



Anzenna Brings AI DLP to Microsoft Purview, Cutting Alert Volume and Closing the AI Gap

Anzenna auto-investigates Microsoft Purview alerts, cutting volume by 90 percent, and covers the browser AI, IDE, MCP, and OAuth channels Purview cannot see.

REDWOOD CITY, CA, UNITED STATES, August 12, 2026 /EINPresswire.com/ -- Anzenna today announced [AI DLP](#) & Usage Control for Microsoft Purview environments, a way for Microsoft shops to keep Purview for content inspection and compliance and add the AI visibility and investigation it lacks. Anzenna connects to Microsoft 365, ingests every Purview alert, and works each one automatically, cutting alert volume by 90 percent. It also covers the channels data now leaves through that Purview cannot inspect: a prompt pasted into a browser AI tool, a coding agent reaching a credential in the IDE, files moved through an MCP server, and OAuth grants to AI apps.

A busy Purview tenant produces thousands of DLP and audit alerts a day, with no built-in way to rank the one that matters against the hundreds that do not. Security teams work each alert by hand across the Purview portal, Entra, and Defender, or they pay an outside firm to keep the policies tuned. At the same time, the fastest-growing path for data to leave the enterprise, employee AI use, sits outside Purview's content inspection entirely.

Anzenna closes both gaps at once. When a Purview alert fires, an investigative agent gathers the surrounding context, the identity, the data, the destination, and the person's HR and departure status, and assembles a finished case. The benign majority closes on its own, and only the few that need a human are escalated, already worked up with the evidence attached. Alongside that, Anzenna discovers every AI tool, agent, and MCP server across the fleet, audits what each one does at the prompt level, and classifies and blocks sensitive data at the point it moves in the browser and the IDE. Every action ties back to the identity behind it, human or non-human, and to the person who owns it.

The same data movement can mean three different things, and Anzenna is built to tell them apart. A support agent using a sanctioned assistant inside normal working patterns is allowed through. Someone reaching for a personal account to move faster is coached in the moment. A departing engineer pulling source into a personal AI tool and pushing repositories to a personal account, four days after HR flagged the departure, is escalated to the security team with the whole chain assembled as one case. Purview, on its own, would surface at most a single alert, if a labeled file happened to cross a monitored boundary.

Anzenna preserves the existing investment rather than replacing it. It requires no Purview license and works with any Microsoft 365 tier. Remediation routes through the tools already in place, Microsoft Defender and Intune, Entra ID, and, for multi-vendor stacks, Okta and CrowdStrike, with no second agent to deploy. Setup is agentless and takes about 15 minutes.

"Purview is strong at classifying and protecting data inside Microsoft 365, but it cannot see the prompt, the agent, or who is driving it, and in a busy tenant it buries the team in alerts," said Chinmaya Sharma, Head of Product at Anzenna. "We take those alerts, investigate every one, and hand back only the handful that matter. Then we cover the AI channels the data actually leaves through now, and we tie all of it to the person behind it. Teams keep Purview and finally get the AI visibility and the quiet queue they have been missing."

A security leader at a large healthcare organization running Microsoft 365 described the change this way: "Purview was firing more than a thousand alerts a day and still missing our AI exposure. Anzenna cut the volume by 90 percent and, for the first time, showed us what our people were actually doing with AI."

AI DLP & Usage Control for Microsoft Purview environments is [available now](#). Read the full Microsoft Purview solution brief [here](#), or see it run against a live environment at <https://www.anzenna.ai>.

About Anzenna

Anzenna is an activity governance & DLP solution for AI and human identities, the first to unify AI Usage Control and Insider Risk Management in a single product. It monitors and governs what every AI agent, non-human identity, and employee is actually doing across the full enterprise surface: browser extensions, MCP servers, IDE and coding agents, OAuth-authorized AI apps, and file and data movement. Its Identity-Data-Action graph links AI activity back to the human behind it, fusing behavioral signals with HR context, data sensitivity, and departure status to catch threats that access control alone cannot see. Anzenna remediates in real time through existing endpoint and identity infrastructure like CrowdStrike, Microsoft Defender, Okta, and others with no second agent to deploy, and generates autonomous, evidence-backed investigation reports instead of alerts in under two minutes. For more information visit <https://www.anzenna.ai>

Marketing Team

Anzenna Inc.

[email us here](#)

This press release can be viewed online at: <https://www.einpresswire.com/article/933157944>

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable

in today's world. Please see our Editorial Guidelines for more information.
© 1995-2026 Newsmatics Inc. All Right Reserved.