FRONTIER.

FRONTIER / NEW ANALYSIS



Palo Alto Networks' Consortium Wins U.S. Backing. But Secure by Design Must Earn Its Keep.

Nine suppliers are joining a government-backed push to export complete AI stacks with security built in. For CIOs and CISOs, the prize is less risk and less integration work. The danger is paying for more assurance without getting a safer business.

FRONTIER / Frontier Research / November 19, 2026

NEW ANALYSIS

Built to change a decision.

Nine suppliers are joining a government-backed push to export complete AI stacks with security built in. For CIOs and CISOs, the prize is less risk and less integration work. The danger is paying for more assurance without getting a safer business.

EDITION	**New Analysis**
AUTHOR	**Frontier Research**
PUBLISHED	**November 19, 2026**

This PDF is a publication rendering of the canonical Frontier News article. The web article remains the live source of truth; analysis, figures, and source citations travel with this edition.

What changed

Secure by design should make the enterprise harder to compromise, not simply more expensive to approve.

That is the test facing the Secure SI by Design Consortium, whose U.S. government approval Palo Alto Networks announced on November 17. The coalition brings together nine companies: Palo Alto Networks, NVIDIA, IBM, OpenAI, Snowflake, Cognizant, Reflection, Factory and ServiceNow. Its ambition is to supply international buyers with complete AI stacks in which cybersecurity is engineered across infrastructure, data pipelines, models and applications.[1]

This is significant because the enterprise AI buying problem is becoming a systems problem. A capable model is only one ingredient. The business also needs trustworthy data access, controlled execution, usable evidence and a way to recover when something goes wrong. Buying those ingredients separately often leaves the customer responsible for making their security assumptions agree.

The consortium has received designation under the Department of Commerce's American SI Exports Program. The official directory lists Palo Alto Networks as its anchor member, the company organizing and leading the coalition. The program uses "SI" for Super Intelligence; that is its terminology, not evidence that a particular deployment has achieved a scientifically established level of intelligence.[2]

What approval actually brings

The commercial significance is access to coordinated federal support. ExecutiveBiz's coverage highlights government advocacy, faster export-license reviews and potential financing through the Export-Import Bank and U.S. International Development Finance Corporation.[3]

The government's own FAQ puts boundaries around those benefits: designated packages may receive one or more forms of support when officials determine that doing so serves U.S. interests. Licensing review remains subject to applicable law and agency procedure; financing referrals remain subject to the relevant agency's requirements and underwriting. Priority consideration is not guaranteed approval or guaranteed money.[4]

For enterprise buyers, designation should therefore be treated as a meaningful route to market, not as a blanket technical certification or assurance that every proposed configuration is secure. The FAQ identifies cybersecurity architecture and standards alignment among considerations in selection, but does not establish a customer-specific security acceptance test.[4]

Palo Alto Networks describes automated agent registries, standardized risk assessments and real-time monitoring as practical elements of the consortium's approach. Its announcement does not provide a versioned interoperability matrix, independently reproduced security results or a contractual allocation of responsibility for failures between suppliers. Those are the next deliverables buyers should request.[1]

Why it matters

The biggest opportunity is at the boundaries

The nine-company roster matters because the problem extends across layers. Security needs to survive the journey from a user's instruction through retrieved data, model reasoning, tool selection and an actual business-system change. A strong control at one step cannot automatically make the entire journey safe.

Consider an illustrative procurement agent that reads supplier documents and can update payment details. A malicious instruction inside a document should not acquire the authority of the employee who asked the agent to review it. The model may identify suspicious content, but the system also needs authorization checks at the tool boundary, appropriately limited credentials and an independent approval for sensitive changes. The meaningful outcome is that an unauthorized payment change does not happen—even when one component makes a mistake.

This is where a consortium could earn its place. Suppliers could agree on identity propagation, policy enforcement, event formats and supported upgrade paths before the customer assembles a production service. Joint testing could expose failures that individual product evaluations miss. Shared evidence could reduce the repeated work of proving the same control to multiple internal reviewers.

For a CIO, that could shorten the route from pilot to useful production. For a CISO, it could provide clearer enforcement and investigation across the stack. Both benefits depend on an actual integration commitment. A collection of individually capable products does not become a secure system merely because their vendors share a consortium name.

Security must become an engineering asset

The best economic case for secure by design is reuse. A well-tested identity boundary, hardened deployment pattern or reliable audit mechanism can protect many applications. The cost is incurred deliberately, then spread across deployments instead of being rediscovered during every project or incident.

That principle has an established foundation. CISA's Secure by Demand guidance urges buyers to examine product security before purchase, put appropriate requirements into contracts and continue assessing outcomes afterward. It builds on supplier ownership of customer security outcomes, transparency and leadership accountability.[5]

NIST's SP 800-218A similarly extends the Secure Software Development Framework with practices for AI model development and is intended to help model producers, system builders and acquirers. The consortium has existing foundations to build on; it need not invent another parallel checklist for customers to reconcile.[6]

Our assessment is that a successful package would turn those principles into reusable, operating controls. Default permissions would be narrow. Security-relevant changes would be traceable. Evidence would identify exactly which component versions and configurations were tested. Updating one component would trigger the relevant tests without requiring the buyer to repeat every review from scratch.

That last point is essential to cost control. Reusing evidence is sensible only while its assumptions still hold. A changed model, data source, tool permission or execution environment can alter the risk. Vendors should make those dependencies visible and automate targeted retesting, while retaining periodic end-to-end exercises to catch interactions that narrower tests miss.

How the bill can grow without the protection improving

The failure mode is easy to imagine: each supplier adds its own scanner, policy engine, dashboard and evidence format. Every application then passes through several queues that ask similar questions but produce incompatible answers. Engineering slows, licensing grows and security teams inherit the work of resolving conflicts.

Runtime protection can also impose latency, consumption charges and false blocks. Excessive friction creates an incentive to bypass approved systems, moving work into places the enterprise can see less clearly. Conversely, optimizing only for throughput can encourage controls that wave risky actions through. Buyers need to measure both effects.

A low incident count during a short pilot proves little on its own. The deployment may have had little exposure or weak detection. Compare candidate architectures against the same representative workflows and controlled attack scenarios. Measure successful unauthorized actions, sensitive-data leakage and containment time alongside legitimate task completion, false interventions, latency and total operating cost.

The relevant business measure is the cost of delivering an acceptable, protected outcome. More tests executed or more attacks claimed to be blocked do not establish that result.

A package can reduce integration work and increase dependence

A tightly integrated stack could simplify deployment while making substitution expensive. A buyer may discover that changing a model or security component invalidates the package's support commitments or forces substantial requalification.

That trade-off is manageable if exposed before purchase. Ask which components are optional, what happens when a supplier outside the consortium is used, and whether policies, logs and test evidence can be exported in usable forms. A commercially open membership list is not the same thing as a technically interchangeable architecture.

International buyers have an additional responsibility: distinguish export facilitation from operational sovereignty. They still need precise answers about deployment location, administrative access, encryption-key control, telemetry destinations and continuity if a supplier or service becomes unavailable. Government support can help a transaction happen; it does not answer those deployment questions for the customer.

Nor can security controls alone guarantee that an agent makes a sound business decision. An authorized action can still be inaccurate, inappropriate or costly. Business evaluation, financial controls and operational recovery remain necessary alongside cybersecurity.

Frontier take

This is worth watching closely because it could move the AI market toward suppliers taking responsibility for the assembled system. Enterprises have spent years buying individually defensible components and then funding the work required to make them safe together. Agentic execution makes the gaps in that arrangement harder to tolerate.

The coalition's breadth and government-backed distribution give it an opportunity to make secure deployment patterns repeatable across customers and countries. If it delivers, the benefit could extend beyond individual security features: fewer integration exceptions, faster acceptance of well-understood configurations and more confidence to delegate useful work to agents.

But the commercial incentive cuts both ways. A consortium can standardize protection, or standardize a larger bundle of products that customers must buy and repeatedly validate. Buyers should make the distinction with operational evidence and contractual responsibility.

Our view is that secure by design should earn its premium in three ways: remove meaningful attack opportunities, reduce the customer's repeated engineering and assurance work, and preserve the ability to operate and recover when something fails. Some upfront cost is justified—even substantial cost for high-impact systems. The burden is to explain what risk that investment changes and what recurring work it eliminates.

The harder governance question is who owns a failure between members. The consortium's anchor role does not, by itself, establish responsibility for every integration defect. A customer needs an accountable delivery party, a common incident process and clear remediation obligations. Nine suppliers cannot be allowed to turn one enterprise incident into nine separate arguments about scope.

The right response is neither automatic endorsement nor dismissal as another industry club. Put the coalition on the shortlist for a bounded, demanding evaluation. Require it to demonstrate that security travels with the business workflow, survives supported changes and costs less to operate than an equivalent collection of disconnected controls.

Government backing opens a door. A safer, more workable enterprise is what must come through it.

Three moves for CIOs and CISOs

1. — Make security outcomes an acceptance condition Select one valuable workflow with real integration complexity and define its allowed actions, prohibited outcomes and recovery needs. Test the proposed package against the same legitimate tasks and controlled attacks as your current or alternative architecture. Require independent review of the results, including failures and configuration limits, rather than relying only on the suppliers' demonstration.
   - Decision trigger: Use these acceptance gates before granting production privileges or committing to a broad deployment contract.
   - Why now: The designation creates a purchasing opportunity. Defining success before the sales and procurement process hardens prevents a government-backed label from becoming a substitute for evidence.
2. — Put a budget around assurance and eliminate duplicated work Have engineering, security and finance baseline build effort, review delays, runtime costs and remediation effort. Require reusable control implementations and versioned test evidence mapped to existing frameworks. Agree which changes trigger targeted retesting and which require a full reassessment. Track false blocks and business-task completion alongside security performance.
   - Decision trigger: Expand only when measured protection improves within the workload's agreed cost, latency and reliability limits; redesign the control set when repeated reviews add effort without finding meaningful new risk.
   - Why now: Early architecture choices determine whether verification becomes a shared capability or a permanent tax on every team shipping AI.
3. — Contract for the seams and rehearse an exit Name the delivery party accountable for cross-vendor failures. Specify supported combinations, incident coordination, vulnerability notification, patch ownership and evidence access. Exercise a supplier outage and a component replacement, including policy and log export, independent emergency access and restoration of critical work.
   - Decision trigger: Resolve material gaps before accepting exclusive dependencies, retiring existing safeguards or expanding to additional countries.
   - Why now: The benefits of an integrated stack arrive with new dependencies. Clear responsibility and tested recovery prevent those dependencies from becoming the enterprise's next unmanageable risk.

Sources

- Palo Alto Networks: Secure SI by Design Consortium officially approved, November 17, 2026 [1]
- American SI Exports Program: Designated consortia [2]
- American SI Exports Program: Scope, designation and support FAQ [4]
- ExecutiveBiz: Palo Alto Networks-led consortium receives government approval, November 18, 2026 [3]
- CISA: Secure by Demand guide for software buyers [5]
- NIST SP 800-218A: Secure development practices for generative AI and foundation models [6]

SOURCES / CITATIONS

Trace the argument.

Every source below is linked. Source numbering follows first appearance in the report.

01 **Palo Alto Networks: Secure SI by Design Consortium officially approved, November 17, 2026**
https://www.paloaltonetworks.com/blog/2026/11/secure-si-by-design-consortium-officially-approved-by-u-s-government/

02 **American SI Exports Program: Designated consortia**
https://siexports.gov/consortia-menu

03 **ExecutiveBiz: Palo Alto Networks-led consortium receives government approval, November 18, 2026**
https://www.executivebiz.com/articles/palo-alto-networks-secure-si-by-design-consortium

04 **American SI Exports Program: Scope, designation and support FAQ**
https://siexports.gov/faq

05 **CISA: Secure by Demand guide for software buyers**
https://www.cisa.gov/sites/default/files/2024-08/SecureByDemandGuide_080624_508c.pdf

06 **NIST SP 800-218A: Secure development practices for generative AI and foundation models**
https://csrc.nist.gov/pubs/sp/800/218/a/final

FRONTIER.

TAKE THE NEXT STEP

Push this further.

Research is useful when it changes a decision. Bring us the question, the constraint, or the argument you want tested.

RESEARCH / BRIEFINGS

inquiries@thisisfrontier.ai

Briefings, advisory work, executive discussions, and research inquiries.

REPRINTS / LICENSING

reprints@thisisfrontier.ai

Reprints, enterprise distribution, excerpts, and licensing.

FRONTIER

www.thisisfrontier.ai

Research, Data Insights, News, and the live Frontier corpus.

All rights reserved. Contact Frontier for redistribution or reprint rights.

SEE FURTHER. MOVE SOONER.