sentra

The Dirt on DSPM POVs

What Vendors Don't Want You to Test

Discover What DSPM Vendors Try to Hide

Your goal in running a data security/DSPM POV is to evaluate all important performance and cost parameters so you can make the best decision and avoid unpleasant surprises. Vendors, on the other hand, are looking for a 'quick win' and will often suggest shortcuts like using a limited test data set and copying your data to their environment.

On the surface this might sound like a reasonable approach, but if you don't test real data types and volumes in your own environment, the POV process may hide costly failures or compliance violations that will quickly become apparent in production. A recent evaluation of Sentra versus another top emerging DSPM exposed how the other solution's performance dropped and costs skyrocketed when deployed at petabyte scale. Worse, the emerging DSPM removed data from the customer environment — a clear controls violation.

⚠ **If you want to run a successful POV and avoid DSPM buyers' remorse you need to look out for these "dirty little secrets".**

Dirty Little Secret #1

'Start small' can mean 'fails at scale'

The biggest 'dirty secret' is that scalability limits are hidden behind the 'start small' suggestion. Many DSPM platforms cannot scale to modern petabyte-sized data environments. Vendors try to conceal this architectural weakness by encouraging small, tightly scoped POVs that never stress the system and create false confidence. Upon broad deployment, this weakness is quickly exposed as scans slow and refresh cycles stretch, forcing teams to drastically reduce scope or frequency. **This failure is fundamentally architectural, lacking parallel orchestration and elastic execution, proving that the 'start small' advice was a deliberate tactic to avoid exposing the platform's inevitable bottleneck.**

In a recent POV, Sentra successfully scanned **10x more data** in approximately the same time than the alternative:

sentra

9 PB Scanned

<72 hrs

A Top Emerging DSPM Vendor:

0.9 PB scan failed

~60 hrs

© Sentra 2026. All rights reserved.

Dirty Little Secret #2

High cloud cost breaks continuous security

Another reason some vendors try to limit the scale of POVs is to hide the real cloud cost of running them in production. They often use brute-force scanning that reads excessive data, consumes massive compute resources, and is architecturally inefficient. This is easy to mask during short, limited POVs, but quickly drives up cloud bills in production. The resulting cost pressure forces organizations to reduce scan frequency and scope, quietly shifting the platform from continuous security control to periodic inventory. Ultimately, tools that cannot scale scanners efficiently on-demand or scan infrequently trade essential security for cost, proving they are only affordable when they are not fully utilized.

In a recent POV run on 100 petabytes of data, Sentra proved to be **10x more operationally cost effective** to run:

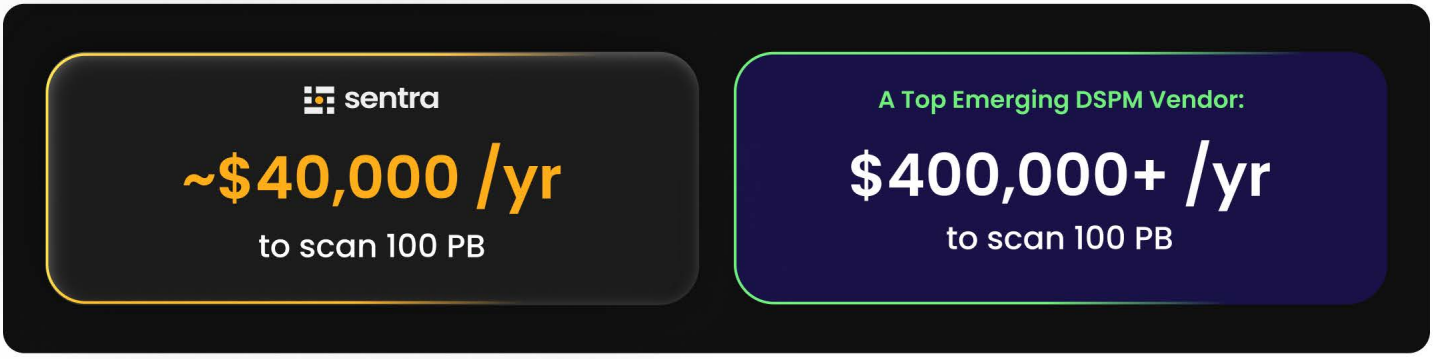

Dirty Little Secret #3

AI is quietly exploding your data exposure

The most overlooked dirty secret is how dramatically **AI increases data security risk**, not by creating new data, but by expanding access to existing sensitive data at machine speed. Copilots, agents, and AI pipelines inherit permissions that were designed for humans, not autonomous systems, and can traverse vast data estates in seconds. Many DSPM vendors downplay this risk by treating AI as just another application, failing to model how AI systems access, move, summarize, and reuse sensitive data. **This blind spot remains hidden during POVs that focus on static discovery or small datasets, never testing AI access paths or AI-driven data movement.**

In production, the weakness becomes obvious. Overshared data that seemed low risk suddenly becomes high impact when AI systems can aggregate and expose it across teams, regions, or external outputs. Without AI-aware data discovery, classification, and governance, organizations lose control over what data AI can see and how it can be used. The result is unmanaged shadow data, regulatory exposure, and AI systems interacting with sensitive or regulated information without enforceable guardrails — proving that ignoring AI in a DSPM evaluation is not simplifying the POV, but concealing one of the largest data risk accelerants in the enterprise.

© Sentra 2026. All rights reserved.

POV Best Practices

1. Data stays in customer environment
2. Test coverage of different file formats
3. Test customization of policies and classification
4. Test real time monitoring (DDR)
5. Test scale and cloud cost efficiency

POV Worst Practices

1. Copy data to vendor environment so they can hide costs and 'fix' the results
2. Limit features and capabilities
3. Limit size of scanned data
4. Restrict integrations to avoid "complications"
5. Limit the use of API

Four DSPM POV Requirements That Expose the Truth

Scalability	Run discovery and classification on at least 1 petabyte of real data including unstructured object storage. Completion time must be measured in hours or days — not weeks
Cost Efficiency	Operate scans continuously at scale and measure actual cloud resource consumption. If cost forces reduced frequency or scope, the model is unsustainable.
Align POV with real deployment	Require the POV to mirror how the platform will actually run in production; same architecture, integrations, permissions, and operating model.
Unstructured Data Depth	Test long-form, heterogeneous, real-world unstructured data including audio, video, etc. Classification must demonstrate contextual understanding, not just keyword matches.

A DSPM solution that only performs well in a limited POV will lead to painful, costly buyer's regret. Once in production, the failures in scalability, cost efficiency, accuracy, and unstructured data depth quickly become apparent.

Schedule a Demo

Gartner Peer Insights
Highest Recommended DSPM Platform

Gartner Peer Insights™

4.9 ★★★★★

By Sentra in Data Security Posture Management (DSPM)

- **4.7** Evaluation & Contracting
- **4.6** Integration & Deployment
- **4.8** Service & Support

98.5%
Recommendation Rate

© Sentra 2026. All rights reserved.