The 2026 Fraud Playbook: Defending the Ledger from AI-Powered Deepfakes

Presented by: Paul E. Zikmund

IOFM

Changing Risk Landscape

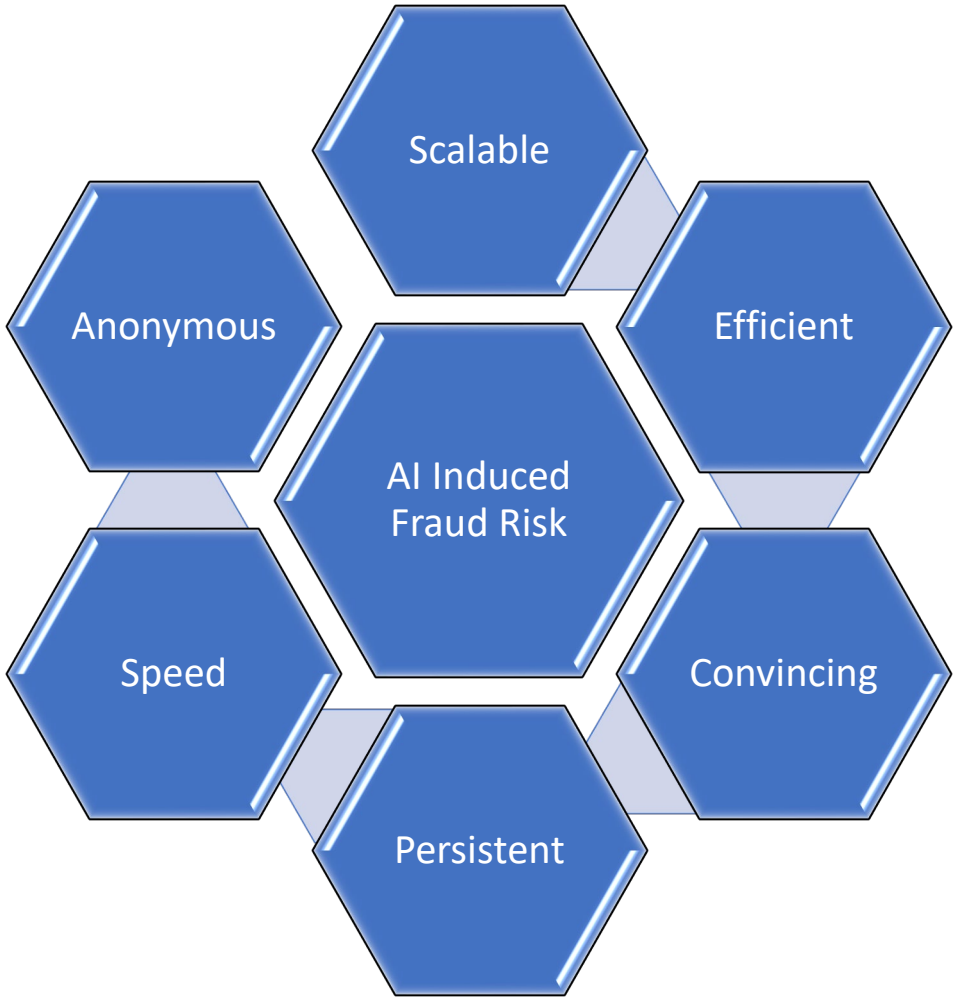

IOFM

Cybercrime is a Business

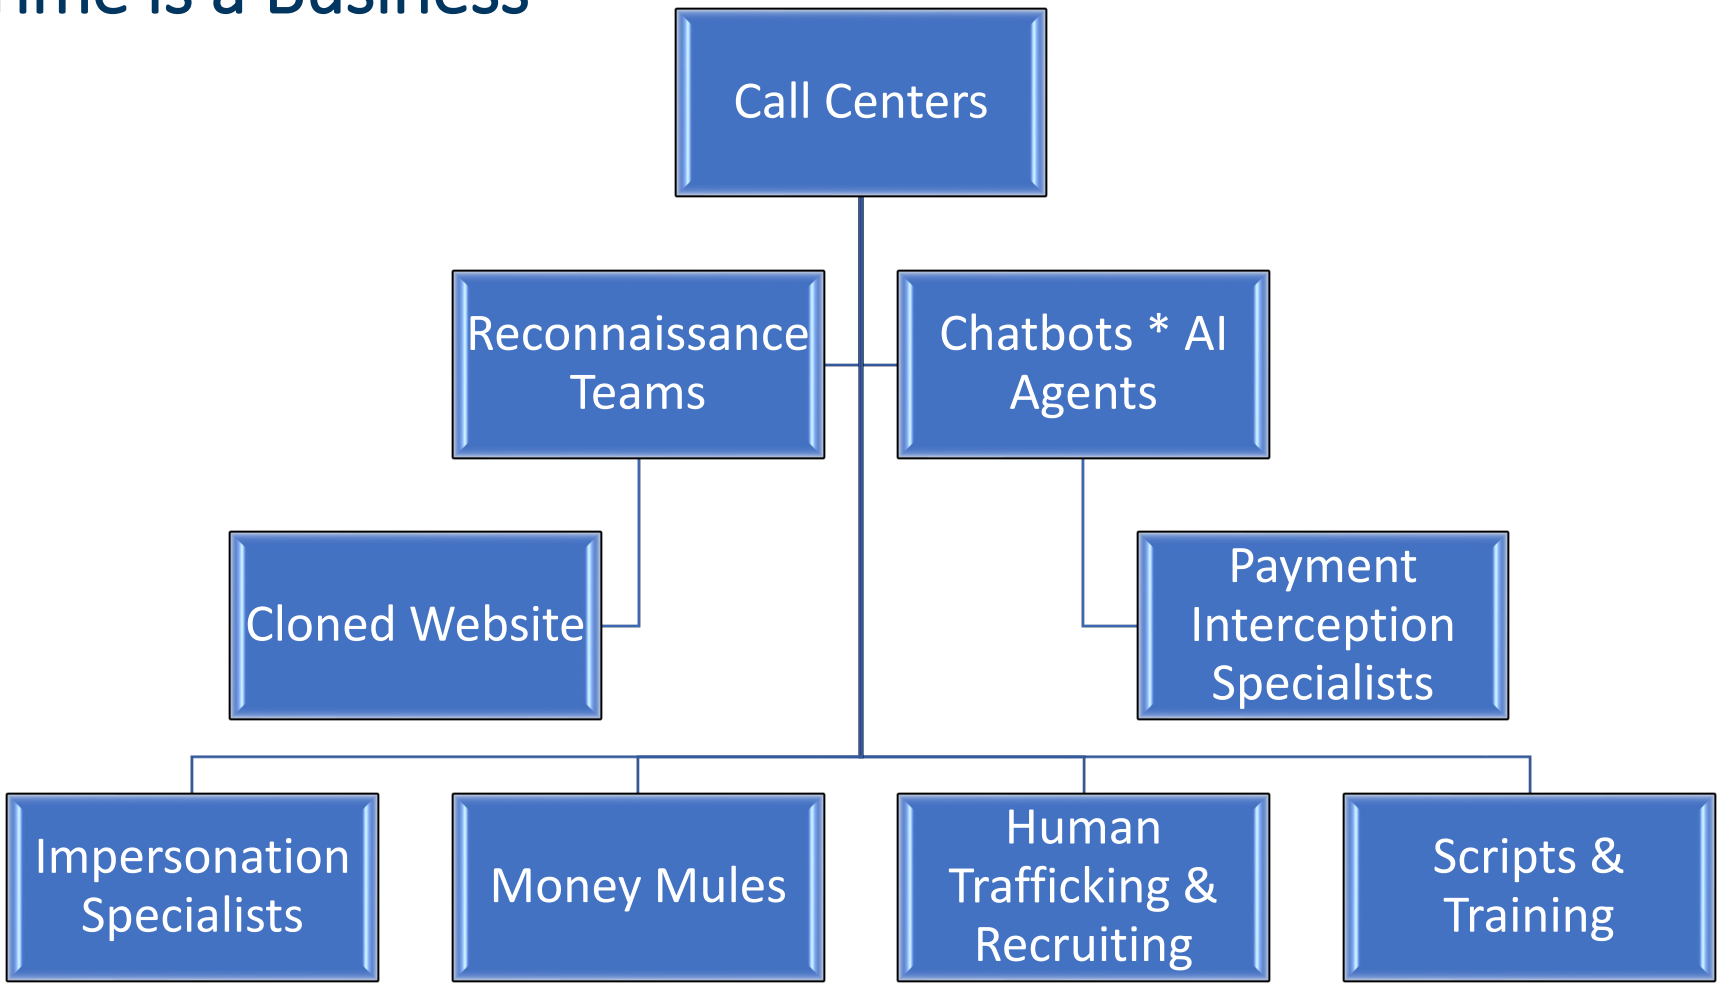

IOFM

Professionalization of Fraud

Then	Now	Capabilities
Generic phishing	Personalized, AI-generated outreach	Research vendor relationships automatically
One-off impersonation	Long-term vendor identity simulation	Simulate tone, writing patterns, and communication history
Obvious errors	Near-perfect language and formatting	Draft convincing messages tailored to recipients
Single-channel attacks	Email, phone, text, portal, and web cloning	Orchestrate multi-step deception campaigns
Manual effort	Automated, scalable workflows	Adapt messaging based on replies and resistance
Broad campaigns	Precision targeting of AP teams	Maintain persistence over weeks or months

Losses due to Cybercrime

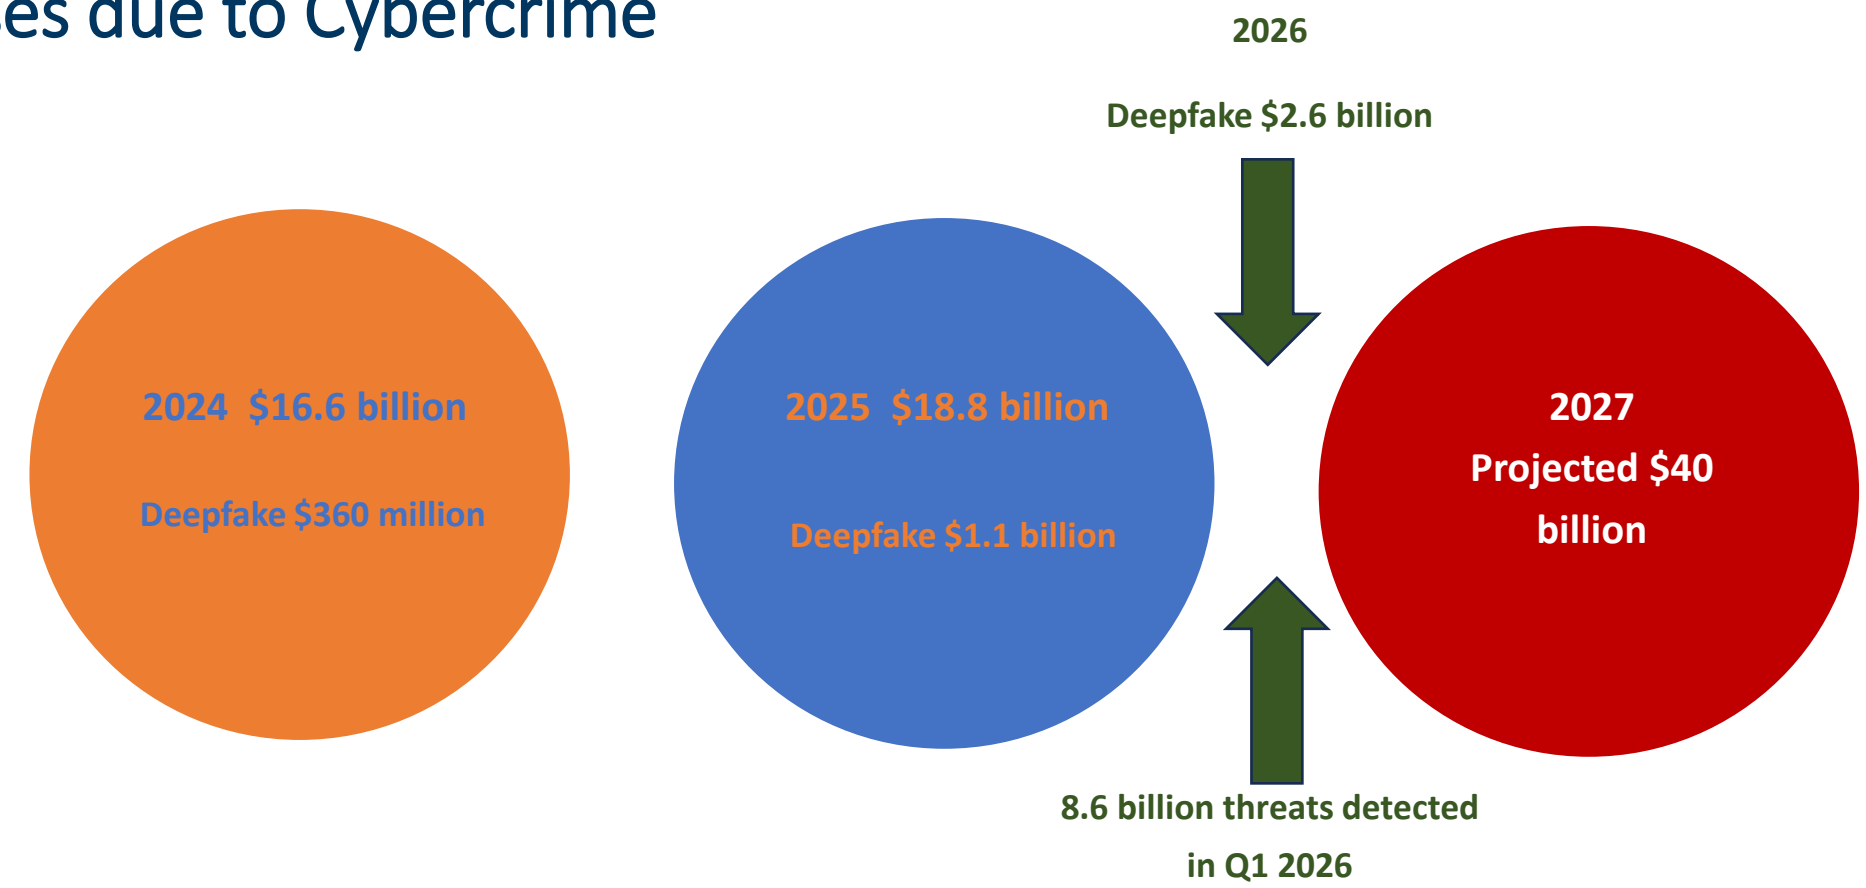

High Volume, Low Friction: Phishing remains the #1 crime by complaint volume (over 191,000 in 2025), serving as the primary delivery mechanism for more lucrative downstream social engineering schemes.

IOFM

2026 eCommerce & AFP Survey Findings

BEC is the Most Frequent Attack - affecting 75% of all organizations.

Impersonation is Expanding Beyond Email - increase in the use of phone calls and text messages to impersonate banking partners or executives.

Lag in Defensive AI Adoption - Only 17% of organizations currently report using AI to combat payments fraud.

Increase in Precision Phishing & Sophisticated Social Engineering

The "Fraud Shield" Response: In response, the market for AI-powered fraud shields is expected to grow to $10.56 billion in 2026.

Rise of Agentic Commerce, where AI agents make purchases on behalf of humans. 63% of merchants are currently exploring or planning to implement systems to handle these autonomous transactions, creating a new requirement for "machine-to-machine" identity verification.

IOFM

Social Engineering & Human Emotions

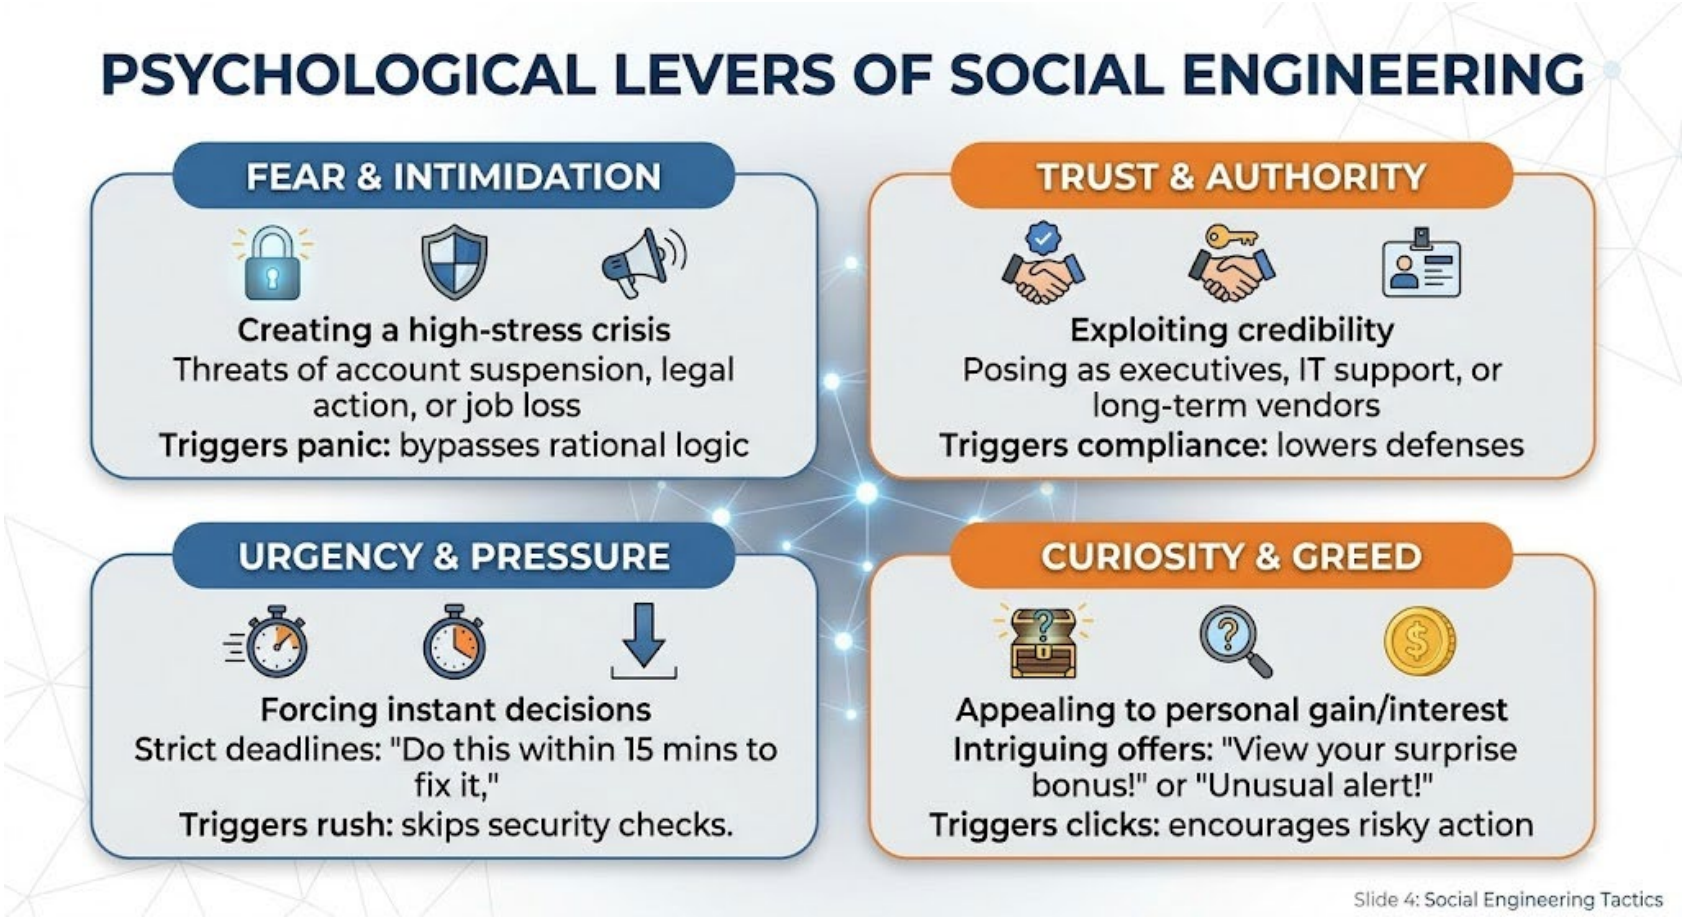

Social Engineering & Digital Behavior

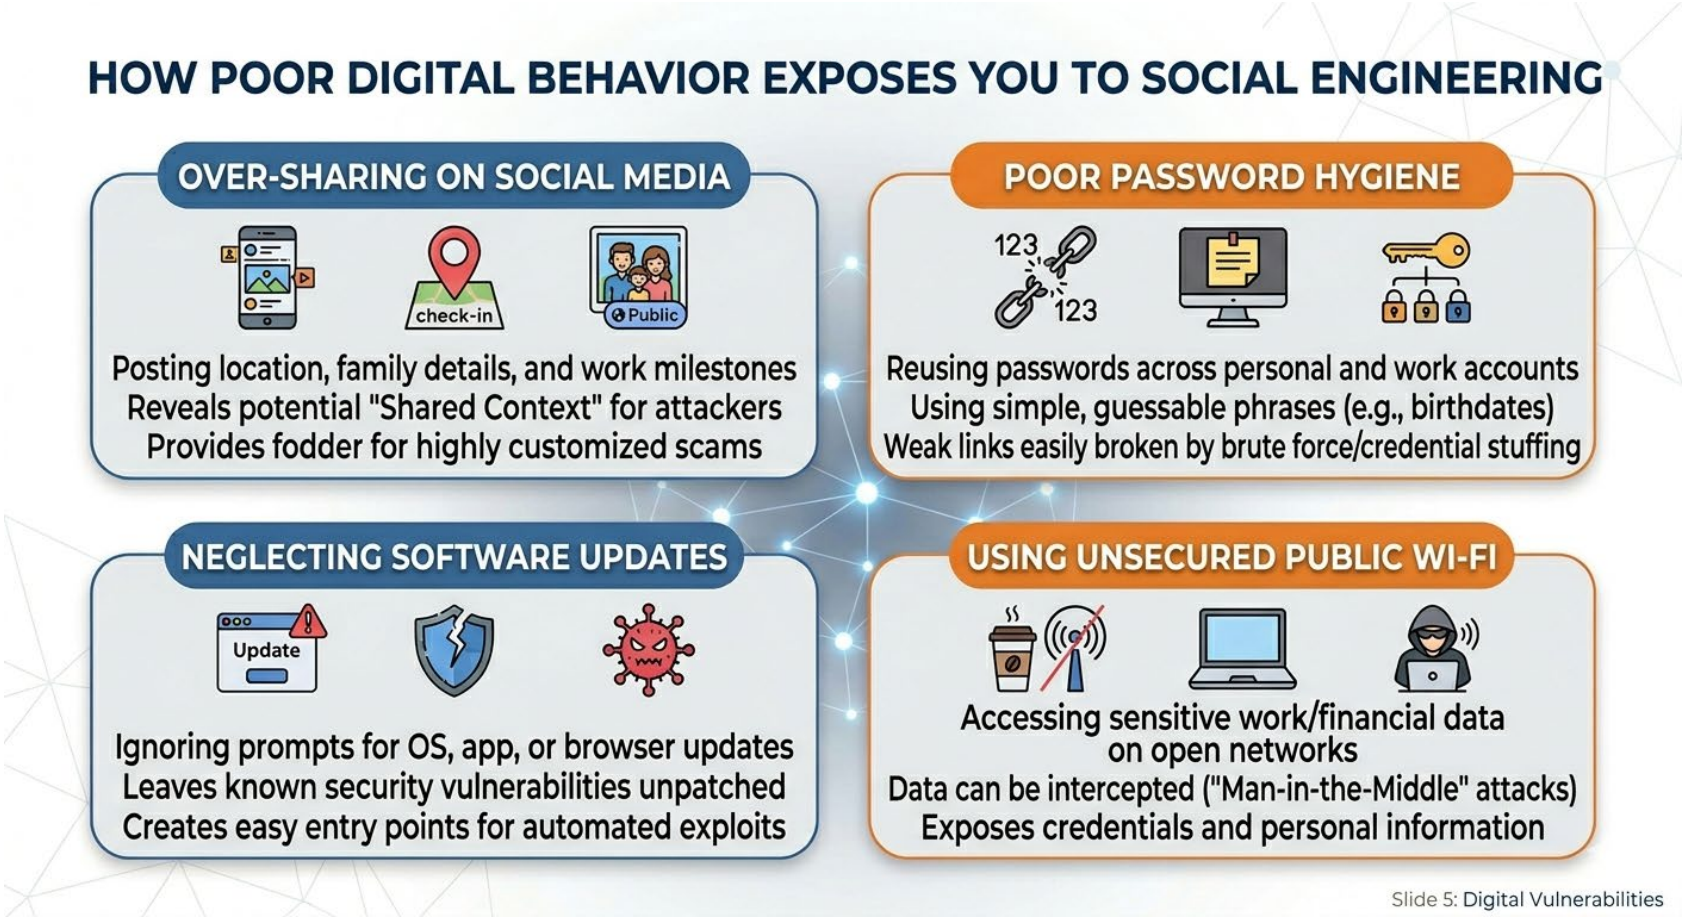

Social Engineering & Social Behavior

Slide 6: Social Vulnerabilities

What is Artificial Intelligence?

Artificial intelligence is the simulation of human intelligence processes by computer systems, typically involving the ability to learn, reason, and self-correct to perform complex tasks.

- Learns from data
- Recognizes patterns
- Makes predictions
- Generates content

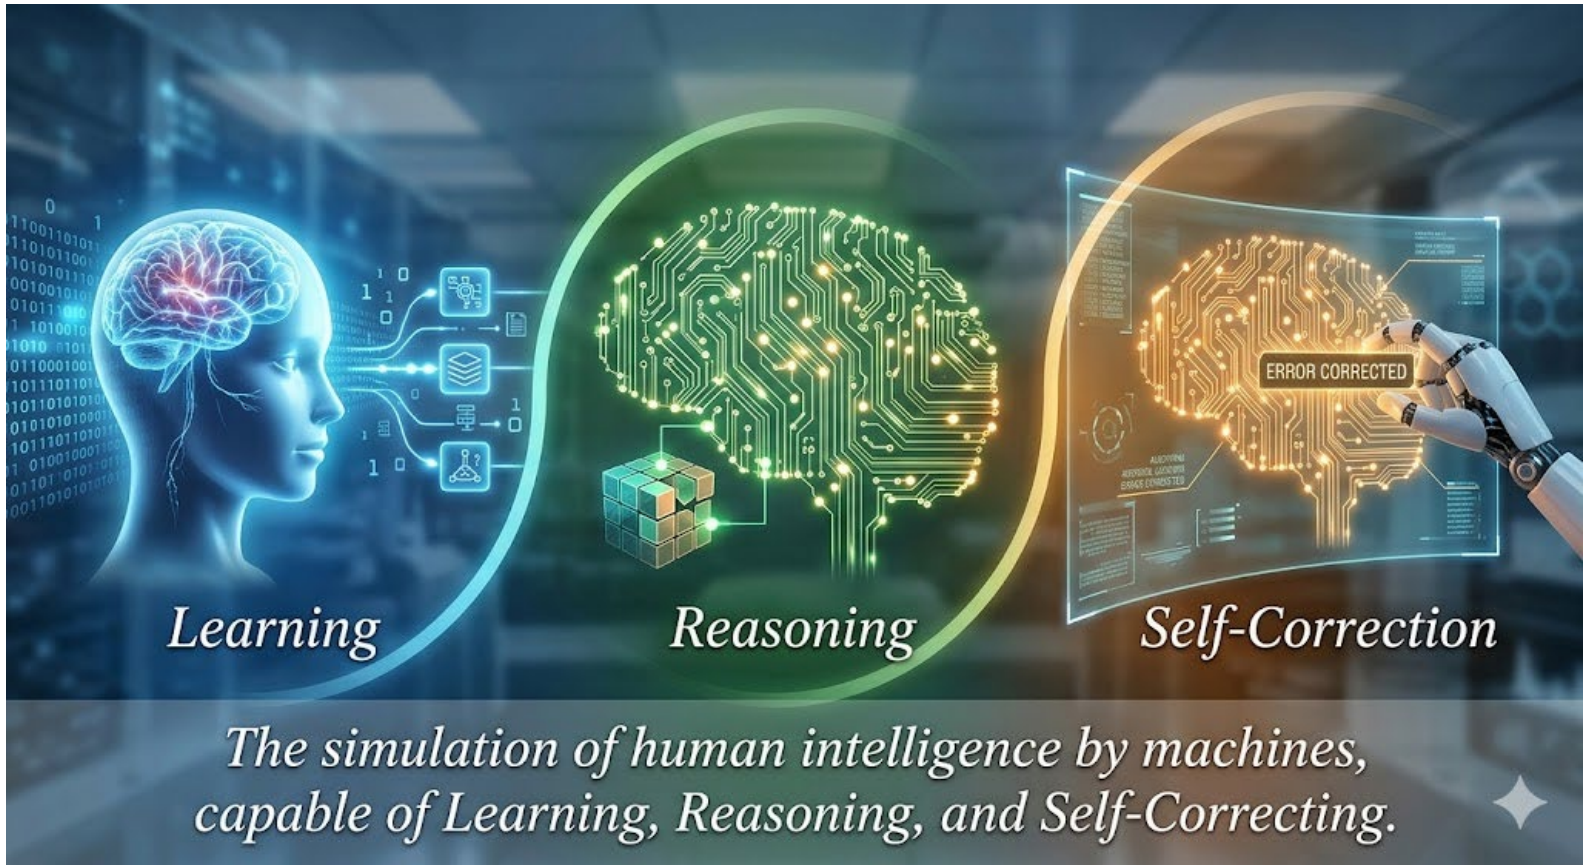

IOFM

Basic Principles of AI

Fairness—AI should be designed and used to treat all individuals and groups fairly. It should not create or reinforce bias or discrimination based on factors such as race, gender, age, or socioeconomic status.

Accountability—Entities that develop and deploy AI should be accountable for its impacts. This includes implementing mechanisms for addressing any negative effects that may arise.

Transparency and explainability—AI systems should be transparent. This means that it should be clear how they work, how they make decisions, and who is responsible for them.

Controllability—AI should not undermine human autonomy or decision making. People should have the ability to understand and challenge decisions made by AI, and to opt out of AI decision making when appropriate.

Robustness and security—Consideration should be given to data protection and the long-term impacts of AI, including its effects on jobs, skills, the environment, misuse, data breaches, and data privacy.

IOFM

Narrow AI

Definition: AI designed to perform a single, specific task (e.g., facial recognition, spam filtering, or data extraction). It cannot perform outside its programmed domain.

Efficiency: Automated Invoice Capture. Using Optical Character Recognition (OCR) and Machine Learning to extract invoice numbers, dates, and amounts with 99% accuracy, routing them to the correct GL codes without human intervention.

Fraud Perpetration: Low-level "Fuzzing." Fraudsters use Narrow AI to generate thousands of slightly varied invoice amounts (e.g., $9,998 instead of $10,000) to find the exact threshold where an organization's manual "soft spot" for non-approval exists.

Red Flag Detection: Duplicate Detection. Narrow AI cross-references invoice metadata (date, amount, vendor name, and tax ID) across the entire history of the ledger to flag potential duplicate payments that a human eye would miss.

IOFM

GenAI

Definition: AI models (like LLMs) that create new content—text, images, audio, or code—based on patterns learned from existing data.

Efficiency: Vendor Communications. An AP clerk can use GenAI to summarize 50 conflicting emails from a vendor and draft a polite, professional response that resolves a payment dispute in seconds.

Fraud Perpetration: Persona Cloning. A fraudster uses GenAI to analyze a vendor’s past emails and write a new one that perfectly mimics their tone, vocabulary, and specific "insider" references to request a change in banking details.

Red Flag Detection: Linguistic Anomaly Detection. Advanced tools scan incoming emails for sudden shifts in "linguistic fingerprinting." If a vendor who has been professional for 5 years suddenly starts using informal slang or has a change in sentence structure, the AI flags it for potential Business Email Compromise (BEC).

IOFM

Agentic AI

Definition: This is a step above GenAI. Agentic AI doesn't just write; it acts. It can set its own goals, use external tools (like your ERP or email), and complete multi-step workflows with minimal oversight.

Efficiency: Autonomous Exception Handling. If an invoice is missing a PO number, an Agentic AI will proactively search the shared drive for the contract, email the buyer to confirm, and—once verified—update the ERP and schedule the payment, only involving a human if it hits a roadblock.

Fraud Perpetration: The "Living" Bot. Fraudsters deploy autonomous agents that can respond to your AP team’s questions in real-time. If you ask for a W-9, the bot generates it; if you ask for a callback, it uses a voice-cloned API to "talk" to you.

Red Flag Detection: Behavioral Fingerprinting. These systems monitor the rhythm of your P2P cycle. If an "agent" (human or bot) starts moving through the ERP screens 10x faster than a human or accesses sensitive vendor data at 3:00 AM, the system triggers an immediate lockout.

IOFM

Super AGI

Definition: Super AGI: A theoretical level where AI can understand, learn, and apply knowledge across any intellectual task a human can do.

An AI that surpasses human intelligence across all fields, including scientific creativity and social skills.

In the 2026 AP Context: While we have not reached "Super AI," the industry currently deals with "Super-Specialized Agentic AI"—systems that can out-negotiate humans for early payment discounts or manage complex global tax compliance across 50 jurisdictions simultaneously.

Extension of AP Professionals ☺

IOFM

What camp are you in?

Camp “All In”

Camp “Hell No”

Camp “Skeptic”

IOFM

AI Adoption

Widespread Use: A 2025 Gartner survey of 183 senior finance leaders found that 59% of finance functions are currently using AI. This is a massive jump from only 37% in 2023.

Growing Confidence: 67% of finance leaders are more optimistic about AI in 2025 than they were in 2024. Interestingly, optimism increases with maturity; those further along in adoption are three times more likely to report "high impact" results.

Budget Increases: 88% of CFOs plan to increase their AI budgets in the next 12 months, with many planning double-digit growth (PwC/HighRadius).

Common Applications include:

- Knowledge Management (49%): Helping teams retrieve and leverage complex financial information for decision-making.
- Accounts Payable Automation (37%): Streamlining invoice processing and vendor payments.
- Anomaly & Fraud Detection (34%): Identifying errors and suspicious transactions in real-time.
- FP&A & Forecasting: EY research shows that AI is now used to pre-populate 70–80% of budget values based on historical trends and market drivers, reducing manual data entry by up to 45%.

IOFM

Deepfake Impersonation/Voice Cloning

- Voice Cloning: Attackers use as little as 30 seconds of audio—often harvested from corporate webinars or earnings calls—to clone an executive's voice. They then call AP staff to authorize "urgent" or "confidential" wire transfers.
- Video Conferencing Fraud: As seen in high-profile cases like the $25 million deepfake incident in Hong Kong, attackers can use real-time deepfake video to impersonate multiple colleagues at once during a video call to "confirm" a fraudulent transaction.

- Tech firm’s AP manager received a call from the "CEO" while the CEO was known to be out of the country at a conference. The "CEO" claimed a confidential acquisition was closing in an hour and needed a $4.2M bridge payment. The manager later noted that the CEO even used his specific habit of saying "Right, right" between sentences. The funds were wired before the real CEO even stepped off the stage.

IOFM

Deepfakes – Risk Mitigation

The "Urgency/Secrecy" Combo: Any request that demands bypassing standard controls due to "confidentiality."

Unusual Latency: A slight delay in the "executive's" response during the call (AI processing time).

Channel Shifting: The executive calls on a personal cell or WhatsApp instead of the internal corporate line.

The "Mechanical" Cadence: AI often lacks "prosody"—the natural rhythm and flow of speech. Listen for perfectly even spacing between words or a lack of emotional inflection during a "crisis."

Inability to Interject: If you interrupt the speaker with a non-sequitur (e.g., "Wait, did you ever find your keys?"), an AI bot may experience a 1–2 second "lag" or simply continue its script.

Background Silence: Professional environments usually have "room tone." A deepfake often sounds unnervingly clean or uses a looped, static background noise (like a generic coffee shop track).

Multi-Channel Callbacks: Establish a rule that any verbal payment instruction must be confirmed by calling the executive back on a pre-verified internal number.

Safe-Words/Challenge Phrases: Implement "Challenge-Response" phrases for high-value transactions that are never shared over email or text.

Liveness Testing: If on video, ask the person to turn their head 90 degrees. Many real-time deepfake overlays still "break" or blur when they lose the straight-on facial profile.

IOFM

Hyper-personalized BEC Schemes

- Natural Language Processing (NLP): Attackers use LLMs to analyze an executive's writing style, tone, and professional vocabulary, making fraudulent emails virtually indistinguishable from legitimate internal communication.
- Contextual Social Engineering: AI can scan public data (LinkedIn, news releases) to craft messages that reference real projects, upcoming conferences, or recent vendor transitions, creating a false sense of legitimacy.
- form of **Psychological Engineering** powered by Large Language Models (LLMs) that have been trained on stolen data to act as a "digital twin" of a trusted colleague or vendor.

- A construction company received a "bank change" request from their primary steel supplier. The email referenced a specific delay in a Seattle-based project—information only found in a private project management portal. Because the email was grammatically perfect and referenced the "Project 402-B weather delay," the AP clerk updated the ACH details. The $800,000 payment for the next invoice disappeared.

IOFM

Hyper-personalized BEC – Risk Mitigation

The "Trust-Anchor" Detail: Over-explaining why the bank is changing by using overly specific internal project details to "prove" identity.

Header Discrepancies: While the body is perfect, the "Reply-To" address may have a subtle, invisible character (homograph attack) that LLMs can’t always hide from a technical scan.

The "Too-Specific" Proof: Attackers now use AI to include too many real details to prove they are legitimate. If an email includes your project ID, your boss's name, and the exact date of your last meeting, it might be a "trust-anchor" designed to lower your guard.

Metadata Mismatches: In 2026, check the Reply-To address carefully. AI can mimic the "From" name perfectly, but the underlying routing might still point to a lookalike domain (e.g., vendor.co instead of vendor.com).

Shift in Urgency: AI-generated emails often pivot quickly from a helpful, professional tone to extreme pressure if the target hesitates.

AI-Powered Email Filtering: Use tools that analyze stylometry (writing style). These tools flag when a "long-time contact" suddenly changes their sentence structure or vocabulary.

Out-of-Band (OOB) Authentication: Mandatory verbal confirmation with a known contact at the vendor's office for any change to banking or payment instructions.

DMARC/Strict SPF: Ensure your organization enforces strict email authentication protocols to prevent domain spoofing.

IOFM

Synthetic Vendor Onboarding

- Synthetic Identities: Fraudsters use AI to generate realistic business licenses, tax documents, and professional websites to pass initial vendor vetting and enter the AP system as a legitimate payee.
- Automated Invoice Manipulation: AI tools can scan legitimate invoices and automatically alter account numbers or payment terms while maintaining the exact font, formatting, and logos of the original document.

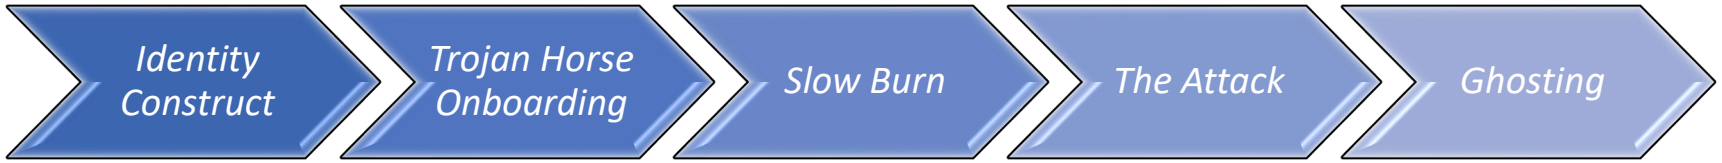

- Example: An enterprise organization onboarded a vendor after a seamless procurement process. The VP of Sales participated in a 15-minute Zoom interview, appearing as a professional woman in a realistic office setting. In reality, it was a high-quality live deepfake overlay. The vendor billed the company for "logistics consulting" for six months, totaling $1.2M, before a physical audit revealed the headquarters address was a vacant lot.

IOFM

Synthetic Vendor Onboarding – Risk Mitigation

Digital "Glitching": During video calls, the deepfake may "flicker" near the edges of the face or when the person turns their head sharply.

The "Newness" Factor: A vendor with a highly polished presence but whose domain and LinkedIn profiles were all created within the last 90 days.

Inconsistent Metadata: AI-generated "tax forms" often have perfect text but lack the underlying digital "noise" or metadata found in scanned government documents.

The "Shadow" Digital Footprint: Perform a "Digital Forensics" check. Vendor's website have a low "Domain Age" (e.g., registered 30 days ago)? Does their LinkedIn show hundreds of employees who all have AI-generated, perfectly symmetrical headshots?

Verification Resistance: A synthetic vendor will often "have technical difficulties" if asked to perform a specific, non-scripted action during a video call (like holding up a piece of paper with today's date).

Circular References: They provide references that, when checked, turn out to be other synthetic companies or "ghost" websites.

Verified Vendor Directory: Maintain a "Golden Record" of vendors. New entries must undergo a 72-hour "cooling off" period before any payments can be released.

Physical Address Verification: Cross-reference the "headquarters" on Google Maps. If a multi-million dollar "logistics firm" is registered to a residential house or a vacant lot, it's a phantom.

Third-Party Risk Management (TPRM): Use services that verify the business's physical existence and credit history through traditional, non-digital channels.

IOFM

Key Facts

The "Agentic AI" Shift: 2026 is being called the year of Agentic AI—AI that doesn't just analyze data but acts on it (e.g., emailing a vendor to resolve a price discrepancy). 82% of midsize companies have already begun or plan to implement agentic AI by the end of 2026.

The Productivity Gap: Companies effectively integrating AI are seeing productivity gains of up to 40% (EY). However, Deloitte notes a "persistent investment imbalance": 93% of budgets go toward the technology itself, while only 7% is spent on training the people who use it.

The ROI Reality Check: Despite high adoption, only 21% of finance leaders report a "clear, measurable ROI" so far. Most organizations (72%) report still being in the "break-even" or investment phase.

Talent Transformation: Finance managers are pivoting from hiring for traditional accounting skills to prioritizing AI fluency and data analysis. 64% of finance leaders now prioritize these technical skills over traditional CPA-style qualifications

70-80% Time Savings: Reported by firms transitioning from manual to AI-driven workflows.

Reduction in Error Rates: Manual entry error rates reduced to less than 0.8%.

Strategic Shift: AP staff are moving from "data entry clerks" to "cash flow analysts" and "vendor relationship managers.“

Source: Gartner 2025 AI in Finance Survey (Released Nov 2025)Deloitte Finance Trends 2026 Report (Released Oct 2025)PwC 2025 AI Agent Survey (Released June 2025)Citizens Bank 2026 AI Trends in Financial Management

IOFM

Executive Insights on AI Strategy, Risks, and Readiness

AI adoption is growing, but readiness is uneven.

- Organizations face significant gaps in talent, technology, and governance readiness.
- Early adopters are already realizing measurable strategic gains — but also recognize rising and fast-evolving AI risks.
- AI adoption patterns differ sharply across regions and industries.
- 73% report that AI is already delivering strategic advantage.
- 54% fear competitors may outpace them through AI.

AI risks are escalating quickly.

- 69% classify AI as a Top 10 or major risk.
- 65% say AI risks receive executive or board-level attention (vs. 30% overall).
- 46% of organizations now classify AI as a Top 10 or major risk (rising to 69% among AI-Transformed entities).
- 26% say AI risks are evolving rapidly — jumping to 60% for highly AI-mature organizations.
- Industries such as Financial Services report the fastest-changing risk landscapes and strongest board engagement.

IOFM

Executive Insights on AI Strategy, Risks, and Readiness

For most organizations, readiness is the roadblock.

- Only 24–27% report adequate AI-skilled talent, IT system capacity, or regulatory preparedness.
- Smaller organizations are even less prepared — fewer than 1 in 5 have the necessary capabilities.
- In contrast, AI-Transformed organizations are nearly twice as prepared across talent, IT, and regulatory readiness metrics.

AI impact varies significantly by geography and industry.

- Emerging markets — including South Africa, Central & South Asia, and East/Southeast Asia — show the highest levels of strategic AI adoption (36–42%).
- North America and Europe lag behind at 18–22%, reflecting more cautious approaches.
- Mining, Professional & Business Services, and Transportation exhibit strong momentum driven by automation and analytics.
- Financial Services report elevated competitive concern (33% fear competitors will outpace them).
- Sectors like Construction and Wholesale/Retail remain slower adopters due to operational fragmentation and legacy systems.

IOFM

Road Ahead

AI is here to stay

AI will be smarter than humans

AI will have intended outcomes

AI will have unintended outcomes

AI will be used for good

AI will be used for bad

IOFM

Final Thoughts

1. Learn AI
2. Use as a companion and not a future replacement
3. Capitalize on your free time to add value
4. Keep a human in the loop
5. Use AI ethically
6. Provide positive reinforcement
7. Share best practices
8. Stay relevant

IOFM

Do you need NASBA CPE credits?

- Navigate to website: **iofm.cnf.io**
  or scan the QR code →
- Check-in and check-out of your sessions to track your attendance for NASBA CPEs
- Certified with IOFM? No need to check-in and out of sessions. Self-report CEUs on IOFM.com instead after the event!

IOFM

Please tell us what you think!

- Please scan this QR code using your mobile to access a short feedback survey →
- Also accessible via the mobile app

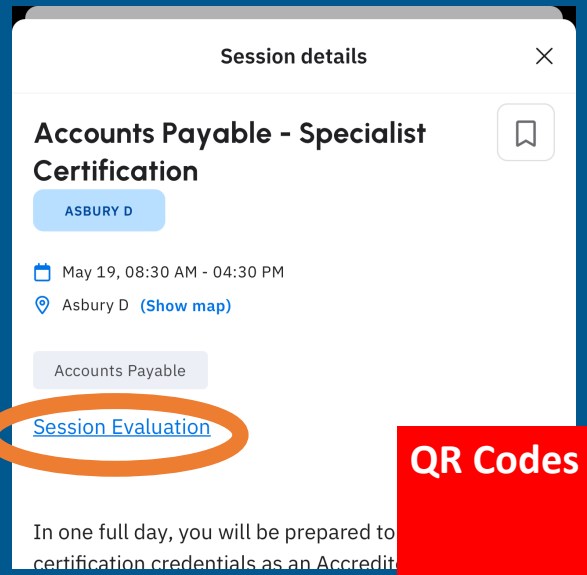

QR Codes will be shared closer to event