The Cyber Coalition: Why Rival AI Companies Are Suddenly Calling for Collective Defense.

Vladimir Tsakanyan, Ph.D. · Center for Cyber Diplomacy and International Security

August 31, 2026

More than 100 organizations have joined an unusual appeal for a global surge in cyber defense. The important development is not another warning about dangerous AI. It is the emergence of a new public-private security architecture around the technology.

By Vladimir Tsakanyan, Ph.D. Center for Cyber Diplomacy and International Security (CCD-IS)

Something unusual happened in the technology industry on August 27.

Companies competing intensely for leadership in artificial intelligence agreed on something important.

OpenAI, Anthropic, Google, Microsoft, Amazon Web Services and more than 100 other organizations joined an open letter calling for a global surge in cyber defense.

The coalition extends far beyond AI developers. Its signatories include cybersecurity companies, telecommunications providers, cloud and software companies, financial institutions, critical-infrastructure specialists and other major technology organizations.

Their warning is direct: as increasingly capable AI models spread, AI-enabled cyberattacks are expected to become more widespread and sophisticated.

But the warning itself is not the most important development.

Cybersecurity officials have been discussing AI-enabled attacks for years.

What is new is the proposed response.

The companies are effectively arguing that neither governments nor private organizations can manage the next stage of cyber risk independently.

They are proposing something closer to collective cyber defense.

That represents a significant change in the politics of AI security.

From Competition to Cooperation

The AI industry is currently defined by competition.

Companies are competing for models, computing infrastructure, enterprise customers, developers, investment and technological leadership.

Governments are simultaneously treating advanced AI as an element of economic and national power.

Against that background, the breadth of the August 27 coalition is notable.

Competitors that disagree over models, safety approaches, openness and commercial strategy are aligning around a common cybersecurity proposition:

the existing defensive environment may not be prepared for the speed at which AI cyber capability is developing.

The letter says organizations have a limited window to strengthen their defenses.

Its authors are not predicting a particular cyber catastrophe on a particular date.

Their argument is more structural.

As models improve, capabilities that previously required experienced cyber operators may become faster and cheaper to perform.

That could increase the number of actors capable of conducting meaningful cyber operations.

And it could happen while enormous amounts of digital infrastructure remain dependent on longstanding vulnerabilities, weak authentication, excessive privileges, insecure configurations and legacy software.

AI is therefore colliding with decades of accumulated technical debt.

The Weaknesses Are Already Here

This distinction is essential.

The emerging cyber problem is not simply that artificial intelligence creates new vulnerabilities.

Many of the vulnerabilities already exist.

Organizations routinely operate outdated systems.

Known vulnerabilities remain unpatched.

Identity systems grant excessive permissions.

Critical infrastructure sometimes depends on technology that cannot easily be upgraded without interrupting essential services.

Security teams remain understaffed.

Municipal governments, hospitals and utilities frequently operate with cybersecurity resources dramatically smaller than those available to major technology companies.

AI changes the equation because it can potentially reduce the cost of finding and exploiting those weaknesses.

The attack surface does not have to become larger for the danger to increase.

The ability to search that attack surface can become more powerful.

That is why the coalition is emphasizing remediation rather than simply calling for more AI regulation.

Its central message is essentially:

fix the weaknesses before increasingly capable machines become better at finding them.

The Defender's Window

This creates what the industry increasingly describes as a defender's window.

The same AI capabilities that could assist attackers can also help defenders.

AI can analyze software.

It can assist vulnerability research.

It can review configurations.

It can help investigate incidents.

It can process security telemetry.

It can support patch development and remediation.

The strategic question is therefore not simply whether AI favors offense or defense.

It is whether defenders can use advanced capability to reduce the existing attack surface before comparable capability becomes broadly available to attackers.

That turns cybersecurity into a race against diffusion.

If defensive AI reaches vulnerable organizations quickly, society may be able to eliminate large numbers of weaknesses before attackers can exploit them at machine speed.

If defensive capability remains concentrated among a handful of sophisticated organizations, the opposite may occur.

Advanced attackers could gain increasingly powerful tools while hospitals, utilities, municipalities and smaller businesses continue operating vulnerable infrastructure.

The outcome therefore depends not only on AI capability.

It depends on access.

The Most Important Proposal Is About Distribution

This may be the most consequential aspect of the August 27 letter.

The signatories are not merely calling for better cybersecurity products.

They are arguing that cyber-capable AI should be placed into the hands of more defenders.

That includes organizations that historically have not had access to elite cybersecurity capabilities.

The letter specifically calls for support for hospitals, water utilities, local governments and critical-infrastructure supply chains.

This represents an important shift.

Frontier cyber capability has until recently been discussed largely as something that may need to be restricted.

That remains a legitimate concern.

A model capable of discovering vulnerabilities or assisting sophisticated penetration testing can potentially be misused.

But excessive restriction creates another risk.

If only the largest companies and government agencies have access to the strongest defensive systems, then the organizations most vulnerable to attack may remain the least capable of defending themselves.

Cybersecurity policy therefore faces a difficult distribution problem.

Capability must be restricted enough to reduce misuse while being distributed widely enough to produce meaningful defensive advantage.

That balance will become one of the central questions of AI security governance.

Government Is Being Assigned a New Role

The letter also places substantial responsibility on governments.

It calls for coordination at local, national and international levels.

Governments are encouraged to strengthen threat-intelligence sharing, fund cyber defense, support critical infrastructure, expand trusted-access programs and improve incident-response coordination.

This is important because it moves the discussion beyond conventional regulation.

The government is not being asked simply to establish rules for AI companies.

It is being asked to become an organizer of defensive capability.

That means identifying vulnerable sectors.

Funding security improvements.

Connecting intelligence with private defenders.

Determining who qualifies for trusted access to advanced cyber models.

Coordinating response during major incidents.

And helping smaller institutions acquire capabilities they could not independently afford.

This is closer to national resilience policy than conventional technology regulation.

Cybersecurity Is Becoming a Collective-Action Problem

There is a deeper reason this approach is emerging.

Digital security is interconnected.

A vulnerable supplier can expose its customers.

A compromised software library can affect thousands of organizations.

A telecommunications weakness can create risk across entire regions.

A compromised identity provider can become an entry point into otherwise well-defended networks.

One organization's cybersecurity therefore affects others.

AI increases the importance of this interdependence because machine-assisted vulnerability discovery can potentially identify common weaknesses across many organizations more rapidly.

The security model consequently begins to resemble collective defense.

The objective is not simply to protect one institution.

It is to raise the cost of attack across the ecosystem.

That requires information sharing, coordinated remediation and the rapid distribution of defensive knowledge.

If one organization discovers a serious vulnerability, thousands of others should not have to rediscover the same problem independently.

AI potentially makes that model much more powerful.

A verified defensive finding could be translated into detection logic, remediation guidance or automated testing and distributed rapidly across many organizations.

One defender's discovery could strengthen many defenders.

That is a fundamentally different security model from every organization defending its own perimeter independently.

The Cyber-Diplomatic Consequence

The letter also explicitly calls for international coordination.

That matters.

AI cyber capability will not remain within national borders.

Models developed in one country can be accessed from another.

Software supply chains are international.

Cloud infrastructure crosses jurisdictions.

Cybercriminal organizations operate transnationally.

Critical infrastructure depends on foreign technology.

A purely national defensive strategy therefore has obvious limits.

The emerging challenge for cyber diplomacy is not simply establishing norms against malicious cyber behavior.

It is building mechanisms through which defensive capability can move across borders.

That could include faster international vulnerability disclosure.

Joint AI-assisted security testing.

Shared threat intelligence.

Emergency access to defensive models during major incidents.

Support for countries with weaker cybersecurity capacity.

Common standards for trusted cyber-AI access.

And coordination when an AI-enabled campaign affects infrastructure across multiple jurisdictions.

This moves cyber diplomacy from norm-building toward operational cooperation.

That transition is significant.

For much of the last decade, international cyber diplomacy has concentrated on defining responsible state behavior.

The AI era may increasingly require states to develop shared defensive capacity as well.

Private Companies Are Becoming Security Institutions

There is also an uncomfortable political reality behind the coalition.

Many of the capabilities required for this defensive strategy are not controlled by governments.

They belong to private companies.

Frontier AI models are privately developed.

Cloud infrastructure is largely privately operated.

Many of the world's strongest cyber-response teams work for private cybersecurity companies.

Telecommunications infrastructure is often commercial.

Much critical software is developed outside government.

The result is a redistribution of security power.

Governments retain sovereignty and legal authority.

But private companies increasingly control infrastructure and capabilities necessary for national cyber defense.

The August 27 coalition makes this relationship unusually visible.

The companies are not simply selling security products.

They are proposing roles for themselves within a broader security architecture.

That creates important governance questions.

Who determines which defenders receive frontier capability?

Who evaluates whether access is safe?

What information must companies share with governments?

What happens when commercial priorities conflict with national-security priorities?

And how much of society's cyber defense should depend on technology controlled by a relatively small number of corporations?

Those questions will become increasingly difficult to avoid.

The Coalition Also Has a Conflict of Interest

The industry's warning should not be accepted without scrutiny.

Many of the companies calling for increased investment in AI-enabled cyber defense also develop the technology that would benefit from that investment.

AI companies may gain customers.

Cybersecurity firms may sell new defensive products.

Cloud providers may host them.

Consulting firms may implement them.

That does not make the threat assessment wrong.

But it means policymakers should distinguish evidence from advocacy.

Predictions that AI-enabled attacks will become dramatically more sophisticated within months remain forecasts.

Timelines are uncertain.

Different models possess different capabilities.

And not every cyber operation becomes substantially more dangerous merely because AI is involved.

The strongest case for action therefore does not depend on accepting the most dramatic prediction.

Many of the measures proposed by the coalition—patching vulnerabilities, reducing excessive privileges, strengthening authentication, replacing insecure legacy systems, improving incident response and supporting under-resourced critical infrastructure—would improve cybersecurity even if AI capability develops more slowly than expected.

That makes the core defensive argument relatively robust.

From the Cyber Arms Race to the Cyber Coalition

This development also marks an important continuation of an issue previously examined by CyberCenter.

On August 14, CyberCenter published “The Cyber Arms Race Has Entered the AI Model,” examining how frontier AI systems themselves were becoming strategic cybersecurity assets.

The central question then was control:

Who controls increasingly powerful cyber-capable models?

The August 27 coalition introduces the next question:

Who should have defensive access to them?

These are related but distinct stages of the same transition.

The first concerns capability.

The second concerns governance.

The third—now beginning—is institutional organization.

Governments, AI laboratories, cybersecurity companies and infrastructure operators are starting to define how responsibility should be divided when cyber capability becomes

increasingly automated.

That institutional architecture may ultimately matter as much as the models themselves.

What Success Would Actually Look Like

The open letter is voluntary.

It creates no binding obligations.

It establishes no international institution.

It provides no enforcement mechanism.

And signing a declaration is considerably easier than replacing vulnerable infrastructure.

Its importance should therefore not be exaggerated.

The real test comes next.

Will frontier AI companies provide meaningful defensive access?

Will governments fund cybersecurity for organizations that cannot afford it?

Will cybersecurity companies share useful threat intelligence rather than simply marketing new products?

Will critical-infrastructure operators actually eliminate legacy vulnerabilities?

Will countries cooperate during AI-enabled cyber incidents?

And will progress be measurable?

A successful coalition should eventually be able to demonstrate outcomes:

shorter remediation times,

fewer exploitable legacy systems,

better protection for hospitals and utilities,

faster cross-sector threat sharing,

and greater defensive capacity among organizations that previously lacked it.

Without those results, the letter risks becoming another declaration of cybersecurity urgency.

Bottom Line Assessment

The August 27 appeal is important not because more than 100 companies agree that AI creates cybersecurity risks.

That conclusion is increasingly conventional.

The important development is that organizations normally separated by commercial competition and institutional boundaries are beginning to describe cyber defense as a shared responsibility.

That is a meaningful change.

Artificial intelligence is making cyber capability easier to reproduce, faster to deploy and potentially available to a wider range of actors.

The response being proposed is therefore equally distributed:

government coordination,

private-sector capability,

shared intelligence,

critical-infrastructure investment,

and controlled access to advanced defensive AI.

This is the beginning of something larger than another cybersecurity initiative.

It is the early architecture of collective defense for the AI era.

Whether that architecture becomes operational—or remains an industry declaration—will depend on what happens after the signatures.

Vladimir Tsakanyan, Ph.D. Center for Cyber Diplomacy and International Security (CCD-IS)

Sources and links

[1] Vladimir Tsakanyan, Ph.D.

https://cybercenter.space/vladimir-tsakanyan/

[2] Center for Cyber Diplomacy and International Security (CCD-IS)

https://cybercenter.space/about/

Read the original article online

https://cybercenter.space/2026/08/31/the-cyber-coalition-why-rival-ai-companies-are-suddenly-calling-for-collective-defense/