June 2025

ICIT

The Digital Immune System:
How AI Can Outpace Cyber Threats

Jim Routh
Fellow, Institute for Critical Infrastructure Technology

www.icitech.org

Table of **Contents**

About ICIT

The Institute for Critical Infrastructure Technology (ICIT) is a nonprofit, nonpartisan, 501(c)3 think tank with the mission of modernizing, securing, and making resilient critical infrastructure that provides for people's foundational needs. ICIT takes no institutional positions on policy matters. Rather than advocate, ICIT is dedicated to being a resource for the organizations and communities that share our mission. The views and opinions expressed in this essay are solely those of the author(s) and do not necessarily reflect the official policy or position of ICIT. Any assumptions made within the analysis are not reflective of the position of any entity other than the author(s).

To learn more, please visit www.icitech.org

CyberRisk ALLIANCE | Thank you to our Strategic Partner **CyberRisk Alliance** | cyberriskalliance.com

Introduction

Cyberattacks can now compromise critical infrastructure faster than humans can perceive or respond. As adversaries grow more sophisticated and the cost of disruption rises, human-led defenses alone are no longer sufficient. We must re-engineer cybersecurity to match the speed and complexity of modern threats, starting with core operations.

We can build an enterprise digital immune system by using models that adapt in real time, like the human immune system. This system augments Security Operations Centers (SOCs) by autonomously detecting and defending against active threats in milliseconds. The result: greater productivity, lower costs, and minimized business impact.

With AI agents advancing rapidly, the moment to shift toward a proactive, adaptive model of cyber resilience is now.

The Immune System Approach

Our bodies—divided across levels of organization that build on each other—are not unlike many of our networks. Not only do they share many interconnected systems that depend on constant communication to complete critical functions, but they also need to be defended from foreign invaders.

In the body, our immune system protects us against external threats. It does so by recognizing the presence of a pathogen through pattern recognition receptors on immune cells, and in turn triggering a response to neutralize the threat. Moreover, this process occurs automatically and offloads most of the day-to-day blocking of harmful intrusions from our conscious mind, only escalating (through fever or pain, for example) when a real threat is present.

Security Operations Centers (SOCs) play a similar protective role in the context of our networks. However, they don't share many advantages with our immune system. Most SOCs today are overwhelmed with alerts. Similar to how allergies impact our bodies, SOCs are inundated with false positive indicators, which consume time and attention that would otherwise be best applied to more critical events and incidents. This is because—unlike our immune system—SOCs do not have automatic response mechanisms on which we can rely to protect us.

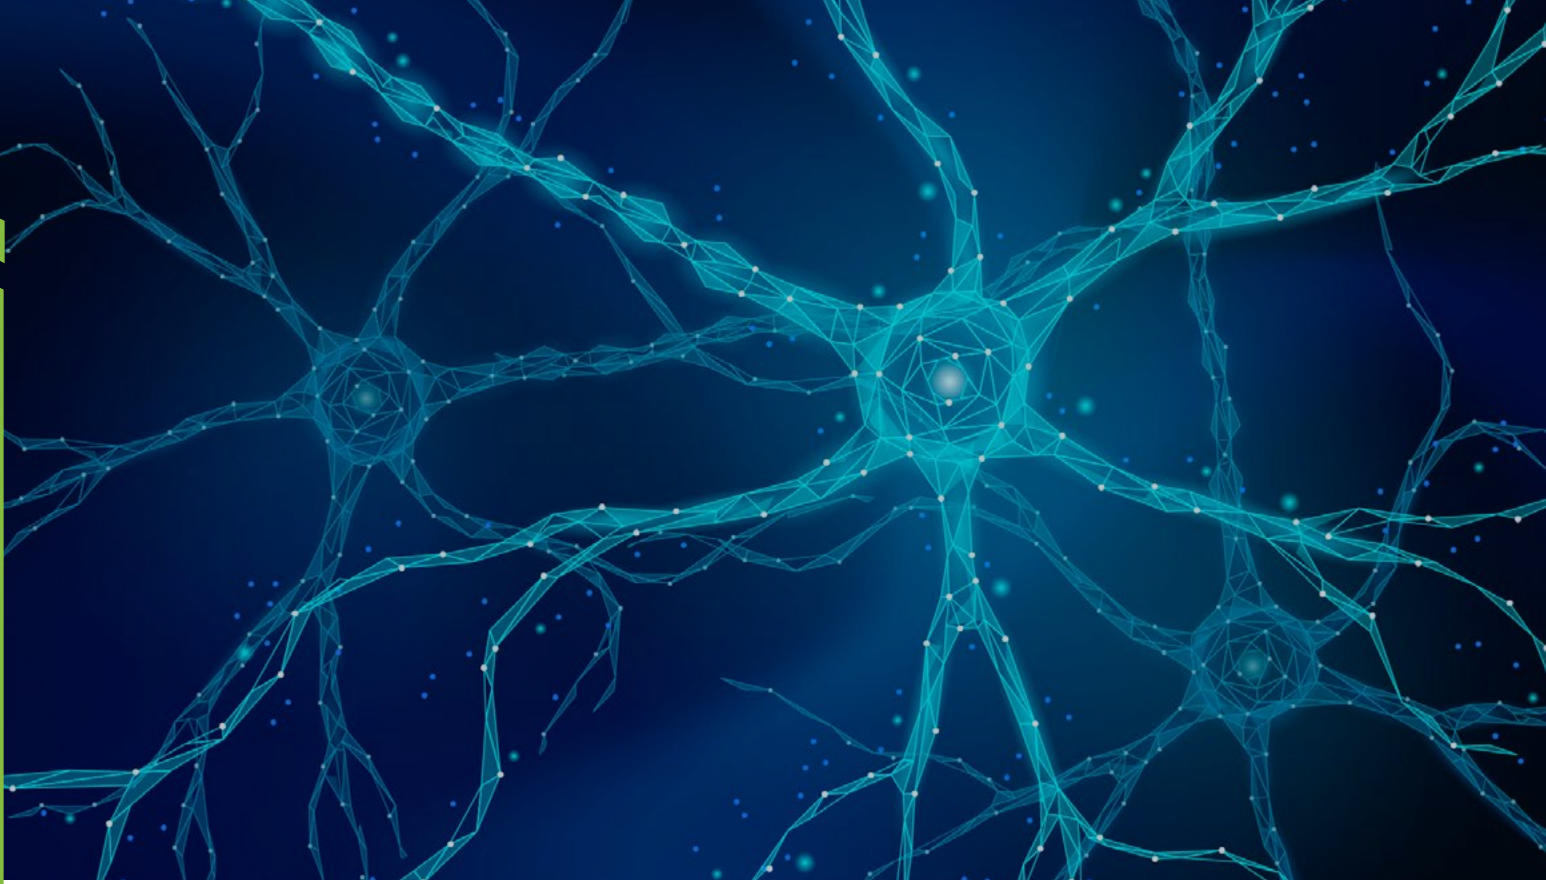

The Digital Immune System

It's time for enterprises to deploy a digital immune system that operates in real-time that protects customers, employees, and other stakeholders against cyberattacks without relying exclusively on humans to initiate protective measures. The need for this change to how we approach cybersecurity operations is highlighted by three compelling drivers:

Humans are Slower than Models & Machines

Detecting and protecting an enterprise from cyberattacks that occur in milliseconds isn't feasible when humans are required to process information, apply context, and then take appropriate action for all transactions (events, incidents, workflows).

Machine Learning Enables Real-Time Response

Using machine learning (ML) systems with deterministic models that trigger based on pattern deviations enables real-time response. This results in a significantly lower unit cost per transaction.

Large Language Models Improve Accuracy

Large language models (LLMs) and AI agents are growing in availability and maturity for reasoning in inference, in turn improving the accuracy of results.

How to Implement a Model: The Basics

What would this "digital immune system" (DIS) look like? Let's consider a relatable example.

Assume you order a Latte at your local coffee shop every day. You indulge in a Caramel Macchiato on weekends and non-workdays as a special treat. Using the DIS approach, a simple model with a few variables could be deployed to place your coffee order daily, requiring no more effort than picking up your drink at the coffee shop.

In this straightforward scenario, a deterministic model is particularly effective at delivering a consistent result, especially considering the limited attributes involved:

1. **Is it a workday or a non-workday?**
2. **Is it the right time for your coffee order?**

The model ensures that your order is placed exactly when you arrive at the coffee shop in the morning by triggering an automated workflow. This gives you some time in your day to focus on other tasks without sacrificing the convenience of having your preferred coffee ready for pickup at your convenience.

What if your coffee preferences shift over time? You might suddenly crave a Pumpkin Spiced Latte in autumn, instead of your usual Latte. This presents a challenge, as you may need to visit a different coffee shop to satisfy your new craving. You could revert to your old manual ordering system or use a more advanced model that considers more factors:

1. **The change in season (calculated based on weather variables and the calendar)**
2. **The availability of Pumpkin Spiced Lattes at your local coffee shops (monitored through advertisements and inventory checks)**
3. **The specific locations of the new coffee shops (as some shops may not carry it)**

In this new scenario, the existing deterministic model is no longer sufficient to handle the increased complexity of these additional attributes, their varying weights, and the substantial volume of information required. This now becomes an ideal task for generative AI, which can consider multiple variables with different weightings to automate the task of placing the order.

Just as you could choose between different models in the example depending on your coffee needs and preferences, enterprises today have choices for applying models and LLMs in more sophisticated ways. Particularly, they can leverage various models to handle time-consuming and tedious information collection and correlation tasks in seconds to dramatically increase mean time to respond (MTTR).

As demonstrated in the example, the implementation of a Digital Immune System for cyber operations starts with determining what tasks and outcomes can be performed without a human being. Although a human could remain in the loop, they would only be required to act when convenient and necessary. In all other situations, the DIS would operate autonomously, retaining a near-instantaneous response capability and becoming as comprehensive as its processing demands would require. This DIS would be empowered today by models that recognize and respond to pattern deviation, triggering automated workflows in milliseconds.

Applying the Model: PAM Example

Let's focus on a more practical example: enhancing our privileged access management (PAM) capability to monitor the real-time online behavior of every privileged user when they have entitlement access. By doing so, we can detect and respond to potential security threats promptly.

Here's how it works and how I have seen it be successfully implemented: when a privileged user logs in and accesses sensitive systems, their online behavior is continuously monitored in real-time. This includes attributes such as login location, device type, and browsing habits. The system then compares this real-time data to the user's previously established pattern, which is stored in a baseline profile.

If the user's online behavior deviates significantly from their normal pattern, their entitlement access is automatically revoked. This is a critical capability in preventing ransomware-as-a-service (RaaS) attacks, which can spread rapidly through privileged accounts. By monitoring user behavior in real-time, we can identify and isolate potential threats before they cause harm.

The system generates a security event notification, which is sent to the SOC analyst for review. However, the analyst is not required to take immediate action, as the model has already taken the necessary steps to protect the organization. This approach enables the enterprise to build a Digital Immune System that can detect and respond to threats in real-time, reducing the risk of data breaches and cyberattacks.

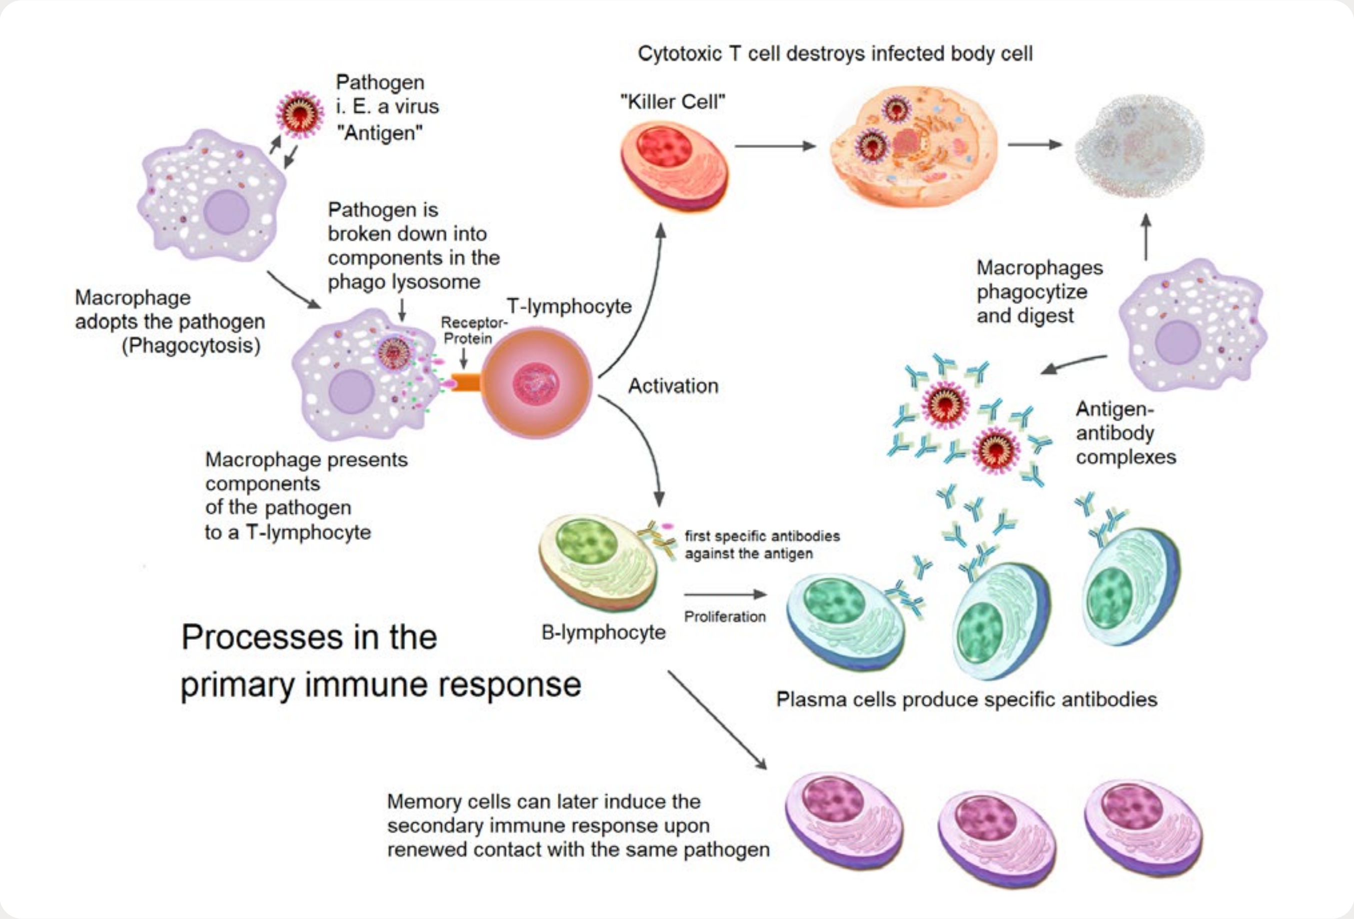

A DIS mirrors a human immune system with models triggering automated workflows from pattern deviation.

Benefits of an AI-Driven Immune System

There are three key benefits to incorporating Large Language Models (LLMs) or agents into SOC operations:

1. **Productivity Gain**: LLMs reduce the time analysts spend collecting and analyzing data across disparate systems.
2. **Cost Savings**: When humans are required to contextualize and act on every alert, each incident carries a high transaction cost. Automated workflows can reduce or eliminate this burden.
3. **Business Impact Reduction**: AI agents operating in milliseconds can detect and contain threats like ransomware before significant damage occurs.

The benefits of incorporating generative AI into the Security Operations Center (SOC) operations are multifaceted. Let's examine them in more detail.

Productivity Gain

Generative AI grants SOC analysts' additive capacity on higher-value activities such as threat hunting. This requires analysts to allocate their time efficiently and prioritize tasks that require human expertise. The aggregate benefit across all analysts cumulates in more compelling Key Performance Indicators (KPIs), such as increased threat detection rates and reduced mean time to detection (MTTD).

Cost Savings

When existing workflows require SOC analysts to understand context before acting on an event/incident, the transaction (event/incident closure) incurs a high cost (typically measured by the hourly rate multiplied by the number of hours spent completing the task). However, AI SOC agents can be leveraged in the following ways to eliminate the cost of the human in the transaction:

- **To gather contextual data and correlate information**
- **To form a conclusion and decide on a desired outcome**
- **To implement it in an automated workflow**

As a result, the unit cost per transaction decreases substantially, often by 50% or more. This means that the operating cost within the SOC for this type of event is half of that when using a SOC analyst exclusively.

Business Impact Reduction

Under the DIS approach, your enterprise can avoid the business impact and related incident response costs associated with a specific cyber-attack (such as a disruptive ransomware attack). The SOC can implement a continuous monitoring capability by leveraging models to compare actual data to established patterns and trigger automated workflows and agents (when necessary). The AI-augmented SOC can keep up with the flood of incoming alerts as AI SOC agents give every alert a consistently thorough investigation. With AI automation, the SOC will be able to quickly and automatically determine false positives from true positives. This, in turn, grants a more rewarding experience for SOC analysts and engineers, as their skills are applied to a higher level of analytic and design capabilities. Beyond empowering SOC staff to leverage their expertise to drive improved cyber resilience at a lower cost, this approach also reduces dependence on people. It creates a more fulfilling work environment for SOC professionals who don't have to stare at a screen full of alerts all day.

Human Immune System

The human immune system operates based on an unconscious response combined with a conscious response.

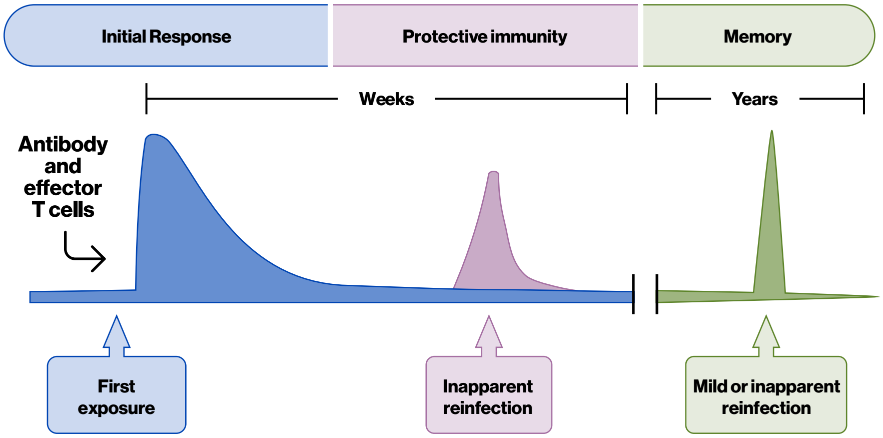

Digital Immune System (DIS)

The Digital Immune System operates in an automated response, followed by a human action.

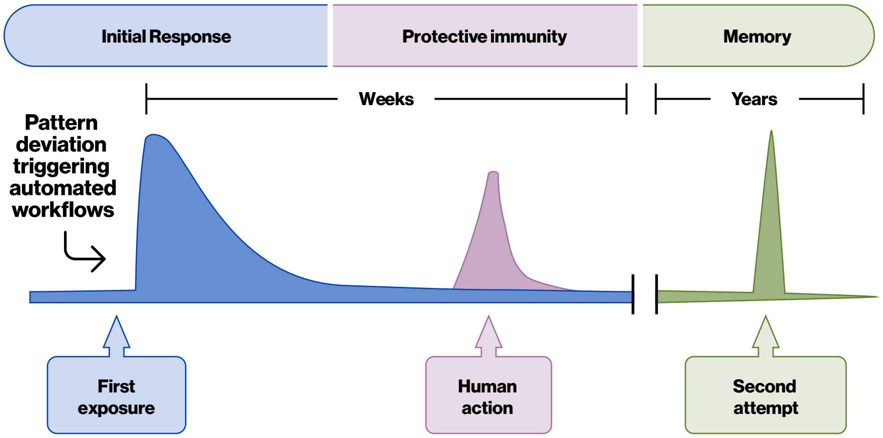

Doing More with Less

For Chief Information Security Officers (CISOs), this approach results in a desirable outcome that:

1. **Enables them to acquire the technology and models required to implement an Enterprise Immune System, and**
2. **Offsets the costs within a budget year by reducing operating costs.**

This allows CISOs to "do more with less," a mandate often shared with the Chief Information Officer's (CIO) objectives. The first step in deploying a Digital Immune System requires an appetite to change the way we work to be more efficient and effective. Cyber threats now operate at machine speed—your defenses should, too. Deploying a DIS isn't just a smart strategy; it's operational survival.

About

ICIT

The Institute for Critical Infrastructure Technology (ICIT) is a nonprofit, nonpartisan, 501(c)3 think tank with the mission of modernizing, securing, and making resilient critical infrastructure that provides for people's foundational needs. ICIT takes no institutional positions on policy matters. Rather than advocate, ICIT is dedicated to being a resource for the organizations and communities that share our mission. The views and opinions expressed in this essay are solely those of the author(s) and do not necessarily reflect the official policy or position of ICIT. Any assumptions made within the analysis are not reflective of the position of any entity other than the author(s).

To learn more, please visit www.icitech.org

Jim Routh

ICIT Fellow Jim Routh is a board member, advisor and investor with specific expertise as a transformational security leader focused on applying risk management discipline to a converged security function for global enterprises to achieve enterprise resilience. He has a demonstrated track record of designing security control using innovation and data science to align senior executives to deliver world-class level security capabilities to drive positive business results in a digital world.

CyberRisk Alliance

CyberRisk Alliance provides business intelligence that helps the cybersecurity ecosystem connect, share knowledge, accelerate careers, and make smarter and faster decisions. Through our trusted information brands, network of experts, and innovative events we provide cybersecurity professionals with actionable insights and act as a powerful extension of cybersecurity marketing teams. Our brands include SC Media, the OfficialCybersecurity Summits, TECHEXPO Top Secret, Security Weekly, InfoSec World, Identiverse, Cybersecurity Collaboration Forum, Cybersecurity Collaborative, ChannelE2E, MSSP Alert, LaunchTech Communications, CyberRisk TV and Execweb.

INSTITUTE FOR CRITICAL INFRASTRUCTURE TECHNOLOGY

ICIT

www.icitech.org