DON'T LET "INHERENT MAINFRAME SECURITY" MAKE YOU COMPLACENT ABOUT MYTHOS

Anthropic's Claude Mythos has put mainframe security teams on notice. As a frontier AI model capable of autonomously identifying and chaining vulnerabilities at machine speed, Mythos represents a new class of risk for z/OS environments, dramatically accelerating the ability of adversaries to exploit any weakness present in the environment. While Mythos is the most visible example today, it represents a broader trend toward increasingly capable AI systems that are transforming vulnerability discovery and exploitation across the enterprise. Traditional assumptions about the inherent security of the mainframe overlook operational risks such as misconfigurations, expired certificates, unmonitored privileged sessions, and delayed threat detection. To be Mythos-ready, organizations need to take active, end-to-end measures to reduce their attack surface and detect threats in real time.

Why Claude Mythos poses a threat to the enterprise

By acting as an autonomous agent capable of independent, multi-step actions, Mythos can identify and link multiple security flaws into a programmatic chain of exploitation. As a result, time-to-exploit can collapse from weeks or months to hours or days, making vulnerability discovery and exploitation effectively simultaneous. These capabilities lower technical barriers and enable motivated actors with modest resources to perform multi-step exploitation that historically required elite specialists. Mythos has already found thousands of high-severity vulnerabilities across major operating systems and web browsers, including one vulnerability that had been present in a system for 27 years.

Global leaders are taking this threat seriously. Canadian Finance Minister François-Philippe Champagne described Mythos as an "unknown unknown" that warrants the attention of all finance ministers. Bank of England Governor Andrew Bailey has said his institution is carefully examining what it means for the risk of cybercrime.

Project Glasswing is a temporary measure at best

Formed in response to the risk of Mythos weaponization, Project Glasswing is a narrow, defensive consortium led by Anthropic with more than 50 member organizations. The group's intention is to use Mythos to find and patch vulnerabilities in critical software while the model remains withheld from general availability. While Project Glasswing aims to give security teams time to reorganize their defenses, this window won't hold forever, and its partners may not be able to fix all vulnerabilities in time.

In fact, reporting indicates that an unauthorized group has already accessed Mythos through a data breach at an AI training startup combined with contractor access to Anthropic's systems, demonstrating that access restrictions alone cannot be relied upon. Rivian CISO Mike Johnson has warned that "by the end of the year, Mythos-level capabilities will be in the hands of any attacker."

Organizations can't rely solely on rollout restrictions imposed by AI developers. Mainframe security teams have to work proactively to reduce the risks posed by Mythos.

Why platform strength isn't enough

The mainframe has traditionally offered security advantages such as physical isolation, proprietary architecture, and granular access controls. But unless this platform strength is complemented with the right operational security capabilities, including automation, monitoring, and lifecycle management, that strength won't translate into real-world security outcomes.

To actively protect the mainframe end-to-end, mainframe security teams must address risks including:

- **API-based threats:** Attackers can hijack trusted API connections between distributed systems and the mainframe, exploiting machine-to-machine authentication to reach mainframe data directly.
- **Identity-based threats:** Credential theft, insider threats, and privilege escalation, already among the most common attack vectors across platforms, are turbocharged by AI-driven phishing and social engineering; machine identities including service accounts and API keys represent an additional and often under-protected attack surface.
- **Software supply chain attacks:** Third-party involvement in breaches doubled in a single year. Compromised mainframe development tools, vendor software, and shared CI/CD pipelines all represent potential entry points.
- **Lateral movement:** Once a distributed system is compromised, existing trusted connections including database links, shared credentials, and network pathways can be used as a stepping stone to the mainframe.

Operational gaps such as expired or misconfigured certificates, unmonitored privileged sessions, outdated software versions, and manual security processes can make these threats more exploitable. As AI-driven attack capabilities continue to evolve, staying current on supported BMC AMI Security software versions and adopting new protections quickly becomes increasingly important.

Preparing Mainframe Security for AI-Accelerated Threats

AI-accelerated threats increase the cadence of attacks, expand the scope of vulnerability discovery across hybrid environments, and compress the time available for security teams to respond. As Mythos-level capabilities proliferate, ISACA, a global, independent non-profit association focused on digital trust, calls for organizations to adopt continuous exposure management. With this approach, security teams correlate findings with runtime topology and business criticality in real time to close the operational gaps that advanced threats exploit.

This shift also reflects what organizations such as FNTS are seeing in the field: compliance reporting alone is no longer enough. Security teams need operational security capabilities that provide continuous visibility, real-time monitoring, and rapid response across the hybrid enterprise. In the Mythos era, organizations must move beyond periodic audits and static controls to actively detect, investigate, and contain threats before automated attacks can escalate

BMC AMI Security helps close these operational gaps across key threat categories facing the mainframe:

- **Delivering faster forensic investigations** with real-time visibility and session auditing**:** BMC AMI Command Center for Security provides an affordable mainframe-native SIEM for immediate visibility into z/OS security incidents. BMC AMI Datastream streams mainframe security events into enterprise SIEMs in real time, eliminating the mainframe security silo and enabling coordinated detection of lateral movement across the hybrid environment. BMC AMI Security Session Monitor delivers continuous auditing of application and data access to detect session-level anomalies before they escalate. Together, these capabilities turn enriched, real-time security data into the speed of response that the Mythos era demands.

- **Protecting against API-based threats** with certificate lifecycle management: Certificate-based authentication provides stronger protection against API-based threats than API keys alone, but only if certificates are kept current and correctly configured. As certificate lifetimes shorten toward 47 days by 2029, manual certificate management creates an unacceptable operational risk. BMC AMI's integration with enterprise certificate infrastructure automates certificate lifecycle management across the hybrid environment, preventing vulnerabilities resulting from expired or misconfigured certificates.
- **Protecting against identity-based threats** with behavioral monitoring, MFA, and privileged access controls: As threats accelerate, organizations must be able to identify stolen credentials and rogue users immediately before an automated attack can escalate. BMC AMI User and Entity Behavior Analytics (UEBA) monitors both human and machine identities under an assumed-breach posture, using AI-driven behavioral baselining to detect anomalies such as unusual login times, abnormal data access, and anomalous query patterns that signal credential misuse. MFA support via Okta extends enterprise authentication standards to mainframe access. Privileged Access Management further reduces risk by enforcing controlled, auditable access to critical systems, commands, and sensitive data.
- **Protecting against supply chain attacks** with DevSecOps integration and file integrity monitoring: BMC AMI DevSecOps integration supports code signing and hashing verification, ensuring that mainframe deployments carry the same supply chain security controls as distributed systems. File Integrity Monitoring (FIM) capabilities detect and remove malware from recovery sites during restoration, with change management controls providing a full forensic audit trail of what was deployed, when, and by whom.

- **Reducing exploitable gaps with AI-driven penetration testing and exposure assessment:** Traditional periodic testing is no longer sufficient against autonomous, AI-accelerated threats. AI/ML penetration testing helps organizations identify vulnerabilities across LLMs, APIs, applications, identities, and hybrid infrastructure before attackers can chain them together into real-world exploits, enabling security teams to proactively remediate weaknesses and strengthen operational resilience.

Limiting the blast radius: Zero Trust and resilience

As mainframe teams work to strengthen their security posture against Mythos-enabled attacks, they must also ensure the resilience to mitigate the impact of any incident that does occur. The Cloud Security Alliance (CSA) advises organizations to verify and enable mitigating controls such as segmentation, egress filtering, Zero Trust architectures, and phishing-resistant MFA to limit post-exploitation impact.

AMI Security provides essential capabilities to address these requirements.

- **BMC AMI Enterprise Connector for Illumio** brings the mainframe into enterprise Zero Trust with automatic micro-segmentation, restricting lateral movement even when a distributed system has been compromised.
- **Immutable Cloud Vaults** prevent data tampering and ensure forensic recovery. Transaction logs are streamed continuously off-platform—even to on-premises object storage—to enable no-data-loss recovery even if the mainframe environment is compromised.
- **Automated Forward Recovery** restores lost transactions post-attack, supporting the tight recovery windows required under frameworks such as DORA.

What AI-resilient security means in practice

While Mythos may be the first widely discussed example of autonomous AI-driven cyber exploitation, it is part of a broader wave of increasingly capable models that are accelerating the pace of cybersecurity risk. This makes established best practices even more critical. The UK AI Security Institute has noted it cannot confirm whether Mythos Preview would be able to attack well-defended systems, reinforcing that strong, active defenses remain the most effective response. As autonomous AI models continue to emerge and proliferate, organizations must shift from static audits to continuous, real-time defense, with active management of certificates, identities, sessions, and supply chain integrity across the full hybrid environment.

However the Mythos story evolves, it's clear that mainframe teams will continue to see new types of threats and rising risk. By strengthening security across your mainframe attack surface with end-to-end technologies and operational capabilities, you can protect your organization more effectively against whatever the future holds.

Explore how BMC AMI Security helps identify vulnerabilities, monitor privileged activity, and reduce exposure before threats become incidents.