Session Title	**Description**
From Retrieval to Reasoning: Building Defensible Legal AI with Structured Knowledge	Generative AI is quickly becoming embedded in legal workflows, but the quality of its outputs depends entirely on how well it is grounded in organizational knowledge. Systems built primarily on document retrieval can surface relevant text, yet often fall short when asked to produce answers that reflect legal reasoning, context, and precedent. This session explores what it truly means to ground AI in your organization's data and why structured knowledge models are essential for trustworthy legal systems. Attendees will learn how knowledge graphs capture the relationships that matter in legal work - matters, parties, clauses, outcomes, jurisdictions, and playbooks - and how this structure improves accuracy, traceability, and institutional learning. We'll compare retrieval-centric architectures with graph-grounded approaches, examine real litigation and transactional workflows, and outline a practical roadmap for moving from document search to connected, governed legal intelligence.
AI for Brains That Work Differently	Neurodivergent professionals bring powerful strengths to legal work, including creativity, pattern recognition, and deep focus - but they may also face challenges with organization, cognitive load, and information density. This session focuses on how thoughtfully deployed AI can serve as a practical support layer rather than a replacement for human judgment. Attendees will explore concrete, real-world use cases for applying AI to structure complex writing, manage deadlines, break down dense legal analysis, and improve comprehension through summarization and text-to-speech tools. The session will also provide a practical toolkit of prompts, applications, and inclusive implementation practices that legal IT teams can adopt to support neurodiverse colleagues—and to improve productivity and accessibility across the organization.

From Hype to FOMO to Fully Funded: Building a Defensible AI Investment Plan and Roadmap that is Boardroom Ready	It's imperative that organizations navigate the AI hype and FOMO so they can build an AI investment plan and roadmap that adheres to their core values and matches or exceeds their client requirements in order to not just survive the Boardroom, but leave the Board asking for more. This session will tackle: Avoiding the classic legal tech project pitfalls and clearly identifying where your AI strategy can add real qualitative benefit Translating AI/GAI from technical curiosity into a clear business case that will resonate with managing partners, practice leaders and finance committees Building tangible roadmaps for thoughtful AI adoption Developing best practices to measure and maximize the AI ROI story Aligning on the organization’s strategy to realize matter profitability, reduce write downs, increase cycle time and client value, confidence and loyalty Attendees should walk away from this session with concrete and viable steps they can take to combat the pervasive “AI is everywhere but our strategy is nowhere” conundrum. Rather than defending investments in the Boardroom, master how to quickly garner leadership support, accelerate adoption and develop your organization’s AI/GAI strategy ultimately gaining a clear competitive advantage.
Dungeons, Data & Decisions – Building your AI Adventure Party	AI adoption isn’t a solo quest—it’s a campaign that demands a balanced party of skills and roles. In this gamified session, participants will learn how to “roll up” the right mix of AI agents and human leaders to tackle organizational challenges. Through interactive exercises, attendees will build adventuring parties using archetype cards, then face simulated quests like defending against a data breach or navigating an ethical dilemma. The goal? To demonstrate that successful AI integration requires teamwork, leadership, and critical thinking—not just technical know-how.

It’s the Start That’s Stopping You (or: How I Learned to Stop Worrying and Love Generative AI)	This session explores AI not as a threat to identity or relevance, but as a practical ally for creativity, learning, and well-being—both inside and outside of work. Anchored in ILTA’s commitment to mental health and support of learning advanced technology, the session reframes AI as a tool that lowers the barrier to starting: starting a new skill, a creative practice, a hobby, or even a difficult professional transition. Rather than focusing solely on productivity or efficiency, the session addresses the emotional realities of AI adoption—fear of obsolescence, imposter syndrome, and anxiety around changing roles—while grounding the conversation in governance, ethics, and intentional use. Participants will explore how AI can support personal growth (such as writing, music, learning new skills, or reflective practices) and how those same principles translate back into healthier, more confident professional engagement. With a tone that acknowledges our collective tendency to catastrophize new technology (and gently invites us to stop worrying), the session offers a humane, empowering perspective: AI doesn’t replace purpose or creativity—it helps us begin. And sometimes, beginning is the hardest part.

How data governance drives successful AI	AI adoption in law is accelerating faster than governance frameworks can keep up. It's also exposing existing issues with data proliferation. While this enthusiasm for innovation is promising, it creates serious risks: breaches of client confidentiality, ethical violations, and even malpractice exposure. Attendees will hear how strong governance practices create the conditions needed for reliable AI: accurate data, controlled access, clear accountability and consistent compliance with privacy and regulatory obligations. Attendees will learn: How data lifecycle management can reduce cyber risk and cost. Why strong governance is essential for trustworthy AI. How security, governance and AI work best when aligned rather than treated as separate initiatives.
Finding the balance between risk and innovation.	The big question in legal when it comes to AI seems to be - "Can we use it safely and effectively?" This session will show you the steps necessary to implement the governance policies we discussed in Part 1 and be able to answer that question in the affirmative. Attendees will learn how to align AI initiatives with organization objectives—enhancing efficiency, client service, and risk management—while establishing robust governance and ethical frameworks for responsible AI use. We'll cover creating a culture of safe AI usage, actionable training for legal teams, how to evaluate shiny new AI tools, managing AI usage with third-parties, etc. Attendees will walk away with a governance structure that balances innovation with safety ensuring GenAI enhances legal work without introducing risk.

The Evolving AI Policy & Legal Landscape	This presentation will review some of the challenges and questions related to AI law and policy in the last year. It will provide an overview of key legal and policy developments as well as a discussion of "what's next" in terms of notable proposed legislation and regulations. It will cover Federal and State efforts to regulate AI law and drive policy, as well as global developments in Asia, the EU, and the UK.
The AI-Driven Legal Team: Practical GenAI Tools Lawyers Can Safely and Easily Use Today	Legal professionals are eager to adopt Generative AI but are overwhelmed by the noise and worried about confidentiality, accuracy, and ethics. This session cuts through the hype. We showcase real, responsible GenAI workflows that organizations can use right now across both the business and practice of law, with topics running from invoice processing to contract comparison to knowledge extraction from large document sets. With a strong focus on auditability, privacy, and secure model use, attendees will learn how to roll out GenAI that is reliable, defensible, and aligned with professional responsibility. Attendees will leave with a set of high-ROI, low-risk ideas for implementing AI in their organizations today.
License to Automate: Recruiting Your Own Secret Agents	Discover how to clone your best productivity by building custom AI agents that serve as tireless extensions of your legal team. This session offers a hands-on look at the tools required to automate workflows in client intake, document drafting, and practice management without a computer science degree. Join this session and walk away with the gadgets and guidance needed to launch your first mission-ready assistant.
Synthetic Evidence: The Next Frontier in eDiscovery	As generative AI tools proliferate, the likelihood of encountering synthetic evidence in litigation is no longer hypothetical. This session explores how eDiscovery professionals can identify, authenticate, and challenge AI-generated content. We’ll examine technical detection methods, legal standards for admissibility, and strategies for mitigating risk when synthetic data infiltrates corporate systems or becomes part of a production set.

From Prompting to Context Engineering: The New KM Superpower	GenAI is only as good as the context you give it. In legal organizations, that context is exactly where KM shines. This session reframes KM professionals not as “prompt writers” but as architects of the information environments that make AI useful, safe, and trusted. We will explore how to move from one-off prompting tricks to systematic context engineering: shaping data, metadata, matter profiles, playbooks, precedents, and policies so that agents consistently produce grounded, auditable, and repeatable outcomes. Using real-world scenarios from research, drafting, and matter management, we will examine how to design AI-ready knowledge assets, structure retrieval around legal work types, and bake risk posture directly into AI workflows. The focus is practical and role-specific: how KM can own the fabric that GenAI runs on, partner with IT and practice leaders, and turn “KM as a library” into “KM as an AI performance engine.”

Does GenAI Supercharge or Eliminate KM?	As legal teams move beyond experimentation with generative AI, knowledge management is emerging not as an endangered function but as a critical competitive advantage. This session examines how leading organizations are integrating GenAI to transform how knowledge is captured, accessed, shared, and applied across legal delivery and business operations. Attendees will explore real-world implementations and measurable outcomes, including improved matter efficiency, faster access to precedent and expertise, stronger collaboration, and more informed decision-making. The session moves beyond theoretical possibilities to demonstrate what successful integration looks like in daily practice and why GenAI is expanding rather than replacing the core discipline of knowledge management. Key Topics: Integration strategies: How leading organizations are embedding AI into existing knowledge management frameworks Practical use cases with demonstrated operational and financial impact Day-to-day realities of successful implementation Strategic considerations for planning the next phase of AI-enabled knowledge management

Safeguarding Legal Data: Advanced Data Protection Strategies	Legal teams are racing to adopt generative AI tools like Microsoft Copilot to boost efficiency and deliver client value. But with this innovation comes heightened security and compliance risks. Managing vast volumes of highly sensitive client data in your document management system (DMS), makes strategies like Data Loss Prevention (DLP) and data segregation mission-critical. In this session, we'll explore how you can embrace AI responsibly without compromising confidentiality or compliance. Discover strategies to reduce risk and prevent accidental data leaks while enabling practitioners to work smarter. Key Takeaways: •How to leverage integrations within your legal tech stack to create a secure foundation for AI •Techniques to minimize false positives and alert fatigue while safeguarding sensitive client data •How to get the most out of ethical walls/information barriers to prevent data exposure •Best practices for securing boundaries for generative AI tools to accelerate productivity without sacrificing security

The Hidden Toll: Navigating the Emotional Side of a Data Breach	Data breaches are often framed as technical failures or compliance headaches. Yet behind the forensics and crisis calls lies a profound human challenge: the emotional toll on the very teams tasked with protecting client confidentiality and restoring trust. For CIOs, CISOs, and organization leaders, the pressure to respond swiftly while safeguarding reputation can create an environment of relentless stress, fear of professional failure and even long-term psychological strain. This session goes beyond the technical playbook to address the human dimension of breach response. Drawing from real-world incidents, we'll explore how leadership can recognize and mitigate the emotional impact on IT teams and executives, foster resilience during high-stakes recovery and build a culture that prioritizes wellbeing without compromising operational resilience. Key Takeaways: •Understand the psychological toll of breach response on IT and leadership teams •Practical steps for Security leaders to reduce stress and prevent burnout during prolonged incidents •How to foster a culture of care and support without compromising security or compliance

From Cloud Chaos to AI Driven Control: Rethinking Data Classification & Governance	As legal organizations accelerate cloud adoption and generate vast amounts of unstructured data—emails, documents, and more—traditional governance frameworks are struggling to keep pace. The exponential growth of data, combined with unclear accountability in multi-cloud environments, is fueling a governance crisis marked by compliance blind spots and security vulnerabilities. Manual classification is no longer sustainable; instead, AI-powered governance is becoming essential, automating document classification, metadata tagging, and retention policy enforcement. This session will explore why current governance models are insufficient and how enterprises can rethink their architectures to clarify shared responsibility and build resilience. Attendees will learn how AI transforms governance from a reactive process into a proactive, real-time system that reduces errors, accelerates compliance, and frees resources for strategic priorities. Special emphasis will be placed on the critical role of automated data classification in mitigating risk and ensuring compliance in the legal sector. Key Takeaways: Understand the risk landscape created by unstructured data and multi-cloud proliferation. Discover how AI can automate document classification supporting retention policies, reducing manual effort and errors. Learn strategies to clarify accountability and shared responsibility across providers. Gain insights into emerging trends and proactive steps to avoid compliance failures and security breaches in 2026 and beyond. Explore implementation strategies for AI-driven governance and real-time compliance monitoring.

AI-Ready Security: Evolving Your Controls	Introducing AI into legal operations can feel like inviting an unpredictable force into a delicate environment. This session provides a practical roadmap for modernizing your existing security program to confidently manage AI risk. As AI rapidly permeates legal workflows—from research and drafting to client collaboration—security and IT leaders face the challenge that traditional information security controls weren't built with AI in mind. Fortunately, many of the controls your organization already relies on can be expanded and evolved to meet emerging AI risks head-on. Learn how to transform familiar security disciplines into powerful AI-governance mechanisms that protect client confidentiality, ensure responsible AI adoption, and reduce organizational risk. Attendees will leave with actionable strategies for adapting their security control ecosystem to support AI adoption—without reinventing their entire program. This session is designed for leaders who want to strengthen their AI defenses using tools they already know and trust.
Beyond the Black Box: Legal Leadership in AI Ethics	AI is transforming legal practice, but its lack of granular controls and sweeping data access leaves individuals vulnerable to client data disclosures and amplifies insider risk. This fireside chat explores how leadership can respond with ethical guardrails, governance frameworks, and training programs that balance innovation with responsibility. Attendees will gain practical strategies for fostering accountability, embedding ethical grounding into daily workflows, and guiding people to make sound decisions in an era defined by AI-driven legal practice.

I’d Give Away My Password and MFA—And Still Sleep at Night	As remote work and cloud migration expand the attack surface, identity has become the new crown jewel — and the prime target for compromise. This session explores how organizations can harden defenses through different identity provider's controls, conditional access policies, and zero-trust architectures that ensure stolen credentials alone can’t breach systems. Experts will also highlight how layered solutions such as SASE, SIEM, and advanced email protections build resilience in SaaS-first environments. Attendees will leave with a practical roadmap for securing their digital identity and protecting client trust.
Rethinking the Fort: Building Resilience in the Modern Legal Organization	As cyber threats evolve and client expectations heighten, organizations are re-evaluating their approach to cybersecurity. Once viewed primarily as a technical or compliance obligation, cybersecurity is now recognized as a foundational business practice that directly impacts firm reputation, operational efficiency, and client trust. True protection extends beyond tools and policies; it depends on governance, readiness, and clarity established before disruption occurs. By embedding a culture of security across operations; supported by clear leadership ownership, disciplined governance, and well-practiced incident response capabilities, organizations are better positioned to prevent disruption and respond effectively when incidents arise. In this session, participants will explore a practical roadmap for building cybersecurity readiness and maturity. The discussion examines the key pillars of a resilient cybersecurity framework, including leadership alignment, governance models, vendor risk management, incident readiness, and continuous improvement. Attendees will gain actionable insights to assess their current level of readiness, identify common maturity gaps, and prioritize improvements that reduce risk and support secure collaboration; without impeding day-to-day legal work.

Detecting Security Stack Bypass Before Breaches Occ	Organizations continue to expand their security stacks, yet many defensive strategies rely on the assumption that endpoint detection and logging tools alone are sufficient to protect their environments. Modern attackers increasingly exploit this assumption by establishing hidden compromise that degrades visibility, suppresses alerts, or weakens defensive controls before ransomware or extortion activity begins. Rather than focusing on specific products, this session examines the trust-model gaps that allow security controls to be bypassed. It explores how attacker techniques have evolved to operate beneath traditional monitoring layers, and how continuous evaluation of detection capabilities, response mechanisms, and control tuning can improve confidence in security telemetry. Attendees will learn how to recognize early indicators of hidden compromise and apply practical detection strategies that strengthen existing security investments, keeping vendors accountable and improving resilience, without increasing operational overhead.
Implementing Effective DLP Strategies in Legal Organizations	Law firms handle vast amounts of highly sensitive client data, making Data Loss Prevention (DLP) a critical but complex undertaking. Implementing DLP solutions in a legal environment presents unique challenges: the sheer volume and sensitivity of information can lead to frequent alerts and false positives, creating "noise" that overwhelms IT teams and disrupts workflows. This session will explore the specific obstacles law firms face when deploying DLP, including balancing security with attorney productivity, managing the flood of alerts, and ensuring compliance without hindering daily operations.

Don't Bore the Board: Presenting CyberSecurity Metrics that Matter	When presenting the success of your cybersecurity program to leadership, it’s essential to share metrics that are not only SMART (Specific, Measurable, Achievable, Relevant, Time-bound), but also clear and meaningful to lawyers focused on business risk. This session will cover best practices for selecting and communicating cybersecurity metrics that resonate with leadership, demonstrate alignment with business objectives, and highlight progress under standards like ISO 27001. Key takeaways include: What cybersecurity metrics matter most to the board and why How to make metrics relevant and understandable for non-technical audiences Strategies to connect metrics to business value and risk reduction How to use current metrics to build a compelling case for additional security budget and resources
Confidential Computing: Protecting Sensitive Data in the Ag	As adoption of Generative AI and cloud-based workflows increases, the risk surface expands, especially when sensitive data is processed in memory. Confidential computing is rapidly becoming one of the most important technologies for legal data protection. This session demystifies it. We explain how modern secure enclaves, encryption-in-use, and hardware-backed isolation can protect client files, evidence, contracts, and PHI even while models analyze them. We contrast traditional security with AI-era requirements and provide a roadmap for implementing confidential computing in practice. Attendees will leave with a clear understanding of why this technology is essential for the future of legal AI.

Real Incidents, Ready to Run Tabletop Exercises	In the last 15 years of both running and participating in Tabletop Exercises, I've seen patterns emerge. Too many exercises reuse the same scenarios and ask the same questions, missing key decision points and organizational gaps. There are better ways. This talk takes real incidents handled by Inversion6 and crafts them into Tabletop Exercise scenarios that surface hidden weaknesses within organizations. Both moderators and participants will learn techniques to get better results in exercise and increase their value to clients and organizations. Key Takeaways for Tabletop Exercise moderators: - Techniques to turn real incidents into effective scenarios - Useful prompts and injects that reveal hidden gaps Key Takeaways for Tabletop Exercise participants: - New areas to dive into in exercises - Methods to ensure you get the most out of your exercise
Tools of the Trade - Turning a Pentesting Tool into Proactive Security	Many free and open-source tools are used by attackers and pentesters alike to find weaknesses in networks, AD environments, and machines. Before the pentest occurs or the hacker gets a hold of your environment, why not use these tools yourself to find those weaknesses before they do? This session would cover tools like Bloodhound, SCUBA, Purple Knight, Certify, Nmap, WPScan, SharpShares, and more!