

Teleskope vs Varonis

Overview

As organizations accelerate cloud adoption and AI initiatives, selecting a data security platform that moves beyond visibility to sustained risk reduction is critical.

Varonis offers deep visibility in regulated environments, but often at the cost of

high operational overhead and slower adaptation to modern data usage.

On the other hand, [Teleskope](#) provides continuous, measurable reduction of data risk with significantly less operational effort, even as data usage, sharing, and AI adoption accelerate.

Why Teleskope

DSPM tools, such as Varonis, show you posture, but rely on people to prioritize and resolve risk. Teleskope reduces risk by intelligently enforcing your policies. Here is why it's different:



Relevant, Prioritized Risk

- Finds risks across petabytes of any type of data 10X faster
- Prioritizes exposure that matters specifically to your business
- Enables high-confidence automation instead of manual triage



Native, Automated Remediation

- Resolves risk directly instead of routing to tickets or external tools
- Mitigates existing exposure and new risk as it occurs
- Shortens time to risk reduction, not just time to insight



Auditable, Intelligent Controls

- Enforces your existing policies and follows your workflows
- Humans stay involved where validation is required
- Actions are safe, auditable, and reversible

How Teleskope Compares to Varonis

Teleskope and Varonis both address data risk, but from different starting points. Varonis is built around visibility and access analytics; Teleskope is built to continuously reduce risk through prioritization and automated policy

enforcement. The comparison below focuses on where those differences matter most: Risk Assessment & Prioritization, Risk Reduction & Policy Enforcement, Risk Prevention & AI Agents, Deployment, and Customer Support.

RISK ASSESSMENT & PRIORITIZATION



Highly accurate classification (low false positive or negatives) that scales across structured and unstructured data of various formats, which enables high-confidence automation of risk reduction.

Wide regex-based classification coverage, but tends to generate **high false positives**, especially when data patterns change or when deployed across large, heterogeneous datasets. Effective for access controls and certain legacy file systems

- Multi-stage AI pipeline: lightweight ML models route to specialized SLMs and LLMs, which combine contextual reasoning, semantic understanding, and business-specific policy signals
- Classifies a broad array of entity types (PII/PHI/PCI, credentials/secrets, custom data types)
- Classifies document types (e.g., IP, payroll, etc.) and provides redacted AI-generated summaries
- Differentiates dormant sensitive data from actively used or broadly accessed data, enabling prioritization
- Associates classified data with business relevance (e.g., customer vs employee data) to drive enforcement.

- Uses behavior-based threat models to detect abnormal behavior
- Detects weak identity and system configuration
- Dictionary keywords and regex-driven pattern matching
- Behavioral analytics and access metadata
- Requires hours or days of tuning per environment to reach acceptable precision
- Classification outputs require human validation before action
- Only catalogs data that it deems sensitive. This prevents real data lifecycle management and limits risk visibility
- Limited structured data monitoring.

RISK REDUCTION & POLICY ENFORCEMENT



Reduces risk **automatically** and **continuously**, without tickets, external tools, or manual follow-through, resulting in materially faster time to risk reduction and lower operational burden on security teams.

- Native actions: access revocation/scoping, sharing restriction, redaction/masking, encryption, cleanup of overexposed/stale sensitive data
- Embedded into existing workflows
- Works for existing exposure and in real time
- Supports full automation or human-in-the-loop based on risk and confidence



Actions are **manual** and **reactive**. Native remediation and redaction limited to some controls around access and permissions. In most cases, automation is impossible due to a high rate of false positives.

- Native automation is restricted to narrow scopes: remove excessive permissions, fix risky misconfigurations, apply labels
- Remediation often triggered after alerts or investigations
- Human review frequently required before action
- Time to risk reduction depends heavily on team capacity

RISK PREVENTION & AI AGENTS

Organizations can **safely adopt AI tools and agentic workflows** without increasing data exposure, because sensitive data is cleaned up in **real time** as it is used by AI systems.

- Detects and controls access to data at rest and in use by AI tools, copilots, and agents
- Prevents overexposed or non-compliant data from being used in AI workflows
- Enables AI adoption without requiring manual approvals or separate governance tools
- Supports continuous enforcement as AI usage evolves

Customers gain **visibility into sensitive data that could be used by AI tools**, but must rely on manual controls and existing access governance to prevent AI-related risk.

- Observes copilot activity from prompt to response and alerts based on risky behaviors
- Response is largely limited to alerting and permission controls: it can revoke/adjust access within M365, but it cannot remediate data in real time as it is ingested by or emitted from copilot, nor can it sanitize, redact, or block content

DEPLOYMENT



Deploys in hours to days across cloud, SaaS, and on-prem without agents, enabling faster time to value and consistent risk reduction across hybrid environments.

- Agentless, API-driven architecture across SaaS, cloud, and AI data flows
- Flexible deployment (SaaS or self-hosted) with on-prem scanning
- Single control plane across hybrid environments



Longer deployments and slower time to value—can take several weeks to months. Additionally, no option to air-gap the data as Varonis can't be self-hosted.

- SaaS deployment only, no on-prem
- Agents are used for activity monitoring on Windows systems only
- A fairly robust coverage of SaaS products
- Limited structured data monitoring

CUSTOMER SUPPORT

Support model is tied to **business outcomes**: every engagement starts with the outcome alignment.

- Every customer is paired with a hands-on customer engineer who actively helps define use cases, configure policies, and drive outcomes.
- Customer work is structured around business objectives (e.g., insider threat reduction, IP protection, audit streamlining) and measured through POC discovery, kickoff plans, and quarterly outcome reviews.

Support model is tied to **professional services** needed to fine-tune and support classification models.

- Support model is tightly coupled with professional services
- Significant time spent on tuning, suppression, and alert management
- Ongoing changes (new data stores, new users, new permissions) frequently require renewed engagement
- Supports troubleshooting, but ownership of outcomes remains with the customer