AI Trends

2026-W43 · reporting 19 Oct – 25 Oct 2026
Makani Lab · AI Trends weekly
DoD · China / Indo-Pacific · Cyber · Space · Frontier labs & policy

BLUF — AI-found software flaws turned into live attacks this week: PHP2Shell leaked before patches were out, and criminal AI-agent swarms were caught hitting US firms. The Pentagon made 'unbiased AI' rules mandatory in LLM buys and threatened export controls on labs that don't comply. Attention item: AI-speed exploitation of web-facing systems.

By the numbers · 7d (vs prior) DoD / military AI **16** (+4) · China / Indo-Pacific **13** (+2) · Cyber + AI **10** (+2) · Space + AI **5** (+2) | **Rising** Funding & markets +6, Military & defense +4, Policy & regulation +4 | **Most-covered labs** Anthropic 15, OpenAI 13, Claude 12 | 39 items from 9 sources

DoD / military AI

adoption, programs, C2 and targeting

Pentagon makes 'unbiased AI' rules mandatory in every LLM procurement
An Oct. 20 DPCAP memo gives the services, SOCOM, CYBERCOM and TRANSCOM templates for solicitations and SOW language. Vendors must document that their LLMs comply with the 'Preventing Woke AI' order. (DefenseScoop · 23 Oct)
So what: Every new AI buy, including space C2 and intel tools, now carries extra compliance paperwork.

Marine Corps opens Maven Smart System accounts at every echelon
A message effective Oct. 1 gives authorized personnel their own Palantir MSS accounts on networks from unclassified to top secret. MARFOR headquarters get accounts through their combatant command. (DefenseScoop · 20 Oct)
So what: MSS is becoming the joint common picture that space effects will have to feed into.

China / Indo-Pacific

PRC AI-military, US-China AI diplomacy

Beijing's draft export controls cover frontier models and wargaming AI
MOFCOM's Announcement 48 adds 40 items: 14 bans and 26 licensed exports. They include models over 10^26 FLOP or 2T parameters and 'AI decision-making for strategic interactions'. Pre-export registration is in force now. (ChinaTalk · 21 Oct; Interconnects · 20 Oct)
So what: Beijing now treats wargaming and decision AI as strategic technology, the same way DoD does.

DHS orders federal civilian agencies to remove DeepSeek models and tools
The DHS notice bars new purchases and requires removing DeepSeek, including third-party tools built on it, from civilian systems. Interconnects asks whether models trained on DeepSeek outputs are covered. (The Neuron · 20 Oct; Interconnects · 20 Oct)
So what: Chinese open models are leaving US government systems; expect DoD to follow suit.

Cyber + AI

AI-enabled intrusion and defense

AI-found PHP flaw leaks before patching and is exploited millions of times
GPT 5.6 Sol and Wiz found PHP2Shell, which may affect up to 70% of websites. It leaked ten hours before disclosure. Cloudflare saw millions of exploit attempts, and the NSA advised taking systems offline until patched. (AI Daily Brief · 22 Oct; Transformer · 23 Oct)
So what: Disclosure-to-exploit time is down to hours; web-facing mission systems need same-day patching.

Exposed panels show criminal AI-agent swarms breaking into US firms
NorForge found 20+ open Cairn control panels; one Chinese-language cluster ran 1,100+ attack projects. It estimates a 10% hit rate at about $5 a try. The Verge named John Deere, SharkNinja and Cummins as targets. (The Neuron · 20 Oct)
So what: Cheap autonomous intrusion lets small crews probe at scale; defense-industrial suppliers are exposed.

Frontier labs & policy

lab and government decisions that bear on the force

Pentagon's Emil Michael threatens export controls on labs that defy the AI task force
The Information reported, via AI Daily Brief, that he would use 'the full range' of executive powers, including export controls and limits on foreign nationals. He also named Chinese open models as a top concern. (AI Daily Brief · 20 Oct)
So what: DoD is the lever on frontier labs, and access to their models may follow from compliance.

Nothing significant this week: Space + AI

WATCH NEXT WEEK

- Kimi K3.1 Pro open weights, due by Nov. 4, with offensive cyber held back by design.
- D.C. Circuit on rehearing Anthropic's supply-chain-risk designation en banc.
- Next AI-found flaw disclosure and whether the patch lands before the leak.