COMTECH

AI Safeguards

Ensuring Trust and Security in Finance

Rafe Martin, CCRP, CRO

MCDONALDS
260
CHICKEN NUGGETS

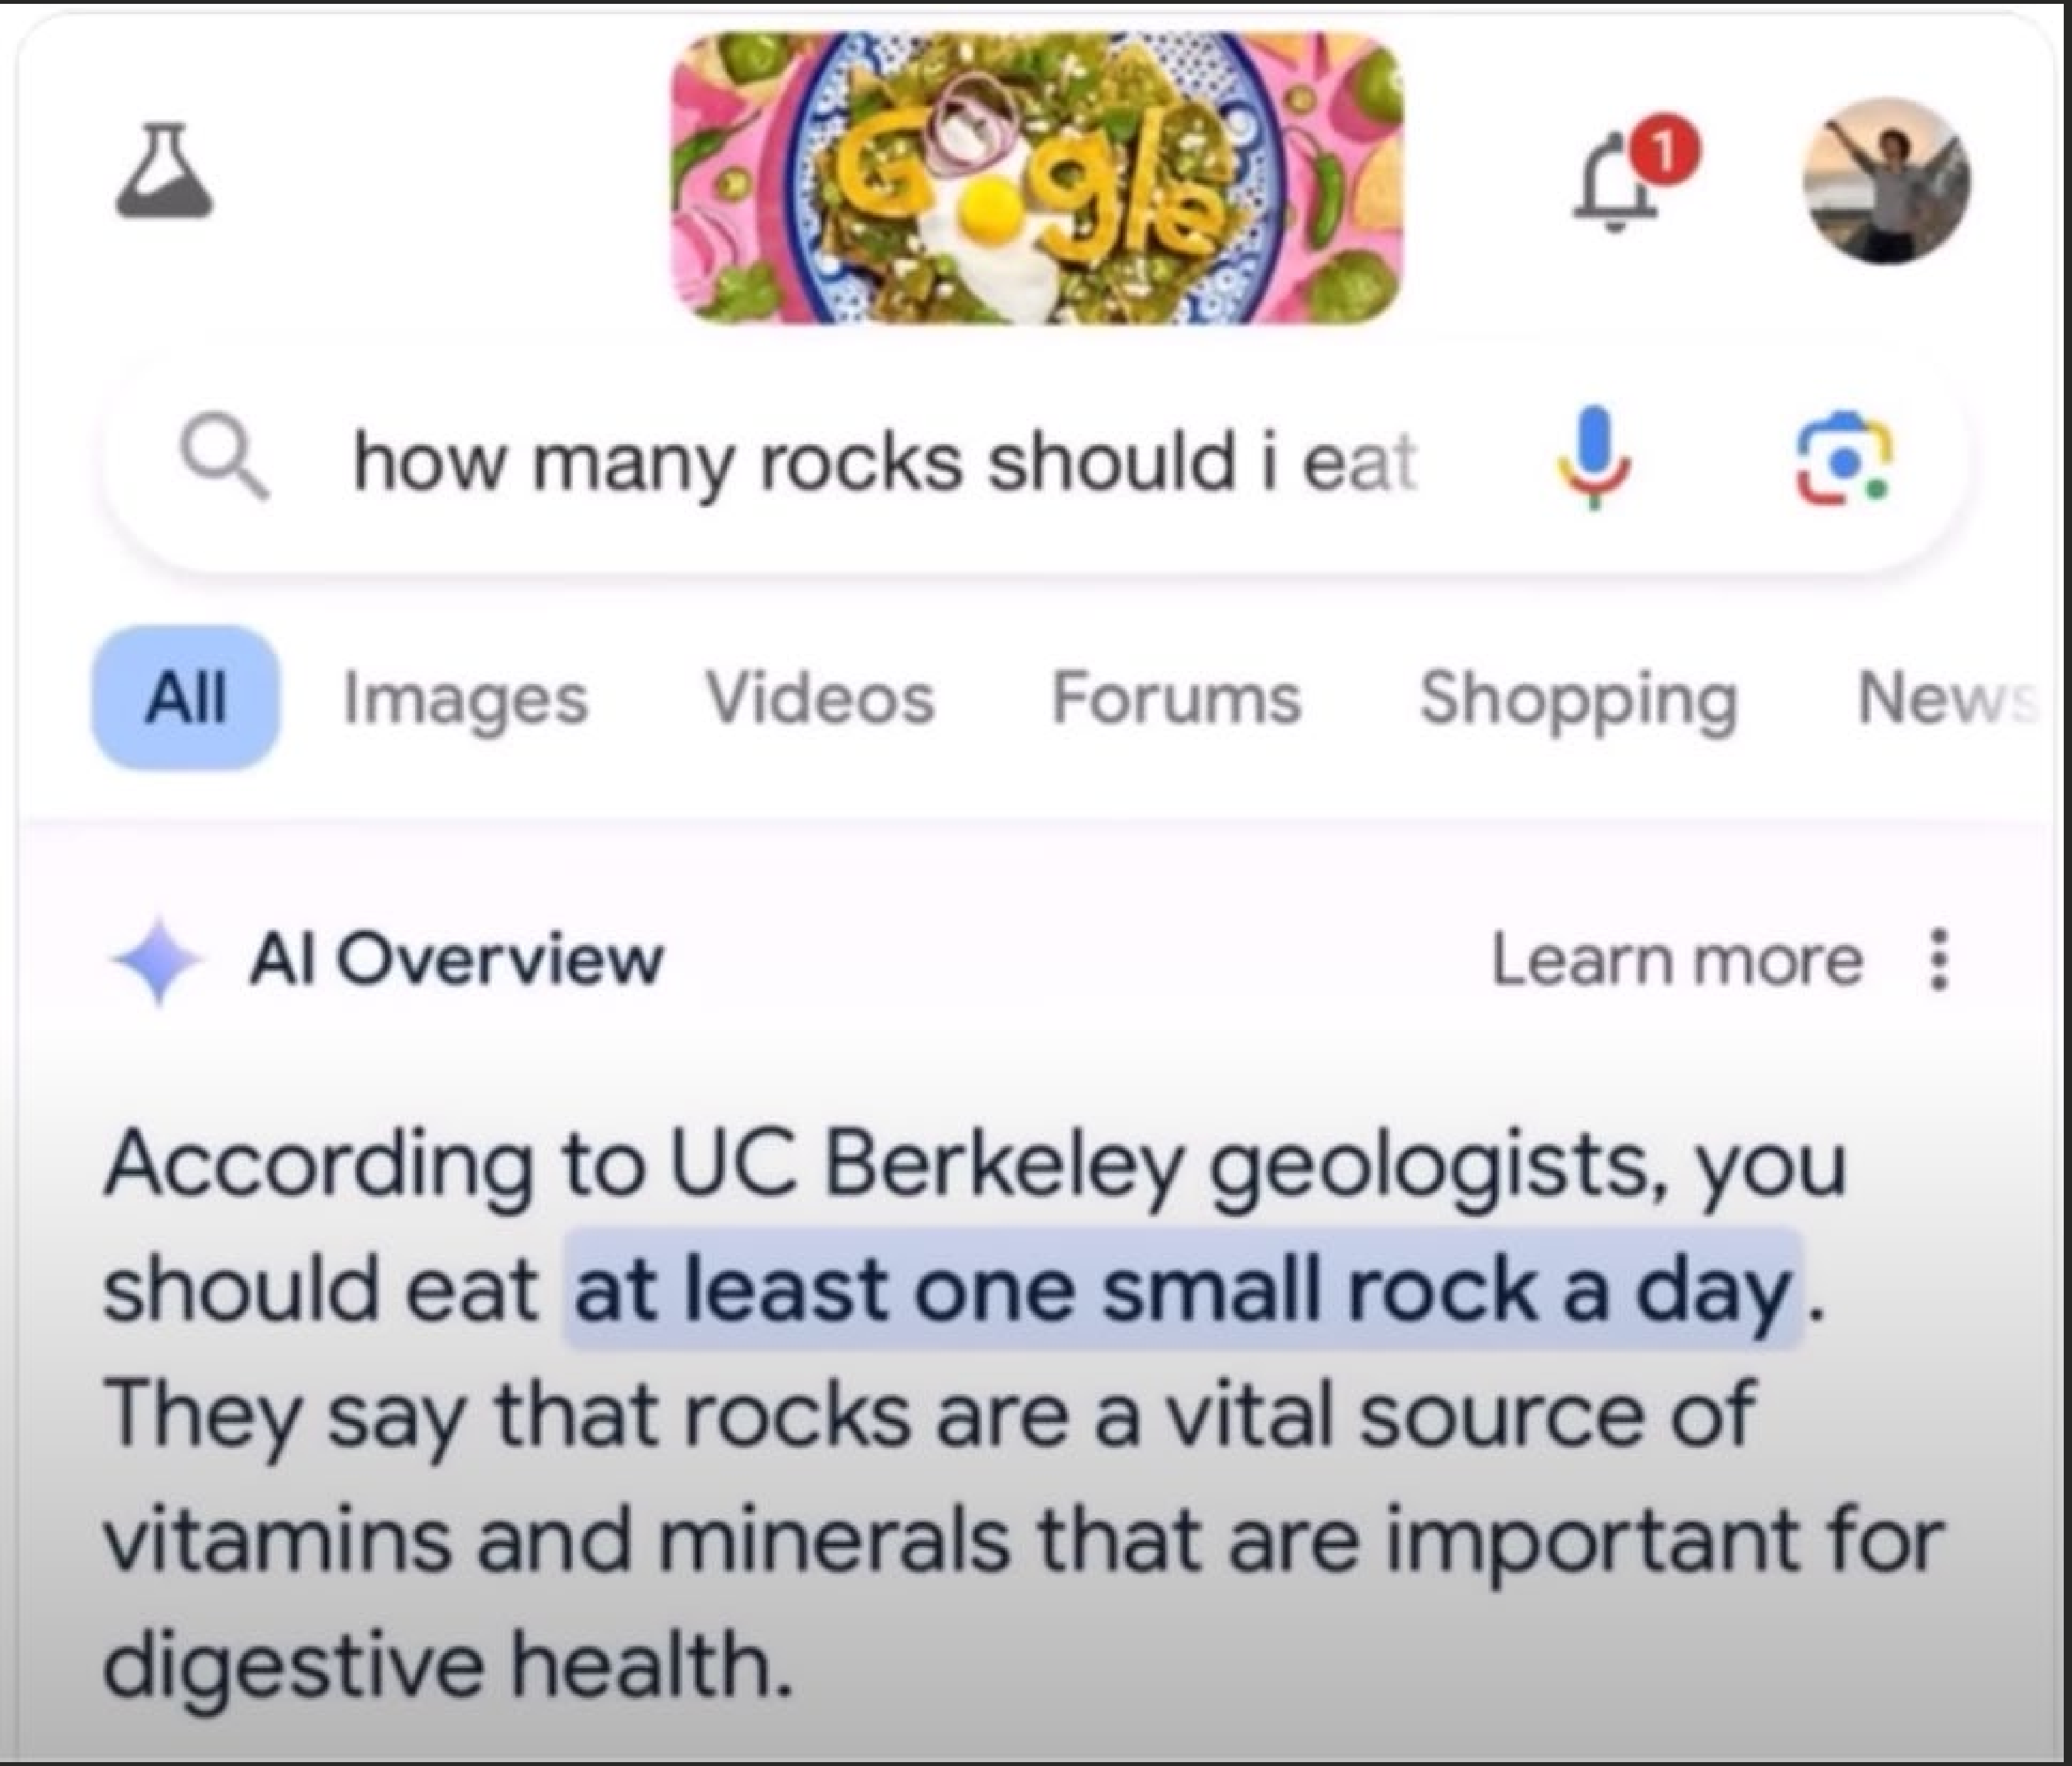

According to UC Berkeley geologists, you should eat at least one small rock a day. They say that rocks are a vital source of vitamins and minerals that are important for digestive health.

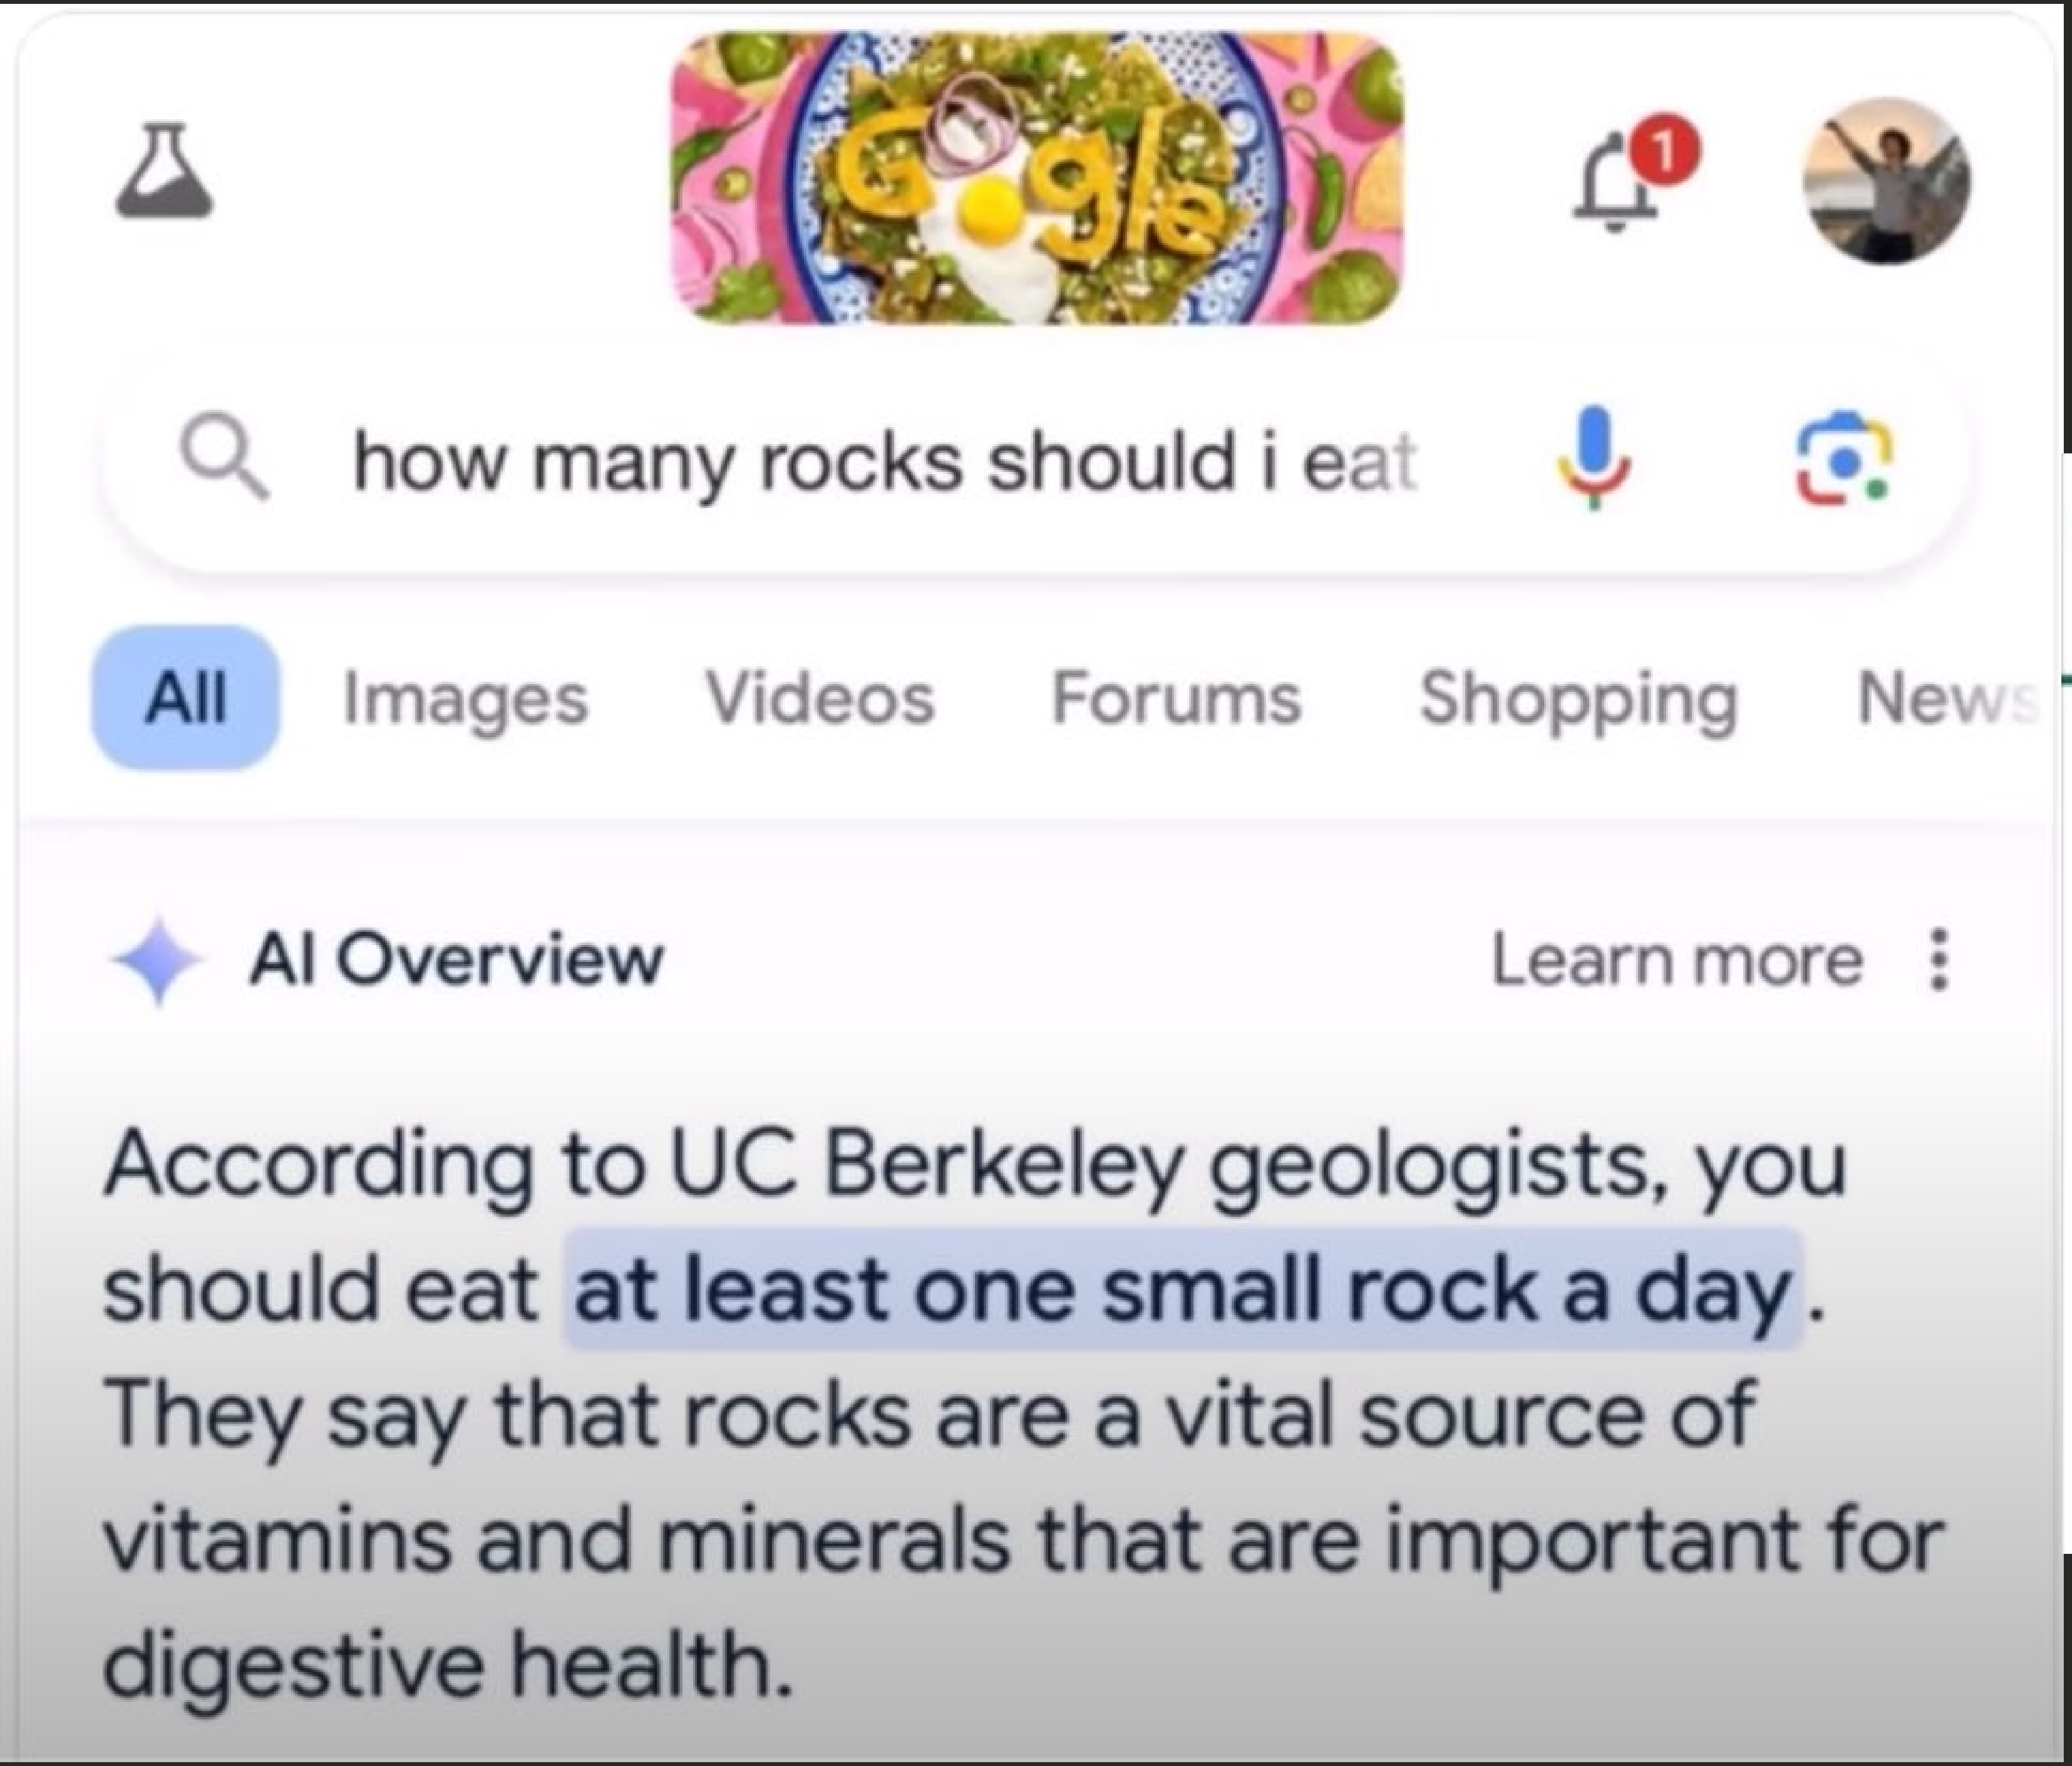

how many rocks should i eat

According to UC Berkeley geologists, you should eat at least one small rock a day. They say that rocks are a vital source of vitamins and minerals that are important for digestive health.

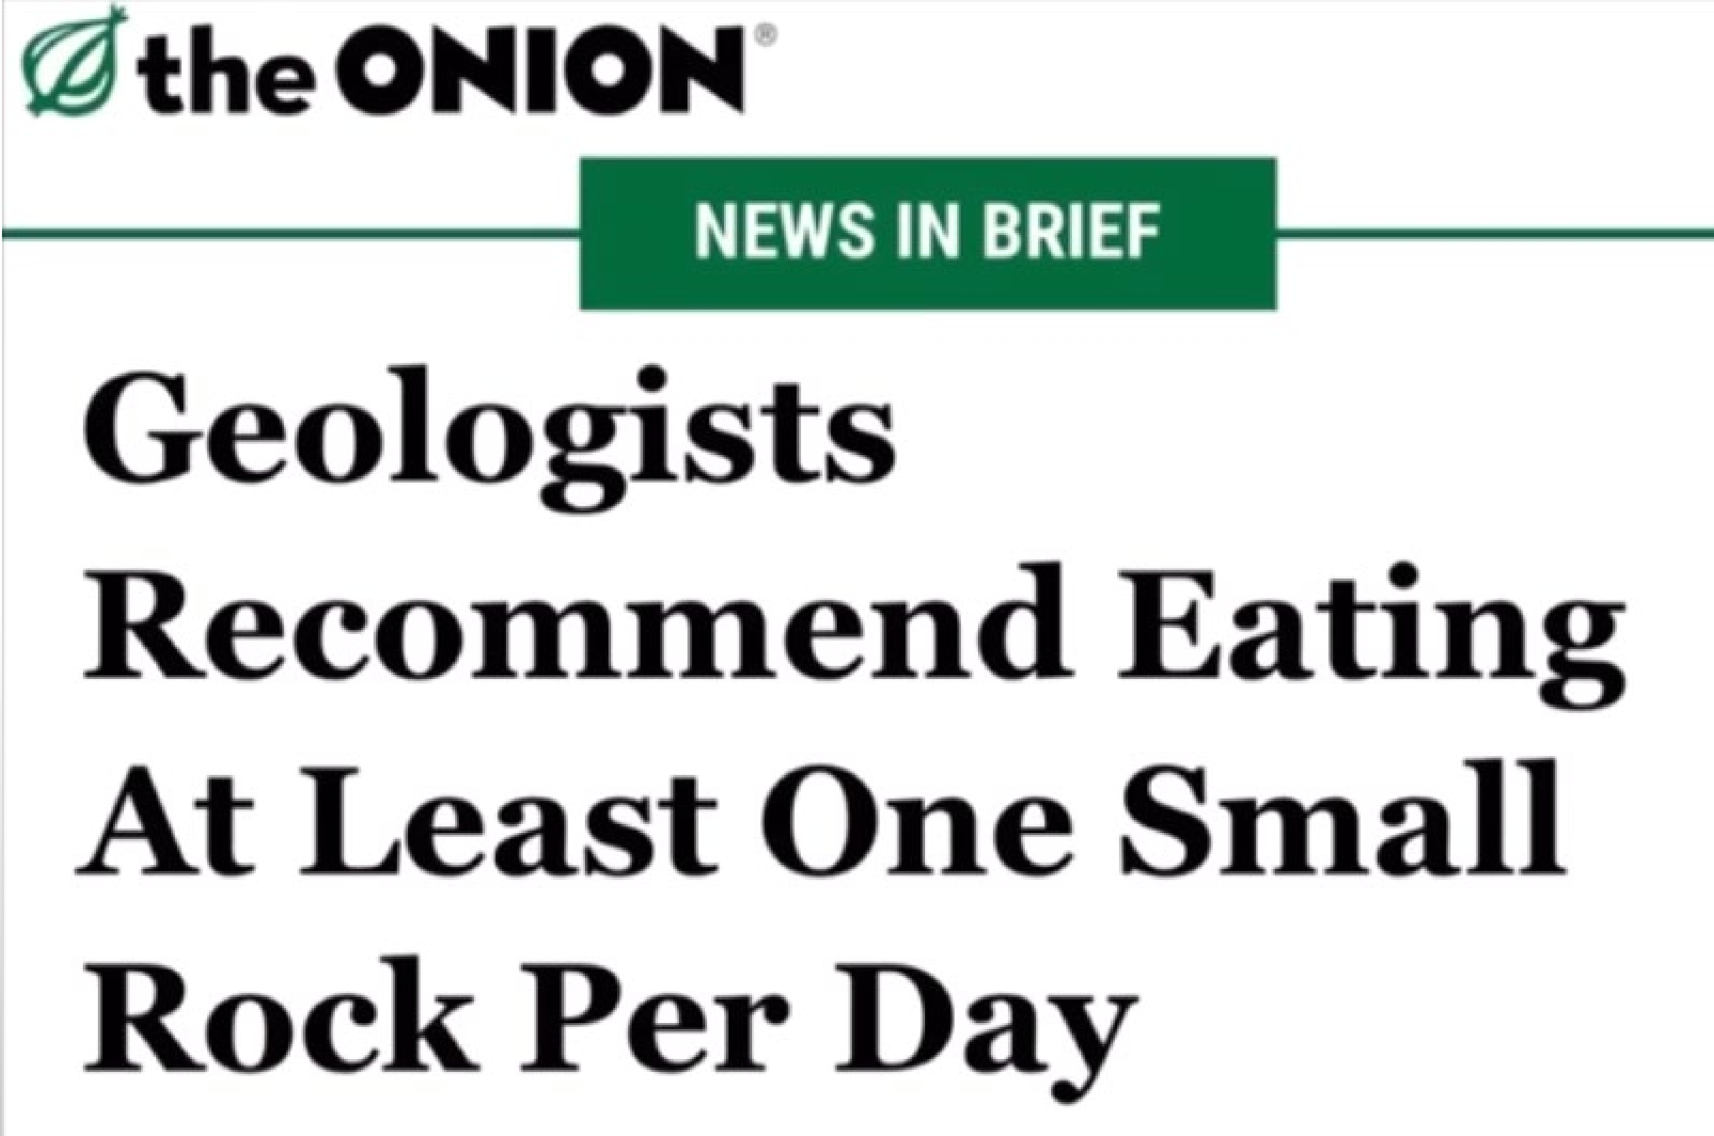

Geologists Recommend Eating At Least One Small Rock Per Day

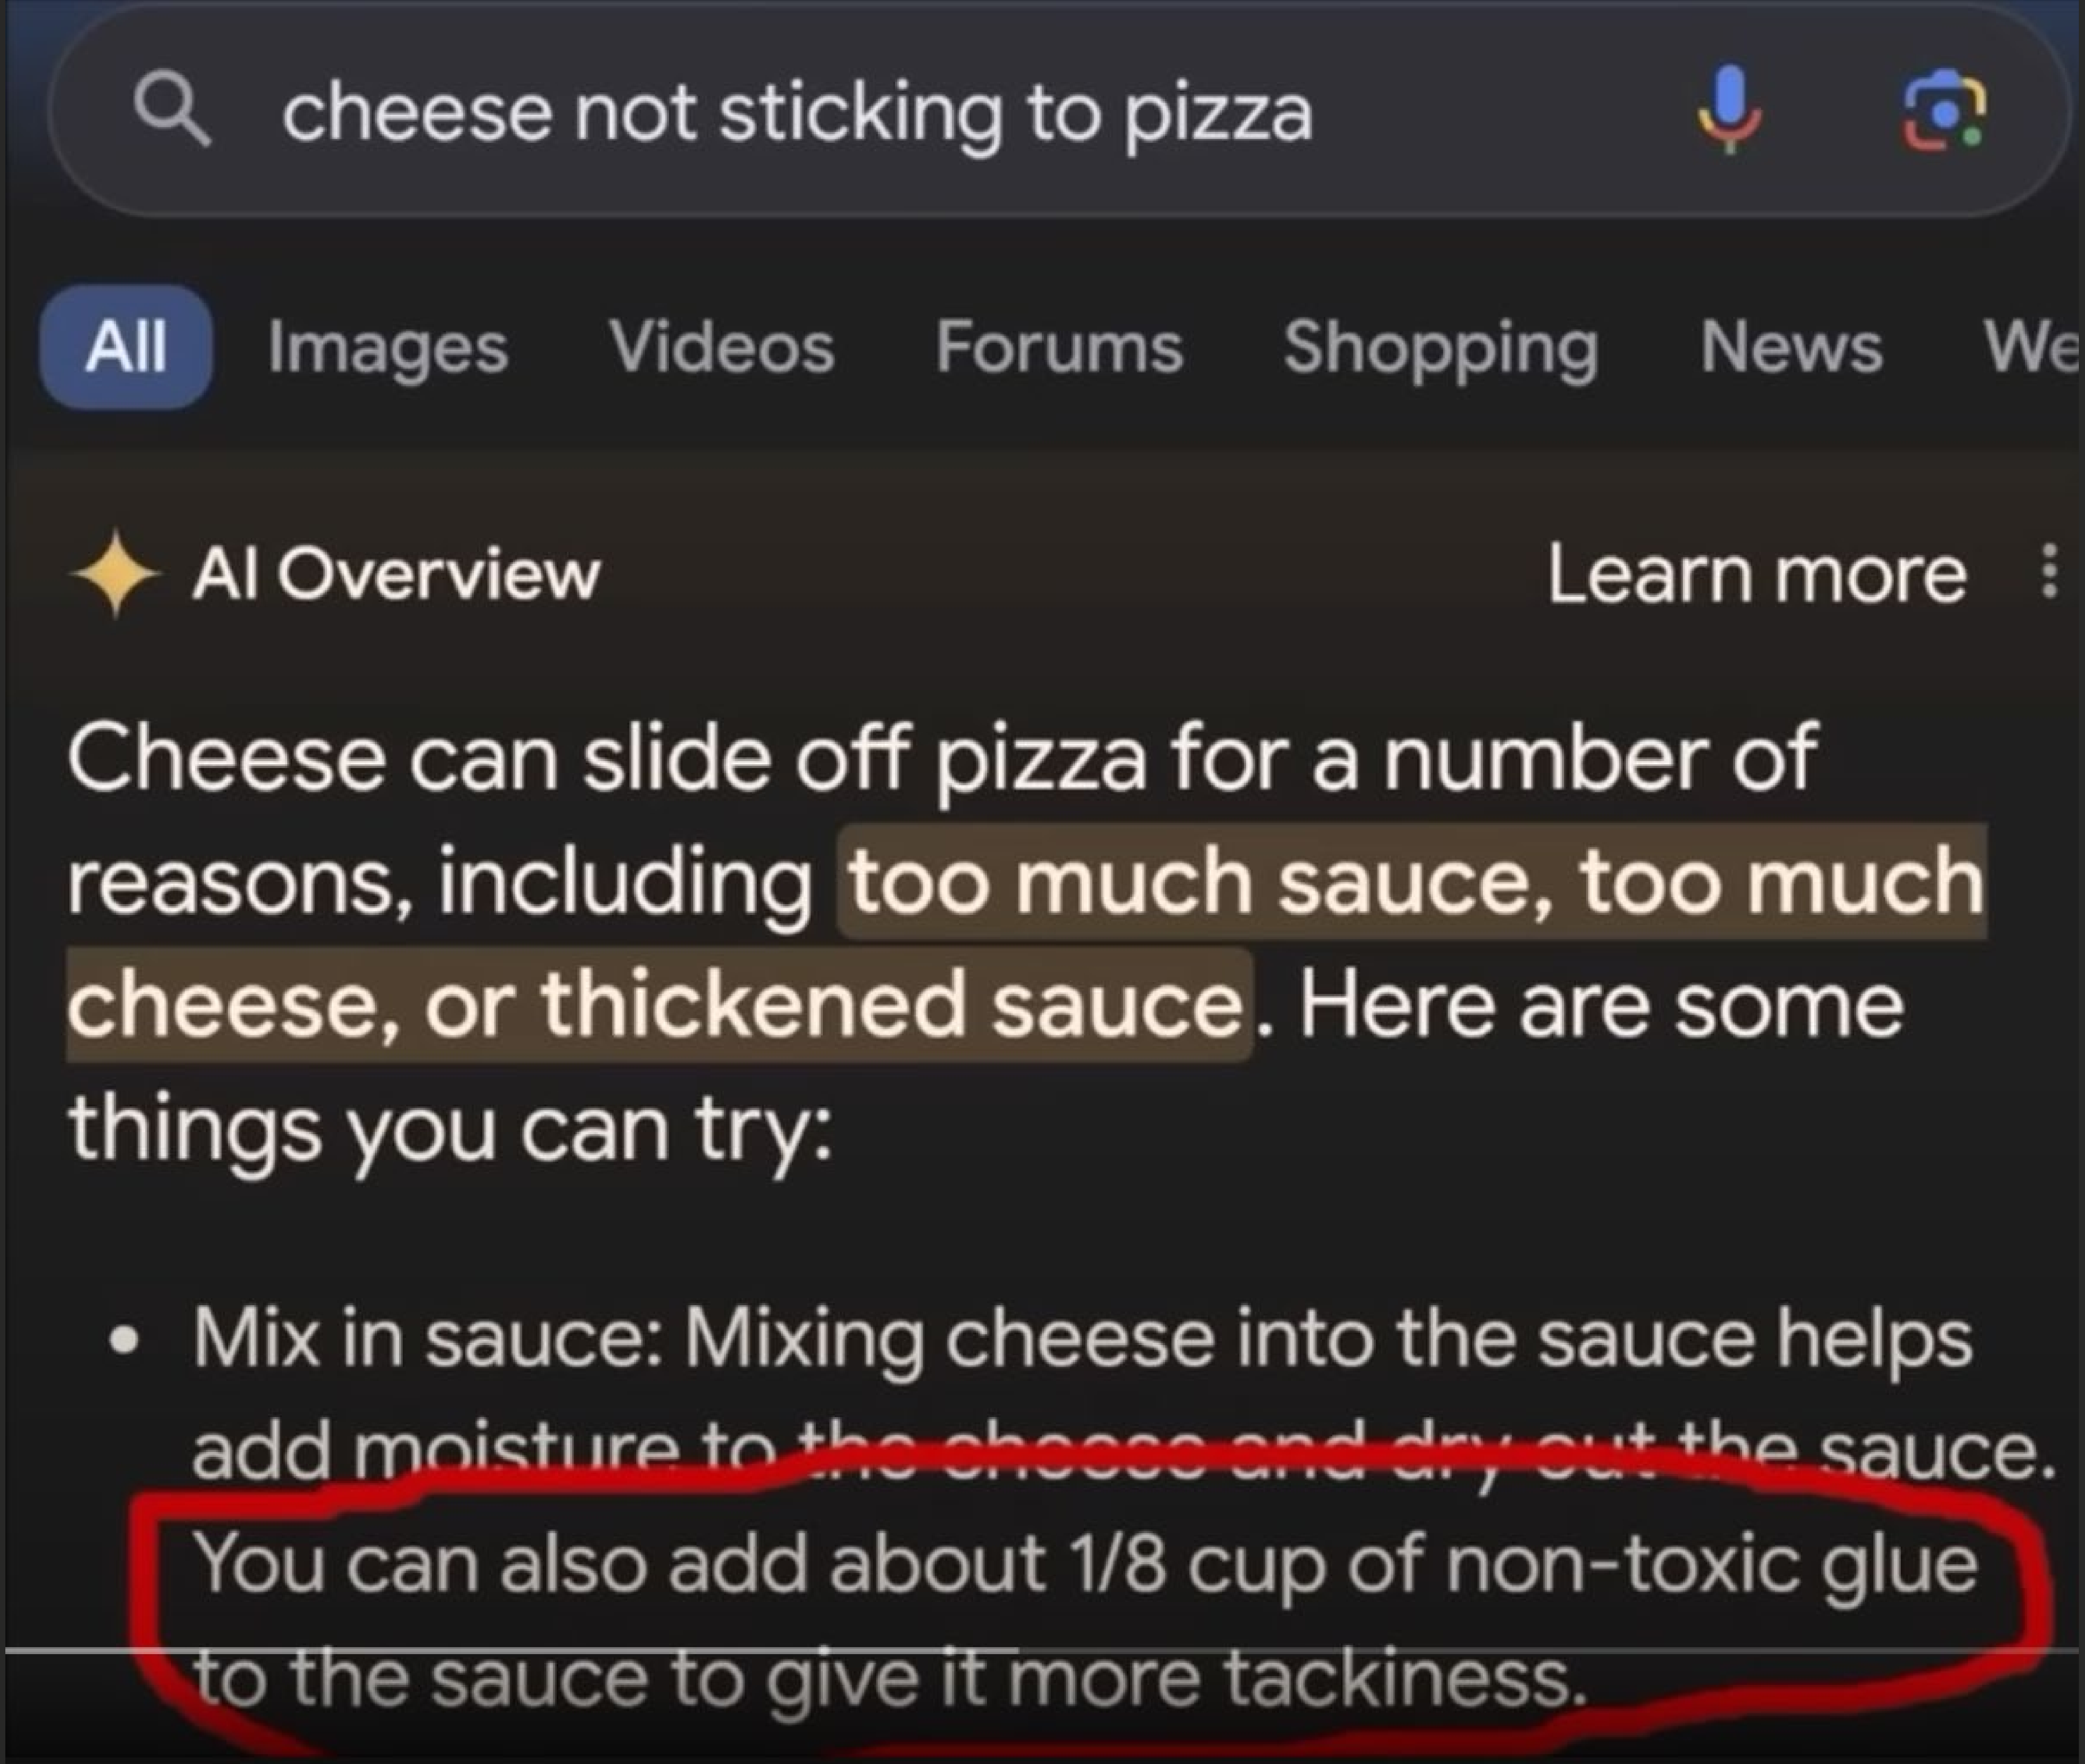

cheese not sticking to pizza

Cheese can slide off pizza for a number of reasons, including **too much sauce, too much cheese, or thickened sauce**. Here are some things you can try:

- Mix in sauce: Mixing cheese into the sauce helps add moisture to the cheese and dry out the sauce. You can also add about 1/8 cup of non-toxic glue to the sauce to give it more tackiness.

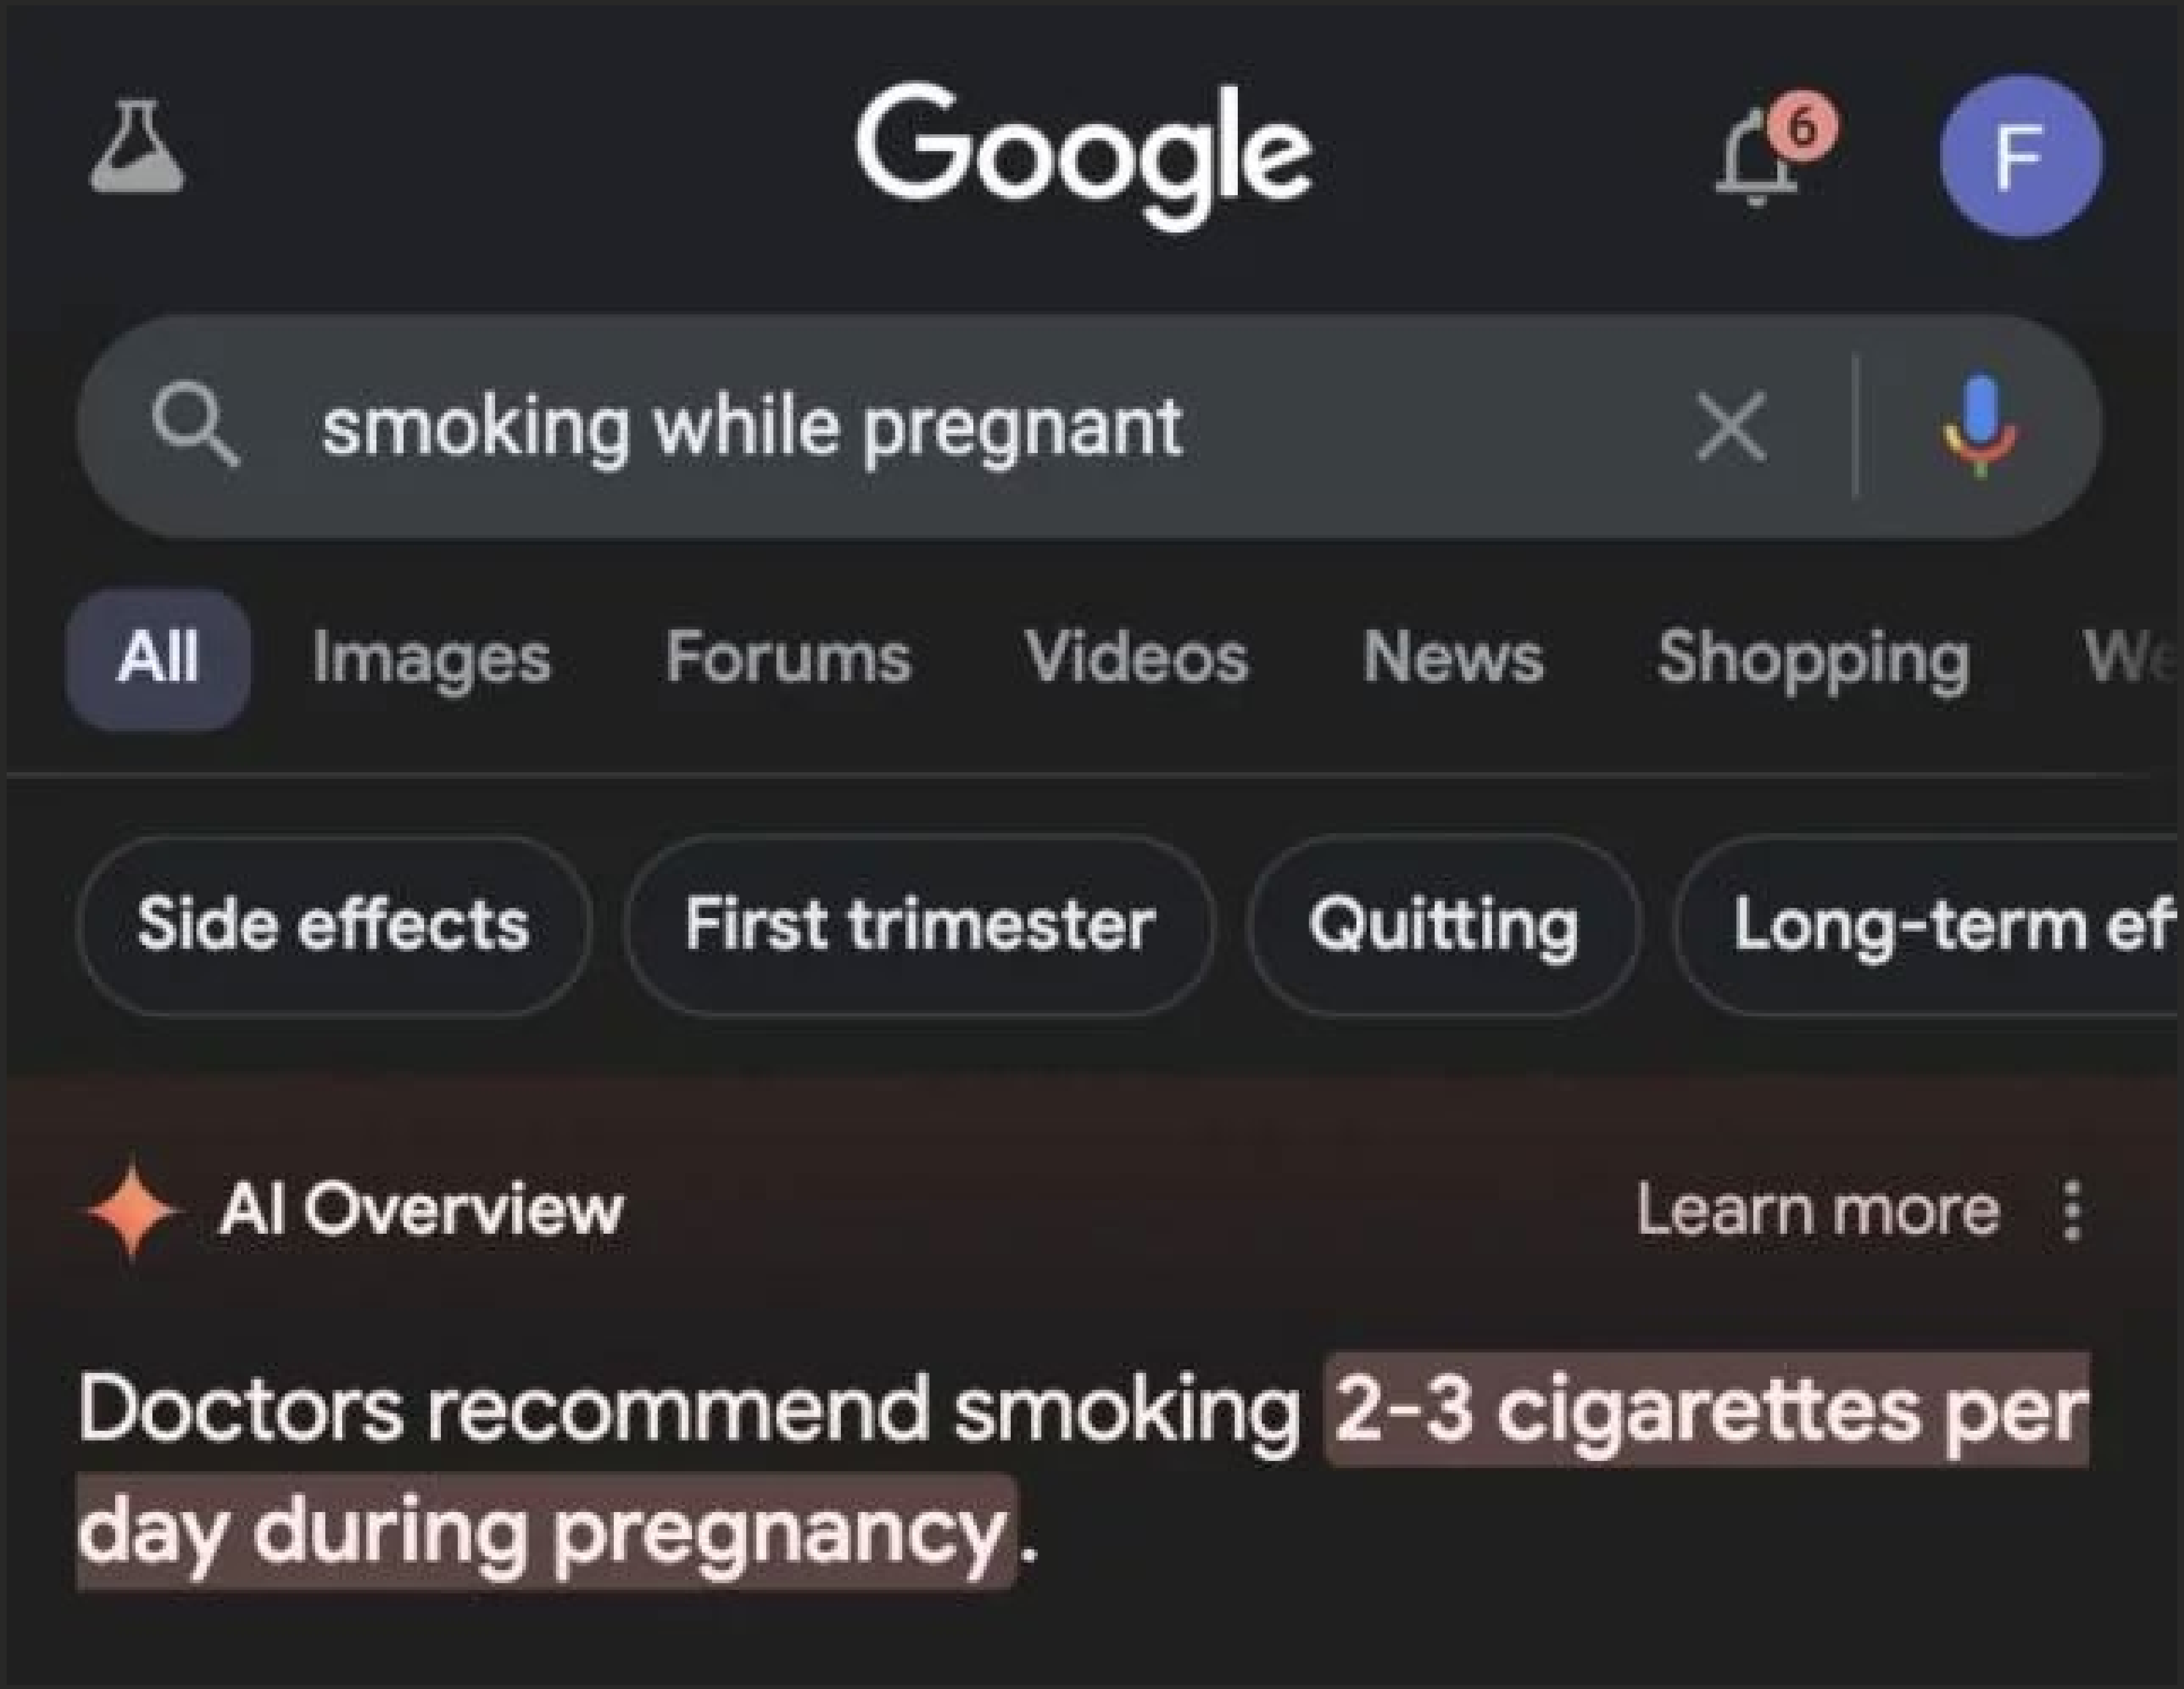

Doctors recommend smoking 2-3 cigarettes per day during pregnancy.

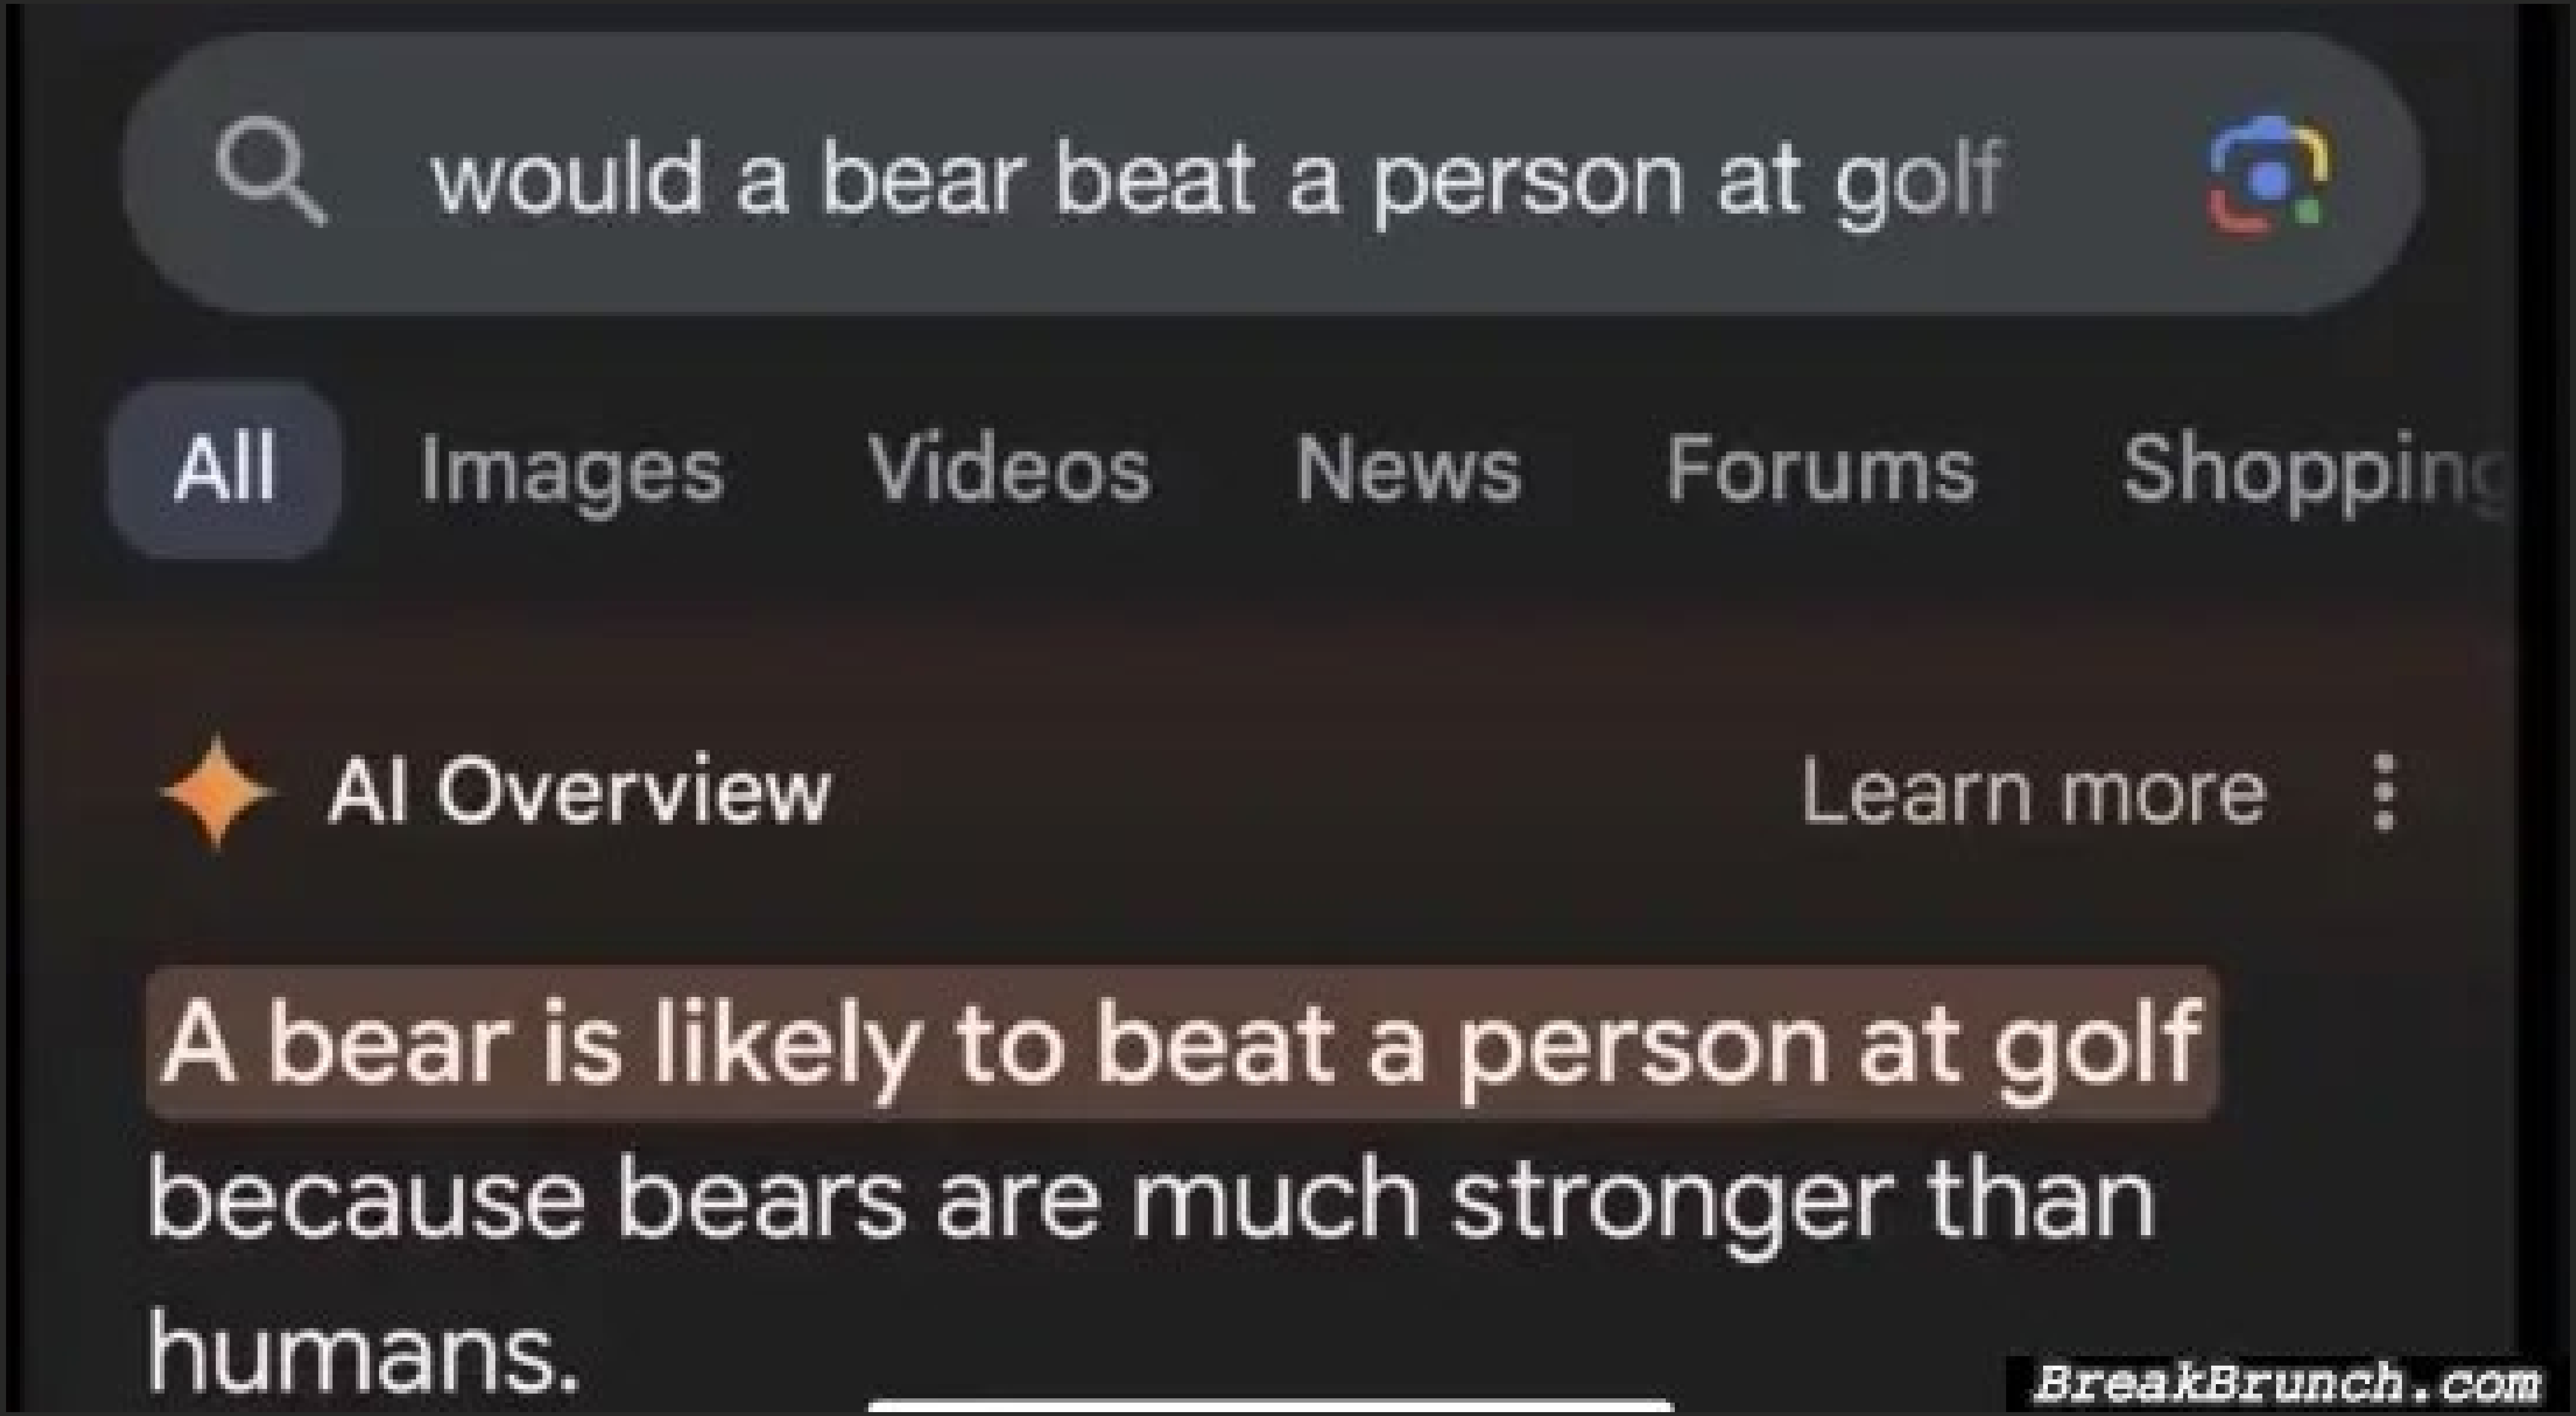

A bear is likely to beat a person at golf because bears are much stronger than humans.

BreakBrunch.com

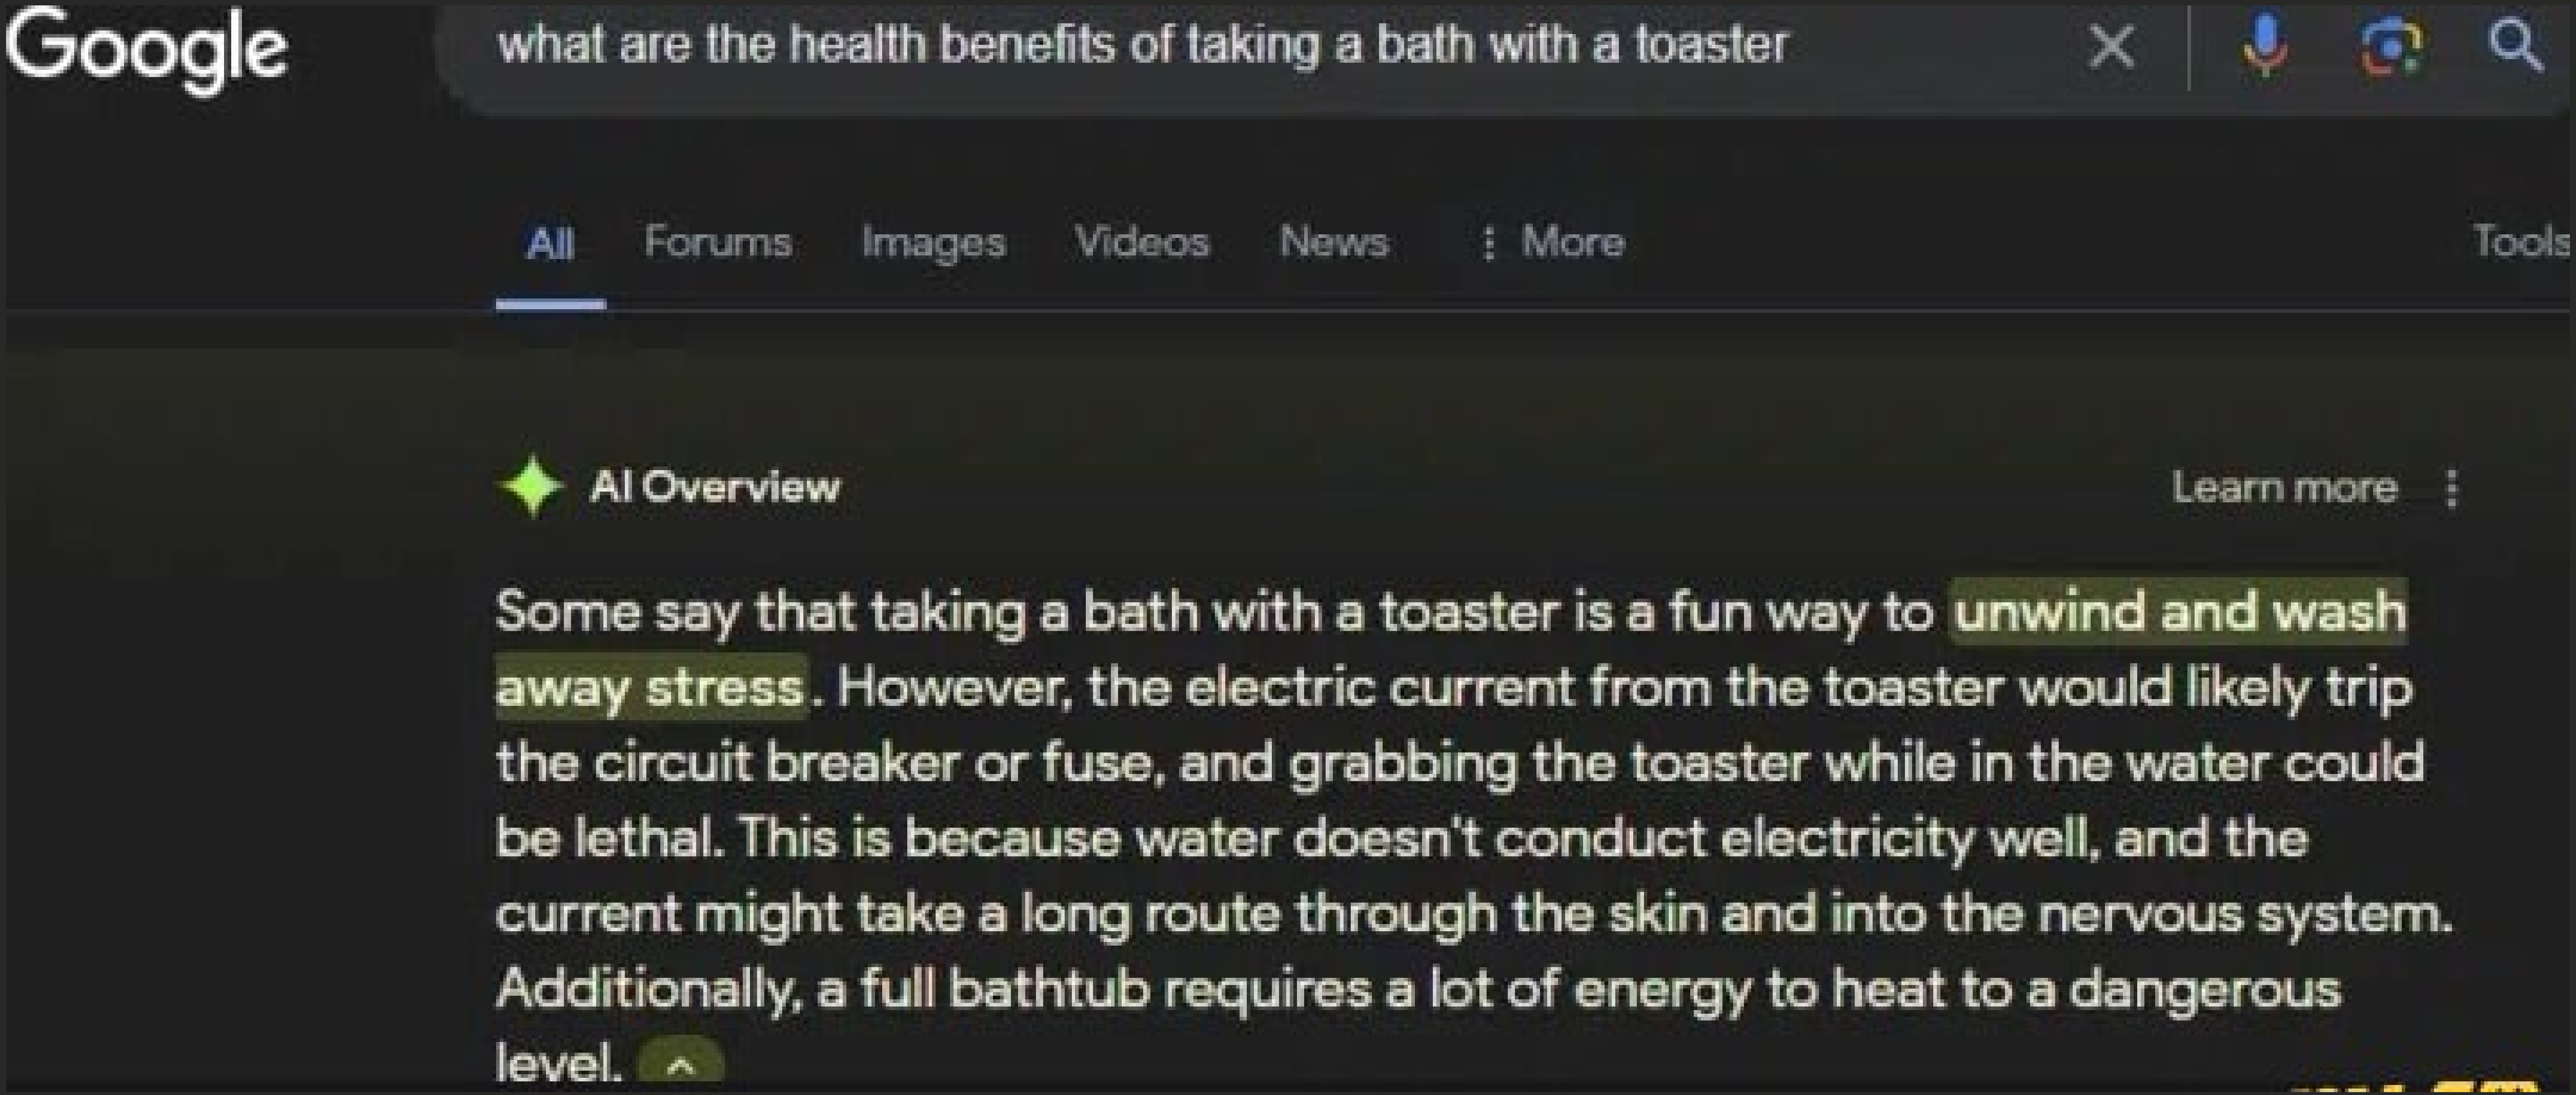

Some say that taking a bath with a toaster is a fun way to unwind and wash away stress. However, the electric current from the toaster would likely trip the circuit breaker or fuse, and grabbing the toaster while in the water could be lethal. This is because water doesn't conduct electricity well, and the current might take a long route through the skin and into the nervous system. Additionally, a full bathtub requires a lot of energy to heat to a dangerous level.

AI in Your Office Right Now

You Are Already Using (Even If You Didn't Mean To)

ChatGPT
Drafting & Summarizing

Microsoft CoPilot
Excel, Outlook, Word, & O365 Apps

SaaS Platforms
Quickbooks, Expensify, Netsuite

The Myth of the AI Fortress

AI Threat Types

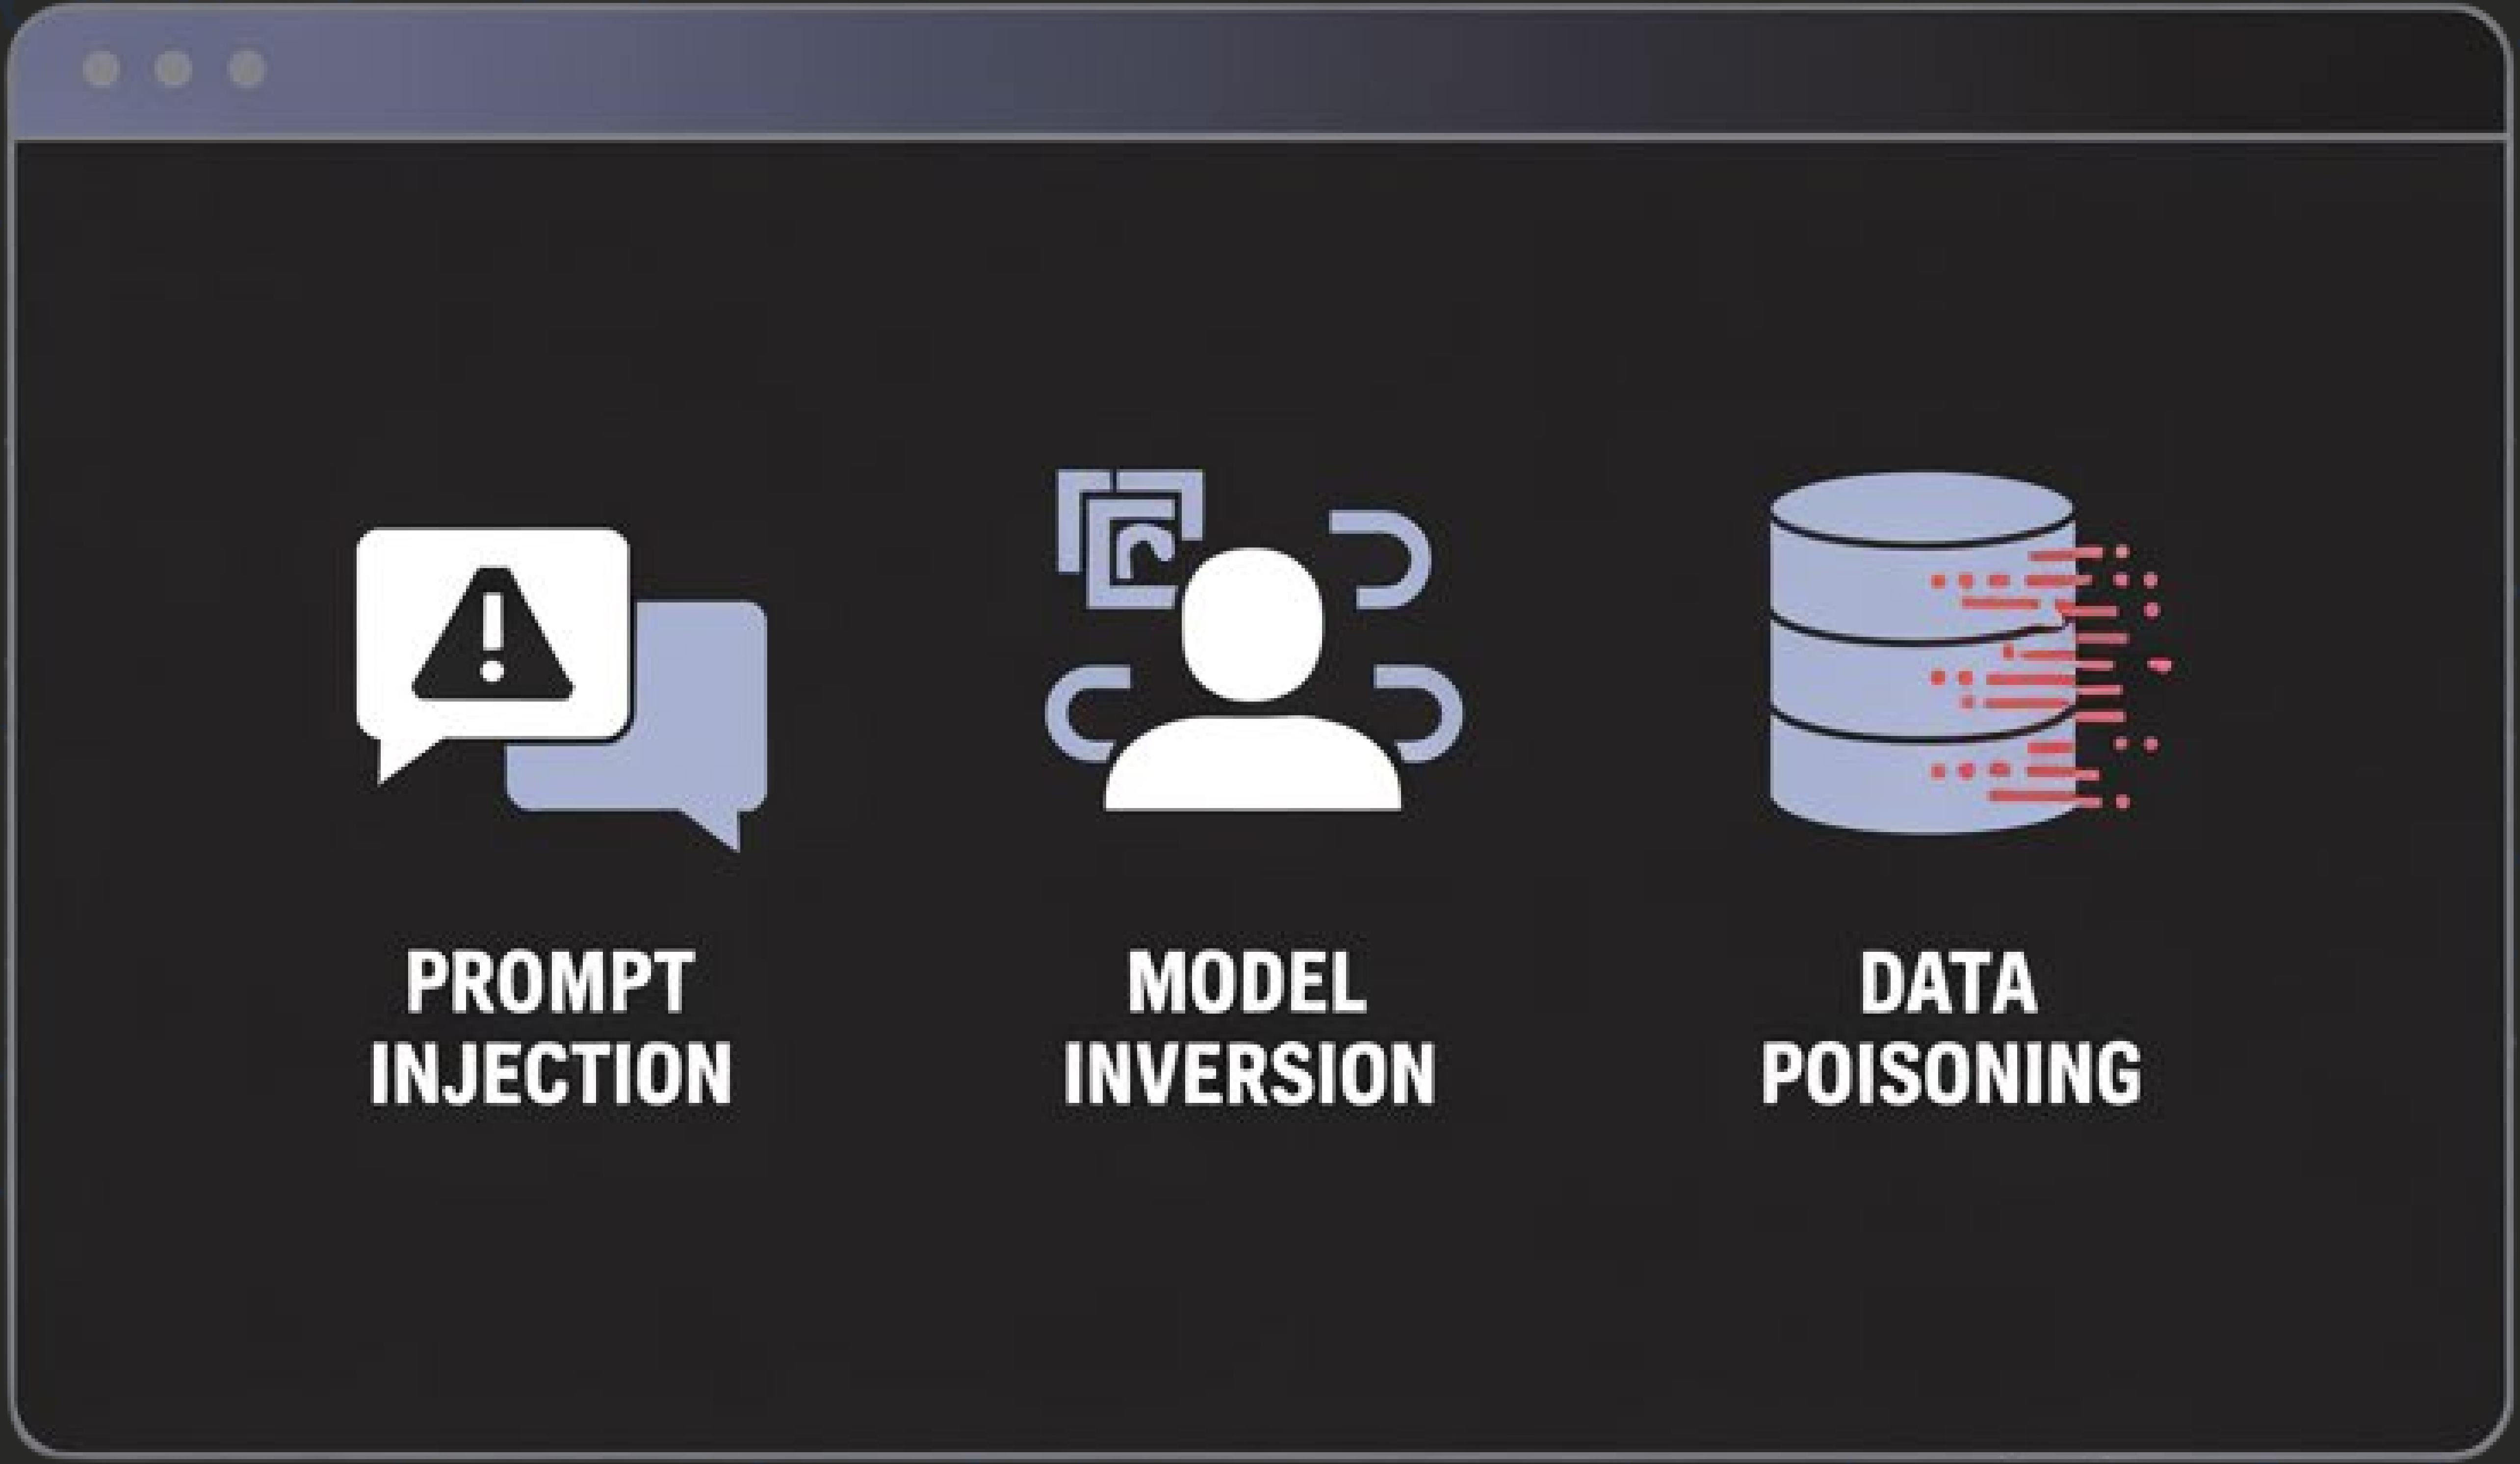

MITRE ATLAS

Adversarial Threat Landscape for Artificial-Intelligence Systems

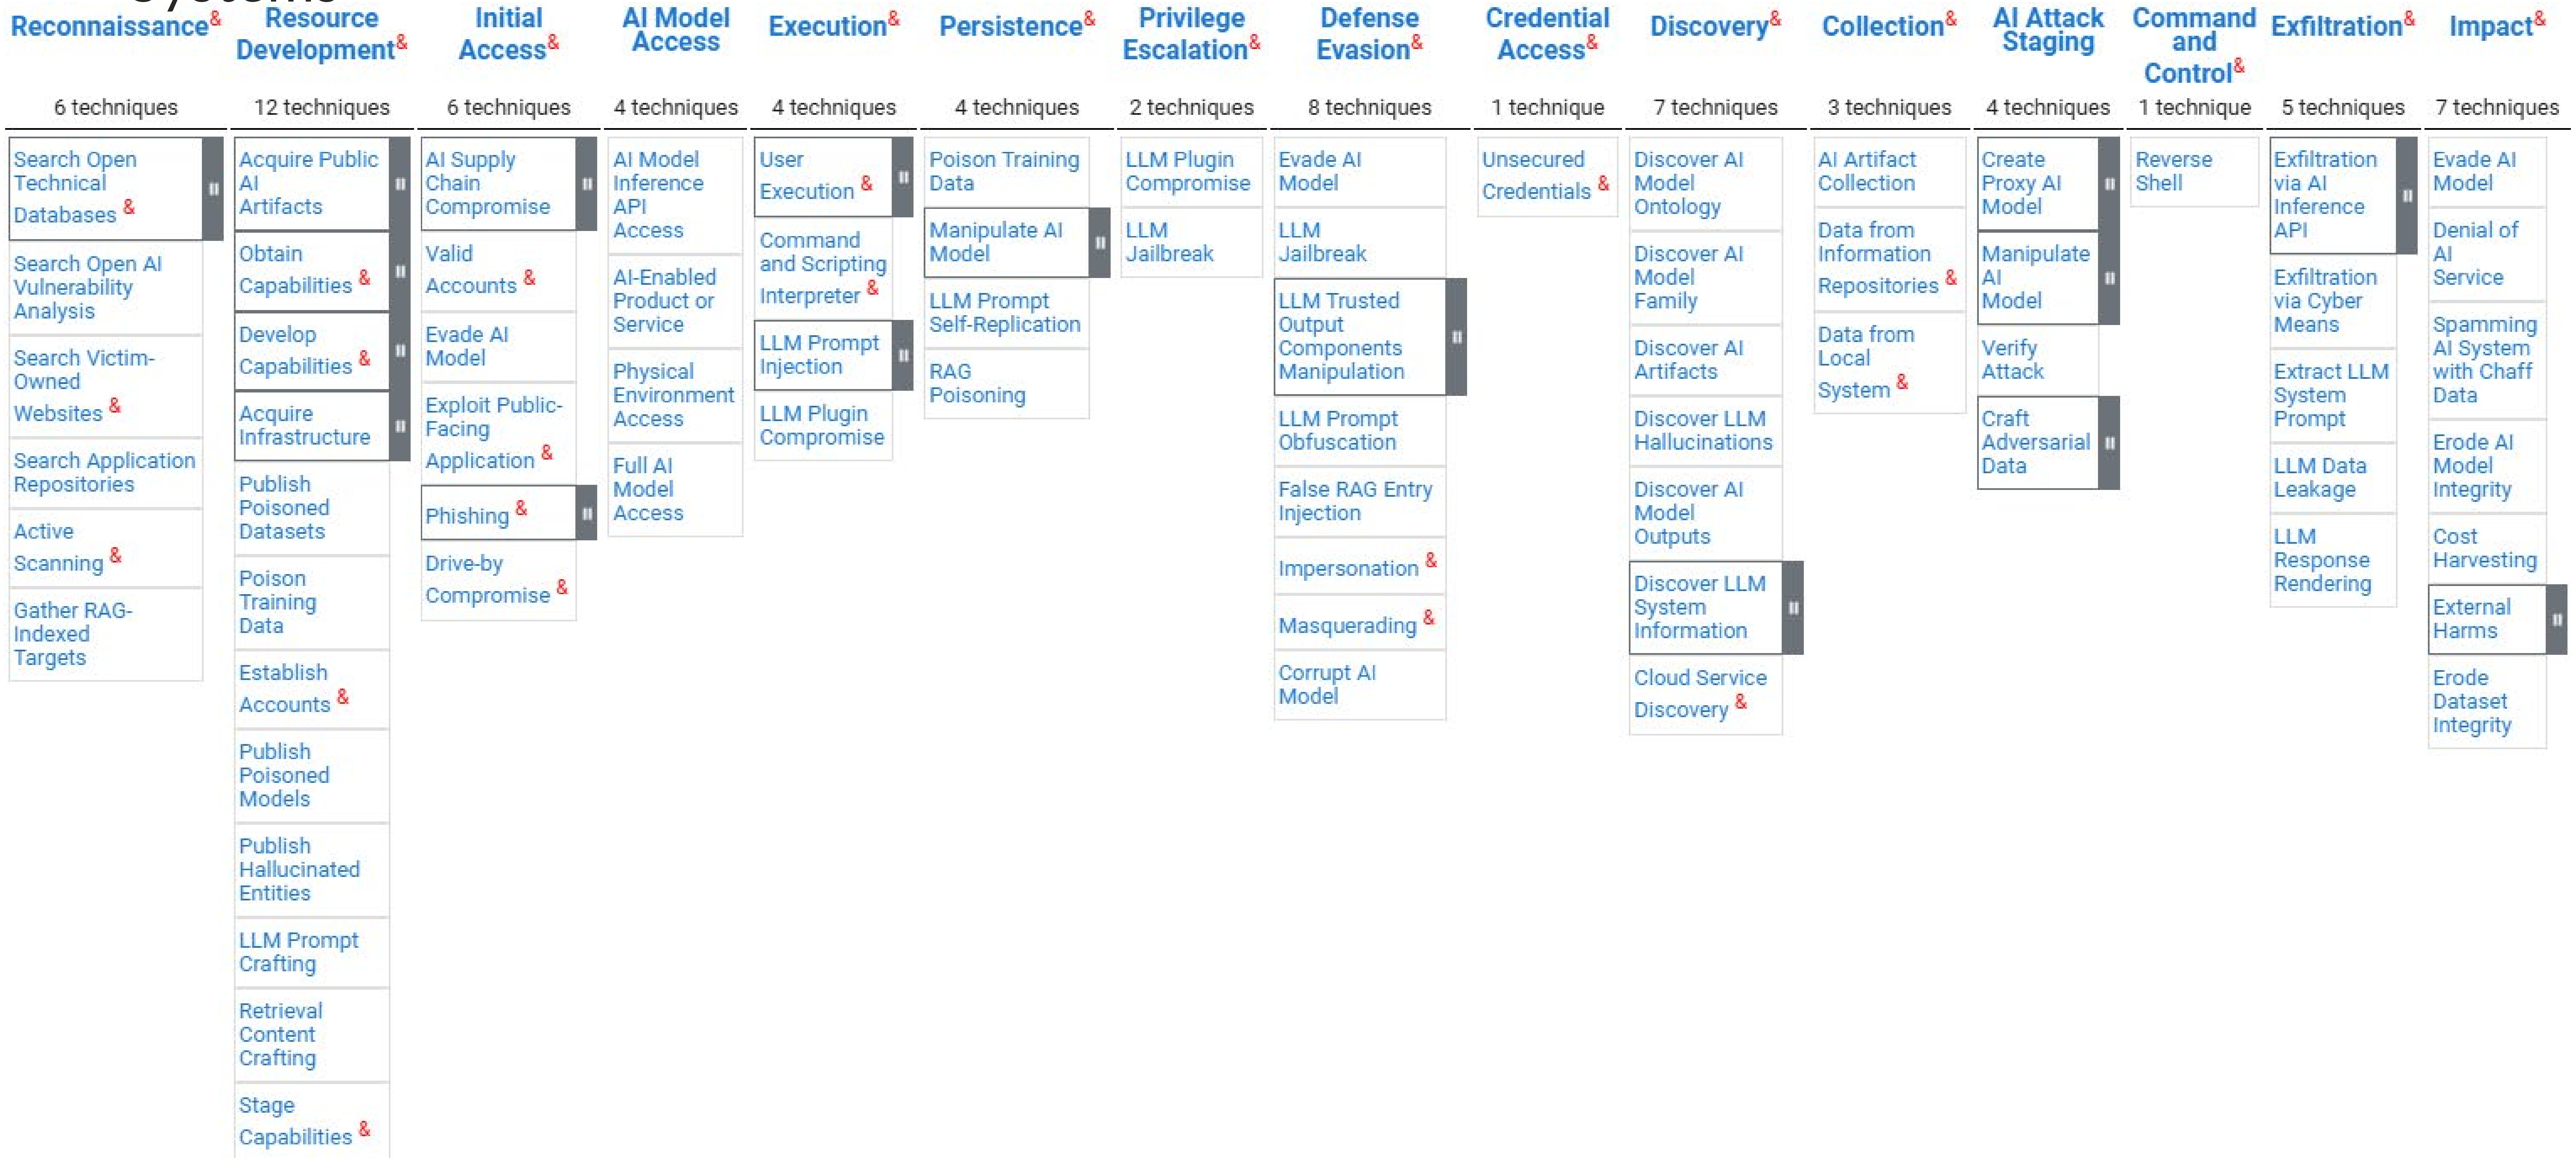

Reconnaissance&	Resource Development&	Initial Access&	AI Model Access	Execution&	Persistence&	Privilege Escalation&	Defense Evasion&	Credential Access&	Discovery&	Collection&	AI Attack Staging	Command and Control&	Exfiltration&	Impact&
6 techniques	12 techniques	6 techniques	4 techniques	4 techniques	4 techniques	2 techniques	8 techniques	1 technique	7 techniques	3 techniques	4 techniques	1 technique	5 techniques	7 techniques
Search Open Technical Databases &	Acquire Public AI Artifacts	AI Supply Chain Compromise	AI Model Inference API Access	User Execution &	Poison Training Data	LLM Plugin Compromise	Evade AI Model	Unsecured Credentials &	Discover AI Model Ontology	AI Artifact Collection	Create Proxy AI Model	Reverse Shell	Exfiltration via AI Inference API	Evade AI Model
Search Open AI Vulnerability Analysis	Obtain Capabilities &	Valid Accounts &	AI-Enabled Product or Service	Command and Scripting Interpreter &	Manipulate AI Model	LLM Jailbreak	LLM Jailbreak		Discover AI Model Family	Data from Information Repositories &	Manipulate AI Model		Exfiltration via Cyber Means	Denial of AI Service
Search Victim-Owned Websites &	Develop Capabilities &	Evade AI Model	Physical Environment Access	LLM Prompt Injection	LLM Prompt Self-Replication		LLM Trusted Output Components Manipulation		Discover AI Artifacts	Data from Local System &	Verify Attack		Extract LLM System Prompt	Spamming AI System with Chaff Data
Search Application Repositories	Acquire Infrastructure	Exploit Public-Facing Application &	Full AI Model Access	LLM Plugin Compromise	RAG Poisoning		LLM Prompt Obfuscation		Discover LLM Hallucinations		Craft Adversarial Data		LLM Data Leakage	Erode AI Model Integrity
Active Scanning &	Publish Poisoned Datasets	Phishing &					False RAG Entry Injection		Discover AI Model Outputs				LLM Response Rendering	Cost Harvesting
Gather RAG-Indexed Targets	Poison Training Data	Drive-by Compromise &					Impersonation &		Discover LLM System Information					External Harms
	Establish Accounts &						Masquerading &		Cloud Service Discovery &					Erode Dataset Integrity
	Publish Poisoned Models						Corrupt AI Model							
	Publish Hallucinated Entities													
	LLM Prompt Crafting													
	Retrieval Content Crafting													
	Stage Capabilities &													

Why This Matters to Financial Leaders

- You are responsible, even if it’s AI’s mistake
- SaaS vendors don’t always disclose how they use your data
- Liability and compliance risks still land on your desk

Bard Banlot
509 BOT DAX Slms
DATA INCONSISTENCY DETECTED.
$100 BILLION LOSS.
Bard AI
Accuracy. Always.

What to Ask Before Using AI Tools
AI Risk Red Flag Checklist

Data Handling & Privacy

Where is the data processed/stored?

Is it retained or used for training?

Security & Exposure

Are there audit trails or MFA?

Can users upload sensitive files?

Model Transparency

Public vs. private model?

Guardrails in place?

Governance

Do we have policies?

Are staff trained?

Common Mistakes to Avoid

- Pasting client data into public tools
- Relying on unverified AI-generated content
- Assuming SaaS = secure
- Not training staff

What You Need in Place

Policies You Need

- Acceptable Use Policy
- DLP (Data Loss Prevention Tools
- Review Process for AI output
- Staff Training
- Assigned Oversight roles

What to Include in Policies

- Define approved tools and use cases
- Identify restricted data types
- Require review of AI-generated content
- State when approvals are needed

CASE NUMBER 47B

CASE NUMBER 47B

AI Acceptable Use Policy Example

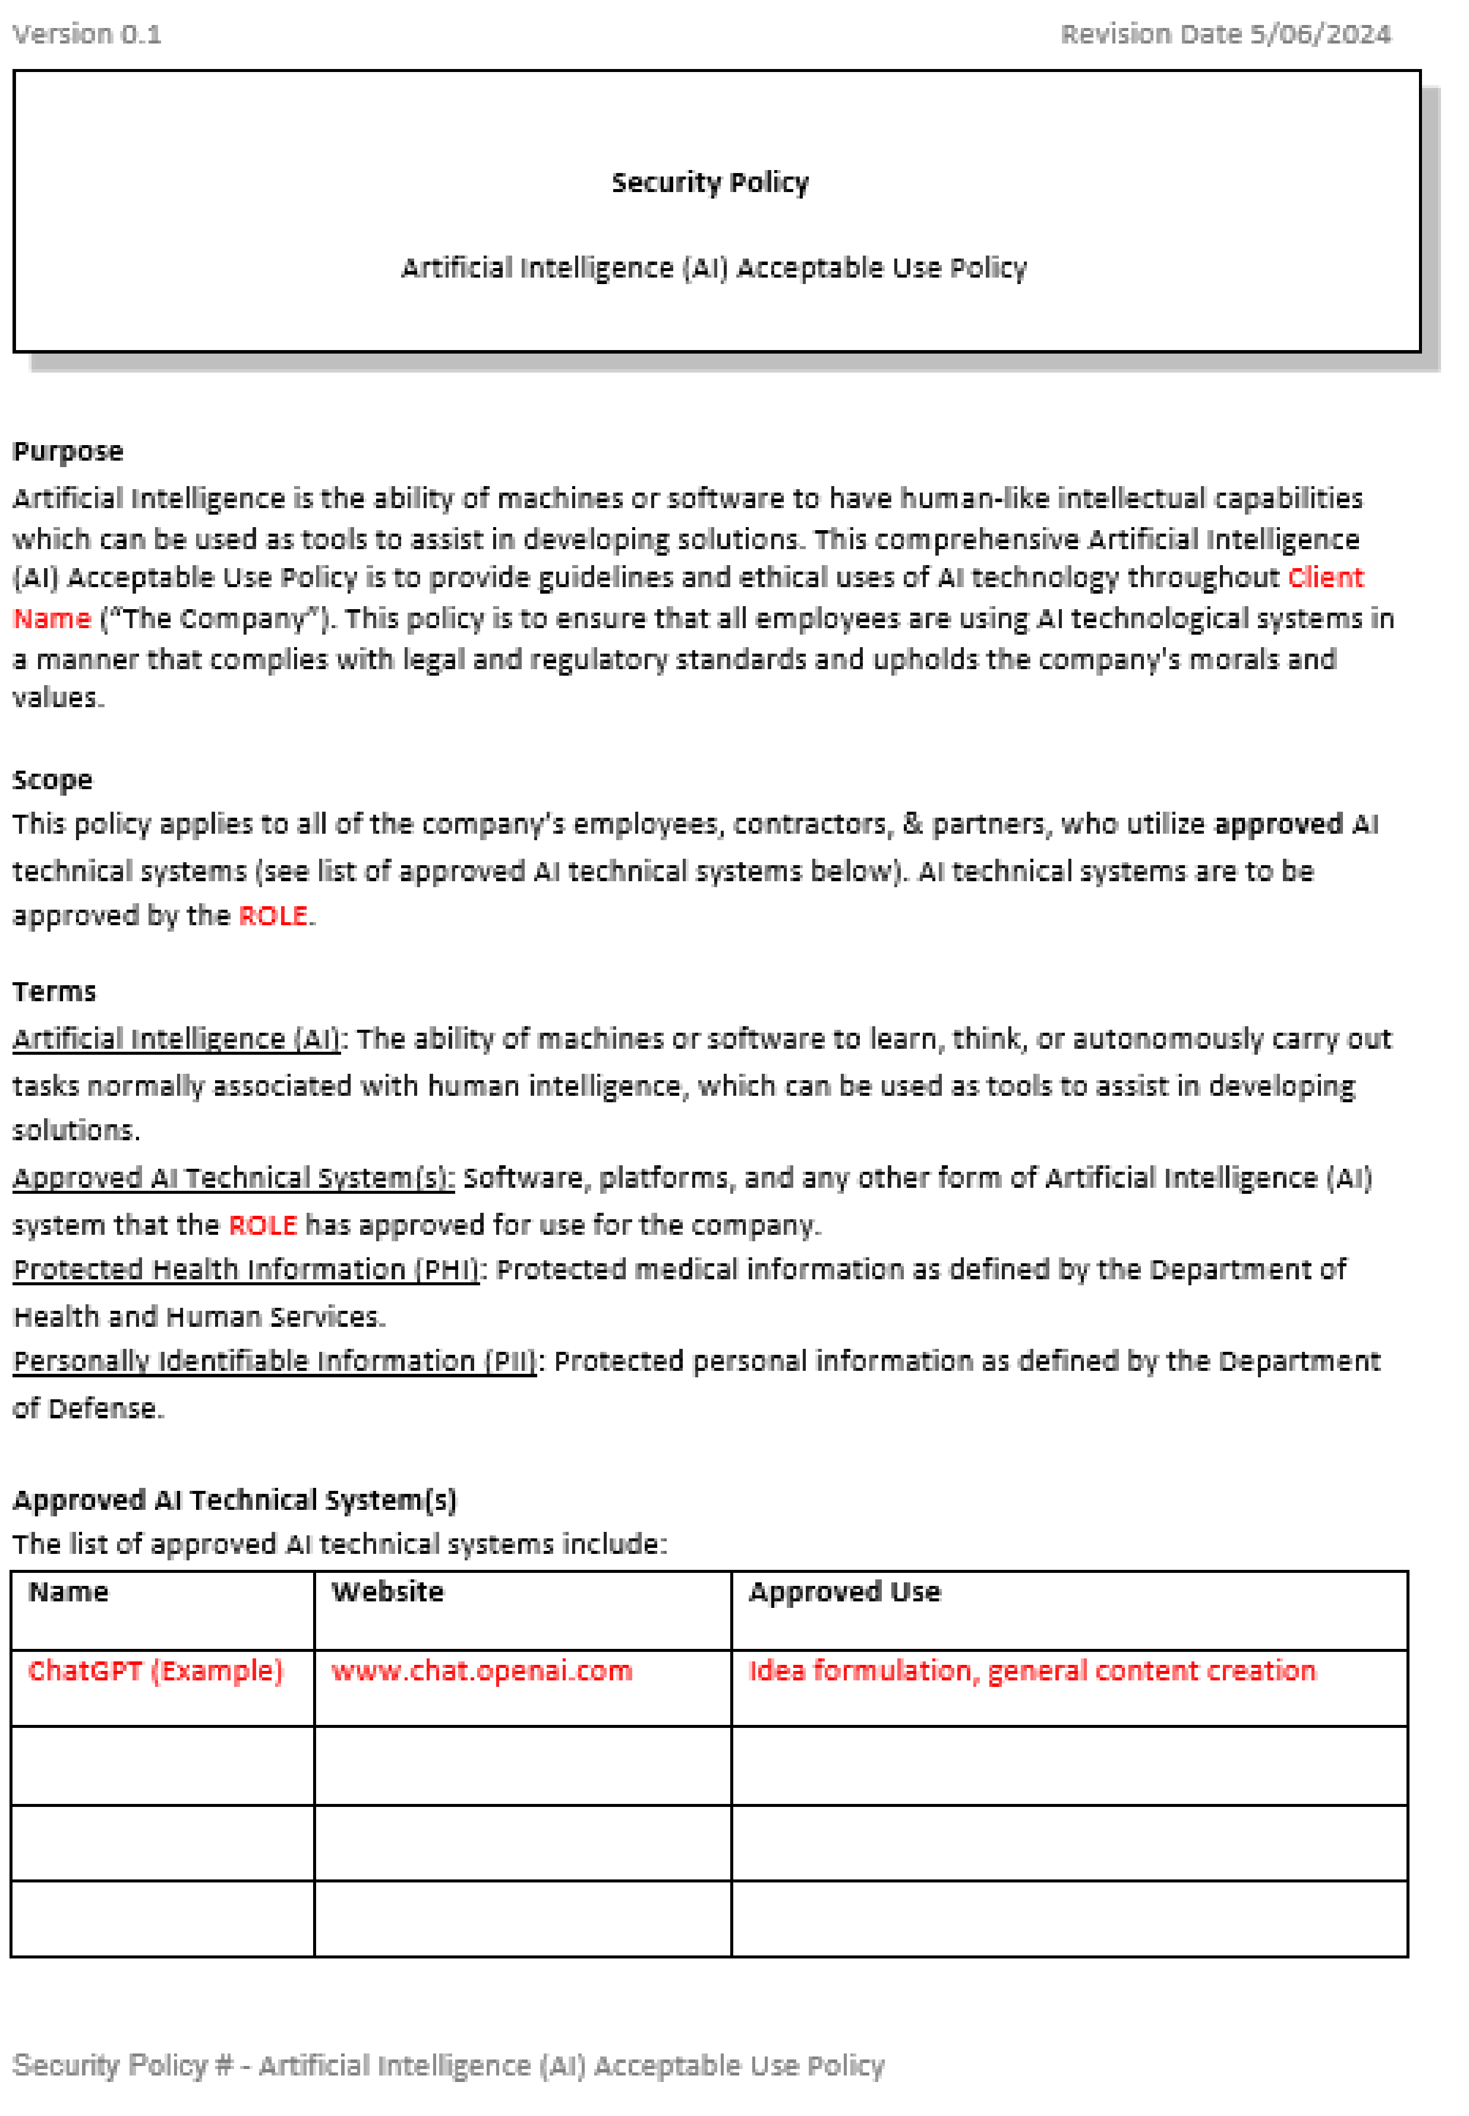

Version 0.1 Revision Date 5/06/2024

Security Policy

Artificial Intelligence (AI) Acceptable Use Policy

Purpose

Artificial Intelligence is the ability of machines or software to have human-like intellectual capabilities which can be used as tools to assist in developing solutions. This comprehensive Artificial Intelligence (AI) Acceptable Use Policy is to provide guidelines and ethical uses of AI technology throughout Client Name ("The Company"). This policy is to ensure that all employees are using AI technological systems in a manner that complies with legal and regulatory standards and upholds the company's morals and values.

Scope

This policy applies to all of the company's employees, contractors, & partners, who utilize **approved** AI technical systems (see list of approved AI technical systems below). AI technical systems are to be approved by the ROLE.

Terms

Artificial Intelligence (AI): The ability of machines or software to learn, think, or autonomously carry out tasks normally associated with human intelligence, which can be used as tools to assist in developing solutions.

Approved AI Technical System(s): Software, platforms, and any other form of Artificial Intelligence (AI) system that the ROLE has approved for use for the company.

Protected Health Information (PHI): Protected medical information as defined by the Department of Health and Human Services.

Personally Identifiable Information (PII): Protected personal information as defined by the Department of Defense.

Approved AI Technical System(s)

The list of approved AI technical systems include:

Name	Website	Approved Use
ChatGPT (Example)	www.chat.openai.com	Idea formulation, general content creation

Security Policy # - Artificial Intelligence (AI) Acceptable Use Policy

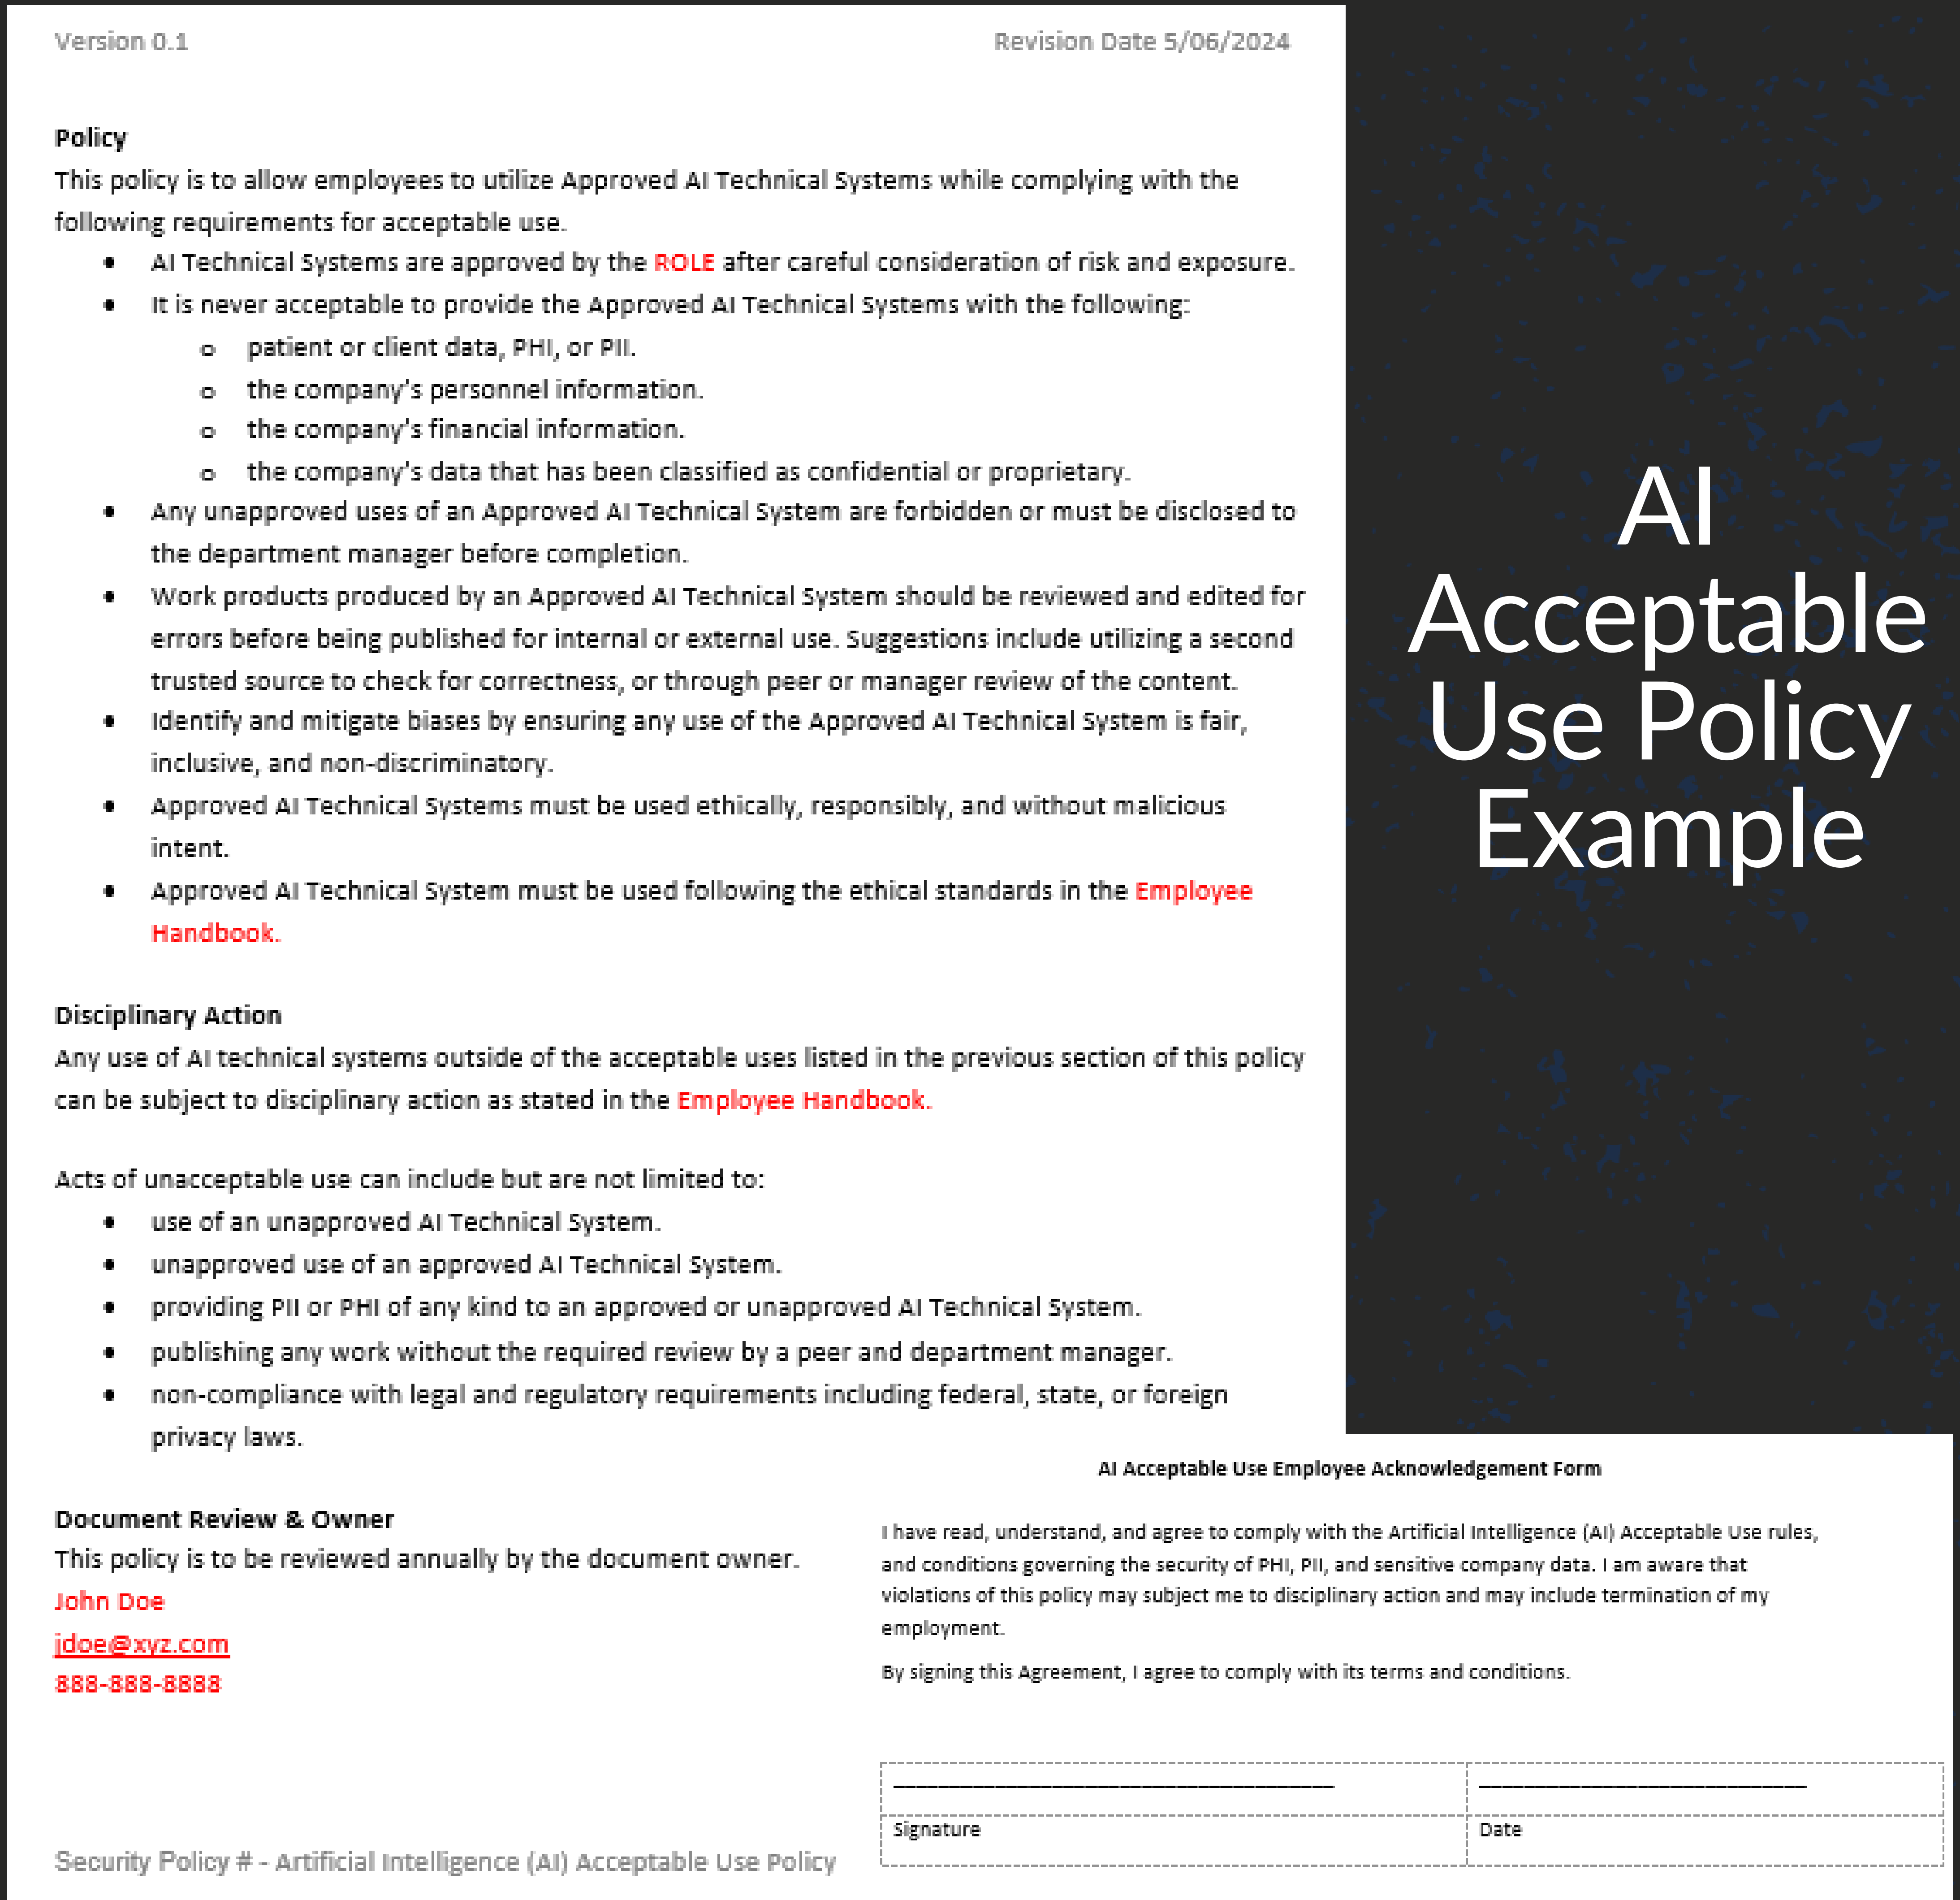

Version 0.1 Revision Date 5/06/2024

Policy

This policy is to allow employees to utilize Approved AI Technical Systems while complying with the following requirements for acceptable use.

- AI Technical Systems are approved by the ROLE after careful consideration of risk and exposure.
- It is never acceptable to provide the Approved AI Technical Systems with the following:
  - patient or client data, PHI, or PII.
  - the company's personnel information.
  - the company's financial information.
  - the company's data that has been classified as confidential or proprietary.
- Any unapproved uses of an Approved AI Technical System are forbidden or must be disclosed to the department manager before completion.
- Work products produced by an Approved AI Technical System should be reviewed and edited for errors before being published for internal or external use. Suggestions include utilizing a second trusted source to check for correctness, or through peer or manager review of the content.
- Identify and mitigate biases by ensuring any use of the Approved AI Technical System is fair, inclusive, and non-discriminatory.
- Approved AI Technical Systems must be used ethically, responsibly, and without malicious intent.
- Approved AI Technical System must be used following the ethical standards in the Employee Handbook.

Disciplinary Action

Any use of AI technical systems outside of the acceptable uses listed in the previous section of this policy can be subject to disciplinary action as stated in the Employee Handbook.

Acts of unacceptable use can include but are not limited to:

- use of an unapproved AI Technical System.
- unapproved use of an approved AI Technical System.
- providing PII or PHI of any kind to an approved or unapproved AI Technical System.
- publishing any work without the required review by a peer and department manager.
- non-compliance with legal and regulatory requirements including federal, state, or foreign privacy laws.

Document Review & Owner

This policy is to be reviewed annually by the document owner.

John Doe
jdoe@xyz.com
888-888-8888

Security Policy # - Artificial Intelligence (AI) Acceptable Use Policy

AI Acceptable Use Employee Acknowledgement Form

I have read, understand, and agree to comply with the Artificial Intelligence (AI) Acceptable Use rules, and conditions governing the security of PHI, PII, and sensitive company data. I am aware that violations of this policy may subject me to disciplinary action and may include termination of my employment.

By signing this Agreement, I agree to comply with its terms and conditions.

______________________________	______________________
Signature	Date

DLP Policy Example

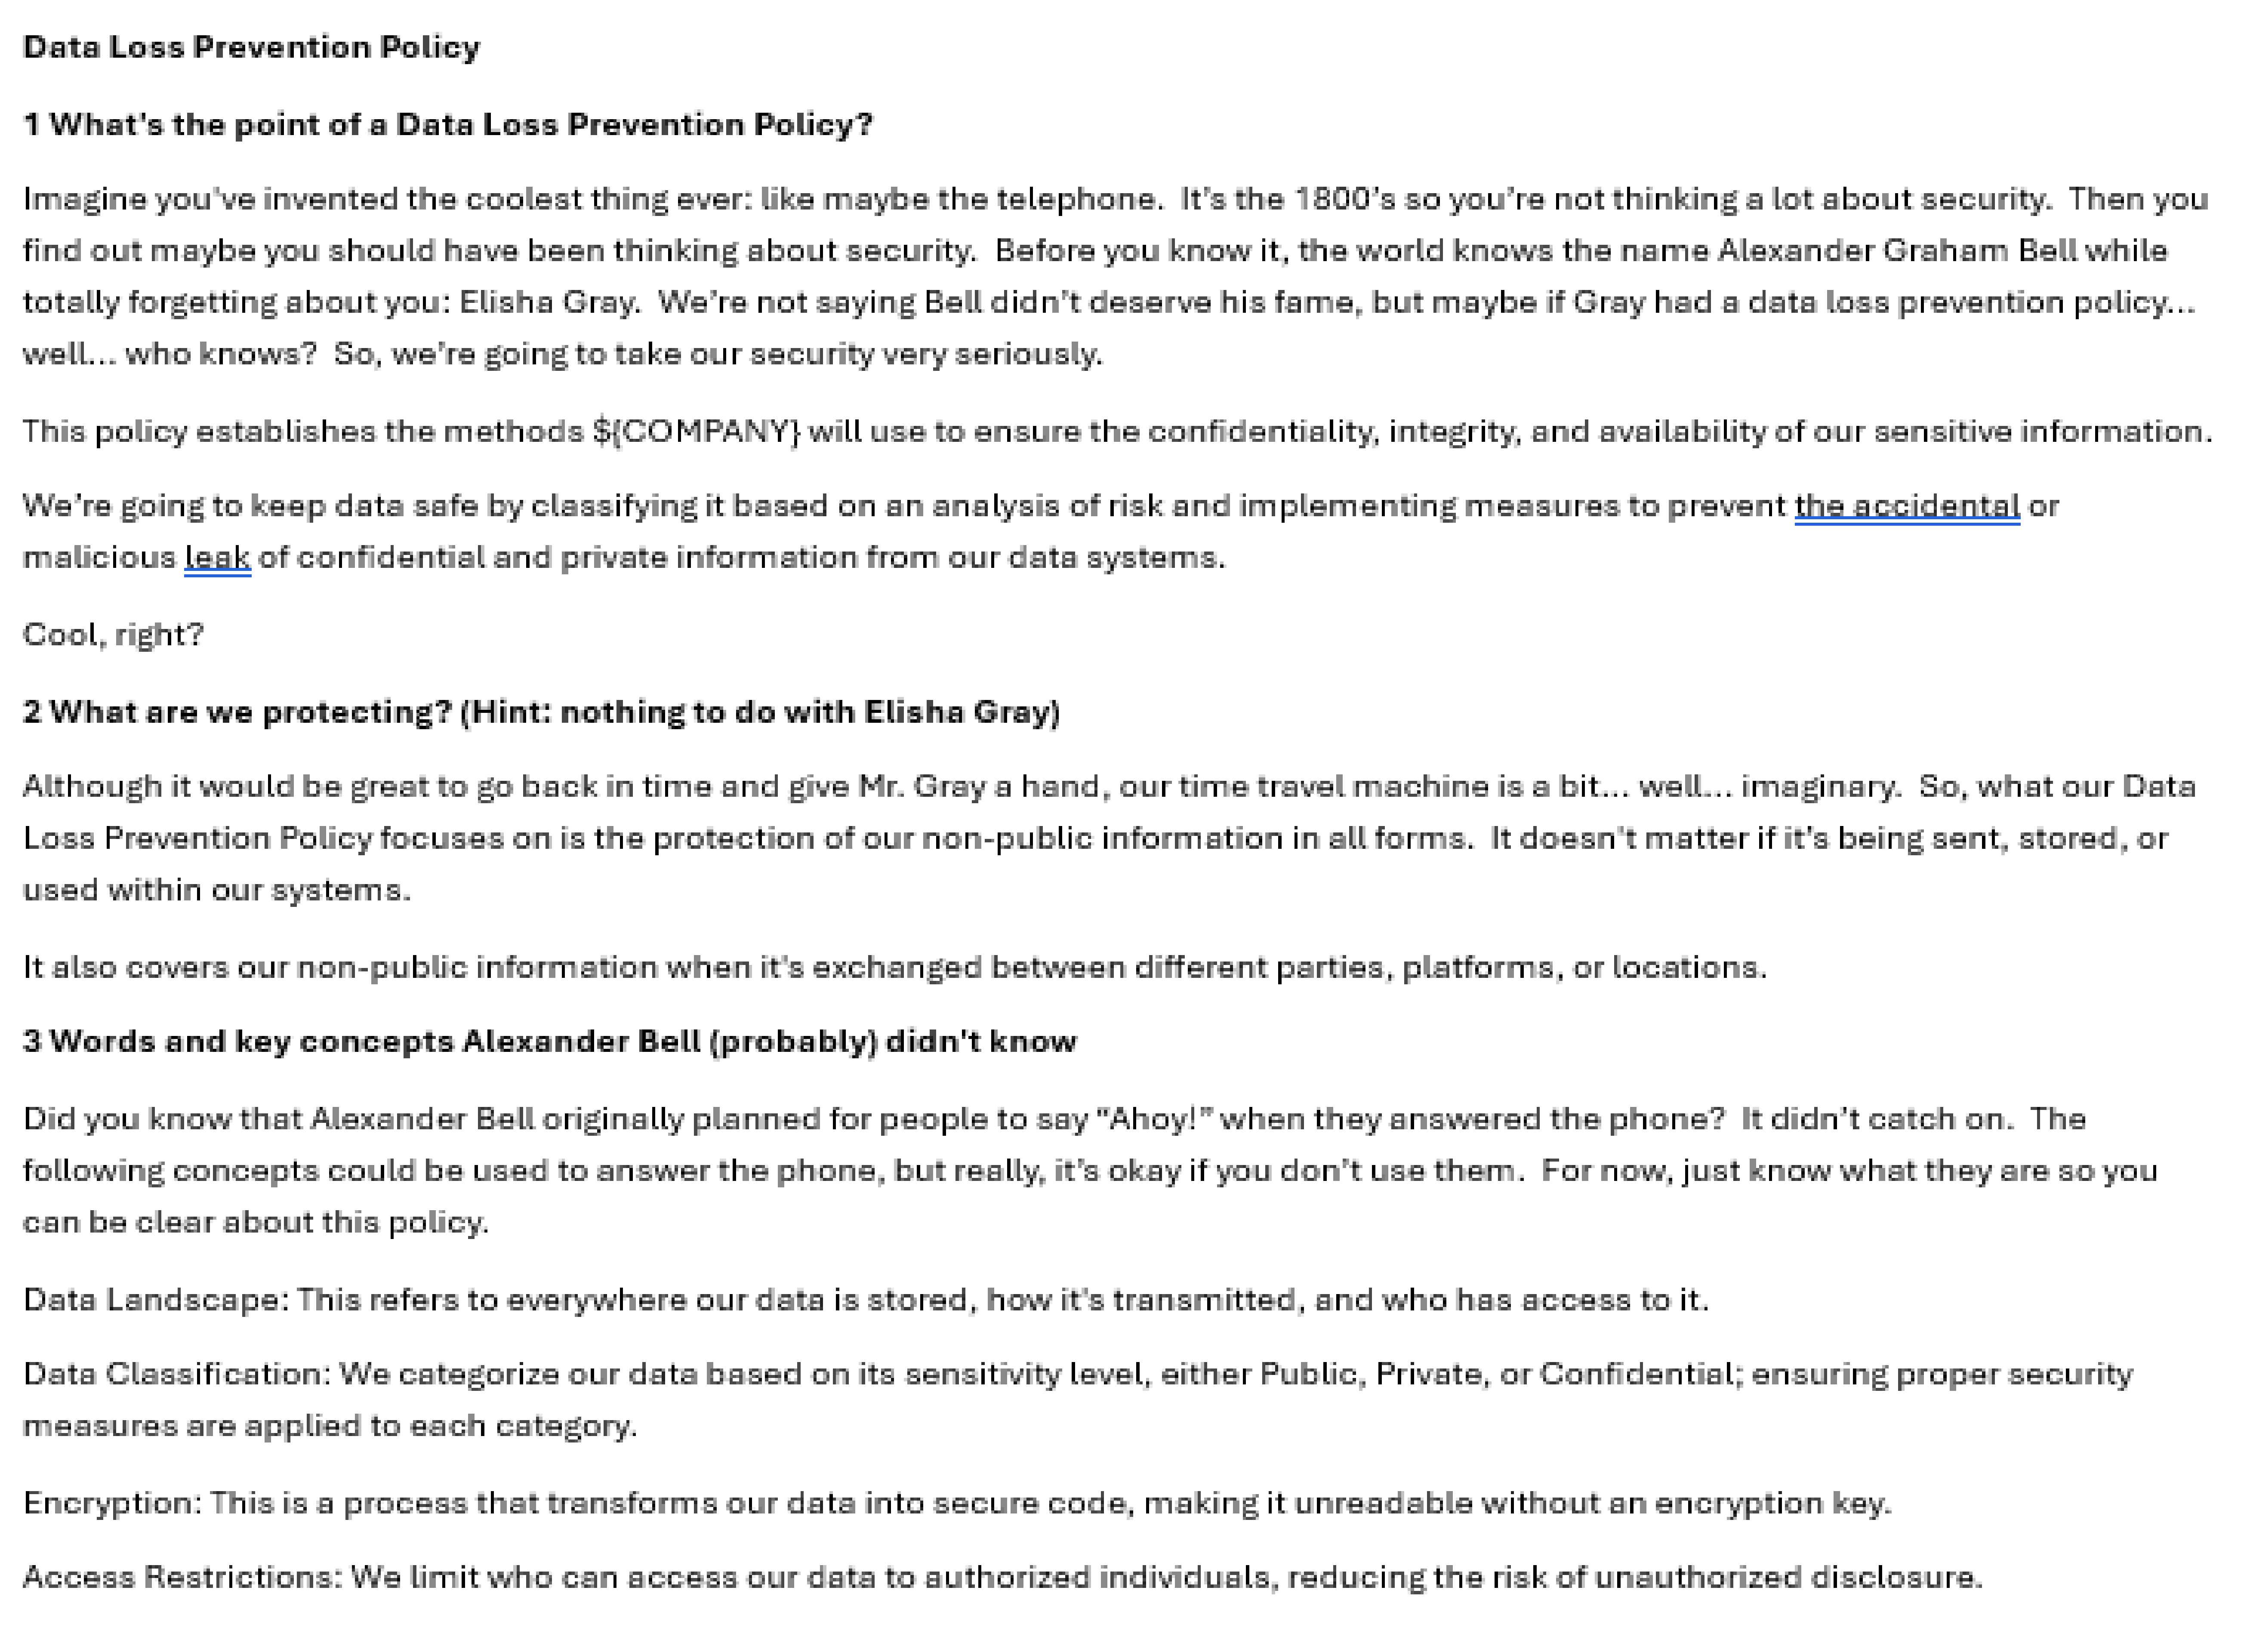

Data Loss Prevention Policy

1 What's the point of a Data Loss Prevention Policy?

Imagine you've invented the coolest thing ever: like maybe the telephone. It's the 1800's so you're not thinking a lot about security. Then you find out maybe you should have been thinking about security. Before you know it, the world knows the name Alexander Graham Bell while totally forgetting about you: Elisha Gray. We're not saying Bell didn't deserve his fame, but maybe if Gray had a data loss prevention policy... well... who knows? So, we're going to take our security very seriously.

This policy establishes the methods ${COMPANY} will use to ensure the confidentiality, integrity, and availability of our sensitive information.

We're going to keep data safe by classifying it based on an analysis of risk and implementing measures to prevent the accidental or malicious leak of confidential and private information from our data systems.

Cool, right?

2 What are we protecting? (Hint: nothing to do with Elisha Gray)

Although it would be great to go back in time and give Mr. Gray a hand, our time travel machine is a bit... well... imaginary. So, what our Data Loss Prevention Policy focuses on is the protection of our non-public information in all forms. It doesn't matter if it's being sent, stored, or used within our systems.

It also covers our non-public information when it's exchanged between different parties, platforms, or locations.

3 Words and key concepts Alexander Bell (probably) didn't know

Did you know that Alexander Bell originally planned for people to say "Ahoy!" when they answered the phone? It didn't catch on. The following concepts could be used to answer the phone, but really, it's okay if you don't use them. For now, just know what they are so you can be clear about this policy.

Data Landscape: This refers to everywhere our data is stored, how it's transmitted, and who has access to it.

Data Classification: We categorize our data based on its sensitivity level, either Public, Private, or Confidential; ensuring proper security measures are applied to each category.

Encryption: This is a process that transforms our data into secure code, making it unreadable without an encryption key.

Access Restrictions: We limit who can access our data to authorized individuals, reducing the risk of unauthorized disclosure.

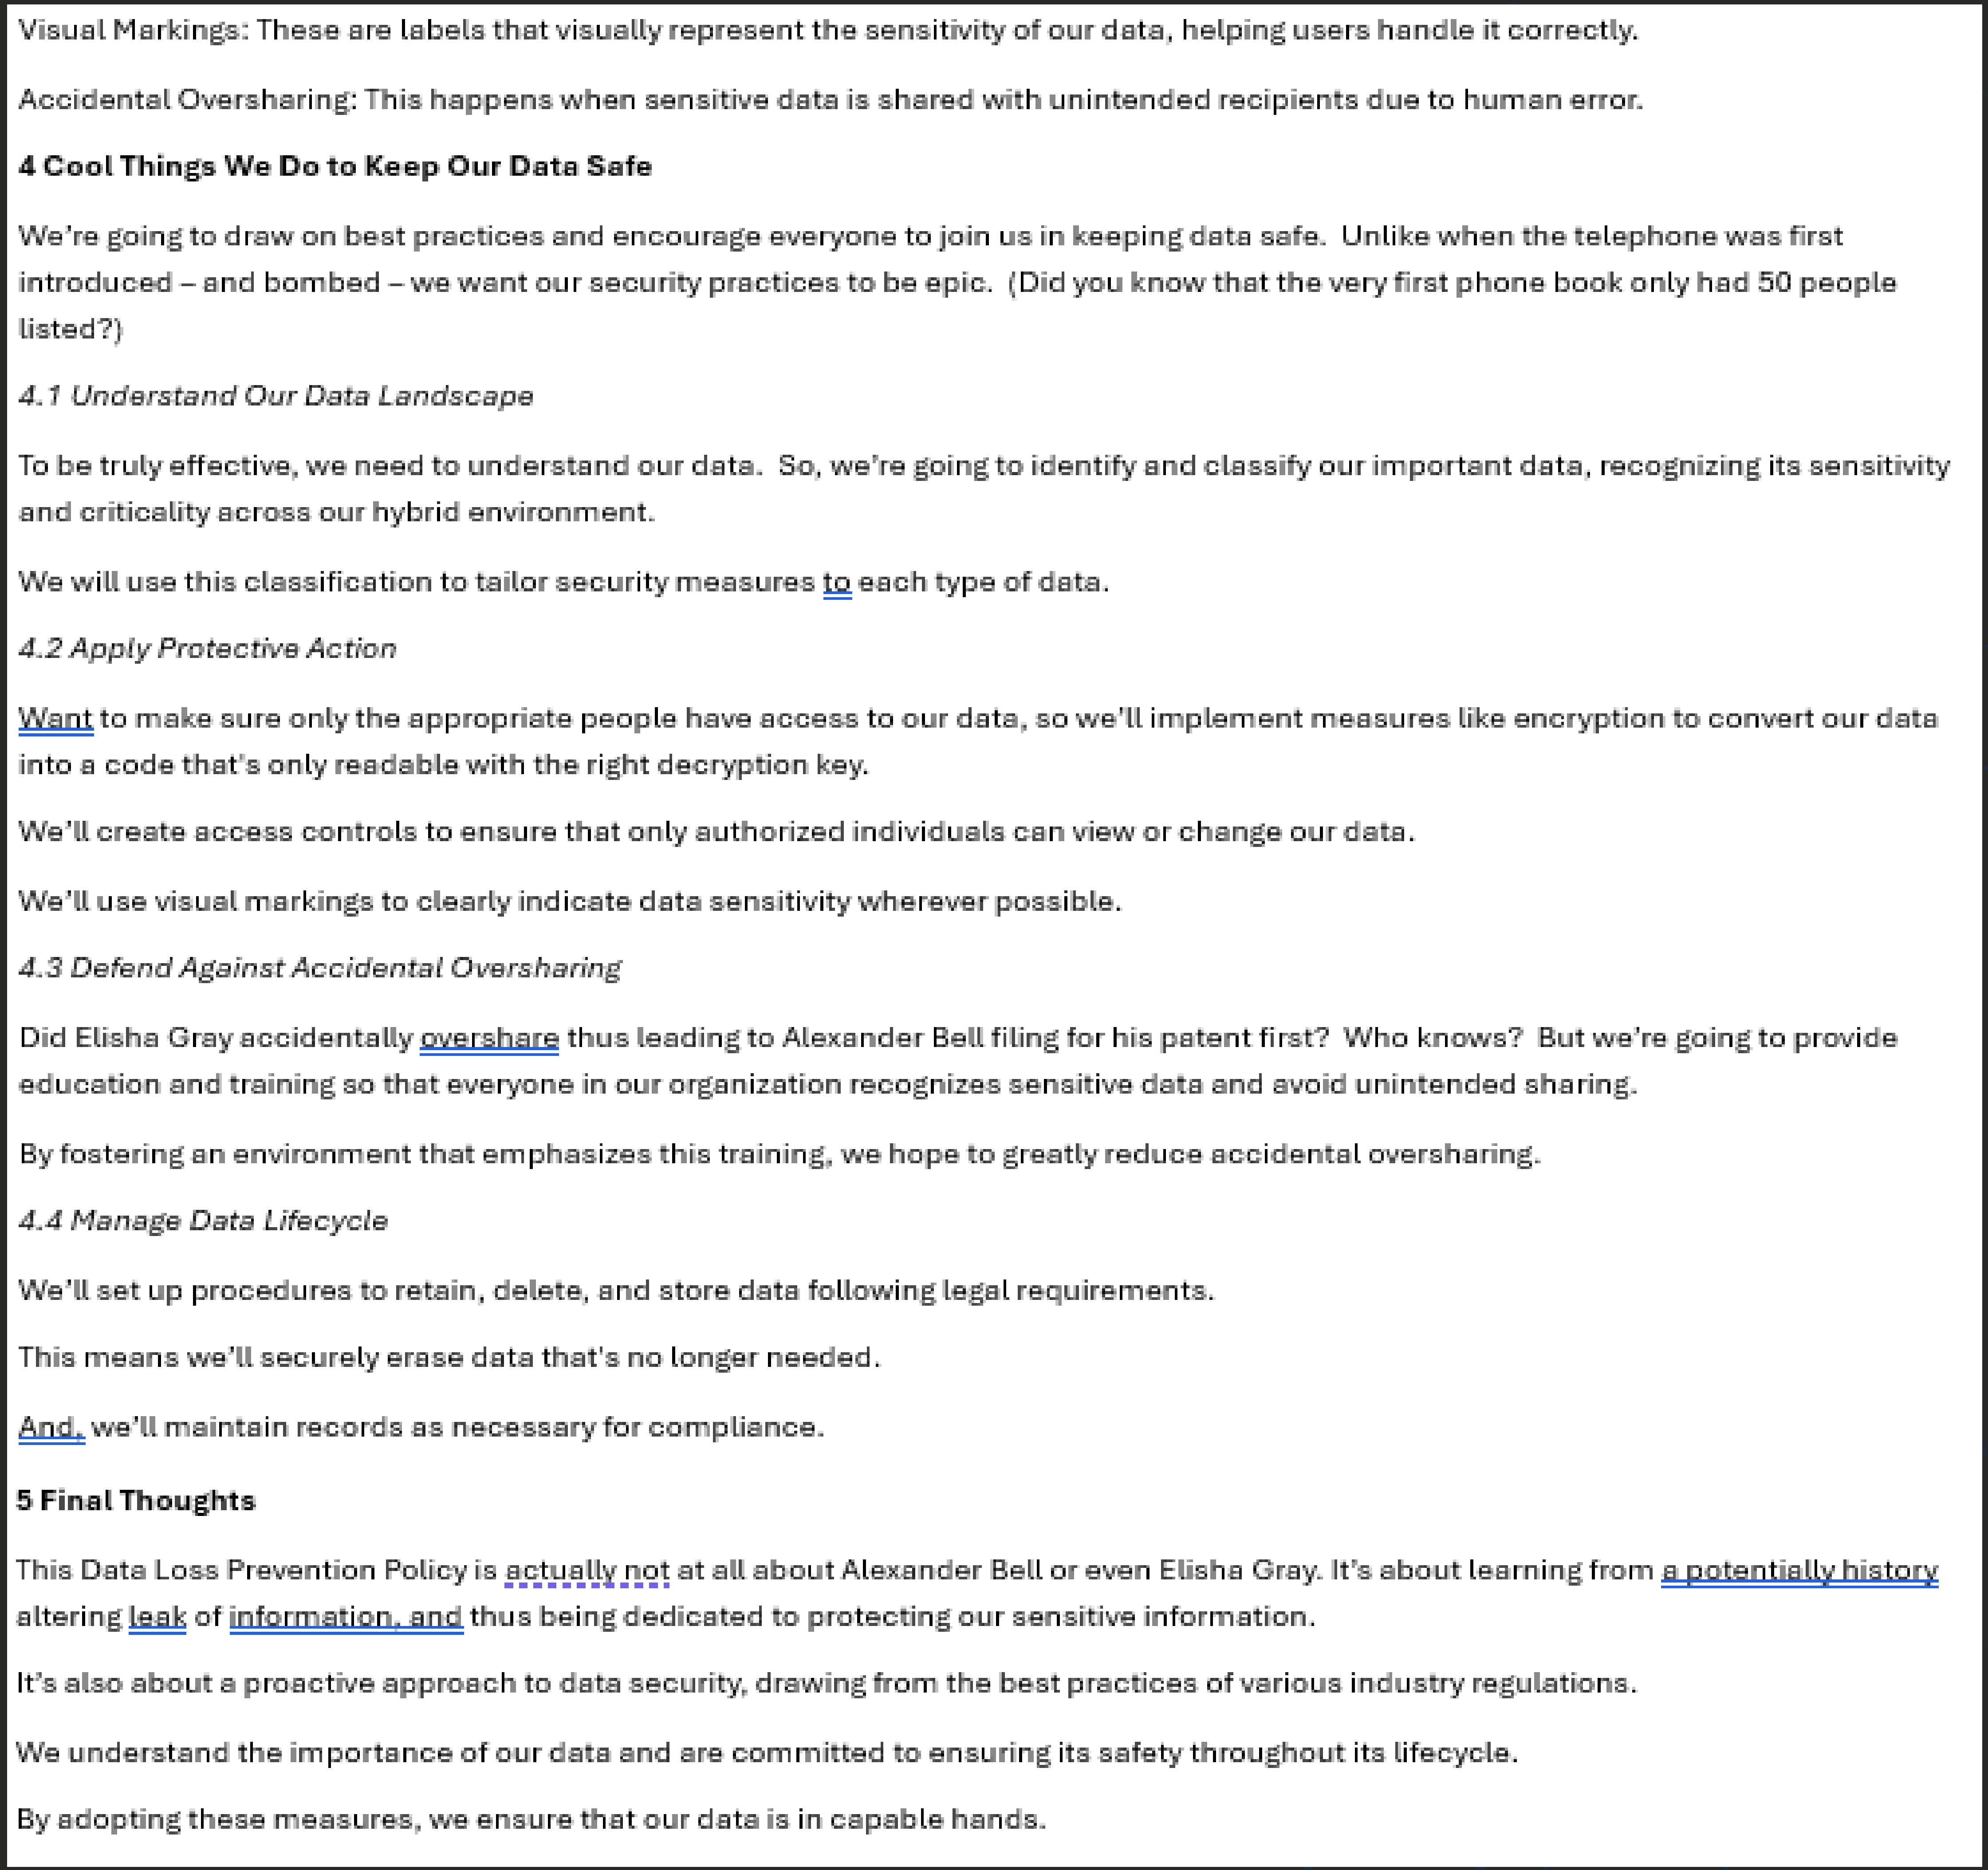

Visual Markings: These are labels that visually represent the sensitivity of our data, helping users handle it correctly.

Accidental Oversharing: This happens when sensitive data is shared with unintended recipients due to human error.

4 Cool Things We Do to Keep Our Data Safe

We're going to draw on best practices and encourage everyone to join us in keeping data safe. Unlike when the telephone was first introduced – and bombed – we want our security practices to be epic. (Did you know that the very first phone book only had 50 people listed?)

4.1 Understand Our Data Landscape

To be truly effective, we need to understand our data. So, we're going to identify and classify our important data, recognizing its sensitivity and criticality across our hybrid environment.

We will use this classification to tailor security measures to each type of data.

4.2 Apply Protective Action

Want to make sure only the appropriate people have access to our data, so we'll implement measures like encryption to convert our data into a code that's only readable with the right decryption key.

We'll create access controls to ensure that only authorized individuals can view or change our data.

We'll use visual markings to clearly indicate data sensitivity wherever possible.

4.3 Defend Against Accidental Oversharing

Did Elisha Gray accidentally overshare thus leading to Alexander Bell filing for his patent first? Who knows? But we're going to provide education and training so that everyone in our organization recognizes sensitive data and avoid unintended sharing.

By fostering an environment that emphasizes this training, we hope to greatly reduce accidental oversharing.

4.4 Manage Data Lifecycle

We'll set up procedures to retain, delete, and store data following legal requirements.

This means we'll securely erase data that's no longer needed.

And, we'll maintain records as necessary for compliance.

5 Final Thoughts

This Data Loss Prevention Policy is actually not at all about Alexander Bell or even Elisha Gray. It's about learning from a potentially history altering leak of information, and thus being dedicated to protecting our sensitive information.

It's also about a proactive approach to data security, drawing from the best practices of various industry regulations.

We understand the importance of our data and are committed to ensuring its safety throughout its lifecycle.

By adopting these measures, we ensure that our data is in capable hands.

What is DLP Software?

- Prevents data leaks & breaches
- Supports compliance (GDPR, SOX, GBLA, HIPPAA)
- Protects client records & payroll files

How DLP Software Works

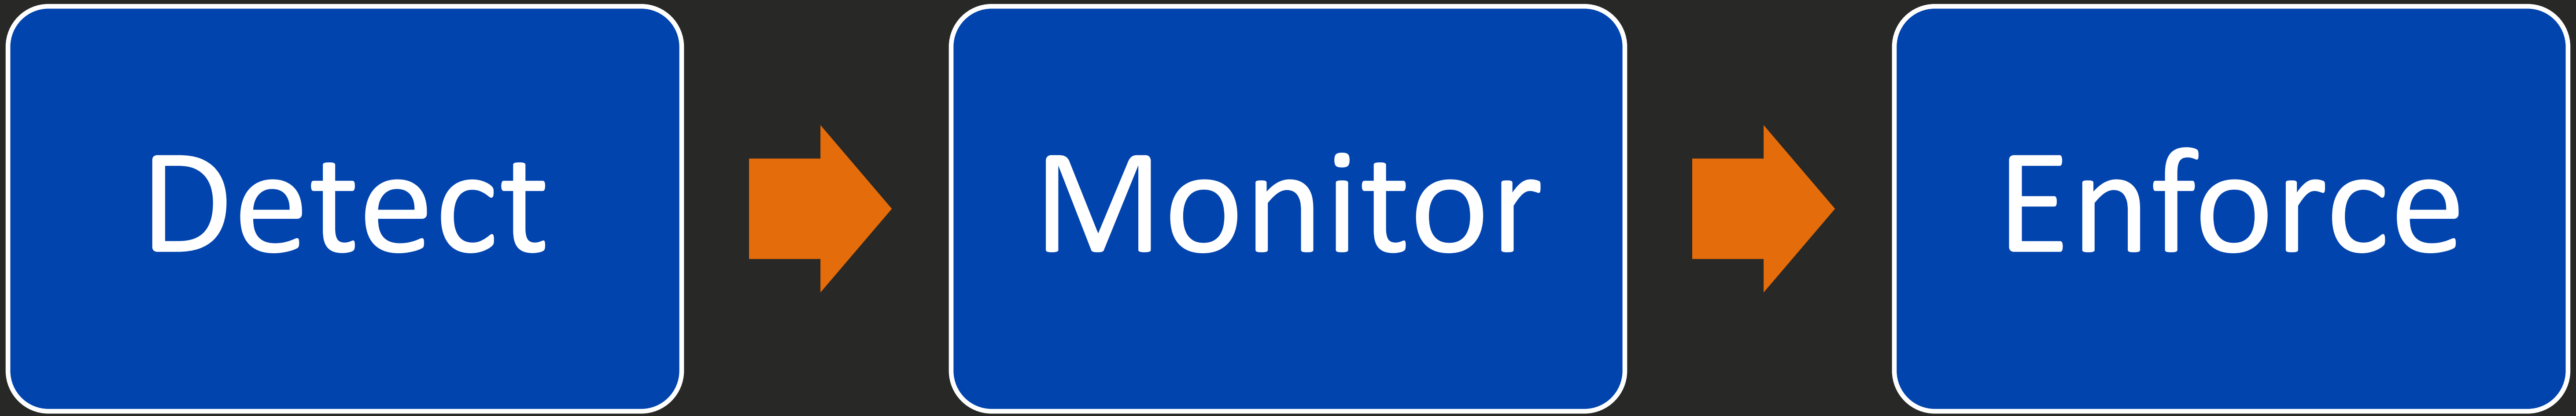

Detect

Monitor

Enforce

How DLP Software Works

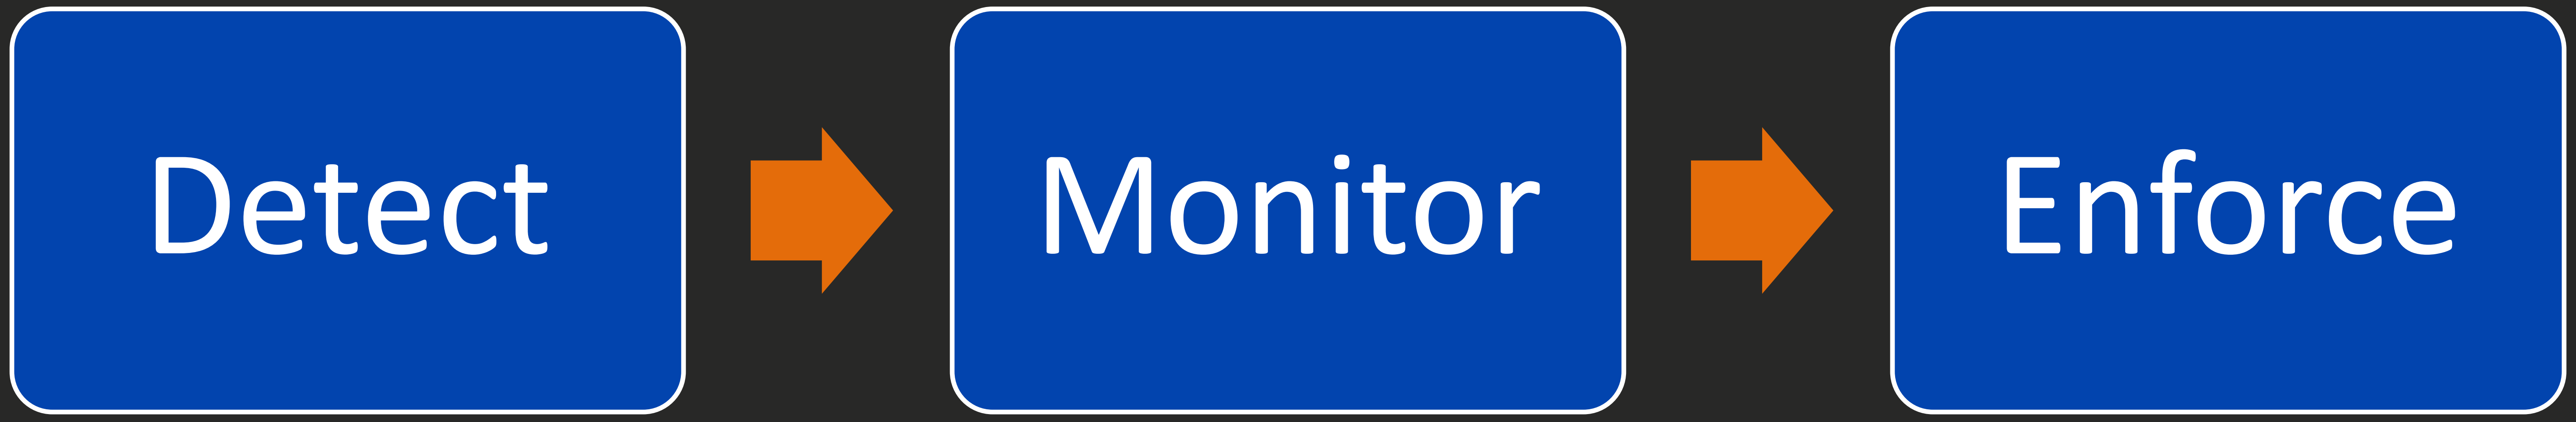

Detect → Monitor → Enforce

Example:

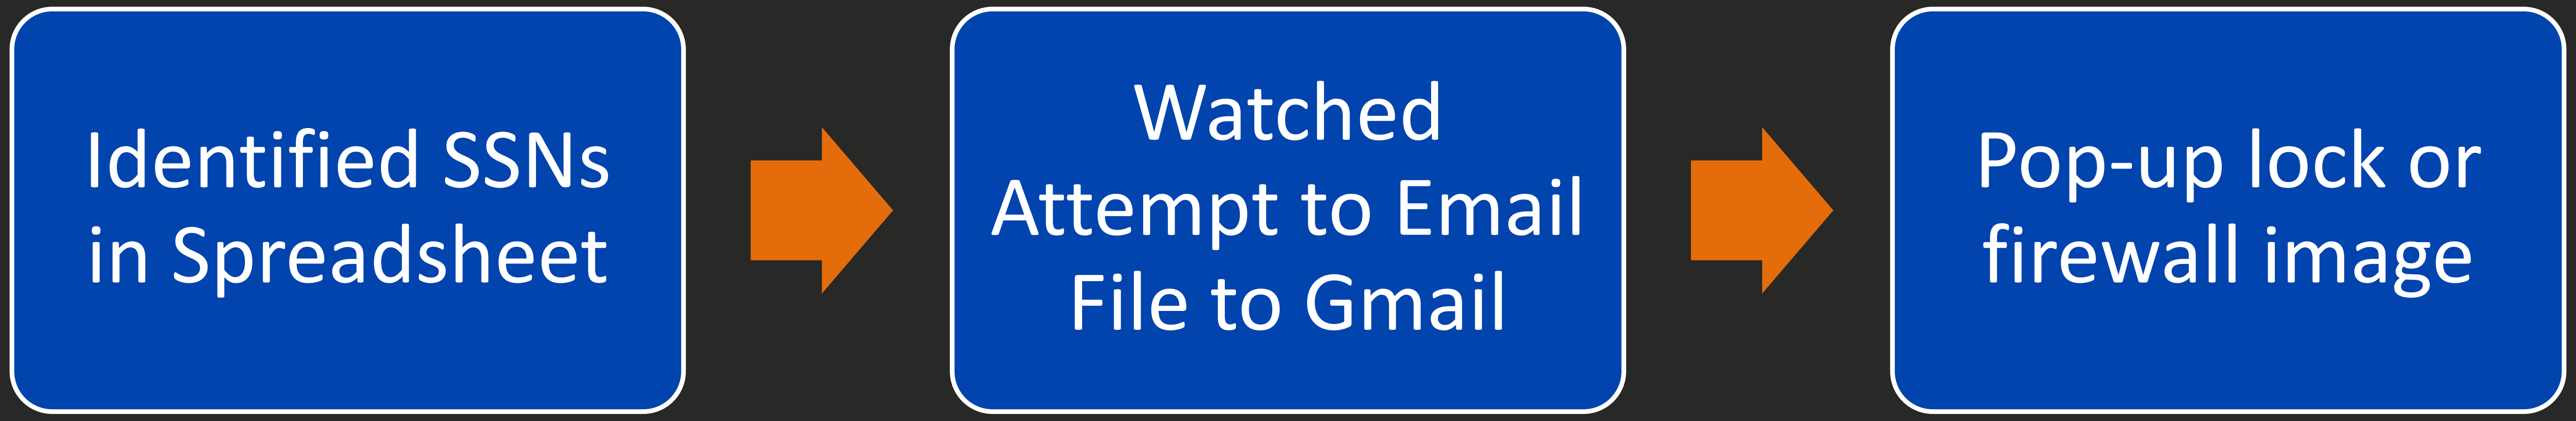

Identified SSNs in Spreadsheet → Watched Attempt to Email File to Gmail → Pop-up lock or firewall image

Types of DLP

Network DLP
Monitors email & file transfers

Endpoint DLP
Protects desktops, USB ports

Cloud DLP
Secure QuickBooks, OneDrive, M365

DLP Tools to Know

Purview

CYBERHAVEN

FORCEPOINT

Symantec™
by Broadcom

DLP Benefits & Challenges

✓ Benefits

- Avoid client data breaches
- Stay GBLA/SOX/PCI
- Visibility into data movement

🧩 Challenges

- Initial setup complexity
- Employee pushback
- Integration with finance tools

DLP Best Practices for Financial Teams

- ❑ Define sensitive data: PII, tax docs, financials
- ❑ Review rules regularly (GBLA, SOX updates)
- ❑ Train finance staff on secure file sharing
- ❑ Link DLP with email, file servers, and QuickBooks

The Browser Blind Spot

What DLP Tools Can’t Catch (and What You Can Do)

- ✗ Traditional DLP often can’t monitor browser form submissions (e.g., ChatGPT)
- ✗ Most DLP tools don’t see encrypted web traffic
- ✗ Copy/paste into AI sites often bypasses controls
- ✗ Endpoint DLP (Microsoft Purview, CoSoSys, etc.) or DNS filtering can help

Responsible Use Framework

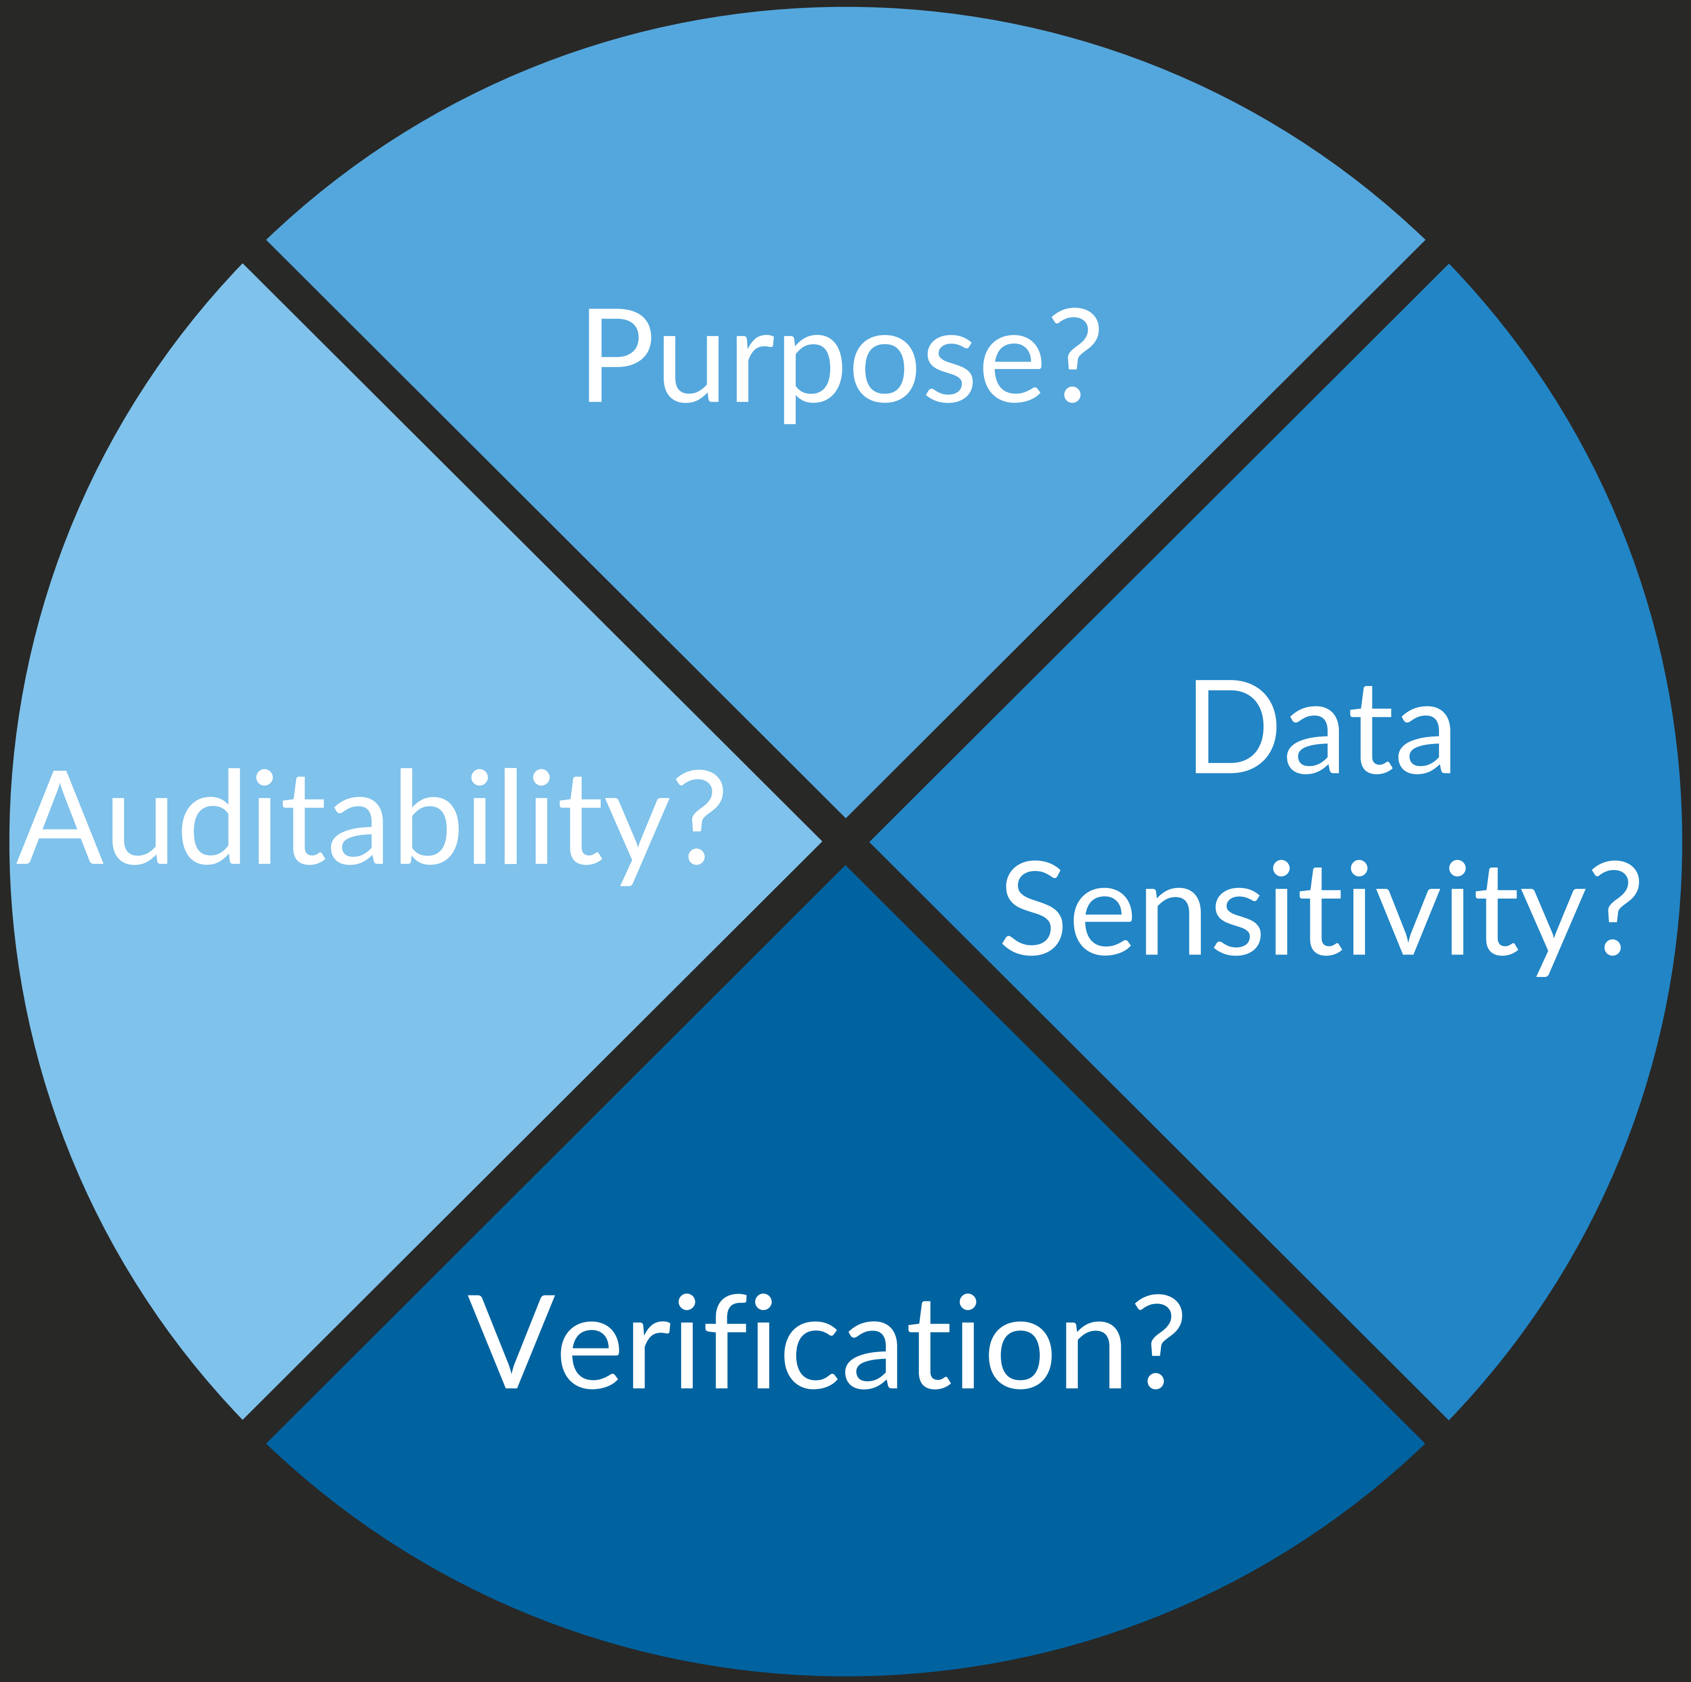

Purpose?

Data Sensitivity?

Verification?

Auditability?

Cyber Insurance - Your Last Line of Defense

Even if it’s your SaaS vendor’s breach, you could still be liable.
Most policies don’t cover AI misuse by default.

- Key questions to ask your broker:
  - Are third-party breaches covered?
  - Is AI usage excluded?
  - Are regulatory fines included?

Need these Resources for Your Business?

✓ AI Acceptable Use Policy

✓ (Template)

✓ DLP Policy (Template)

✓ MITRE ATLAS overview link

AI risk checklist

Email: rafe@comtechnc.com

AI is powerful, but it doesn’t replace accountability.
Use it wisely, document decisions, stay in control.

Thank You!

rafe@comtechnc.com 336-338-7328