Kavryl

Built for the AI agent era

Kavryl Executive AI Agent Permission Risk Report

A board-ready synthetic report showing how Kavryl summarizes cross-cloud, SaaS, identity, and agent exposure.

Assessment metadata

Sample customer	Northstar Health Cloud
Assessment ID	KAV-ASM-2026-05-NSHC-001
Tenant IDs	Azure tenant 7a62d5b4-91d0-4cf5-ae83-4dd9e3d8a210; AWS org o-a9f2xk4p71
Subscriptions / accounts	Azure sub 4b812c7f-301d-45fb-8a4d-8a7750c4c19f; AWS acct 391742608155
Scanner coverage	Azure, AWS, GitHub, Microsoft 365 / Entra, Google Workspace, MCP/local agent config
Data note	Synthetic data modeled after kavryl-scan output. Resource IDs are fake and safe for public demo use.

Executive summary

Kavryl found a High exposure environment where AI-enabled support and data agents can reach sensitive SaaS, cloud, and repository systems. The highest risk is not one isolated misconfiguration. The highest risk is the combined path: broad OAuth scopes, privileged cloud roles, public AI endpoints, and agent workflows that can read sensitive data and perform outbound actions.

Top business risks

- A support agent can read ticket context, retrieve clinical documents, and post outbound Slack messages without consistent approval gates.
- An Azure service principal used by agent workflows has Owner scope at subscription level, increasing cloud blast radius.
- AWS Bedrock automation role includes wildcard IAM signals and can be invoked by a public API Gateway endpoint.
- Microsoft Graph delegated grants include Mail.Read, Files.Read.All, and offline_access on an agent-like OAuth app.
- GitHub repositories contain agent and MCP configuration signals, but repository ownership and token rotation evidence is incomplete.

Findings snapshot

KAV-001	Critical	Azure	Agent service principal has subscription Owner role	Replace with least-privilege custom role scoped to rg-ai-agents-prod

KAV-002	High	M365	Support Agent OAuth App has broad Graph permissions	Remove unused Graph scopes and split read/write actions into separate apps
KAV-003	High	AWS	agent-bedrock-admin-role has wildcard permission/trust signal	Restrict trust policy and remove wildcard actions
KAV-004	High	MCP	Local MCP server exposes filesystem and shell tools in same agent context	Separate tools, add allowlist, disable shell for support workflows
KAV-005	Medium	Google	Workspace service account can access Drive folders used by RAG	Restrict Drive access by folder, group, and data classification

Priority action plan

0-7 days	Disable shell-capable MCP tools for support agents and require approval before outbound Slack/email actions.	Reduce immediate data exfiltration and unsafe action risk.
0-14 days	Replace Azure Owner assignment and AWS wildcard role with least-privilege roles.	Reduce cloud blast radius from compromised agent identity.
15-30 days	Review OAuth grants, remove offline_access where not required, split high-risk Microsoft Graph scopes.	Reduce SaaS data exposure and delegated access persistence.
30-60 days	Add recurring Kavryl scans and owner review for each agent/tool identity.	Create operating rhythm for agent exposure management.

Framework interpretation

The findings map primarily to OWASP API Security API2, API5, API8, and API9, plus MITRE ATLAS-style paths around AI-enabled system discovery, credential misuse, excessive agency, and public agent/API exposure. For executives, the plain-English meaning is: agent identities should be governed like privileged users, but current controls do not yet show the full access path.