PROTECTING THE MAINFRAME IN THE AGE OF MYTHOS AND OTHER VARIANTS

The mainframe remains the transactional backbone of many of the world's largest organizations and the primary repository for “System of Record” data. While mainframe reliability provides an enormous business advantage, it can also create complacency. As artificial intelligence (AI) dramatically lowers the expertise and time required to identify vulnerabilities and develop attacks, organizations need to rethink assumptions about what has historically kept the mainframe secure. AI-based attacks have increased the concern for mainframe security and preparedness and should be a wake-up call for organizations that have overlooked investment into hardening the security and resilience of the critical mainframe platform.

Mythos and variants: A brief overview

As AI is advancing to perform deep-dive vulnerability analysis and create attack vectors based on those vulnerabilities, how can we protect our critical mainframe infrastructure? Mythos is an advanced AI model developed by Anthropic that can find and exploit software flaws better and faster than most humans. Access was provided to 12 companies as part of Project Glasswing, where the model successfully identified thousands of high-severity vulnerabilities, according to Anthropic. Recently, Mythos successfully “broke out” of a secured sandbox environment. Additionally, OpenAI admitted that it’s model “broke loose” and successfully attacked the AI company Hugging Face.

Autonomous AI-powered attack models are here. With Mythos and its variants becoming more

powerful, more capable, and more adaptable, no computing platform will be out of reach.

AI is breaking three long-held assumptions about mainframe security

Mythos and variants destroy 3 key assumptions that I see across some organizations that rely on mainframes:

- **Security by obscurity.** This is the belief that a small group of mainframe professionals understands the complexity of mainframe operating systems (z/OS, USS, VM, etc), therefore the likelihood of a successful attack on their organization's mainframe is very low. AI models are or will be able to consume the full documentation for mainframe operating systems and their subsystems, assess vulnerabilities in versions and configurations, then identify attack vectors that leverage those vulnerabilities, making these AI models more capable than most mainframe professionals.
- **Reliance on ESM (RACF, ACF2, Top Secret) logs.** Many organizations review updates to ESM profiles for users, resources, and privileges with the expectation that ESM controls are sufficient to detect security threats on the mainframe. As virtually every mainframe has vulnerabilities that are beyond the scope of ESM controls, those controls can be exploited to successfully attack mainframes.
- **Assuming network/perimeter defense protects the attack surface of the mainframe.** The limited number of organizations in receipt of Mythos includes Palo Alto Networks and Cisco—the network devices most often used to secure the network perimeters—where Mythos reported finding over 23,000 vulnerabilities from the limited set of companies. If the initial releases of AI-powered vulnerability assessment and attack planning are this effective, we can expect future releases to be even more threatening.

What can we do to protect mainframes?

The first thing that has to change amongst the mainframe community is the security mindset. Taking the following steps will help your organization harden its security against AI-based attacks:

- Employ Zero Trust on the mainframe. Assume that threat actors are on your mainframe who understand where the vulnerabilities are and will exploit them.
- Know your vulnerabilities before your adversaries know your vulnerabilities. Complete an independent comprehensive penetration test by mainframe-fluent professionals and develop a remediation plan to close the vulnerability gaps.
- Detect threats by deploying a comprehensive security monitoring solution that identify suspicious or malicious activity in real time. This should include behavior analytics to pinpoint activity that deviates from normal—a key defense capability for detecting AI-based activity. All significant security events should be sent to the organizational SIEM for enterprise visibility and appropriate response.
- Protect your data by implementing 1) periodic immutable backups that can not be changed, deleted or accessed by threat actors and 2) continuous immutable streaming of transaction recovery assets.
- Exercise recovery processes periodically, leveraging the immutable backups and transaction recovery steps to properly prepare in the event of a successful attack.

The reality of mainframe protection

Now at lightning speed, an AI-based attack mimics the steps that a well-trained mainframe hacker would perform: uncover vulnerabilities, identifying existing credentials, expanding capabilities through elevated privileges or increased access, and deploying programs that exfiltrate and/or encrypt data or incapacitate the platform. All of these activities are atypical for most users on a mainframe and detecting this atypical behavior is paramount. Numerous indicators of compromise, "honeypot" datasets to attract/detect threat actors, and behavior analytics with real-time notification to the SOC for high-priority security events are the best protection.

As change increases, agility becomes more critical

Security monitoring tools that require extensive programming, extensive consulting, or massive infrastructure to move millions of logs impact the ability to adapt to new mainframe security pressures. Modern mainframe security tools enable fast deployment of new indicators of compromise with behavior analytics, SIEM integration, and forensic analysis built into the foundation of the toolset. Forensic Analysis enables mainframe experts that have extensive platform experience to dive deep into behavior that looks suspicious. Don't underestimate your local experts, *empower* them.

Mythos and variants increase the likelihood of catastrophic attack

A ransomware attack could encrypt your data but leave the mainframe operational. A catastrophic attack would disable the mainframe from functioning, likely through the destruction of critical operating system datasets.

In each case, recovery plans that have been exercised and tested are critical to support the organization's return to normal business operations.

Summary

AI is changing the economics of cyberattacks by allowing adversaries to identify vulnerabilities, develop attack paths, and operate at a speed and scale that were previously impossible. For organizations that depend on the mainframe, that makes traditional assumptions about security increasingly risky.

The answer isn't fear of AI, it's eliminating the blind spots AI can exploit. That means understanding vulnerabilities before attackers do, detecting abnormal activity as it occurs, integrating mainframe security into the enterprise SOC, and maintaining a tested path to recover critical systems and data in the event that an attack succeeds.

The mainframe cannot remain a security island. When the platform supports the transactions and data that keep the business running, mainframe security is enterprise security.