NYC - Thursday 25 June 2026, Brown Rudnick

Executive Summary

The June 2026 New York City CxO Security Forum brought together senior cybersecurity, technology, risk, and business leaders for candid, practitioner-led discussions on how AI is reshaping enterprise security and governance.

A recurring theme was that the greatest challenges surrounding AI are no longer technical—they are organizational. Speakers explored how AI agents should be governed as digital workers, why identity and access management must evolve for AI-native environments, and how automation should eliminate repetitive work while improving explainability and trust.

Participants also examined new approaches to measuring cyber risk. From third-party risk scoring and insider threat detection to OT security and identity governance, discussions highlighted the limitations of manual processes and static controls, pointing instead toward AI-driven models that enable faster, more informed risk decisions.

The forum concluded with lessons from the SolarWinds incident, reinforcing that major cyber events quickly become executive leadership challenges. Today's CISO must balance technology expertise with governance, communication, and business judgment.

Across every session, the message was consistent: organizations that modernize governance, embrace AI responsibly, and make risk measurable will be best positioned to navigate the next generation of cybersecurity challenges.

Find more details and all of the presentation materials in the CxO Security Forum Online Community Portal at engagez.net/cxo.

The web page for this Forum is CxOSecurityForum.com/NYC

Discussion Leaders & Featured Sessions

Agentic AI Meets Zero Trust: What Actually Breaks—and How to Fix It

Discussion Pod #1

Josh Woodruff
Author, Agentic AI + Zero Trust

"AI agents are the employees you forgot to interview."

The forum opened with a thought-provoking discussion that challenged one of the most common assumptions about enterprise AI: that the biggest obstacles are technical. Instead, Josh argued that successful AI adoption is primarily a leadership and governance challenge. Organizations that struggle with AI often attempt to automate broken processes, while those seeing meaningful business value first rethink how work should be performed and then apply AI to accelerate it.

A central theme was that AI agents should no longer be viewed as software tools but as digital workers operating inside the enterprise. Unlike employees, however, these "workers" are often deployed without interviews, onboarding, training, or supervision—despite receiving broad access to enterprise systems. This changes the security problem from preventing unauthorized access to managing authorized misuse.

The discussion explored how Zero Trust principles must evolve to address non-human identities and autonomous decision-making. Rather than slowing innovation, participants agreed that governance enables organizations to move AI projects from perpetual pilots into production with greater confidence and lower risk. The conversation also emphasized that security leaders have an opportunity to reposition themselves from the department that says "no" to the function that enables responsible AI adoption.

Key Takeaways

- AI initiatives fail more often because of leadership and governance failures than technology shortcomings.
- AI agents should be governed like employees—with identities, defined responsibilities, supervision, and accountability—not simply treated as software.
- The emerging risk is no longer unauthorized access, but authorized misuse by autonomous agents acting within their permissions.
- Organizations need visibility into which AI agents can perform irreversible actions, just as they track privileged human administrators.
- Clear guardrails and governance improve AI performance rather than limiting it, allowing organizations to deploy AI faster and with greater confidence.

- Zero Trust provides a practical operating model for securing AI through identity, behavioral monitoring, segmentation, data governance, and rapid response when agents behave unexpectedly.

Notable Quotes

"AI doesn't fail because of technology. It fails because of leadership."

"The threat model has changed. It's not unauthorized access. It's authorized misuse."

"Governance doesn't slow AI. Uncertainty slows AI."

"Security stops saying 'no.' It starts saying, 'Here's how we make this happen.'"

"Agentic AI is where Zero Trust stops being optional."

Start-Up Showcase #1

OT Security's Missing Control Layer?

Josh Ross – Founder & CEO, IronLoop

"There's an opportunity to take DevOps principles and apply them to the OT space."

Josh Ross challenged attendees to rethink a longstanding assumption in operational technology (OT) security: that visibility alone is enough. While most OT security investments focus on asset discovery, monitoring, and vulnerability detection, he proposed that organizations are missing a configuration control layer capable of validating intended system states and understanding the impact of changes before they occur.

The discussion highlighted the differing priorities of IT and OT teams. While IT emphasizes protecting data, OT prioritizes safety and operational continuity, making organizations reluctant to patch systems that could introduce downtime. Ross argued that applying DevOps-style configuration management principles to industrial environments could enable organizations to modernize OT more safely and confidently.

Key Takeaways

- OT security has focused on monitoring rather than validating intended configurations.

- Many OT vulnerabilities remain unpatched because operational disruption is seen as a greater risk than cyber threats.
- A configuration control layer could improve resilience by assessing the impact of changes before deployment.
- DevOps-inspired configuration management offers a promising new approach for securing critical infrastructure.
- As cyber attacks increasingly target physical systems, IT and OT security must become more closely aligned.

Notable Quotes

"We shouldn't be scared of change in OT."

"Most of the OT industry has been sold network monitoring solutions."

"Patching in OT can be dangerous."

The Missing Metric in Third-Party Risk Management?

Discussion Pod #2

Jeremy Huval
Chief Innovation Officer, HITRUST

"We don't need another assessment—we need better measurement."

Jeremy Huval shared an early look at a new approach for measuring third-party cyber risk that moves beyond today's fragmented assessment process. Rather than introducing another framework or questionnaire, the concept normalizes existing assurance artifacts—including questionnaires, SOC 2 reports, ISO certifications, HITRUST certifications, audit reports, and inherent risk factors—into a standardized Information Risk Score designed to quantify residual cyber risk and support executive decision-making.

The discussion centered on a common frustration shared by both vendors and enterprise TPRM teams: today's process measures activity rather than risk. Executives need to know how much third-party risk they are actually accepting, whether it exceeds their organization's risk appetite, and when it should be mitigated, accepted, or financially transferred. Participants also explored an emerging cyber insurance model that could transfer vendor-specific residual risk through contract-specific insurance policies.

Key Takeaways

- Current third-party risk programs generate significant operational overhead but often fail to quantify actual business risk.

- A normalized Information Risk Score could enable apples-to-apples comparisons across vendors regardless of the assurance framework used.
- Financial quantification of residual risk would allow boards and executives to make more informed governance decisions.
- Standardized scoring may also enable new forms of vendor-specific cyber risk transfer through insurance.
- HITRUST invited a small number of enterprises to participate in a confidential pilot program evaluating the model.

Notable Quotes

"How much third-party vendor risk are we actually accepting?"

"We want to move from activity-based reporting to governance-based decision making."

"Think of it as a cyber FICO score for third-party risk."

Beyond the Headlines: Executive Accountability, Regulatory Pressure, and the Reality of CISO Risk

Discussion Pod #3

Tim Brown

Former Chief Information Security Officer, SolarWinds
Current Partner at Team8 VC

"Security leaders today are responsible for communicating cyber risk to the business—not just running security technology."

When the SolarWinds SUNBURST attack became one of the most consequential cyber incidents in history, Tim Brown found himself leading the technical response while simultaneously navigating an unprecedented personal legal battle. As the first CISO ever charged individually by the SEC—a case ultimately dismissed—Brown offers a candid, deeply personal account of what happens when technical, business, legal, regulatory, and human crises collide.

Rather than focusing on the attack itself, Brown examines what follows: building cross-functional crisis teams, communicating under extraordinary uncertainty, managing boards, customers, regulators, media, and law enforcement simultaneously, and maintaining leadership while facing

the possibility of losing both career and reputation. He also reflects on how the incident reshaped Secure by Design initiatives, software supply chain security, executive accountability, and the evolving role of today's CISO.

This session is less about SolarWinds and more about what every security executive should understand before facing their own defining moment.

Three Key Takeaways

- **Major cyber incidents quickly become business, legal, and leadership crises—not simply security events.** Successful response requires coordinated leadership across executives, legal, communications, engineering, customer success, and government agencies.
- **Transparency, preparation, and culture determine long-term outcomes.** Organizations that communicate honestly, prioritize customers, empower cross-functional teams, and maintain trust recover faster than those focused primarily on reputation management.
- **The CISO role has fundamentally changed.** Today's security leaders must understand governance, disclosure obligations, board communications, personal liability, regulatory expectations, and executive decision-making—not just technology. Brown argues that CISOs are increasingly enterprise risk executives whose influence extends well beyond cybersecurity.

Notable Quotes

> "This happened on my watch."
>
> "Customers came before everything else."
>
> "Don't just be good—be exemplary."
>
> "Focus on the facts you know to be true."
>
> "The written word isn't enough. People want to talk to the CISO."
>
> "You will be out-numbered, out-marketed, and out-communicated."
>
> "The CISO shouldn't be running the crisis alone. This is an executive leadership event."

AI-Native Identity: Why IAM Is Being Rewritten

Discussion Pod #4

Barak Perelman
Founder & CEO, Opti

"Identity governance is no longer one program—it's a matrix of jobs."

Barak Perelman challenged attendees to rethink identity governance for the AI era, arguing that today's Identity and Access Management (IAM) platforms were designed for a workforce of human users—not AI agents, non-human identities, and increasingly dynamic enterprise environments. Drawing on interviews with nearly 100 CISOs and CIOs, he observed that most identity teams remain overwhelmed by manual processes, professional services engagements, and repetitive operational work, leaving little time to reduce actual risk.

Rather than layering generative AI onto legacy IAM platforms, Perelman advocated for purpose-built, AI-native identity models capable of understanding entitlements, context, business roles, and access behavior. He argued that the goal of AI should not be to replace security professionals, but to eliminate the repetitive work that prevents them from focusing on governance, risk, and business enablement.

The discussion also highlighted the growing urgency created by AI agents and non-human identities, which are rapidly expanding the identity attack surface while exposing the limitations of today's governance models.

Key Takeaways

- Identity governance has evolved beyond human users to include service accounts, machine identities, and AI agents.
- The primary challenge is no longer collecting identity data—it's acting on that data efficiently at enterprise scale.
- Generic LLMs are insufficient for identity governance; organizations need purpose-built AI models trained specifically on identity relationships and access patterns.
- AI should automate repetitive operational work, allowing identity teams to focus on reducing business risk rather than managing workflows.
- Explainability and auditable decision-making are essential if AI is to be trusted with identity governance.

Notable Quotes

> "The bottleneck isn't collecting data anymore. It's acting on it."

"Don't use AI to do the work you enjoy. Use it to eliminate the work you don't."

"The right AI isn't a generic LLM wrapped around IAM—it's purpose-built for identity."

"Architecture is the bet."

"The window is now—not the next audit cycle."

Start-Up Showcase #2

The Insider Threat Problem We Still Haven't Solved

Casper Neo - Founder & CEO, Palace Cybersecurity (nee Google)

"Access abuse hides in an ocean of legitimate activity."

Casper Neo presented a novel approach to one of cybersecurity's most persistent challenges: identifying when users with legitimate access begin behaving in ways they should not. Drawing on technology he developed and deployed at Google, Neo demonstrated how deep recommendation models—the same techniques that power platforms like YouTube and TikTok—can learn what "normal" access looks like across an enterprise and identify anomalous behavior without relying on static rules or manually maintained detection logic.

Rather than asking whether a user has permission to access a resource, the model predicts whether they *should* be accessing it based on historical patterns, organizational context, and continuously evolving behavior. The approach enables security teams to prioritize insider threat investigations, identify excessive entitlements, and surface suspicious activity from millions of legitimate access events with remarkably low alert volumes. Research presented at Black Hat and accepted for USENIX Security demonstrated strong results when evaluated against simulated insider attacks at Google.

Key Takeaways

- Traditional IAM and detection rules struggle to keep pace with constantly changing enterprise environments.
- Recommendation-system AI can learn normal access patterns and continuously adapt as organizations evolve.
- Behavioral scoring can prioritize insider threat investigations, identify excessive entitlements, and improve access governance.

- High-precision detection dramatically reduces analyst workload by focusing attention on the most suspicious activity.
- AI offers a promising new approach to insider threat detection by learning expected behavior rather than relying on manually written rules.

Notable Quotes

"Like YouTube finds videos you should watch, we find resources employees should not access."

"The system learns what normal looks like—and continuously relearns as the business evolves."

"The question isn't whether someone *can* access a resource. It's whether they *should*."

"We can catch insider access abuse with fewer than one alert per 10,000 employees per day."

References & Additional Discussion

CxO Security Forum Online Community Portal: www.engagez.net/CxO

More information on this agenda: CxOSecurityForum.com/NYC

Details & Summaries from Gartner Security & Risk Management Summit 2026

- Summary & Notes:
  - Tuesday Executive RoundTable
  - Wednesday Healthcare Executive RoundTable
- Article: Key Summit Sessions June 2026
- Article: Gartner Keynote @ Summit 2026
- Article: Third Party Cyber Risk Management (TPCRM)