How AI Expands Risk Across Enterprise

Veeam's Brad Linch & Emilee Tellez on AI-Driven Cyber Risk and Data Exposure

Brad Linch

Linch drives strategic product and go-to-market programs while championing Veeam's Safe AI vision by keeping data resilient and trusted. He has more than 10 years of experience in helping organizations recover their data from both disaster and cyber incidents.

Emilee Tellez

Tellez has 13 years of experience working with organizations to design, implement and maintain BC/DR and cyber recovery strategies. This includes guiding teams in developing comprehensive data resiliency programs and facilitating incident response and cyber recovery workshops focused on rapid, reliable restoration after any disaster. As field CTO, Tellez works with global technology executives and collaborates closely with Veeam's product management teams to advance the company's data protection portfolio and explore emerging technologies across cloud, security and AI.

Generative artificial intelligence has introduced new layers of enterprise risk by enabling automated actions, increasing data exposure and expanding opportunities for tool exploitation. "The tactics are very similar, but the rate at which these threat actors are executing things is much faster. AI has amplified things fivefold, tenfold, a hundredfold," said Brad Linch, director of enterprise strategy at Veeam.

As organizations deploy AI systems, they face uncertainty around whether automated decisions align with intended outcomes or expose sensitive data. These risks make it critical for organizations to clearly understand, evaluate and quantify their exposure before deploying AI at scale.

Without defined guardrails, AI systems with broad permissions and access to sensitive data can create cascading failures. "We have non-human identities that can now take action based off of broad sets or permissions," said Emilee Tellez, field CTO at Veeam. "Now that becomes another area in which we have to make sure we're identifying those risks and making sure that we're building disaster recovery planning."

In this video interview with Information Security Media Group at RSAC Conference 2026, Linch and Tellez also discussed:

- The need for structured AI governance and cross-functional oversight;
- The importance of data quality and hygiene for reliable AI outcomes;
- How weak AI committee oversight can lead to enterprise data exposure.

"Data is that foundational layer. If you feed it incorrect or poor-quality data, the output will be a poor outcome."

— Brad Linch, Director, Enterprise Strategy, Veeam

AI Amplifies the Threat Landscape

ANNA DELANEY: What risks does AI introduce when it comes to automated actions, data exposure and tool exploitation?

BRAD LINCH: The attack tactics themselves haven't changed so much, but AI has amplified them 50-fold to 100-fold. The rate at which threat actors are executing is much faster.

Take software vulnerabilities as an example. In 2023 or 2024, the time between a vulnerability becoming public and being exploited was five months. Now, AI has accelerated that through automated toolkits. A vulnerability goes public and the average time to exploitation is a day and a half. It's an example of how AI is accelerating things we've always known about.

Data sprawl is another example. Data sprawl has existed since I started in this industry, but it wasn't taken very seriously because the associated risk was low. Now, when you pair it with AI, data is that foundational layer, like the first layer of a cake. If the data feeding these AI models and pipelines is inaccurate or stale, the output will be incorrect. The same problems we've always dealt with are now significantly amplified by AI.

EMILEE TELLEZ: Now that we've introduced all those risks around data and AI, we have to ask whether we're going to allow AI to take action, and whether the actions it takes will produce the outcomes we want. Could it expose sensitive company data to the outside world? Could it accidentally delete or corrupt the data you have?

There's a lot that goes into knowing your data state and data foundation. Knowing what actions an AI system will take, and whether those actions will be the right ones, is something we discuss constantly with organizations. Many are excited to implement AI in some way within their industry, but the real focus needs to be on the outcome you're trying to drive, what data you're feeding in and how you ensure the actions it takes are the correct ones.

Quantifying AI Risk

DELANEY: Why is understanding, evaluating and quantifying risks so important to building effective AI strategies?

TELLEZ: Everything needs a guardrail. We're at RSAC to talk about security and the threat landscape, and AI is just another dimension of that landscape. But we also have to account for operational resilience. We used to plan for tornadoes, natural disasters, fires and floods. Those were relatively straightforward to plan for. Then cyberthreats arrived, and we had to account for external factors: threat actors gaining access to data, as well as encrypting, exfiltrating and corrupting that data.

Now AI introduces another dimension, and the risk isn't only internal. We have non-human identities that can take action based on broad sets of permissions, and these identities have access to different APIs executing across different workflows. That becomes another area where we must identify risks and build disaster recovery planning based on what we're seeing from the AI landscape today.

LINCH: The reason clients implemented protection, restoration and disaster recovery plans was because humans make mistakes. Now multiply that by 82, because research shows there are 82 non-human identities per human identity in organizations, each capable of making mistakes.

There is a lot of focus on cyberattackers using AI for malicious activity, and that is certainly happening. But the accidental mistakes are what you will read about in the news every day: agents causing unintentional damage that costs organizations time, resources and money. Understanding those risks, understanding the privileges these agents hold and understanding their access to sensitive data is what will better prepare you to respond.

Data Quality Drives AI Outcomes

DELANEY: Why does the success of an AI initiative depend on the context, quality and accessibility of data?

LINCH: To give you a simple analogy, it's no different from the ingredients you give a chef. If you give a chef spoiled ingredients, the dish won't taste good because they can only work with what they're given. AI works the same way. Data is that foundational layer. If you feed it incorrect or poor-quality data, the output will be a poor outcome.

TELLEZ: There is a real opportunity to leverage data. The amount of data humans consume and create is enormous, and there is a competitive advantage to be gained from it. The problem is that if you never invested in good data hygiene, meaning tagging, qualifying and adding contextual metadata, then building on that foundation will not produce the best outcomes. Garbage in, garbage out. You cannot expect excellent results if the data you're providing is old, stale or irrelevant to the objective you're trying to achieve.

The Role of AI Committees

DELANEY: How can AI committees help evaluate projects before they go into production?

TELLEZ: An organization I met with two weeks ago has an AI committee that operates every quarter. They wanted to give employees a voice, ensuring that any AI agent project has multiple stakeholders involved from the start. If they're building something for the accounting department, they ensured the lead person for accounting is involved throughout the entire project, aware of what is being built, what outcome it will achieve and how it will help that specific team.

They then layer in additional guardrails, such as third-party testing. If the system handles financial data, including credit cards, account numbers and monetary values, they bring in a third party to conduct penetration testing and verify there are no exposures. Every time a risk or gap is identified, they move immediately to fix it. Every quarter they review new risks and new projects and bring in external help.

"Many are excited to implement AI in some way within their industry, but the real focus needs to be on the outcome you're trying to drive, what data you're feeding in and how you ensure the actions it takes are the correct ones."

— Emilee Tellez, Field CTO, Veeam

They are still moving fast on innovation, but they've put best practices in place to provide digital checks along the way.

LINCH: That's a great example. And I'm genuinely impressed by how many organizations have AI committees and a process in place. But I do see a couple of common blind spots. I don't often hear the question: What's the blast radius? If this AI system or pipeline goes wrong or causes an accident, what's the potential impact on the organization, and how do we recover from that? There are additional questions these committees could be feeding into their monthly or quarterly meetings that would let them prepare more effectively.

Where AI Committees Fall Short

DELANEY: What other ways do you see these committees fall short?

TELLEZ: Some are not including the right stakeholders. They're not bringing in the heads of different departments. I was sitting at a table with two organizations recently. One described their quarterly AI committee. The other said they also have a committee, but they also gave access to Claude Code to all 10,000 employees and told them to start building. Some of those employees had never written a line of code in their lives, and now they were rapidly creating agents with access to unknown data.

There is a need to balance a strike between innovation speed and safe innovation. The takeaway from that conversation was that with the right quarterly meetings and the right stakeholders involved, you can still drive toward ambitious AI projects but with more control, rather than giving everyone unrestricted access and letting them do what they want.