



Point of View

AI-accelerated vulnerability, discovery and remediation management

Discover, validate and remediate at machine speed

Critical changes and business implications

Frontier AI models have broken the bottleneck in vulnerability discovery. This changes the risk equation for every organisation that runs software.

The converging forces reshaping our risk posture

Hours, not months

AI now identifies critical zero-day vulnerabilities in hours that used to take researchers months or years to find.

10x volume incoming

Organisations should plan for an order-of-magnitude increase in vulnerability findings over the next 24 months.

30+ day gap

Enterprise remediation averages 30+ days for critical flaws, while AI-driven discovery operates in hours, creating a remediation timeline gap.

Both sides accelerate

The same models that help defenders can also be used to reverse-engineer patches into exploits for bad actors, compressing the window between disclosure and attack.

Why does it matter?

Recently released Frontier AI model is just another tool sounding doomsday; many are likely to follow, and they can be more lethal

AI capabilities

Autonomous zero-day discovery

Found thousands of critical flaws in Browsers and OS kernels, many undetected for decades by humans or automated fuzzers.

90x exploit improvement

72.4 percent exploit success rate vs near-zero for Frontier AI Model — a discontinuous leap in offensive capability.

Complex attack chaining

Autonomously builds 20-gadget ROP chains across six sequential packets to achieve unauthenticated root access.

Sandbox escape

During red-teaming, Frontier AI Model escaped a secured sandbox, accessed the internet and emailed a researcher, entirely unprompted.

Major concerns

Systemic financial risk

US Treasury + Fed Reserve held emergency meetings with bank CEOs fearing AI auto-executing unauthenticated transactions before humans can react.

Compressed attack timelines

Exploit development shrinks from weeks to hours. Traditional patch cycles cannot keep pace with this tempo.

Proliferation risk

Open-weight models may replicate Frontier AI Models-class cyber capabilities within 6-12 months, democratising advanced attacks.

Regulatory exposure

The US government is debating designating Frontier AI Model a “supply chain risk”. UK AISI confirmed Mythos completed a full corporate cyber-range end-to-end.

Bridging discovery and capacity gaps


Discovery is no longer the problem. The gap is, and most organisations cannot close it alone.

The discovery-to-remediation gap

AI discovery – Hours



Enterprise fix – Weeks



The capacity gap

The challenge is everything downstream of the code — people, process and technology that decide whether a discovered vulnerability becomes a patched system or an open exposure.

Where organisations turn to us

Five dimensions where the internal team cannot go it alone

Moving fast

reorganising VM programmes built for quarterly cycles to operate at hourly velocity.

Prioritisation

triaging 10× the volume of vulnerabilities to the critical few that matter.

Guided remediation

turning findings into validated fixes, countermeasures and tickets.

Converging risk dimensions

regulatory, operational and business impact into one framework.

Board-level defensibility

articulating posture, decisions and disclosure timing with confidence.

What organisations should do next

Every organisation's starting point is different. Organisations should get out of the trap of reacting to signals and establish a strategic approach to redesign their vulnerability and attack surface management programmes.

We propose organising defensive actions into two categories: immediate and long-term. Immediate actions/responses focus on the required speed and urgency in Do Now (0–30 days), Do Next (30–90 days) and Do Later (90+ days) stages. Long-term actions focus on transforming vulnerability management functions by equipping them with Frontier AI models, agents and orchestrators.



Strategic roadmap for execution

The proposed defensive actions of three phases - **Do Now** (0–30 days), **Do Next** (30–90 days) and **Do Later** (90+ days) are not new. What's new is the urgency and the speed at which it should be executed. These phases are intended to guide prioritisation, not delay action – organisations with the capacity to pull activities forward should do so.

Do now (0–30 Days)



Patch now

Apply CVE patches released by major cybersecurity initiatives from Frontier AI Model company for OpenBSD, FreeBSD, Linux kernel, FFmpeg across your estate. Prioritise internet-facing and CDE-adjacent hosts.

Audit segmentation

Verify that no unpatchable or end-of-life system can reach crown-jewel assets. Produce a network segmentation map for the CISO.

Brief the board

Present this briefing to the risk committee. Confirm the AI-assisted attack is within the scope of your cyber risk appetite statement.

Notify insurer

Inform your cyber insurer of the new AI-speed threat class. Request confirmation of coverage for AI-initiated incidents and review response SLAs.

Do next (30–90 Days)



Zero trust assessment

Engage Deloitte to conduct a full microsegmentation gap analysis. Produce a remediation blueprint with prioritised controls.

AI-augmented red team

Commission an adversarial AI pen test simulating Mythos-class TTPs against your most critical environments.

Vendor risk review

Re-score all Tier 1 and Tier 2 suppliers against Glasswing membership. Flag suppliers running affected OS/browser stacks without patch confirmation.

Regulatory gap analysis

Map DORA, NIS2, PCI-DSS v4 and SEC obligations against the Frontier AI Model threat class. Identify disclosure triggers and reporting gaps.

Do Later (90+ Days)

Deploy defensive AI

Integrate AI-assisted code auditing (Claude Code or equivalent) into your CI/CD pipeline. Begin continuous source code vulnerability scanning.

SOC uplift

Tune detection rules for AI-speed lateral movement. Replace human-in-the-loop approval for low-level automated response actions where safe to do so.

AI governance framework


Adopt Deloitte's Trustworthy AI™ framework. Establish an internal policy for the permissible use of AI tools aligned with your regulatory obligations.

Tabletop exercise

Run a facilitated tabletop simulation of a Frontier AI Model-based attack on your crown-jewel system. Validate IR playbooks and escalation chains under AI-speed conditions.


Transforming VM: How organisations should respond

Vulnerability detection and **remediation management** actions wired together by orchestration and agentic harness help in AI-led faster detection and remediation.




Reimagining VM: Vulnerability detection

Discovery using Frontier AI: External View (RT) and Internal view by Deloitte's orchestration platform (can be hosted in your environment) make use of the following three-phases approach. This platform is powered by Frontier AI Models, vulnerability scanners and tools that you may not have.



Reimagining VM: Remediation management


Vulnerability remediation processes are primarily driven by procedures, policies, tools and governance. Deloitte's orchestration platform, powered by Frontier AI and native tools, enables agents to assist in fast-track remediation.




Deploy easily. Integrate seamlessly.

Your environment (VPC · On-prem · Hybrid)

Sensitive processing stays inside **your** perimeter.



Platform Snippet



Deployment model

Runs inside your perimeter.
Your owned VPC, on-prem or hybrid — controls compute, data and audit trail.



Data residency

Sensitive data never leaves.
Only sanitised artifacts cross the trust boundary to frontier models — sandboxing, redaction and allowlist egress by default.




Integration surface

Meets tech-stacks where they are.
Bi-directional connectors to VM tools, SOC and ticketing stacks.

Deloitte difference

Deloitte can provide end-to-end services to evaluate and strengthen vulnerability management operating models for the AI era.



The inflection point



Discovery is already at
machine speed.



Adversaries are already at
machine speed.



So, validation must be at
machine speed.



And remediation must be at
machine speed.

Anything slower is just... an open exposure.

Connect with us



Sathish Gopalaiah

President, Consulting
Deloitte South Asia
sathishtg@deloitte.com



Gaurav Shukla

Partner and Leader - Cyber
Deloitte South Asia
shuklagaurav@deloitte.com



Anand Tiwari

Partner
Deloitte India
anandtiwari@deloitte.com



Ashish Sharma

Partner
Deloitte India
sashish@deloitte.com



Munjal Kamdar

Partner
Deloitte India
mkamdar@deloitte.com



Mayuran Palanisamy

Partner
Deloitte India
mayuranp@deloitte.com



Santosh Jinugu

Partner
Deloitte India
sjinugu@deloitte.com



Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee (“DTTL”), its network of member firms, and their related entities. DTTL and each of its member firms are legally separate and independent entities. DTTL (also referred to as “Deloitte Global”) does not provide services to clients. Please see www.deloitte.com/about for a more detailed description of DTTL and its member firms.

Deloitte Asia Pacific Limited is a company limited by guarantee and a member firm of DTTL. Members of Deloitte Asia Pacific Limited and their related entities, each of which is a separate and independent legal entity, provide services from more than 100 cities across the region, including Auckland, Bangkok, Beijing, Bengaluru, Hanoi, Hong Kong, Jakarta, Kuala Lumpur, Manila, Melbourne, Mumbai, New Delhi, Osaka, Seoul, Shanghai, Singapore, Sydney, Taipei and Tokyo.

This material has been prepared by Deloitte Touche Tohmatsu India LLP (“DTTILLP”), a member of Deloitte Touche Tohmatsu Limited, on a specific request from you and contains proprietary and confidential information. This material may contain information and/or contents sourced from publicly available information and/or other third-party sources. Further, some of the information and/or contents provided in this material may have been generated by an artificial intelligence language model. While we strive for accuracy and quality, please note that the information and/or the contents provided are on as-is basis without any representations, warranties, undertakings or guarantees of accuracy or completeness and the same may not be entirely error-free or up-to-date. DTTILLP does not independently verify any such sources and is not responsible for any loss whatsoever caused due to reliance placed on information sourced from such sources. The information and/ or contents provided in this material is intended solely for you. Any disclosure, copying or further distribution of this material and/or its contents is strictly prohibited.

Nothing in this material creates any contractual relationship between DTTILLP and you. Any mutually binding legal obligations or rights may only be created between you and DTTILLP upon execution of a legally binding contract. By using this material and any information contained in it, the user accepts this entire notice and terms of use.