The next AI arms race isn't about chips. It's about trust.

Semiconductors and compute dominate the strategic debate. But the decisive contest will be over who can make autonomous systems accountable - and the window to act is closing.

The world is preparing for the wrong AI war. Governments obsess over semiconductor supply chains. Technology companies race to build ever-larger models. Intelligence agencies monitor computing power, data-centre capacity and access to advanced chips as though these were the decisive determinants of future advantage.

They matter. But they are no longer the centre of gravity. The next strategic contest will be decided not by who builds the most powerful AI model, but by who solves the problem of trust before autonomous systems become embedded throughout the global economy.

Consider an autonomous AI agent working for a defence contractor. It negotiates contracts in Singapore, accesses cloud infrastructure in the US, verifies inventory through European logistics systems and coordinates shipping across the Indo-Pacific. No human approves each step. The agent acts independently - exactly as it was designed to do.

Now imagine something goes wrong. A critical system fails. Sensitive information is exposed. Financial losses cascade across multiple

jurisdictions. The first question will not be how many parameters powered the model. It will be: who is responsible?

If that question isn't answered instantly and with certainty, the foundations of trust that underpin modern commerce, government and critical infrastructure begin to erode. That moment is far closer than most policymakers appreciate.

Today, organisations take comfort from the fact that most AI systems remain confined to controlled environments. Enterprise identity systems, internal permissions and access controls create a reassuring sense of oversight. That security is largely an illusion. It exists because autonomous agents are still in their infancy, not because we have solved the accountability problem.

Indeed, the entire economic promise of AI depends on dismantling those barriers. The value of these systems does not come from keeping agents locked inside corporate networks; it comes from letting them deal with customers, suppliers, banks and governments - invoking external tools, exchanging information and making decisions with limited human supervision. The very behaviour that makes AI economically transformative is the behaviour that makes accountability exponentially harder.

Yet governments continue to treat this challenge as one that can be deferred. History suggests precisely the opposite. The internet was built before security was understood. Social media scaled before democracies grasped its political consequences. Global software supply chains expanded before organisations appreciated the vulnerabilities they introduced. Each time, governance arrived after deployment. Each time, retrofitting trust cost vastly more than building it into the architecture from the outset. We now risk making the same mistake with autonomous AI.

As agents move across corporations, sectors and national borders, traditional identity and access management loses its authority. Systems designed to establish trust within a single network were never designed to establish it across thousands of interconnected environments operating at machine speed. This is the strategic blind spot at the heart of the AI revolution.

> Leading security frameworks are converging on a simple principle: every autonomous agent needs a persistent, cryptographically verifiable identity that stays attached throughout its operational life.

Every action attributable; every delegation auditable; every permission revocable; every claim of authority independently verifiable. Above all, accountability must travel with the agent. Without that, organisations will

find themselves deploying autonomous systems whose actions span multiple networks while ownership, responsibility and liability become progressively harder to determine.

The implications reach far beyond corporate risk. Autonomous agents will soon be embedded in defence supply chains, energy grids, financial markets and telecommunications networks - negotiating transactions, allocating resources, coordinating logistics and shaping operational decisions at a speed and scale beyond meaningful human supervision.

That creates extraordinary opportunity. It also creates an extraordinary new attack surface. Hostile actors do not need to build better AI than the Western democracies; they need only exploit weaknesses in the trust architecture surrounding it. A compromised agent. A forged identity. A manipulated delegation chain. A misattributed action inside critical infrastructure. Any one of these could produce consequences far beyond the initial technical failure.

The danger is not simply cybersecurity. It is strategic ambiguity. In a crisis, governments must be able to determine immediately what acted, who authorised it, who owns it and who bears responsibility. If those answers require weeks of forensic analysis, the damage is already done.

Worse, a single catastrophic incident involving an unverified agent could trigger a collapse in public and institutional confidence. The sequence is predictable: governments impose emergency restrictions, boards pause deployments, investors retreat, critical programs stall. Adoption slows precisely when democratic nations need every available productivity advantage to compete with increasingly capable authoritarian rivals. That would be a strategic self-inflicted wound.

The answer is not to halt autonomous AI, nor to force every decision back through human approvals that destroy the speed and scale these systems offer. That would simply guarantee others move faster. Instead, democracies must treat portable accountability as strategic infrastructure and move rapidly towards a globally recognised open standard for agent identity: a common framework that applies regardless of platform, cloud provider or jurisdiction.

Crucially, that foundation must remain open and non-proprietary. If a handful of technology companies entrench competing identity architectures, the democratic world will inherit a fragmented landscape of incompatible trust systems precisely when collective coordination matters most. Railways needed standards. Telecommunications needed standards. The internet succeeded because common protocols created the trust on which extraordinary innovation was built. Autonomous intelligence needs

something more fundamental still: accountability that travels with the agent.

The window is open now because the agent ecosystem is still taking shape. It will not stay open. Once agents are deeply embedded across government, defence, finance and critical infrastructure, fragmented trust models will harden into software architectures, contracts and operational dependencies that are extraordinarily difficult to unwind.

The question is no longer whether autonomous agents will become integral to economic and national power. That future has already arrived. The question is whether democracies will build the trust architecture to govern these systems before the first major crisis exposes its absence. Because that crisis is coming - not because AI will suddenly turn malevolent, and not because the technology will fail, but because societies are racing towards a world in which autonomous systems negotiate, transact and act across borders without any universally accepted way of proving who authorised them and who answers for them.

History teaches a simple lesson: trust is easiest to build before a crisis and hardest to build after one. The first great crisis of the autonomous age will not be caused by artificial intelligence itself. It will be caused by our failure to know who controls it, who authorised it and who must answer for its actions. By then, the time to build trust will have passed.