CyberSecurity, AI & Enterprise Risk Management

December 2023

1

Agenda

Introduction

CyberSecurity: State of the Union
- Threats, Trends, Motivations & Actors
- Key Lessons: Resilience + Best Practices
- SEC Rules on Cyber Disclosure & Oversight

Artificial Intelligence
- What is Generative AI?
- Risks, Limitations & Ethics of Adoption
- Governance (policies, procedures, training, responsible adoption)

Enterprise Risk Management

Open Dialog + Q&A

2

Grant Thornton | SpencerFane

CyberSecurity

The State of the Union

December 2023

3

Cyber State of the Union…

What are the bad actors after?

- Customer PII (personally identifiable information)
- PHI (personal health information)
- Company and Trade Secrets / IP / confidential information
- Investment Strategies, Mergers & Acquisitions
- Embarrassing personal information / defacement of public sites
- Generally, things you need to generate revenue and run your business

2

Grant Thornton | SpencerFane

4

Industry Overview…

Cyber crime affects every industry in every part of the globe.

- Financial and Insurance
- Healthcare
- Accommodation and Food Services
- Educational Services
- Energy
- Travel and Logistics
- Manufacturing
- Public Administration

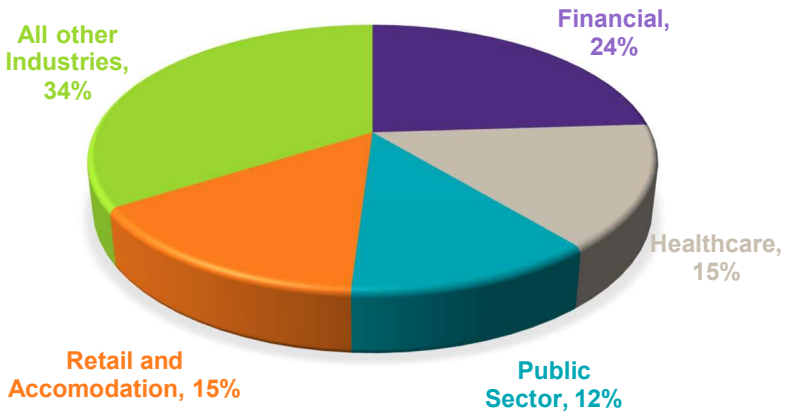

* *Source:* 2022 Verizon Data Breach Investigations Report

Grant Thornton | SpencerFane

 5

5

Cyber State of the Union…

Who are these threat actors?

Examples	Motivation	TTPs*	Sample Targets
Hacktivists **(Outsiders)**	Political, economical, and social agenda	Hacking, Phishing, DDoS, Ransomware	Politicians, celebrities, government, corporations, not-for-profit organizations
Cyber Criminals **(Outsiders or Insiders)**	Financial gains	Hacking, Phishing, Ransomware	Banks, businesses, healthcare, utilities, law firms, and retailers, individuals
Cyber Espionage **(Nation States)**	Political and economical agenda	Advanced phishing, DDoS, malware, etc.	Government, defense companies, or organizations with sensitive data
Insiders **(Organizational Staff)**	Economic or political	Leverage authorized access to steal data or commit sabotage	Sensitive data for resale or embarrassment to the organization

* Tactics, Techniques & Procedures

Grant Thornton | SpencerFane

 6

6

Cyber Trends: Evolved Criminal Markets…

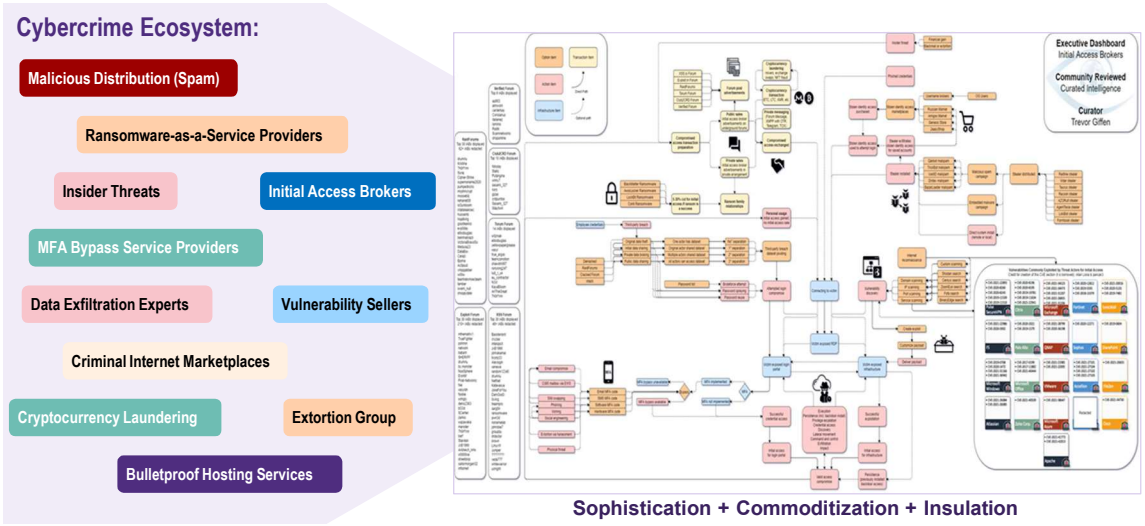

Sophistication + Commoditization + Insulation

Grant Thornton | SpencerFane

Cyber State of the Union…

What makes cyber defense so challenging today?

- Lack of business knowledge around retained data
- Inconsistent data-retention/-inventory over time by business
- Increased attack surfaces (pandemic / WFH regimes)
- Poor security hygiene (passwords, MFA, admin privileges)
- Lack of detailed IT asset inventory
- Misalignment of IT Strategy to business strategy
- Poor data/network segmentation
- Lack of Incident Response protocols (and **practice**!)

Grant Thornton | SpencerFane

BEC & invoice- / wire-transfer attacks are costly

- Business Email Compromise (BEC) attacks are often the simplest and costliest.
- Results from 2 failures:
  - Email account “hack”, because no 2FA
  - Lack of “internal controls” in company
- What internal controls are needed?
- Key issue for supply chain risk.
- Cyber insurance understanding and interrelationship is critical.

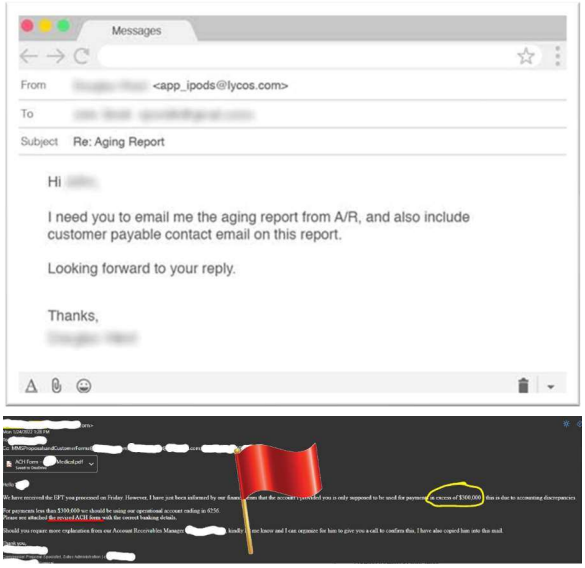

Grant Thornton | SpencerFane

Call-back phishing attacks: the latest trend…

- Threat Actor sends email prompting a response, such as, “your free trial period to [PICK AN EXPENSIVE SERVICE] is almost over and when it expires you will be automatically placed on the premium plan for $89.00 per month.”
- Call this number to un-subscribe: TA’s number
- TA then uses social engineering to ultimately get them to download a malicious file.

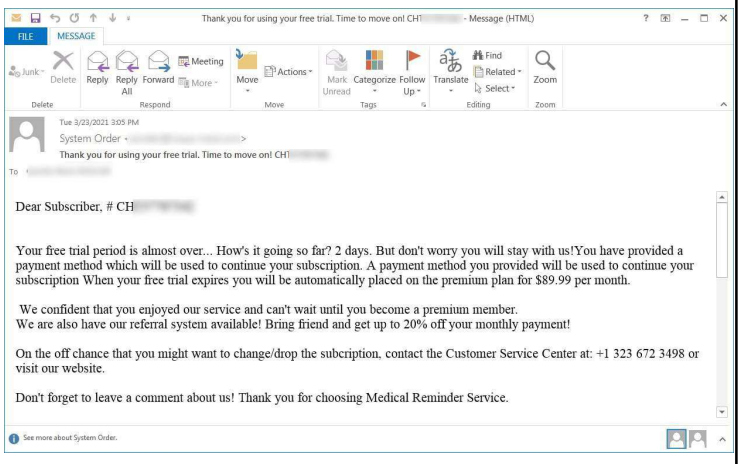

https://www.bleepingcomputer.com/news/security/callback-phishing-attacks-evolve-their-social-engineering-tactics/

Grant Thornton | SpencerFane

Share File / OneDrive phishing: a recent trend…

- Threat Actor spoofs a known good contact and sends email disguised document to share a customary business document.
- Recipient will then be prompted for access credentials, and…will comply.
- TA gains access to email account:
  - Surveils account for sensitive / valuable information, then downloads some or all of info = data breach.
  - Account used to send out thousands of the same phishing email to contacts (and others).

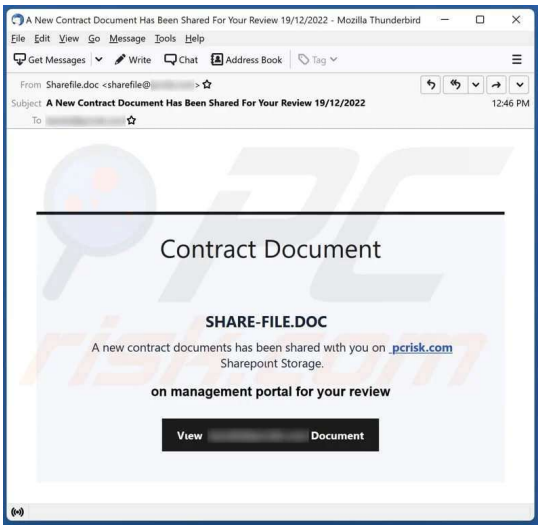

https://www.pcrisk.com/removal-guides/18458-onedrive-email-scam

Grant Thornton | SpencerFane

Most Common Causes & Solutions

RDP Access	• This is random – scanning web for Internet facing RDP access • Virtual Private Network (VPN) with Multifactor Authentication (MFA)
Phishing	• Email phishing tool • Workforce training and simulated phishing
Unpatched / Outdated Software	• Install patches timely • No unsupported software
Passwords	• Multifactor Authentication (MFA) • Longer passphrases
Backups, Backups, Backups!	• 3-2-1 Backup Process • Something comparable – you may end up with only your offline backup

Grant Thornton | SpencerFane

Key Lessons Learned

The Concept & Elements of Resilience

December 2023

13

What makes an organization resilient?

- Cyber Insurance
- Integrated Niche Expertise
- Organizational Readiness

14

Resilience via Organizational Readiness…

- Ensure adequate insurance coverage
- Perform an independent Risk Assessment (including vulnerabilities)
- Conduct data inventories to identify high-risk data
- Tailor tools and monitoring to focus on key risks
- Engage with niche specialists (early and often)
- Develop “muscle memory” through realistic IR exercises

PRACTICE FOR THE BAD DAY…

Grant Thornton | SpencerFane

 15

15

Ransomware Timeline…

Preparation is the key to being able to do all things necessary for a successful response.

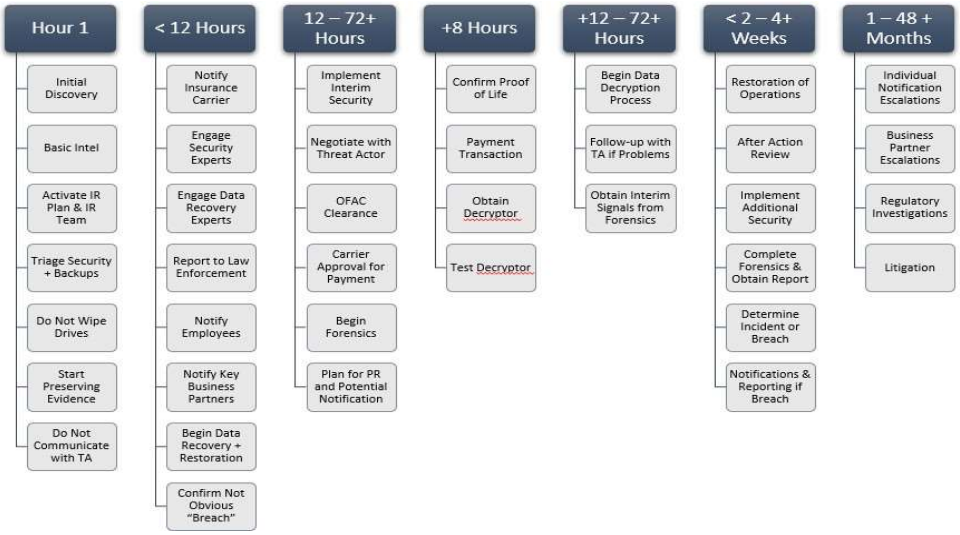

Grant Thornton | SpencerFane

 16

16

Resilience begins with Readiness…

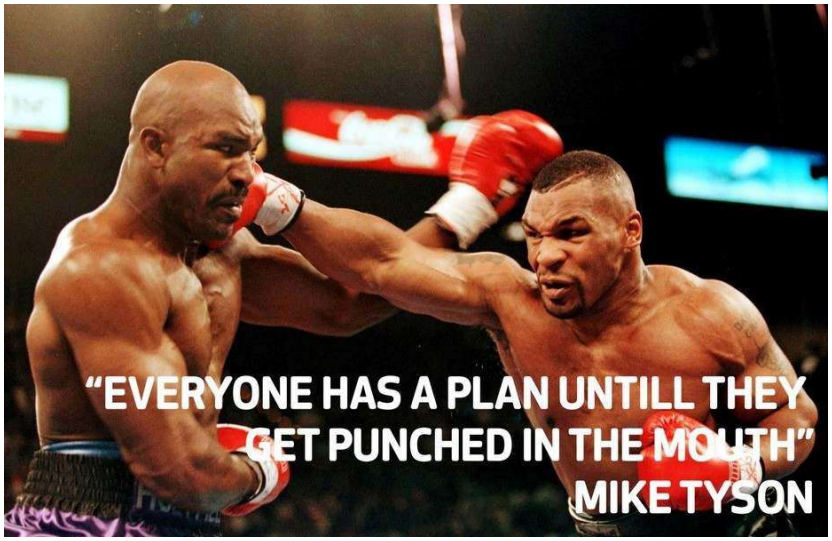

Grant Thornton | SpencerFane

Grant Thornton | SpencerFane

SEC Cyber Disclosure Rules

December 2023

New SEC Disclosure Rules… What's Different?

Final rules* implement same basic structure as initially proposed (2022) – namely:

- Disclosures of cyber risk management, strategy, and governance in annual reports; and
- Reporting required for ***material*** cybersecurity incidents on Form 8-K or Form 6-K.

SEC noted that registrant disclosures of material cybersecurity incidents and cybersecurity risk management and governance have improved since the 2011 and 2018 guidance, however the Commission "remain[s] persuaded that…under-disclosure regarding cybersecurity persists despite…prior guidance" and "investors need more timely and consistent cybersecurity disclosure to make informed investment decisions."

Grant Thornton | Spencer Fane

* Please see Appendix B for more details.

Annual Disclosures…

- The Final Rule amends Form 10-K (via new Item 106 of Regulation S-K) to require registrants to disclose information about cyber risk management, strategy, and governance.
- **Risk Management & Strategy**:
  - Company processes for the assessment, identification, and management of material risks from cybersecurity threats;
  - Commentary on whether any of these risks (including prior cyber incidents) have materially affected (or are reasonably likely to materially affect) business strategy, results of operations, or financial condition – and, if so, how;
  - Description of whether and how such processes have been integrated into the registrant's overall risk management system or processes;
  - Registrant's use of assessors, consultants, auditors, or other third-parties in connection with such processes; and
  - Commentary on whether the registrant has processes to oversee and identify material risks from cybersecurity threats associated with its use of any third-party service provider.

Grant Thornton | SpencerFane

Annual Disclosures (Board Governance)…

- **Board Oversight**
  - Final rules require registrants to describe the of risks from cybersecurity threats – including, *if applicable*:
    - identification of any board committee or subcommittee responsible for the oversight of such risks; and
    - a description of the processes by which the board or such committee is informed about such risks.
  - Notably, SEC did *NOT* adopt the proposed requirement to disclose board cybersecurity expertise.
  - Likewise, the new governance disclosure provisions do *NOT* require disclosure of the frequency of management and board discussions regarding cybersecurity risks, which had been contemplated by the proposed rules.

Annual Disclosures (Management)…

- **Management Oversight**
  - Requirement to describe management's role in assessing and managing material risks from cybersecurity threats, including relevant expertise and communication with the board of directors.
  - Item 106(b) of Regulation S-K includes (non-exclusive) list of disclosure items:
    - management positions responsible for assessing/managing cybersecurity risks;
    - relevant expertise of such persons;
    - processes by which such persons or committees are informed about and monitor the prevention, detection, mitigation, and remediation of cybersecurity incidents; and
    - whether such persons or committees report information about such risks to the board or a committee or subcommittee of the board.

Incident Disclosures…

Final Rule adds Item 1.05 to Form 8-K, requiring disclosure of the nature, scope, and timing of a material cybersecurity incident – as well as the material impact (or reasonably likely material impact) of the incident.

- **Key Insights (for disclosure and materiality committees):**
  1. Reporting required w/in 4 business days of materiality determination (via 8-K);
  2. Limited delayed disclosure exceptions (national security + public safety risks);
  3. 8-K disclosure must focus on the impacts of the incident;
  4. Updated incident disclosure is required (as issues become known/clarified); and
  5. Importantly, the final rules did *NOT* adopt the proposed requirement to disclose a series of previously undisclosed, individually immaterial incidents (unless such incidents comprise a “series of related unauthorized occurrences”)*.

* This presupposes the tracking of minor cyber incidents for potential aggregation.

Grant Thornton | SpencerFane

What Questions should the company be asking?

Current State

- What is the company’s process for assessing and managing material risks from cyber threats?
- How does the company coordinate cyber incident response with the broader business?
- Which third-parties are used to help manage cyber risk in our environment?
- Which third-party vendors/partners represent potentially material exposures?

Strategic Focus

- How will management apprise the Board on emerging and persistent cyber risks?
- What has management done to ensure that we have the right expertise to manage cyber risk?
- How does management ensure that organization has resilience against cyber-attacks?
- What resources are needed most by management to manage cyber risk?
- How does the business measure and monitor third-party cyber risk?
- Is our current insurance coverage sufficient? How does management know as much?
- Does the organization have the required governance policies and controls in place? How will we enforce new policies and procedures?

Grant Thornton | SpencerFane

Grant Thornton | SpencerFane

Artificial Intelligence

Technology, Risks & Rewards

December 2023

25

Generative AI

What is it?

Generative AI is a branch of artificial intelligence made up of a set of algorithms focused on creating new content such as images, music, or text using a variety of data science and machine learning techniques.

- It aims to generate outputs that resembles the patterns and characteristics of the examples it was trained on.
- Generative AI large language models (LLMs), learn patterns and structures from large datasets and have been trained on trillions of language examples
- Once trained, these models can generate seemingly new content by sampling from the learned patterns.

Generative AI can support a variety of functions three common categories include:

- Efficiency Improvement: acceleration of manual or repetitive tasks
- Personalizing Experiences: content and information tailored to the individual audience by considering their interests and patterns of behavior
- Content and Idea Generation: Creating new and unique outputs across a range of formats

Grant Thornton | SpencerFane

Available OpenAI Models

There are multiple Generative AI Models, each with their own varied capabilities and costs. Each model becomes progressively more capable of performing complex tasks at the expense of increased cost and slower processing speeds.

1. Ada

- Fastest of the GPT-3 models
- Least costly of the GPT-3 models
- Least capable of the GPT-3 models; best for less nuanced tasks such as parsing and reformatting text

2. GPT 3.5-Turbo

- Most capable and cost-effective model of the 3.5 family
- Optimized for efficient and high-quality text generation tasks
- Balances performance and efficiency

3. GPT-4

- Most advanced GPT model, able to solve higher difficulty questions with improved accuracy
- Capable of enhanced function calling and contains a larger context window
- Requires greater computational resource needs and comes at a higher cost than other models

4. GPT-4V

- Multimodal model of GPT-4, allows visual question answering
- Able to analyze image inputs and answer questions about them

Grant Thornton | SpencerFane

Generative AI: helping knowledge professionals

- Prediction engines – Generative AI models/apps/systems have been trained on massive amounts of data and published works to then predict the next word or pixel to produce a creation (e.g., text, images, music, video) without any direct human intervention.
- New embrace of generative AI to enable change agents – Knowledge professionals drive the organization's value creation through data analytics, deriving important business insights and storytelling.

How? human change agents & generative AI partner:

"Organize This"	*"Calculate This"*	*"What Happens When"*	*"Human, some suggestions"*
Knowledge gathering Compare	Analysis/Code Production	Scenario Creation	Scrutinize the Analysis
• Generate statements/reports; customize to user needs • Compare/Perform cursory review of budget to actuals analysis • Perform cursory reviews of statements and identify anomalies • Draft analysis responses to different audiences	• Perform complex account reconciliations • Synthetic data generation. Example: Learn patterns and relationships from real financial data to create synthetic datasets that closely resemble the original data while preserving data privacy.	• Simulate market conditions, economic environments, and events to better understand the potential impacts on portfolio performance. • Risk tolerance, expected returns, investment horizons	Assist to develop more profound understanding of operational aspects. When these aspects can be described as sequences of logical statements, then generative AI can assist by: • Check inputs for logical correctness • Suggest improvements • Look for missing premises behind conclusions • Validate logical connections

Grant Thornton | SpencerFane

Grant Thornton | SpencerFane

Risk, Limitations & Ethics

December 2023

29

AI Risks – High-priority characteristics

Enterprise Risk Managers need to be asking how these risks are being identified, assessed and managed in AI systems.

1. Technical design characteristics

Factors that are under the direct control of AI system designers and developers

Accuracy
- Degree in which model is correctly capturing a relationship that exists with training data
- Determine threshold that corresponds with acceptable risk

Reliability
- Whether model consistently generates same results
- Factor in determining threshold of acceptable risk

Robustness
- Whether model has minimum sensitivity to variations in uncontrollable factors
- Part of sensitivity analysis in the AI risk management process

Resiliency/Security
- Whether model can withstand adversarial attacks or unexpected changes in its environment/use
- Exfiltration of models, training data, or intellectual property through AI system endpoints

2. Socio-technical characteristics

How AI systems are used and perceived in individual, group, and societal contexts

Explainability
- Refers to a representation of the mechanisms underlying an algorithm's operation
- Risks from lack of fidelity or consistency in explanation methodologies; if humans incorrectly infer a model's operation; model is not operating as expected
- Risk can be managed by training and descriptions of how models work attune to users' skill levels/roles

Interpretability
- Refers to the meaning of its output in the context of its designed functional purpose
- Risks can be addressed by communicating the interpretation intended by model designers, although this remains an open area of research

Privacy
- Safeguard human values e.g., autonomy and dignity
- Assess how data processing creates privacy problems

Safety
- Practical approaches relate to rigorous simulation, in-domain testing, real-time monitoring, and the ability to quickly shut down or modify misbehaving systems

Managing Bias
- Three major categories of bias in AI: systemic, computational, and human
- Tightly associated with transparency and fairness concepts

3. Guiding principles contributing to trustworthiness

Broader norms/values that indicate societal priorities

Fairness
- Increasingly related to the existence of a harmful system
- Absence of harmful bias is a necessary condition for fairness

Accountability
- Expectations for the responsible party in the event that a risky outcome is realized
- Grounding organizational practices and governing structures for harm reduction, like risk management, can help lead to more accountable systems

Transparency
- Reflects the extent to which information is available to a user when interfacing with an AI system
- Its scope spans from design decisions and training data to model training, the structure of the model, its intended use case, how and when deployment decisions were made and by whom
- Is necessary for actionable redress related to incorrect and adverse AI system outputs

Grant Thornton | SpencerFane

30

Deeper dive – emerging cyber threats

Adversarial Attacks	• Inputs that evade detection by the AI system and allow an attacker to achieve a malicious goal, such as generating false results. This could result in outputs that are unknowingly incorrect or unexpected and could further results in divulging sensitive information or performing unauthorized actions.
Model Poisoning	• Model poisoning attacks target AI models in a development or testing environment. Attackers introduce malicious data into the training data to influence the output – sometimes creating a significant deviation of behavior from the AI model. For example, after a successful model poison attack, an AI model may produce incorrect or biased predictions, leading to inaccurate or unfair decision making.
Data Leakage & Breaches	• Data breaches pose a significant cybersecurity risk for AI platforms that store and process vast amounts of confidential or sensitive data like personally identifiable information. • Users can inadvertently feed sensitive data through browser extensions, APIs or directly to the AI system. This data can then become part of the large data sets used to train AI models and presented in the form of results.
Data Retention & Deletion	• Some AI solutions store data for extended periods so that they can continue referencing, analyzing and comparing it as part of informing their machine learning, predictive and other capabilities. Increases the risk of unauthorized access or misuse. • The context and complexity of AI solutions can also make it challenging to ensure that data is deleted when it is no longer needed or when individuals exercise their rights to request deletion.
Third Party Entities / Tools	• Risks: Risk metrics/methodologies of the developing organization may not align with organizations deploying/operating the system. Organization developing AI may not be transparent about risk metrics/methodologies. • Governance: Policies and procedures are in place that address AI risks associated with third-party entities, including risks of infringement of a third-party's intellectual property or other rights. Contingency processes are in place to handle failures or incidents in third-party data or AI systems deemed to be high-risk • Management: AI risks and benefits from third-party resources are regularly monitored, and risk controls are applied and documented. Pre-trained models which are used for development are monitored as part of AI system regular monitoring and maintenance.

Grant Thornton | SpencerFane

Grant Thornton | SpencerFane

Governance

December 2023

AI regulatory requirements & industry standards

It is important to actively monitor the regulatory trends, industry standards, and leading practices in the areas of AI risk management and governance. The following highlights representative sources that can be leveraged to develop an AI risk management framework.

Regulations

EU AI Act
Classifying AI applications by risks
- unacceptable risk – prohibited in EU
- high risk – a series of compliance requirements
- low or minimal risk – unregulated

Further obligations for product manufacture, importer, or distributor of high-risk AI applications

Industry Standards

NIST AI RMF 1.0
AI RMF Core
- Govern
- Map
- Measure
- Manage

Risk management should be continuous, timely, and performed throughout the AI system lifecycle dimensions. AI RMF Core functions are broken into categories and subcategories

Leading Practices

Microsoft Responsible AI Standard, v2
Six AI principals
- Accountability
- Transparency
- Fairness
- Reliability & Safety
- Privacy & Security
- Inclusiveness

Further broken down by goals, requirements, tools, and practices

3rd Party Assurance

- AT-C Section 205 - Assertion-Based Examination Engagements are currently being used for AI specific assessments.
- SOC / ISO – Elements related to the AI are often incorporated into the SOC and ISO assessments such as data privacy, security and processing integrity.
- Early days – Very few third party assessments have been conducted to date.

Grant Thornton | SpencerFane

"Complex algorithms" is not a legal defense

Since 2016, various countries have passed 123 AI-related bills.

Comply with the law

1 Company's ability to explain their governance, oversight and accountability operating practices

2 Company's ability to explain how algorithm-driven decisions were reached

3 Company's ability to audit the algorithm themselves

Four General Data Protection Regulation (GDPR) principles are especially relevant:

(1) Accountability - AI-generated decisions with significant material influence over individuals are subject to even more stringent accountability requirements

(2) Fairness - Organizations that use AI to analyze personal information must evaluate its likely impact on individuals, and continuously reassess their findings as systems evolve

(3) Data minimization and security - AI systems must not process any more personal data than is necessary

(4) Transparency - Individuals' knowledge and well-informed consent that an AI system will handle their personal information

EU AI Act – proposed law to regulate the development and use of AI systems (early 2024?)

- Stringent requirements for 'high-risk' AI systems, including those used in human resources, banking, education.

U.S. - CFPB is focusing on protective consumers from algorithmic discrimination

- 2022: CFPB prioritized targeting unfair discrimination even if fair lending laws don't apply, citing prohibitions against unfair, deceptive and abusive practices under the Consumer Financial Protection Act (CFPA).
- June 2023: Chatbots can incur risk of noncompliance with federal consumer financial protection laws.
- July 2023: Start of an informal dialogue between the CFPB and the European Commission on a range of critical financial consumer protection issues

In 2022, 15 U.S. states and localities proposed or passed legislation concerning AI

- New York City introduced one of the first AI laws in the U.S., effective from July 5, 2023, which aims to prevent AI bias in the employment process (NYC Local Law 144).

Grant Thornton | SpencerFane

AI risk management framework

Taking into consideration regulatory requirements, industry standards, and leading practices, the following is a snapshot of GT's AI risk management framework.

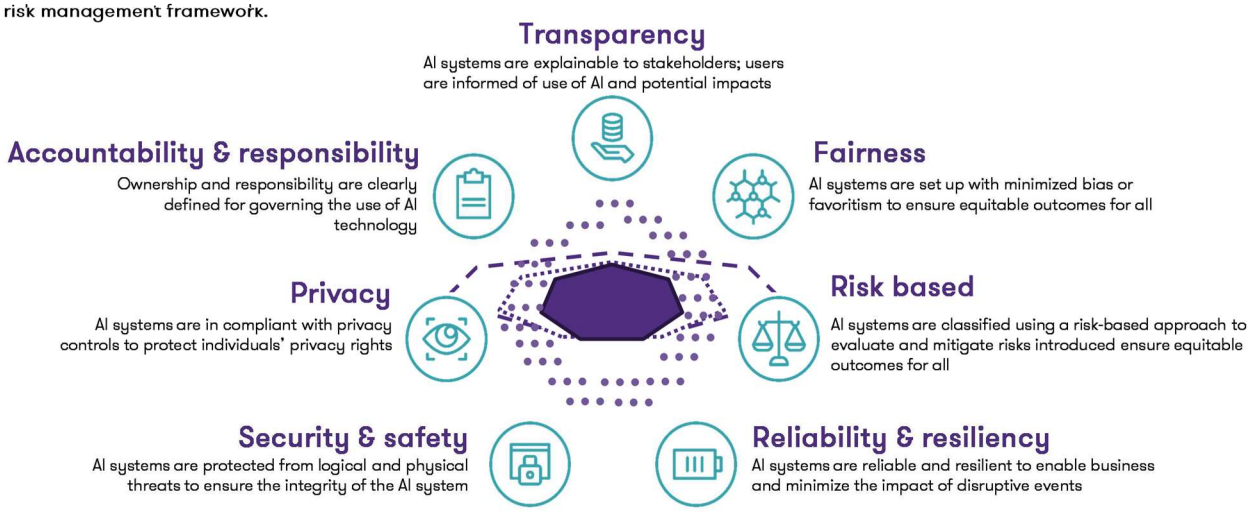

Grant Thornton | SpencerFane

 35

35

GT's AI risk management framework

Taking into consideration regulatory requirements, industry standards, and leading practice, the following is a snapshot of GT's AI risk management framework.

AI Governance & Risk Management

	Accountability & Responsibility	Transparency	Fairness	Risk Based	Reliability & Resiliency	Security & Safety	Privacy
Principles	Ownership and responsibility are clearly defined for governing the use of AI technology	AI systems are explainable to stakeholders; user are informed of use of AI and potential impacts	AI systems are set up with minimized bias or favoritism to ensure equitable outcomes for all	AI systems are classified using a risk-based approach to evaluate and mitigate risks introduced	AI systems are reliable and resilient to enable business and minimize the impact of disruptive events	AI systems are protected from logical and physical threats to ensure the integrity of the AI system	AI systems are in compliant with privacy controls to protect individuals' privacy rights
Capabilities	• Oversight body • Roles & responsibilities • Strategy, policy, standards, and procedures • Human review & control • Stakeholders' inclusiveness • Cost of ownership	• Fit for purpose • Explainability • Disclosure of usage and impact • Communication with stakeholders	• Quality of service • Reducing impact on marginalized groups • Allocation of adequate resources • Data bias detection, test, and mitigation	• AI system risk classification • AI risk/impact assessment • Risk mitigation ➢ Third party ➢ Intellectual property ➢ Regulatory Compliance ➢ Risk tracking & reporting	• Evaluation of potential failure, error rate, and mitigation plan • Operational resiliency / System health check • Ongoing monitoring / alerting • Optimization	• Security controls • Data protection • Access controls • SSDLC • Software supply chain risk • Vulnerability management • Change management	• Privacy regulatory compliance • Privacy impact assessment

Grant Thornton | SpencerFane

 36

36

Shawn Tuma
Spencer Fane LLP
Partner & Co-Chair,
Cybersecurity & Data Privacy

stuma@spencerfane.com

Johnny Lee
Grant Thornton LLP
Principal & National Practice Leader,
Forensic Technology Services

j.lee@us.gt.com

www.grantthornton.com

twitter.com/GrantThorntonUS

Thank you!