WHITE PAPER



BUILDING AN ENTERPRISE CYBERSECURITY FORTRESS TO COMBAT ADVANCED MALWARE IN THE AI AGE

Advanced malware has emerged as one of the most consequential threats facing modern enterprises. With Artificial Intelligence (AI) accelerating both the creation and evasion capabilities of malicious software, traditional detection paradigms are proving dangerously insufficient. This study examines the evolving advanced malware landscape, its impact across industry verticals, the current state of AI-driven detection and response capabilities, and the persistent gaps that remain unaddressed. It further outlines how organisations like Infosys are investing in next-generation approaches to close these gaps, and the strategic priorities enterprise leaders must adopt to build resilient, explainable, and adaptable malware defence programmes in the AI era.

Infosys
Navigate your next

1. The Cybersecurity Threat Landscape

Cybersecurity has evolved from a discipline focused on perimeter firewalls and rule-based intrusion prevention into a dynamic, intelligence-driven field in which defenders must continuously adapt to adversaries who exploit the same technological advances. The modern enterprise attack surface spans cloud infrastructure, mobile endpoints, operational technology networks, software supply chains, and an intricate web of third-party integrations. Within this broad landscape, ransomware has consistently been one of the most prevalent and costly threat categories, found in 44% of all confirmed data breaches, according to the Verizon Data Breach Investigations Report 2025.

Kaspersky's Security Bulletin 2024 documented more than 467,000 new malicious files detected each day, representing not total infections but the raw rate of new variant creation. IBM's Cost of a Data Breach Report 2025 placed the average US breach cost at $10.22 million, a figure rising nine percent year-over-year, driven both by the growing sophistication of attacks and expanding regulatory consequences for inadequate data protection.

The Cybersecurity Threat Landscape — Narrowing to Advanced Malware

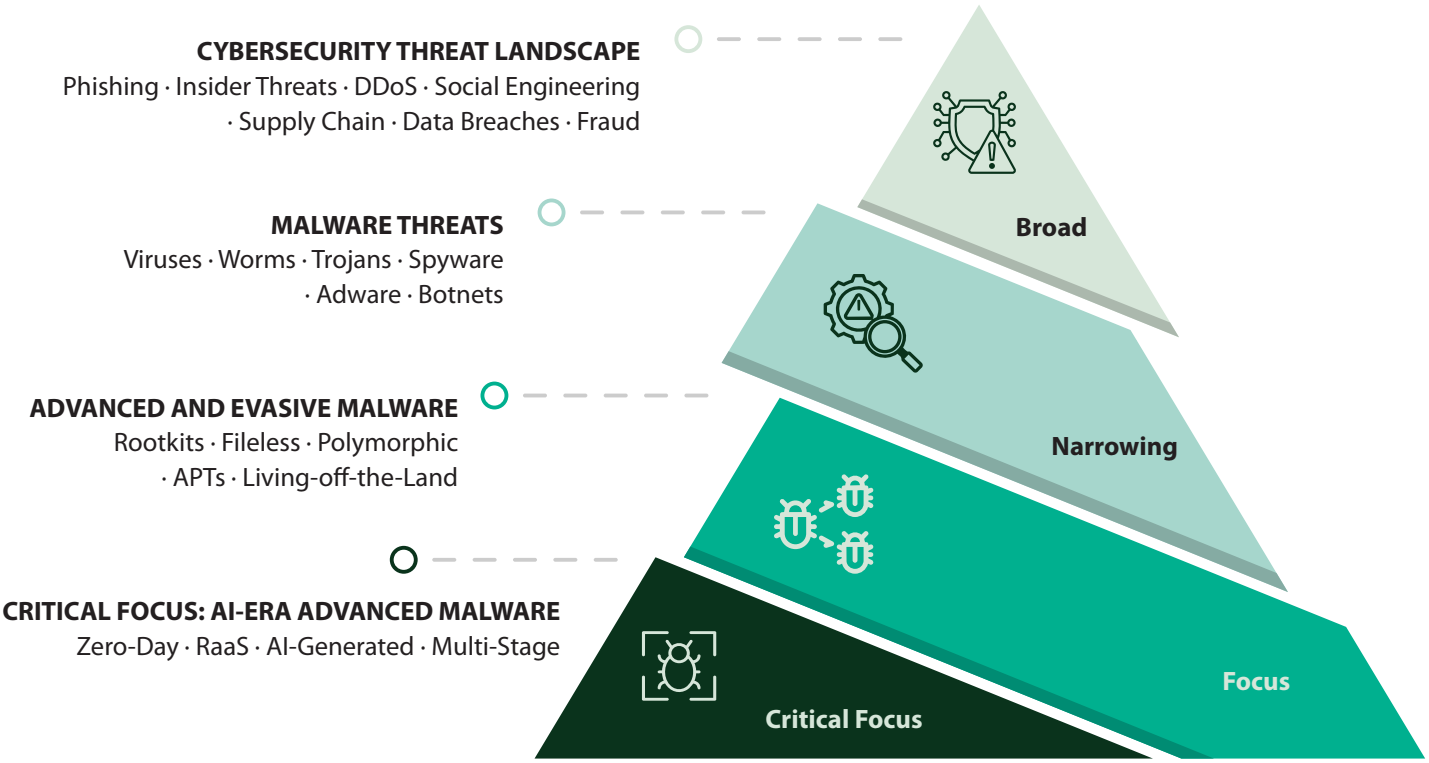


Figure 1: The cybersecurity threat hierarchy narrowing from the broad cyber security landscape to the critical focus area of advanced and evasive malware in the AI era.

2. Advanced Malware in the AI Era: Characteristics and Escalation

Advanced malware differs from conventional malicious software across several critical dimensions. While traditional malware relies on known techniques that can be catalogued and blocked via signatures, advanced malware is engineered specifically to circumvent those mechanisms.

Key characteristics of the threat category include:

Polymorphism and metamorphism:
Code that rewrites itself at each infection, generating unique file hashes that defeat signature-based detection while preserving malicious functionality.

Fileless execution:
Operation entirely within system memory, leveraging legitimate tools such as PowerShell, WMI, and living-off-the-land binaries (LOLBins), leaving no artefacts on disk for conventional antivirus to inspect.

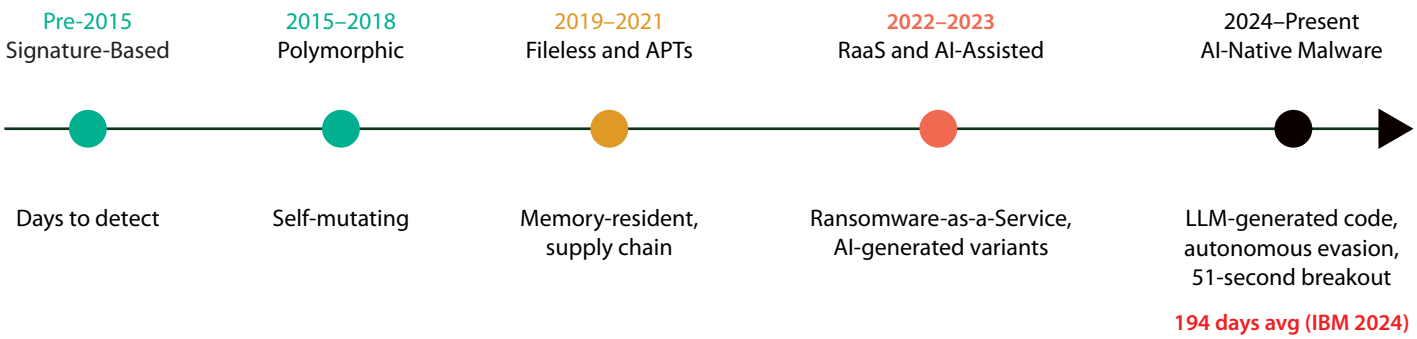
Zero-day exploitation:
Targeting vulnerabilities unknown to vendors or the security community, for which no patch or signature exists at the time of attack.

Ransomware-as-a-Service (RaaS):
A criminal business model in which ransomware toolkits are leased to affiliates, dramatically lowering the technical barrier to entry and increasing attack volume exponentially.

AI-assisted generation:
Adversaries using Large Language Models (LLMs) to write novel malware code, craft convincing phishing lures, and iteratively test evasion techniques against public detection engines before deployment.

Multi-stage and modular architectures:
Initial access tools, loaders, and payloads delivered in separate stages across extended time windows, making campaign attribution and detection significantly harder.

Figure 2: Evolution of Advanced Malware — From Signatures to AI-Native Threats



467,000+ new malware variants detected daily (Kaspersky 2024) · Fastest eCrime breakout: 51 seconds (CrowdStrike 2025)

Figure 2: Evolution of advanced malware from signature-era variants to AI-native threats. Detection gaps widen at each generational shift as adversary sophistication outpaces defender tooling.

External Document © 2026 Infosys Limited

2.1 The AI Amplification Effect

The integration of AI into offensive security has fundamentally altered the threat calculus. AI enables malware authors to automate previously labour-intensive tasks: generating thousands of functional variants, identifying the precise combination of obfuscation techniques most likely to evade a specific target's security stack, and producing contextually accurate social engineering content that bypasses human vigilance at scale.

The consequences are measurable. CrowdStrike's Global Threat Report 2025 recorded the fastest observed eCrime breakout time at 51 seconds — the interval between initial access and lateral movement to other systems.

The average breakout time dropped to 48 minutes in 2024, from 62 minutes the previous year. This compression of the attack window dramatically narrows the time available for automated detection and human response, placing unprecedented demands on the speed and precision of security operations.

"Generative AI has fundamentally democratised advanced malware creation. Capabilities that previously required nation-state resources or elite criminal groups are now accessible to a much broader range of threat actors." — CrowdStrike Global Threat Report 2025

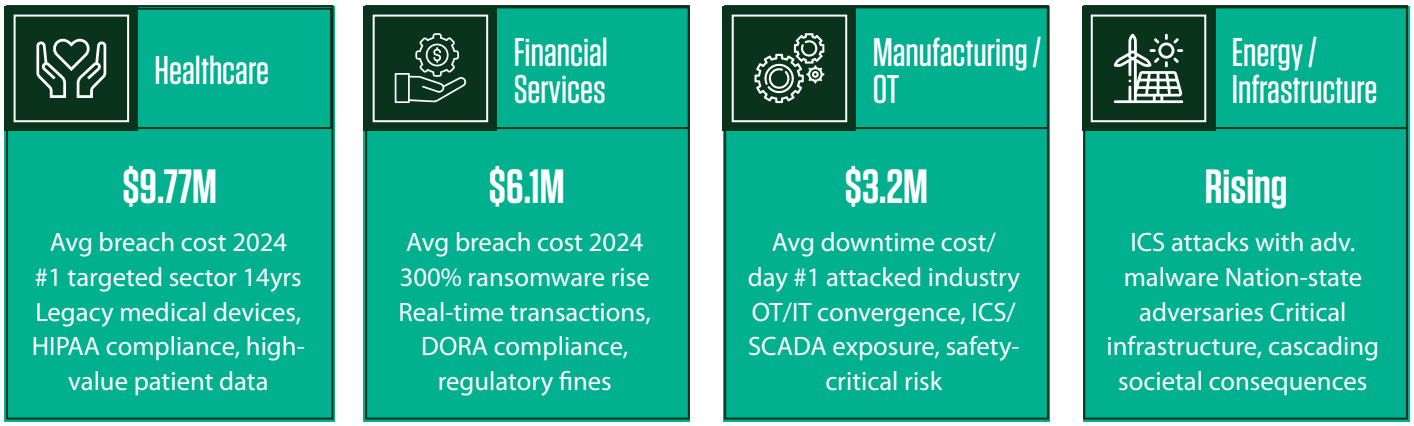
3. Real-World Impact: Enterprise Case Studies

The financial and operational consequences of advanced malware are no longer theoretical. A consistent pattern emerges from recent high-profile incidents: organisations with sophisticated, enterprise-grade security infrastructure were nonetheless compromised — not because they lacked security investment, but because evasive malware successfully circumvented the tools deployed. Table 1 summarises five defining incidents between 2021-2024.

Organisation	Date	Threat Vector	Financial Impact	Key Lesson
Change Healthcare	Feb 2024	ALPHV/BlackCat Ransomware	**$3.09B total FY2024**	Single credential, no multi-factor authentication on Virtual Private Network
CDK Global	June 2024	BlackSuit Ransomware	**$1B+ dealer losses**	15,000 dealerships offline; pen-and-paper fallback
MGM Resorts	Sept 2023	Scattered Spider (Vishing)	**$100M+ total impact**	Social engineering bypassed all technical controls
Colonial Pipeline	May 2021	DarkSide Ransomware	**$4.4M ransom + billions**	Dormant VPN account with no MFA enabled
MOVEit Mass Breach	June 2023	Clop Ransomware (Zero-Day)	**2,600+ orgs affected**	Zero-day in widely used file-transfer platform

Table 1: Major advanced malware incidents in 2021-–2024. All affected organisations had enterprise security infrastructure in place at the time of breach. Sources: UnitedHealth Group FY2024 Earnings; IBM Security; MGM SEC Filing 8-K; US Department of Justice.

Industry Sectors Most Impacted by Advanced Malware (2023–2024)

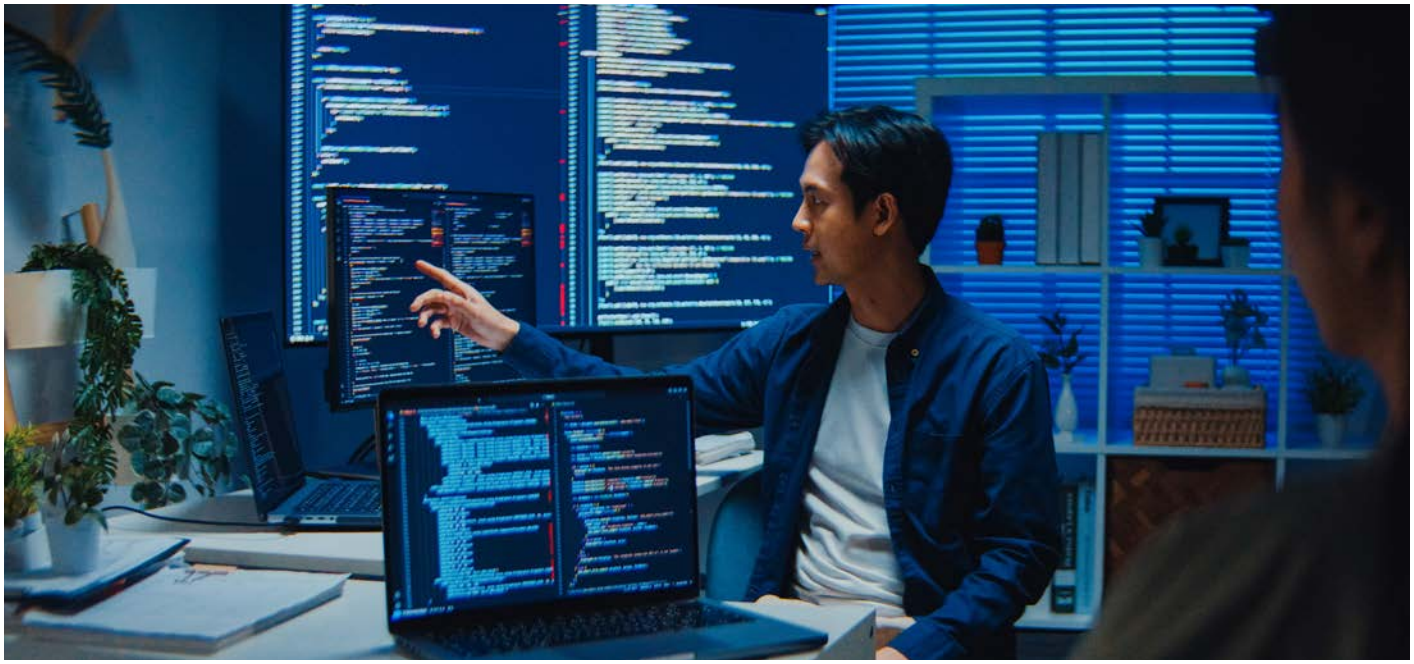


Sources: IBM Cost of a Data Breach 2024 · IBM X-Force Threat Intelligence Index 2024 · Dragos ICS/OT Year in Review 2024

Figure 3: Industry verticals most severely impacted by advanced malware, with representative financial and operational metrics. Sources: IBM Cost of a Data Breach 2024; IBM X-Force Threat Intelligence Index 2024.

3.1 The Hidden Costs Beyond the Ransom

The figures in Table 1 frequently under-represent the true cost of an advanced malware incident. Direct costs — ransom payments, incident response retainers, forensic investigation — constitute a fraction of the total impact. Indirect costs include reputational damage and customer attrition, regulatory fines under the EU's General Data Protection Regulation (GDPR), and Digital Operational Resilience Act (DORA), the US' Health Insurance Portability and Accountability Act (HIPAA), and the EU's Network and Information Security Directive 2 (NIS2), the operational cost of reverting to manual processes during recovery, legal liability from affected third parties, and the long-term cost of board-mandated security uplift programmes. For critical infrastructure operators, the calculus is more severe still: a successful ransomware attack on an operational technology (OT) environment can trigger safety shutdowns, cause physical equipment damage, and carry consequences extending into public safety and national security domains.



External Document © 2026 Infosys Limited

4. Industry-Specific Threat Scenarios

4.1 Healthcare and Life Sciences

Healthcare organisations manage an extraordinary concentration of high-value data: protected health information (PHI), medical device firmware, drug discovery datasets, and insurance claims infrastructure. Legacy systems — including medical devices running end-of-life operating systems — create persistent vulnerabilities that cannot be addressed through conventional patching cycles without disrupting patient care. Advanced ransomware targeting healthcare has evolved beyond data encryption to include exfiltration and threatened publication of patient records, a double-extortion model creating acute exposure under HIPAA and international privacy regulations. IBM's 2024 Cost of a Data Breach report found healthcare has held the top position for breach cost for 14 consecutive years, averaging $9.77 million per incident.

4.2 Financial Services and Fintech

The financial sector presents both the highest motivation for adversaries — direct access to monetary assets and transaction infrastructure — and the most stringent regulatory requirements for demonstrating security posture. DORA, effective January 2025, requires financial entities to document not only security incidents but the technical rationale behind detection and response decisions. Advanced malware in financial services frequently targets the window between transaction initiation and settlement, while AI-generated deepfake audio and video are increasingly used to impersonate executives and bypass authentication controls in high-value transfer authorisation processes.

4.3 Manufacturing and Operational Technology

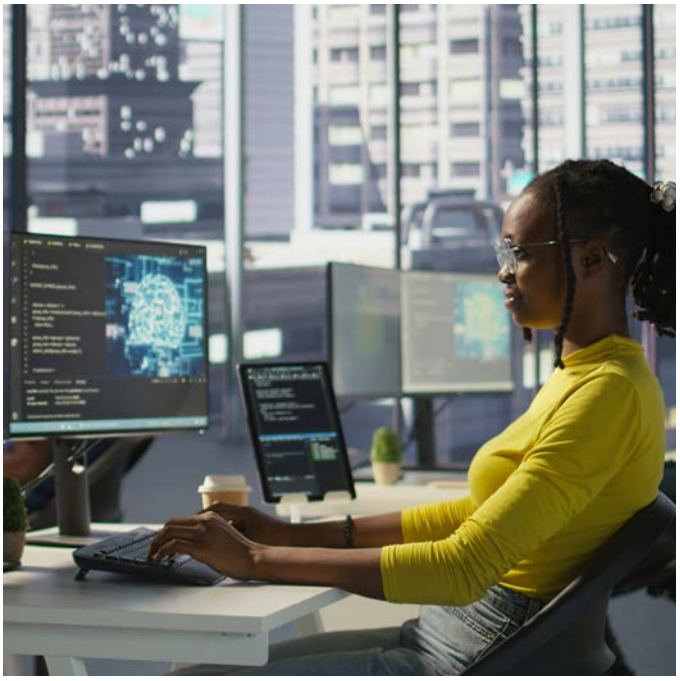
Driven by the convergence of IT and OT environments, manufacturing has become the most targeted industry globally, according to IBM's X-Force Threat Intelligence Index 2024. Historically air-gapped industrial control systems are increasingly connected to enterprise networks and cloud- management platforms, creating new pathways for malware to propagate into industrial settings. Ransomware encrypting historian databases and human-machine interfaces can halt production lines without touching a single programmable logic controller. More sophisticated attacks manipulate control signals directly, risking physical equipment damage and, in extreme cases, environmental or safety incidents. The average cost of unplanned manufacturing downtime attributable to cyberattacks has been estimated at $3.2 million per day.

4.4 Energy, Infrastructure, and Government

Energy and critical infrastructure operators face a distinct threat profile characterised by nation-state adversaries with long-term strategic objectives. Advanced persistent threats in these sectors are designed not for immediate financial gain but for persistent access, intelligence collection, and the potential to cause disruption at a time and manner of the adversary's choosing. The Dragos ICS/OT Cybersecurity Year in Review 2024 found that a growing share of Industrial Control Systems-targeted attacks now involve advanced malware with specific OT awareness — capabilities once exclusive to nation-state actors but now increasingly available through criminal ecosystems.

4.5 MSSPs, Cloud, and Supply Chain

Managed Security Service Providers (MSSPs) represent a high-value supply chain target: compromising a single MSSP potentially provides adversaries with access to dozens or hundreds of client environments simultaneously. The SolarWinds and Kaseya supply chain attacks demonstrated the multiplicative effect of targeting security tooling vendors. For cloud-native organisations, advanced malware increasingly targets cloud configuration, identity providers, and Continuous Integration and Continuous Delivery (CI/CD) pipelines — areas where traditional endpoint-focused security tools have limited visibility and where a single compromised privileged identity can enable access at enterprise scale.





5. Current Approaches: AI-Enhanced Malware Detection

Marking a significant departure from signature-dependent approaches, the security industry has responded to escalating advanced malware sophistication with a generation of AI-driven detection and response capabilities. The modern enterprise security operations centre (SOC) increasingly functions as an intelligence platform, correlating signals across multiple data sources and applying Machine Learning (ML) models to prioritise analyst attention on the highest-fidelity threats.

Figure 4: Next-Generation AI-Driven Security Operations Centre Architecture

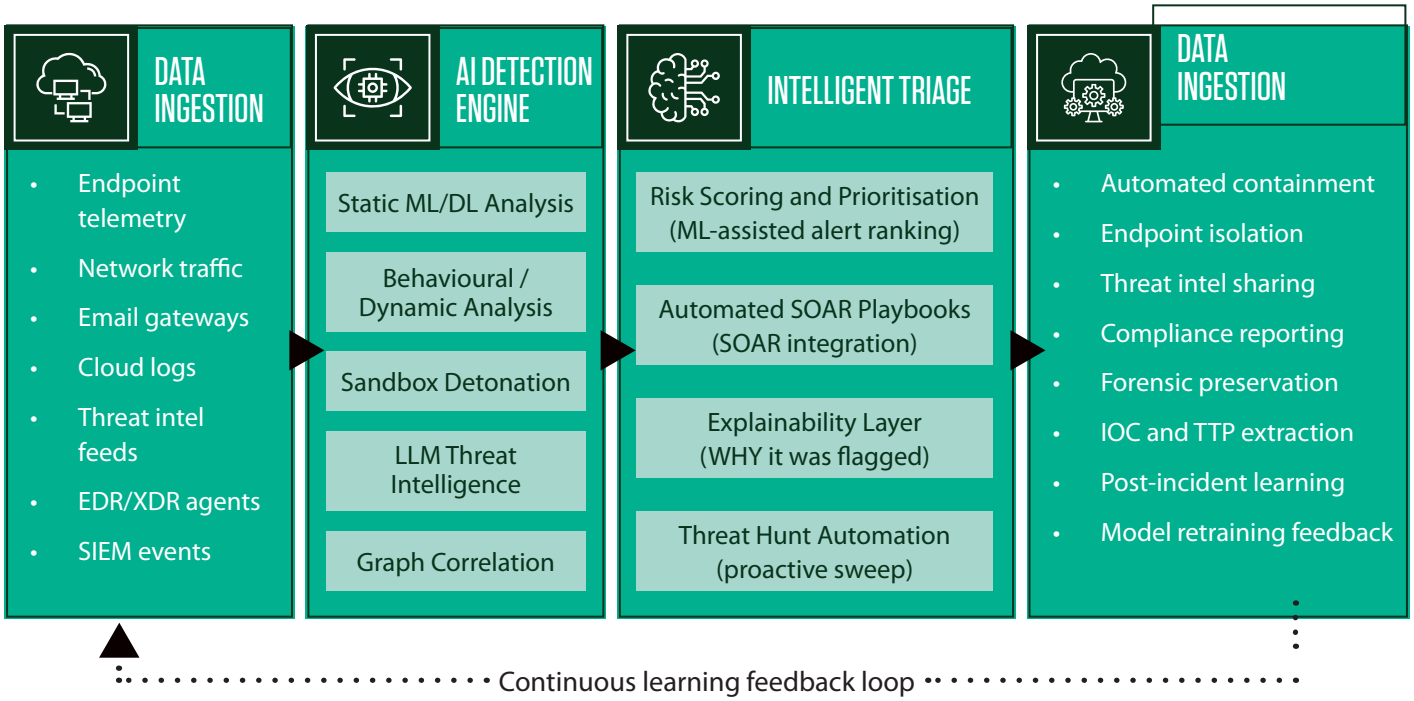


Figure 4: Next-generation AI-driven SOC architecture, showing the four-layer pipeline from data ingestion through detection, triage, and response, with a continuous learning feedback loop connecting outcomes back to detection models.

External Document © 2026 Infosys Limited

5.1 Machine Learning and Deep Learning for Static Analysis

Static analysis ML models examine file properties — byte distributions, structural metadata, imported function signatures, and string characteristics — without executing the file, enabling rapid screening of large volumes at scale. Gradient-boosted tree models have demonstrated high detection accuracy on Portable Executable (PE) malware benchmarks; deep learning approaches, including convolutional neural networks applied to byte sequence visualisations and transformer models applied to assembly code, have extended detection to previously unseen variant families by learning generalised representations of malicious intent rather than specific signatures. The critical advantage is throughput: modern implementations classify tens of thousands of files per second on commodity hardware, enabling real-time screening at email gateway and endpoint scale.

5.2 Behavioural and Dynamic Analysis

Behavioural analysis addresses the fundamental limitation of static approaches by observing what a programme does rather than what it looks like. Endpoint Detection and Response (EDR) and Extended Detection and Response (XDR) platforms instrument operating system kernels to capture process creation, file system modifications, network connections, registry changes, and memory allocation patterns. ML models trained on normal and malicious behavioural sequences identify malicious patterns even when the underlying file has been obfuscated to defeat static classifiers. Sandbox-based dynamic analysis detonates suspicious files in isolated virtual environments, with modern commercial platforms employing hardware-level virtualisation to detect sandbox-aware malware that modifies its behaviour when it suspects an artificial analysis environment.

5.3 LLM-Augmented Threat Intelligence and SOC Automation

LLMs have been applied to cybersecurity in several high-value ways. Security vendors have integrated LLMs into analyst interfaces to enable natural-language queries against Security Information and Event Management (SIEM) data, translate complex Indicators of Compromise (IOC) reports into actionable detection rules, and summarise security events for non-technical stakeholders. Autonomous investigation agents — combining LLMs with tool-use capabilities — are beginning to automate



portions of the Tier 1 analyst workflow: triaging alerts, correlating events across data sources, and drafting incident reports. LLMs have also been demonstrated to map malware behaviour automatically to MITRE ATT&CK technique identifiers (alphanumeric codes), enabling faster threat hunt pivot points and more consistent incident documentation for regulatory compliance.

5.4 Multi-View, Graph-Based, and Ensemble Approaches

Advanced persistent threats and multi-stage ransomware campaigns are characterised by complex activity sequences spanning multiple endpoints, identities, and time periods — patterns invisible to point-in-time detection but visible when the attack sequence is modelled as a temporal graph. Graph neural networks applied to security event data identify attacker behaviour patterns — lateral movement, privilege escalation, data staging etc. — that are individually innocuous but collectively diagnostic of a breach in progress. Multi-view analysis, which combines file properties, behavioural telemetry, network traffic, and threat intelligence feeds, provides a richer and more evasion-resistant representation of a potential threat than any single data source alone. Research in multi-view learning for malware detection is particularly promising: by distributing analysis across orthogonal feature spaces, detection systems maintain accuracy even when adversaries successfully manipulate individual feature dimensions.

Detection Approach	Key Strength	Known Limitation	Maturity
Signature-based antivirus AV	Fast, low overhead	Useless against variants and zero-days	Legacy
Behavioural Analysis (EDR/XDR)	Detects post-execution malice	High false positives; LOLBin evasion	Current
ML Static Analysis	No execution required; scalable	Concept drift; single-feature evasion	Current
Sandbox Detonation	Detailed behavioural insight	Slow; malware detects and lies dormant	Current
Graph and Network Analysis	Detects lateral movement and C2	Encrypted traffic; compute cost at scale	Emerging
LLM-Augmented Threat Intel	Natural-language IOC correlation	Hallucination risk; context limits	Emerging
Multi-View Ensemble ML	Evasion-resilient; explainable verdicts	Drift resilience still maturing	Ongoing research at Infosys

Table 2: Comparative overview of advanced malware detection approaches. Research maturity indicates active academic and early-stage industry investigation.

6. Persistent Gaps and Open Challenges

Despite significant advances in AI-driven detection, substantial gaps remain between the threat posed by advanced malware and the defence capabilities available to most enterprises. These gaps are not primarily a function of technology availability but of how detection systems are architected and how their outputs are consumed by security operations teams.

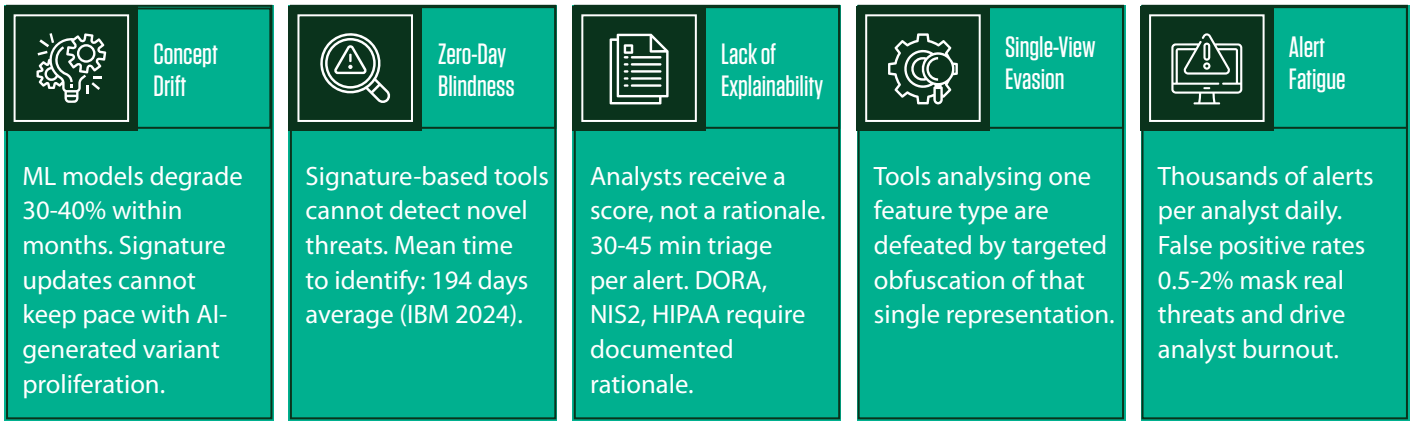


Figure 5: Five persistent gaps in current advanced malware detection. Each represents both a technical challenge and an operational risk that current enterprise tooling does not adequately address.

External Document © 2026 Infosys Limited

6.1 Concept Drift

Concept drift in malware detection refers to the continuous evolution of malware behaviour, where new variants change their structure, code patterns, and execution techniques over time. As attackers adopt strategies like obfuscation, polymorphism, and fileless attacks, the feature distributions learned by a model during training no longer remain valid. This causes a degradation in detection performance, as previously reliable indicators of malicious activity may become obsolete or misleading. For example, Application Programming Interface (API) usage patterns or byte signatures that once indicated malware may also appear in benign software. Hence, models must be continuously updated and designed to handle such shifts to remain effective in real-world environments.

6.2 Zero-Day and Novel Family Detection

Zero-day vulnerabilities and genuinely novel malware families present no prior examples for detection systems to learn from. While behavioural anomaly detection can surface suspicious activity, the false positive rate of purely anomaly-based approaches at enterprise scale renders them operationally impractical without significant human analyst involvement. IBM's 2024 Cost of a Data Breach report found the average time to identify a breach was 194 days globally. This gap partly reflects the period during which advanced malware operates undetected because it does not match any known pattern.

6.3 The Explainability Deficit

When ML-based detection systems generate an alert, they typically provide a risk score rather than a rationale — 'Malware: 94% confidence' — without explaining which properties triggered the assessment. This opacity creates two problems. First, analysts cannot efficiently triage alerts without understanding the basis for the classification, leading to the 30–45-minute average triage time per alert. Second, regulatory frameworks, including DORA, HIPAA, and NIS2, need organisations to provide documented technical rationale for security decisions — a requirement that confidence scores alone cannot satisfy. Building explainability directly into detection architectures is an increasingly urgent requirement.

6.4 Alert Fatigue and Analyst Capacity

The volume of security alerts generated by modern enterprise security stacks has grown far beyond what analyst teams can meaningfully investigate. Ponemon Institute's 2024 research found that large enterprise SOCs generate thousands of alerts per analyst daily, with false positive rates of 0.5–2% across commercial tools. At enterprise scale, even a 1% rate generates tens of thousands of erroneous alerts per day, consuming analyst attention that could be directed towards genuine threats. Alert fatigue is a primary driver of analyst burnout and turnover, compounding the global cybersecurity talent shortage.

6.5 OT, IoT, and Air-Gapped Environment Coverage

Cloud-connected, telemetry-dependent detection solutions address the needs of conventional IT environments well, but are structurally inadequate for air-gapped OT networks, industrial control systems, and medical device environments where files must be analysed offline, without cloud-based threat intelligence, and persistent endpoint agents. The growing convergence of IT and OT environments is creating new attack paths precisely in areas where detection coverage is thinnest — and where the consequences of a successful attack are most severe.



7. How Infosys Aims to Address Advanced Malware Detection

Infosys Cybersecurity Practice has built a broad portfolio of capabilities to help enterprise clients understand, detect, and respond to advanced malware threats, spanning managed security services, proprietary tooling, ongoing applied AI research, and regulatory compliance enablement. These capabilities are currently being developed with an explicit awareness of the gaps described above and are aimed to design and provide measurable value at each layer of the detection and response lifecycle.

7.1 Intelligence-Led Security Operations

Infosys operates globally distributed SOCs that integrate threat intelligence from commercial feeds, open-source repositories, and proprietary telemetry to provide contextualised and current threat awareness to client environments. SOC operations are structured around intelligence-led hunting methodologies — proactively searching client environments for indicators consistent with active threat actor campaigns — rather than relying solely on automated alert generation. This approach is particularly effective for detecting advanced malware that may evade automated classifiers but leaves behavioural artefacts that experienced analysts can identify through structured hypothesis-driven investigation.



7.2 AI-Augmented Detection and Response

Infosys is working towards embedding AI and ML capabilities across its managed detection and response offerings. Automated correlation of security events across multiple data sources, ML-assisted alert prioritisation that reduces the volume of alerts requiring analyst attention, and LLM-aided investigation workflows that accelerate the documentation and escalation of confirmed incidents are currently the focus. The integration of AI into SOC operations will be treated as an augmentation of analyst capability rather than a replacement for expert judgement, with human-in-the-loop architecture ensuring that AI outputs are reviewed and acted upon by experienced practitioners.

7.3 Research into Next-Generation Malware Analysis

Infosys invests in applied research into advanced malware detection techniques that aim to address the structural limitations of current commercial tooling. This includes investigation into multi-dimensional feature analysis that maintains detection accuracy even when individual feature channels are deliberately obfuscated by adversaries, architectures designed to maintain performance over time as the malware landscape evolves, and explainability frameworks that translate detection decisions into analyst-actionable intelligence. Research outcomes shall be applied directly to client engagements and, where appropriate, contribute to the broader academic and industry knowledge base. The company does not publicly disclose the specifics of proprietary methodologies under active development prior to formal publication, consistent with responsible disclosure and Intellectual Property (IP) protection practices.

7.4 Offline and OT-Capable Malware Screening

Recognising the specific challenges of operational technology environments, the firm's ongoing research is focused around offering malware analysis capabilities designed for offline deployment — operating without cloud connectivity, without persistent agent installation on target devices, and with detection models compact enough to deploy on standard gateway hardware. This will enable clients in manufacturing, energy, and healthcare to apply consistent malware screening at IT/OT boundaries without requiring changes to OT network architecture or violating air-gap requirements that are fundamental to the safety and integrity of those environments.

External Document © 2026 Infosys Limited

7.5 Regulatory Compliance and Documentation

Infosys supports clients across regulated verticals in meeting the documentation and reporting requirements of frameworks, including DORA, HIPAA, NIS2, and National Institute of Standards and Technology Cybersecurity Framework version 2. This includes the development of automated incident documentation capabilities that capture detection rationale, response actions, and containment outcomes in formats suitable for regulatory submission, reducing the compliance burden associated with malware incidents while improving the consistency of post-incident reporting. As regulatory requirements increasingly mandate not just that security incidents are reported but that the technical basis for detection and response decisions is documented, the ability to generate this documentation automatically from detection system outputs has become a meaningful competitive differentiator.

8. Strategic Recommendations for Enterprise Leaders

Based on the current threat landscape analysis, the documented impact of advanced malware incidents, and the state of detection technology, the following strategic priorities are recommended for enterprise cybersecurity leadership.



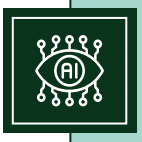
Treat model drift as a security risk with quantifiable impact. Establish quarterly performance assessments of AI-based detection tools against recent malware samples, define drift thresholds that trigger model refresh cycles, and hold vendors contractually accountable for maintaining stated detection rates on contemporary threats.



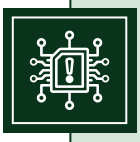
Demand explainability as a procurement requirement. Require security vendors to demonstrate per-alert feature attribution, not just confidence scores. Explainability is both an operational efficiency tool that reduces analyst triage time and a compliance enabler under DORA, NIS2, and HIPAA.



Layer detection approaches rather than consolidating to a single tool. No individual detection paradigm addresses all advanced malware vectors. A mature defence architecture combines static ML analysis, behavioural monitoring, network-level detection, and threat intelligence correlation, with results fused through an XDR or SIEM platform.



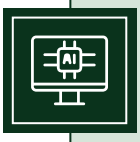
Invest in AI-augmented SOC capacity rather than attempting to replace human analysts. The goal of AI in security operations is to focus analyst attention on the subset of signals that genuinely require human judgement. Organisations that deploy AI without corresponding analyst training see limited operational benefit.



Address the OT/IT boundary explicitly in threat modelling. Map all data flows across IT/OT boundaries, assess the feasibility of offline malware screening at boundary points, and include OT assets in tabletop exercises. The most significant unaddressed exposure in many manufacturing, energy, and healthcare organisations lies precisely at this interface.



Build regulatory documentation capability before an incident occurs. Organisations unable to rapidly produce technically coherent incident documentation face compounded consequences: regulatory fines, reputational exposure, and additional scrutiny separate from the breach itself. Automated compliance reporting should be a standard feature of any enterprise incident response platform.



Monitor the adversarial AI landscape as a strategic risk input. The pace of AI capability development means that attack techniques available to sophisticated nation-state actors today will be available to lower-tier criminal actors within 18–24 months. Threat modelling should account for the democratisation of advanced malware capabilities that AI will continue to accelerate.

External Document © 2026 Infosys Limited

Conclusion

Advanced malware represents a defining challenge for enterprise cybersecurity in the AI ecosystem. The convergence of AI-assisted malware generation, RaaS (ransomware-as-a-service) criminal economies, and an expanding enterprise attack surface has elevated the stakes of detection failure to levels that are existential for some organisations and systemically significant for others. The incidents documented in this study — affecting healthcare delivery, automotive retail, casino operations, energy infrastructure, and thousands of organisations through supply chain compromise — share a common thread: sophisticated adversaries defeated enterprise-grade security tooling not through brute force but deliberate exploitation of architectural limitations in detection systems built for a simpler threat environment.

The path forward lies not in any single technological solution but in a coherent strategy that acknowledges the limitations of current approaches, invests in architectures designed to maintain performance as threats evolve, builds explainability into detection systems as a first-class requirement alongside raw accuracy, and treats analyst empowerment as the primary objective of AI adoption in security operations. Organisations that can answer not just whether a threat was detected, but why it was detected and what specifically in the file or behaviour triggered that assessment, will be significantly better positioned to reduce dwell time, satisfy regulatory obligations, and respond effectively when — not if — advanced malware attempts to breach their defences.

Organisations most resilient to advanced malware are not those with the highest security budgets, but those that have built detection architectures designed for adaptability, invested in analyst capability alongside tooling, and treated explainability and drift resilience as structural requirements rather than optional features.

References

1. IBM Security. Cost of a Data Breach Report 2024. IBM Corporation, 2024. ibm.com/reports/data-breach
2. IBM Security. X-Force Threat Intelligence Index 2024. IBM Corporation, 2024.
3. Verizon. Data Breach Investigations Report 2025. Verizon Business, 2025.
4. CrowdStrike. Global Threat Report 2025. CrowdStrike Holdings, 2025.
5. Kaspersky. Security Bulletin 2024. Kaspersky Lab, 2024.
6. Ponemon Institute. The Economics of the Insecure Remote Work Environment, 2024.
7. Dragos. ICS/OT Cybersecurity Year in Review 2024. Dragos Inc., 2024.
8. Jordaney, R. et al. Transcend: Detecting Concept Drift in Malware Classification Models. USENIX Security 2017.
9. UnitedHealth Group. FY2024 Annual Earnings Report. UHG Investor Relations, 2024.
10. EU Regulation 2022/2554. Digital Operational Resilience Act (DORA). Official Journal of the EU, 2023.
11. MITRE Corporation. MITRE ATT&CK Enterprise Matrix v14. attack.mitre.org, 2024.
12. Sophos. State of Ransomware in Manufacturing 2025. Sophos Ltd., 2025.
13. CISA. Cybersecurity Workforce Study 2023. Cybersecurity and Infrastructure Security Agency, 2023.
14. Anderson, H. & Roth, P. EMBER: An Open Dataset for Training Static PE Malware ML Models. arXiv:1804.04637, 2018.
15. Pendlebury, F. et al. TESSERACT: Eliminating Experimental Bias in Malware Classification. USENIX Security 2019.

External Document © 2026 Infosys Limited

About the Authors



Malvika Singh
Researcher

Malvika Singh is a Researcher at the Cyber Applied Research Centre, Infosys, specializing in cybersecurity research, AI-driven security analytics, and behavioural intelligence. Her work focuses on User and Entity Behaviour Analytics (UEBA), Insider Risk Management, Agent Behaviour Analysis, and Malware Detection to enable proactive threat detection and cyber resilience.

For more information, contact askus@infosys.com

Infosys
Navigate your next

© 2026 Infosys Limited, Bengaluru, India. All Rights Reserved. Infosys believes the information in this document is accurate as of its publication date; such information is subject to change without notice. Infosys acknowledges the proprietary rights of other companies to the trademarks, product names and such other intellectual property rights mentioned in this document. Except as expressly permitted, neither this documentation nor any part of it may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, printing, photocopying, recording or otherwise, without the prior permission of Infosys Limited and/ or any named intellectual property rights holders under this document.