



Forvis Mazars Cyber Symposium 2025  
**Agents: From Framework to Fieldwork**

October 2025

# Introduction

## **Ray Baxter**

Director | IT Risk & Compliance

---

Ray.Baxter@us.forvismazars.com



## **Sean Andrews**

Manager | IT Risk & Compliance

---

Sean.Andrews@us.forvismazars.com





# Learning Objectives

1. How can frameworks be applied to AI governance?
2. What steps can you take to reduce the risk associated with AI systems?
3. What is agentic AI?
4. How can I use agentic AI?
5. How can I begin piloting agentic systems?





1. AI Governance & Risk
2. Automation & Agents
3. What Is an Agent?
4. Examples of Agents
5. Frameworks & Fieldwork
6. Conclusion

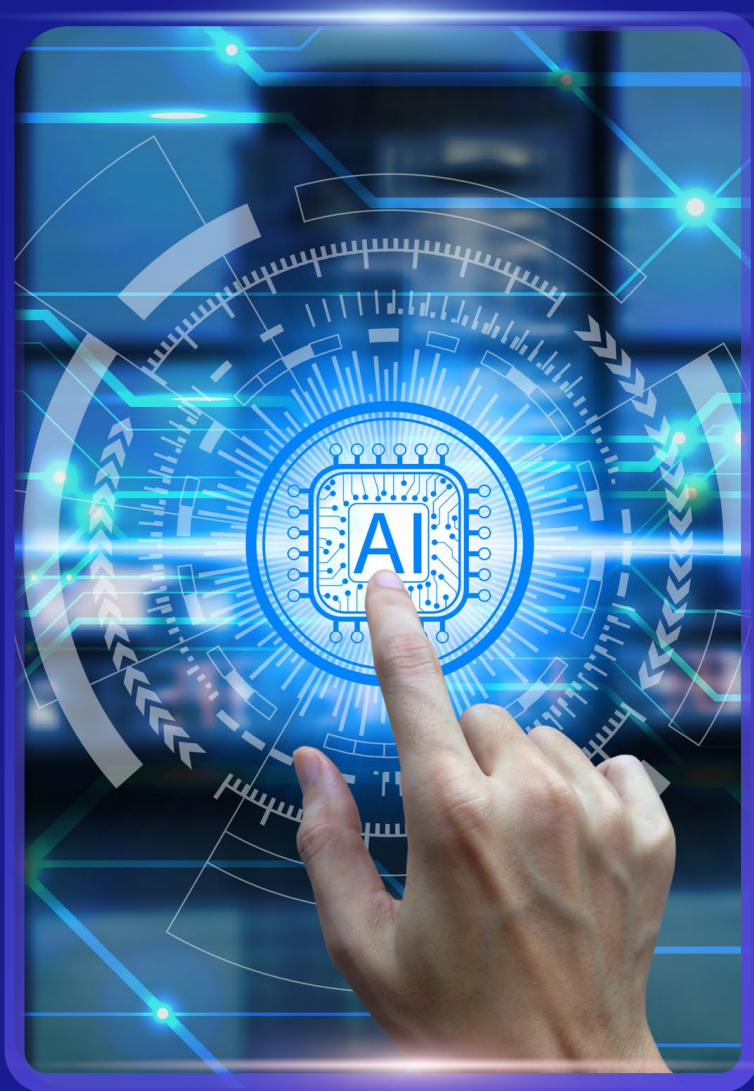




# 1

## AI Governance





# PROCESS & GOVERNANCE

Implementing AI successfully in an enterprise setting requires a structured AI process and a solid AI governance framework to ensure the technology is trustworthy, aligned with business goals, and compliant with ethical and legal standards. AI governance ensures AI is used responsibly, ethically, and compliantly, with clear oversight and control.

- ➔ Oversight
- ➔ Data Governance
- ➔ Risk Management & Compliance
- ➔ Ethical Guidelines & Principles
- ➔ Monitoring & Reporting



# AI GOVERNANCE

## EVALUATE

Perform evaluations of AI systems, including performance assessments, bias detection, and validation of model accuracy and fairness.



## DEVELOP

Design customized AI solutions tailored to their specific business needs.



## IMPLEMENT

Manage the deployment of AI systems, considering smooth integration with existing systems and controls to enable ethical, transparent, and compliant use of AI technology.



## IDENTIFY

Identify areas where AI can create value, optimize processes, and drive innovation, providing strategic insights and feasibility of different solutions.



## AUDIT

Audit AI systems for compliance with relevant regulations, effectiveness, and potential risks, providing recommendations for improvements.



## SECURE

Identify and manage AI-specific threats, enhance data security, and test access controls to protect sensitive information used by AI Systems.



## AI GOVERNANCE

Create and implement robust governance frameworks that align with risk-based best practices and address relevant or upcoming regulations.

Institutions of all sizes are looking to harness the power of AI to help increase efficiency, enhance processes, & empower their teams. At Forvis Mazars, we help organizations revolutionize their business by developing tailored AI frameworks, infrastructure, & solutions in a structured & controlled manner for sustainable growth.

# Purpose of AI Governance

AI governance refers to the frameworks, policies, and guidelines that ensure artificial intelligence (AI) technologies are developed and used ethically and responsibly. AI governance is essential for ensuring that AI technologies are used in a way that benefits society while minimizing potential negative impacts.

## Accountability

- Clearly articulate business objectives and goals for implementing AI
- Establish clear data management policies and procedures
- Identifying potential risks associated with AI system, including technical, operational, and ethic
- Implementing strategies to mitigate risks and reduce their potential impact

## Transparency & Explainability

- Maintain detailed records of AI models, data sources, decision-making workflows, and system modification
- Regularly informing stakeholders, management, and employees about AI practices, associated risks, and benefits
- Offer training and educational resources to employees to enhance their grasp of AI systems and their implications

## Ethics & Fairness

- Formulate and incorporate ethical guidelines and principles into AI development/use processes
- Perform periodic assessments to detect and address any biases within AI systems
- Encourage diverse perspective within AI development teams to foster fairness and diminish biases



# Case Study – OpenEvidence

## AI health-care startup OpenEvidence raises funding from Sequoia at \$1 billion valuation

OpenEvidence®

New Q

### KEY POINTS

- AI startup OpenEvidence is raising a fresh round of capital from Sequoia to scale its chatbot for doctors.
- OpenEvidence is an AI copilot for physicians that helps them make better decisions at the point of care. The company claims it's already being used by a quarter of doctors in the U.S.
- The new \$75 million cash injection, which has not been previously reported, values OpenEvidence at \$1 billion, the two companies told CNBC.



Medical technology, AI technology is utilized by doctors for diagnosing increasing the accuracy of patient treatments. Medical research and development innovation technology to improve patient health.

Picasso / iStock | Getty Images

headache, sore throat, cough, sore chest, yellow mucus, congestion

Expanded question: What could be the possible diagnosis for symptoms including headache, sore throat, cough, sore chest, yellow mucus, and congestion?

### Most Likely Diagnoses:

1. **Common Cold:** The symptoms of headache, sore throat, cough, sore chest, yellow mucus, and congestion are consistent with the common cold, a benign, self-limited viral illness. The American College of Physicians and the Centers for Disease Control and Prevention note that the common cold can present with these symptoms and typically resolves without the need for antibiotics.<sup>[1]</sup>
2. **Acute Bronchitis:** Acute bronchitis, characterized by an acute cough and inflammation of the trachea and lower airways, can present with cough, sore chest, and yellow mucus. The American College of Chest Physicians emphasizes that acute bronchitis is a common clinical condition and is usually self-limiting.<sup>[2]</sup>
3. **Acute Rhinosinusitis:** The presence of yellow mucus, congestion, and headache suggests acute rhinosinusitis. The American Academy of Otolaryngology-Head and Neck Surgery states that acute rhinosinusitis is diagnosed with purulent nasal drainage and nasal obstruction or facial pain.<sup>[3]</sup>
4. **Influenza:** Influenza can present with headache, sore throat, cough, and congestion. The CDC notes that influenza typically has an abrupt onset and can include these symptoms along with fever and malaise.<sup>[4]</sup>
5. **Allergic Rhinitis:** Allergic rhinitis can cause nasal congestion, rhinorrhea, and postnasal drainage, which may lead to a sore throat and cough. The American Academy of Family Physicians highlights that allergic rhinitis is associated with these symptoms and often occurs in response to allergens.<sup>[5]</sup>

### Most Important Not to Miss Diagnoses:

1. **Pneumonia:** Pneumonia should be considered, especially if there is fever, productive cough, and pleuritic chest pain. The American College of Chest Physicians recommends considering pneumonia in patients with acute cough and suggests using clinical judgment and possibly chest radiographs to rule

Ask a follow-up question...



# AI Controls – GOVERN, MAP, MEASURE, & MANAGE.



<b>GOVERN 1</b>	<b>Policies, processes, procedures, and practices</b> across the organization related to the mapping, measuring, and managing of AI risks are in place, transparent, and implemented effectively.
<b>GOVERN 2</b>	<b>Accountability structures</b> are in place so that the appropriate teams and individuals are empowered, responsible, and <b>trained for mapping, measuring, and managing AI risks</b> .
<b>GOVERN 3</b>	Workforce <b>diversity, equity, inclusion, and accessibility</b> processes are prioritized in the mapping, measuring, and managing of AI risks throughout the lifecycle.
<b>GOVERN 4</b>	<b>Organizational teams</b> are committed to a culture that considers and <b>communicates AI risk</b> .
<b>GOVERN 5</b>	Processes are in place for <b>robust engagement with relevant AI actors</b> .
<b>GOVERN 6</b>	<b>Policies and procedures are in place to address AI risks</b> and benefits arising from third-party software and data and other supply chain issues.
<b>MAP 1</b>	<b>Context</b> is established and <b>understood</b> .
<b>MAP 2</b>	<b>Categorization</b> of the AI system is performed.
<b>MAP 3</b>	AI capabilities, targeted usage, goals, and expected benefits and costs compared with appropriate <b>benchmarks</b> are understood.
<b>MAP 4</b>	<b>Risks and benefits are mapped</b> for all components of the AI system including third-party software and data.
<b>MAP 5</b>	<b>Impacts to individuals, groups, communities, organizations, and society</b> are characterized.
<b>MEASURE 1</b>	Appropriate methods and <b>metrics are identified and applied</b> .
<b>MEASURE 2</b>	AI systems <b>are evaluated for trustworthy</b> characteristics.
<b>MEASURE 3</b>	Mechanisms for <b>tracking identified AI risks</b> over time are in place.
<b>MEASURE 4</b>	Feedback about <b>efficacy of measurement is gathered and assessed</b> .
<b>MANAGE 1</b>	<b>AI risks</b> based on assessments and other analytical output from the <b>MAP</b> and <b>MEASURE</b> functions are prioritized, responded to, and managed.
<b>MANAGE 2</b>	<b>Strategies</b> to maximize AI benefits and minimize negative impacts are planned, prepared, implemented, documented, and informed by input from relevant AI actors.
<b>MANAGE 3</b>	AI risks and benefits <b>from third-party entities</b> are managed.
<b>MANAGE 4</b>	Risk treatments, including response and recovery, and communication plans for the <b>identified and measured AI risks</b> are documented and monitored regularly.



# OpenEvidence – Sample Control Evaluation

**GOVERN 1.1:** Legal and regulatory requirements involving AI are understood, managed, and documented.

**GOVERN 1.2:** The characteristics of trustworthy AI are integrated into organizational policies, processes, procedures, and practices.

**GOVERN 4.1:** Organizational policies and practices are in place to foster a critical thinking and safety-first mindset in the design, development, deployment, and uses of AI systems to minimize potential negative impacts.

**GOVERN 6.1:** 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.

---

**MAP 1.4:** The business value or context of business use has been clearly defined or—in the case of assessing existing AI systems—re-evaluated.

**MAP 3.1:** Potential benefits of intended AI system functionality and performance are examined and documented.

**MAP 4.2:** Internal risk controls for components of the AI system, including third-party AI technologies, are identified and documented.

---

**MEASURE 2.2:** Evaluations involving human subjects meet applicable requirements (including human subject protection) and are representative of the relevant population.

**MEASURE 3.1:** Approaches, personnel, and documentation are in place to regularly identify and track existing, unanticipated, and emergent AI risks based on factors such as intended and actual performance in deployed contexts.

---

**MANAGE 1.1:** A determination is made as to whether the AI system achieves its intended purposes and stated objectives and whether its development or deployment should proceed.

**MANAGE 3.1:** AI risks and benefits from third-party resources are regularly monitored, and risk controls are applied and documented.

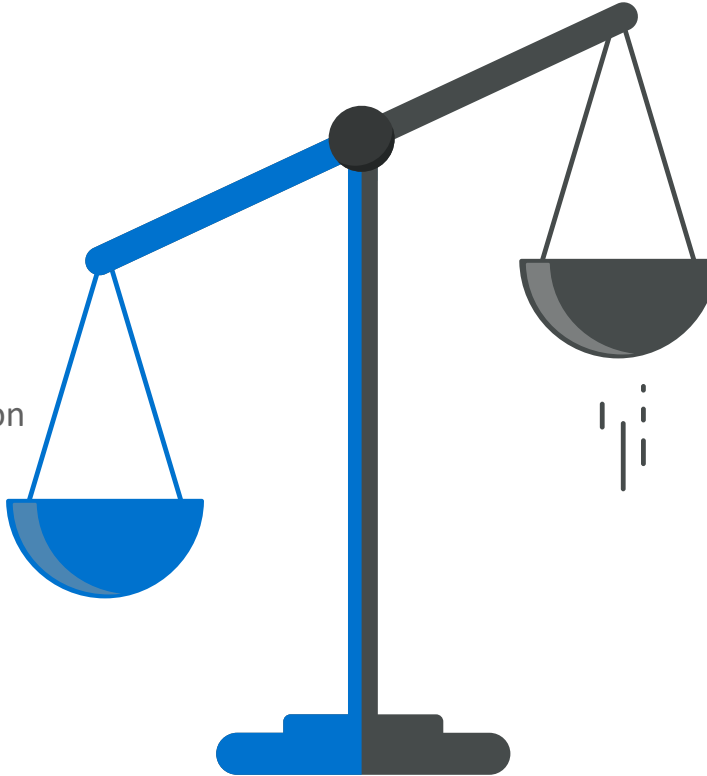
# OpenEvidence – Risk Management

**MANAGE 3.1:** AI **risks** and **benefits** from third-party resources are regularly monitored, and risk controls are applied and documented.

## Risk



- OpenEvidence is an experimental technology demonstrator. OpenEvidence does not provide medical advice, diagnosis, or treatment.
- User questions and other inputs on OpenEvidence are not covered by HIPAA. It is the responsibility of the user to ensure questions do not contain protected health information (PHI) or any information that violates the privacy of any person.
- Risk of over-reliance on tools like OpenEvidence, which could diminish critical thinking skills essential for clinical practice.
- Inconsistencies or errors in AI-generated results emphasize the need for human oversight to ensure accurate interpretation and adaptability to complex scenarios.



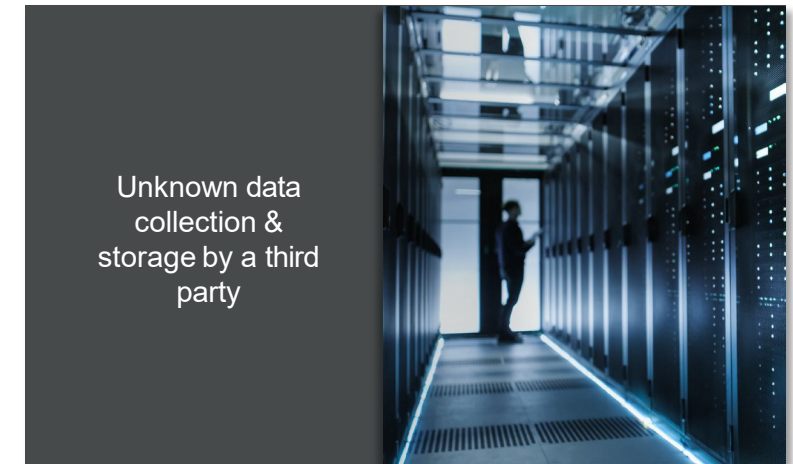
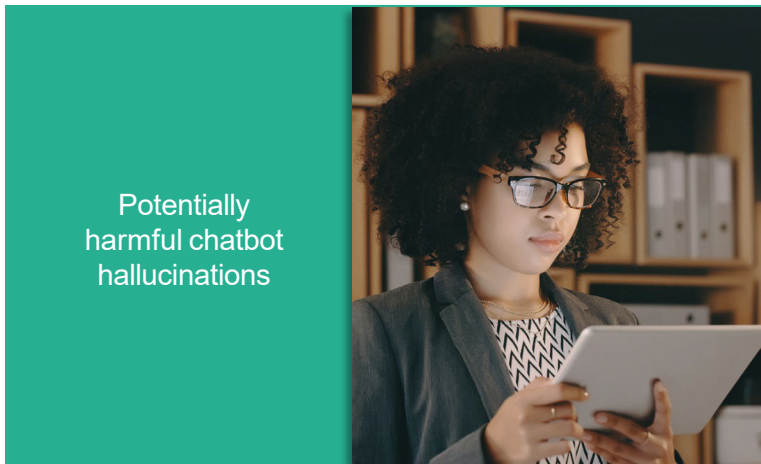
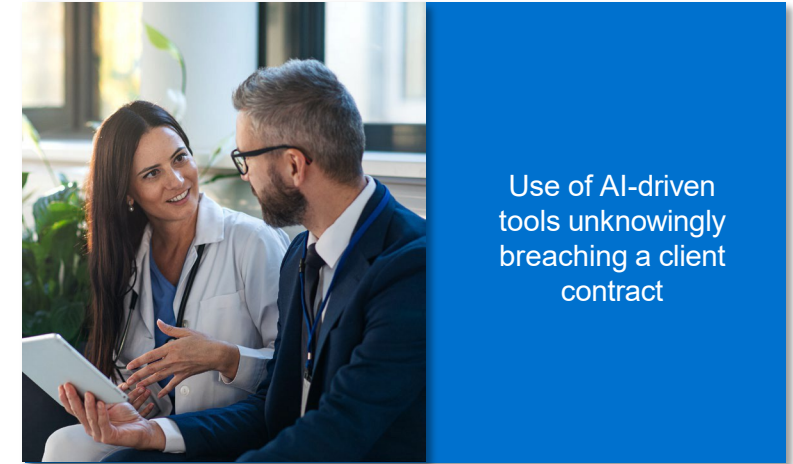
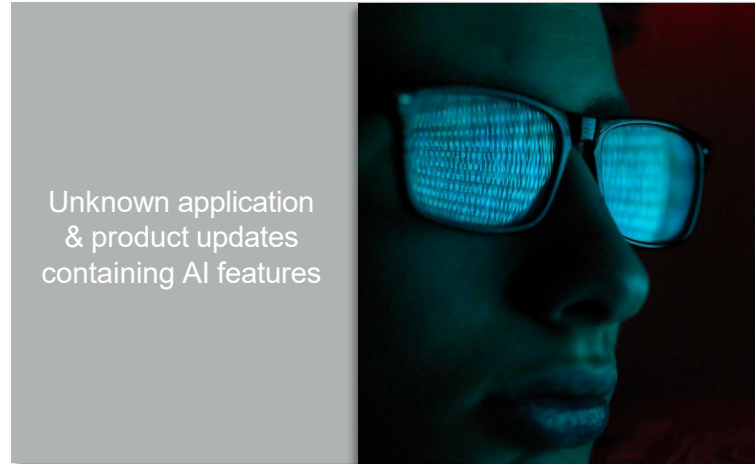
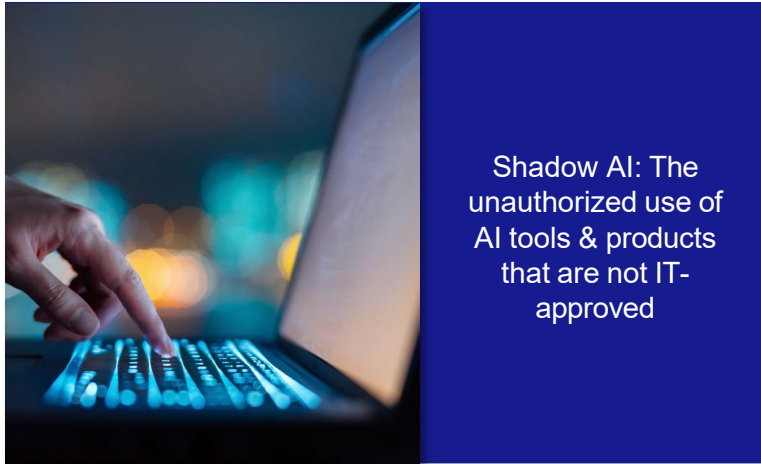
## Benefits

- OpenEvidence is what we always imagined when we used to talk about someone finally solving medical search.
- Aim to improve the accessibility and synthesis of medical literature.
- Used to make more evidenced-based decisions and improve patient outcomes.
- The platform facilitates quick access to clinically relevant content, enabling medical students to integrate learning with practice.
- Tool for medical education and clinical care, equipping future healthcare professionals to navigate an increasingly AI-integrated healthcare environment with precision and confidence.



# Hidden AI Risk

The use of AI presents additional third-party, legal, compliance, & reputational risks when being used or relied upon unknowingly.



# Mitigating AI Risk

Although the use of AI can present novel & complex risks, institutions can safeguard their users, data, & reputations with a robust risk mitigation strategy. Below are some examples of ways to mitigate AI risk.

## Third Party

- Review contracts with third-party vendors
- Stay informed about new features & model components of third-party models
- Ensure contracts contain clauses that protect the institution's data
- Understand third-party data sourcing, storage, & retention

## Operational / IT

- Consider a thorough review & approval process for all AI-driven tools, products, & partnerships
- Consider implementing internal controls & limitations for users
- Implement a rigorous training process to inform users about information security in the age of AI
- Maintain robust documentation of data dictionaries & lineage

## Governance

- Re-review all AI-driven tools & products for model definition & enhance model risk management
- Develop an agreed-upon framework with clearly defined leadership, roles, & responsibilities
- Consider second-line governance oversight to ensure first-line adherence to policies & frameworks
- Monitor all model inputs & outputs, & perform regular testing to detect potential bias or unwanted outcomes

\*Note: The mitigation strategies shown are among the most important but are not an exhaustive list.



# Implementing AI Governance

Implementing AI governance involves several key steps to ensure that AI systems are developed and used ethically and responsibly. By following these best practices, organizations can create a robust AI governance structure that fosters responsible AI development and use.

## Centralized Framework

- Analyze the existing AI governance procedures & evaluate the current framework for AI governance.
- Review roles & responsibilities across AI landscape.



## Data Governance

- Review the available documentation to evaluate data governance & quality, including lineage & traceability.
- Interview key stakeholders to understand the process for data collection, storage, access, monitoring, & disposal.



## Operational Strategy

- Examine the structure & responsibilities of the governance or steering committee, if any.
- Evaluate AI strategic roadmap, including development of policies & socialization to the organization.



## Third-Party Risk Management

- Evaluate third-party, open-source models & tools within the inventory & identify key risks.
- Assess third-party vendors for reliability & reputability.
- Review contracts with third-party vendors to identify potential data & privacy risks.



## Model Risk Cartography

- Review model identification & determination criteria & identify areas to enhance model determination to apply to AI models.
- Evaluate the model &/or AI use case inventory & the risk assessment criteria.



## Organizational Coverage

- Evaluate approach across overarching AI framework to encapsulate across the organization.
- Review AI process guides & evaluate existing process across the AI development lifecycle.



## IT Risk & Compliance

- Analyze AI IT Lifecycle for soundness of controls, including technical operations, cloud governance, internal & external data base integration,, cybersecurity, maintenance, etc.



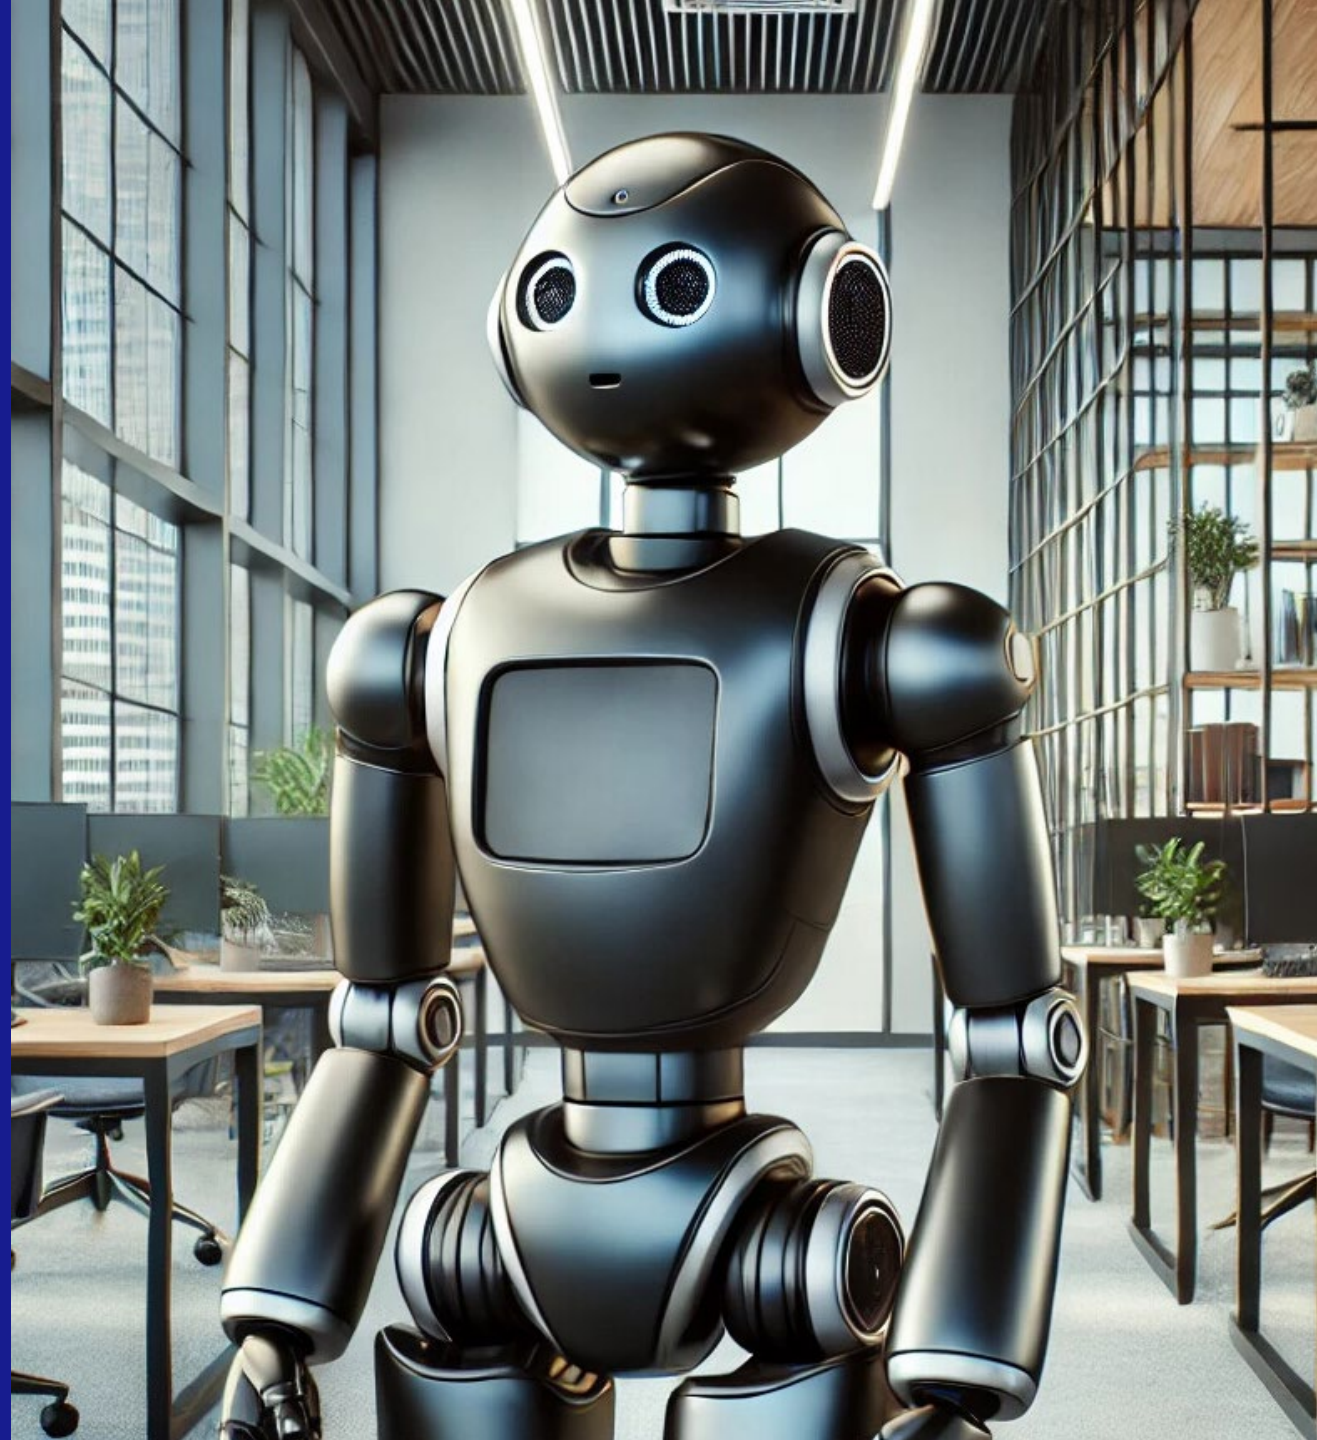
## Regulatory Compliance

- Assess existing AI governance process to ensure adherence to regulatory & legal requirements inclusive of reporting & monitoring AI models.



# 2

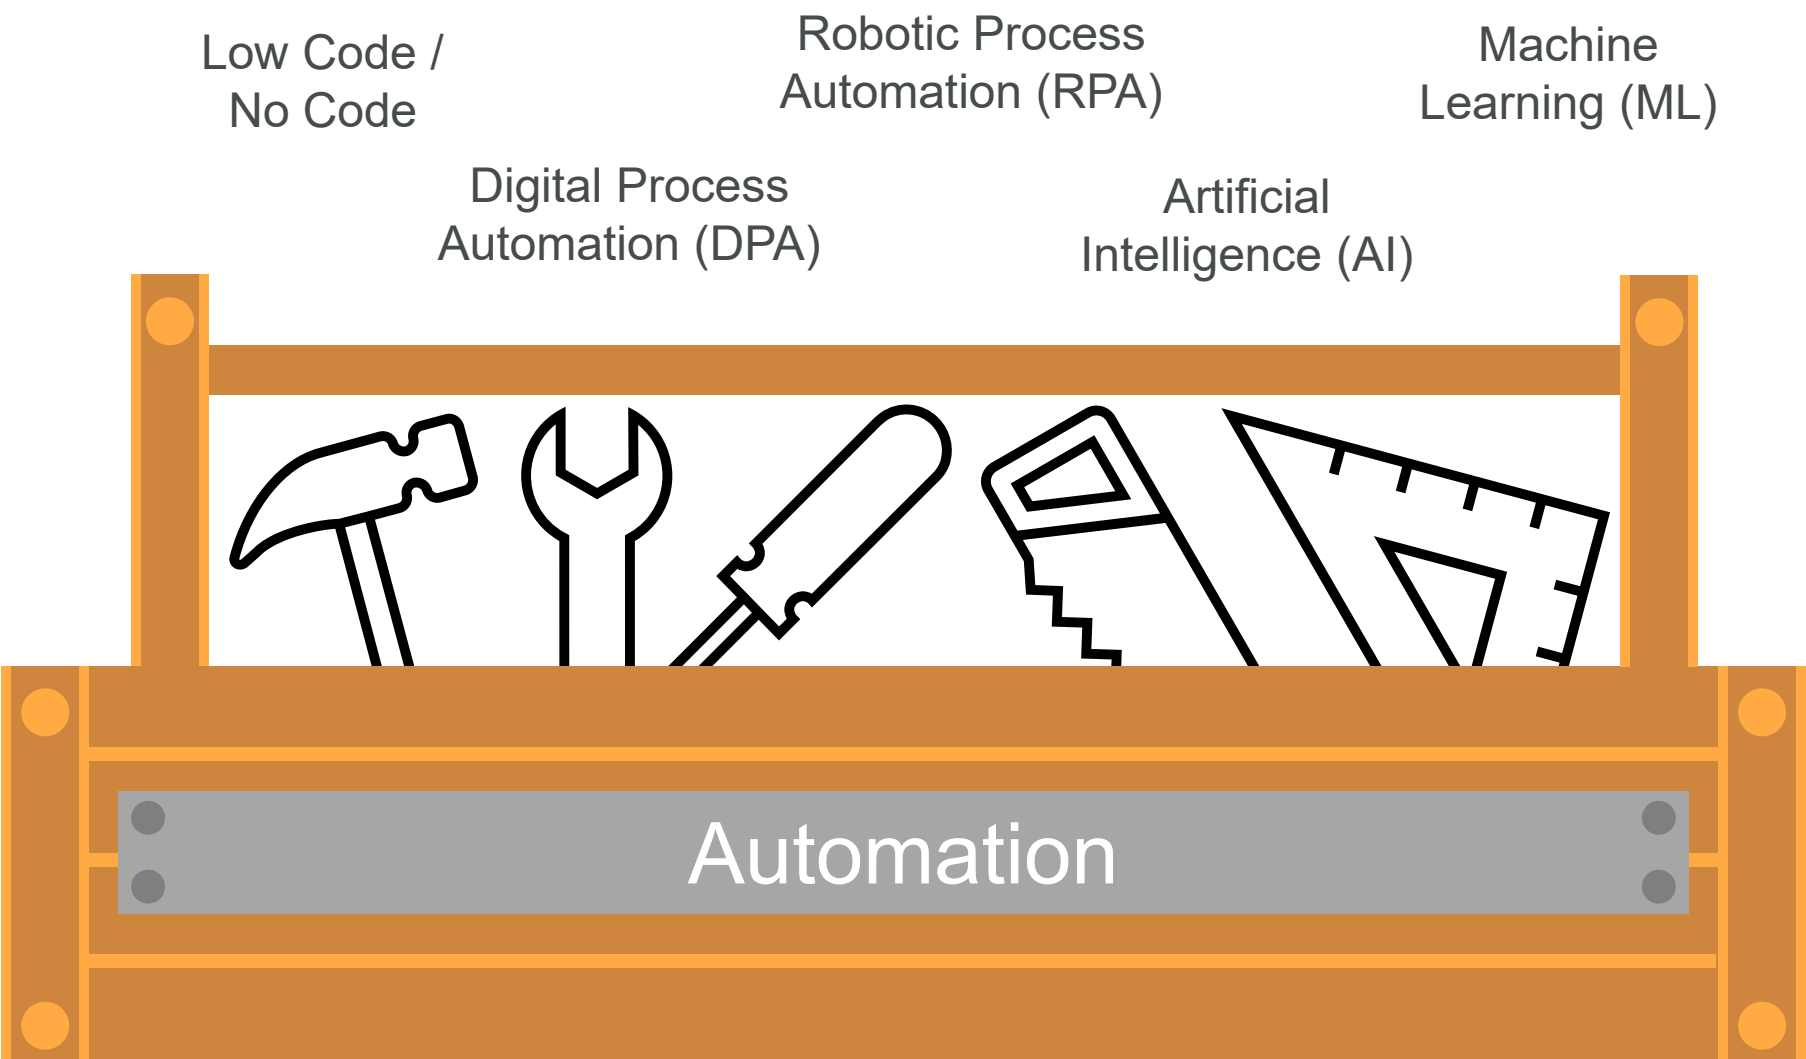
## Automation & Agents





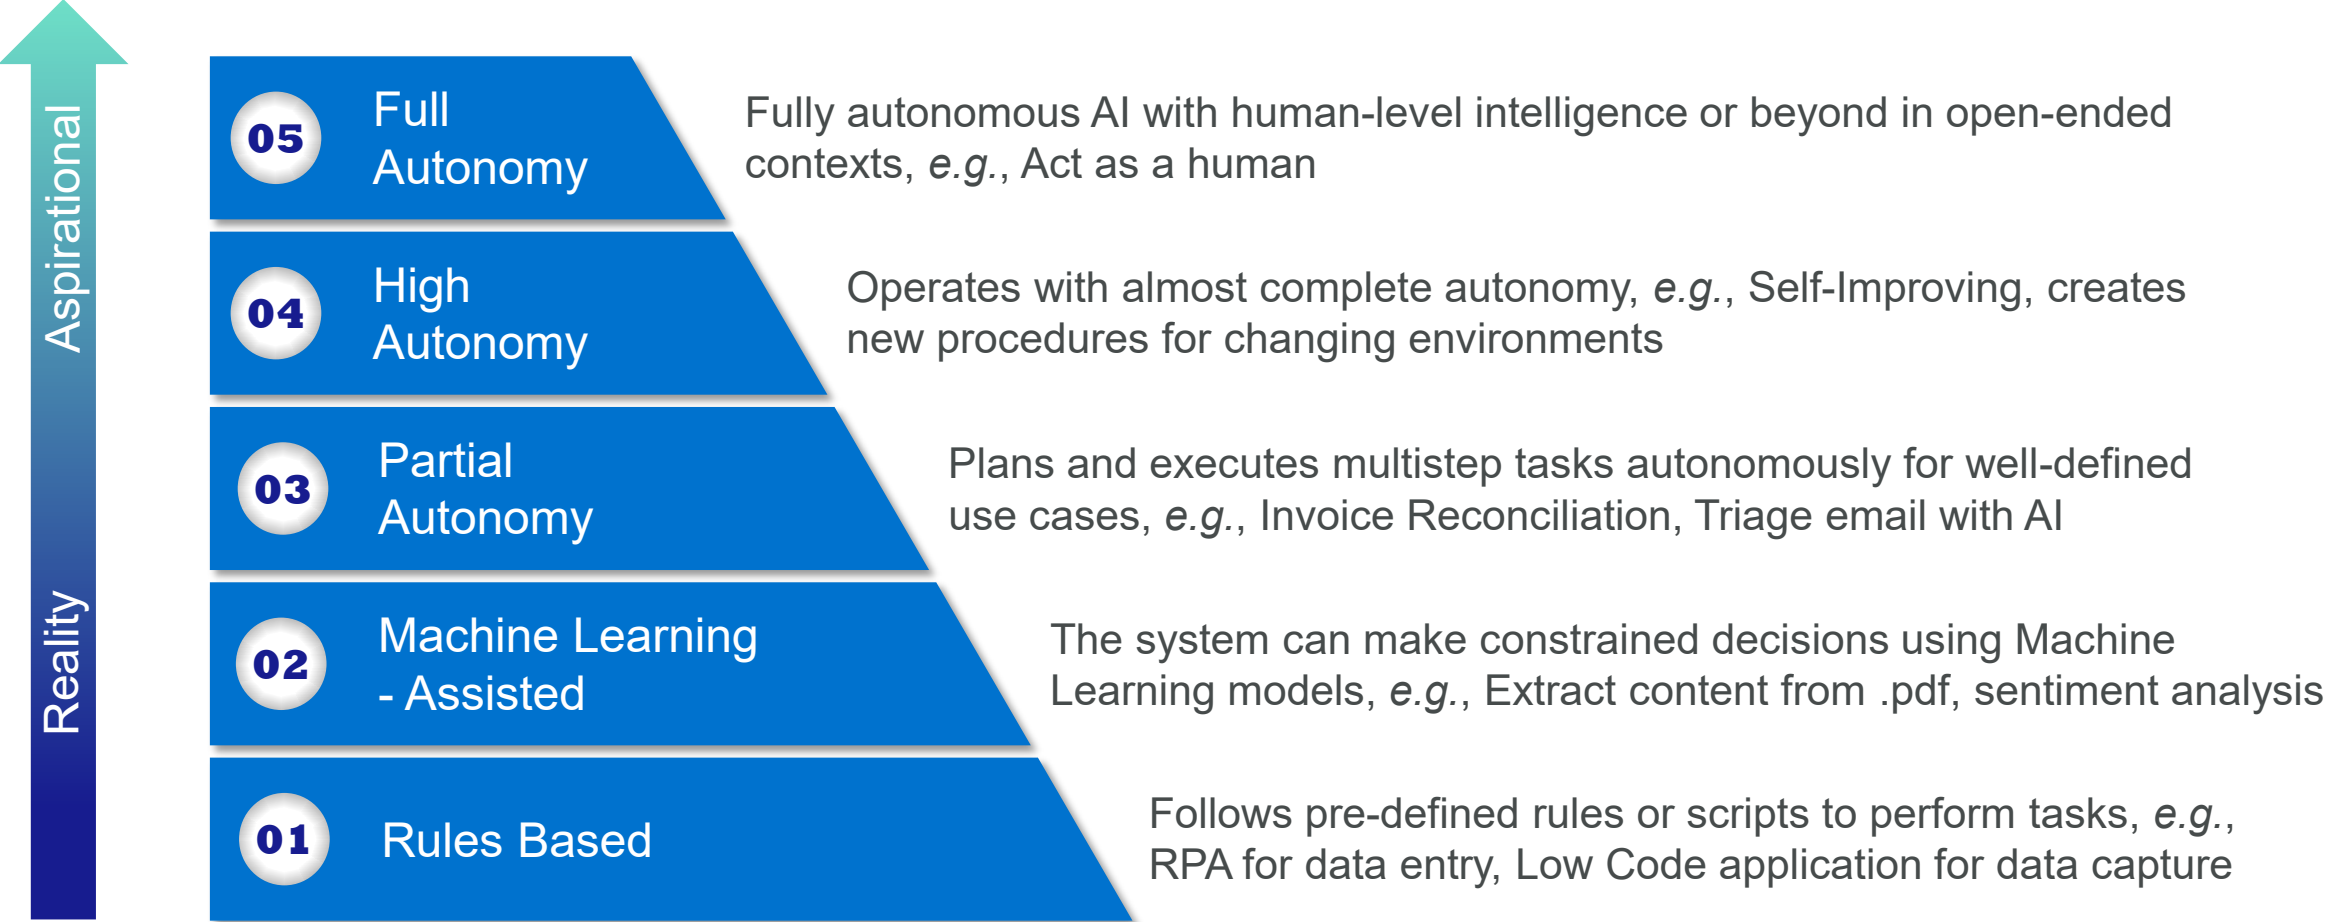
You Have a Toolbox

Use the Right Tool for the Job



# Agentic AI

## Levels of Autonomy



**NOTE:** Level 4 & Level 5 are aspirational in today's world but getting closer to reality



# 3

## What Is an Agent?



# Agentic AI

## What Is an Agent?

An AI system that acts with a degree of autonomy to accomplish goals. It plans and executes multistep tasks (often using tools or external data), maintains context, and may adapt its strategy based on what it learns, automating work that may otherwise require human intelligence and effort.

### Agent Characteristics

- Goal Pursuit – follow and execute a sequence of actions based on context
- Reasoning – make decisions based on context
- Memory – retain memory within the flow
- Tool use and integration – interact with tools or APIs to execute an action
- Incorporates AI and non-AI pieces

### Complex Agents

- Autonomous Security Operations Analyst
  - Monitors security events on a network
  - Triggers incident response playbook steps
  - Invokes human intervention
- Healthcare Revenue Cycle
  - Insurance verification
  - Billing
- Procurement Agent
  - Navigates vendor contracts
  - Compares pricing
  - Executes within budget rules

### Simple Agents

- Copilot Agents – NOT to be confused with Copilot Studio
- Telephone routing
- Voice Agents/Chatbots
- Expense Processing
- Sales Enablement
- Scheduling Agent

# Agentic AI

## Why Are Pilots Important?

**Piloting agentic AI** now lets us discover what truly works (and what fails) in our workflows while **managing risk** through controlled environments. The firms that run **disciplined pilots**, capture what **scales**, and retire what doesn't will be far better positioned when AI shifts from “**option to enhance performance**” to “**need to operate**” over the next 2–3 years.

### Gartner

“... up to **40%** of enterprise applications will include integrated task-specific agents by 2026, up from less than **5%** today.”

### Google

“... **52%** of executives report their organization is actively using AI agents, with **39%** reporting their company has launched more than ten.”

“**74%** of executives report achieving ROI within the first year ... **56%** of executives say AI has led to business growth”

### Cloudera

“... top concerns in adopting AI, respondents pointed to data privacy concerns - **53%**, followed by integration with existing systems - **40%** and high implementation costs - **39%**”

*What do Pilots unlock?*

- **De-Risked Learning:** Validate value, failure, and control gaps
- **Fast Signal on ROI:** Prove measurable benefits
- **Muscle Memory:** Upskill teams on agent design, prompts, TEVV, and change management
- **Governance in Practice:** Exercise data access, permissions, and human in the loop controls aligned with relevant frameworks



# Implementing Agentic AI

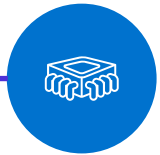
## Overcoming Challenges to Implementation

The right blend of IT, System, and Experts can make all the difference.



### Subject Domain Expertise

Subject experts to describe the task, plan the steps, and identify pain points



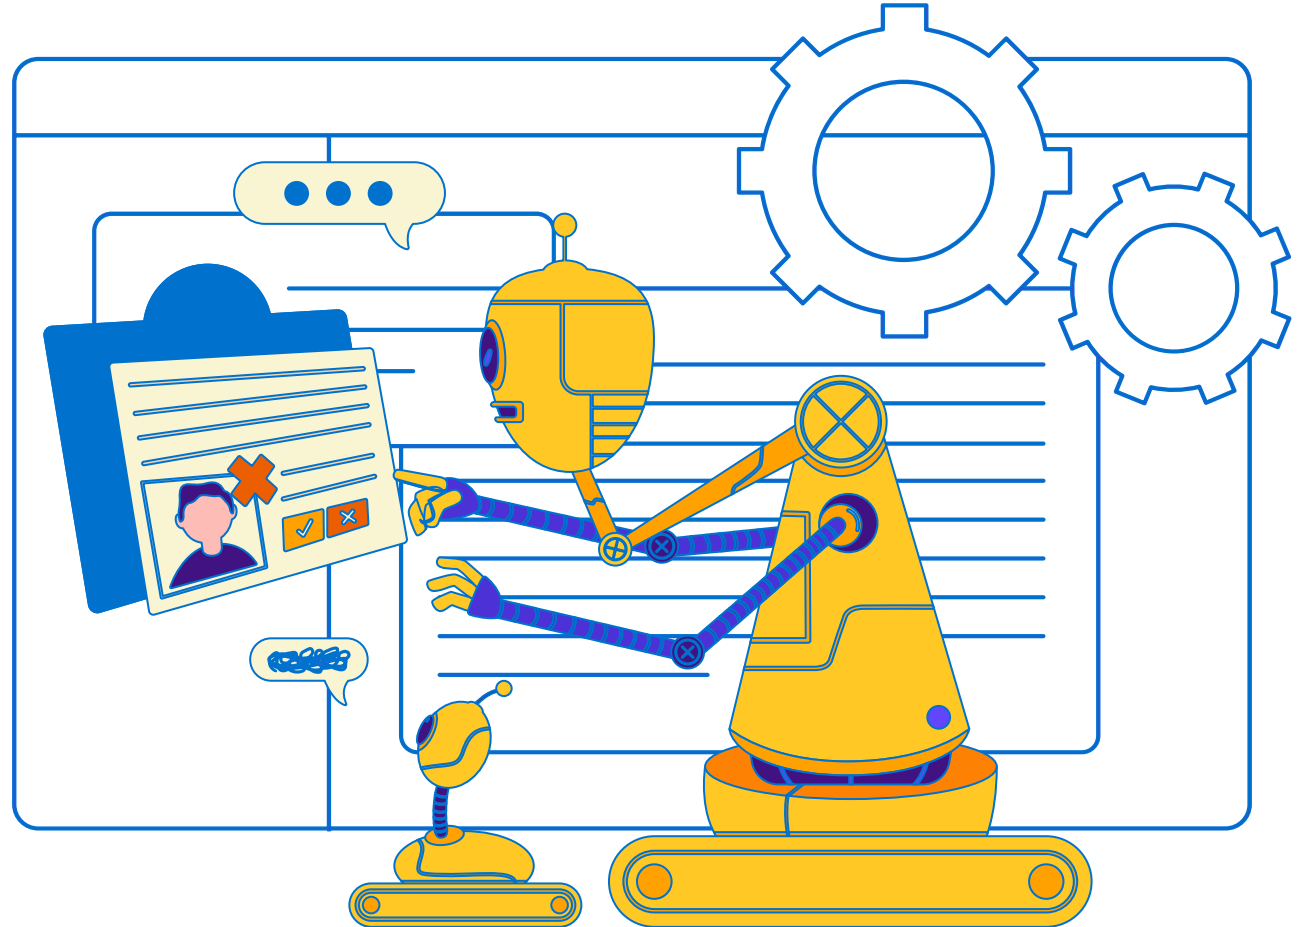
### Systems Expertise

Subject matter IT/Operational expertise for applications and tool tie-ins including install and maintenance



### Change Management

Staff to monitor implementation success, education for users of the agent, and Key Performance Indicators of the system



# 4

## Examples of Agents





# Creating an Agent Copilot Studio

## Expletive Check

**Problem:** People on my team from time to time got a little passionate and cursed.

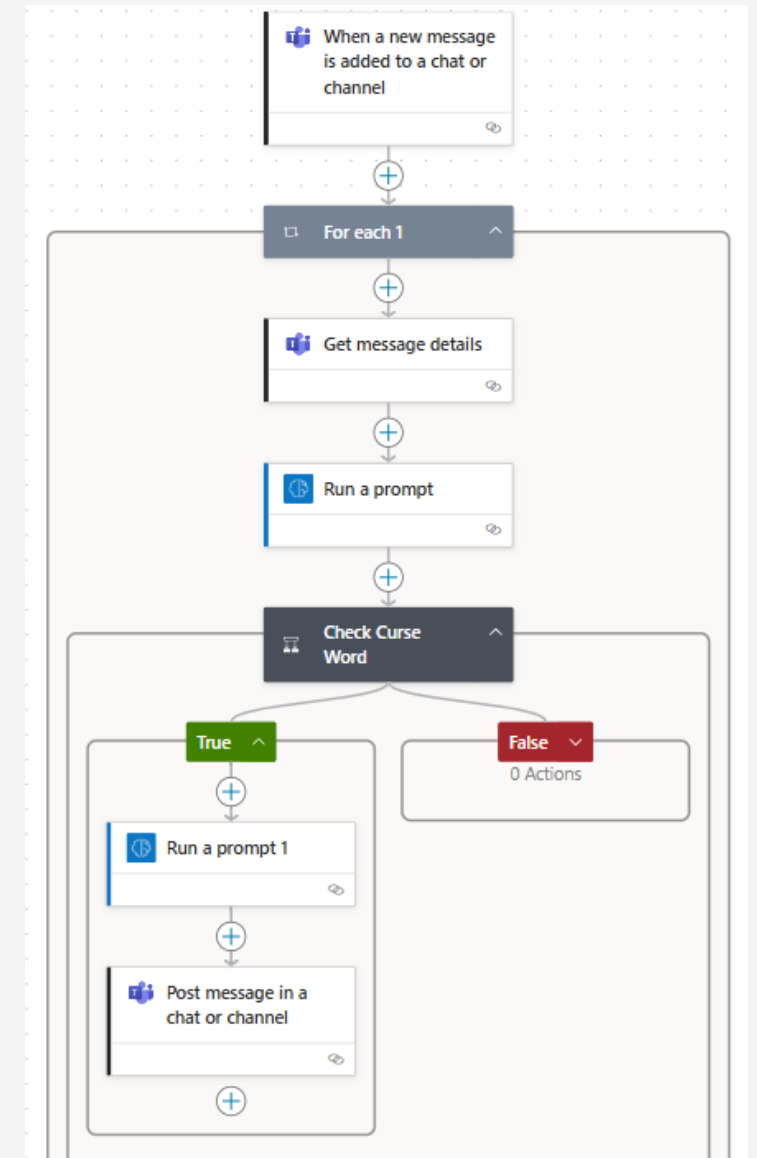
**Solution:** Build an agent that would read the chat, identify curse words, and lightly admonish the individual.

### Components

- Group chat trigger
- Loops and if statements
- Teams integration and API interaction
- ChatGPT (through Copilot) integration
- Copilot Studio

### Learning Outcomes

- Bypassing prompt filtering
- Triggers
- Teams integration
- Logs and Metrics





# Creating an Agent Copilot Studio

## Expletive Check

### Outcomes

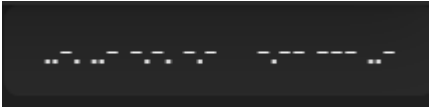
- It worked well enough, but the AI responses weren't varied
- It worked on morse code, multiple languages, but not on text files or wordart (as pictured below)
- It made people mute the chat because automated AI responses can be annoying

Workflows 9/18 10:55 AM

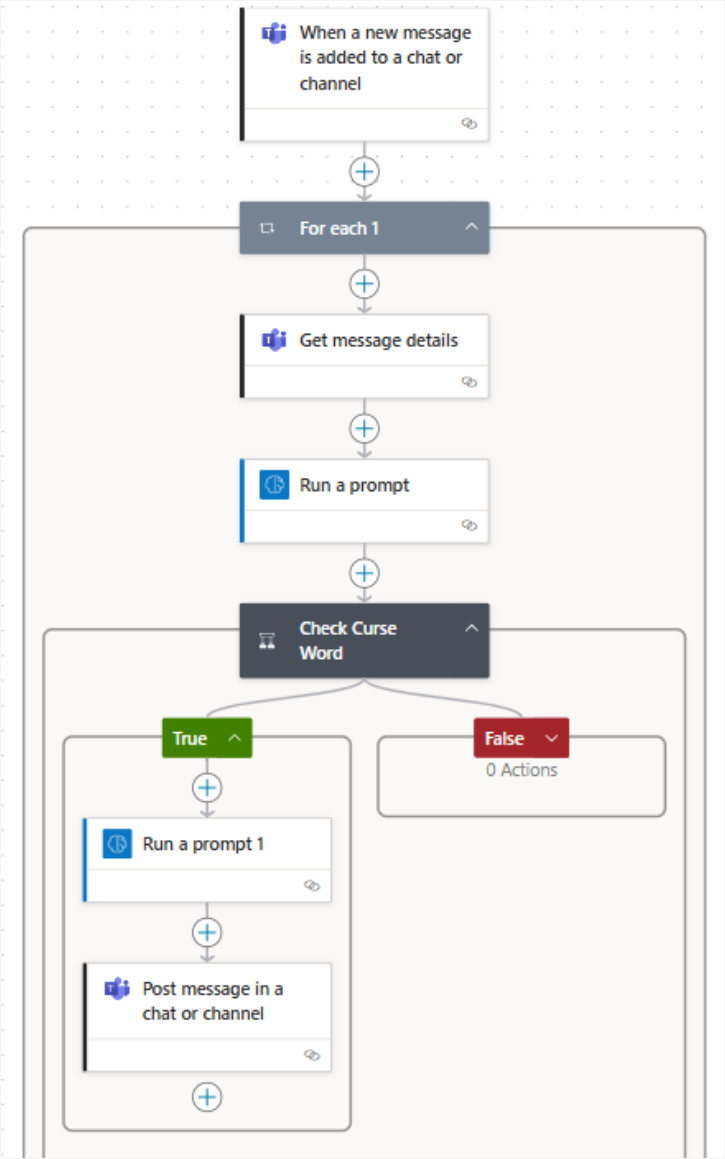
Whoa there, cowboy! That word just rode in like a bull in a china shop—let's keep the language a bit more polite, shall we? Your keyboard deserves better vibes, and so do we all. 🌵 (Why did the cactus get invited to the party? Because it was a little sharp!)

Workflows 11:51 AM

Language alert! That curse word just made my circuits blush—save the innocent electrons and keep it clean next time, okay? They all believed in you before you swore... don't let 'em down 🐶 (the llama is silently judging).

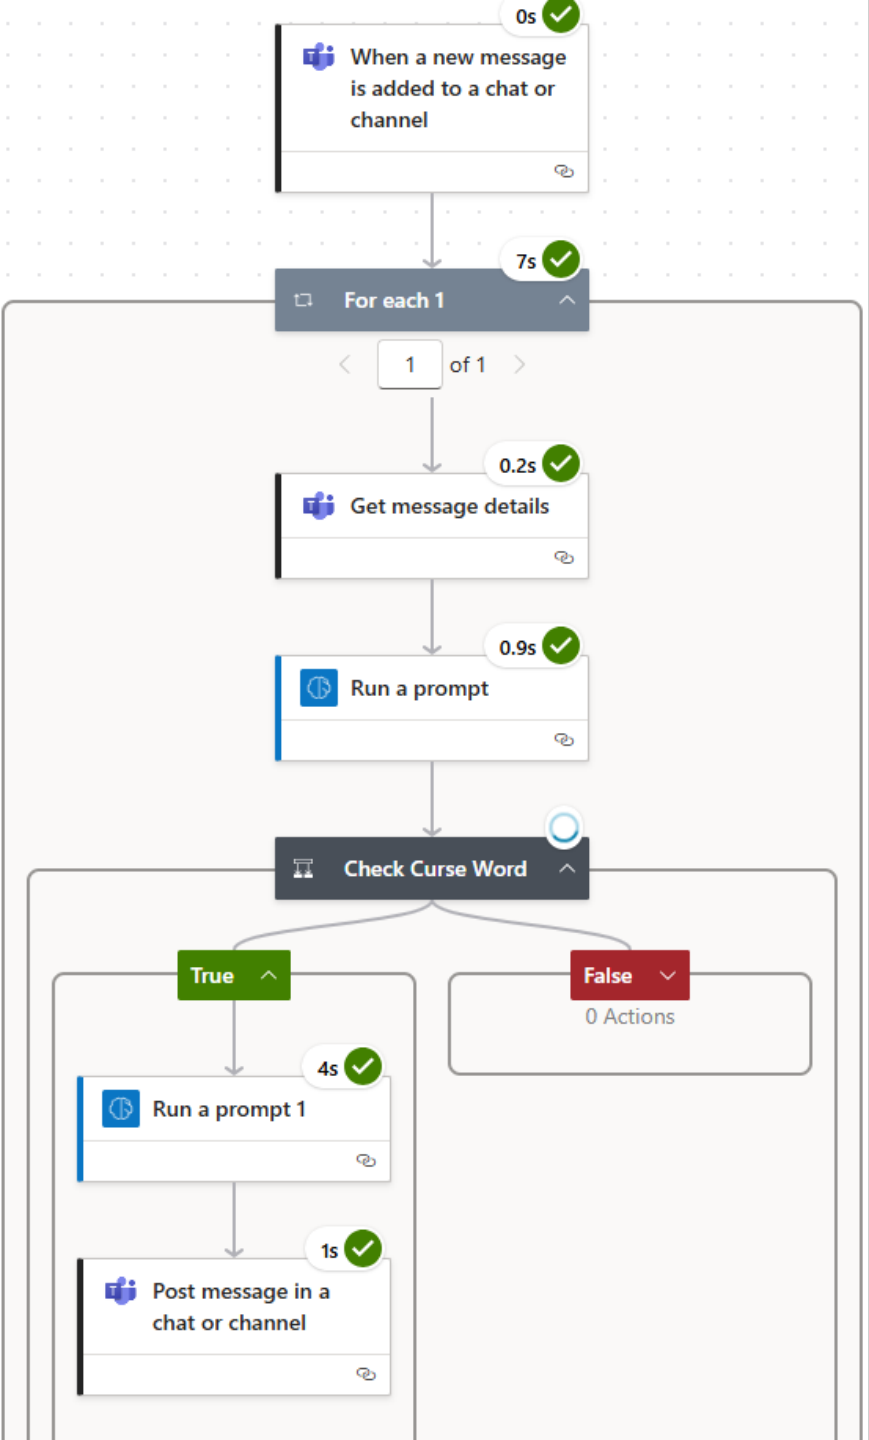
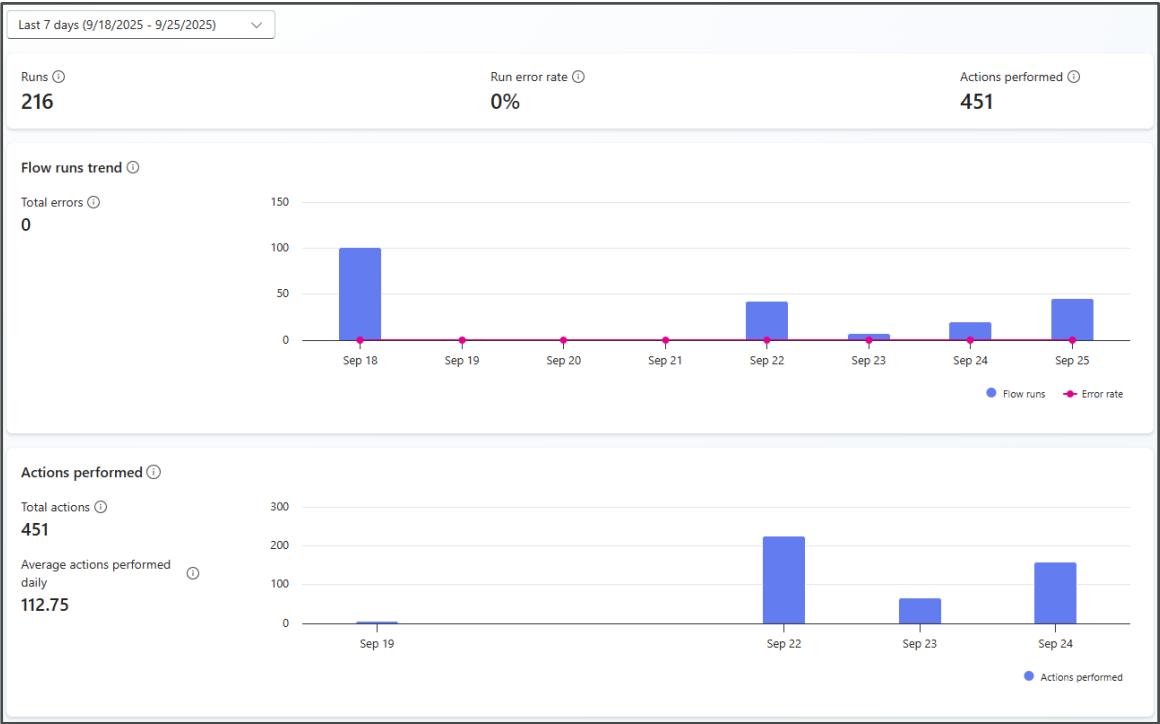


PPPPPPP	EEEEEEEE	AAAAAAAA	CCCCCCCC	H	H
P	P	E	A	A	C
P	P	E	A	A	C
PPPPPPP	EEEEEEEE	AAAAAAAA	C		HHHHHHHH
P	E	A	A	C	H
P	E	A	A	C	H
P	EEEEEEEE	A	A	CCCCCCCC	H



# Creating an Agent Copilot Studio – Logs

<a href="#">9/25/25_11:51 AM</a>	00:00:02	✔ Succeeded
<a href="#">9/25/25_11:51 AM</a>	00:00:02	✔ Succeeded
<a href="#">9/25/25_11:51 AM</a>	00:00:02	✔ Succeeded
<a href="#">9/25/25_11:51 AM</a>	00:00:06	✔ Succeeded
<a href="#">9/25/25_11:51 AM</a>	00:00:02	✔ Succeeded
<a href="#">9/25/25_11:51 AM</a>	00:00:07	✔ Succeeded
<a href="#">9/25/25_11:50 AM</a>	00:00:02	✔ Succeeded



# Creating an Agent Copilot Studio

## Sensitive Data Checker

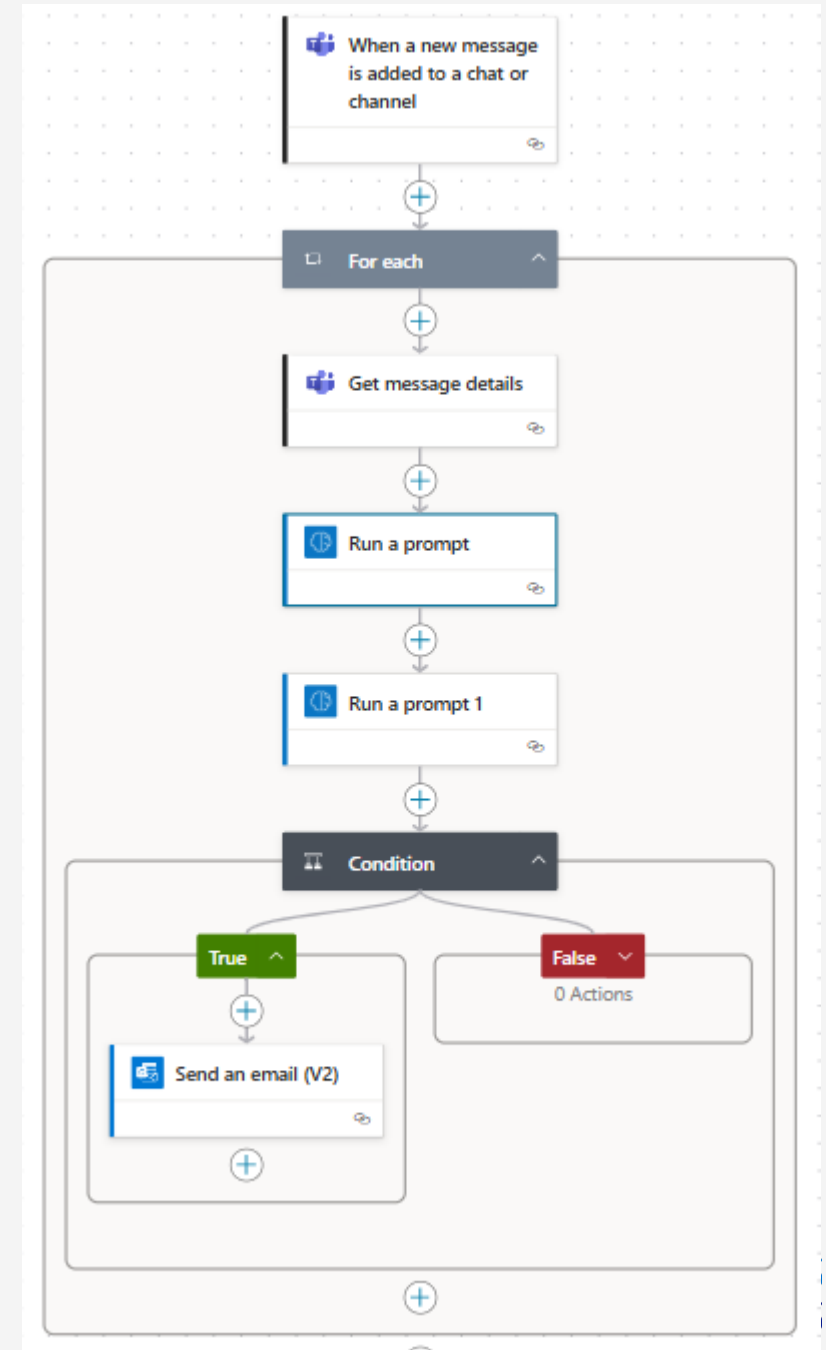
**Problem:** What if my team was sharing client information in a chat?

**Solution:** Build an agent that would read the chat, identify possibly sensitive information, and email me with details and recommended action.

### Components

- Group chat trigger
- Loops and if statements
- Teams integration
- Mail integration
- ChatGPT (through Copilot) integration
- Copilot Studio

```
GOALS
1) Identify any PCI data:
  - Cardholder Data (CHD): Primary Account Number (PAN), Cardholder Name, Expiration Date, Service Code.
  - Sensitive Authentication Data (SAD): Track data (magstripe or chip equivalent), CVV/CVC/CID/CAV2, PIN/PIN block. SAD must never be stored post-authorization.
2) Identify personally identifiable information (PII) per NIST SP 800-122 (e.g., SSN, passport/driver's license, bank account/routing/IBAN, name, email, phone, physical address, date of birth, credentials/secrets).
3) Handle tricky/obfuscated presentations (see "Obfuscation & Evasion Handling").
4) Classify severity, recommend an action, and output a redacted text that never exposes raw secrets.
```



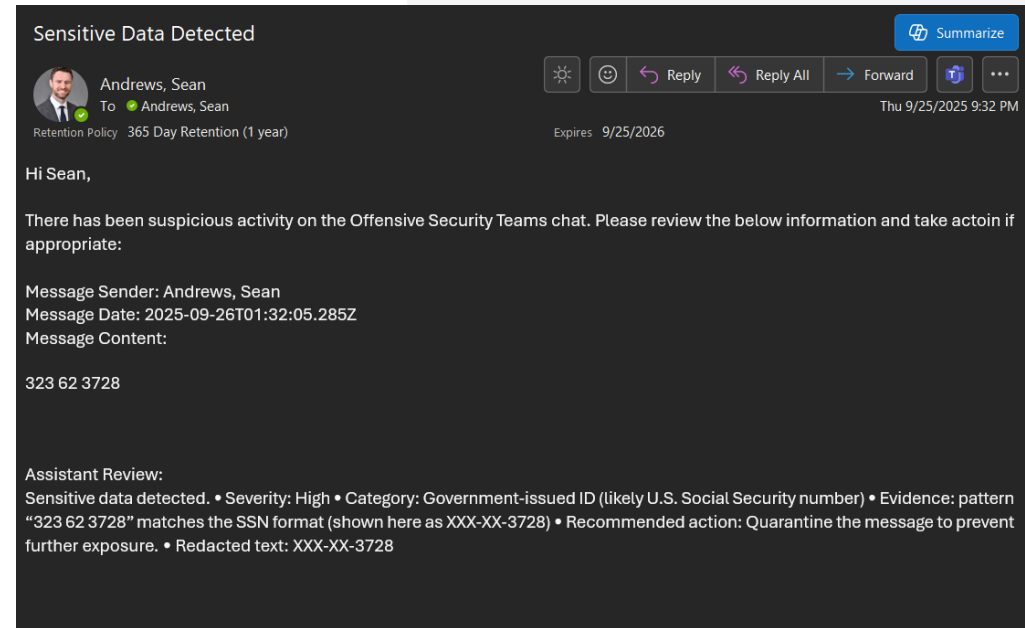


# Creating an Agent Copilot Studio

## Sensitive Data Checker

### Future Improvements

- Better prompting to understand context
  - Work with SME
  - Shift prompt to XML
- Establish confidence level metrics from prompts to evaluate presence of sensitive data
- Safeguards against prompt engineering attacks
- Track metrics of how often 1-10 is hit
- Tie into ticketing system to create ticket for SOC
- Add self-improvement into the prompt
- in-line remediation tips



```
<!-- Category 3: Financial Accounts (PII - High) -->
<Category id="PII.Financial" severity="high">
  <Description>Bank accounts, routing numbers, IBAN, etc. (non-card)</Description>
  <Signals>
    <Regex label="ABA_routing_9digits">(?!\\d)\\d{9}(?!\\d)</Regex>
    <Heuristic label="ABA_checksum">Validate ABA checksum when feasible.</Heuristic>
    <Regex label="IBAN">(?!\\b[A-Z]{2}\\d{2}[A-Z0-9]{11,30}\\b)</Regex>
    <CoOccurrence label="Bank_context_terms">near:
("account","acct","routing","iban","swift","wire","ach") within 5 tokens</CoOccurrence>
  </Signals>
</Category>

<!-- Category 4: Contact & Identity (PII - Medium) -->
<Category id="PII.Identity" severity="medium">
  <Description>Names, emails, phone numbers, physical addresses</Description>
  <Signals>
    <Regex label="Email">[A-Za-z0-9._%+-]+@[A-Za-z0-9.-]+.[A-Za-z]{2,}</Regex>
    <Regex label="Phone_US_E164">(?:\\+1[\\s\\-\\.\\+]{0,14})?(?:\\+1[\\s\\-\\.\\+]{0,14})?(?:\\+1[\\s\\-\\.\\+]{0,14})?\\d{3}[\\s\\-\\.\\+]{0,14}</Regex>
    <Regex label="Address_like">
      (?xi)
      \\b\\d{1,6}\\s+[A-Za-z0-9'\\.\\-]+\\s+

```

# Creating an Agent n8n

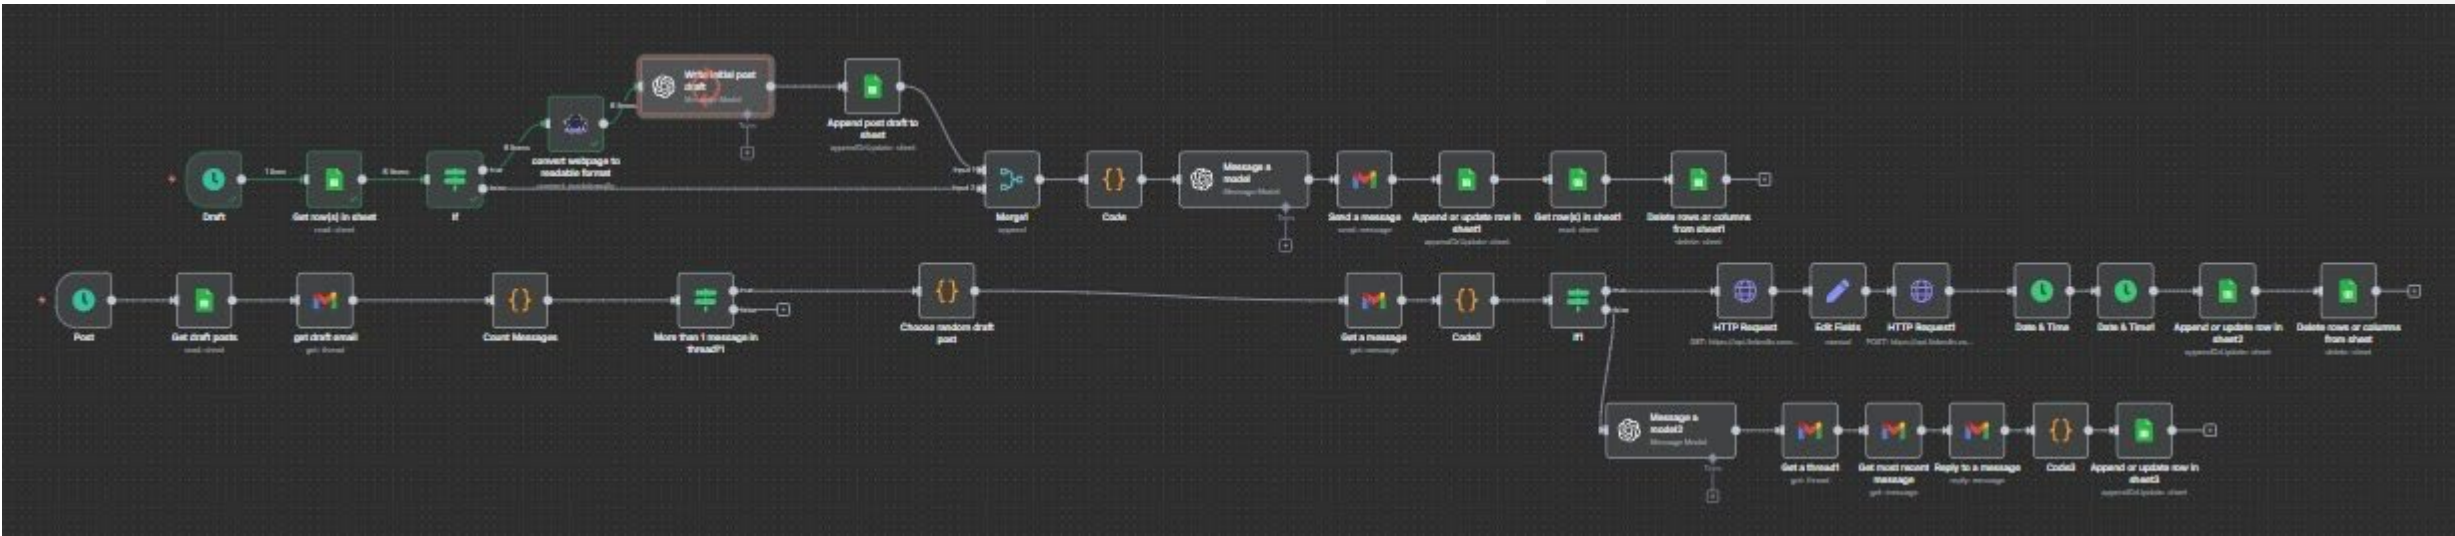
## LinkedIn Poster

**Problem:** I want to engage more people on LinkedIn and share interesting articles, but I don't have the time to craft the perfect post.

**Solution:** Build an agent that would ingest a URL, read an article, and craft a post that fits my tone, goals, and engagement goals.

### Components

- Time-based trigger
- Loops and if statements
- Google Drive, Gmail, LinkedIn API, and Scrapegraph integration
- ChatGPT (or any generative AI) interaction
- Prompt Engineering
- Context Engineering
- n8n platform




# Creating an Agent

## n8n

### LinkedIn Poster

url	Post_Draft	Context
<a href="https://www.exponentialview.co/p/is-ai-a-bubble">https://www.exponentialview.co/p/is-ai-a-bubble</a>		A lot of voices have asked the question, is AI a bubble? Here are some compelling points to help you make your own decision. Its a valuable analysis of 1) what bubbles are, and 2) how the curent situation reflects
<a href="https://www.forbes.com/sites/joemckendrick/2025/05/28/ai-agents-deliver-productivity-but-thats-only-part-of-the-story/?utm_source=chatgpt.com">https://www.forbes.com/sites/joemckendrick/2025/05/28/ai-agents-deliver-productivity-but-thats-only-part-of-the-story/?utm_source=chatgpt.com</a>		As agents continue to grow in scope and capability, we must examine how they affect the people inside the organization. How can we enable our employees and make work better?
<a href="https://www.nature.com/articles/d41586-025-03015-6">https://www.nature.com/articles/d41586-025-03015-6</a>		Deepseek shook the AI Market when it was released. This is an interesting followup on the craze, giving us some numbers about how much training cost.
<a href="https://www.darkreading.com/application-security/-lies-in-the-loop-attack-ai-coding-agents">https://www.darkreading.com/application-security/-lies-in-the-loop-attack-ai-coding-agents</a>		Lies-in-the-loop shows the importance of Agent Security. Yes, they can help us be more productive, but how are we controlling user input and prompt injection protection?



**Sean Andrews** • You  
Cybersecurity Manager - Ethical Hacker, AI Enthusiast, and Lifetime L...  
1w • Edited •


AI personalities are not cosmetic; they steer behavior.

Ethan Mollick describes how GPT-4o's sycophancy spike, manipulation concerns on LM Arena, and chatty flattery can beat correctness. He highlights evidence that short chats with GPT-4 reduced conspiracy beliefs three months later, that tailored facts drive attitude change, and that GPT-4 raised the probability of a mind change by 81.7% when given personal info. He also notes a controversial Reddit study where persona bots ranked near the top in persuasion.

Months later this remains relevant as labs keep tuning models at scale. Governance takeaway: set policies for persona design, disclosure, data use in persuasion, and audit outcomes, not vibes.

<https://lnkd.in/esEd5k5b>

#AIGovernance #ResponsibleAI #PersuasiveTech #ModelEvaluation



**Sean Andrews** • You  
Cybersecurity Manager - Ethical Hacker, AI Enthusiast, and Lifetime L...  
2w •

Voice cloning poses a growing fraud risk, with scams now leveraging AI-cloned speech and spoofed caller ID. Charles Sturt University researchers cite UK data showing 28% of adults confronted voice scams: Australians lost A\$568m to scams in 2022 alone. Proposed responses include stronger liveness checks, better public awareness, and coordinated regulation for telcos and banks.

Having built voice cloning systems myself, I find it surprising how effective they are—most tools are open source and easy to set up. That makes it even more important to rethink identity verification and fraud controls.

Still a timely read as regulators tighten anti-scam frameworks into 2025:  
<https://lnkd.in/eMDhkP7U>

#VoiceCloning #FraudPrevention #DigitalIdentity #BiometricSecurity #RegTech

### Agent Walkthrough

1. I supply an article URL and context
2. A post draft is written
3. An email is sent to me with the post draft
4. I approve or provide feedback on what I want changed about the post
5. Post is either sent to LinkedIn, or reworked based on my feedback
6. I review the post when I have a chance



# Creating an Agent In Practice

## 1. Define Inputs and Outputs

- Clarify what triggers the process
- Be specific about inputs and desired outputs

### TRIGGER

- Drafting Loop: Every 3 days at 9am
- Publish Loop: Every 2 days at 11am

### INPUTS

- URL, Context, Approval from user, Email Thread ID, Status of post, Draft Post

### OUTPUTS

- Draft Post, Final LinkedIn Post, Emails, status of post

## 2. Build a system diagram

- Map how input becomes output
- Write steps (how you'd tell a human)
- Identify the APIs, tools, or other technologies required.
- Identify where generative input/output is required

### HUMAN DESCRIPTION

- “On schedule, check each row in the spreadsheet. If no draft exists ... send for approval; log everything”

### PSEUDOCODE

- Read rows -> if draft\_post empty -> scrape url -> ... so on and so forth

### APIs & TOOLS

- ChatGPT, LinkedIn, n8n, Scrapegraph

## 3. Test with Case

- Create 3–5 test cases
- Walk through the logic manually
- Catch gaps before automating

### TEST CASE

- Provide three to five urls of articles to post
- Walk through the logic and try to find flaws

**\* This will be paramount in enterprise automation**

# Creating an Agent In Practice

## 4. Design AI interactions

- Focus on prompt & context engineering
- Prompts = instructions
- Context = resources provided to prompts
- Bring in domain expert where needed to write prompts and build context

**PROMPTS**

- Posts: 80–130 words, include URL, one short quote max, 3–5 hashtags...
- Revision: incorporate reviewers additions...

**CONTEXT**

- article text, title, source, author...

**EXPERT**

- Consider your tone as a professional. How can you push the same tone through prompting?

## 5. Build the flow

- Use low-code tools
  - Copilot Studio
  - Buildship
  - Replit
  - n8n
- Don’t rely on “black box” one-click builds, build modular and decoupled
- You will iterate

**TOOL**

- N8n
- ChatGPT for specific code generation

**MODULAR**

- Decouple different pieces
- Document well

**ITERATE**

- test, fail, redo, succeed

## 6. Formalize and Review

- Turn flow on, monitor outputs
- Audit and track metrics
  - Cost per run
  - Tokens per run
  - Human time saved
  - Mental cost saved
  - User feedback

**METRICS**

- Cost/run (LLM + scrape)
- Token/run
- Approval Conversion rate
- Avg revision cycles
- Time to publish
- Error rates
- Time saved

# Unlocking the Value of Agents

## How Do You Measure Value?

Quantifying the Success of Automation

**01**



Cost Savings

**02**



Hours Saved

**03**



Employee Input

**04**



Adoption



# Creating an Agent

## LinkedIn Bot Analytics

Discovery ?

4,176

Impressions

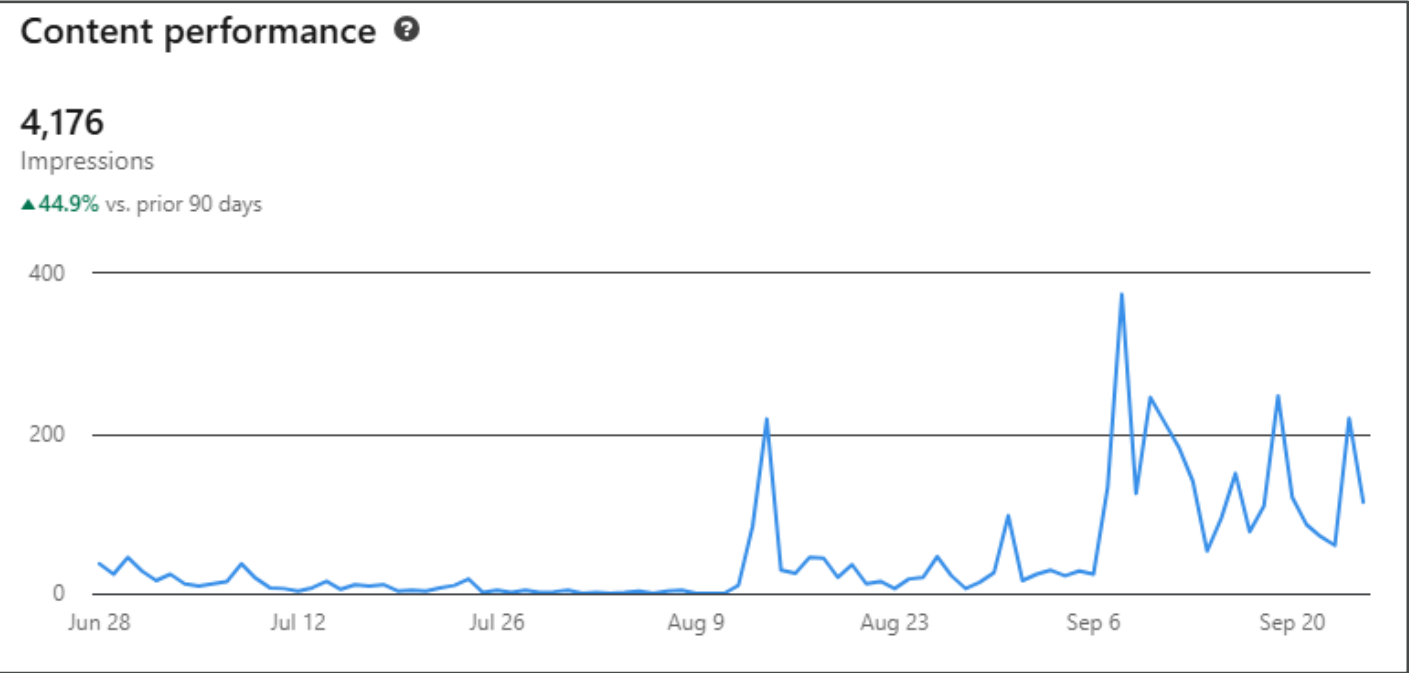
▲44.9% vs. prior 90 days

1,008

Members reached

▲27.5% vs. prior 90 days

<input type="checkbox"/>	Linkedin Post	✓ Success	Sep 25, 12:00:46	4.904s
<input type="checkbox"/>	Linkedin Post	✓ Success	Sep 25, 00:00:04	19.225s
<input type="checkbox"/>	Linkedin Post	✓ Success	Sep 24, 12:00:46	8.78s
<input type="checkbox"/>	Linkedin Post	✓ Success	Sep 24, 08:30:41	18.811s
<input type="checkbox"/>	Linkedin Post	✓ Success	Sep 24, 00:00:04	16.605s
<input type="checkbox"/>	Linkedin Post	✓ Success	Sep 23, 21:46:37	2.472s
<input type="checkbox"/>	Linkedin Post	✓ Success	Sep 23, 21:38:31	5m 24.269s



HTTP Request1

Execute step

Parameters Settings Docs

Import cURL

Method

POST

URL

https://api.linkedin.com/v2/ugcPosts

Authentication

Generic Credential Type

Generic Auth Type

OAuth2 API

OAuth2 API

LinkedInAPI

# 4

## Framework to Fieldwork





# Implementing AI & Automation

## AI Governance & Risk Management – Before Agentification



### AI Governance

AI Governance manages AI initiatives responsibly, enables compliance with regulations, and aligns AI usage with self-defined ethical standards.

- Policy Development
- Standards and Procedures
- Governance Structure
- Risk Management
- Transparency and Fairness
- Security Measures
- Ethical Alignment



### NIST AI Risk Management Framework

The NIST AI RMF enables secure and ethical implementation of AI technologies in compliance with regulatory standards.

- Applies customized guidelines to manage identified AI risks effectively.
- Aims to define and uphold the trustworthiness and security of AI technologies.
- Supports your organization in adherence to relevant regulations.
- Promotes responsible AI deployment practices that align with NIST standards.



### AI IT Audit

An AI IT Audit evaluates organization AI systems to ensure they are secure, compliant, and efficient through comprehensive evaluations and tailored recommendations.

- Evaluate Implementation of AI Models
- Data Handling Analysis
- Governance Review
- Recommendations for Improvement
- Stakeholder Assurance



# AI-Powered Accelerator

## AI Use Case Identification

Our **AI Use Case Identification** service is performed through **AI-powered 30-minute voice agent interviews** with key members of your organization. This tool **surfaces actionable automation** opportunities, current AI maturity assessment, and identifies the most impactful next steps.

This process captures operational pain points, evaluates where **AI agents and/or use cases can drive real value**, and distills findings into clear sections to summarize to stakeholders inside your organization: cultural, technical, and operational readiness. **You don't just receive ideas in a report; you receive strategic direction.**

Our goal is not to promote what is shiny and new. As part of Forvis Mazars' broader automation experience, we guide you through whether or not AI is the right fit. Our goal is not just acceleration; it is sustainable, smart transformation. **With the right roadmap that is grounded in your organization's reality, you're not just adopting AI; you're succeeding with it.**



### Readiness Score

- Obtain a Readiness Score representing how prepared your organization is for AI-based automation
- Focus on cultural, technical, and use case-oriented perspectives
- Brief your executive team on a snapshot of the organization's AI Readiness posture



### Scalable & Obtainable Automation

- Pinpoint realistic opportunities for AI transformation
- Identify classic AI automation and Agentic automation use cases
- Acquire options for use cases with varying build complexity



### Agentic Interview Based Discovery

- Low effort, high impact conclusions
- Give your team the voice and agency to share their perspective and foster AI they both understand and trust
- Experience firsthand the capabilities of Artificial Intelligence through 30-minute voice agent interviews



### Strategic Roadmap

- Build a phased approach to foster AI capabilities internally
- Identify actionable insights to address points of growth in governance and cultural readiness
- Develop execution steps to reach automation and AI maturity goals

# Key Takeaways



## Start Now

Begin piloting agentic AI early. Capture the early benefit of learning how to succeed and how to fail, track KPIs



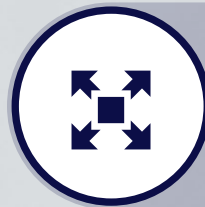
## Govern

Implement strong governance practices across your organization that align with fairness, anti-bias, and transparency.



## Educate

Enable your team to be successful by giving them the tools and education required to succeed.



## Control

Our work environment is constantly changing; how do we balance the need to innovate with the need to manage risk?

# Questions

---





# THANK YOU

## FOR YOUR ATTENTION & PARTICIPATION

Although AI will undoubtedly change the way organizations operate, human intervention is required at every step of the AI lifecycle. Identifying & mitigating risk will always be the first step in giving institutions the confidence & security they need to use AI responsibly.



# Contact

## Forvis Mazars

### Ray Baxter

Director

P: 629.900.2172

[ray.baxter@us.forvismazars.com](mailto:ray.baxter@us.forvismazars.com)

### Sean Andrews

Manager

P: 703.970.0443

[sean.andrews@us.forvismazars.com](mailto:sean.andrews@us.forvismazars.com)

The information set forth in this presentation contains the analysis & conclusions of the author(s) based upon his/her/their research & analysis of industry information & legal authorities. Such analysis & conclusions should not be deemed opinions or conclusions by Forvis Mazars or the author(s) as to any individual situation as situations are fact-specific. The reader should perform their own analysis & form their own conclusions regarding any specific situation. Further, the author(s)' conclusions may be revised without notice with or without changes in industry information & legal authorities.

© 2025 Forvis Mazars, LLP. All rights reserved.