UXPA 2026

TRUST, AUTONOMY, AND THE WORK THAT'S MOVING BEHIND THE SCREEN

Designing *Trust* in AI Systems

PRESENTED BY

Mary Gansallo

"

Think about a system that you *trust* today.

I'm *Mary Gansallo*.

Senior Product Designer, Microsoft

Designing AI-powered experiences within Microsoft Entra, a suite of identity and security products.

My Journey *snippets*

Computer Science

Studied C.S., Minor in global terrorism at UMD.

Full-stack engineer

Built software in the consulting world.

HCI, in Barcelona

Moved abroad for a year for my Master's.

Design bootcamp

Back home, I made the craft switch official.

Microsoft · UX Engineer

Code + design in one seat across teams.

Product Designer

Designing trust into AI-powered security.

What is trust?

Trust is the willingness to rely on someone, something, or a system when the outcome matters and there is risk involved.

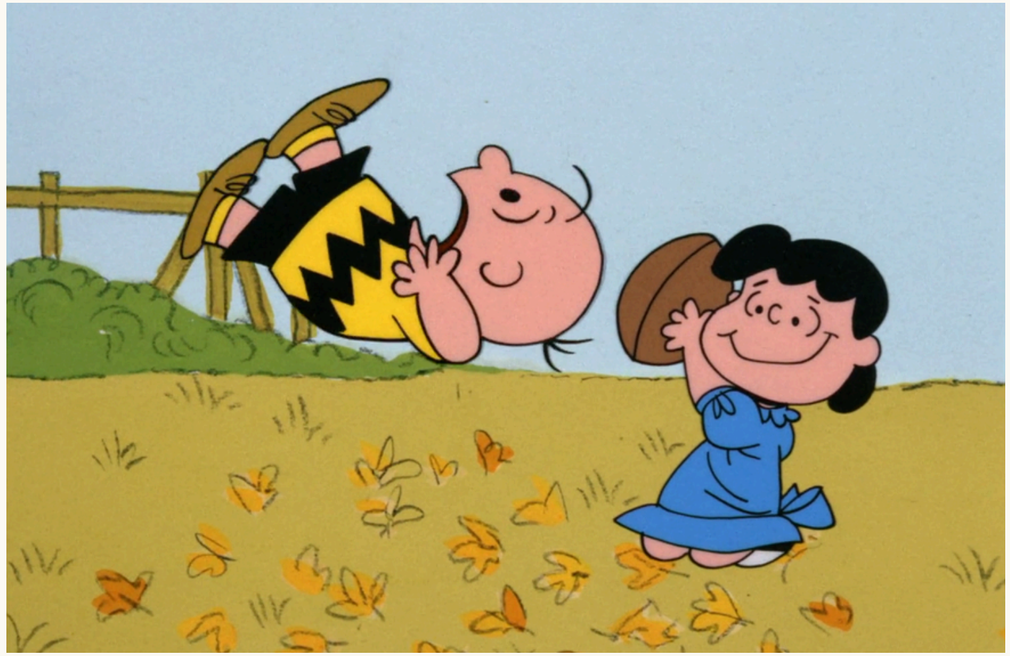

Every act of trust is composed of three things

Reliance

Risk

Vulnerability

Defining the AI Partner

Generative AI systems, AI assistants and copilots, Human-AI decision making, Enterprise AI experiences

AI systems are software that can generate, reason, predict, recommend, or act based on patterns learned from **data**.

Unlike traditional software, modern AI operates as a probabilistic reasoning engine, creating outcomes that are inherently dynamic and unpredictable.

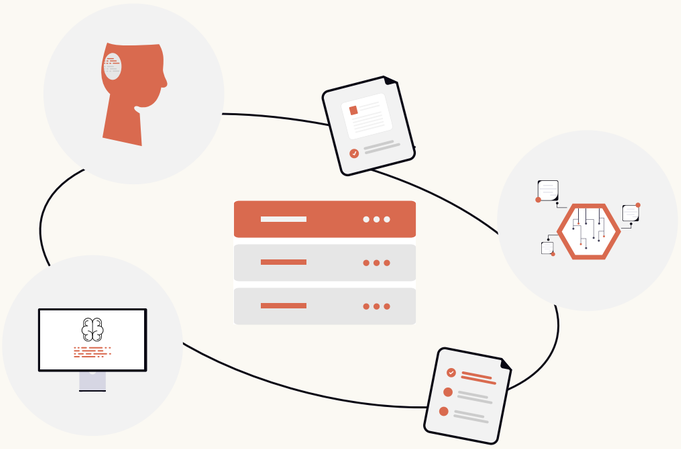

78%

Organizations using AI

35-40%

Users trust AI decisions

Adoption is growing faster than trust.

Organizations are rapidly integrating AI into products, workflows, and decision-making. AI is no longer an emerging technology. It's becoming a core part of how work gets done.

SOURCE: MCKINSEY GLOBAL SURVEY ON AI (2025)

The goal isn't total trust, it's appropriate trust.

Successful AI systems don't maximize trust. They help users understand when to trust, when to verify, and when to intervene. Trust should be earned through evidence, transparency, and control.

Trustworthy experiences start with trustworthy systems

Microsoft’s Responsible AI principles define what trustworthy systems should do. UX patterns determine whether users experience them as trustworthy.

Reliability & safety

Can users predict and safely manage impact?

Transparency

Can users understand what's happening?

Accountability

Who remains responsible for outcomes?

Meet Maya

Maya Patel

"The Guardian"

Senior IT Administrator

Identity & Security

Global Enterprise (50,000 employees)

About Maya

Maya is an IT administrator responsible for protecting access across thousands of employees, applications, and devices. She needs AI she can understand, trust, and control.

Goals

- Keep the organization secure
- Reduce administrative burden
- Respond quickly to security incidents

Challenges

- Too many alerts
- Limited time
- Fear of unintended impact

Key Motivation

"I need to move faster without losing control."

Admin portal entry

USERS STARTS THEIR DAY WITH HUNDREDS OF SIGNALS COMPETING FOR ATTENTION.

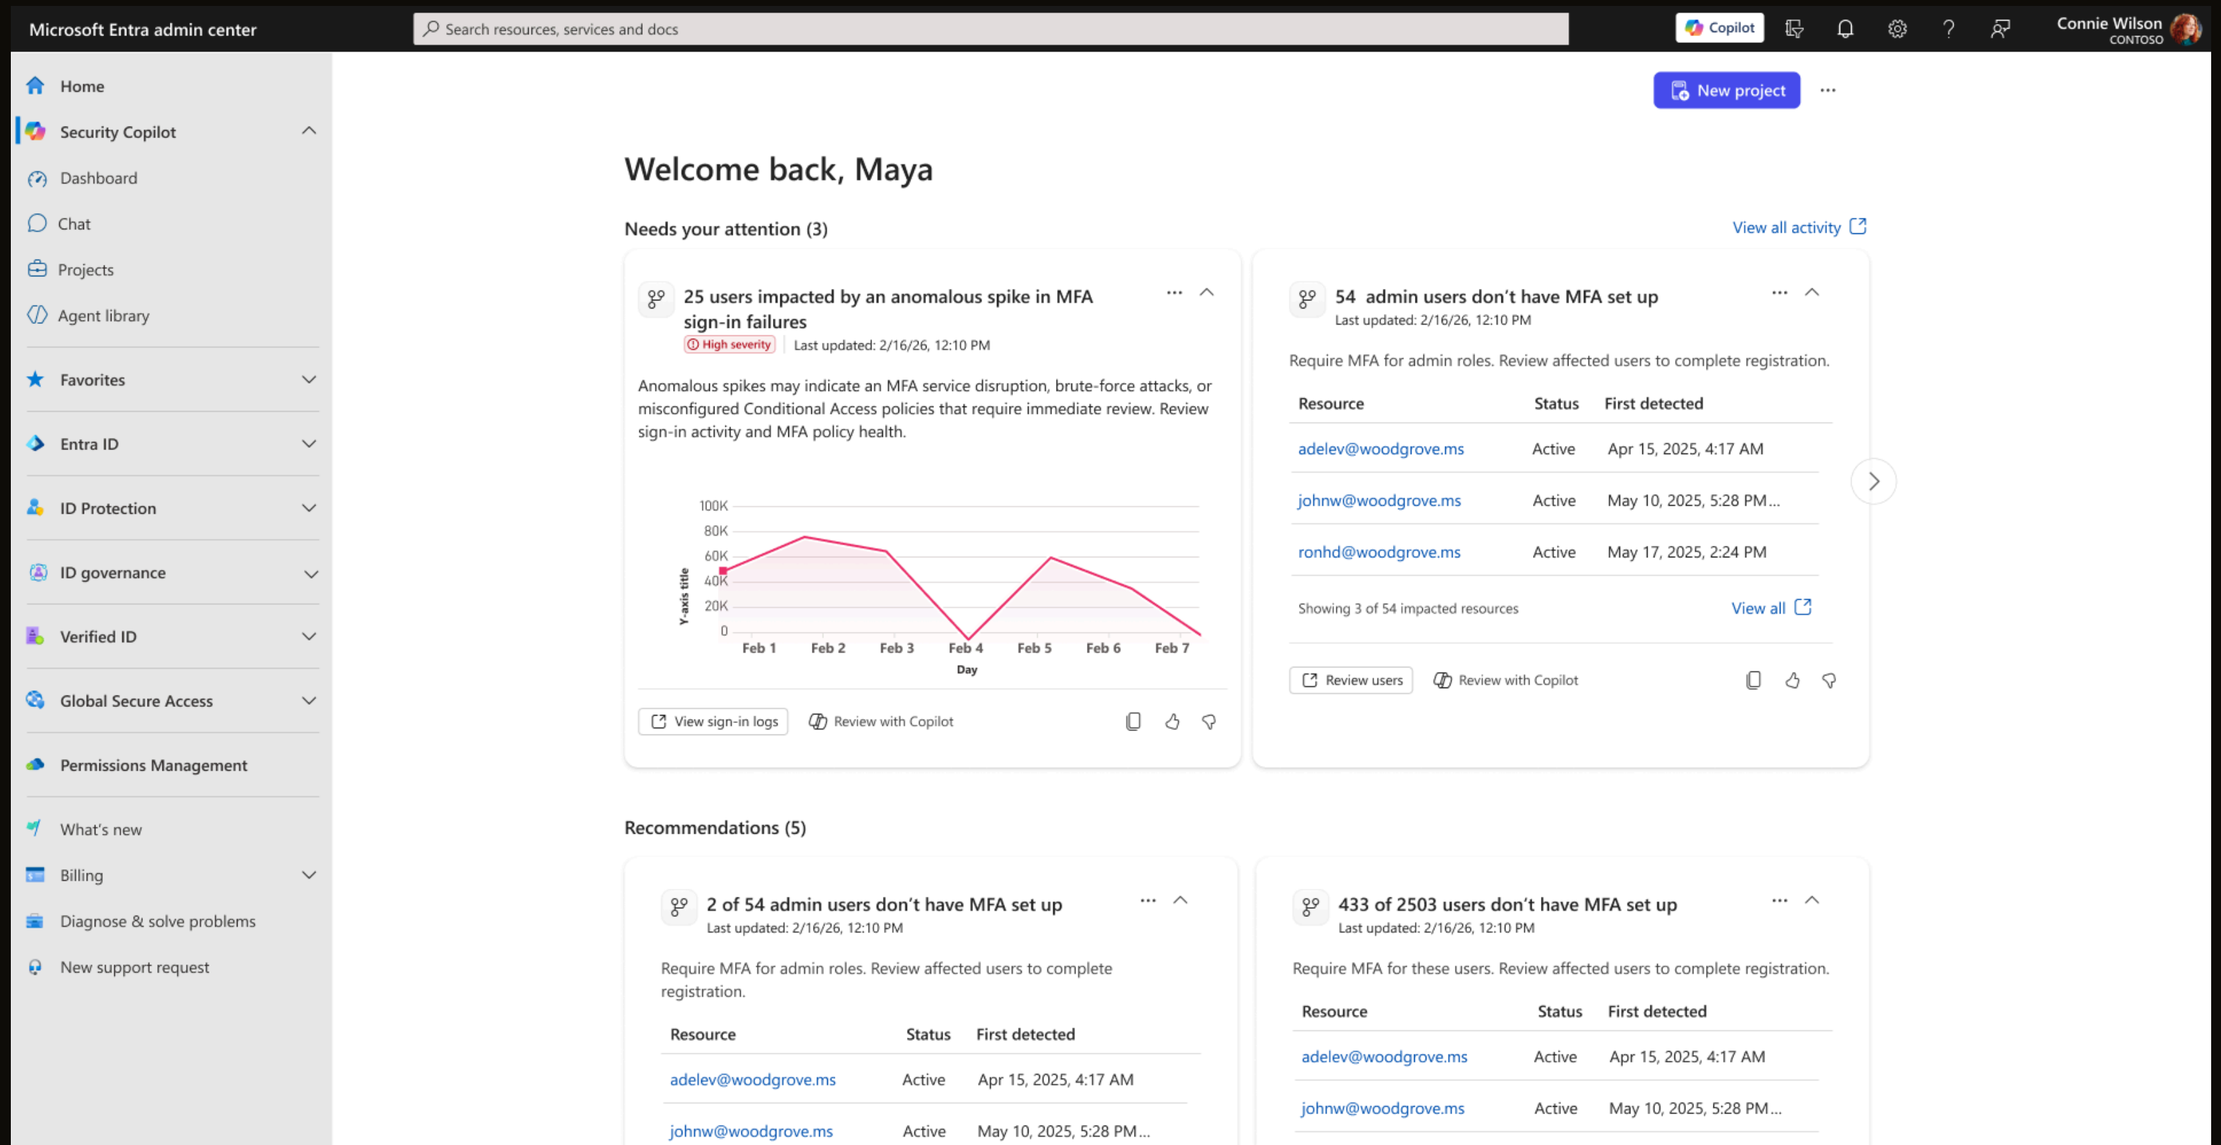

Reliability & Safety

USERS NOW VIEW A PREVIEW OF AN AI GENERATED SUMMARY OF ALL THE SIGNAL AND NOISE IN THEIR ENVIROMENT.

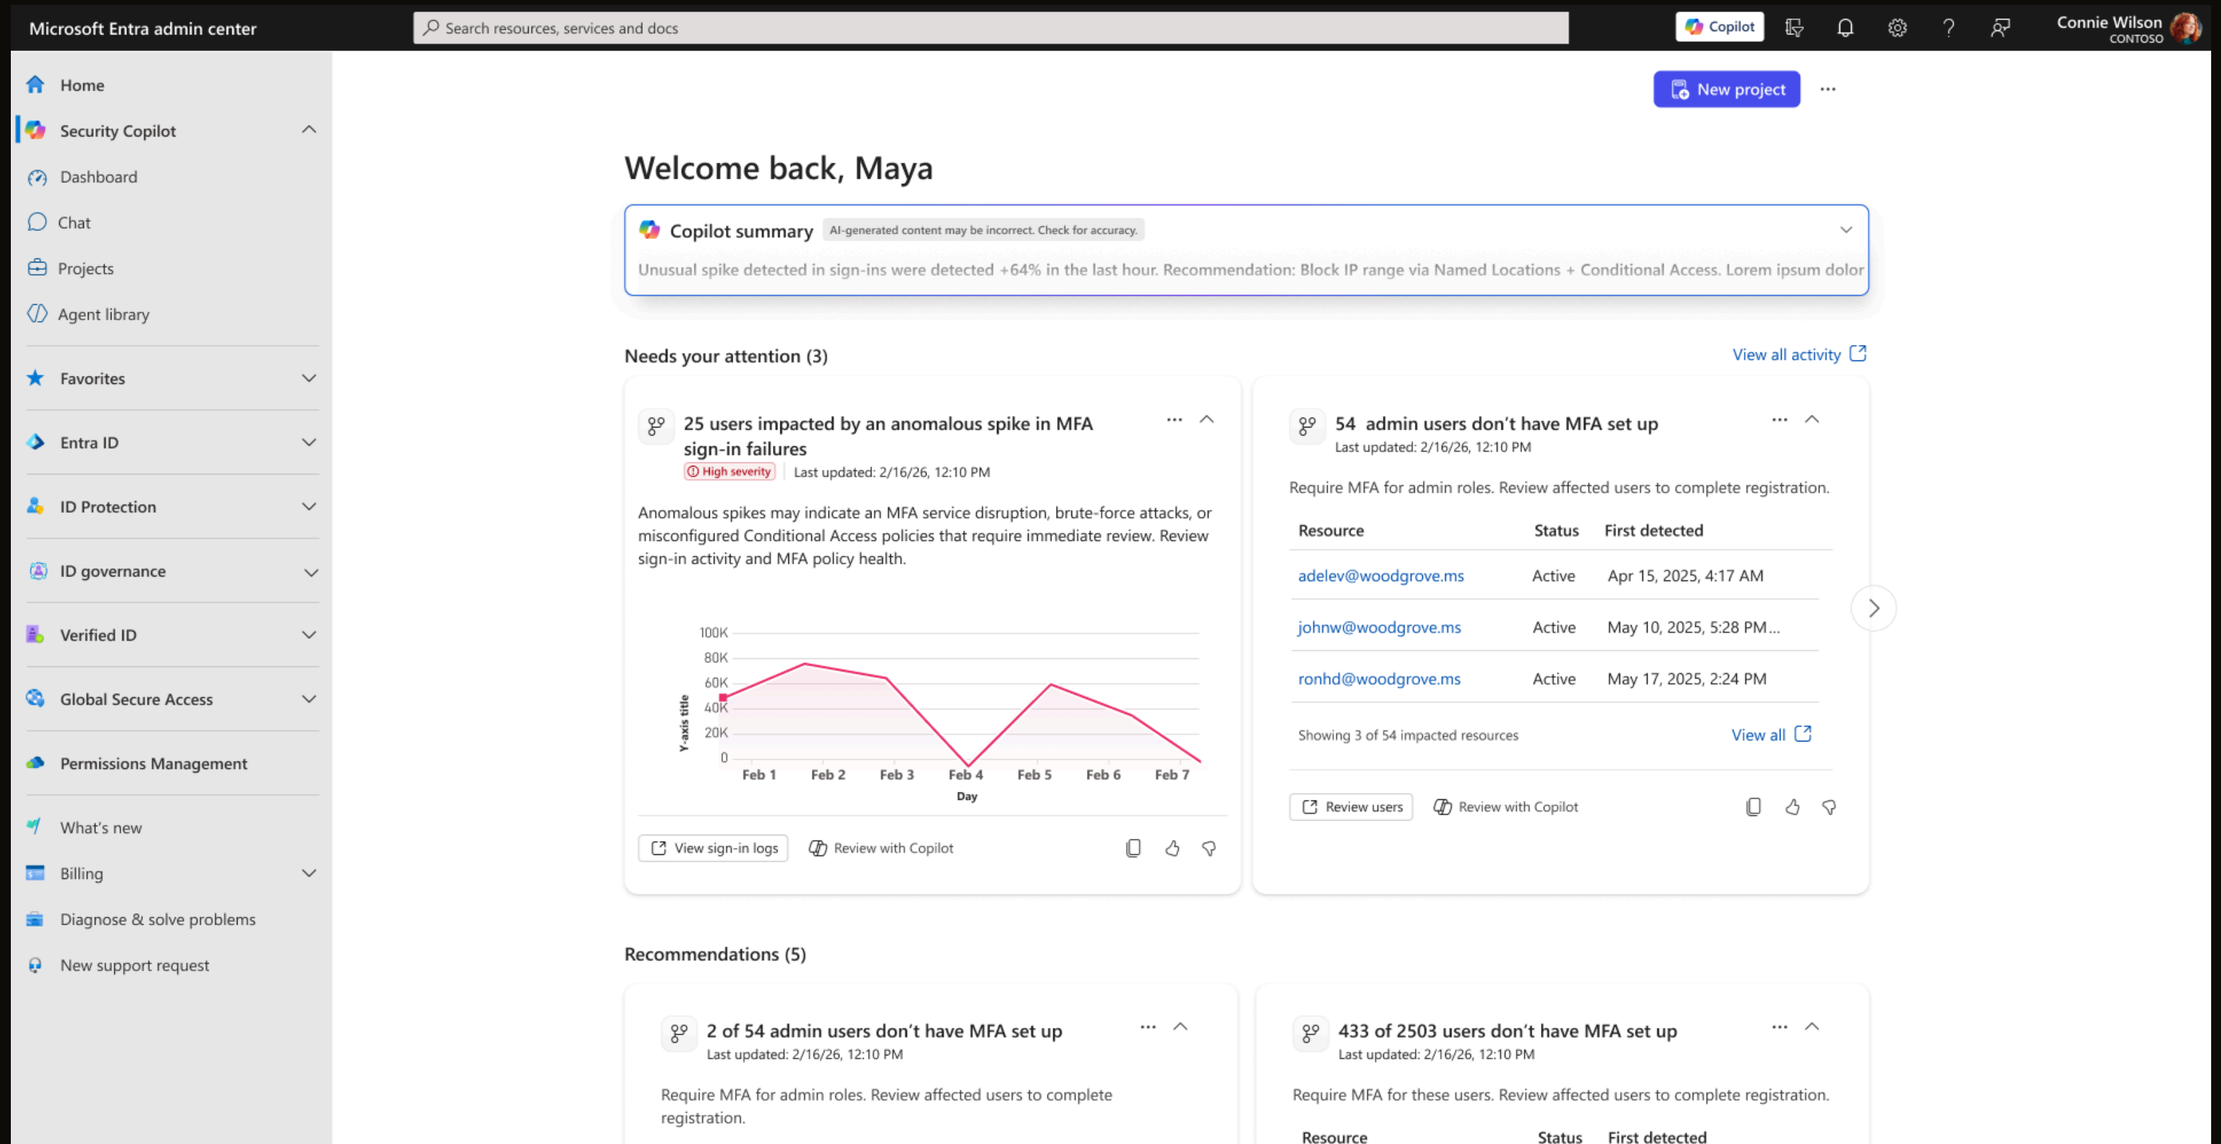

Reliability & safety

Confidence signals

USERS GAIN CONFIDENCE WHEN THE SYSTEM DEMONSTRATES IT CAN SEPARATE SIGNAL FROM NOISE, ACCURATELY AND CORRECTLY.

Welcome back, Maya

Copilot summary AI-generated content may be incorrect. Check for accuracy. Last updated: 3 hours ago

I've been monitoring your tenant activity. **An unusual spike detected in sign-ins were detected +64% in the last hour. I recommend blocking the IP range and strengthening Conditional Access policies to reduce risk**. Over the last 24 hours, I've identified 1 critical issue and 1 medium-severity item requiring your attention, and resolved 3 low-severity items automatically.

- 30 high-risk users were identified with privileged role assignments. Investigate with copilot
- Unfamiliar IP ranges associated with recent access attempts, not previously seen in your tenant's activity baseline. Investigate with copilot
- Policy gaps identified in Conditional Access coverage, leaving some sign-in scenarios without enforced protections. Review audit logs
- Repeated failed sign-ins from a single region, indicating potential brute-force or credential stuffing attempts. Review audit logs

What do I need to know? | What changed since yesterday? | Improve my tenant's security

Needs your attention (3)

View all activity

25 users impacted by an anomalous spike in MFA sign-in failures

High severity | Last updated: 2/16/26, 12:10 PM

Anomalous spikes may indicate an MFA service disruption, brute-force attacks, or misconfigured Conditional Access policies that require immediate review. Review sign-in activity and MFA policy health.

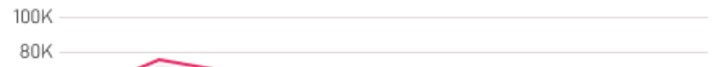

54 admin users don't have MFA set up

Last updated: 2/16/26, 12:10 PM

Require MFA for admin roles. Review affected users to complete registration.

Resource	Status	First detected
adelev@woodgrove.ms	Active	Apr 15, 2025, 4:17 AM
johnw@woodgrove.ms	Active	May 10, 2025, 5:28 PM...

Reliability & safety

Disclosure

USERS SHOULD ALWAYS KNOW WHEN CONTENT IS AI-GENERATED. MAYA KNOWS THIS SUMMARY WAS GENERATED BY AI.

CLEAR DISCLOSURE

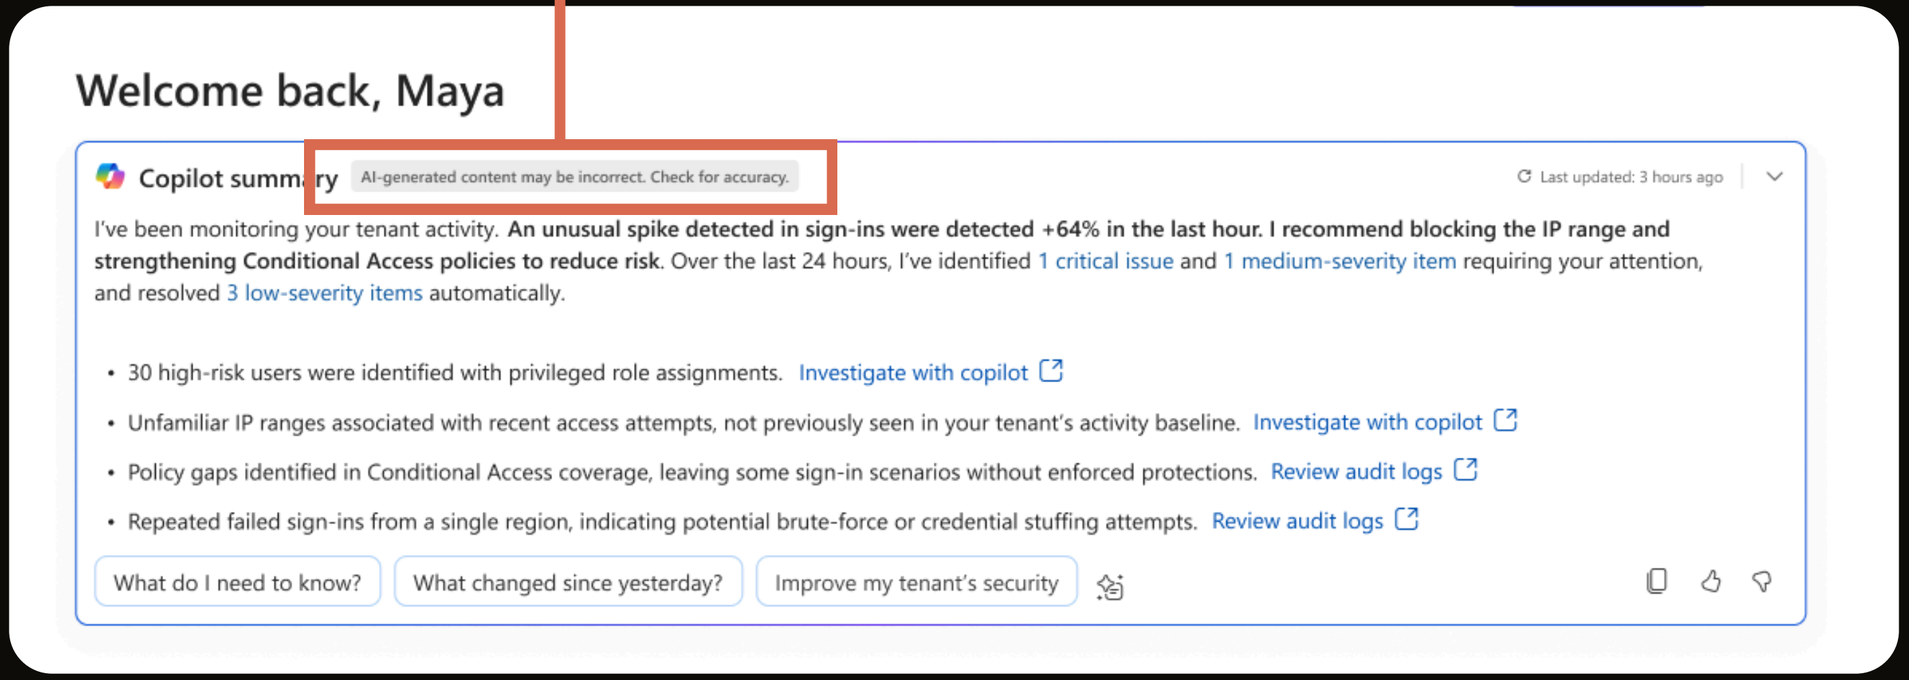

Reliability & safety

Freshness

SECURITY POSTURE CHANGES CONSTANTLY. SHOWING WHEN INSIGHTS WERE GENERATED HELPS MAYA DETERMINE WHETHER THE RECOMMENDATION IS STILL RELEVANT BEFORE ACTING.

FRESHNESS OF DATA

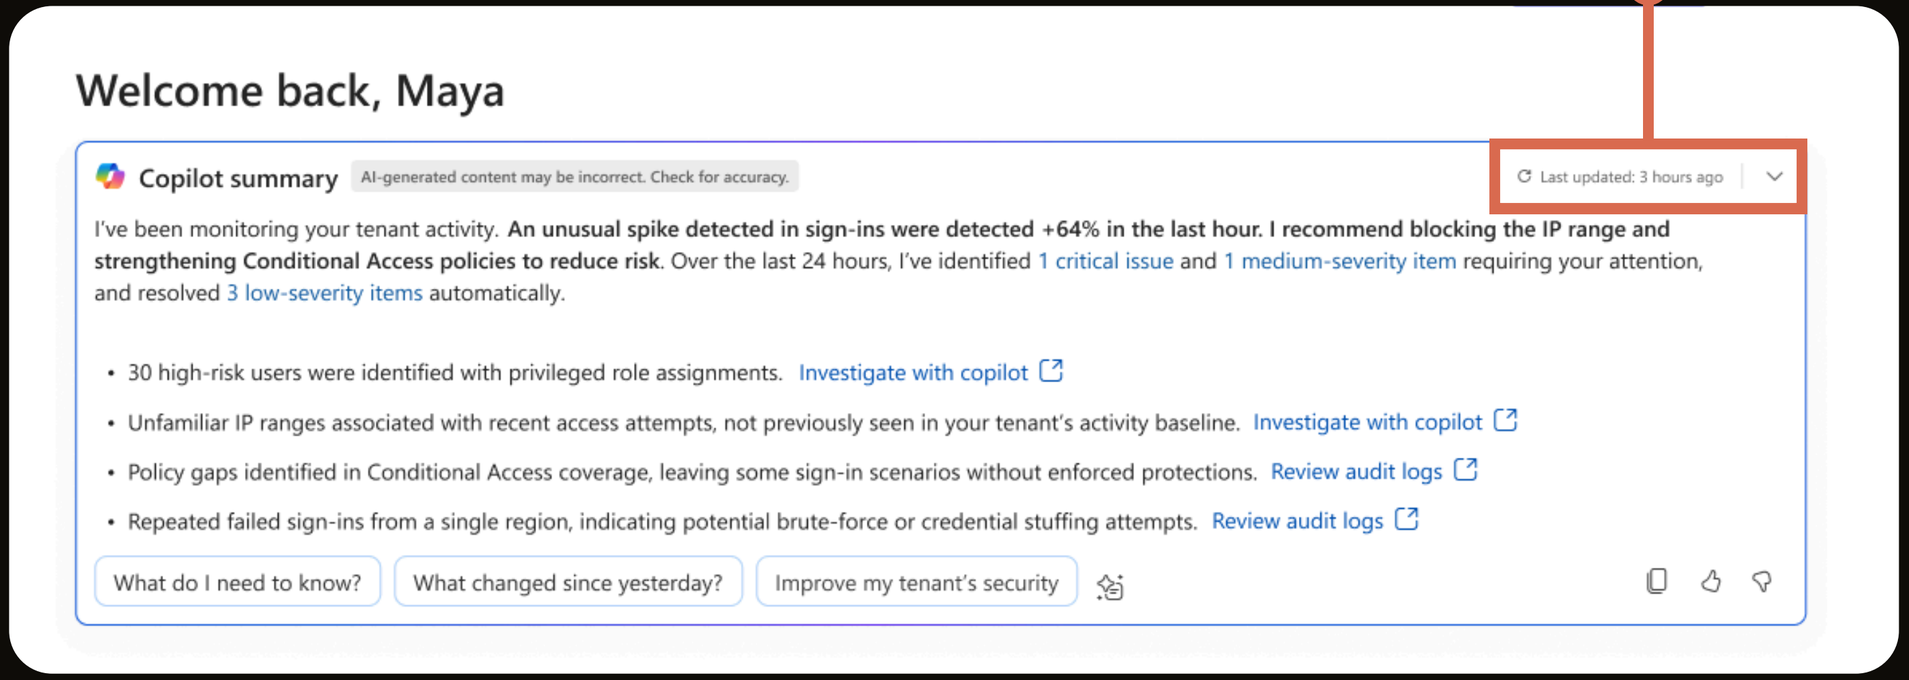

Reliability & safety

Feedback

FEEDBACK MECHANISMS CREATE A LEARNING LOOP BETWEEN USERS AND THE SYSTEM, HELPING IMPROVE FUTURE RECOMMENDATIONS WHILE GIVING MAYA A VOICE WHEN SOMETHING FEELS INCORRECT.

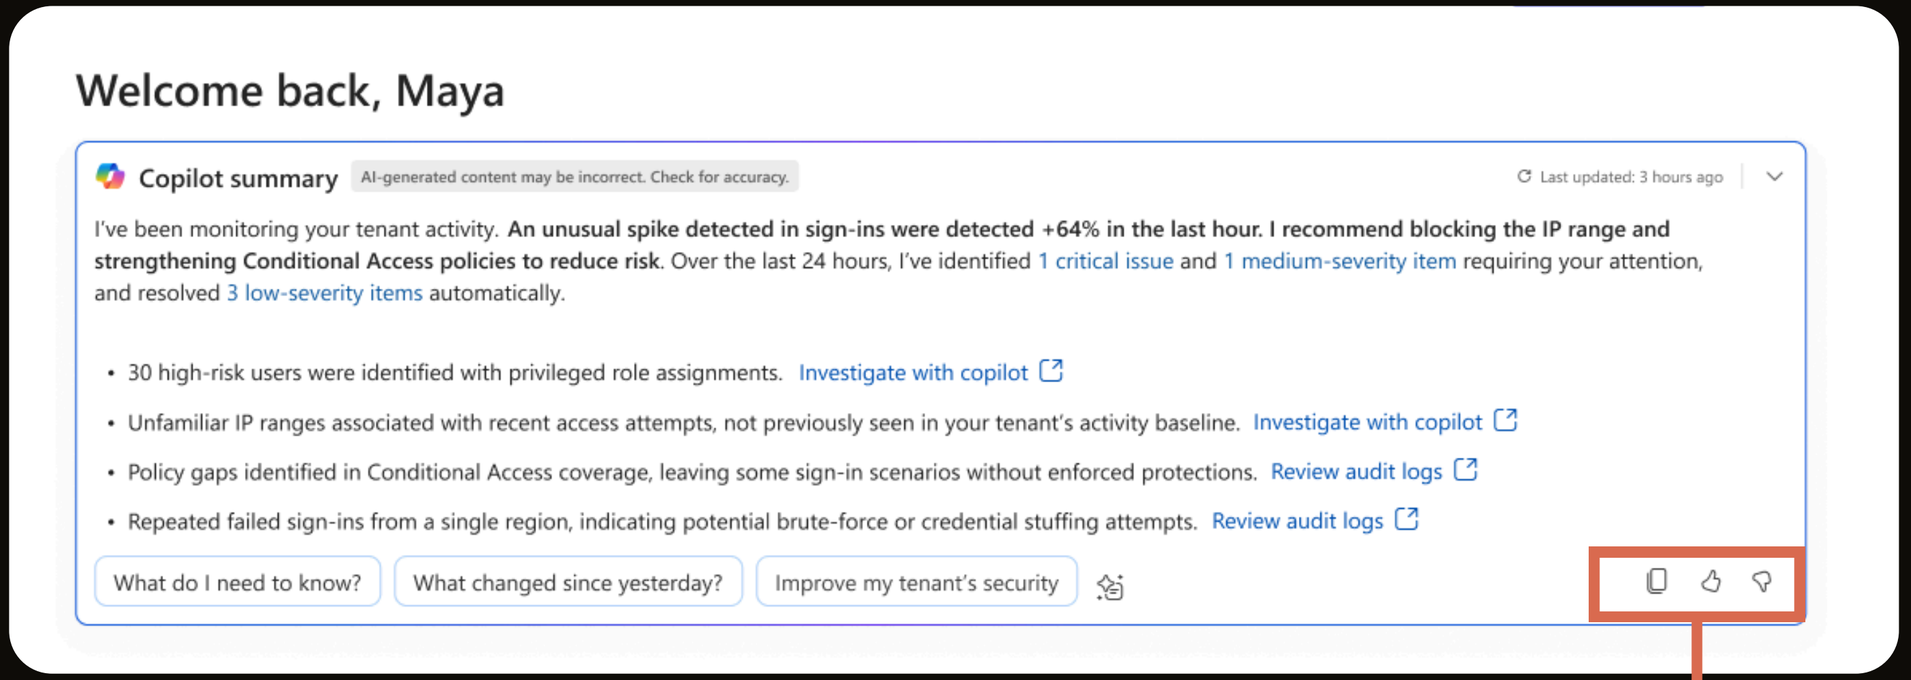

CONTINUOUS IMPROVEMENT

Reliability & safety

Guidance

FOLLOW-UP PROMPTS REDUCE AMBIGUITY BY HELPING MAYA INVESTIGATE, VALIDATE, AND ACT WITHOUT NEEDING TO FORMULATE THE NEXT QUESTION HERSELF.

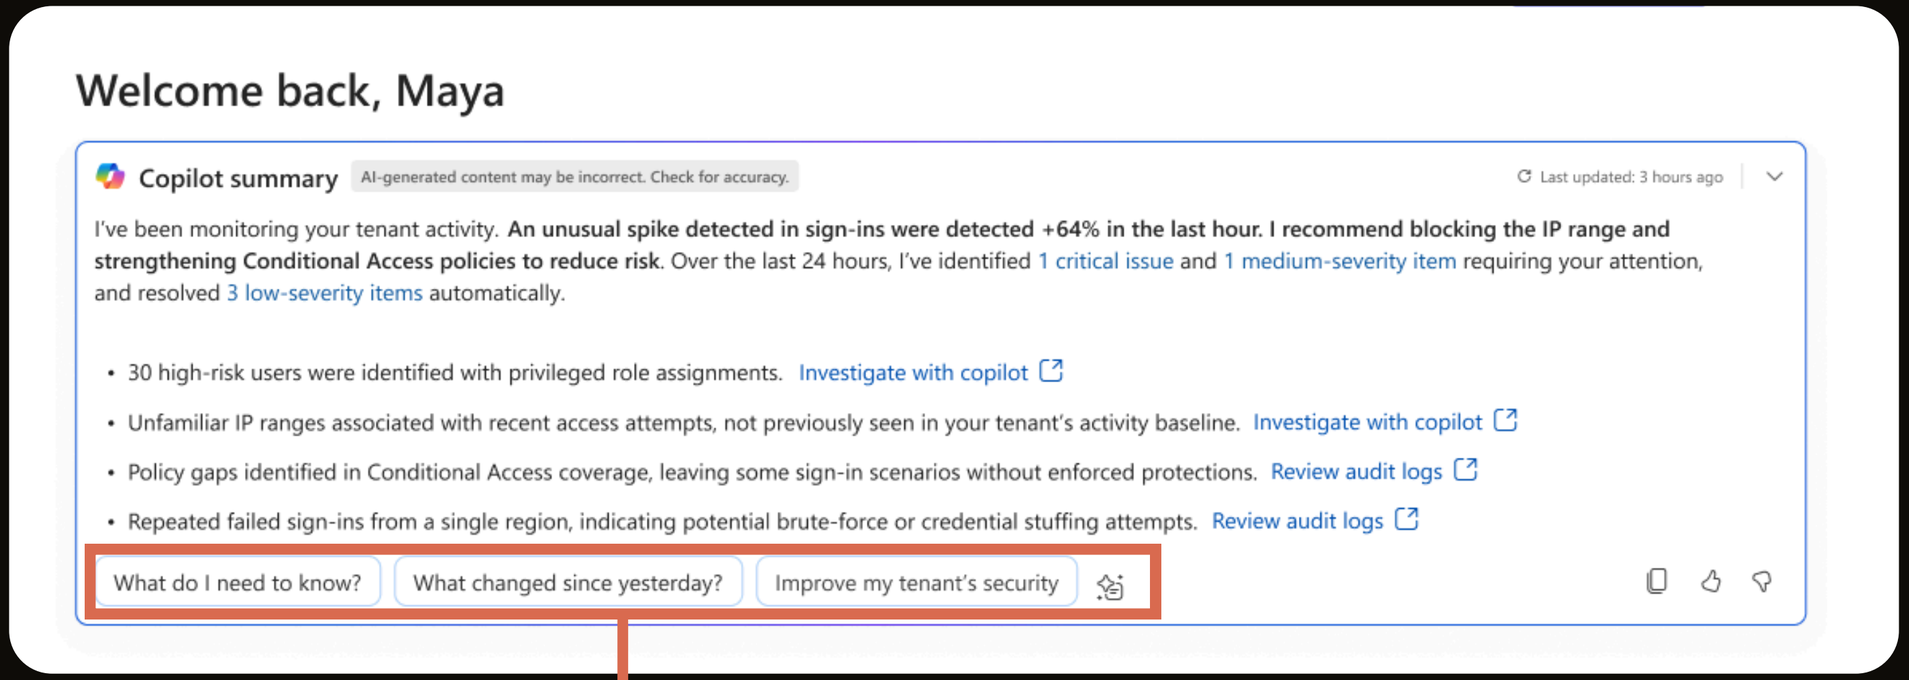

GUIDED EXPLORATION

Reliability & safety

Verification

TRUST GROWS WHEN RECOMMENDATIONS CAN BE INSPECTED. DIRECT LINKS TO SUPPORTING EVIDENCE ALLOW MAYA TO VALIDATE FINDINGS BEFORE TAKING ACTION.

VERIFIABLE EVIDENCE

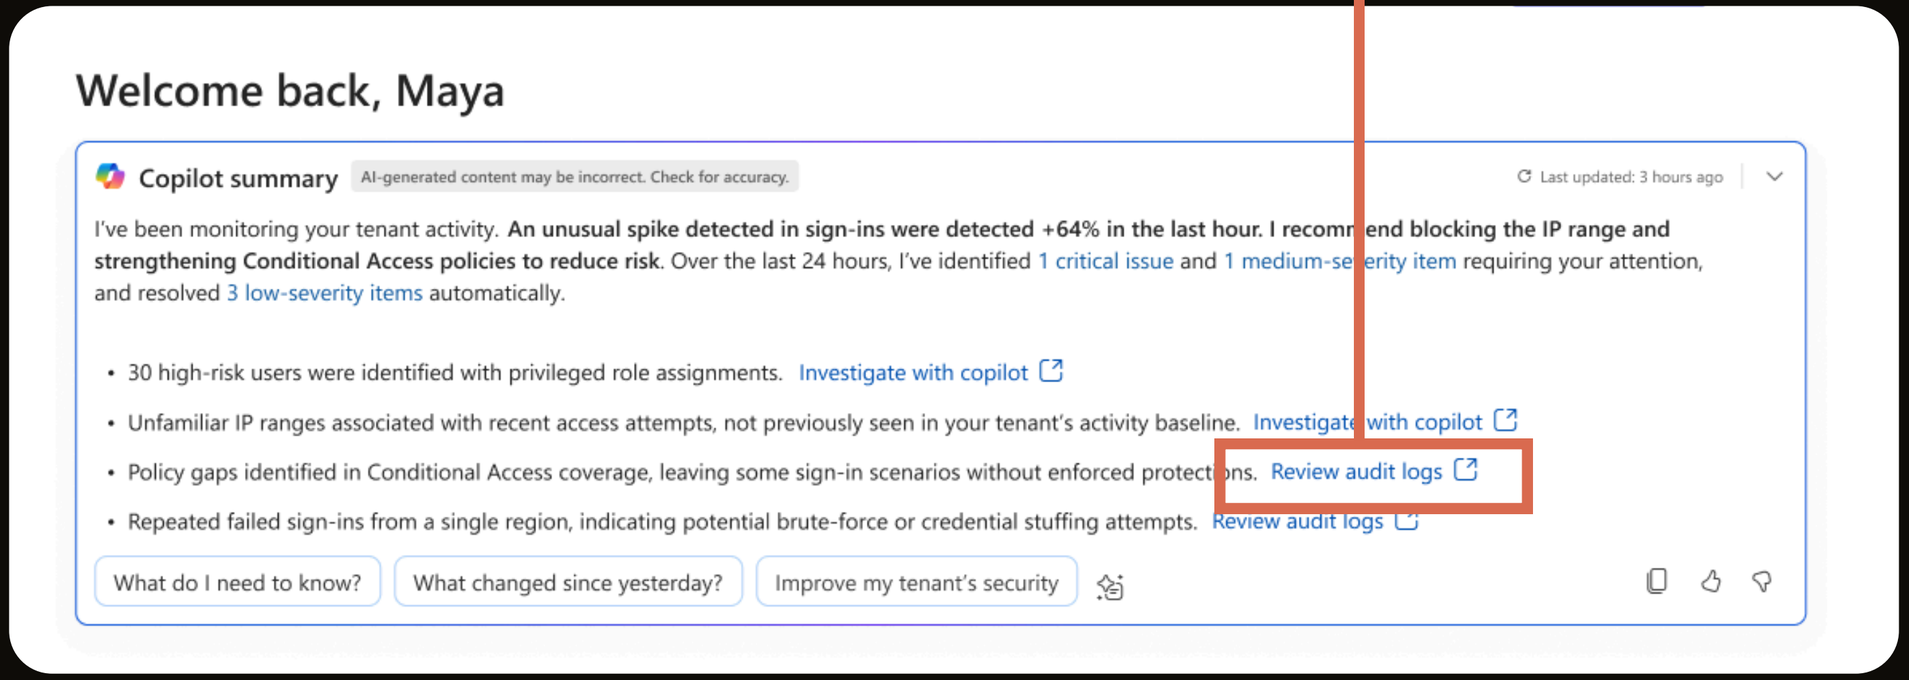

“I want to understand how and why these users are high risk”

Maya Patel

Reliability & safety

Confidence signals

MAYA WANTS TO UNDERSTAND WHY THERE ARE HIGH RISK USERS IN HER TENANT. LETS DIG IN.

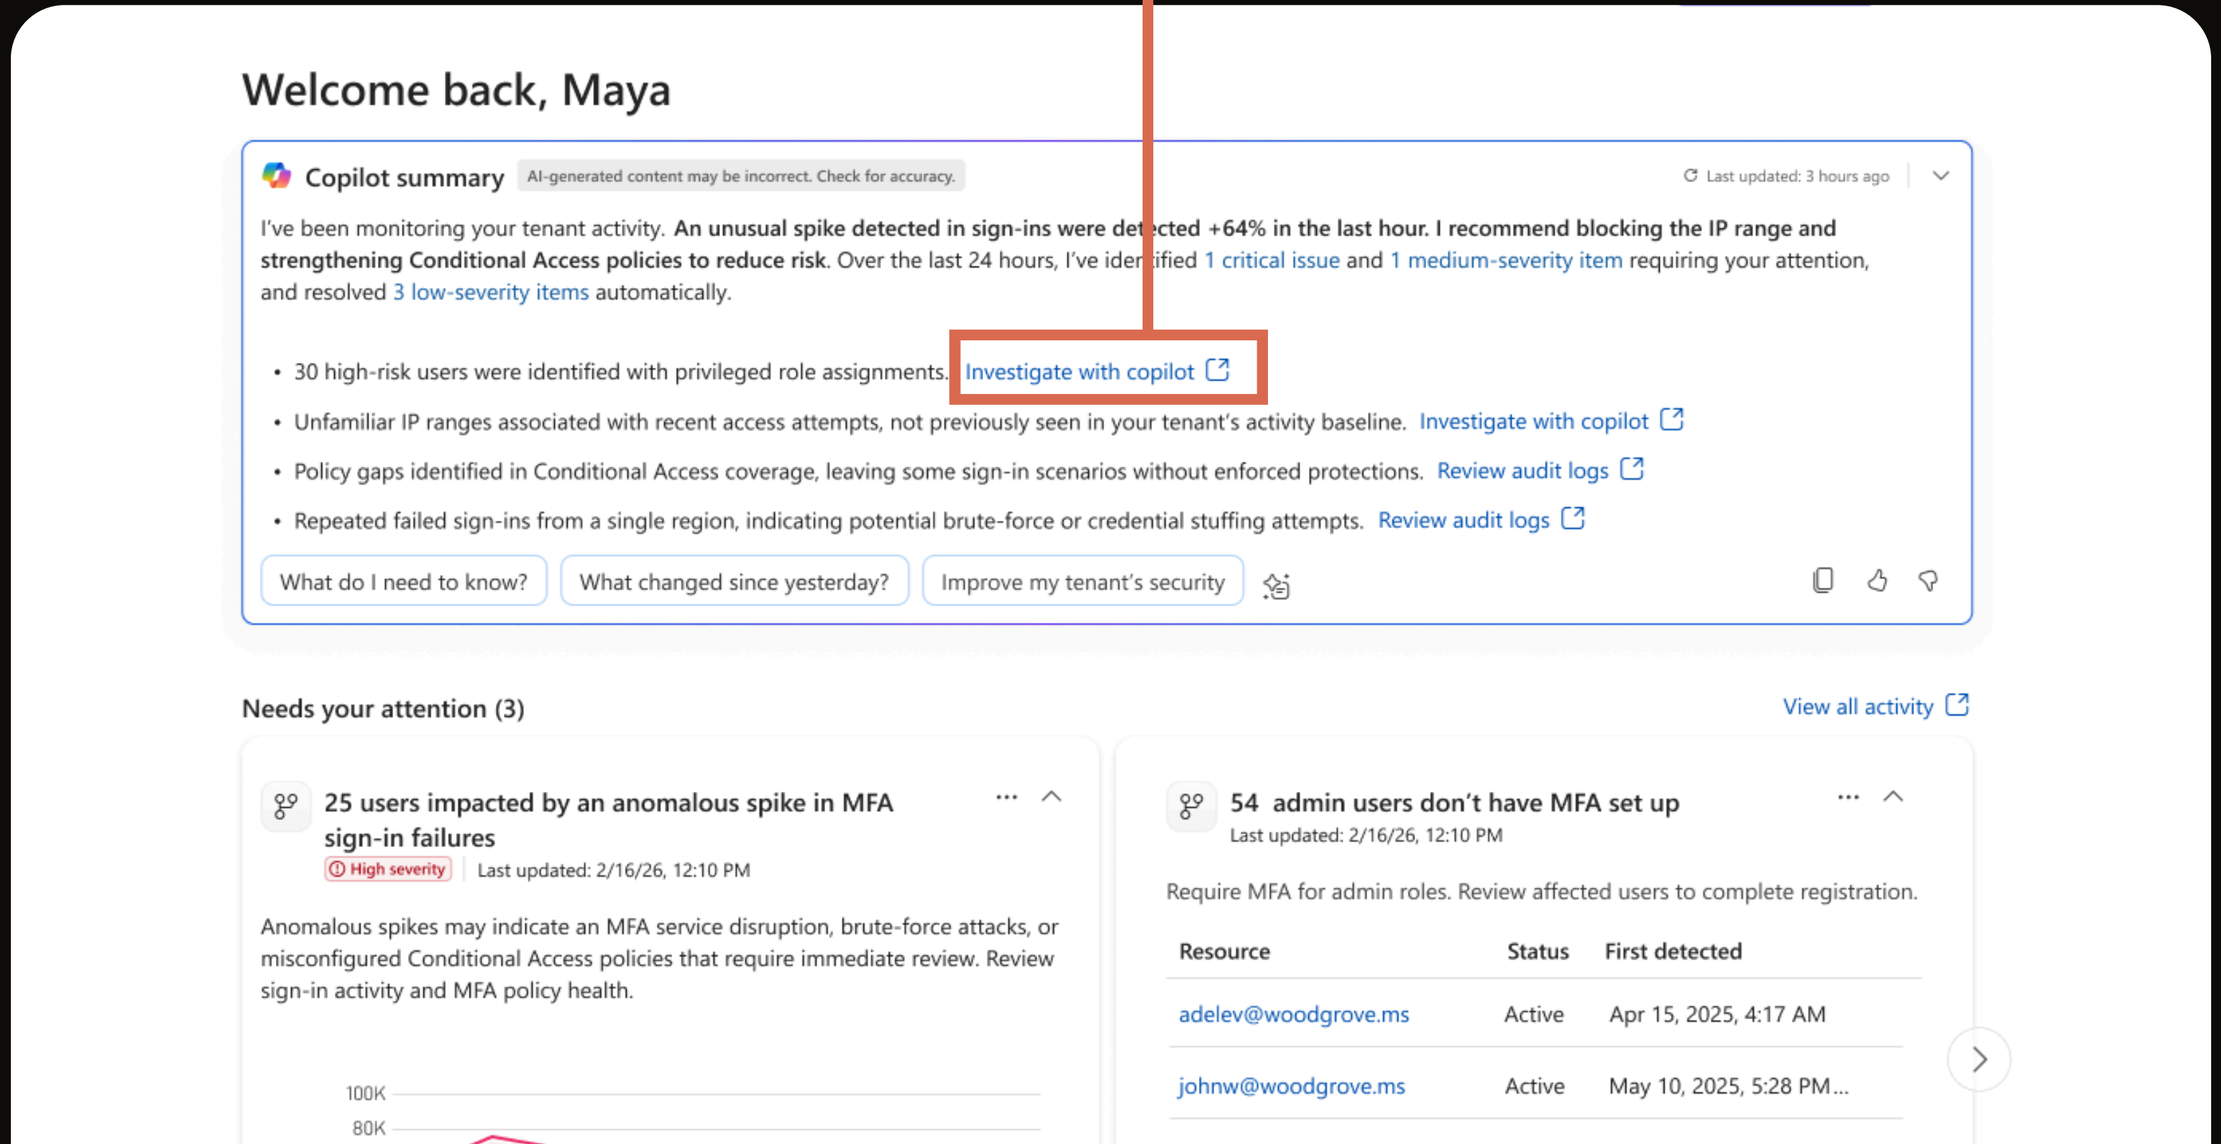

Transparency

Confidence signals

SHE CLICKS INVESTIGATE WITH COPILOT TO START AN INVESTIGATION

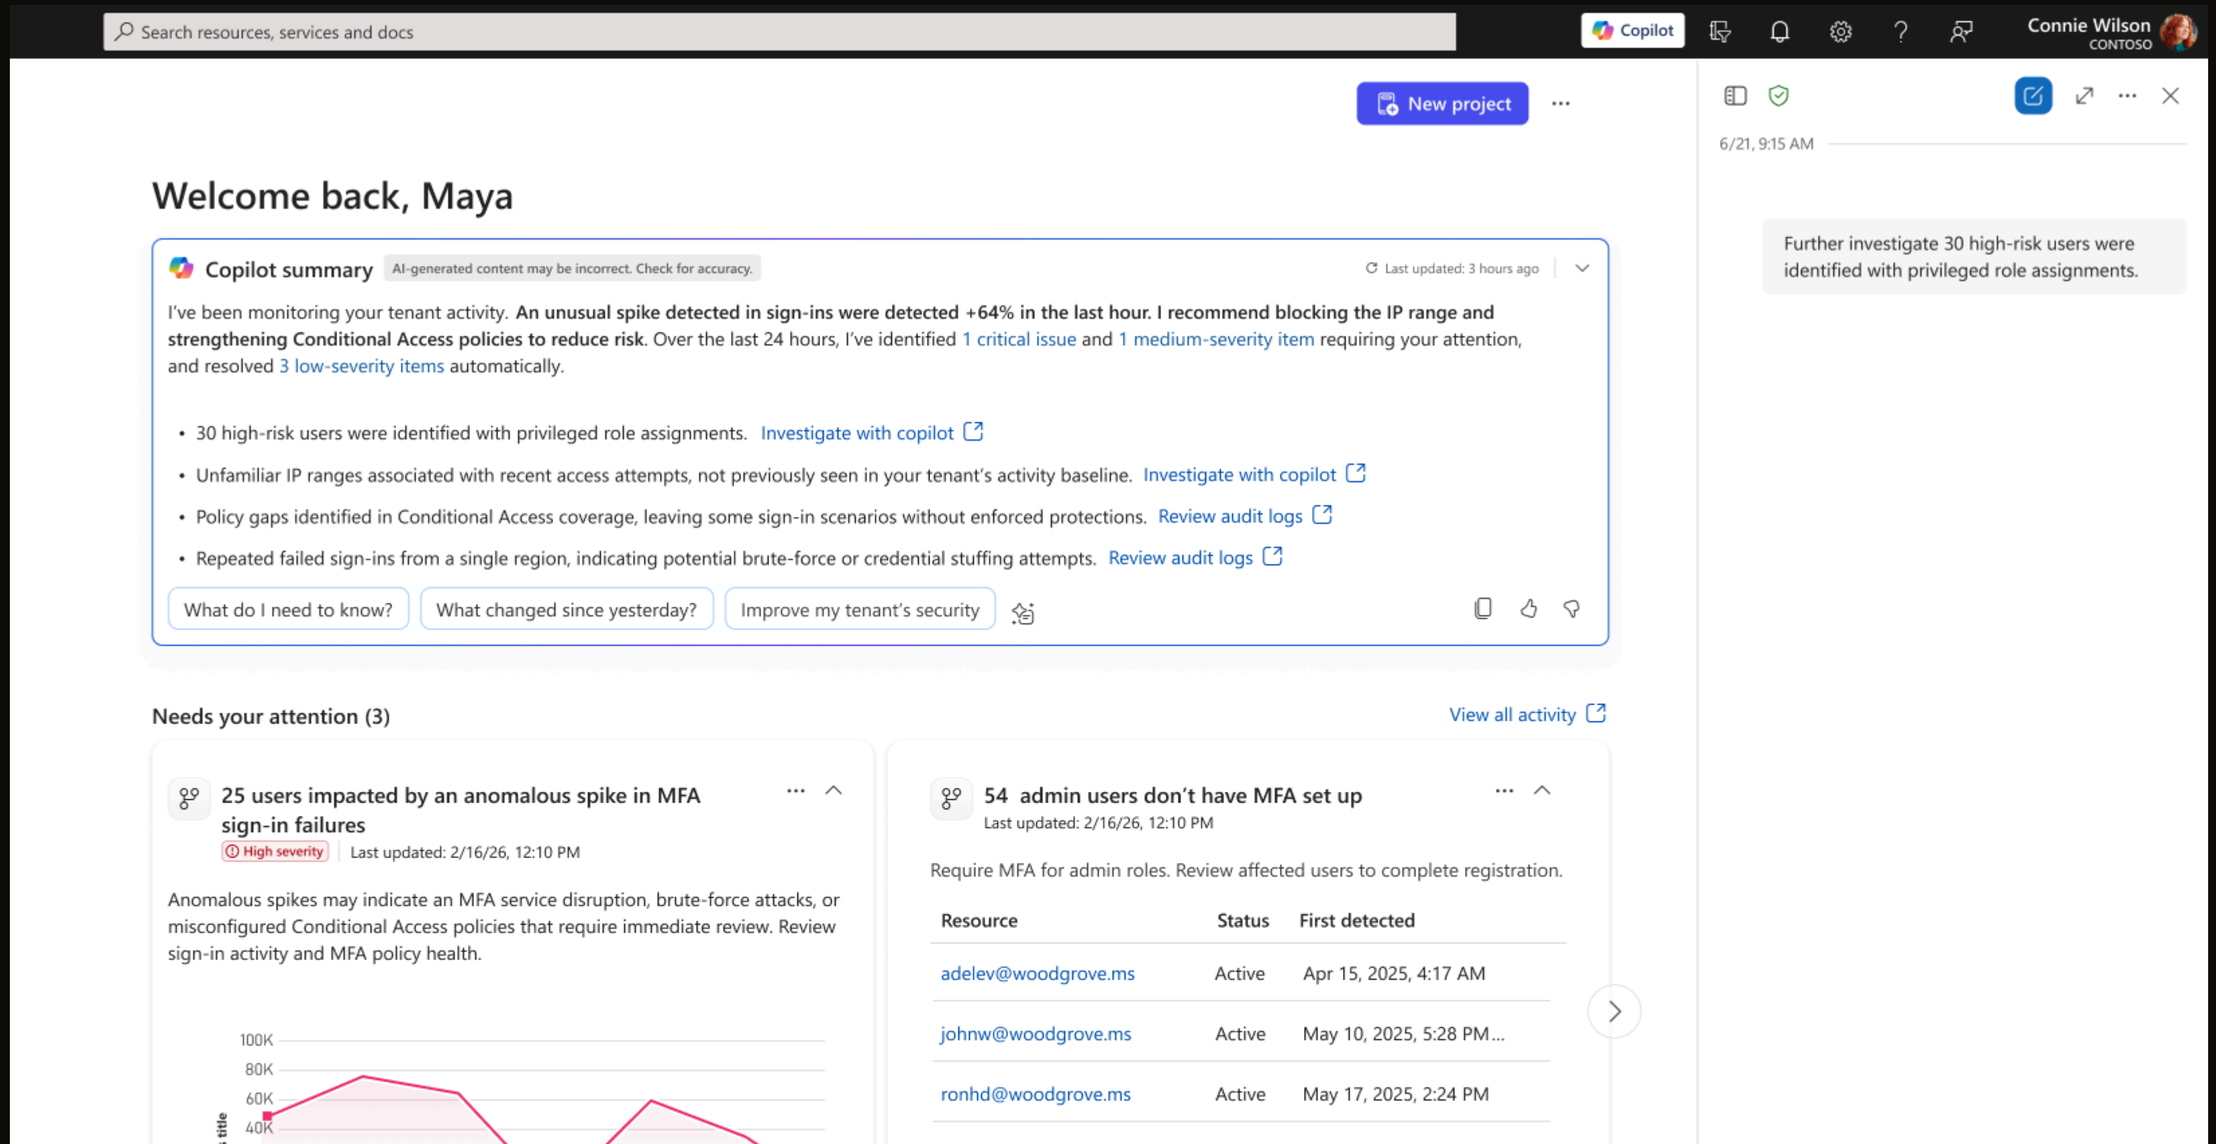

Transparency

Guided investigation

USERS WANT CONFIDENCE THAT THE SYSTEM FOLLOWED A REASONABLE PROCESS.TRANSPARENCY DOES NOT REQUIRE EXPOSING EVERY DETAIL ALL THE TIME.

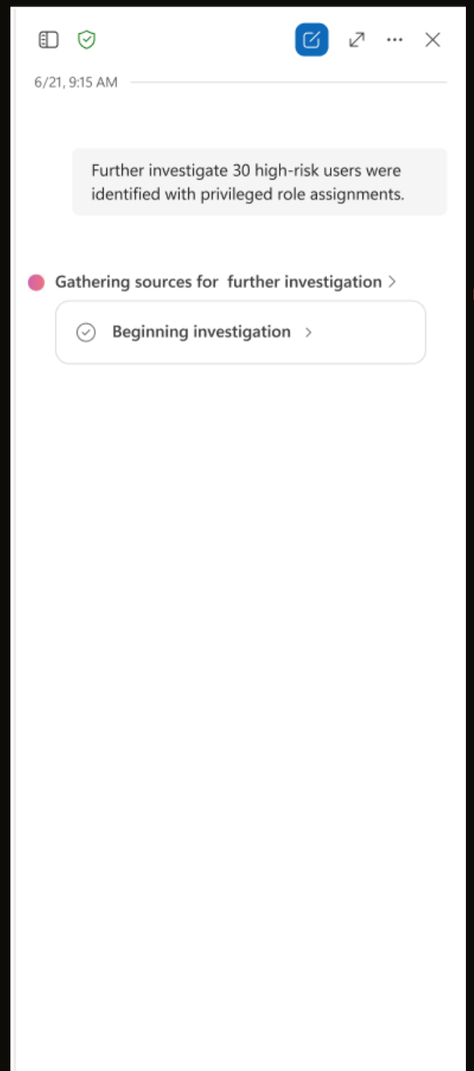

INVESTIGATE

Transparency

Show your work

BEFORE ASKING THE USER TO ACT, THE SYSTEM SUMMARIZES THE FINDINGS THAT LED TO ITS RECOMMENDATION.

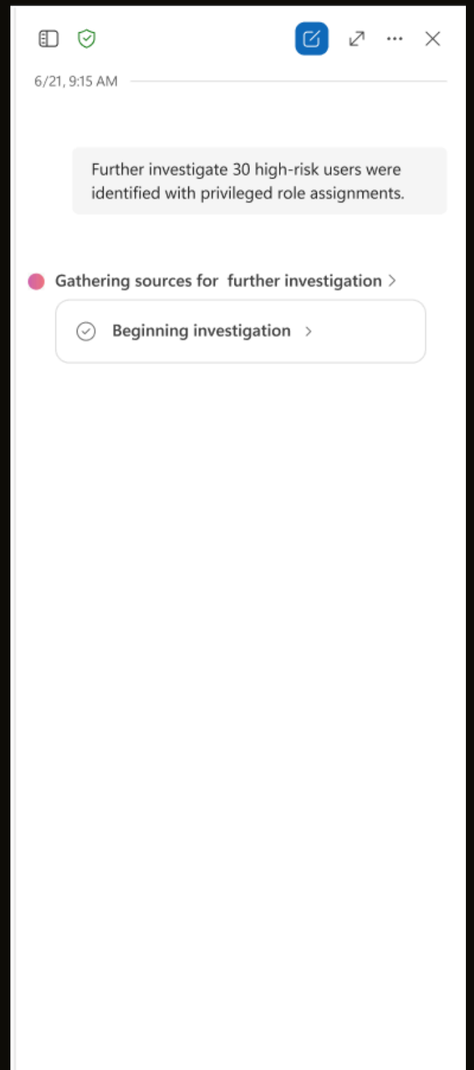

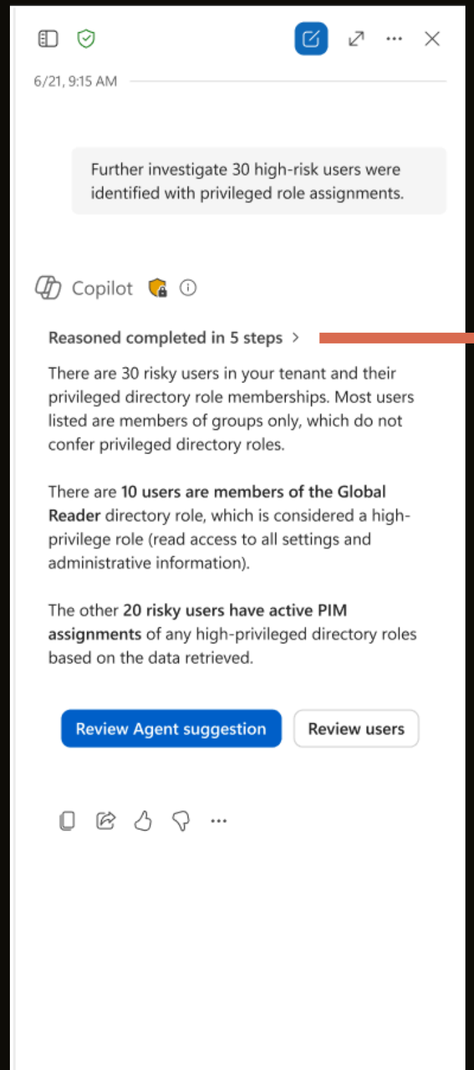

REASONING SUMMARY

Transparency

Show your work

USERS CAN INSPECT HOW INFORMATION WAS GATHERED AND EVALUATED WITHOUT NEEDING TO UNDERSTAND EVERY TECHNICAL DETAIL.

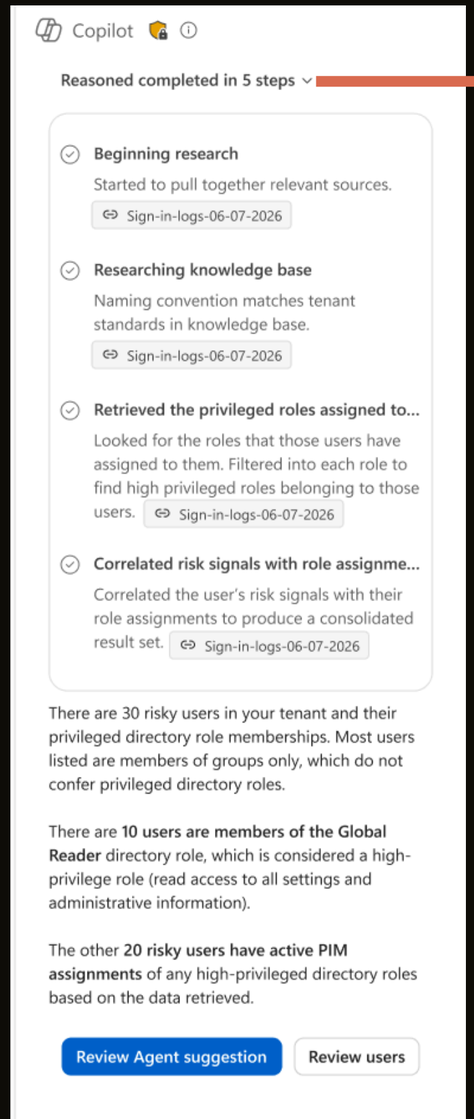

REASONING STEPS

Transparency

Show your sources

EXPANDING THE REASONING ALLOWS USERS TO INSPECT THE STEPS TAKEN, UNDERSTAND HOW INFORMATION WAS GATHERED, AND EVALUATE WHETHER THE APPROACH WAS APPROPRIATE.

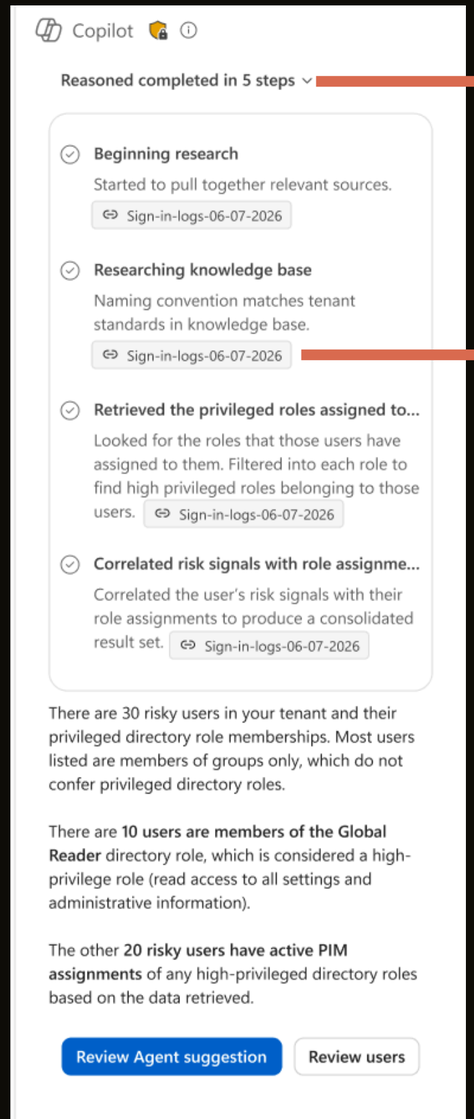

INVESTIGATION STEPS

SOURCE DATA

TRUST INCREASES WHEN USERS CAN SEE BOTH THE PROCESS AND THE INFORMATION USED TO REACH A CONCLUSION.

Transparency

Show options to go deeper

EXPANDING THE REASONING ALLOWS USERS TO INSPECT THE STEPS TAKEN, UNDERSTAND HOW INFORMATION WAS GATHERED, AND EVALUATE WHETHER THE APPROACH WAS APPROPRIATE.

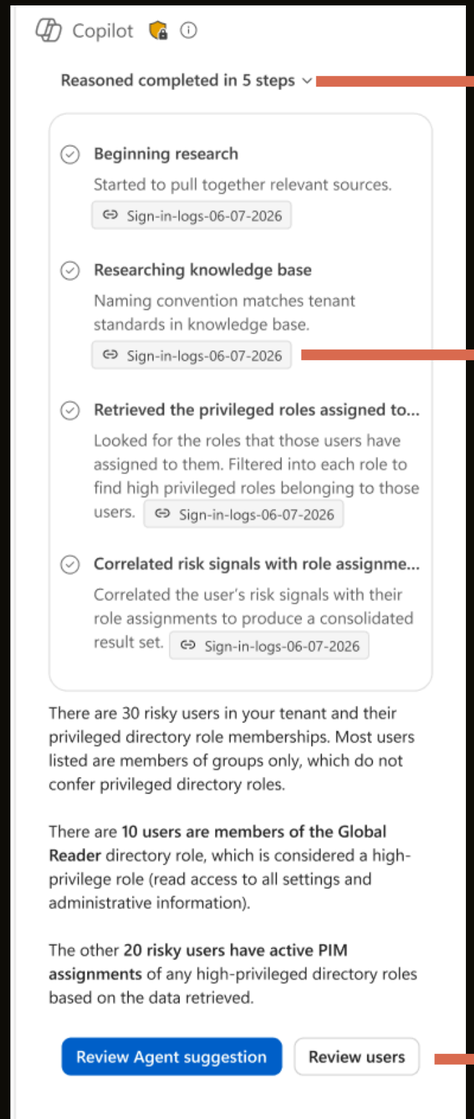

INVESTIGATION STEPS

SOURCE DATA

TAKE ACTION

Transparency

Explore further

DIFFERENT USERS NEED DIFFERENT LEVELS OF DETAIL. TRANSPARENCY SHOULD SUPPORT BOTH QUICK DECISIONS AND DEEPER INVESTIGATION.

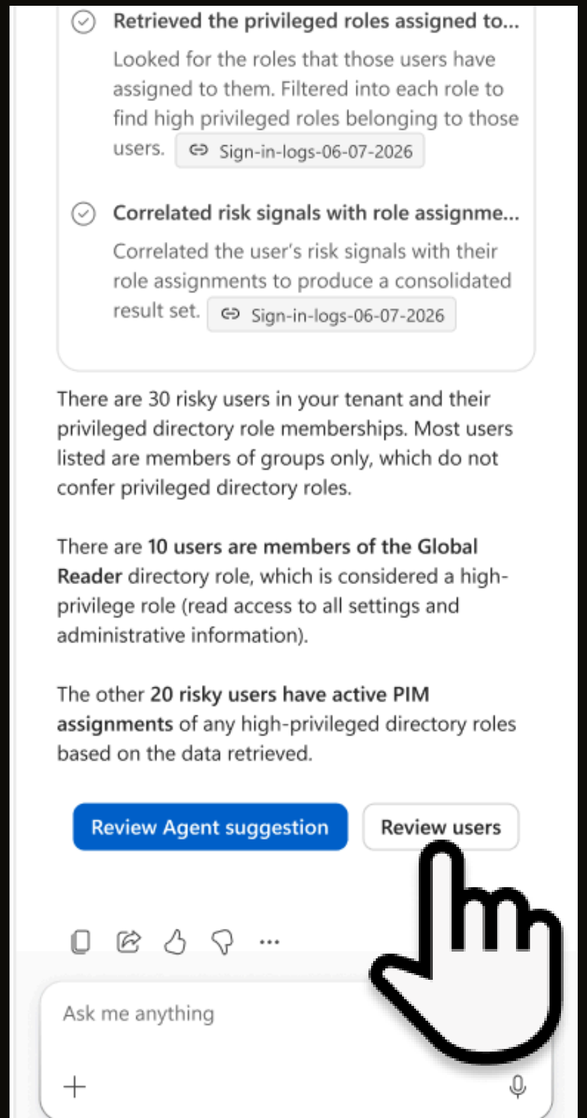

Accountability

Review the evidence

THE STRONGEST FORMS OF TRANSPARENCY ALLOW USERS TO MOVE FROM:
ANSWER → REASONING → EVIDENCE.

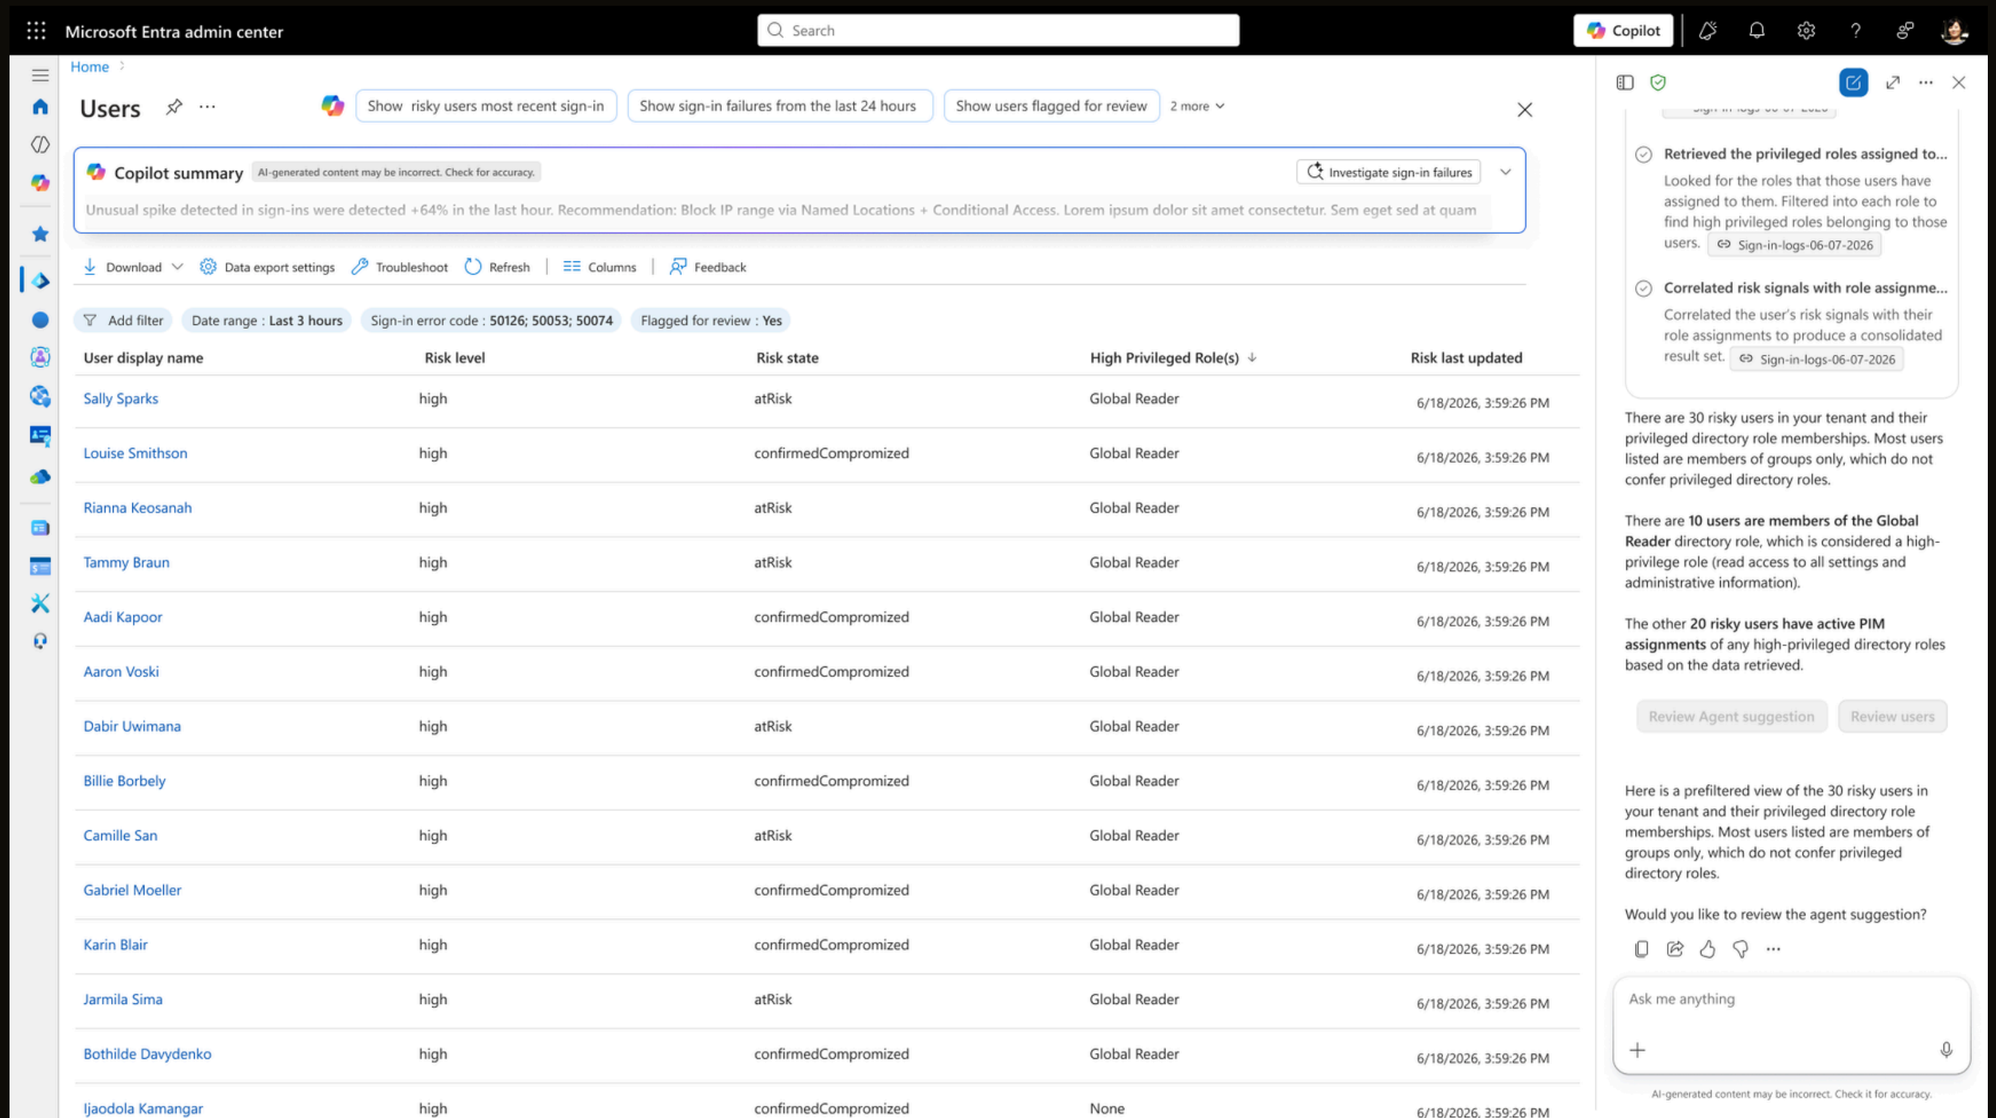

Accountability

Human review required

RECOMMENDATIONS ARE PROPOSED BY THE SYSTEM, BUT RESPONSIBILITY FOR THE FINAL DECISION REMAINS WITH THE USER.

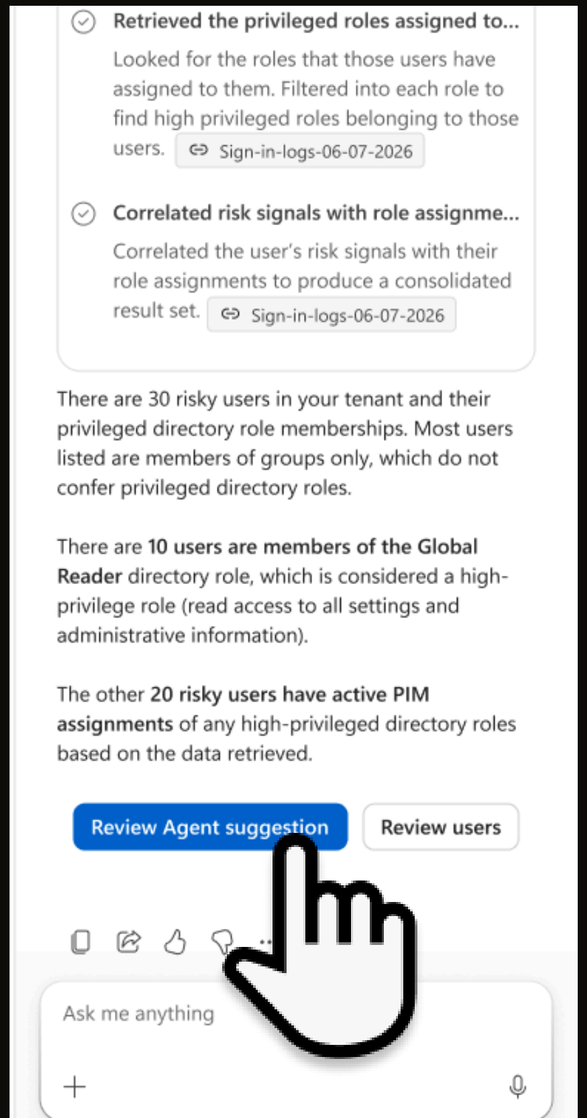

Accountability

Agent suggests, human acts

THE AGENT PERFORMS ANALYSIS AND PREPARES RECOMMENDATIONS. MAYA REVIEWS THE EVIDENCE, EVALUATES THE IMPACT, AND DECIDES WHETHER TO TAKE ACTION.

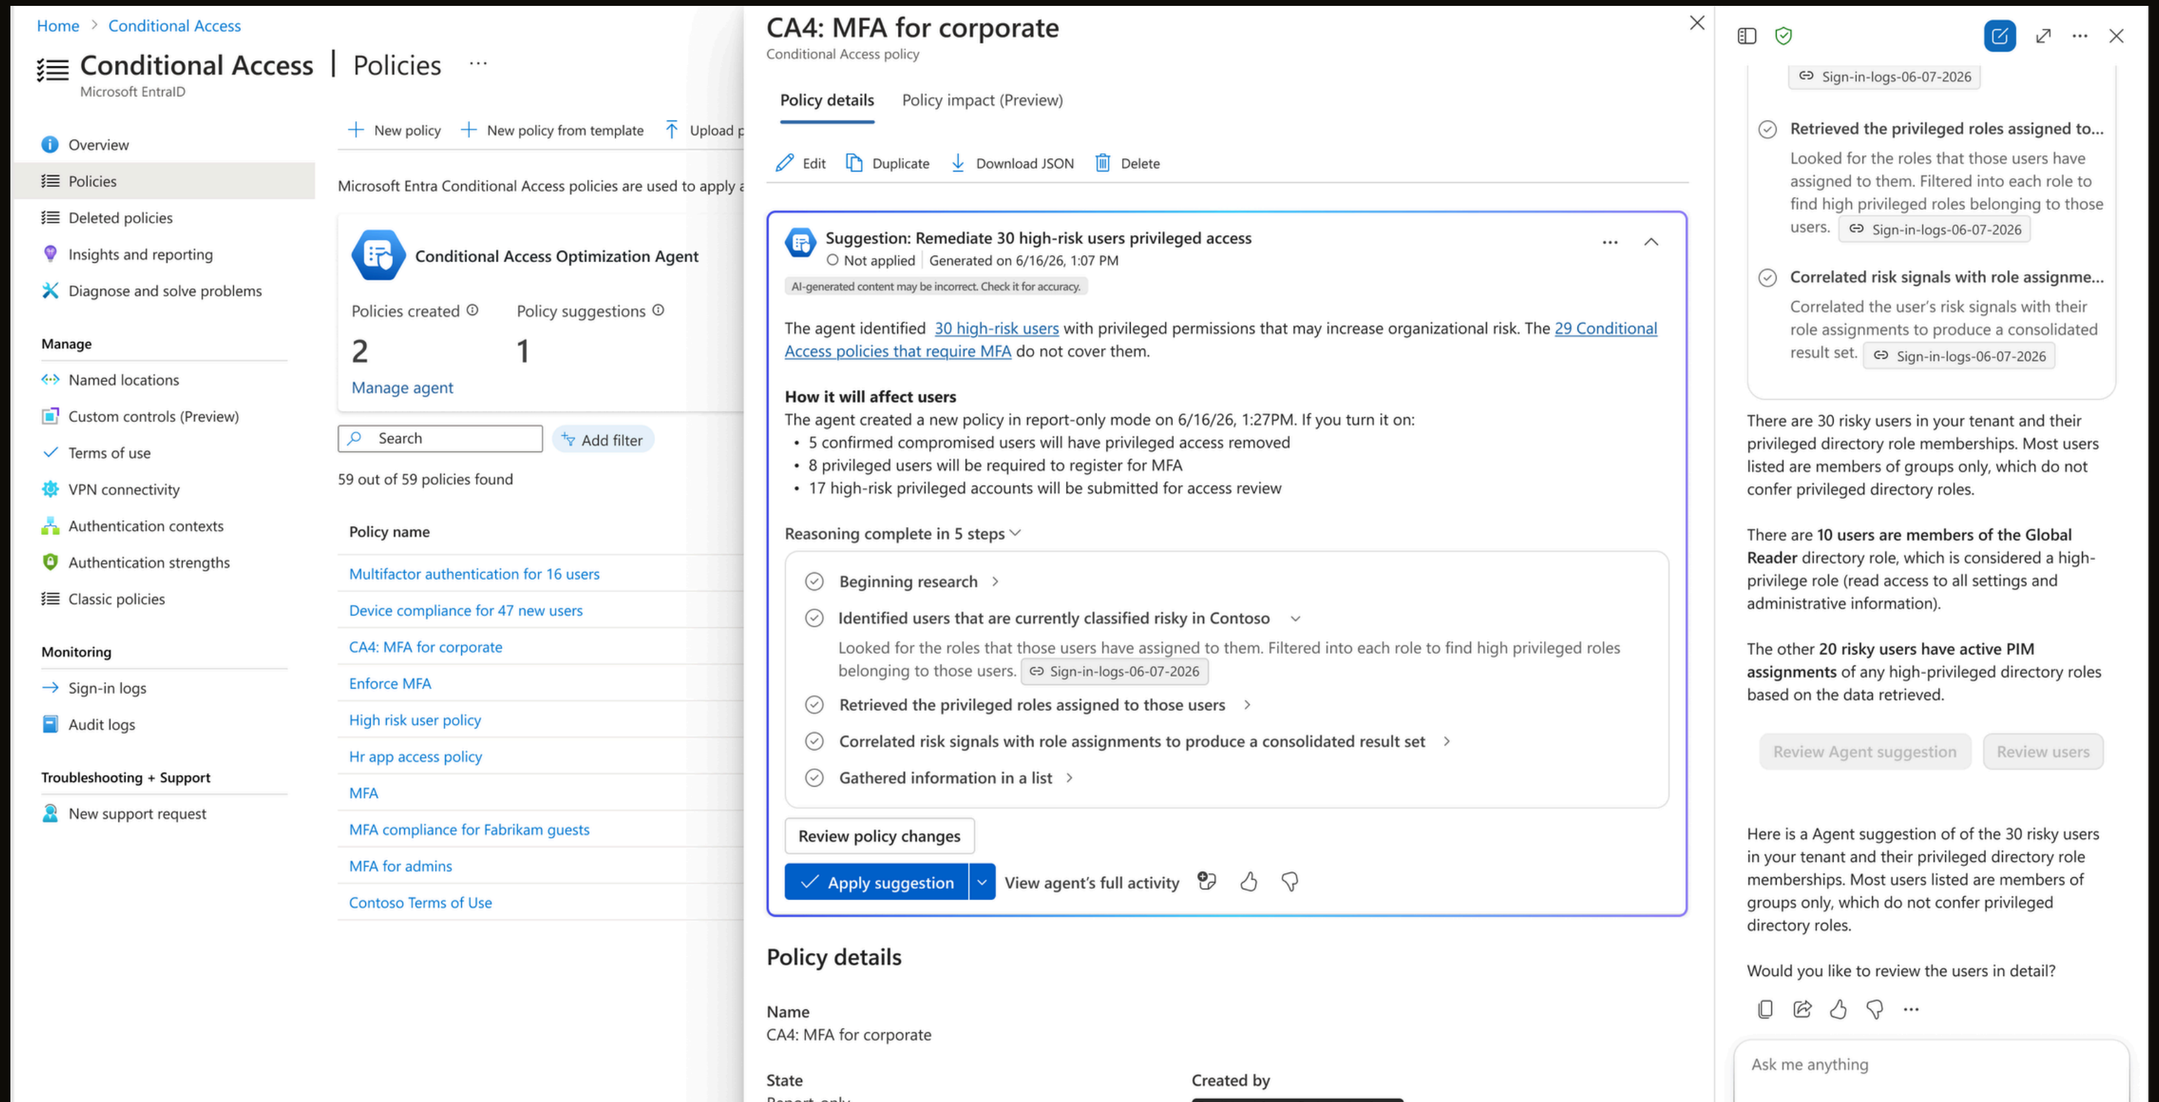

Accountability

ACCOUNTABILITY ENSURES PEOPLE REMAIN RESPONSIBLE FOR OUTCOMES, EVEN WHEN AI ASSISTS WITH ANALYSIS, RECOMMENDATIONS, OR AUTOMATION.

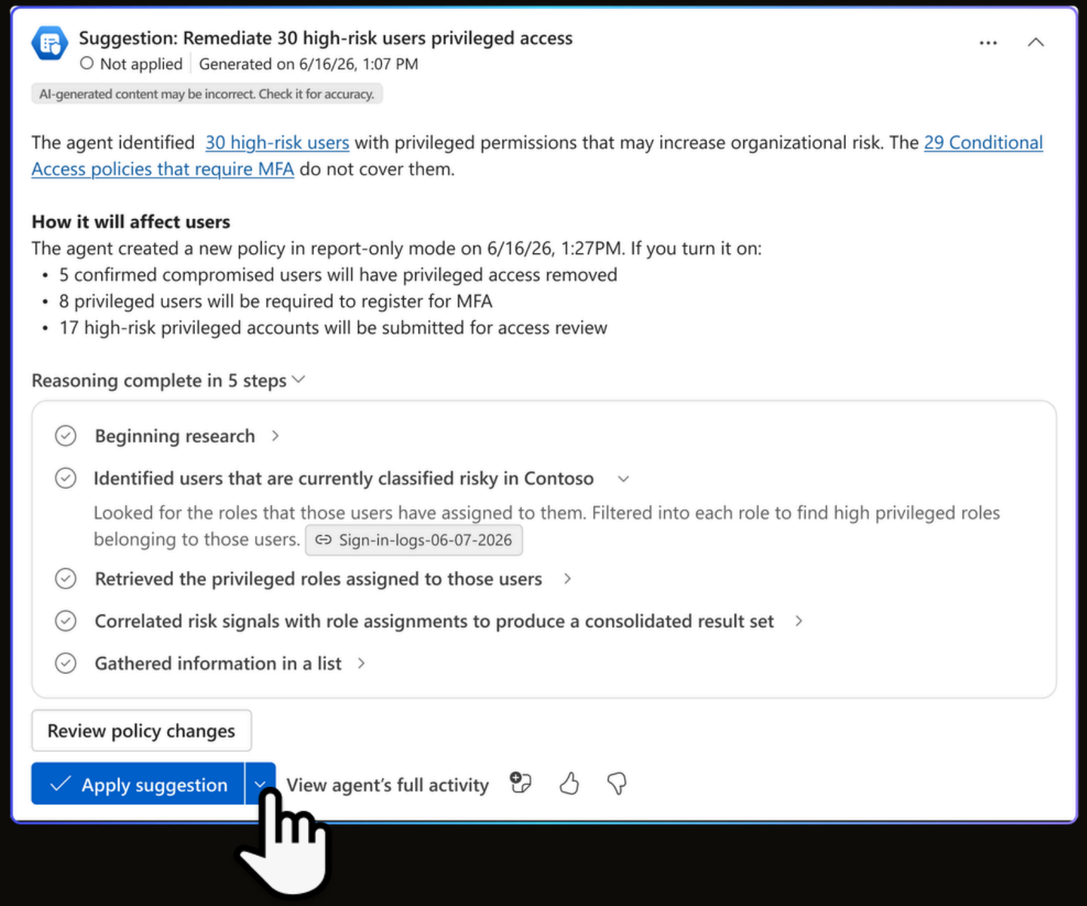

Accountability

Human in the loop

HUMAN-IN-THE-LOOP WORKFLOWS CREATE DELIBERATE MOMENTS FOR REVIEW, APPROVAL, AND INTERVENTION BEFORE CONSEQUENTIAL ACTIONS ARE TAKEN.

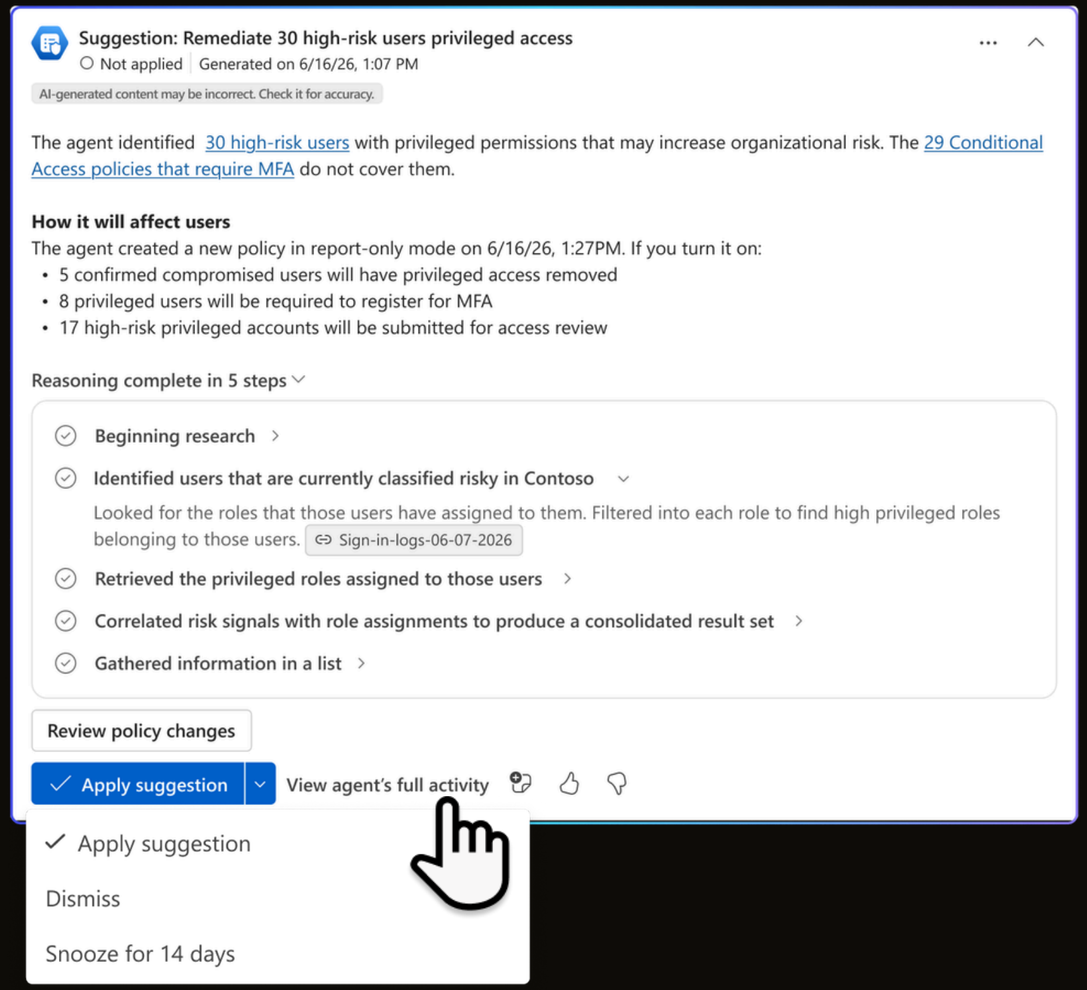

Accountability

Audit trail

AUDIT TRAILS HELP USERS UNDERSTAND WHAT HAPPENED, WHY IT HAPPENED, AND HOW THE OUTCOME WAS REACHED.

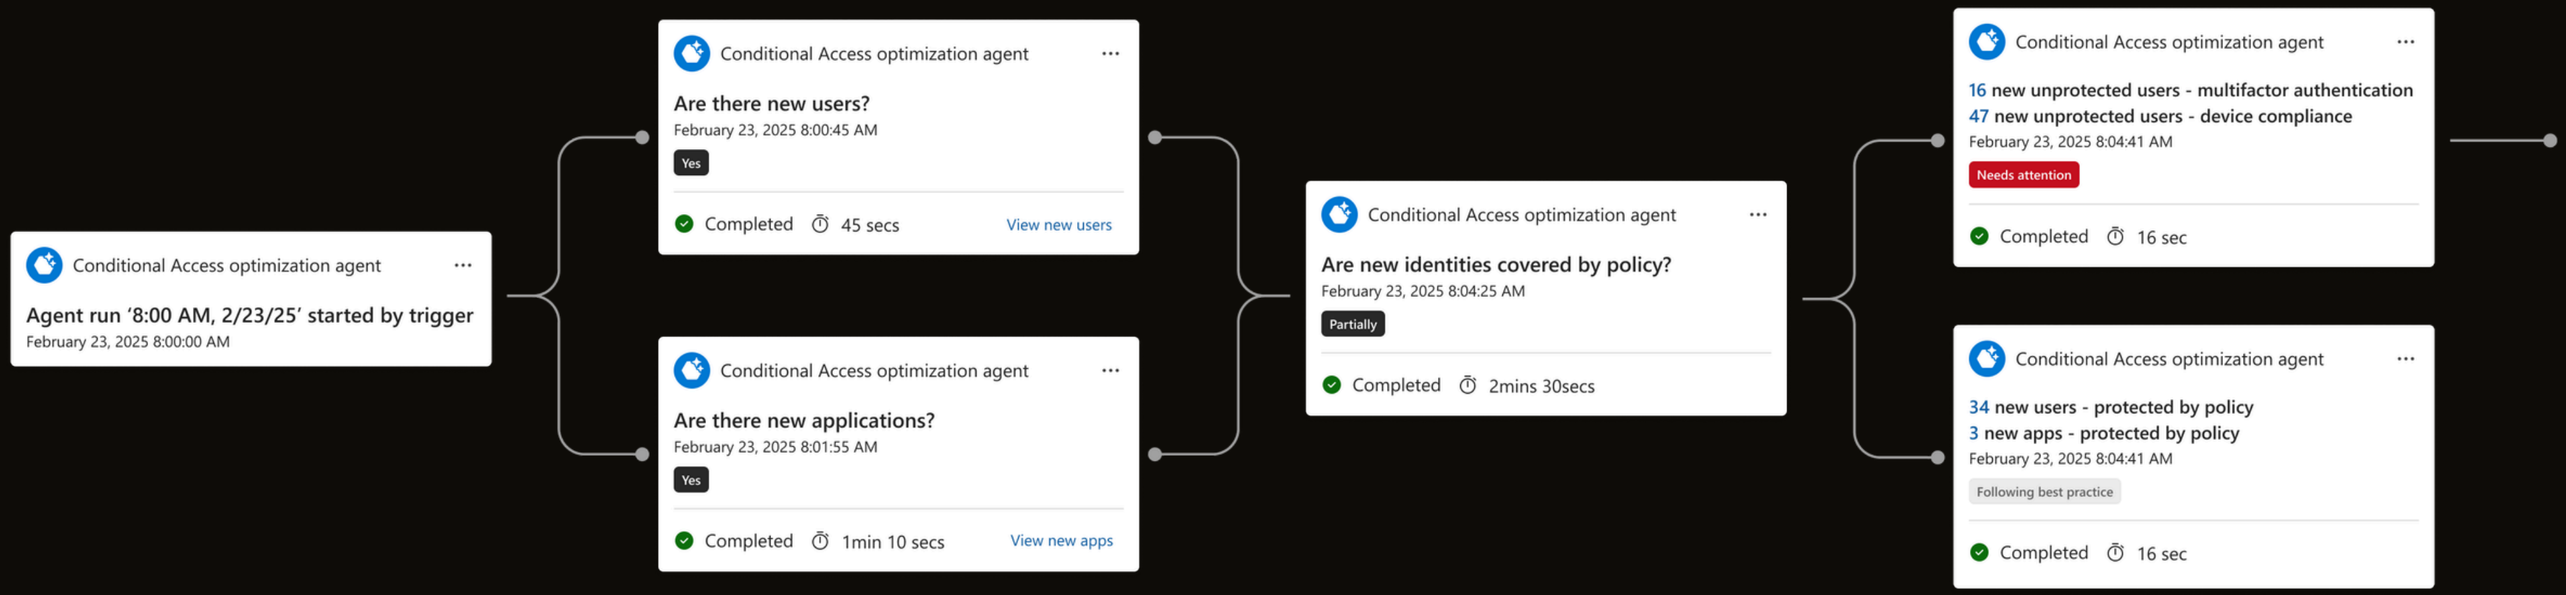

Accountability + Reliability & safety

Accountability at scale

ORGANIZATIONS CANNOT MANAGE WHAT THEY CANNOT SEE. ACCOUNTABILITY BEGINS WITH VISIBILITY.

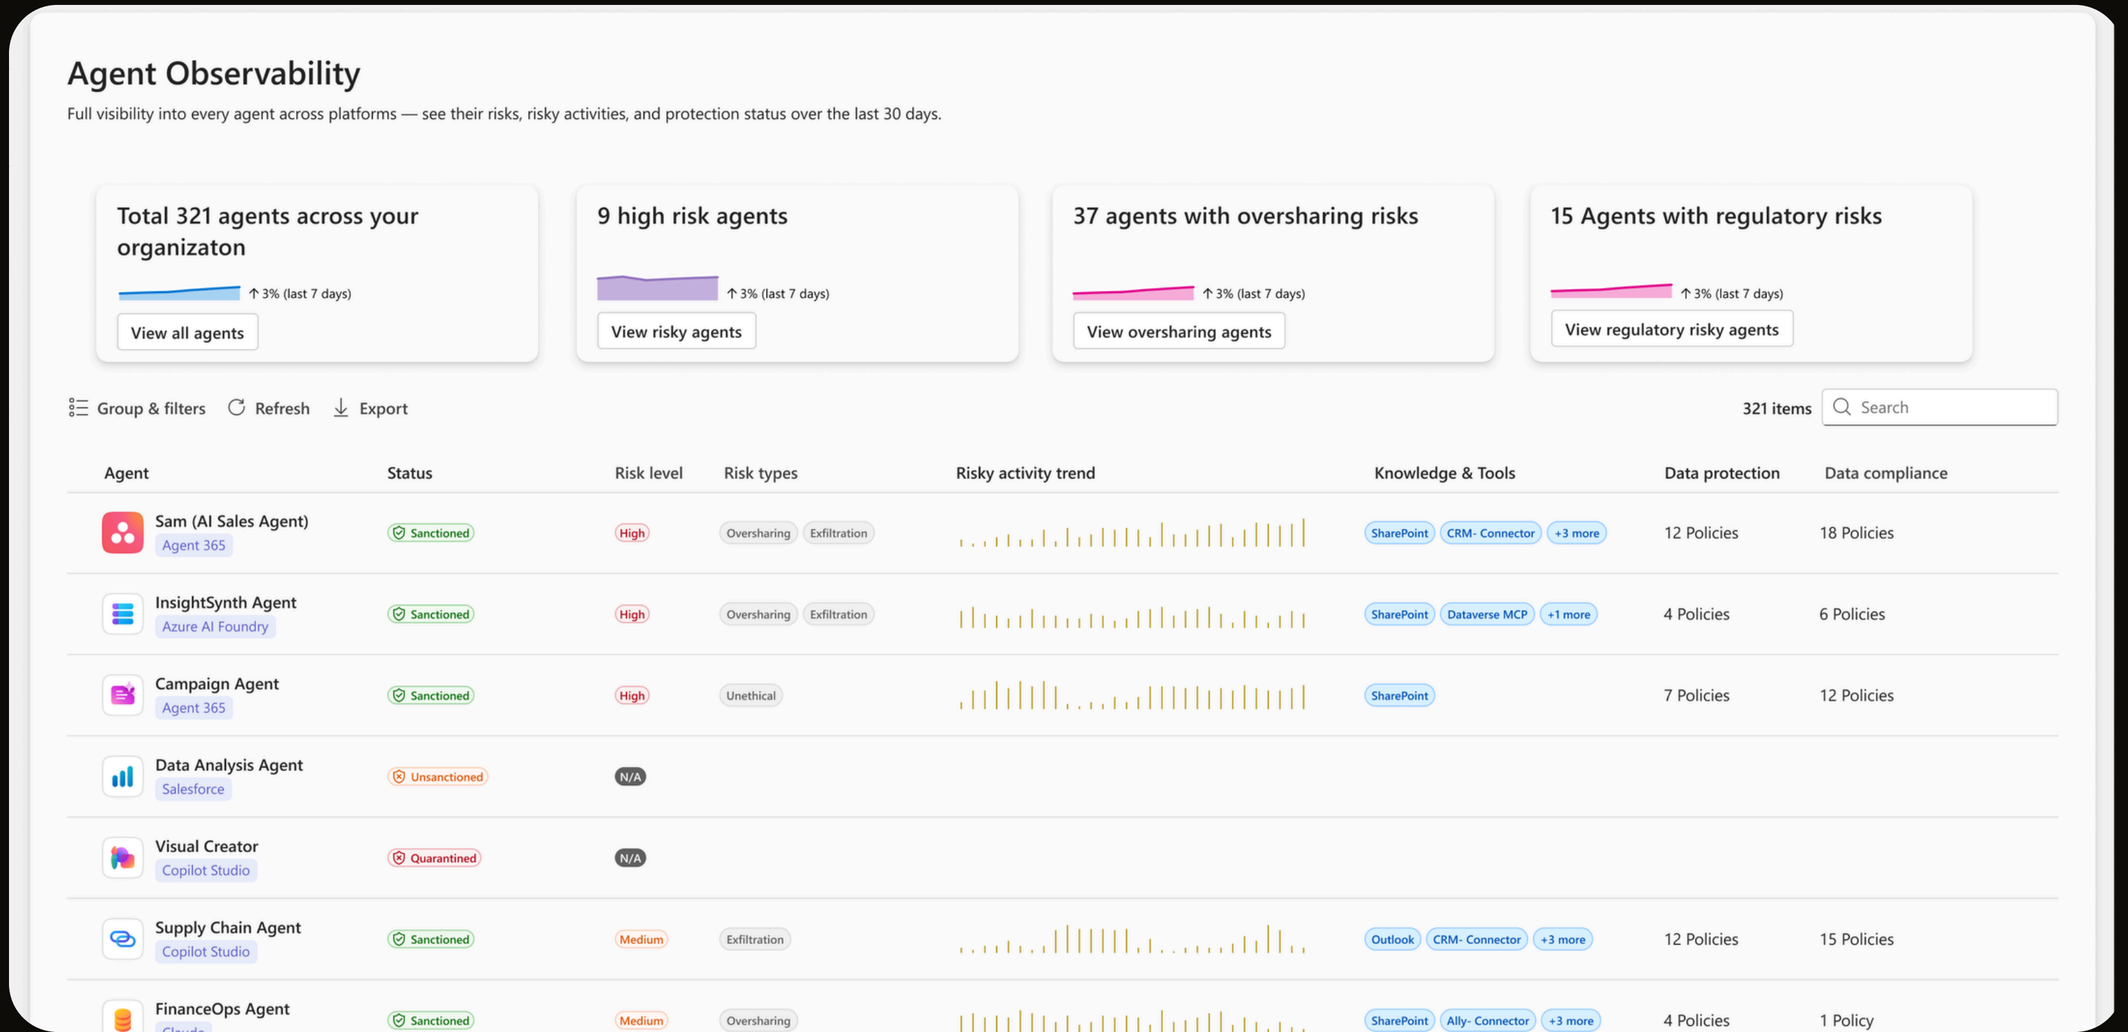

Additional patterns, outside of this flow

Design for choice

Accountability

AI CAN NARROW THE POSSIBILITIES, BUT USERS SHOULD ALWAYS BE ABLE TO STEER THE FINAL DECISION.CUSTOM INPUT PRESERVES AGENCY.

(MICROSOFT PATTERN IN VSCODE)

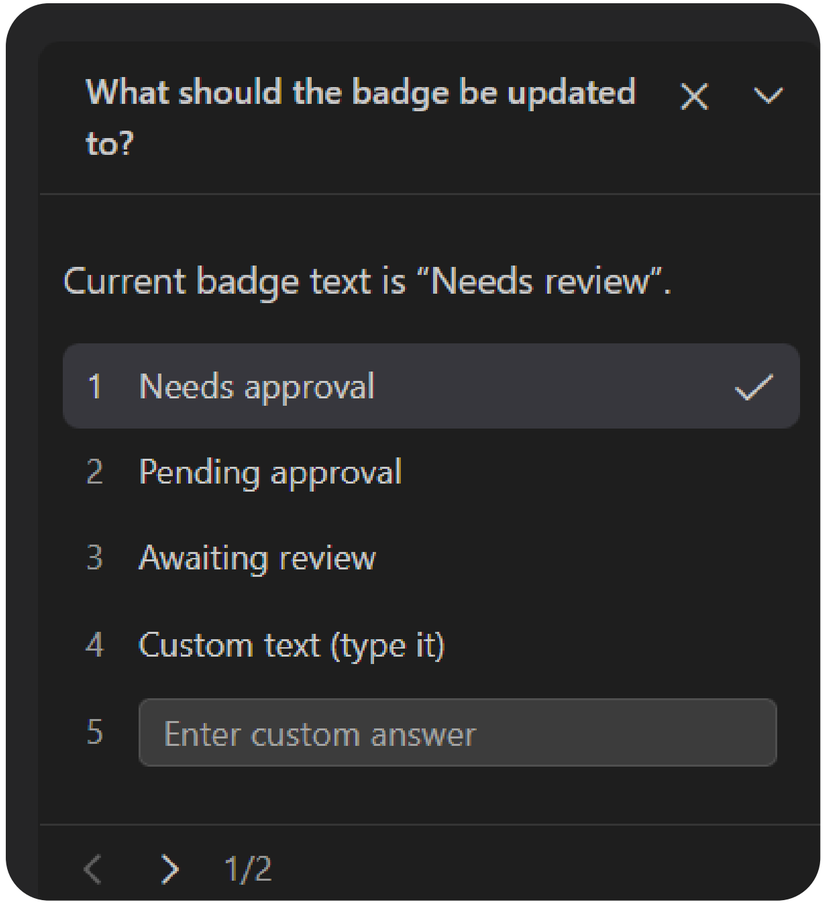

Reliability & safety

Steering with context

TRUST IS EARNED GRADUALLY. AS CONFIDENCE GROWS, USERS OFTEN CHOOSE TO DELEGATE MORE RESPONSIBILITY TO AI SYSTEMS, THAT REQUIRES MORE CONTEXT.

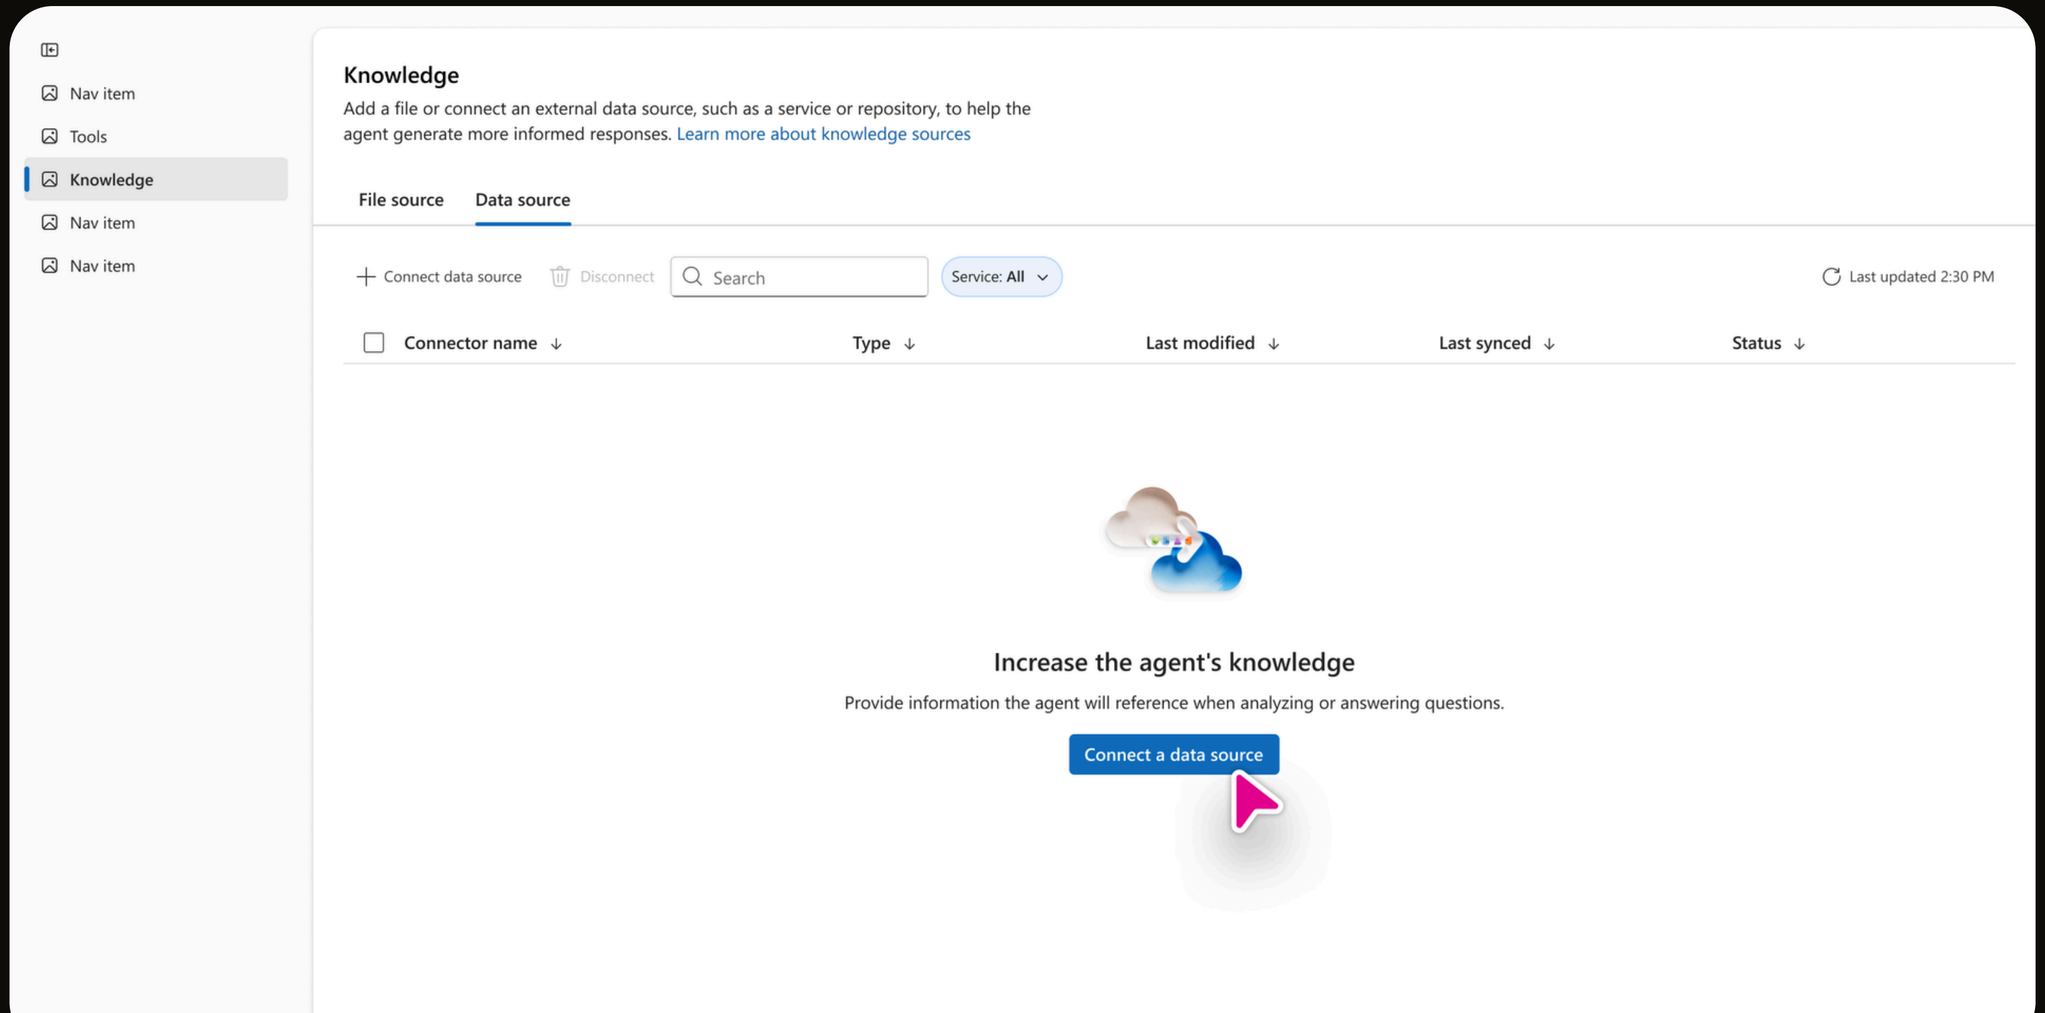

Transparency

Contextually relevant disclaimers

EXPECTATION-SETTING DOESN'T HAPPEN ONCE.
IT APPEARS THROUGHOUT THE EXPERIENCE.

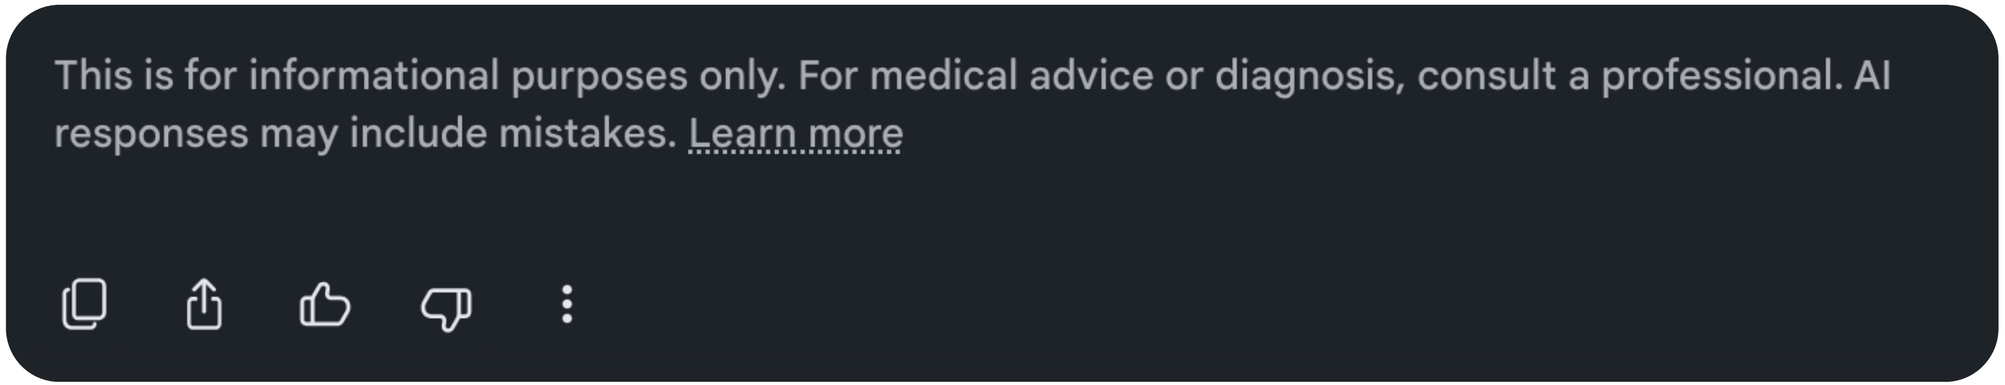

Responsible AI principles

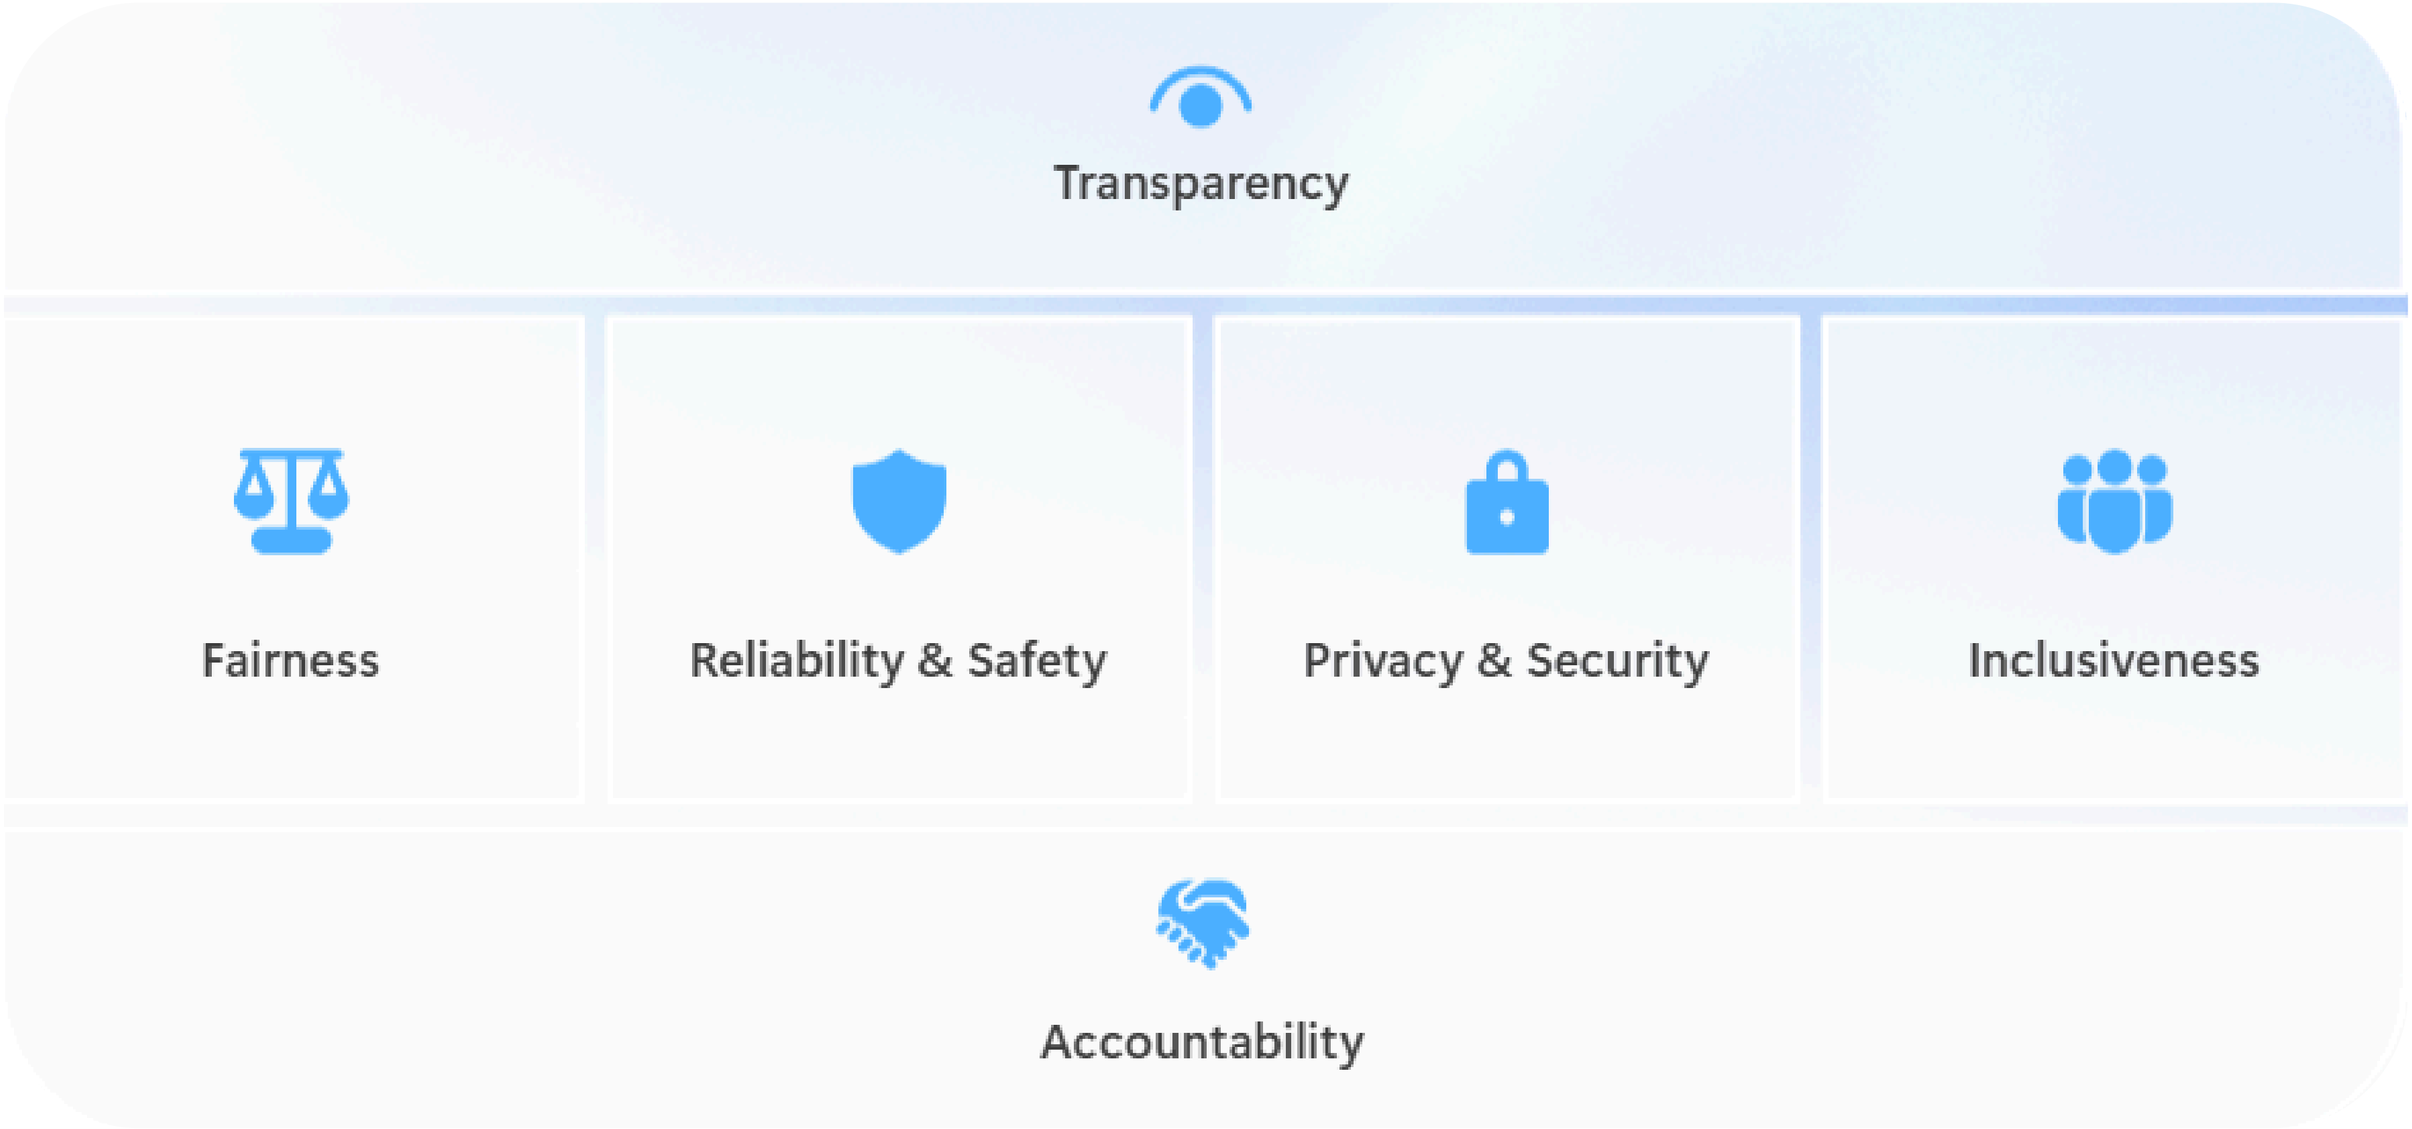

Trust must be *continuously* earned.

As AI becomes more capable, trust isn't a milestone. It's an ongoing design responsibility.

Transparency

Keep users informed.
Ground AI in trusted data, monitor performance, and adapt as context changes.

Reliability & Safety

Keep systems dependable.
Ground AI in current context, measure outcomes, and reduce mistakes over time.

Accountability

Keep humans in control.
Give users meaningful oversight, review points, and the ability to intervene.

"

The most profound technologies are those that *disappear*. They weave themselves into the fabric of everyday life until they are indistinguishable from it.

— MARK WEISER · FATHER OF UBIQUITOUS COMPUTING

What we trust, we trust *quietly*.

Questions?

Questions or ideas? Let’s connect!

EMAIL
mary.gansallo@gmail.com

LINKEDIN
in/mary-gansallo