anzenna

SOLUTION BRIEF

Microsoft Purview × Anzenna

The AI usage layer on top of *your DLP*

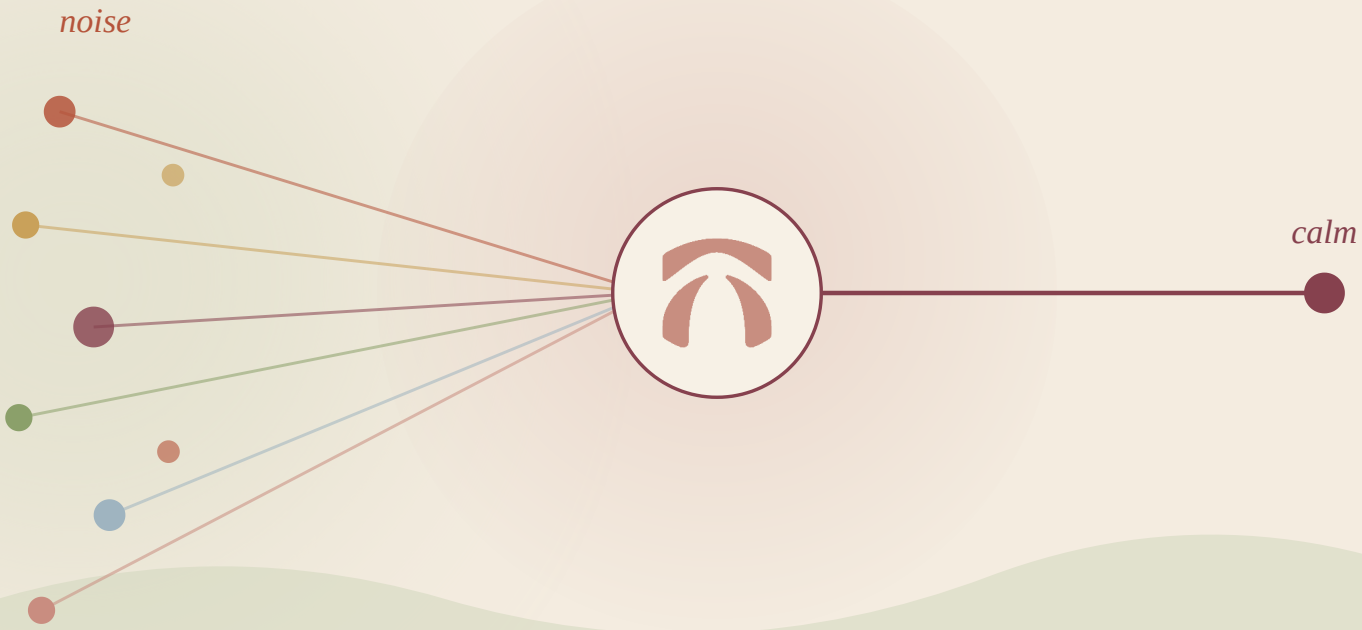

The seam

Purview sees the file. It can't see the AI, or *who's driving it.*

What Purview sees

- A DLP policy match in SharePoint, Exchange, or OneDrive
- File downloads and sharing inside Microsoft 365
- Sensitivity labels and Entra sign-in events

What it can't see

- Source pasted into a browser AI tab on a personal account
- A coding agent reaching a credential, pushing to a personal repo
- An agent moving data through an unsanctioned MCP server
- Whether a person or a non-human identity took the action

11,200+ AI tools, agents, and MCP servers found across Anzenna customers. 38% run on personal accounts, outside the tenant Purview can see.

Every AI security tool sees the agent. Every identity tool sees the person. *Anzenna shows you who's driving the AI,* in one timeline, with response routed through the Microsoft stack you already run.

© 2026 ANZENNA

The problem we hear most in Microsoft shops

Purview catches a lot. In a busy tenant, *thousands of alerts a day drown the signal.*

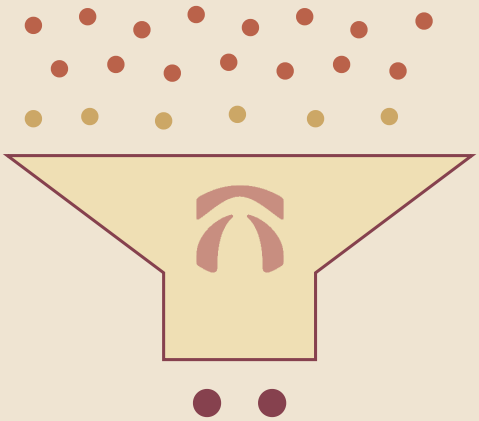

the few that matter

~95%

closed automatically

Anzenna ingests every Purview alert the moment you connect Microsoft 365, then works each one automatically. Only the few that need a person reach your team, already assembled into a case.

AUTO-HANDLED

~470 */ 500*

daily Purview alerts worked without a human

OFF THE TEAM

~130K

alerts a year that never hit the queue

YEAR-ONE VALUE

~$675K

modeled on a 5,000-employee environment

Ingest

every Purview alert, automatically, the moment you connect Microsoft 365.

→

Investigate

each one with an agent that gathers the context and the evidence.

→

Escalate

only the few that need a person, already assembled into a case.

© 2026 ANZENNA

How it works

Discover the AI, audit what it does, *anchor it to a person*

01

Discover

Agentless inventory of every AI tool, agent, and MCP server across the fleet, corporate and personal.

02

Audit

A per-prompt record of what was entered, files moved, MCP destinations, and OAuth grants.

03

Classify & act

Native classification at the point data moves. Warn, redact, or block inline in the browser and IDE.

04

Anchor

Every action tied to the identity behind it, correlated into one case, the benign majority closed on its own.

Alongside your Microsoft stack

- Microsoft 365
- Defender
- Intune
- Entra
- Purview
- Copilot

Resolved, attributed, evidence attached

One integration with Microsoft pulls Purview, Defender, Intune, Entra, and Copilot signal. Remediation routes back through the tools you already run, with no second agent to deploy. You keep Purview for content inspection and compliance, and Anzenna adds the AI visibility, investigation, and response.

Why it holds up

Detection quality

Net-new, high-signal detections tuned to your environment, from the browser and the endpoint, not half the picture.

Cases, not alerts

The investigative agent does the analyst's first pass, so what reaches a person is already worked up.

A forward-deployed engineer

Real detection and remediation engineering alongside your team, not a support tier.

© 2026 ANZENNA

What it looks like · illustrative

One departing engineer. *One coding agent.*

Maya gave notice on Monday. To Purview, her week looks ordinary. Over three days, the AI tells a different story.

MONDAY
Gives notice. HR updates. No DLP flags.

BROWSER AI
Pastes an internal API key into a personal AI tab.

CODING AGENT
Copies two repos to a personal GitHub.

MCP
Moves design docs through an unsanctioned MCP server.

OAUTH
Grants a token to an unused AI file-sync app.

What Purview sees

At most one DLP match, and only if a labeled file crosses a monitored boundary in Microsoft 365. The prompt, the coding agent, the MCP transfer, and the personal repo never register.

What Anzenna sees

The whole chain as one worked case: the pasted API key, the repos copied to a personal GitHub, the MCP transfer, and the OAuth grant, all tied to Maya and her departure four days earlier.

The investigation, auto-assembled

RISK · HIGH

SIGNALS CORRELATED	CONTEXT ADDED	RECOMMENDED
Five, across browser, IDE, MCP, and identity	Departure flagged four days earlier	Revoke OAuth, block MCP, reset session

20 min

From first push to a closed case: revoke the grant through Entra, block the MCP destination, notify legal. Previously a two-day job across four consoles.

The same pattern, at enterprise scale · production deployments

2.1*M*
prompts audited every month across browser, IDE, and SaaS

73,600
sensitive items warned, redacted, or blocked

960+
departing-employee moves flagged before exit

$4*M*
largest theft caught, days before an exit

© 2026 ANZENNA

anzenna

SOLUTION BRIEF

Keep Purview.

Add the layer that sees your AI

marketing@anzenna.ai · www.anzenna.ai

© 2026 ANZENNA