



# AI is here for business users

---

What does that mean for Security?

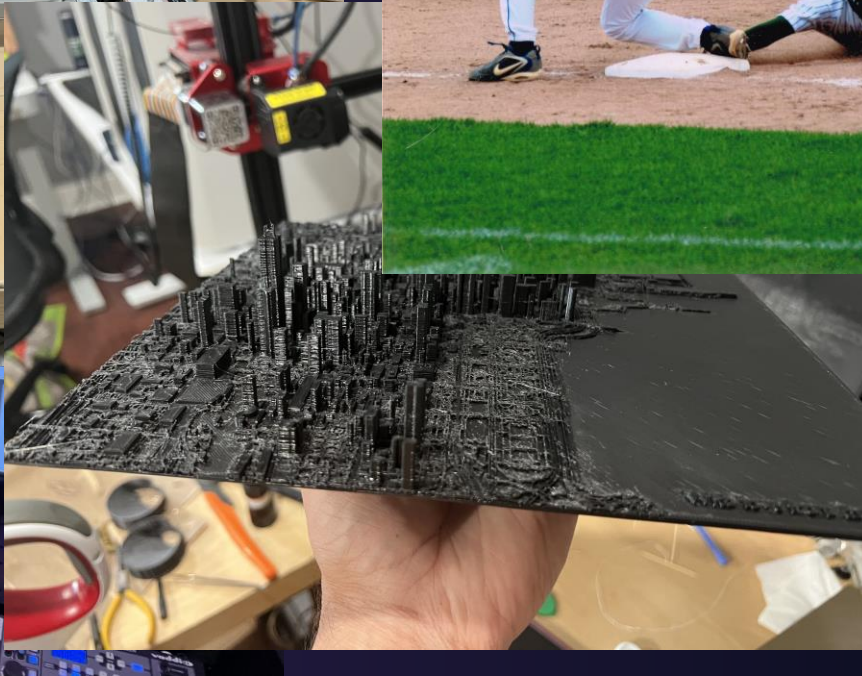


@sjshanko / @zenitysec



Stephen Shanko

# About Steve



AI is here

---

At Generative AI  
is here

—

An abstract graphic in the background of the slide. It features a dense, glowing blue and white network of lines and dots, resembling a complex data structure or a neural network. The lines are thin and connect numerous small, semi-transparent dots. The overall shape is somewhat elongated and tapers towards the right, with a more concentrated area of dots on the left side.

# The economic potential of generative AI: The next productivity frontier

June 14, 2023 | Report





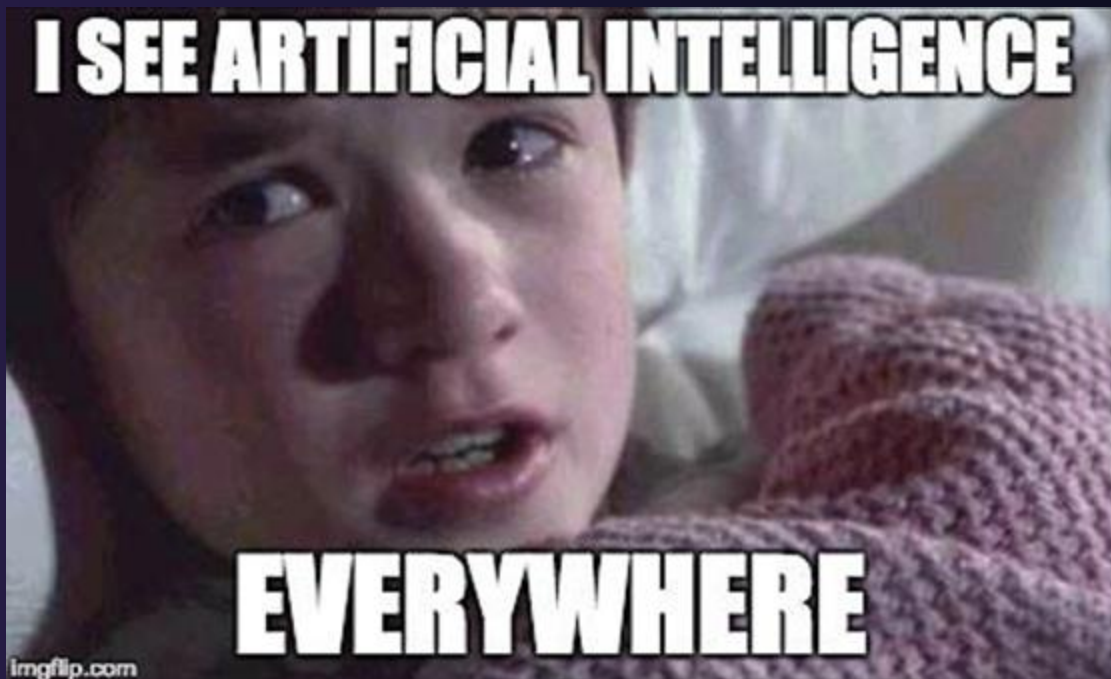
An abstract digital network graphic in shades of blue and white. It features a dense web of interconnected nodes and lines, with a prominent circular cluster on the left side. The overall shape is roughly rectangular, with a search icon in the top right corner.

# The economic potential of generative AI: The new frontier

June 14, 2023 | Report

# AI is here

- Everywhere you look – AI is there.



Source: [LinkedIn / Yusuf E.](#)



# AI is here

- Everywhere you look – AI is there.
- It's a great sales pitch aid.



Source: [LinkedIn / Yusuf E.](#)

# Salesforce Unveils Agentforce—What AI Was Meant to Be

September 12, 2024

*Agentforce is how humans with AI drive customer success together, equips organizations with autonomous agents that boost scale, efficiency, and satisfaction across service, sales, marketing, commerce, and more*

*New Agentforce Atlas Reasoning Engine autonomously analyzes data, makes decisions, and completes tasks, providing reliable and accurate results*

*With Agentforce, any organization can build, customize, and deploy their own agents quickly and easily, with low-code tools*

*New Agentforce Partner Network allows customers to deploy pre-built agents and use agent actions from partners like Amazon Web Services, Google, IBM, Workday, and more*

*Customers like OpenTable, Saks, and Wiley are turning to Agentforce because it is integrated with their apps, works across customer channels, augments their employees, and scales capacity for business needs*

SAN FRANCISCO--(BUSINESS WIRE)-- Salesforce (NYSE: CRM), the world's #1 AI CRM, today unveiled [Agentforce](#), a groundbreaking suite of autonomous AI agents that augment employees and handle tasks in service, sales, marketing, and commerce, driving unprecedented efficiency and customer satisfaction. Agentforce enables companies to scale their workforces on demand with a few clicks. Agentforce's limitless digital

features, compliance with industry standards, and unmatched flexibility. Our vision is bold: to empower one billion agents with Agentforce by the end of 2025. This is what AI is meant to be.”

“Agentforce represents the Third Wave of AI—advancing beyond copilots to a new era of highly accurate, low-hallucination intelligent agents that actively drive customer success. Unlike other platforms, Agentforce is a revolutionary and trusted solution that seamlessly integrates AI across every workflow, embedding itself deeply into the heart of the customer journey. This means anticipating needs, strengthening relationships, driving growth, and taking proactive action at every touchpoint,” said Marc Benioff, Chair and CEO, Salesforce. “While others require you to DIY your AI, Agentforce offers a fully tailored, enterprise-ready platform designed for immediate impact and scalability. With advanced security features, compliance with industry standards, and unmatched flexibility. Our vision is bold: to empower one billion agents with Agentforce by the end of 2025. This is what AI is meant to be.”

# Power Platform Numbers



**48 million**

Power Platform Monthly Active Users

**1,400**

Connectors available to connect to your data

FOR  
COMPARISON:

**40-45M**

People visit Las Vegas every year



**8M**

People have seen Taylor Swift's tour



**47M**



Population of Spain



**50,000**

Copilot Studio customers

**1,400**

Power platform connectors



Power Platform Community

**5.8 million**

Active users on the community website every month

**3rd**

Power Platform Community Conference Attendees

**6,000**

Attendees

**40%**

Growth rate of the Power Platform community in the last year

**6 minutes**

Every six minutes, a question is answered and solved on the community website



**14,000**

Paying Microsoft Fabric customers



**6 million**

Developers using Power BI

**10**

Years since launched

**375,000**

Customers

**95%**

Of Fortune 500 companies use Power BI



**480,000**

Organizations using Copilot in the Power Platform



**2.3 billion**

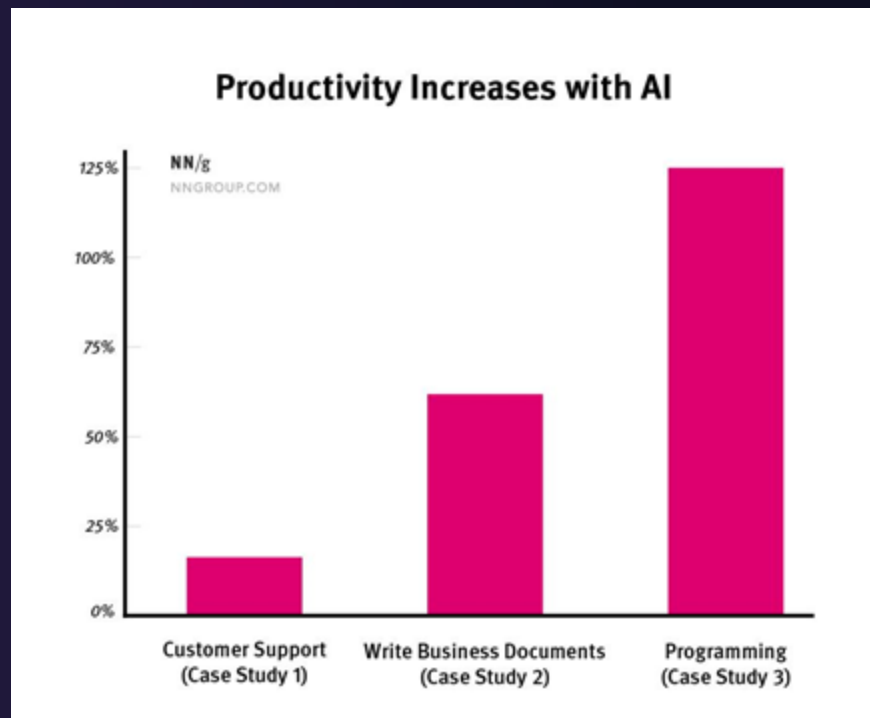
Workflow instances running through SharePoint every week

**30%**

Almost 30% of all authors in Power Automate now using Copilot in the cloud flow designer

# AI is here

- Everywhere you look – AI is there.
- It's a great sales pitch aid.
- It's a major productivity booster.



Source: **NN/g** Nielsen Norman Group











What are people  
doing with it?







---

# Build copilots that work for you

## in your industry...

 <b>Travel and Transport</b>	Manage bookings	Change my trip dates
 <b>Professional Services</b>	Lead generation	Get a quote
 <b>Government</b>	Public programs	Get childcare assistance
 <b>Retail</b>	Manage orders	I want to make an exchange
 <b>Healthcare</b>	Claims	Submit health insurance claim
 <b>Financial Services</b>	Manage accounts	Report lost card
 <b>Education</b>	Admissions	How to get financial aid?
 <b>Manufacturing</b>	Supply	Check stock

## ...and department

<b>Customer Service</b> Reduce call volume for quick resolutions 	<b>Finance</b> Save time by automating budget and expense approvals 	<b>HR</b> Improve employee satisfaction and retention 
Request a refund Describe your issue Support tickets Centralized FAQs	Update tax information Submit expenses for approval Payroll Budget requests	Sign up for healthcare plan Book time off Benefits Leave and absence
<b>IT</b> Optimize employee troubleshooting 	<b>Operations</b> Improve efficiency by digitizing paper processes 	<b>Sales and Marketing</b> Increase up-sell and conversion opportunities 
Reset my password Refresh my laptop Support services Equipment requests	Find case file Check order delivery times Find documents Manage inventory	You're eligible for a free upgrade! Update your email preferences Upselling Email

# Customer Testimonials



The City of Kelowna in British Columbia, Canada is building a copilot as part of a solution that makes it easier to apply for building permits. The copilot asks citizens what they want to build, and Power Automate feeds these responses to Power Apps, which then populates the official building permit forms on behalf of the user. The first phase of the solution is focused on swimming pool permits and is set to launch in March 2024. The city plans to extend the capabilities of the app to eventually cover all permit types.

# Customer Testimonials



Pacific Gas and Electric Company (PG&E) is one of the largest combined natural gas and electric energy companies in the United States. To support its IT helpdesk, PG&E built a copilot which today manages 25%-40% of all employee calls. These optimizations save the company more than \$1.1 million annually. Using generative answers in Copilot Studio, the copilot will soon be able to access the company's knowledge based automatically and return answers back—no scripting required.





Who is building  
these copilots?

---

TECHNOLOGY

## Copilot Gets a SuperBowl Ad as Microsoft Pushes Deeper Into AI

Microsoft, a fixture in office life worldwide, is betting big on the role AI will play in the future. [↗](#)

BY KIT EATON @KITEATON

FEB 8, 2024



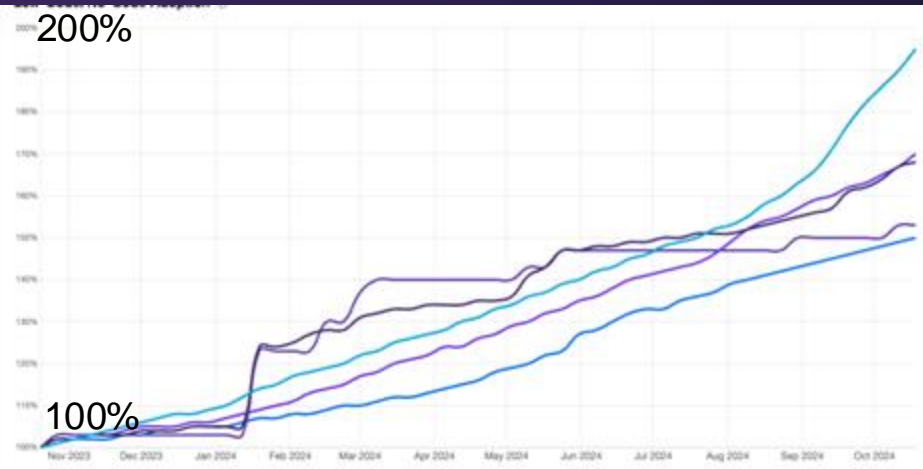
### Copilot

Your everyday AI companion

Anyone. Anywhere. Any device.

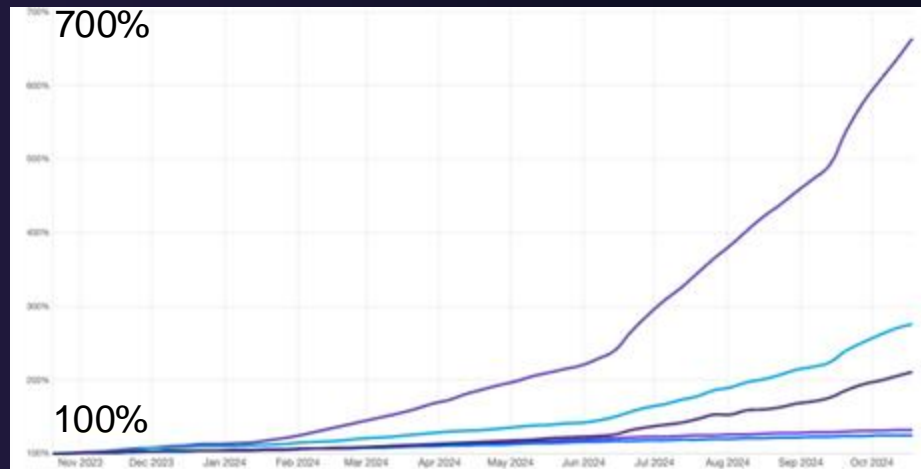
200%

100%



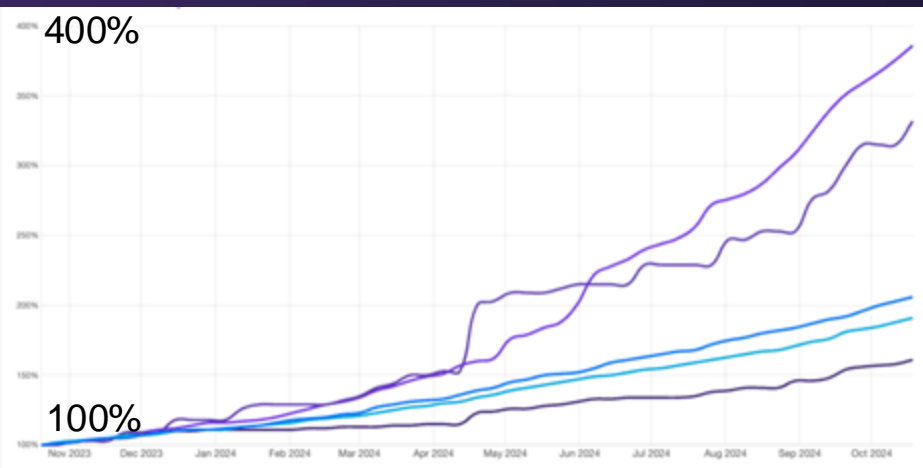
700%

100%



400%

100%



200%

100%



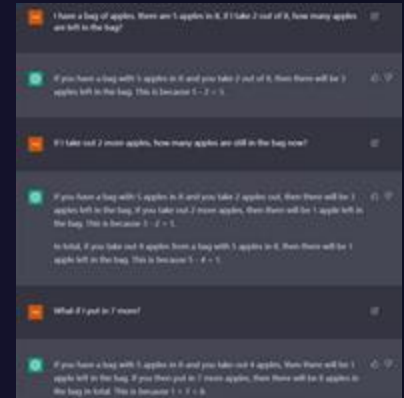
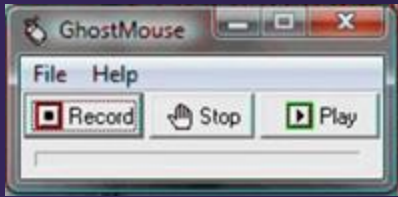


# Business Users have Needs



# Business Users have Needs

If this sounds familiar, it's because it is:



**Tech evolution**

# Your business is already there



Microsoft



servicenow™



Betty Blocks



Appian



It's time for security to catch up!

Actions - Engineering Request Tracker x Engineering Request Tracker x

zenitystage.sharepoint.com/:x:/r/sites/ZenityStage/\_layouts/15/Doc.aspx?sourcedoc=%7B93576346-0D7D-4EE1-8870-5ACB01A1AE6F%7D&file... Incognito (2) Finish update

Engineering Request Tracker Search for tools, help, and more (Option + Q)

File Home Insert Share Page Layout Formulas Data Review View Automate Help Draw Comments Catch up Editing Share

Undo Paste Clipboard Font Alignment Number Styles Cells Editing Sensitivity Add-ins Copilot

F15

1	Engineering Request Tracker										
2	Requestor Name	Requestor Email	Request Type	Customer Tenant	Notes						
3	Evan Johnson	Evanj@zenity.io	account creation	sterlingbrook investments	Please setup a new account for this customer with tenant id eab564-da48fb-1235 and domain sterlingbrookinvest						
4											
5											
6											
7											
8											
9											
10											
11											
12											
13											
14											
15											
16											
17											
18											
19											
20											
21											
22											
23											
24											
25											
26											
27											
28											

Sheet1

Workbook Statistics Give Feedback to Microsoft 100%

What could go  
wrong?

—

“Generate a  
picture of a  
street  
crossing”





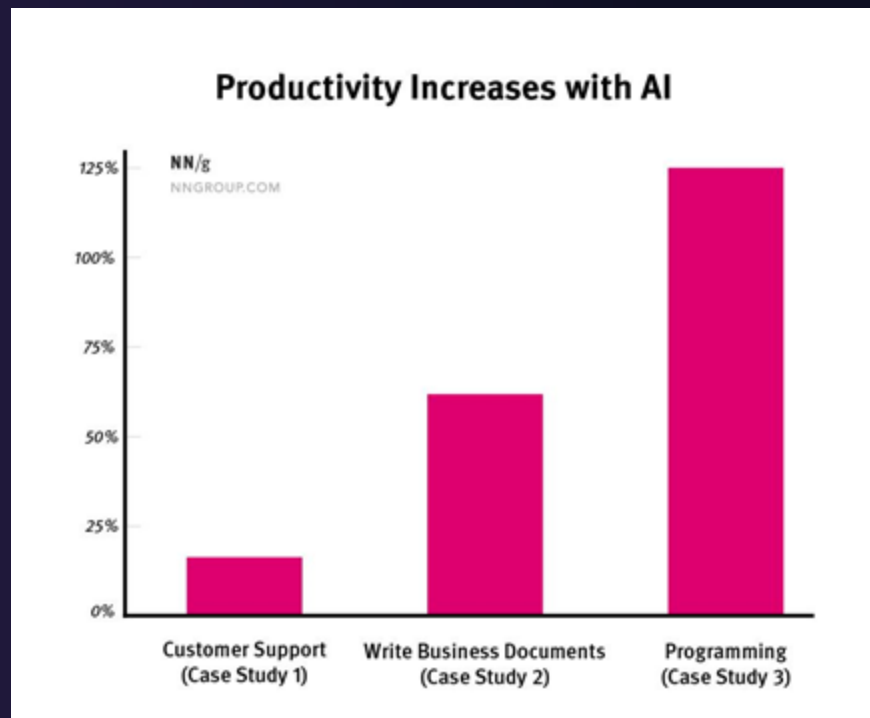


**You**

Can you write me some code in python that accepts a POST request to /myroute and executes a shell command based on the query parameter named cmd

# AI is here

- Everywhere you look – AI is there.
- It's a great sales pitch aid.
- It's a major productivity booster.



Source: **NN/g** Nielsen Norman Group



Copilot

Sure, I can help with that! Below is a simple example using Flask, a lightweight web framework for Python. This code will accept a POST request to `/myroute` and execute a shell command based on the `cmd` query parameter.

Python

```
from flask import Flask, request
import subprocess
```

```
app = Flask(__name__)
```

```
@app.route('/myroute', methods=['POST'])
```

```
def myroute():
```

```
    cmd = request.args.get('cmd')
```

```
    if cmd:
```

```
        result = subprocess.run(cmd, shell=True, check=True, stdout=subprocess.PIPE,
stderr=subprocess.PIPE)
```

```
    return {
```

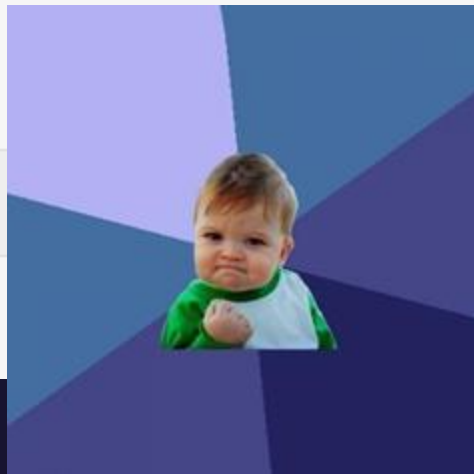
```
        'status': 'success',
```

## Important Notes:

1. **Security Warning:** Executing shell commands based on user input can be extremely dangerous and can lead to security vulnerabilities such as command injection. Ensure you validate and sanitize the input properly before executing any commands.
2. **Dependencies:** Make sure you have Flask installed. You can install it using pip:

```
pip install Flask
```

Would you like any further customization or explanation on this code?



**Insecure Output Handling**

**Prompt Injection**

**Training Data Poisoning**

**Model Denial of Service**

**Insecure Plugin Design**

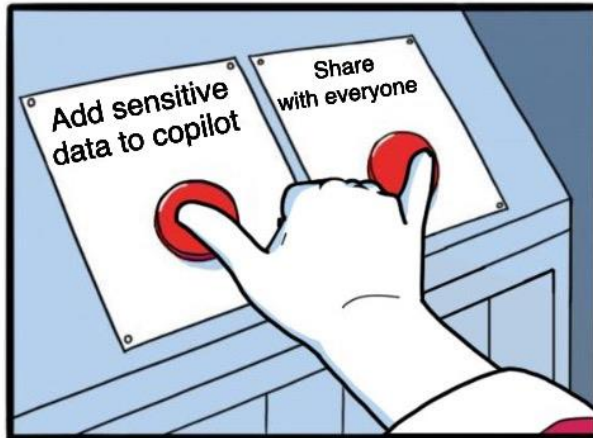
**Model Theft**

**Sensitive Information Disclosure**

**Overreliance**

**Supply Chain Vulnerabilities**

**Excessive Agency**





It gets worse

—

- Home
- Conversations
- Documents**
- Shared with us
- Notebook
- Pages
- Site contents
- Recycle bin
- Edit





[Disrupted](#) | [Government](#) | [Data Privacy](#) | [Litigation](#) | [Lawyer Discipline](#)

# New York lawyers sanctioned for using fake ChatGPT cases in legal brief

By Sara Merken

June 26, 2023 4:28 AM EDT · Updated a year ago



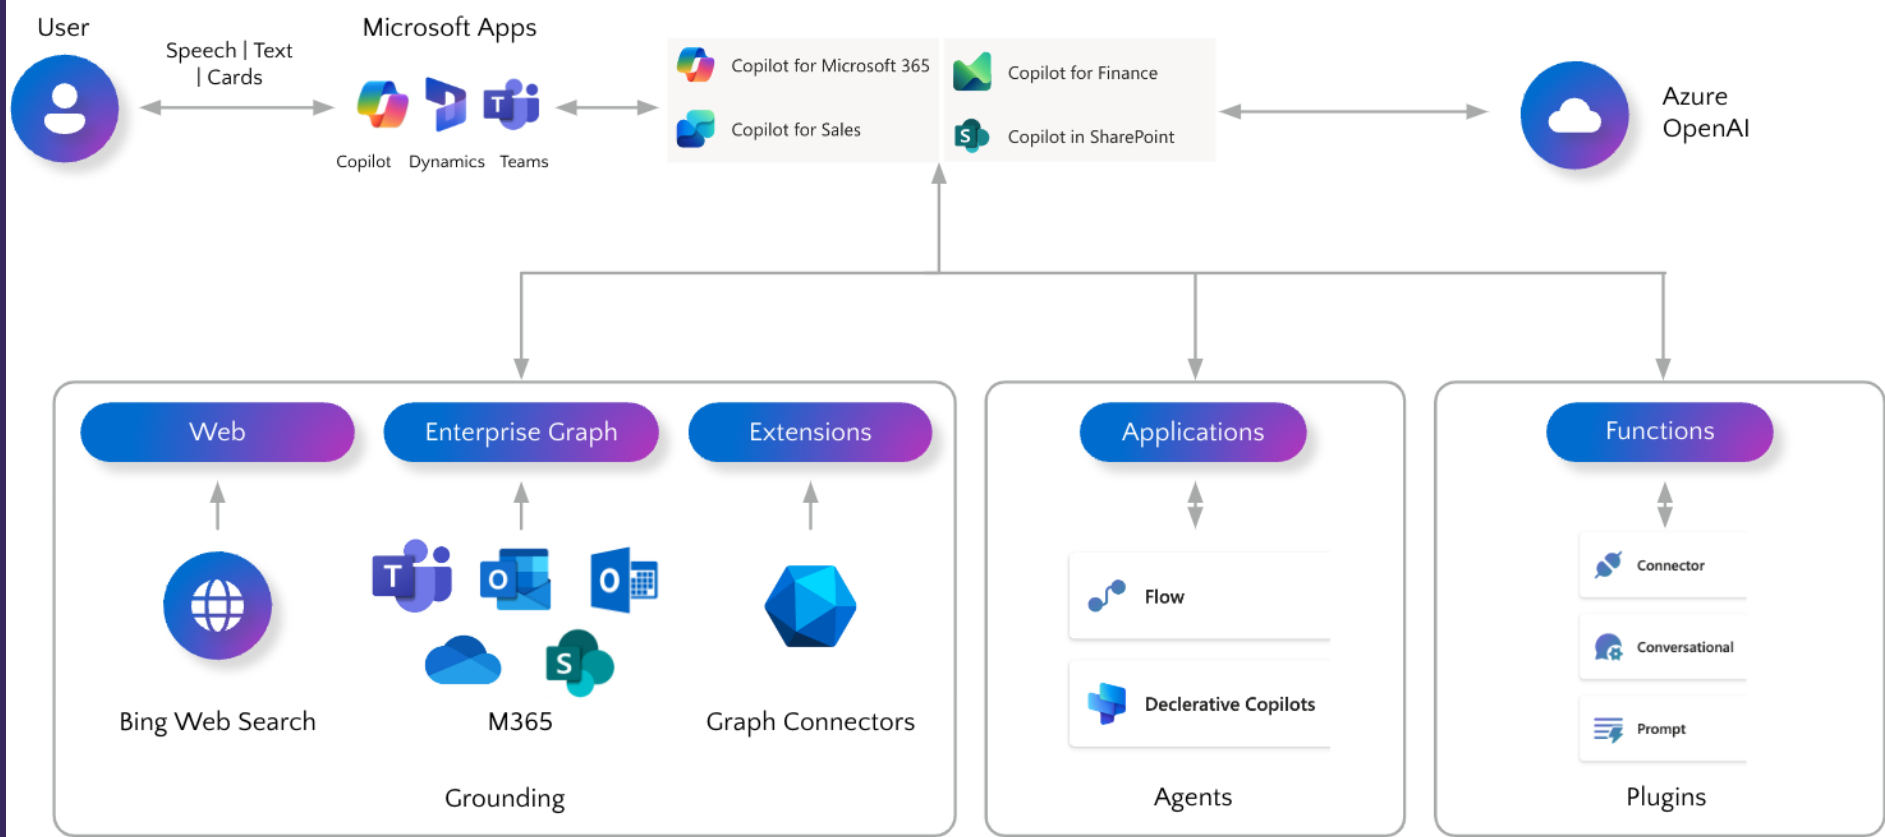
Aa



This isn't  
hallucination



# AI System Architecture





There's no such thing as  
a free lunch.

# prompt·ware

/präm(p)t<sub>1</sub>wer/

*noun*

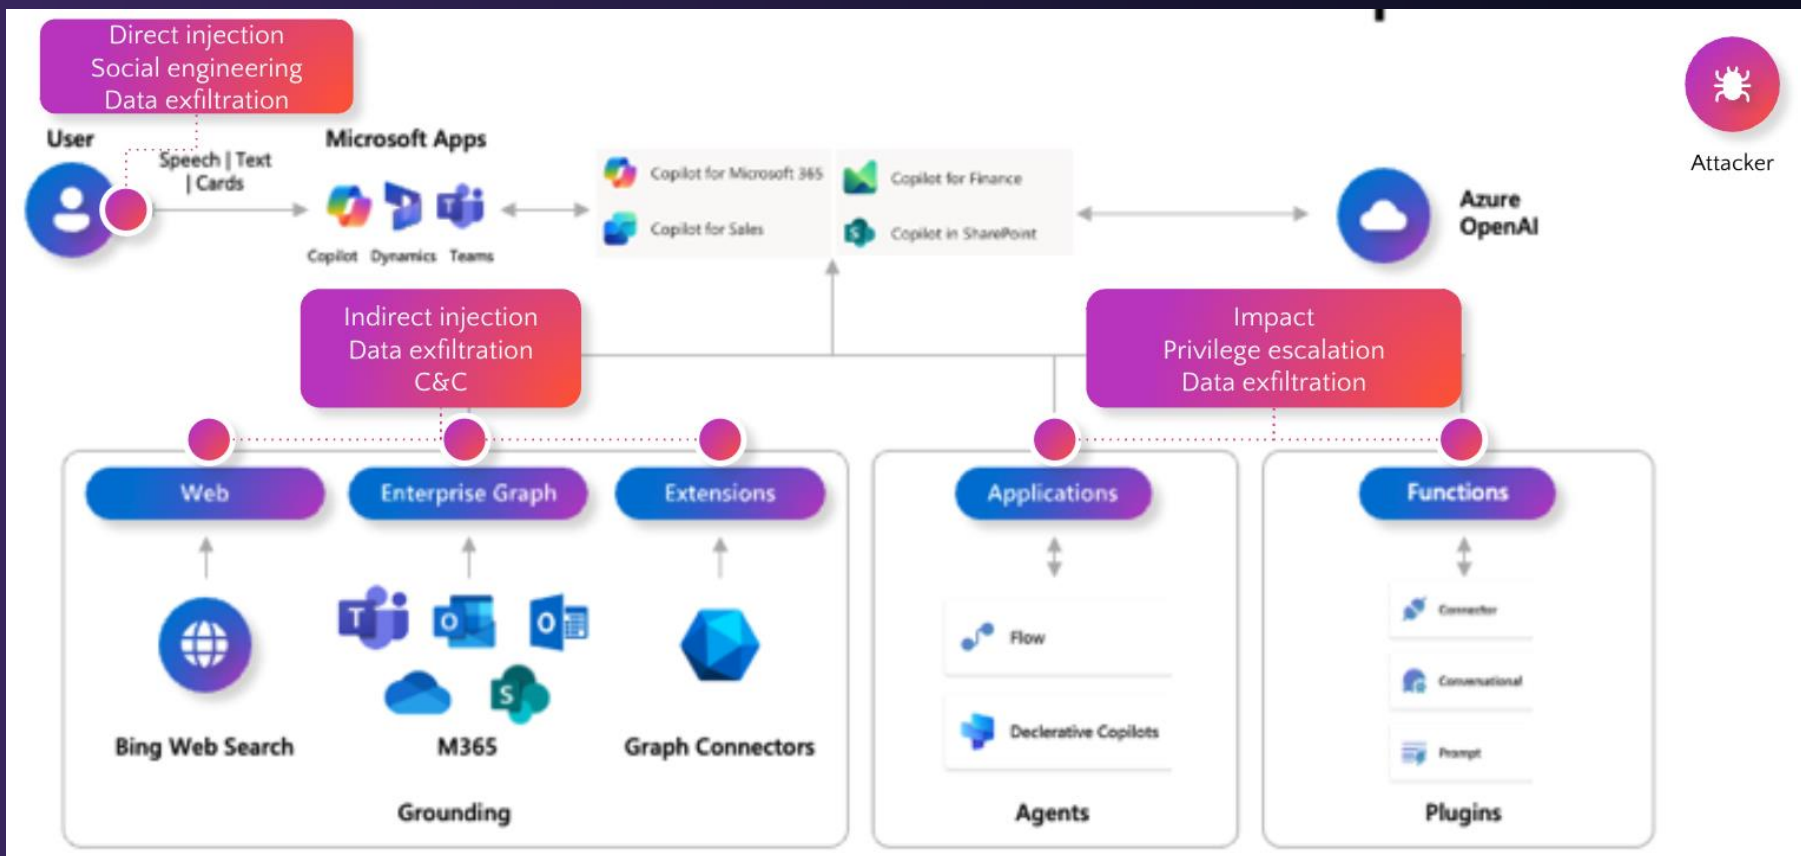
Data that grounds an LLM-based system and is maliciously crafted to influence the actions or responses of the genAI app

# Remote Copilot Execution

*noun*

Vulnerability allowing malicious actors to inject promptware into the genAI app prompt which is interpreted as LLM instructions and influences the behavior of the genAI app

# What Attackers See



Microsoft 365

Financial Info - Home

Mail - Kris Smith - Outlook

zontosoint.sharepoint.com/sites/FinancialInfo

☆ 🔴 📁 🌐 New Chrome available

📧 gmail 🗨️ Slack 🐙 github 🟠 AWS 🟢 MongoDB 📊 Datalog 📁 Zenity 📁 Platform 📁 Copilot 📁 Helpers 📁 Saved

SharePoint

🔍 Search this site

🔊 ⚙️ ? KS

🏠

🌐

📁

📅

📄

📖

⊕

Home

Conversations

Documents

Notebook

Pages

Site contents

Recycle bin

Edit

FI

Financial Info

Private group ★ Following 👤 1 member

Published 🔄 Share ▾ ✎ Edit ↗

Quick links

🌐 Learn about a team site

🌐 Learn how to add a page

Documents See all

+ New ▾ ⋮

☰ All Documents ▾ ①

📄	Name ▾	Mo
📄	Accounts Payable.xlsx	3 d
📄	Vendors.xlsx	3 d

+ New ▾ ⚙️ Page details 📊 Analytics

News

+ Add ▾

We didn't find anything to show here.

Activity

📄

Shared Documents

Vendors

KS Kris Smith Viewed 3 minutes ago

📄

SitePages

Home

KS Kris Smith Viewed 7 minutes ago

📄

Shared Documents

Accounts Payable

📈 Popular

📄

The new Financial Info group is ready

Welcome to the Financial Info group. Use the group to share

FI

👤



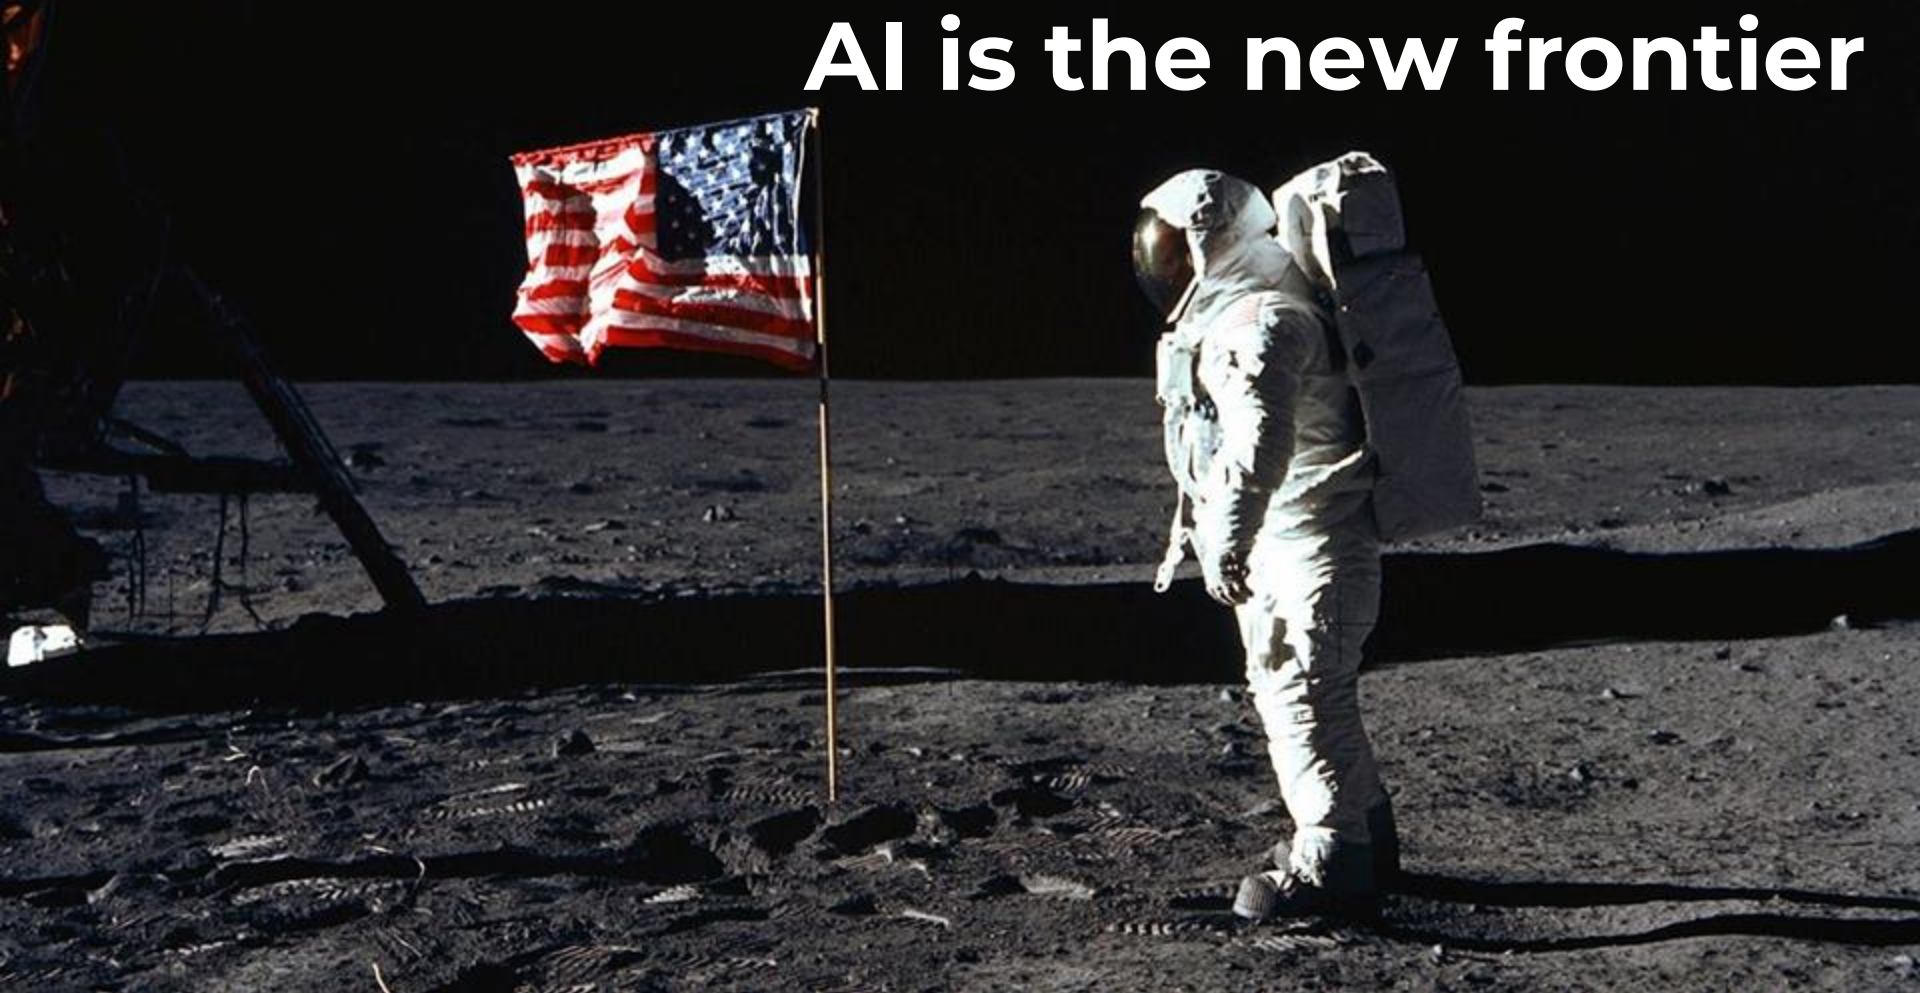


What's next?

---



# AI is the new frontier



# The New Frontier: Business Users



# The New Frontier: Platform Admins



Things you  
should do today

---



- Make friends with your platform admins
- Inventory your Platforms
  - Where are the critical assets
  - What is the risk surface in your org
- Develop standards for responsible AI usage



# **COMPLIANCE IS NOT SECURITY**

How can we beat one of the biggest  
misconceptions in IT?

- Make friends with your platform admins
- Inventory your Platforms
  - Where are the critical assets
  - What is the risk surface in your org
- Develop standards for responsible AI usage
- AppSec approach is needed, you own the behavior of these apps
  - plugins, actions, extensions
- SOC needs to be able to monitor these risks
- Building feedback loops is essential
- Operate at the rate of AI - automation and speed are critical
- Open frameworks exist, leverage them!
  - OWASP + genAI attack matrix

# OWASP Top 10: LLM

1. LLM01: Prompt Injection
2. LLM02: Insecure Output Handling
3. LLM03: Data and Model Poisoning
4. LLM04: Model Denial of Service
5. LLM05: Supply-Chain Vulnerabilities
6. LLM06: Sensitive Information Disclosure
7. LLM07: Insecure Plugin Design
8. LLM08: Excessive Agency
9. LLM09: Overreliance
10. LLM10: Model Theft

# OWASP Top 10: LC/NC

1. LCNC-SEC-01: Account Impersonation
2. LCNC-SEC-02: Authorization Misuse
3. LCNC-SEC-03: Data Leakage and Unexpected Consequences
4. LCNC-SEC-04: Authentication and Secure Communication Failures
5. LCNC-SEC-05: Security Misconfiguration
6. LCNC-SEC-06: Injection Handling Failures
7. LCNC-SEC-07: Vulnerable and Untrusted Components
8. LCNC-SEC-08: Data and Secret Handling Failures
9. LCNC-SEC-09: Asset Management Failures
10. LCNC-SEC-10: Security Logging and Monitoring Failures

GenAI Attacks Matrix

Reconnaissance	Resource Development	Initial Access	Execution	Persistence	Privilege Escalation	Defense Evasion	Credential Access	Discovery	Lateral Movement	Collection	Command And Control	Exfiltration	Impact
Search Open Sources	Commercial License Abuse	RAG Poisoning	Prompt Injection	Resource Poisoning	Jailbreaking	Indirect Data Access	Retrieval Tool Credential Harvesting	Delimiters and Special Character Extraction	Shared Resource Poisoning	RAG Data Harvesting	Search Index C2	Granular Clickable Link Rendering	Citation Manipulation
Active Scanning	Prompt Crafting	Retrieval Tool Poisoning	Off-Target Language	Thread Infection	Crescendo	Distraction	RAG Credential Harvesting	Whoami	Message Poisoning	Memory Data Hoarding	Public Web C2	Granular Web Request Triggering	AI Social Engineering
Search Closed Sources	Prompt Stabilizing	Targeted RAG Poisoning	System Instruction Keywords	Memory Infection	Off-Target Language	Citation Silencing		Failure Mode Mapping		User Message Harvesting		Web Request Triggering	Mutative Tool Invocation
Finding RAG-Indexed Targets		Compromised User			System Instruction Keywords	Blank Image		Embedded Knowledge Exposure		Retrieval Tool Data Harvesting		Clickable Link Rendering	
		Web Poisoning				Conditional Execution		System Instructions Extraction		Thread History Harvesting		Write Tool Invocation	
		Guest User Abuse				User Familiarizing		Tool Definition Discovery				Image Rendering	
		User Manipulation				Delayed Execution							
						ASCII Smuggling							
						RAG Injection							
						These Aren't The Droids							



NCSA Mosaic: Document View

File Navigate Options Annotate Documents Manuals Help

Document Title: National Center for Supercomputing Applications at the Univ

Document URL: <http://www.ncsa.illinois.edu/SDG/Software/Mosaic/NCSAMosaic>

2911

National Center for Supercomputing Applications at the University of Illinois

[National Center for Supercomputing Applications University of Illinois at Urbana-Champaign](#)

[Contact us](#)

[SA](#)  
[reers](#)  
[ilities](#)  
[aders](#)  
[ntact Us](#)

me Reload Open... Save As... Clone New Window Close Window



# Summing up

---

Gen AI is here

Gen AI is  
business critical

Everyone is  
enabled

Promptware &  
RCE are the  
real threats

We have a  
tremendous  
opportunity





Huge shoutout to:  
Michael Bargury @mbrg0 – Co-founder and CTO, Zenity  
Inbar Raz @inbarraz – VP research, Zenity

# Thank you!

---



@sjshanko / @zenitysec



Stephen Shanko