Boston - 23 June 2026, Orrick

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

Executive Summary

The Boston CxO Security Forum brought together senior cybersecurity, technology, risk, and business leaders for a day of candid, practitioner-led discussions focused on some of the most consequential issues facing organizations today: Agentic AI, executive accountability, identity governance, operational resilience, third-party risk, and the future cybersecurity workforce.

Participants moved beyond headlines and vendor messaging to examine the real-world challenges of governing rapidly evolving technologies, managing organizational risk at machine speed, and preparing for a future in which AI fundamentally reshapes both cyber defense and the workforce itself.

A consistent theme emerged across every discussion: technology is advancing faster than traditional governance, operating models, and talent pipelines can adapt, creating both new opportunities and new risks for leaders. The summaries that follow capture the key insights, takeaways, and discussion themes shared by speakers and participants throughout the Forum.

We extend our sincere thanks to **Banc of California**, our Strategic Partner for the event; **Orrick**, our gracious host in Boston; and **Protiviti**, **Robert Half**, and **Opti.ai** for their support of the CxO Security Forum community and their commitment to advancing meaningful dialogue among cybersecurity and technology leaders.

Discussion Leaders & Featured Sessions

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

Discussion Pod #1

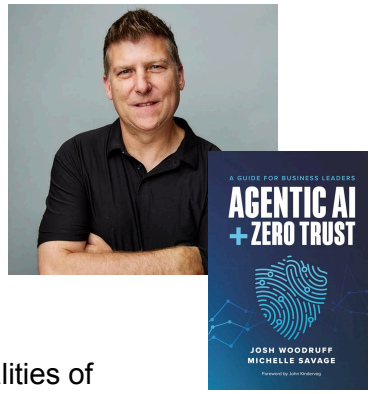

Josh Woodruff

Author, Agentic AI + Zero Trust

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

Josh challenged us to move beyond AI hype and focus on the operational realities of deploying autonomous AI inside the enterprise. His central argument was that AI failures are rarely technical failures—they are leadership, governance, and architecture failures. Organizations that achieve meaningful AI outcomes are not simply deploying tools faster; they are rethinking business processes, governance models, and security controls from the ground up.

A recurring theme throughout the discussion was that AI agents should not be viewed as software tools, but as actors operating inside the organization. Unlike employees, however, these "digital workers" arrive without background checks, judgment, or an understanding of consequences. As Josh noted, *"autonomy without governance is privilege escalation at machine speed."*

The group explored why traditional security approaches are insufficient for Agentic AI and why Zero Trust principles must evolve to address non-human identities, intent, authorization, segmentation, and continuous monitoring. Discussion also focused on the need for architectural guardrails outside the AI model itself, accountability for AI-driven actions, and the growing importance of containment strategies as machine-speed decision making accelerates.

Key Takeaways

- **AI amplifies existing strengths and weaknesses.** Organizations that automate broken processes simply create faster broken processes. Real ROI requires redesigning workflows and governance, not just deploying AI.
- **Agentic AI changes the threat model.** The primary concern is no longer unauthorized access, but authorized misuse by autonomous systems operating at machine speed.
- **Zero Trust becomes mandatory.** Identity, behavioral monitoring, data governance, segmentation, and rapid response controls must extend to AI agents and non-human identities.

- **Guardrails must be architectural, not instructional.** Security controls cannot rely solely on prompts or model instructions; deterministic controls outside the model are required.
- **Security teams have an opportunity to become AI enablers.** Organizations that successfully move AI into production increasingly depend on security leaders to provide governance, trust, and operational guardrails rather than simply saying "no."
- **Containment matters more than enumeration.** As AI accelerates attack and decision cycles, limiting blast radius and rapidly containing failures becomes more important than trying to catalog every possible risk.

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

Discussion Pod #2

Tim Brown

Former Chief Information Security Officer, SolarWinds

"The support of the company, the community, and fellow CISOs is what got me through it."

In one of the most candid discussions of the Forum, Tim Brown shared his firsthand experience leading SolarWinds through the SUNBURST supply chain attack, the years-long recovery effort that followed, and the unprecedented SEC enforcement action brought personally against him as CISO. What emerged was not simply a breach story, but a leadership story about resilience, transparency, governance, and the evolving realities of executive accountability.

Tim described the first days after learning of the compromise as "controlled chaos," requiring rapid coordination across legal, engineering, communications, law enforcement, customers, regulators, and governments around the world. Rather than focusing on blame or public relations, SolarWinds made a deliberate decision to prioritize transparency and customer support, even when doing so created additional legal and reputational risk.

The discussion then shifted to the SEC case and the personal toll of becoming the first sitting CISO charged by the agency. Tim provided a rare look inside the legal process, the importance of executive support, clear role definitions, strong D&O coverage, and the distinction between accountability and authority. He emphasized that CISOs must understand exactly what they own, what they do not own, and how risk decisions are documented and communicated throughout the organization.

Throughout the conversation, attendees repeatedly returned to a central question: as regulatory scrutiny increases, how can CISOs remain effective leaders while managing growing personal

risk? Tim's answer was clear—build trust, document decisions, communicate transparently, and never face major incidents alone.

Key Takeaways

- **Major cyber crises are company crises, not security crises.** Extinction-level events require CEO, board, legal, communications, and operational leadership working as a unified team.
- **Transparency builds long-term trust.** SolarWinds chose to communicate openly with customers, governments, and regulators even when the legal path of least resistance may have been silence.
- **Executive support matters.** Tim contrasted his experience with other high-profile cases, noting that strong support from company leadership, legal counsel, and peers proved critical during years of regulatory scrutiny.
- **Define accountability before a crisis occurs.** CISOs should maintain clear job descriptions, understand disclosure responsibilities, and ensure governance structures accurately reflect decision-making authority.
- **Preparation extends beyond technology.** Crisis management planning, tabletop exercises, alternative communications channels, legal readiness, and trusted external partners are as important as technical controls.
- **Use adversity as a catalyst for improvement.** SolarWinds leveraged the incident to rebuild trust, strengthen Secure-by-Design practices, improve software assurance, and ultimately emerge stronger than before.

Memorable Quotes

> "This is not a security incident. This is an extinction-level event for the company."
>
> "The board may decide to replace the CISO. You can't spend your energy worrying about that. You have to spend it helping customers and fixing the problem."
>
> "Be not just good—be exemplary."

AI-Native Identity: Why IAM Is Being Rewritten

Discussion Pod #3

Barak Perelman

Founder & CEO, Opti

"Identity governance is no longer one program. It's a matrix of jobs across identity types."

Drawing on research conducted with dozens of identity, security, and IT leaders, Barak Perelman argued that identity governance is approaching a breaking point. The challenge is no longer collecting access data—it is acting on it at enterprise scale. Traditional IAM and IGA programs were designed for a world of human users, periodic reviews, and static governance processes. Today's reality includes sprawling SaaS environments, growing populations of non-human identities, and the emergence of AI agents that don't fit existing governance models.

A central theme of the discussion was that most identity teams are overwhelmed by manual work. Access reviews, joiner-mover-leaver processes, audit preparation, exception handling, and entitlement analysis consume the majority of available resources, leaving little time for actual risk reduction. Participants noted that many identity programs still rely on workflows and review processes that were struggling to scale before AI entered the enterprise. AI agents and machine identities are now accelerating that challenge.

Barak challenged the industry's current approach of layering generic AI onto existing IAM platforms. He argued that identity requires purpose-built AI models trained specifically on identity data, access relationships, entitlements, governance policies, and business context. The future of IAM is not simply adding chat interfaces to legacy platforms, but reimagining governance around AI-native automation, continuous decision-making, explainability, and dynamic access models.

The discussion also explored a broader shift in thinking: identity is increasingly becoming a business efficiency and governance problem as much as a security problem. Faster onboarding, automated access decisions, reduced audit burden, and improved governance may ultimately prove more compelling to executive leadership than traditional security metrics alone.

Key Takeaways

- **Identity governance has fundamentally changed.** Human identities are now joined by service accounts, machine identities, and AI agents, creating a governance challenge that existing processes were never designed to handle.
- **Most IAM teams spend the majority of their time on administration rather than risk reduction.** Access reviews, lifecycle management, audit preparation, and exception handling consume resources that could otherwise be focused on security outcomes.
- **The problem is no longer data collection.** Organizations generally have visibility into identity data; the challenge is making timely, informed, and scalable decisions based on that data.
- **Generic AI is not enough.** Effective identity governance requires purpose-built models trained specifically on identity context, entitlement structures, policies, and organizational behavior—not general-purpose LLMs wrapped around IAM workflows.
- **AI's greatest value may be workflow automation.** The biggest opportunity is reducing the enormous amount of manual coordination, context gathering, and repetitive administrative work that currently burdens identity teams.

- **Governance and explainability remain essential.** AI-generated identity decisions must be auditable, explainable, and aligned with business policy if organizations are to trust them in production environments.
- **The next wave of IAM will be dynamic.** Continuous governance, adaptive access decisions, just-in-time access, and identity-aware automation are likely to replace many static review and provisioning processes that dominate today's programs.

Memorable Themes

> "The bottleneck is no longer getting the data. The bottleneck is acting on it."

> "The right AI for identity is not a generic LLM in a wrapper—it is purpose-built identity intelligence."

> "Access sprawl builds between every audit, and AI agents will only accelerate it."

AI Isn't Breaking Security — It's Exposing the Operating Model

Discussion Pod #4

Michael Monday

Managing Director, Protiviti

"The problem isn't the attacker. The problem is how we've organized ourselves to respond."

Drawing on research from more than 1,500 global executives conducted by Protiviti and North Carolina State University's ERM Initiative, Mike Monday challenged attendees to rethink one of cybersecurity's most common assumptions: that security failures are primarily technology or control failures. The data suggests something different. Despite cybersecurity remaining the top-ranked global business risk for three consecutive years, organizations continue to increase spending, headcount, and tooling while breaches continue to rise. The problem, Mike argued, is not a lack of controls—it's the growing gap between how quickly threats evolve and how slowly organizations make decisions.

A recurring theme throughout the discussion was velocity. AI is dramatically compressing the time required for attackers to identify vulnerabilities, chain exploits, and execute attacks. Meanwhile, most organizations still operate on quarterly board reviews, monthly patch cycles, annual vendor assessments, and governance processes designed for a much slower era. As AI accelerates both attack capabilities and business adoption, organizational speed—not technical capability—is increasingly becoming the limiting factor.

The conversation also highlighted the rapid rise of third-party and ecosystem risk. Participants discussed how AI is expanding "transitive trust" relationships, where organizations inherit risk not only from their direct vendors, but also from the vendors, models, and infrastructure providers behind those vendors. The group explored whether future breaches should be viewed primarily as security failures or as broader operating model failures involving governance, accountability, resilience, and organizational design.

Rather than advocating for more tools, Mike encouraged leaders to focus on resilience, decision-making speed, accountability, and organizational adaptability. The organizations best positioned for the AI era may not be those with the largest security budgets, but those that can most effectively absorb and respond to rapid change.

Key Takeaways

- **Cybersecurity remains the top business risk, but more spending alone is not solving the problem.** Organizations continue to increase investment while breaches continue to rise, suggesting deeper organizational issues are at play.
- **AI is accelerating attackers faster than organizations are accelerating governance.** The gap between threat velocity and decision velocity is becoming one of the most significant risks facing enterprises.
- **The next major incident is likely to be an operating model failure.** Governance, accountability, resilience, and organizational alignment may prove more important than any individual security control.
- **Third-party risk is rapidly expanding.** AI-enabled products, SaaS ecosystems, open-source dependencies, and cloud providers are creating increasingly complex chains of trust that most organizations struggle to assess.
- **Traditional security timelines are becoming obsolete.** Quarterly reviews, annual assessments, and lengthy remediation cycles were designed for a slower threat environment than the one organizations face today.
- **Resilience may matter more than prevention.** As attack speeds increase, organizations must focus on detection, containment, recovery, and organizational preparedness rather than assuming every threat can be prevented.
- **Boards are asking new questions.** Executive leadership increasingly wants to understand AI-related risks, accountability, vendor posture, governance maturity, and the organization's ability to safely absorb AI-driven change.

Memorable Quotes

> "AI didn't break your security program. It exposed the gap between how fast threats move and how fast organizations make decisions."

> "The next material incident won't trace back to a missed control. It will trace back to an operating model failure."

"The organizations that win won't necessarily be the ones with the biggest budgets—they'll be the ones that adapt the fastest."

Discussion Pod #5 - Community

Building the Next Generation Cybersecurity Workforce

Chris Zannetos

Founder & Chairman, STEMatch

"The curriculum wasn't the biggest problem. The work skills gap was."

Chris Zannetos brought a practical and provocative perspective to one of the industry's most persistent challenges: finding and developing cybersecurity talent. Drawing on decades of experience building technology companies and hiring security professionals, Chris argued that the cybersecurity workforce shortage is not primarily a skills problem—it's a pipeline problem. Organizations continue to struggle to hire entry-level talent while simultaneously demanding two to five years of experience, creating a gap that has persisted for years and is now being amplified by AI-driven changes in the workforce.

The discussion focused on lessons learned from STEMatch's COMPETE program, which works with community colleges and industry partners to identify, mentor, train, and place students from traditionally underserved backgrounds into technology and cybersecurity careers. Through extensive research and direct engagement with employers, STEMatch found that most community college students possessed adequate technical foundations but lacked professional networks, workplace experience, communication skills, and critical thinking capabilities necessary to succeed in enterprise environments. The program was specifically designed to close those gaps through mentoring, work-skills development, internships, and industry engagement.

A significant portion of the conversation centered on AI's impact on the future workforce. While many industries fear AI will reduce employment opportunities, the consensus among participating CISOs was that cybersecurity will likely require more professionals, not fewer. However, the nature of those jobs will change. Routine tasks may be automated, but demand is expected to increase for higher-order skills such as systems thinking, architecture, governance, project leadership, communication, and problem solving. The challenge facing the industry is ensuring there remains a viable path for new entrants to gain the experience required to eventually fill those roles.

The session ultimately challenged attendees to rethink traditional hiring practices, broaden talent pipelines, and consider whether the industry's future workforce may come from places many organizations have historically overlooked.

Key Takeaways

- **The cybersecurity talent shortage is largely a pipeline problem.** Organizations continue to seek experienced professionals while investing insufficiently in developing the next generation of talent.
- **Technical skills are only part of the equation.** Communication, teamwork, critical thinking, professional networks, and workplace experience were identified as the biggest gaps preventing many students from succeeding in entry-level cybersecurity roles.
- **Community colleges represent an underutilized talent source.** With the right support, mentoring, and industry exposure, students from two-year programs can become highly effective cybersecurity professionals.
- **AI is changing the shape of entry-level work.** Routine tasks are increasingly being automated, raising concerns about how future professionals will gain the practical experience historically required to advance into senior roles.
- **Future cybersecurity roles will emphasize higher-order skills.** Systems design, architecture, governance, project management, solutions engineering, and critical thinking are expected to become increasingly important.
- **Professional networks matter.** Access to mentors, employers, and industry relationships often proves as important as technical education in helping students launch successful careers.
- **Employers may need to rethink traditional hiring requirements.** Limiting recruiting efforts to a small set of universities or relying heavily on automated screening processes can unnecessarily restrict access to qualified talent.

Key Data Point

STEMatch's current COMPETE program includes **104 students across Massachusetts and Maryland, including 45 cybersecurity students, with participating employers rating 92% of interns as "hireable."**

Memorable Quotes

> "We've always talked about a cybersecurity talent shortage, but we've also spent years demanding two to five years of experience for entry-level jobs."
>
> "The future will likely require more cybersecurity professionals, not fewer—but they'll be doing very different work."
>
> "If Draper Labs can hire a community college intern, so can you."