Rocket® software

What Mythos Can't See: The Mainframe Risk Leaders Are Accountable For

Mainframe security in the age of AI and quantum: operational resilience for regulated industries

Introduction

The idea that the mainframe is an isolated fortress is a dangerous myth in today's hybrid IT environment. As you offload data, connect to cloud analytics, and integrate with modern APIs, your attack surface expands far beyond the green screen. For regulated industries, this interconnectivity introduces new variables to the risk equation.

Open-source insight from Mythos is essential, but mainframe security demands layered controls that, include both code and configuration analysis. Mythos uncovers open-source threats that signal risk to the mainframe ecosystem. However, it cannot see inside assembled mainframe binaries, closed-source vendor software, or runtime behavior in production — where many of today's most serious vulnerabilities live.

We understand the immense responsibility you carry in protecting mission-critical IT. Regulators expect robust operational resilience, and frameworks like DORA, OSFI, and BaFin increasingly focus on the integrity of your entire ecosystem. While emerging threats from artificial intelligence (AI) and quantum computing may not target z/OS® directly, they actively exploit the interconnected periphery — potentially reaching the mainframe through trusted pathways that now define today's integrated enterprise.

Here, we explore how AI-driven detection, AI-driven attacks, and quantum risks impact critical systems like your mainframe. We partner with you to map these emerging risks to the controls, evidence, and regulatory expectations your auditors already prioritize, helping you strengthen compliance and security without disrupting business operations.

TRENDS AND STATISTICS

AI-driven intrusion and future breaks in cryptography

When we examine the threat landscape, we focus on concrete risks rather than science fiction doom. The reality of AI and quantum threats is grounded in speed, scale, and long-term data exposure.

First, AI-driven intrusion and model-driven reconnaissance allow attackers to automate vulnerability discovery at an unprecedented pace. Sophisticated actors use AI to craft context-aware social engineering campaigns that can fool even security-conscious mainframe administrators. Furthermore, AI models can rapidly identify obscure vulnerabilities in middleware, compilers, or operating system components that connect to your mainframe.

Second, while quantum computing remains a future threat, it presents an immediate compliance challenge for data protection. Threat actors currently employ "harvest now, decrypt later" strategies. If quantum computers break current encryption standards in the next decade, the sensitive customer data, payment information, and clinical safety cases you store today will become vulnerable.

THE NEW REALITY

How AI-driven threats exploit mainframe vulnerabilities

For years, the prospect of a mainframe compromise was largely dismissed as an edge case, requiring deep institutional knowledge, extended timeframes, and privileged access. Those assumptions are no longer valid. The emergence of AI-driven threats has radically condensed the timeline and removed traditional barriers, fundamentally changing the risk profile for mainframes.

AI has collapsed the mean time to exploit. Advanced models can rapidly analyze complex legacy environments, correlate decades of obscure code patterns, and uncover exploitable conditions in minutes rather than months. Today, vulnerabilities can be identified and weaponized at machine speed, even by attackers who lack prior mainframe expertise.

Consider how a modern attack might unfold:

- An undetected code-level trap door vulnerability resides in a program.
- An attacker, whether external or internal, leverages AI to detect, prioritize, and exploit this weakness.
- Upon exploitation, the attacker can escalate their authority and disable forensic logging and monitoring, erasing all evidence that your security teams depend on for detection and response.
- Security controls may be bypassed or even deactivated, granting unauthorized access to sensitive data and key system functions.

The implications ripple outward from the mainframe's core. Data can be exfiltrated across thousands of connected applications. Trust relationships with downstream systems can be misused to propagate fake credentials or launch further attacks across your ecosystem. What begins as a discrete compromise can escalate quickly into an enterprise-wide incident.

This threat is not hypothetical. Every vulnerability assessment conducted utilizing Rocket® z/Assure® VAP has uncovered issues that, if left unaddressed, could be exploited by AI-enabled attackers. These are not rare exceptions but routine findings, with high-severity risks present in nearly every environment assessed.

For executive and board-level leaders, this makes mainframe security a matter of enterprise risk, not just system hygiene. Traditional ideas about isolation no longer apply. Detection after exploitation is simply not adequate, given the speed and scope of modern threats. Proving integrity, resilience, and robust control — up front, and continuously — must become standard practice and an integral part of your governance model.

AI-driven threat actors exploit all available knowledge. By automating credential testing, privilege escalation attempts, configuration analysis, and behavioral testing, they can expose gaps that were previously shielded by complexity or obscurity. As AI systems grow more capable and organizations themselves adopt AI for business value, new risks can emerge internally as well through misconfigurations, overly broad permissions, or insecure integrations.

In this landscape, experimentation and active exploitation have never been cheaper, but the consequences for failure have never been higher. The mainframe, at its heart, is another computing platform, albeit one running at the very core of business operations. Sustained resilience now hinges not on outdated assumptions, but on continuous validation, disciplined controls, and clear visibility into every potential risk vector. In an AI-driven threat era, proactive measures and enterprise-grade assurance are now non-negotiable.

Addressing Mythos gaps and AI threats

You don't need to discard your current security architecture to face these new realities. Instead, you must adapt your existing controls to account for AI and quantum variables. While tools like Mythos offer valuable open-source intelligence by surfacing threats that could signal risk to the mainframe ecosystem, they are not designed to uncover or remediate the most critical vulnerabilities unique to mainframe platforms. Mythos cannot see inside assembled mainframe binaries, closed-source vendor software, or production runtime behavior — areas where serious exploits often hide and where AI-driven threats are most likely to emerge.

Mainframe security demands a layered approach: open-source insight alone is not enough. Regulators and auditors increasingly expect you to treat the mainframe as part of the AI and quantum threat surface, but they also require that layered controls — code-level and configuration analysis — are in place and continuously monitored. For example, when considering AI-driven attacks and quantum-safe cryptography, you must understand what BaFin, OSFI, and DORA implicitly expect you to do with z/OS. They expect to see how you manage operational resilience when attackers use AI to bypass traditional perimeter defenses.

You need to demonstrate how AI changes insider threat detection on IBM® RACF, CA-Top Secret, or CA-ACF2, and you must prepare the exact evidence your auditors will ask to see.

By mapping these new risks to your existing compliance frameworks — such as SOC2, PCI DSS, SOX, and GDPR — you transform abstract fears into tangible, manageable risk controls. Only with comprehensive, multi-layered visibility can you address the evolving landscape and deliver the evidence required by auditors and regulators, bridging the critical gaps that open-source scans like Mythos leave behind.

Operational risk and data integrity in AI-driven environments

Data integrity is the cornerstone of operational resilience. If an AI-driven attack gains a foothold in a connected cloud analytics platform, it can manipulate or corrupt data after it leaves the mainframe, or even before processing occurs. This scenario creates massive PCI DSS and SOX compliance issues.

Your modernization journey today builds the foundation for protecting that data. You must implement robust monitoring across all API gateways and event streams connected to the mainframe. By applying zero trust principles to every data exchange, you ensure that AI-driven anomalies are caught before they compromise the integrity of your core records.

Privileged access management and identity governance when attackers use AI

A single compromised credential can unravel your entire security posture. AI-powered attackers excel at exploiting dormant accounts and orchestrating credential stuffing at scale. Therefore, privileged access management and identity governance require immediate attention.

To counter AI-driven reconnaissance, you must tighten identity governance across RACF, ACF2, and Top Secret. Auditors will look for automated access reviews and AI-assisted anomaly detection that flags unusual behavior from privileged accounts. By integrating your mainframe identity data with enterprise-wide human resources and identity systems, you eliminate the blind spots that AI attackers love to exploit. You can act now to clean up dormant IDs and enforce strict least-privilege access without causing system downtime.

Preparing for quantum-safe encryption on Z without rip-and-replace

Future-proofing your data's confidentiality requires a pragmatic, phased approach. You cannot afford a massive rip-and-replace project that threatens your uptime. However, regulatory bodies will soon demand concrete plans for quantum-safe cryptography.

Start by conducting a comprehensive cryptographic inventory of your z/OS environment. Identify where you use legacy crypto, evaluate your key management practices, and assess data in transit via TLS and data at rest in backups. We partner with you to design phased rollouts and parallel runs of quantum-safe algorithms. By implementing dual algorithms and piloting quantum-safe crypto for specific, high-risk subsets of archival data, you demonstrate proactive compliance to your risk board and regulators.

Automated evidence collection and audit trails for AI-era threats

Auditors no longer accept manual, point-in-time security snapshots. They require continuous, automated evidence collection that proves your controls are actively mitigating AI and quantum risks.

You must modernize your audit trails to capture the context of API expansions and containerized adjuncts interacting with the mainframe. Automated evidence collection ensures you can quickly respond to complex audit questions, including how your environment detects and addresses evolving access risks. By streamlining how you collect and present security-relevant operational records, you not only satisfy OSFI and DORA requirements but also enable your teams to focus on strategic risk management.

PRESCRIPTIVE RECOMMENDATIONS

Ensuring resilience through continuous detection and compliance

Modern mainframe security requires more than compliance checklists and point-in-time reviews. It demands a layered, proactive approach centered on continuous detection, runtime visibility, and comprehensive risk management. As AI-driven threats accelerate, the gaps left by traditional tools and static code reviews have become too significant to ignore.

While open-source intelligence from tools like Mythos delivers valuable ecosystem visibility, these solutions are not designed to uncover deep-seated vulnerabilities within assembled mainframe binaries, closed-source software, or active runtime environments — the areas where today's most critical and complex risks can hide. Relying solely on this intelligence leaves unacceptable blind spots, both for adversaries seeking an entry point and for regulators assessing enterprise assurance.

It's essential to recognize that static and source-based analysis, although helpful, cannot reveal many high-impact issues — such as race conditions, time-of-check/time-of-use (TOCTOU) vulnerabilities, authorization bypass caused by policy drift, or memory-corruption defects like use-after-free. These risks frequently evade detection until systems are running in production, making runtime observation and live integrity monitoring indispensable for demonstrable control.

To advance operational resilience and close these critical gaps, we recommend the following:

- Deploy code-based vulnerability scanning and live integrity monitoring developed specifically for the mainframe, providing visibility into compiled binaries and active runtime, not just source code.
- Layer advanced runtime assessment and monitoring with open-source threat intelligence, ensuring both ecosystem risks and mainframe-specific threats are identified and managed without reliance on any single method.
- Integrate mainframe security controls with enterprise-wide compliance functions, extending requirements like multi-factor authentication, secure access, ongoing testing, and audit readiness across all environments.
- Automate patch validation, reporting, and audit trails, so compliance is provable on an ongoing basis rather than just at audit time or after incidents.
- Begin now to catalog your cryptographic landscape and pilot quantum-safe algorithms in areas of highest exposure, demonstrating proactive response to emerging risk.
- Establish automated, continuous evidence collection for assurance teams and regulators, supporting both operational and financial resilience.

This holistic and layered approach transforms mainframe security from a reactive necessity into a proactive driver of trust and resilience. By acting on these recommendations, you not only bridge the critical gaps that have emerged in the era of AI and quantum risk, but also position your organization to meet the highest regulatory standards and support the secure continuity of business-critical operations.

AI-driven threats have redefined what's required to secure your mainframe in a connected world. While platforms like Mythos offer valuable open-source intelligence across the broader ecosystem, they cannot detect the deep-seated vulnerabilities that exist within assembled mainframe binaries, closed-source vendor software, or actual runtime operations — where some of today's most serious risks hide. Relying solely on this type of intelligence leaves unacceptable gaps for both adversaries and auditors. True resilience demands a layered approach, integrating open-source insight with mainframe-specific controls that provide continuous, live visibility into code, configuration, and active execution.

Achieving resilience demands layered controls that reach far beyond open-source intelligence or static analysis. Mainframe security now requires integrated code and configuration analysis, active runtime integrity monitoring, and continuous compliance — matched to the standards already expected for cloud and distributed environments. Simply put, every control that protects your broader enterprise must also protect the mainframe. As core systems become open and interconnected, continuous, in-depth visibility and assurance across all security layers provide the foundation for operational resilience in an AI-driven world.

The data makes the urgency clear: over 91% of mainframe organizations have experienced security compromises in the past five years, and every vulnerability assessment uncovers issues, — many previously missed by static or traditional tools. This reality reinforces that static analysis and open-source intelligence alone are not enough. Other serious risks — race conditions, time-of-check/time-of-use bugs, authorization bypass due to policy drift, and memory-related vulnerabilities — often remain hidden until they surface during live operation. Only a layered approach that combines code-based vulnerability management, active runtime observation, and continuous evidence collection can close these gaps and enable true operational resilience.

Conclusion

The pace and scope of AI and quantum threats leave no room for complacency. For risk leaders, the message is clear: legacy assumptions about mainframe security must yield to a proactive, layered approach that closes blind spots before attackers can exploit them. Operational resilience demands more than compliance — it requires continuous detection, live visibility into vulnerabilities, and rigorous evidence of control across every layer of your environment.

Now is the time to treat the mainframe as integral to your broader enterprise security strategy. Implement code-based vulnerability management, automate audit and compliance evidence, and ensure your controls meet or exceed the standards set for cloud and distributed systems. The regulatory and operational stakes have never been higher, but with the right partner, they're manageable.

Rocket Software stands ready to empower you with the clarity, assurance, and actionable insight your organization needs to defend what matters most. Together, we can safeguard your mission-critical systems and build a foundation of trust, resilience, and readiness for the future, — no matter how the risk landscape evolves.

Ready to dive in? Learn more about Rocket Software's Security Solutions, or speak to one of our experts.

About Rocket Software

Rocket Software is a global technology leader in modernization and a partner of choice that empowers the world's leading businesses on their modernization journeys, spanning core systems to the cloud. Trusted by over 12,500 customers and 750 partners, and with more than 3,200 global employees, Rocket Software enables customers to maximize their data, applications, and infrastructure to deliver critical services that power our modern world. Rocket Software is a privately held U.S. corporation headquartered in the Boston area with centers of excellence strategically located throughout North America, Europe, Asia and Australia. Rocket Software is a portfolio company of Bain Capital Private Equity.

Rocket® software

Modernization. **Without Disruption.**™

Visit RocketSoftware.com

MAR-18213_WP_AIThreats_V1