Anzenna

Microsoft Purview + Anzenna

Turn DLP Signals into
People-Centric Security Operations

Microsoft Purview is a strong foundation for data protection within the M365 ecosystem. It provides deep content inspection, DLP policy enforcement, sensitivity labeling, and compliance management. But when an insider threat spans systems beyond Microsoft, Purview's alerts only tell part of the story.

Anzenna is the agentic layer for insider risk management. It extends Purview by combining M365 telemetry with identity, endpoint, SaaS, code, and HR signals from 120+ integrations to build a complete picture of every person's activities. AI agents autonomously investigate every detected threat and trigger remediation instantly. Your security team gets full visibility into user behavior across every system, not just M365, and can act on risk the moment it appears, without adding headcount.

The Reality: Purview Sees M365, Not the Whole Person

Purview delivers strong capabilities within the Microsoft ecosystem: DLP content inspection with 300+ sensitive information types, policy enforcement in SharePoint, OneDrive, Exchange, and Teams, and sensitivity labeling across Office apps. For protecting data inside M365, this coverage works well.

But insider risk doesn't stay inside one ecosystem. A departing employee downloads customer data from Salesforce, shares files externally through Google Drive, pushes proprietary code to a personal GitHub repository, and then triggers a Purview DLP alert when they email an attachment. Purview catches that last step, but the full picture requires context from identity providers, HR systems, code platforms, and endpoints that Purview doesn't monitor.

What Purview Sees

- DLP policy violation in SharePoint or Exchange
- File download and sharing activity within OneDrive and SharePoint
- Sensitivity label applied or modified on an M365 document
- Entra ID sign-in events and Conditional Access results
- Admin actions within the M365 admin audit trail

What Gets Buried or Missed

- The user was flagged as departing in Workday three days ago
- That same user logged in from a new country via Okta this morning
- They exported 200 Salesforce reports in the last 48 hours
- They granted a risky OAuth app full read/write access to Google Workspace
- They pushed proprietary code to a personal GitHub repository last night
- Their CrowdStrike endpoint shows USB file writes to a removable drive

Purview's DLP violation is real and worth investigating. But without behavioral analytics that connect identity, HR context, endpoint activity, and SaaS usage across every platform, your team sees only one alert rather than a coordinated insider threat.

Why 'We're Already Covered' Is a Dangerous Assumption

Security teams often believe insider risk is handled because they've invested in Purview and other strong tools. But each tool operates in its own silo. None provides the full picture.

"We Have Purview. Our Data Is Protected."

Reality: Purview protects data within the M365 ecosystem through content inspection and policy enforcement. It doesn't monitor activity in Salesforce, Google Workspace, GitHub, Snowflake, or dozens of other platforms where your data also lives. A DLP alert in SharePoint shows you a single event, not the employee's full risk profile.

"We Have Defender for Cloud Apps. OAuth Is Covered."

Reality: Defender for Cloud Apps monitors app connections to your M365 tenant. It doesn't assess OAuth grants across Google Workspace, Slack, Salesforce, or GitHub. Anzenna's AI scores every OAuth app across all connected platforms using risk categories, grounded in web research.

"We Have DLP. Data Exfiltration Is Blocked."

Reality: Purview DLP blocks sensitive content from being shared in violation of your policies. But DLP can't catch what it can't see: code pushed to Git repositories, data exported from Salesforce, files copied via AirDrop, or uploads to AI tools like ChatGPT. These channels operate outside Purview's content inspection.

"We Have a SIEM. Everything Is Logged."

Reality: SIEMs aggregate logs without user context. Your analysts manually pivot between Purview, Entra ID, Defender, and external tool consoles to investigate a single alert. Investigations stretch from minutes to hours.

"How Many People Do I Need to Run This?"

Reality: Zero new hires. Anzenna automates the investigation workflow that currently consumes your team's time. Your existing analysts review finished cases instead of manually pivoting across the Purview portal, Entra ID, Defender, and external tool consoles. Organizations handling 50+ insider risk cases per month recover 25-50 analyst hours by automating investigation, employee notification, escalation, and follow-up. One analyst can handle 50+ cases per day when the investigation arrives already completed.

Anzenna's Four Core Pillars

People and Priority First

Every activity is tied to a user identity, enriched with HR context and historical behavior. Your security team focuses on high-risk users exhibiting concerning patterns rather than individual alerts in the Purview portal. A DLP violation becomes: "Sarah Chen (Senior Engineer, departing, 92nd percentile risk score) shared a sensitive document externally, with 6 other risk factors across 4 platforms this week" instead of just "DLP policy match: High severity."

Cross-Platform Behavioral Analytics

Purview monitors the M365 ecosystem. Anzenna extends behavioral monitoring to 120+ integrations spanning identity providers (Entra ID, Okta, Google Workspace), endpoints (CrowdStrike, SentinelOne, Defender), HR systems (Workday, BambooHR, Rippling), code platforms (GitHub, GitLab), CRM (Salesforce), databases (Snowflake), and more. A Purview DLP violation is far more meaningful when correlated with activity across every system the employee touches.

Agentic AI That Investigates and Remediates Autonomously

When a Purview DLP violation fires (or any other detection), Anzenna's AI agents activate autonomously. They examine weeks of cross-platform activity across every connected system, pull the employee's full context, connect individual events into narratives, and produce a structured investigation brief with evidence, confidence scoring, and a recommended response. AI agents investigate every threat, not a subset. For cases that require human judgment, agents deliver the investigation already completed: timelines built, evidence linked, and actions recommended. When threats are confirmed, agents trigger coordinated remediation across the entire stack: reset passwords, revoke OAuth grants, suspend accounts, unshare documents, and notify the employee via Slack, Teams, or email, all at once.

Closing the Loop with Employees

Purview alerts your security team. Anzenna takes it further by managing the full response lifecycle: notifying the employee in 9 configurable tones, sending reminders on schedule, escalating to their manager if unresolved, escalating to skip-level management if still open, and creating tickets in Jira, ServiceNow, Freshservice, or Zendesk. Your team tracks every action item from detection through resolution.

What Anzenna Adds on Top of Purview

Capability	Purview Alone	Purview + Anzenna
DLP response	Policy enforcement (block, restrict) and alert routing to the security team	AI investigation brief with cross-platform evidence, automated employee notification, escalation, and ticketing
Who is affected?	Policy match severity (Low/Medium/High)	Named employee with unified risk score from 20+ categories, peer percentiles, and full risk history
Agentic AI investigation	No built-in AI investigation. Analysts manually review cases in the Purview portal	Every threat is investigated autonomously by AI agents. Analysts receive structured briefs with evidence and recommended actions
Identity monitoring	Entra ID sign-in reports	7+ identity providers: impossible travel, proxy/VPN detection, foreign country, service account misuse
OAuth governance	Defender for Cloud Apps (separate license)	AI-powered risk scoring across M365, Google Workspace, Slack, Salesforce, and GitHub
Endpoint coverage	Defender for Endpoint and Intune	Multi-vendor: CrowdStrike, SentinelOne, Cortex XDR + Jamf, Kandji, Workspace ONE + Intune
Anomaly detection	Policy-based rules you define	Statistical anomaly detection (z-score, Mahalanobis) across 4 comparison groups and 4 time windows
HR context	Manual triggers or limited Entra ID signals	Direct integration with 30+ HR systems for automatic employment lifecycle tracking
Remediation	Policy-based blocking within M365 apps	Coordinated response across M365, identity providers, SaaS, and endpoints, all triggered automatically
Non-Microsoft visibility	Not covered	Salesforce, GitHub, Snowflake, Google Workspace, Slack, Dropbox, Box, and 100+ more

Deep Microsoft Ecosystem Ingestion

Anzenna doesn't just connect to Microsoft. It has deep, native integration across the entire M365 and Azure ecosystem. Here's what Anzenna pulls from your Microsoft environment:

- **Entra ID (Azure AD):** Employee directory, sign-in logs (interactive and non-interactive), MFA enrollment status, Conditional Access results, token protection status, service principal data, OAuth app permissions, and risk detections
- **Office 365 Audit Logs:** 300+ event types including SharePoint/OneDrive file operations, sharing events (direct, anonymous, company-wide, and secure links), site administration, and compliance events
- **DLP Policy Violations:** Every Purview DLP match event becomes a behavior attributed to a named employee, including severity, policy name, and incident ID for investigation continuity
- **Intune:** Managed device inventory, compliance status, app installations, and enrollment status
- **Defender ATP:** Threat detections (remote code execution, USB exfiltration, source code exfiltration via Git), installed software inventory, and real-time response capabilities
- **Microsoft Teams:** Team membership data, native notification delivery via a Teams app, interactive response buttons, and automated message management

Anzenna's AI identifies service accounts and human users based on login patterns and account metadata. Service accounts that authenticate interactively (potential credential misuse) trigger a dedicated detection.

Anomaly Detection Beyond Policy Matching

Purview fires when a predefined rule is triggered. Anzenna fires when behavior deviates from what's statistically normal, even when no rule exists for that specific pattern.

Z-score analysis
Single-dimension anomalies like "This person shared 5x more files than usual today"

Mahalanobis distance
Multi-dimensional correlated anomalies where moderate increases in both file count and file size wouldn't trigger individually, but are statistically significant in combination

Per-event analysis
Login-specific anomalies, including impossible travel, proxy/VPN detection with operator identification, foreign country access, and new IP network

Anzenna compares each employee against four baseline groups: their own historical behavior, peers in the same role, their department, and other employees with the same employment status (departing, new hire, contractor). It evaluates activity across four time windows: 14 days, 8 weeks, 8 months, and 4 quarters.

A departing sales rep who downloads twice as many Salesforce reports as their departing peers, while also sharing SharePoint documents externally at 3x their normal rate, triggers anomaly detection even if neither activity individually crosses a Purview policy threshold.

M365-Specific Detections Powered by Anzenna

Beyond ingesting Purview's signals, Anzenna runs its own detections across Microsoft data:

- **Risky M365 OAuth Apps:** Third-party apps connected to M365 with high AI-generated risk scores (13 risk categories, web-grounded research)
- **Risky and Public M365 Shares:** Documents shared with risky external recipients, excessive permissions, or public anonymous links
- **Departed Employee Shares:** Documents shared by link or externally by employees flagged as departed by HR
- **Impossible Travel Logins:** Entra ID sign-ins from geographically impossible locations, correlated across all identity providers
- **Proxy/VPN and Foreign Country Logins:** Sign-ins through consumer VPN services (NordVPN, TOR exit nodes) or from countries where the employee has never authenticated
- **Admin Threat Detection:** 10 admin action types: impersonation, privilege grants, email access, forwarding, security changes, and more
- **DLP Violation Anomalies:** Volumetric spikes in DLP violations compared to employee and peer baselines

Real-World Scenario: The Departing Sales Manager

DAY 1

Resignation Submitted
Alex, a senior sales manager, submits his resignation. HR updates his status in Workday. Purview sees normal M365 activity. He's accessing the same SharePoint sites, sending regular emails, and working in Teams as usual. No DLP violations, no policy matches.

DAY 4

Behavioral Changes Begin
Alex exports 200 Salesforce reports containing customer contact information and deal pipeline data. He downloads 30 SharePoint documents from a team site he rarely visits. He shares a confidential pricing spreadsheet via OneDrive from his personal Gmail account. He grants OAuth access to a file sync app his team doesn't use. He logs into Entra ID from a hotel in a city he's never visited for work.

What Purview Sees:

- DLP policy match: High severity (pricing spreadsheet shared externally)
- SharePoint file downloads from a team site
- Entra ID sign-in from an unusual location

What Anzenna Sees:

- Alex's HR status changed to 'departing' four days ago
- Salesforce export volume is 10x his baseline: 200 reports in 48 hours vs. the typical 20/week
- SharePoint download volume is 6x his baseline across a site he hasn't accessed in 3 months
- The OneDrive share targets a personal email address outside normal collaboration patterns
- The OAuth app has a high AI risk score (data sync capability, unknown publisher, low adoption)
- Impossible travel: Entra ID login from Chicago at 9 am, Okta login from New York at 10 am
- Cross-platform pattern: departing + bulk exports + external sharing + risky OAuth + impossible travel

Anzenna's Response:

Within minutes of Alex's bulk Salesforce exports, Anzenna flags him as high-risk. An AI agent connects his HR departure notice from Workday to the spike in report exports, the SharePoint downloads from a dormant site, the external OneDrive share to his personal Gmail, the unauthorized OAuth grant with a high risk score, and the impossible travel between Entra ID and Okta. By the time his analyst opens the case, a structured brief is waiting: seven correlated signals, a timeline, evidence links, and a recommended response. With one click, the team revokes the OAuth grant, removes the external share, resets Alex's Entra ID session, and notifies legal. Alex receives an automated Teams message explaining the hold. Total response time: 10 minutes.

Without Anzenna, your team would see the Purview DLP alert and start a manual investigation. The Salesforce exports, risky OAuth grant, impossible travel across two identity providers, and the HR departure status would require pivoting across four separate consoles. The full picture would take hours to assemble, if it was assembled at all.

Business Impact

Anzenna transforms insider risk from a resource-intensive manual process into an automated, scalable operation. The combination of Purview's M365 telemetry with 120+ additional integrations gives your team a level of visibility that would otherwise require dedicated headcount and custom tooling to achieve.

Efficiency Gains

Metric	Purview Alone	Purview + Anzenna
Investigation time per threat	1 to 4 hours (manual cross-console review)	Minutes. Analysts review finished briefs, not raw alerts
Time from detection to response	Hours (manual investigation, outreach, follow-up)	Minutes. Detection, investigation, notification, and escalation happen automatically
Analyst capacity	1 analyst handles 5 to 10 cases/day	Same analyst handles 50+ cases/day
Identity anomaly detection coverage	Entra ID only	7+ identity providers simultaneously
Behavioral baselines	M365 activity signals	120+ integrations, 4 comparison groups, 4 time windows
OAuth app risk assessment	Requires Defender for Cloud Apps	AI-powered across 5 platforms, continuous re-evaluation
Departing employee monitoring scope	M365 activity with manual or limited triggers	All connected systems, automatic HR-triggered activation

Cost Savings

Investigation automation
Eliminates 1 to 4 hours of manual cross-console correlation per threat. Every investigation arrives ready for analyst review.

Analyst force multiplier
At 50 cases/month, your team recovers 50 to 200 hours of investigation time. One analyst covers workloads that previously required three.

Remediation coordination
Notification, escalation, and ticketing are automated, eliminating manual follow-up that typically takes 10 to 15 minutes per action item.

Defender for Cloud Apps dependency
Anzenna's AI-powered OAuth governance may reduce dependency on separate Cloud Apps licensing for app risk assessment.

Insider threat program staffing
AI-powered investigation and remediation supplements or replaces the need for 1 to 2 dedicated insider threat analysts ($150K to $350K/year fully loaded).

How Anzenna Works With Purview

Anzenna connects to your Microsoft environment via a secure API connection. You don't deploy agents to endpoints or modify existing Purview policies. Setup takes about 12 minutes and you see value right away.

- **100% Agentless:** Connects directly via Microsoft Graph API. Zero endpoint performance impact.
- **Quick Setup:** Register Anzenna in Azure AD, grant required Graph API scopes (admin consent), and enter your Tenant ID, Client ID, and Client Secret. Anzenna validates the connection and begins collecting data instantly.
- **No Purview License Required:** Anzenna works with any M365 tier. It does not require Purview licensing to deliver value.
- **Bi-Directional Response:** Reset passwords, revoke OAuth grants, suspend accounts, and revoke sessions via Entra ID while simultaneously acting across all other connected platforms.
- **Multi-Tenant Support:** Organizations with separate M365 environments (acquisitions, business units, regional tenants) can monitor them all from a single Anzenna console.
- **Preserves Your Purview Investment:** Purview continues to support content inspection, policy enforcement, sensitivity labeling, and compliance management. Anzenna adds the behavioral analytics, AI investigation, and operational workflows on top.

Appendix: How Anzenna Works, The OODA Framework

Anzenna operationalizes insider risk through a four-stage methodology that turns fragmented data into coordinated action.

Observe

Collect Without Agents
Collect user activities across SaaS, Cloud, endpoint devices, and data movement to build a complete behavioral baseline without agents.

Orient

Enrich with Context
Contextualize raw signals with identity, HR status, and historical behavior patterns to distinguish normal work from actual risk.

Decide

Investigate with Agentic AI
A fleet of specialized AI agents investigates every alert, each handling a different stage of analysis. Your analysts receive finished briefs, not raw data.

Act

Remediate Immediately
Execute automated remediation actions (block, revoke, notify) directly from the platform with or without human approval.

Appendix: What Purview Does Well (and Anzenna Doesn't Replace)

Purview Capability	Why It Stays
DLP content inspection (300+ sensitive info types, custom classifiers, exact data match)	Purview's content inspection engine is purpose-built for deep packet-level analysis within M365 apps. Anzenna ingests the results, not the raw content.
DLP policy enforcement (block sharing, prevent downloads, restrict printing)	Real-time blocking requires native integration with the application layer that only the platform vendor can provide.
Sensitivity labels (manual and auto-applied labels for M365 documents)	Sensitivity labeling within M365 is deeply integrated with the Office client experience.
Compliance Manager (regulatory assessment, improvement actions, compliance score)	Compliance framework mapping is a specialized capability built for audit preparation.
eDiscovery (legal hold, content search, case management)	eDiscovery is a legal function with specific chain-of-custody requirements.
Endpoint DLP (content inspection on Windows/macOS via Defender)	Endpoint-level content blocking requires kernel-level integration via the Defender agent.