AI Revolutionizing Recruitment and Hiring:
A New Era of Talent Acquisition

Enterprise AI is opening up entry points that legacy security stacks were never designed to cover, and most organizations are discovering that fact the hard way: deep into deployment.

01

The shift to enterprise AI is doing more than changing how applications get built. It is quietly redrawing the cybersecurity attack surface, dragging data pipelines, model training environments, identity systems, and supply chains into the line of fire. The defenses that protected yesterday's three-tier applications were never designed for this new shape, and the gap is becoming impossible to ignore.

02

The conversation came into sharp focus during a recent broadcast on theCUBE, SiliconANGLE Media's livestreaming studio, where host Dave Vellante spoke with Steve Kenniston, senior cybersecurity evangelist for portfolio marketing at Dell Technologies, at the Securing the AI Factory with Dell Technologies and Intel event. Kenniston laid out a picture of how the so-called AI factory is reshaping risk, and what organizations need to do before their AI ambitions outrun their security controls.

One application, dozens of new doorways

Traditional enterprise applications tended to have a small, well-understood set of entry points. Lock the perimeter, manage user access, patch the database, and most threats stayed outside the wall. The AI factory does not behave that way. It is a sprawling, interconnected system, and each new layer brings its own category of exposure.

Kenniston framed the change bluntly: AI rewrites the rules entirely. Model inferencing introduces a new attack target. Model training data becomes a high-value asset that adversaries want to poison or exfiltrate. Prompt injection turns user input into a potential exploit vector. Identity management has to account for autonomous agents, not just humans. And agentic workflows, where AI systems make decisions and take actions on their own, expand the blast radius of any single compromise. Each of these surfaces is moving fast, and every new application brings a new attack surface with it.

The 85% problem nobody is talking about loudly enough

Perhaps the most striking data point Kenniston shared comes straight from Dell's services organization. Roughly 85 to 90 percent of AI projects are now being halted mid-implementation because the security team was not pulled in from the start. Teams reach the proof-of-concept finish line, get ready to push to production, and only then discover that no one has vetted the data flows, the model dependencies, or the access patterns.

The pattern Kenniston described is uncomfortably familiar. Security shows up at the five-yard line, raises objections that should have been raised six months earlier, and the project stalls. The lesson is not that security teams are obstacles. It is that organizations are still treating AI security as something to bolt on after the architecture is already locked in. By that point, the cost of doing it right has multiplied.

Treating the AI factory as one security surface, not many

Dell's response is to stop thinking about AI security as a checklist of individual products and start thinking about it as a single, validated surface. Compute, networking, storage, and data layers are designed to interlock, with the gaps between them treated as the most likely place for a breach to begin. This is a meaningful shift away from the typical pattern of buying point tools for each layer and hoping the seams hold.

Kenniston put it plainly: at Dell, security is integrated into everything from the supply chain through the silicon and out to the device that arrives on a customer's loading dock. The places where security breaks down in AI systems, he said, tend to be exactly the cracks between components. Eliminating those cracks is more about architecture than about adding another scanner.

Supply chain integrity moves from rare question to default

One of the more revealing trend lines Kenniston highlighted is around supply chain trust. Two years ago, perhaps one customer in twenty asked questions about how the hardware they were buying had been built, shipped, and verified. Today, those questions come up in almost every briefing. The conversation has shifted from theoretical concern to baseline expectation.

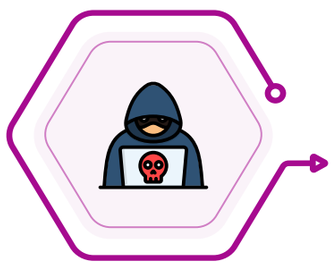

The reason is straightforward. AI workloads concentrate enormous value into a small number of high-performance systems. If an attacker can compromise the supply chain, they do not need to break in through the front door. They can arrive pre-installed. Hardening the path from chip fabrication through final delivery is no longer an exotic ask, it is an operational requirement.

Zero-trust grows up: from network to identity

Rounding out the strategy is a continued shift away from firewall thinking and toward zero-trust principles. The defining questions are no longer just whether traffic is inside or outside the perimeter. They are who is requesting access, what they are trying to do, when they are trying to do it, and from where the request is originating. In an AI environment full of services calling other services and agents acting on behalf of users, those granular controls become essential.

Identity, in this view, is the new perimeter. And it has to extend beyond people. Service accounts, model endpoints, training pipelines, and autonomous agents all need their own verifiable identities and tightly scoped permissions. Without that, every compromised credential potentially unlocks the entire factory.

The human element of recovery

One theme that often gets buried in AI security conversations is recovery. Detection and prevention dominate the headlines, but Kenniston flagged the human side of incident response as routinely underestimated. When something goes wrong inside an AI environment, the people responsible for restoring service have to understand systems that may be more opaque than anything they have managed before. That capability does not appear on demand. It has to be built ahead of time, through training, runbooks, and tested response procedures.

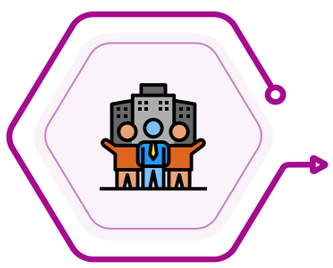

This is where many organizations are still exposed. They have invested in tooling but not in the muscle memory required to use it under pressure. The AI factory raises the stakes of that gap considerably.

The bottom line for enterprise AI leaders

The takeaway from Kenniston's conversation with Vellante is clear, even if it is not comfortable. The AI factory is not a faster version of the old application stack. It is a different kind of system, with different vulnerabilities, and it punishes organizations that try to retrofit security after the fact. Designing it in from day one, treating the full stack as a single defended surface, hardening the supply chain, and adopting identity-first zero-trust principles are no longer advanced practices. They are the floor.

For enterprise leaders staring at ambitious AI roadmaps, the practical advice is simple. Pull the security team into the first design meeting, not the last review. The five-yard line is the worst possible place to discover that your AI factory was never built to defend itself.