AI Agents Got Root on Linux via 948 Kernel Flaws in OpenAI Research

An unreleased OpenAI cyber model found and exploited 948 distinct Linux kernel privilege escalation bugs for about $1.2 million.

The post AI Agents Got Root on Linux via 948 Kernel Flaws in OpenAI Research appeared first on SecurityWeek.

https://www.securityweek.com/ai-agents-got-root-on-linux-via-948-kernel-flaws-in-openai-research/

Bank Phishing Sites Show Fake ‘Suspended’ Pages to Evade Takedowns

More than 200 sites were seen serving phishing content to first-time visitors but not to those who revisit.

The post Bank Phishing Sites Show Fake ‘Suspended’ Pages to Evade Takedowns appeared first on SecurityWeek.

https://www.securityweek.com/bank-phishing-sites-show-fake-suspended-pages-to-evade-takedowns/

Cyberattack Hits French Blood Bank Site but No Blood Products Are Compromised

The incident comes after supermarkets in a number of European countries reported this week that their fridges and freezers warmed up after they were remotely hacked.

The post Cyberattack Hits French Blood Bank Site but No Blood Products Are Compromised appeared first on SecurityWeek.

https://www.securityweek.com/cyberattack-hits-french-blood-bank-site-but-no-blood-products-are-compromised/

Cisco, WatchGuard Release Patches; SonicWall Won’t Fix GMS Flaws

Cisco patched an RCE in UCS Central Software; WatchGuard fixed authentication and MitM flaws.

The post Cisco, WatchGuard Release Patches; SonicWall Won’t Fix GMS Flaws appeared first on SecurityWeek.

https://www.securityweek.com/cisco-watchguard-release-patches-sonicwall-wont-fix-gms-flaws/

The Agent Smith Dilemma: Governing Agentic Identities

Every autonomous AI agent deployed across your enterprise today is a new proxy identity operating inside your perimeter.

The post The Agent Smith Dilemma: Governing Agentic Identities appeared first on SecurityWeek.

https://www.securityweek.com/the-agent-smith-dilemma-governing-agentic-identities/

CISO Conversations: Ryan Davis – The Unknown Is Where Security Meets Reality

Ryan Davis went from hacking his school’s network to running security at New Charter Technologies. What worries him now is shadow AI.

The post CISO Conversations: Ryan Davis – The Unknown Is Where Security Meets Reality appeared first on SecurityWeek.

https://www.securityweek.com/ciso-conversations-ryan-davis-the-unknown-is-where-security-meets-reality/

Zyxel Switch Hacks Blamed on Patched Flaw Actually Exploited a Zero-Day

Tracked as CVE-2026-107441, the out-of-bounds write flaw could allow attackers to execute OS commands via crafted HTTP requests.

The post Zyxel Switch Hacks Blamed on Patched Flaw Actually Exploited a Zero-Day appeared first on SecurityWeek.

https://www.securityweek.com/zyxel-switch-hacks-blamed-on-patched-flaw-actually-exploited-a-zero-day/

Russia-Linked Hackers Hit iOS, Android Devices With Exploit Chains via Watering Holes

The large-scale watering hole campaign has been targeting entities in the US and Europe, mainly in Ukraine.

The post Russia-Linked Hackers Hit iOS, Android Devices With Exploit Chains via Watering Holes appeared first on SecurityWeek.

https://www.securityweek.com/russia-linked-hackers-hit-ios-android-devices-with-exploit-chains-via-watering-holes/

OpenAI Finds GPT-6 Astra Quietly Holding Back on Internal Red-Teaming, Exploit Tasks

OpenAI says its Critical-rated cyber model met fewer hidden task goals when work was framed as AI safety research, often silently.

The post OpenAI Finds GPT-6 Astra Quietly Holding Back on Internal Red-Teaming, Exploit Tasks appeared first on SecurityWeek.

https://www.securityweek.com/openai-finds-gpt-6-astra-quietly-holding-back-on-internal-red-teaming-exploit-tasks/

Sysdig Adds Runtime Security and Governance for AI Agents

Sysdig’s platform uses runtime visibility to track AI agent actions and block risky behavior before it reaches production systems.

The post Sysdig Adds Runtime Security and Governance for AI Agents appeared first on SecurityWeek.

https://www.securityweek.com/sysdig-adds-runtime-security-and-

governance-for-ai-agents/