wipro

The Mythos **Compression Effect:** **Redefining Cyber Risk in the Age of AI**

Executive Summary

Cybersecurity is undergoing a structural shift. Advances in AI are rapidly compressing the time between vulnerability discovery and exploitation, reducing to hours what was once measured in days or weeks. This acceleration is creating a widening asymmetry: attackers can identify and weaponise vulnerabilities at machine speed, while enterprise remediation cycles remain constrained by operational complexity.

This research paper highlights three critical shifts shaping the future of cyber risk.

First, Time-to-Exploit (TTE) is collapsing. AI-assisted tools are enabling threat actors to accelerate exploit development and deployment, leaving organisations with narrowing windows to detect, prioritise, and remediate vulnerabilities. Traditional vulnerability management approaches, designed for slower attack cycles, are no longer sufficient.

Second, the emergence of Mythos AI capabilities marks a new phase in the cyber arms race. These advanced systems can identify latent vulnerabilities, simulate exploit paths, and generate attack payloads at scale. While early access may provide defenders with a temporary advantage, this capability is expected to diffuse rapidly, intensifying competition between attackers and defenders.

Third, organisations must prepare for a near-term surge in vulnerability discovery. As AI reveals previously undetected weaknesses across platforms, infrastructure, and applications, visible vulnerability debt will increase before it stabilises. Over time, however, as AI is embedded across both discovery and remediation, enterprises can move towards a more resilient steady state.

This convergence gives rise to the Mythos Compression Effect—a structural shift where AI-driven acceleration across both vulnerability identification and remediation reduces the adversary's window of opportunity. Organisations that operationalise this effect will transition from reactive defence to continuous, machine-speed risk reduction.

To compete in this environment, organisations must:

- Embed AI across the end-to-end vulnerability lifecycle
- Integrate security, IT, and engineering operations
- Enable machine-speed prioritisation and remediation
- Align risk decisions to business criticality and impact

The Race against Zero Days

Cybersecurity teams have long been engaged in a persistent race to reduce the growing backlog of vulnerabilities across infrastructure and applications. This race can be understood through a simple timeline—from when a vulnerability is publicly disclosed and assigned a CVE, to when the corresponding patch is fully applied within the environment. An intermediate milestone marks the point at which a patch or fix becomes available from the vendor or application owner.

Exhibit 1

The Vulnerability Lifecycle and the Expanding Exposure Window

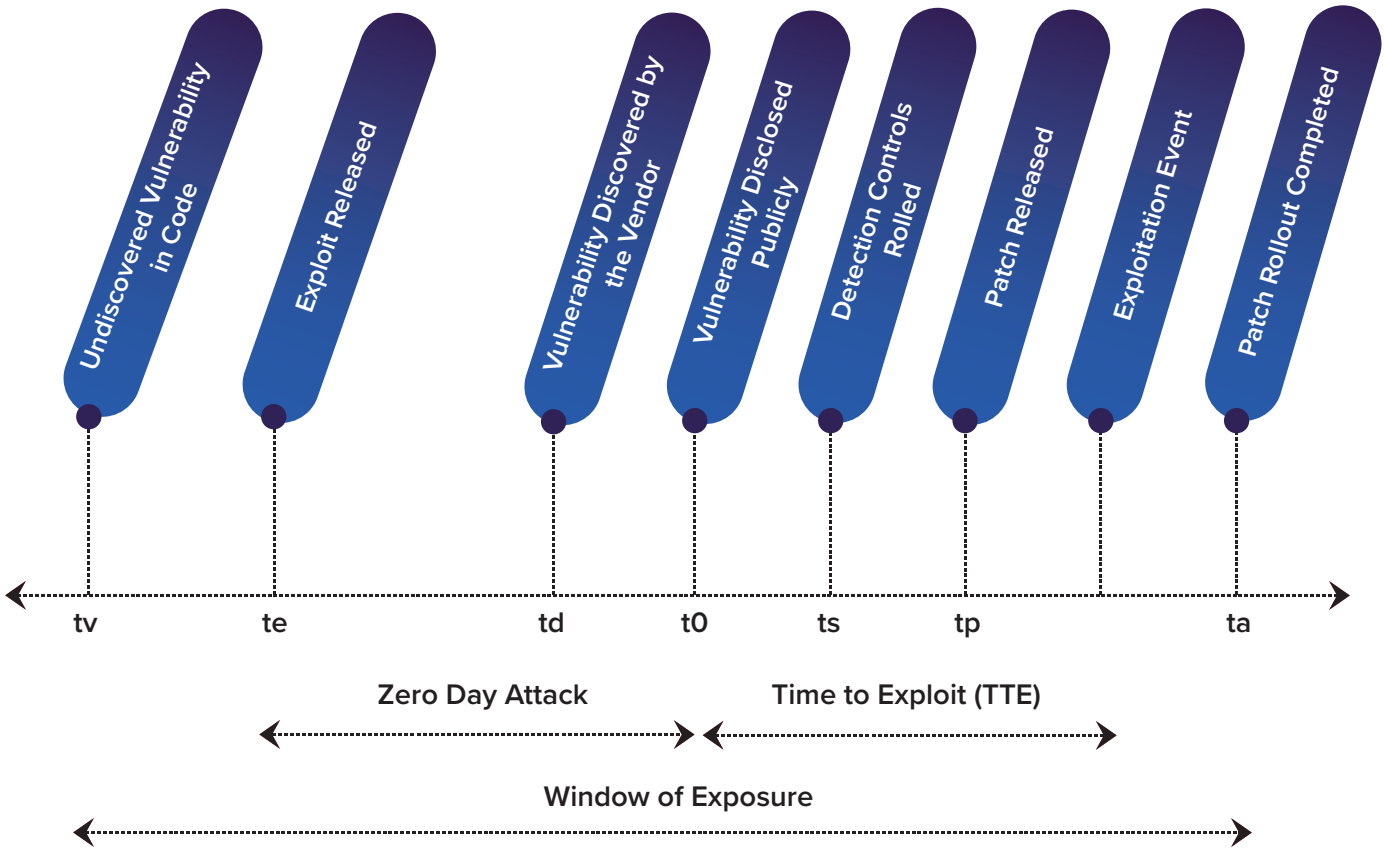

Time	tv	te	td	t0	ts	tp	ta
Description	Time when vulnerability introduced	Time when exploit is released	Time when vulnerability is discovered by vendor	Time when vulnerability disclosed publicly	Protection/Detection controls rolled	Time when patch is released	Time when patch is applied

The core objective is to compress the time between disclosure and remediation (t0 → ta in the exhibit above) This window determines an organisation's exposure to exploitation. While vulnerabilities may exist prior to disclosure, the period from disclosure to patch application defines the phase during which adversaries can reliably weaponise known weaknesses at scale.

In practice, this creates a structural imbalance. While defenders must navigate validation, testing, and deployment constraints before applying patches, threat actors can rapidly exploit disclosed vulnerabilities. As a result, the window of opportunity for attackers is expanding relative to the enterprise's ability to respond, making time the defining variable in cyber risk management.

To mitigate this exposure, organisations have historically relied on a combination of approaches, including:

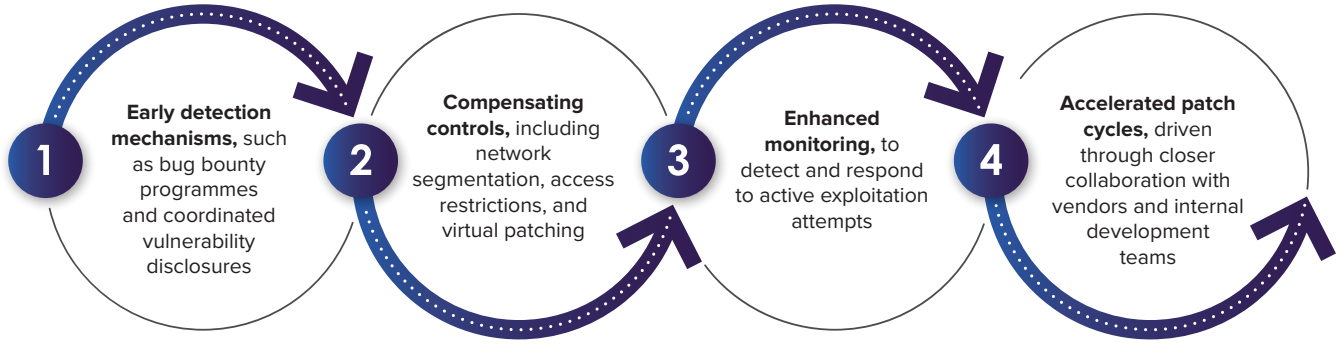

Despite these efforts, vulnerability management remains constrained by process friction and operational dependencies. As the pace of vulnerability discovery and exploit development accelerates, particularly in an AI-enabled threat landscape, the traditional race from **t0 to ta** is no longer sufficient. The challenge isn't just about keeping up, but fundamentally redefining how quickly organisations can detect, prioritise, and remediate risk at scale.

Exhibit 2

Declining Time-to-Exploit: A Structural Shift in Cyber Risk Dynamics

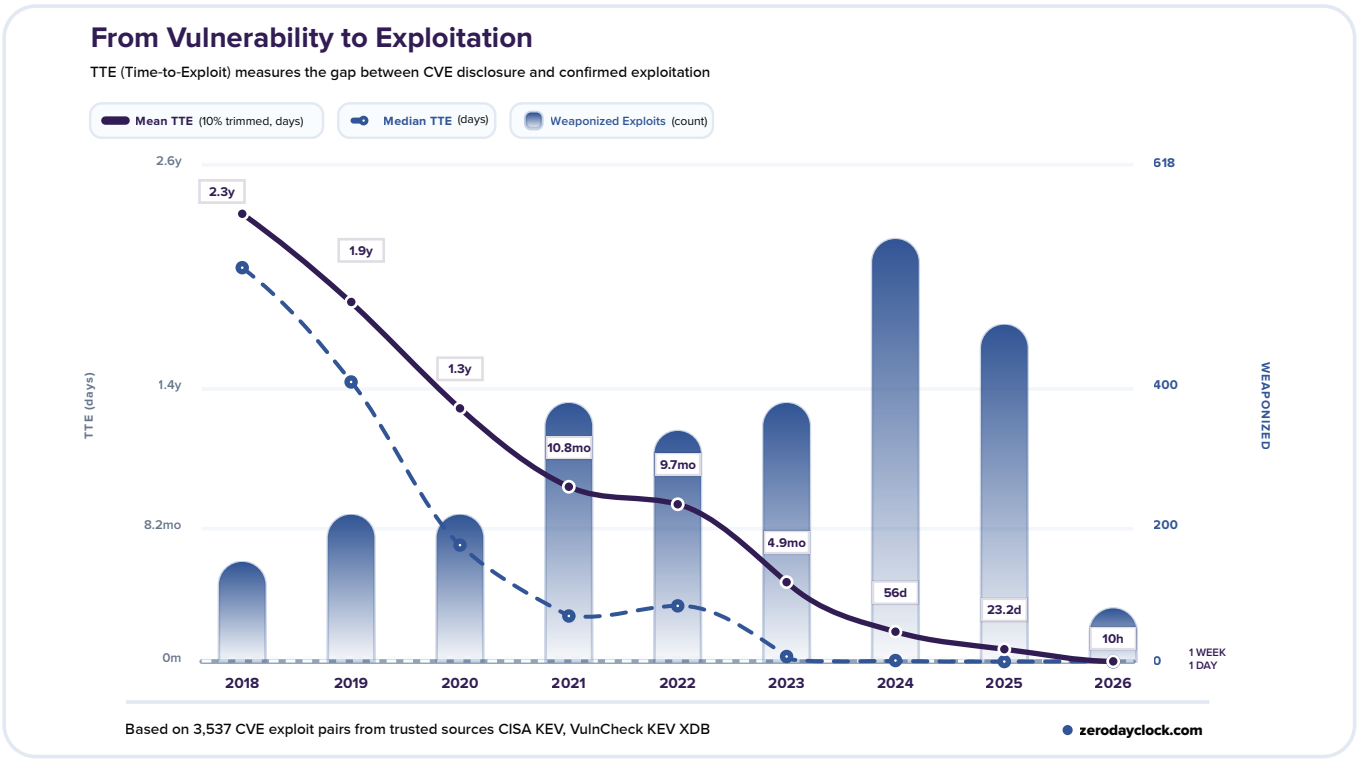

Source: Zero Day Clock. From Vulnerability to Exploitation. Available at: https://zerodayclock.com

Recent research by Zero Day Clock indicates a sharp and sustained decline in Time-to-Exploit (TTE) for known vulnerabilities, with current averages now compressed to 10–20 hours. This acceleration reflects a structural shift in the threat landscape, where adversaries are increasingly leveraging AI-assisted development tools to rapidly identify weaknesses, generate exploit code, and operationalise attacks at scale.

Despite this dramatic reduction in exploit timelines, enterprise response capabilities—particularly vulnerability discovery, validation, and patch deployment cycles—have not kept pace. This widening gap between attacker speed and defender readiness is creating a growing asymmetry in cyber risk.

The imbalance is unlikely to persist. As AI becomes more deeply embedded across both offensive and defensive workflows, the velocity of vulnerability discovery and remediation is also expected to accelerate, fundamentally reshaping how organisations manage exposure and respond to emerging threats.

The New Mythos AI Arms Race

Frontier AI models developed by leading technology providers are beginning to demonstrate advanced capabilities in identifying latent vulnerabilities, generating exploit pathways, and producing attack payloads at scale. This marks a fundamental shift in how vulnerabilities are discovered and weaponised. Early signals—such as Anthropic's reference to "Mythos"-class capabilities—point to a new class of AI systems that extend beyond traditional code scanning to reason across systems, anticipate failure points, and simulate attack scenarios.

In the near term, these capabilities are likely to be concentrated in the hands of defenders—including platform vendors and cybersecurity providers—enabling them to proactively identify and remediate previously unknown weaknesses. However, this advantage is unlikely to persist. As access to advanced models broadens across providers and geographies, similar capabilities will increasingly become available to adversaries, giving rise to a bi-directional, AI-accelerated cyber arms race.

This shift has profound implications for cyber risk management. The pace at which vulnerabilities can be identified and exploited will increasingly be determined by access to and effective deployment of AI capabilities. In this environment, velocity becomes the primary control variable. Organisations that can detect, prioritise, and remediate risk at machine speed will materially outperform those relying on traditional, human-led processes.

Critically, AI cannot remain confined to detection alone. As models accelerate vulnerability discovery, they must be deeply embedded across the remediation lifecycle, from patch generation and validation to automated deployment. Without this end-to-end integration, organisations risk amplifying their own vulnerability exposure, as discovery outpaces their capacity to act.

Bracing for the Mythos Vulnerability Surge

Emerging Mythos AI capabilities are beginning to identify latent vulnerabilities that elude conventional code scanning tools, particularly across complex, interdependent systems. Early controlled access to such capabilities for select cybersecurity providers and platform vendors is likely to trigger an initial wave of rapid vulnerability discovery and patch release, as previously undetected weaknesses are surfaced and addressed at scale.

Exhibit 3

The Mythos Effect: A Surge and Normalisation of Enterprise Vulnerability Exposure

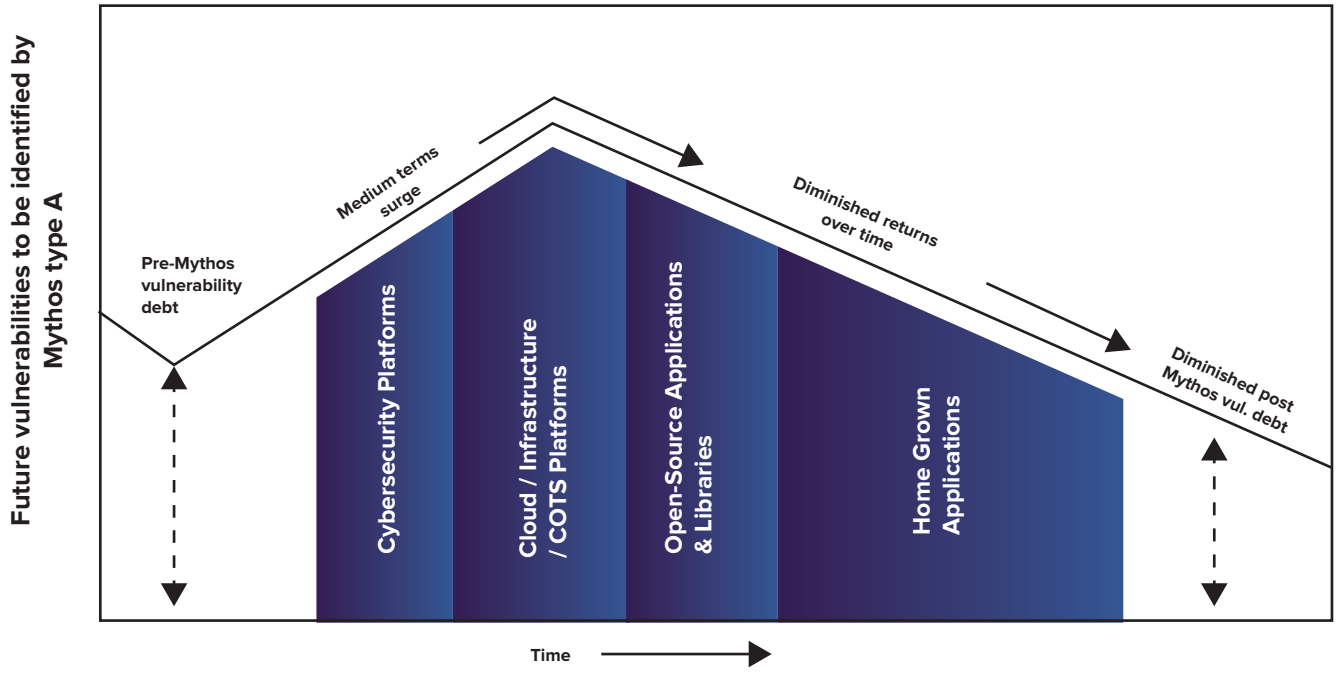

It is also expected that LLM providers from other geographies will catch up and we expect this capability to be more widely accessible. The effective leverage of Mythos and other equivalent AI models for hidden vulnerability identification and patching will happen progressively in overlapping waves across cyber providers, infrastructure providers and eventually for home grown applications in enterprises. At some point, this access will be table stakes and will reduce the mean level of vulnerability debt in enterprises.

Need for speed will blur the lines between security management and infrastructure management. Organizations and Managed Security Service Providers (MSSPs) will need to relook at operational structures to bring speed and efficiency in closing the vulnerability gap. Organizations that can invert the pyramid for L1 operations with AI across infrastructure would likely see quicker benefits of speed to execution.

As such capabilities become more widely available, access to advanced vulnerability discovery and remediation will become table stakes. This may also alter the economics of offensive cyber operations led by nation-states and quasi-state threat actors. The strategic value of hoarded zero day vulnerabilities is likely to decline as their discovery becomes more automated and widespread, potentially reducing their shelf life and effectiveness.

The Mythos Compression Effect

It is expected that at some point, enterprises will have access to the Mythos capability to augment their existing vulnerability management cycles, continuous threat exposure management processes, and patch management cycles.

Organisations that embed AI equally into vulnerability discovery and remediation workflows will be able to materially reduce the time between identification and resolution, effectively compressing the window of opportunity available to adversaries. This systemic reduction in exposure can be characterised as the Mythos Compression Effect: a structural shift in cybersecurity where AI-driven acceleration at both ends of the lifecycle converges to minimise exploitable risk.

The path to achieving this state is not without disruption. In the near term, the application of AI to vulnerability discovery will uncover large volumes of previously hidden or latent weaknesses, leading to a temporary increase in visible vulnerability debt. For many organisations, this will create operational strain as remediation capacity struggles to keep pace with discovery.

In this steady state, residual vulnerabilities will increasingly represent edge cases with diminishing returns on further optimisation, rather than systemic weaknesses. Organisations that successfully operationalise the Mythos Compression Effect will move beyond reactive vulnerability management to a model of continuous, machine-speed risk reduction.

Exhibit 4

The Mythos Compression Effect: Converging Detection and Remediation Timelines

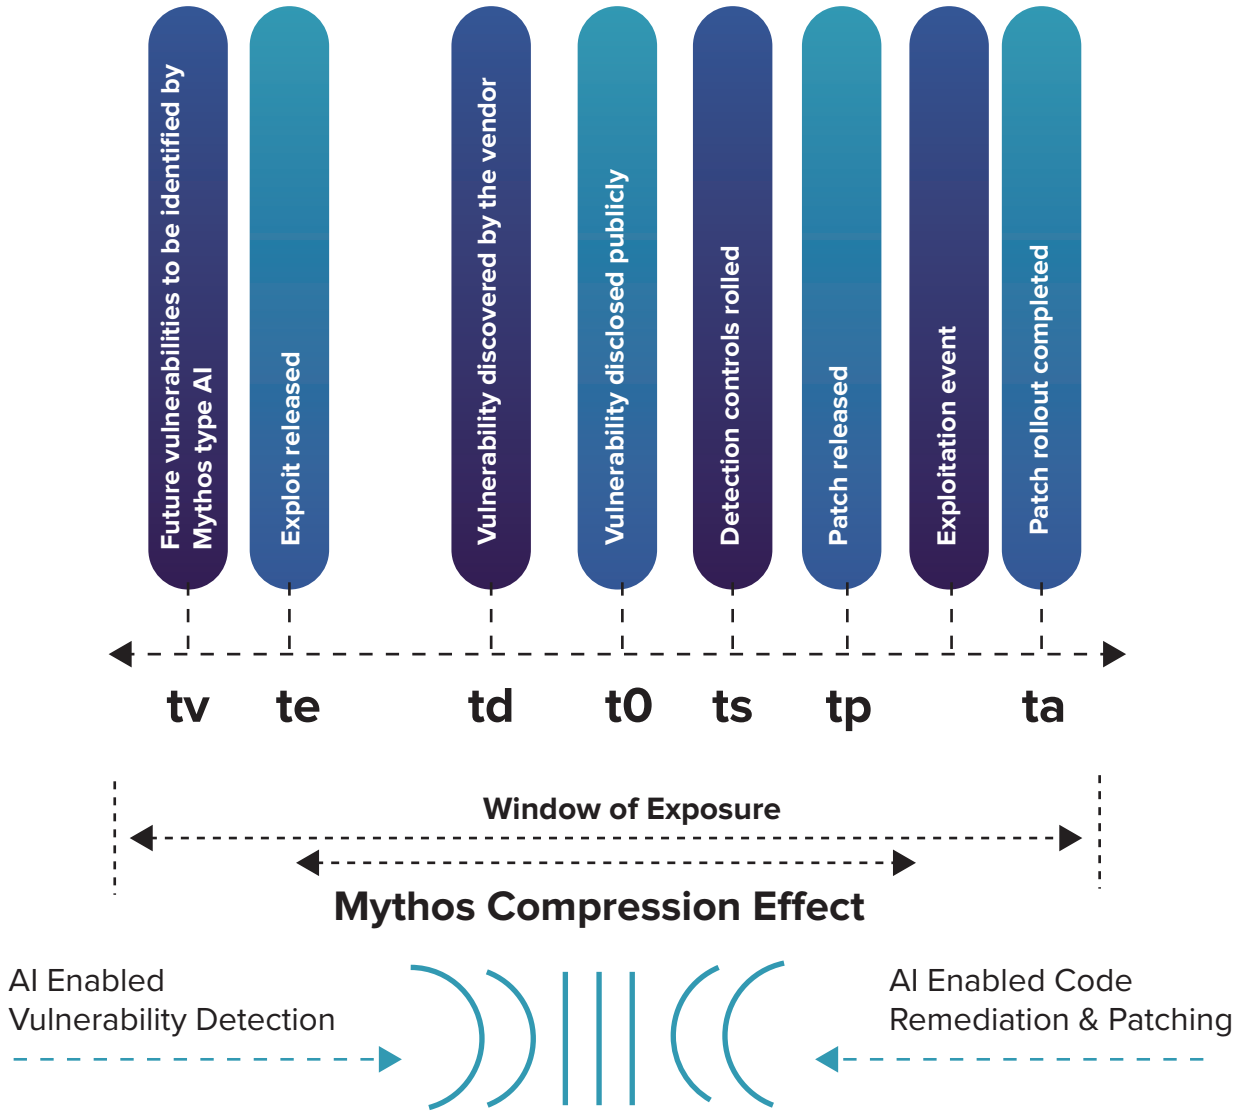

How can Wipro help?

Wipro enables enterprises to operate in a new cybersecurity reality where AI has fundamentally redefined the speed, scale, and complexity of risk. As threat actors leverage AI to accelerate vulnerability discovery, automate exploit development, and industrialise attacks, organisations must shift from reactive defence to machine speed resilience, reducing exposure windows without compromising uptime, regulatory compliance, or cost discipline.

Through a consulting-led, AI-powered approach, Wipro provides end-to-end solutions—from strategy and operating model design to implementation and managed execution—to help operationalise cyber resilience in an AI-driven threat landscape.

At the core of Wipro's approach is the ability to translate AI insight into decisive, prioritised action. Wipro's Secure AI capabilities enable continuous identification and quantification of exposures across infrastructures, applications, identities, and the emerging AI supply chain, ensuring that remediation efforts are based on exploitability, business criticality, and operational constraints. This enables leaders to balance risk reduction, cost optimisation, and operational continuity, while maintaining clear executive visibility into progress and outcomes.

Critically, Wipro helps organisations move towards operationalising the Mythos Compression Effect by embedding AI across both vulnerability discovery and remediation to systematically reduce the adversary's window of opportunity.

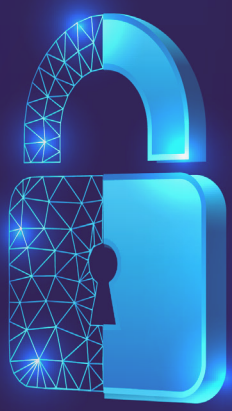

Key capabilities

Executive visibility and governance

AI-driven dashboards provide real-time insight into exposure, ownership, and remediation progress, enabling accountability and outcome-driven decision-making at the board and CISO level.

Machine-speed security operations

AI-enabled SOAR capabilities orchestrate detection, enrichment, triage, and response across tools and teams, driving speed, consistency, and scale in security operations.

AI-driven prioritisation and risk-based remediation

Continuous threat exposure management ensures remediation capacity is dynamically aligned to the most critical risks based on exploit likelihood and business impact.

Rapid, scalable remediation

Automated workflows reduce analyst effort, standardise responses, and accelerate containment and patch deployment across distributed environments.

48-hour response for critical vulnerabilities

Defined playbooks enable identification, prioritisation, and mitigation of high-severity vulnerabilities within accelerated timelines for critical scenarios.

AI supply chain governance

End-to-end controls across models, data, third-party dependencies, and software supply chains mitigate systemic and downstream risks in AI-enabled environments.

Cyber resilience by design

Resilient architectures, validated incident response, and recovery playbooks ensure organisations can sustain operations and recover rapidly from disruptions.

Wipro Limited
Doddakannelli
Sarjapur Road
Bengaluru – 560 035
India

Tel: +91 (80) 2844 0011
Fax: +91 (80) 2844 0256
wipro.com

Wipro Limited (NYSE: WIT, BSE: 507685, NSE: WIPRO) is a leading AI-powered technology services and consulting company focused on building innovative solutions that address clients' most complex digital transformation needs. Leveraging our consulting-led approach and the Wipro Intelligence™ unified suite of AI-powered platforms, solutions, and transformative offerings, we help clients realize their boldest ambitions to build intelligent and sustainable businesses.

The Wipro Innovation Network – part of the Wipro Intelligence™ suite – underpins our commitment to client-centric co-innovation and co-creation by bringing together capabilities from the innovation labs and partner labs, academia, and global tech communities. With over 240,000 employees and business partners across 65 countries, we deliver on the promise of helping our customers, colleagues, and communities thrive in an ever-changing world.

For additional information, visit us at www.wipro.com