use of AI is a practical necessity: while simple methods of automation can streamline traditionally human-enacted workflows and make them faster, they cannot fundamentally add anything new. By contrast, Darktrace's Self-Learning AI—the core technology of both Antigena and the Immune System—does not merely automate legacy security methods, but provides a novel approach to detection, investigation, and response that allows organizations to achieve autonomous cyber defence.

The decade of the 2010's demonstrated repeatedly that most organizations, using the legacy security technologies that were widespread at the time, could not defend themselves against novel and fast-moving attacks or against APTs. This holds in both Enterprise (IT, SaaS, cloud, email) as well as Industrial (OT) environments. And up to this day, organizations across all industry verticals are still repeatedly falling victim to sophisticated cyber-attacks.

Figure 1: Darktrace's Self-Learning AI Closed Loop Diagram

The Industrial Immune System

Darktrace and the NIST publication share a direct analogy between cyber security and the human immune system:

“CYBER-RESILIENT SYSTEMS

Cyber-resilient systems operate somewhat like the human body. The human body has a powerful immune system that absorbs a constant barrage of environmental hazards and provides the necessary defense mechanisms to maintain a healthy state. The human body also has self-repair systems to recover from illnesses and injuries when defenses are breached. But cyber-resilient systems, like the human body, cannot defend against all hazards at all times. While the body cannot always recover to the same state of health as before an injury or illness, it can adapt. Similarly, cyber-resilient systems can recover minimal essential functionality. Understanding the limitations of individuals, organizations, and engineered systems is fundamental to managing risk.”

— NIST SP800-160 Volume 2

Darktrace’s earliest solutions focused entirely on turning this philosophy into cutting-edge technology that is used for the detection of threats. Further, its platform development over time has added new capabilities that align with even more of the NIST resilience concepts.

> “We need to start implementing a cyber “immune system” that learns from its environment to avoid recurring problems and combat new ones. Technologies such as the Immune System work on probabilities and experience, rather than hard-and-fast rules and certainties.”

Nicole Eagan, Chief Strategy Officer and AI Officer, Darktrace in the World Economic Forum

Darktrace first expanded its platform by adding autonomous response capabilities with Antigena. This technology allows organizations to quickly recover from an emerging attack with an AI-driven response that is mathematically precise and occurs in seconds. The speed and extent of recovery depends on how aggressively the threat advances and what specific actions are required to control it. In all scenarios, however, Antigena's autonomous response gives organizations a fighting chance at a time when threats are becoming increasingly automated, and Antigena often contains a threat before it escalates into crisis.

Even if a threat cannot be directly recognized as a threat per se, Darktrace's Immune System focuses on distinguishing changes that are not within its understanding of "self" from the continual wanted changes to the environment. Darktrace understands "self" for devices, users, and the organization as-a-whole. Indeed, many unwanted changes can take place without any intentional malice or active threats being present; for example, a user accidentally transfers sensitive data to an unsecure destination. These forms of unusual behavior that are not intentionally malicious are no less important to find and mitigate, as they can open the door to future threat vectors for attackers.

Darktrace has also introduced Cyber AI Analyst, which reduces the time-to-meaning of the Immune Systems' detections, while at the same time bringing in cyber expertise that the local security analyst might not have. This latter point is not uncommon for OT cyber security, where OT cyber analysts are unlikely to be experts in IT cyber security. In many cases, people overseeing OT security are engineers without extensive training even on the OT side.

The major APTs mentioned earlier were (of course) hybrid IT-OT attacks that blended expertise in both domains. In fact, most sophisticated attacks that disrupt OT are IT-OT, meaning that it is necessary to both close the knowledge gap for people overseeing OT and provide unified protection of both IT and OT. Darktrace's platform allows organizations to achieve these security goals toward which every organization seeking truly robust industrial security should strive.

Cyber Resiliency Objectives

Cyber resiliency objectives are more specific statements of what a system must achieve in its operational environment and throughout its life cycle to meet stakeholder needs for mission assurance and resilient security.

“PREVENT OR AVOID: Preclude the successful execution of an attack or the realization of adverse conditions.

PREPARE: Maintain a set of realistic courses of action that address predicted or anticipated adversity.

CONTINUE: Maximize the duration and viability of essential mission or business functions during adversity.

CONSTRAIN: Limit damage from adversity.

RECONSTITUTE: Restore as much mission or business functionality as possible after adversity.

UNDERSTAND: Maintain useful representations of mission and business dependencies and the status of resources with respect to possible adversity.

TRANSFORM: Modify mission or business functions and supporting processes to handle adversity and address environmental changes more effectively.

RE-ARCHITECT: Modify architectures to handle adversity and address environmental changes more effectively.”

— NIST SP800-160 Volume 2

Compared to a traditional cyber security approach, prevention and avoidance do not dominate the objectives of the resilience objectives. Many of these objectives benefit enormously from a real-time true view of the business’ information activities, displaying network topology according to actual activity and use of any critical external resources such as SaaS. Darktrace provides these capabilities in full. The objective to limit damage from adversity (“constrain”) also obviously benefits greatly from AI-driven investigations and autonomous response technology, both of which empower teams to mitigate damage as an attack emerges.

Detection is also not one of the objectives. However, in practice, detection has an enormous effect on whether and how much resilience is required, that is, should detection be delayed or missing entirely. This makes sense, as NIST likely wanted to create a neater separation between the guidance for resilience and other NIST industrial publications and controls (e.g., SP800-53). At the same time, detection must be factored into any evaluation of the expected effectiveness of resilience objectives.

Cyber Resiliency Techniques and Approaches

The following 14 techniques are part of the cyber resiliency engineering framework. Here is an overview of how Darktrace's solutions correspond with each technique; more detailed explanations of each correspondence can be provided upon request.

1. **Adaptive Response: Implement agile courses of action to manage risks.**

   Darktrace's autonomous response technology, Antigena, can react at machine speed to detections of potentially unwanted activity, that is, without needing to confirm that the activity corresponds to a known or defined threat.

2. **Analytic Monitoring: Monitor and analyze a wide range of properties and behaviors on an ongoing basis and in a coordinated way.**

   The Industrial Immune System achieves this ongoing coordinated monitoring and analysis of a wide range of properties and behaviors. It provides the additional benefit that its analysis is based in AI, is real-time, and can detect novel threats and potentially unwanted activity. Further, Cyber AI Analyst delivers a fully investigated report at machine speed and based on significant OT and IT expertise.

3. **Contextual Awareness: Construct and maintain current representations of the posture of missions or business functions considering threat events and courses of action.**

   Darktrace's Threat Visualizer interface has multiple methods of enhancing awareness including but not limited to the following: asset inventory (passive and optional active); visualization of real-time network activity and user credentials; alert prioritisation; and well-explained reports from Cyber AI Analyst.

4. **Coordinated Protection: Ensure that protection mechanisms operate in a coordinated and effective manner.**

   To provide coordinated protection, Darktrace can integrate with other tools in multiple ways, from delivering alerts and reports to a central SIEM, to comparing asset data with peer technologies, to ingesting logs or alerts from other security tools. Moreover, Darktrace's Antigena components can drive firewalls as part of autonomous response to possible threats.

5. **Deception: Mislead, confuse, hide critical assets from, or expose covertly tainted assets to the adversary.**

   Monitoring of tainted assets or other deliberate security sinks can be performed by the Industrial Immune System.

6. **Diversity: Use heterogeneity to minimize common mode failures, particularly threat events exploiting common vulnerabilities.**

   The Industrial Immune System enters an OT network with no pre-conceived ideas as to what it will see, nor does it assume what certain devices "should" do based on its observance elsewhere. Rather, it learns on-the-job what the unique network normally does and then continues to evolve its understanding as the network changes over time. As a result of this, heterogeneity is inherently handled properly. Differences are not problematic for monitoring, and what commonality does exist is automatically learned and made use of by the AI.

7. **Dynamic Positioning: Distribute and dynamically relocate functionality or system resources.**

   Visibility of where resources currently exist and how they are interacted with can be achieved with the Threat Visualizer interface. Dynamic changes are automatically tracked, and the understanding of the Industrial Immune System's AI also evolves along with them.

8. **Non-Persistence: Generate and retain resources as needed or for a limited time.**

   The bane of manually updated security systems, the Industrial Immune System requires no user input to evolve its understanding alongside changes made to the network. This means that rapid changes and non-persistent existences do not impose extra workloads.

9. **Privilege Restriction: Restrict privileges based on attributes of users and system elements, as well as on environmental factors.**

   The Industrial Immune System models the activities of users (via credentials) as well as devices and groups of devices that it learns are similar. Use of unintended privileges (or maliciously gained privileges) creates unusual activity that Darktrace actively monitors.

10. **Realignment: Structure systems and resource uses to align with mission or business function needs, reduce current and anticipated risks, and accommodate the evolution of technical, operational, and threat environments.**

    This benefits from the same automatic evolution of the security monitoring system described for Non-Persistence. The Industrial Immune System is not a past-signature or prior-intelligence based system; instead, it identifies risks based on its understanding of the business' unique environment and highlights what does not fit in. It has been proven many times that this can find true zero-day and novel attacks as they emerge. The Industrial Immune System comprises a platform encompassing OT networks and related parts of IT networks; it also encompasses cloud services, endpoint, and other areas of both current and future OT-related digital estates. The Industrial Immune System is thus ready to monitor different and growing technologies as well as ongoing changes to current technologies.

11. **Redundancy: Provide multiple protected instances of critical resources.**

    Similar to Dynamic Positioning, visibility of the live environment is extremely useful and so are monitoring the redundant resources for signs of reliability problems. The outputs of an Immune System concept detect possibly unwanted changes regardless of whether they were caused by intentional threat.

12. **Segmentation: Define and separate system elements based on criticality and trustworthiness.**

    The Industrial Immune System provides visualisations of network topology, including ways to automatically specify criticality or other characteristics of systems and view how these separate system elements interact with each other.

13. **Substantiated Integrity: Ascertain whether critical system elements have been corrupted.**

    The Immune System's detection of unknown threats makes use of information about input interactions with critical systems as well as subsequent activities by those systems. Corruption can come from unintentional and non-malicious sources, and the Immune System's approach is ideal to find these as well as targeted threats.

14. **Unpredictability: Make changes randomly or unpredictably.**

    While it cannot perform these actions itself, the Industrial Immune System's visibility can be used to choose where to make changes that would most heavily impact an adversary. It also evolves along with changes to the environment it observes, and so there is no requirement to manually input "the same changes" into the security system to match.

"Using machine learning, Darktrace detects zero-day threats and suspicious insider behaviors, without having to define the activity in advance."

Curator of Disruptive Innovation, City of Las Vegas

"Darktrace has helped us to keep our networks and devices honest about what they should be doing."

Lead SCADA Analyst, City of College Station Utilities

Adaptive Response

- Dynamic Reconfiguration
- Dynamic Resource Allocation
- Adaptive Management

Analytic Monitoring

- Monitoring & Damage Assessment
- Sensor Fusion & Analysis
- Forensic & Behavioral Analysis

Coordinated Protection

- Calibrated Defense-in-Depth
- Consistency Analysis
- Orchestration
- Self-Challenge

Contextual Awareness

- Dynamic Resource Awareness
- Dynamic Threat Awareness
- Mission Dependency & Status Visualization

Deception

- Obfuscation
- Disinformation
- Misdirection
- Tainting

Diversity

- Architectural Diversity
- Design Diversity
- Synthetic Diversity
- Information Diversity
- Path Diversity
- Supply Chain Diversity

Dynamic Positioning

- Functional Relocation of Sensors
- Functional Relocation of Cyber Resources
- Asset Mobility
- Fragmentation
- Distributed Functionality

Non-Persistence

- Non-Persistent Information
- Non-Persistent Services
- Non-Persistent Connectivity

Privilege Restriction

- Trust-Based Privilege Management
- Attribute-Based Usage Restriction
- Dynamic Privileges

Realignment

- Purposing
- Offloading
- Restriction
- Replacement
- Specialization
- Evolvability

Redundancy

- Protected Backup & Restore
- Surplus Capacity
- Replication

Segmentation

- Predefined Segmentation
- Dynamic Segmentation & Isolation

Substantiated Integrity

- Integrity Checks
- Provenance Tracking
- Behavioral Validation

Unpredictability

- Temporal Unpredictability
- Contextual Unpredictability

Resiliency in Action

To demonstrate the effectiveness of the Industrial Immune System in practice, two real examples from anonymized users of the platform are provided below.

The first focuses on how to recognize that threatening behavior is occurring, even when one does not have any prior knowledge of what the threat is and cannot name it. Zero-day ransomware is fast-moving, highly damaging, and usually makes heavy use of encrypted network connections that easily bypass or confuse simple security approaches. However, this example also demonstrates that a security process reliant on human intervention cannot be left unattended as, despite all the early and accurate detections made by the AI, the threat made a lot of progress due to a lack of human attention.

The second example, by contrast, shows the scenario where autonomous response technology (Antigena) is enabled. Here, the attack is brought swiftly to a halt. While in this case the attack was investigated afterwards, there was actually no need to ever know the details of what was blocked.

Example 1: Defending Critical Infrastructure from an Unknown Double-Extortion Ransomware Strain

Darktrace detected every step of an attack as it unfolded over the course of 12 hours against a North American company in the electric grid supply chain. Unfortunately, nobody was watching Darktrace, and autonomous response was not deployed in active mode. And so, no action was taken until the following morning when incident response and remediation began.

The attacker gained entry via an Internet-facing vulnerability and escalated their privileges to admin. Darktrace detection: New Admin Credential on Client.

The attacker then used encrypted connections to download, install, and initiate a common remote management tool. Darktrace detection: Remote Management Tool on Server.

After this, the attacker exploited the Windows file sharing protocol, SMB, to download GBs of sensitive data and exfiltrate it to a public cloud sharing platform, pcloud, using encrypted HTTPS. Darktrace detection: Uncommon 1 GB Outbound.

The attacker then deployed and executed an unknown ransomware strain using administrative Windows tools and encrypted 1000's files including back-ups. Darktrace detection: Sustained MIME Type Conversion.

In total Darktrace produced 23 alerts for the device in question, or 48% of all the alerts produced in the corresponding 24-hour period. And so, the incident stuck out like a sore thumb. If the customer had been using Antigena for autonomous response, the activity would have been neutralized before significant volumes of data were exfiltrated or encrypted.

The customer, however, was able to use the Darktrace Ask the Expert (ATE) service for incident response to mitigate the impact of the attack and aid with disaster recovery. The insights provided by Darktrace helped minimize the damage caused by the attack.

Unaffected systems, including the OT production networks, remained online, while infected systems were isolated to prevent further spread of ransomware. Lists of the files exfiltrated or encrypted were able to be audited and assessed for business risk, and the remote access backdoors used by the attackers were identified and removed to prevent a repeat of the attack.

The customer was able to return to full operations in a relatively short time, notably, without the level of disruption or reputational damage seen in the Colonial Pipeline and JBS ransomware events, which either directly or indirectly lead to shut down of OT systems.

Read the full blog here >

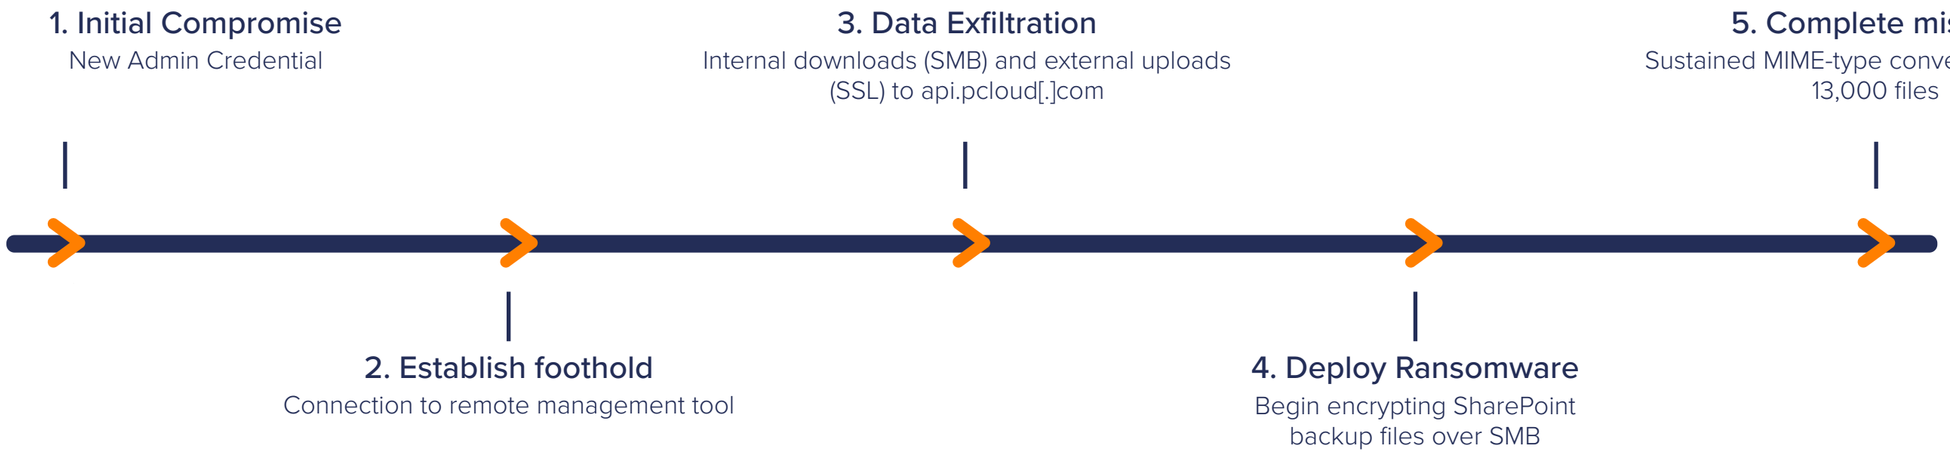

Figure 2: Timeline of Unknown Ransomware Attack

Example 2: How Antigena Autonomous Response Neutralized Zero-Day Ransomware

Darktrace Antigena stopped a previously unknown ‘zero-day’ ransomware attack targeting an electronics manufacturer. This strain of ransomware was not associated with any publicly known indicators of compromise. However, Darktrace was able to detect and autonomously respond to the threat, neutralizing it before it could do damage, illustrating how Darktrace takes a fundamentally assumptive approach, not relying on rules or signatures.

Darktrace was able to detect this never-before-seen attack based purely on its comprehensive understanding of the normal pattern of life for every device and user within the organization. Darktrace’s AI identified a spike in the pattern of regular connections made by patient zero and a series of high-confidence alerts firing in quick succession. These included:

1. Compromise / Ransomware / Suspicious SMB Activity — triggers when a device begins making unusual SMB connections across the organization
2. Antigena Ransomware Block — triggers Antigena to take an action when the behavior is significantly similar to ransomware
3. Device / Reverse DNS Sweep — triggers when a device makes unusual reverse DNS lookups, a tactic often used during reconnaissance

Antigena was in Active Mode, and so it enforced the usual pattern of life by blocking anomalous connections for five minutes, immediately stopping the encryption. This successfully neutralized the threat.

To contain the threat at point of impact, Antigena then stopped the ransomware from spreading by quarantining patient zero for 24 hours, rendering the device unable to connect to the server or any other device on the network.

In this way, Antigena’s autonomous response successfully stopped encryption and prevented further lateral spread that could occur by scanning, using harvested admin credentials, or performing internal reconnaissance. Thus, Antigena surgically contained the threat while allowing normal business operations to continue as usual, effectively mitigating damage while maintaining business continuity

Read the full blog here >

.

Conclusion

“APPENDIX C.2 DISTINGUISHING CHARACTERISTICS OF CYBER RESILIENCY

Any discussion of cyber resiliency is distinguished by its focus and a priori threat assumptions. These are reflected in cyber resiliency constructs and engineering practices:

Focus on the mission or business functions.

Assume a changing environment.

Focus on the effects of the advanced persistent threat.

Assume the adversary will compromise or breach the system or organization.

Assume the adversary will maintain a presence in the system or organization.”

— NIST SP800-160 Volume 2

The Industrial Immune System implements many of the techniques involved in cyber resilience in concert with detection methods that are also based on the fundamental concept of a human immune system. Indeed, the original aspects of the technology’s design included all the distinguishing characteristics of resiliency mentioned in this NIST publication.

Deploying Darktrace’s cyber security platform is a direct way to introduce these concepts into new or existing cyber security operations. The solution has proven capable of detecting unknown novel threats as they emerge, and where its autonomous response components are active, it has been shown to defeat these threats without ever having to know what they were.

Organizations seeking to achieve truly resilient cyber defence will thus find a robust solution in Darktrace’s platform. This is why organizations across all 16 critical infrastructure sectors defined by CISA actively rely upon Darktrace’s AI to safeguard their mission critical assets.

About Darktrace

Darktrace (DARK:L), a global leader in cyber security AI, delivers world-class technology that protects over 6,500 customers worldwide from advanced threats, including ransomware and cloud and SaaS attacks. Darktrace's fundamentally different approach applies Self-Learning AI to enable machines to understand the business in order to autonomously defend it. Headquartered in Cambridge, UK, the company has 1,700 employees and over 30 offices worldwide. Darktrace was named one of TIME magazine's 'Most Influential Companies' for 2021.

For More Information

- Visit darktrace.com
- Book a demo
- Visit our YouTube channel
- Follow us on Twitter
- Follow us on LinkedIn