cyberhaven

Data Loss Prevention Buyer's Guide

8 Criteria *for* Evaluating DLP Solutions

The Technology We Use Is Evolving, and DLP Must Evolve Alongside It

We stand at an inflection point in enterprise data security. AI has not only democratized access to knowledge, it has integrated itself into the core workflows of every organization, reshaping how data is created, shared, and consumed. At the same time, the way people work has fundamentally changed. Sensitive information now moves constantly across SaaS platforms, endpoints, messaging tools, and collaboration environments, often reaching humans and systems that were never meant to have access to it. Exfiltration paths that didn't exist a few years ago are now routine.

Data has broken free from the systems designed to contain it.

Traditional data loss prevention (DLP) was built for a different era. It assumed a well-defined perimeter, a manageable number of data flows, and the ability to inspect traffic at the network edge. None of those assumptions hold today. The modern enterprise is cloud-first, boundaryless, and increasingly AI-driven. Data no longer lives quietly in repositories. It flows continuously and is consumed by AI systems both inside and outside organizational boundaries. The tools built to protect data must reflect this new reality.

These changes also means the threat landscape has grown more complex. Device-to-device transfer technologies like Apple AirDrop and Windows Nearby Share move large volumes of data from managed work devices to personal ones in seconds. Generative AI tools create a different kind of exposure: When employees paste sensitive content into genAI tools, that data may be processed or retained by the model provider in ways the organization can't control. Agentic AI introduces risks that are harder still to see, as these agents are capable of accessing and exfiltrating data residing on the endpoint without the user ever knowing it happened.

Meanwhile, a growing number of cloud applications use certificate pinning and end-to-end encryption that renders network-based inspection blind to what's flowing through them. And even where strong technology exists, many organizations find themselves managing a patchwork of point solutions with different policy engines, separate management interfaces, and no shared context, leaving gaps that sophisticated threats are quick to exploit.

These challenges share a common thread: They all have outpaced the assumptions baked into legacy DLP.

Whether you're evaluating DLP for the first time or replacing a solution that hasn't kept pace with your organization's needs, this guide is designed to help you ask the right questions and recognize the right answers. The sections that follow outline the key challenges driving the need for modern DLP, and the criteria you should use to evaluate any solution you're considering. Modern DLP capabilities address risks that simply didn't exist a decade ago. They also fix the structural shortcomings that have made traditional tools more painful to operate than they are effective at protecting data.

In this new era, data has never been more valuable, and protecting it has never been more complex. The goal of this guide is to help you find a solution equal to that challenge.

Section 01

The Challenges with Traditional DLP

Legacy DLP was designed for a data landscape that no longer exists. These are the limitations that organizations run into, and why they matter.

"Traditional approaches to DLP were reactive, preventing data loss only at the corporate boundary, rather than analyzing user risk and adapting controls to secure data throughout its life cycle." — Gartner, Inc*

Challenge	Why It Matters
Classification built on keywords and RegEx	Rules-based content inspection generates a high volume of false positives, drowning analysts in noise and eroding confidence in the tool. Security teams spend more time tuning policies than responding to real threats.
Important data without text content goes unprotected	Source code, design files, recorded meetings, and other high-value assets often contain no recognizable content pattern. If a tool can only classify what it can read, it will miss some of the data that matters most.
Visibility is limited to what you already know to look for	Traditional DLP requires policy definition before it can protect anything. That means unknown data flows, newly created sensitive files, and data that moves in unexpected ways fall outside coverage zones by default.
User-applied labels are inconsistent and don't follow the data	Manual tagging depends on employees making the right call every time. Labels also fail to persist when data is copied, extracted, or moved to a new file, leaving derived content unprotected even when the source was classified correctly.
Agentic AI operates outside user visibility	Autonomous AI systems can access, process, and move data residing on the endpoint without any user action triggering the event. Traditional DLP has no visibility into these actions because it was built around monitoring human behavior, not AI behavior.
Endpoint protection is difficult to get right without disrupting operations	The endpoint is where legacy DLP most often falls short. Poorly built agents slow down devices, interfere with applications, and create enough friction that security teams face pressure to roll back the controls they worked to deploy. Getting endpoint DLP right requires an agent built from the ground up for that environment, not one retrofitted onto an older architecture.
Policies are complex to build and maintain	Many traditional DLP tools require security teams to write JSON, custom code, or elaborate rule logic to define and manage policies. That complexity raises the barrier to effective deployment and makes it difficult to iterate as the business and threat landscape evolve.

* Gartner Market Guide For DLP

Section 02

8 Criteria For Evaluating Data Loss Prevention Solutions

CRITERIA 1

One Product, One Interface, One Policy Engine for All Exfiltration Channels

Data exfiltration doesn't follow a single path. An employee might copy a CAD file to a USB drive, print a strategy document from Google Drive, or push source code to an unsanctioned cloud app. Each of these represents a different channel, a different mechanism, and in most legacy DLP deployments, a different tool with a different policy engine managing it.

That fragmentation is one of the most significant structural weaknesses in traditional DLP programs. When cloud, email, and endpoint protection are handled by separate products, you end up with inconsistent policy coverage across environments, alert data spread across multiple consoles, and no reliable way to know whether a single incident spans more than one channel. Even with a SIEM aggregating alerts, the underlying policy logic remains siloed. What gets caught in one environment may go undetected in another.

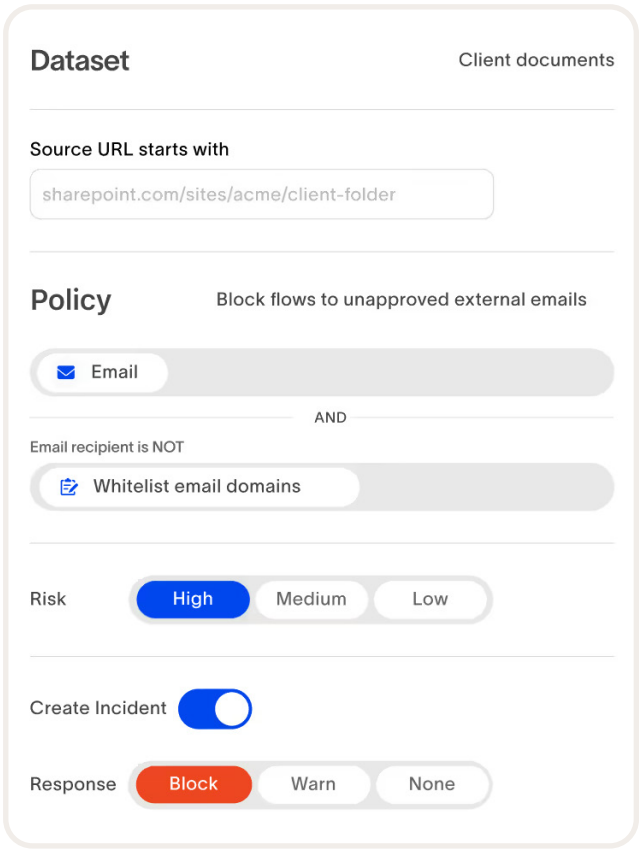

Inconsistent file type support compounds that fragmentation. The DLP capabilities bundled into platforms like Google Workspace and Microsoft Purview are optimized for the file types native to those ecosystems. Microsoft's built-in DLP scans Office documents effectively. It has limited or no ability to classify CAD files, Figma designs, source code repositories, video files, or the dozens of other file formats that represent high-value intellectual property in most organizations. Policies that can't be applied consistently across file types aren't really comprehensive policies. They're partial coverage with the appearance of complete coverage, which in some ways is more dangerous than no coverage at all.

"By 2027, 70% of CISOs in larger enterprises will adopt a consolidated approach to address both insider risk and data exfiltration use cases."
— Gartner, Inc*

Modern DLP should eliminate both forms of fragmentation. Look for a solution that provides a single policy engine enforcing a consistent set of rules across every exfiltration channel, a single interface where policies, classifications, exceptions, users, and watchlists are all managed in one place, and the file type breadth to cover the formats that matter most to your specific organization without requiring custom configuration for each one. The measure of effectiveness shouldn't be which tool caught which incident. It should be whether the data was protected, regardless of how or where the attempt occurred.

Features to look for:

- A single DLP policy engine that enforces standardized policies consistently across all exfiltration channels
- One management interface for creating and maintaining policies, data classification categories, exceptions, users, and watchlists
- Comprehensive file type support, including CAD files, design files, source code, video, and other high-value formats beyond standard office documents

How to leverage this functionality:

Consider a common scenario: An employee downloads a client contract from Salesforce, copies a section into a generative AI prompt to redraft the language, pastes the revised version into a personal web-based document, and then exports it to a USB drive before their last day. Each action crosses a different channel. Without unified policy coverage, any one of those steps could fall outside the visibility of your DLP program entirely.

Configure policies to protect data across every channel where loss can occur, including:

- Personal or unsanctioned SaaS applications
- Copy/paste
- USB storage devices and removable media
- Encrypted messaging channels such as Signal and WhatsApp
- Ingress and egress from generative AI tools and AI agents

* Gartner Market Guide For DLP

CRITERIA 2

Data Lineage That Follows Data Across Applications, Files, and Formats

Most DLP solutions are built around a file-centric model of data. They classify files, tag files, label files, and enforce policies based on what those files contain. That model has a fundamental blind spot: A significant and growing portion of sensitive data never exists as a file at all, and even data that starts in a file rarely stays there.

Sensitive information lives in SaaS application records, structured databases, messaging threads, and collaboration platforms. When a user copies a paragraph from an embargoed press release and pastes it into a WhatsApp conversation, no file has moved, and a file-level DLP tool doesn't see it occur.

The data has moved, the risk is real, and the controls never fired.

The problem is compounded by how freely data moves between contexts in a modern workflow. A user might copy a passage from a confidential strategy document, paste it into a notes application, revise it, paste a portion into an email draft, and ultimately share a fragment via Slack. At each step, any file-level classification or label attached to the original document is left behind. By the time the data reaches its final destination, it has been fully decoupled from the policies that were supposed to protect it.

Content inspection alone can't solve this. Scanning files for keywords and alphanumeric patterns works well in narrow cases: credit card numbers, for example, follow a consistent 16-digit format that can be validated algorithmically. But outside that narrow band, content-only classification breaks down in two directions simultaneously. It generates false positives on data that matches a pattern but isn't sensitive, and it misses data that is genuinely sensitive but contains no recognizable pattern at all. Go-to-market strategy documents, financial models, and investment plans contain no consistent alphanumeric pattern. Recorded meetings, product design files, and marketing assets contain no text whatsoever. For this category of data, content inspection is not a partial solution.

Data lineage addresses both problems mentioned above. By tracking data from its point of origin through every action taken on it within the organization, including where it was created, how it moved, who handled it, and what systems it touched, lineage maintains a continuous record that survives copy/paste, format changes, and application transitions. An extract from a customer database carries the sensitivity context of that database whether it ends up in a spreadsheet, a Slack message, or a prompt entered into a generative AI tool. The origin of the data is what defines its risk, and that context needs to be durable enough to survive the workflows employees actually use.

When evaluating DLP solutions, data lineage should not be a feature bolt-on. It should be a core capability, embedded in the classification engine, available as a policy condition in its own right, and seamlessly follow data wherever it goes.

Features to look for:

- A DLP platform that integrates data lineage as a core capability, tracking data from its origin through every step it takes within the organization
- Data lineage that operates at a granular level, including when data moves from a file into a SaaS application, from one application to another, or through copy/paste actions across contexts
- Policies that support combining content and lineage-based rules to improve classification accuracy and reduce false positives
- The ability to classify data on lineage attributes alone, covering sensitive data with no consistent content pattern or no text content whatsoever
- Policy enforcement that applies to data movement regardless of whether a file is involved in the transfer

How to leverage this functionality:

File-level policies are a starting point, not a complete strategy. Extend coverage to the data itself:

- Classify data by combining content detections with lineage attributes, including where the data originated, how it moved throughout the organization, and who edited or handled it along the way
- Create policies that enforce classification based on data origin or applied labels even after content has been copied out of its source file or application
- Define policies that block or alert on pasting of sensitive content into messaging applications such as Slack or Teams, or into AI tools, where the destination carries its own governance implications
- Treat copy/paste as a first-class data movement event in your policy model. It is one of the most common and least controlled exfiltration vectors in everyday workflows

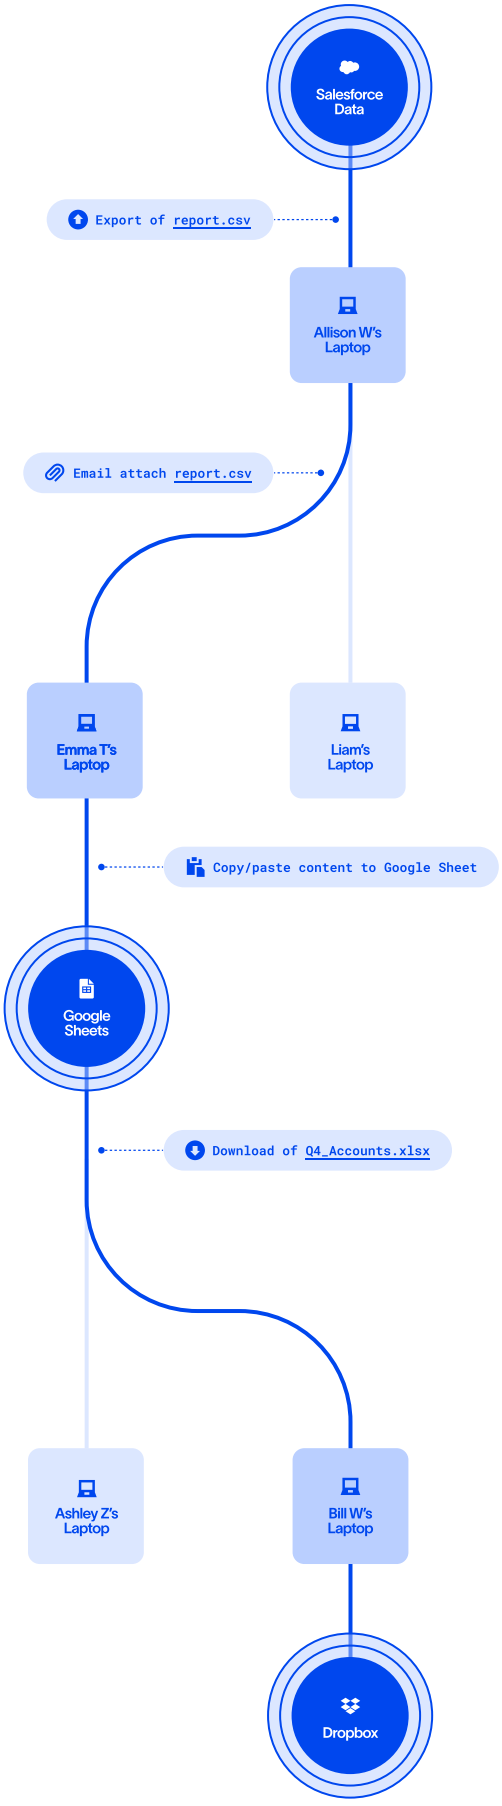

CRITERIA 3

Rapidly Calibrate Policies by Testing on Historical Events

One of the most persistent criticisms of DLP programs is how long they take to deliver value. Part of that is inherent to the problem DLP solves for: Before you can write effective policies, you need to understand what data your organization has, where it lives, and how it moves. But a significant portion of the delay has historically come from the tools themselves, specifically from the way legacy DLP applies and tests policy changes.

With older products, every policy update is a forward-looking exercise. You make a change, deploy it, and then wait for new events to flow in before you can evaluate whether the policy is working as intended. If it generates too many false positives, or misses the risky behavior it was designed to catch, you won't know for days or weeks. Then you adjust and wait again. For organizations trying to stand up meaningful coverage quickly, this feedback loop is one of the most significant barriers to a functional DLP program.

Modern DLP eliminates that delay by maintaining a complete record of historical user actions and applying policy changes against that record in real time. When you update a policy, you can immediately see what alerts it would have generated, which events it would have blocked, and where it would have created unnecessary friction, all before the policy goes live in production. What previously took weeks of observation can be validated in minutes.

That same historical record also unlocks a capability that goes beyond policy testing: genuine visibility into how data actually moves through your organization. Rather than relying on assumptions about where sensitive data lives and how employees use it, security teams can explore actual data flows, identify patterns they weren't aware of, and build policies around real behavior rather than hypothetical risk. That visibility also creates a foundation for more productive conversations with business stakeholders about acceptable use, because the discussion is grounded in evidence rather than conjecture.

Features to look for:

- A DLP platform that stores a complete record of events for every piece of data across the organization
- The ability to preview policy results against historical events before deploying changes to production
- Visibility into data flows across the organization, including who accesses data, how it moves, and where it's stored

How to leverage this functionality:

Before writing policies, use historical data visibility to build a clear picture of how your organization actually uses sensitive information. The right questions to explore:

- Who is accessing sensitive data, and for what purpose?
- Where does data flow as part of normal business activity versus where does it go that seems anomalous?
- What does acceptable use look like for different data types and user groups?

Once you have that understanding, test new policies against the historical record to preview what alerts or preventive actions would have been triggered, and iterate on configuration before any policy touches a live environment.

DSPM + DLP: Better Together

Most DLP programs take weeks or months to deliver meaningful coverage. The reason is rarely the tool itself. It's the gap between knowing what policies you need and actually knowing where your sensitive data lives and how it moves.

Data security posture management (DSPM) closes that gap. By discovering sensitive data at rest across cloud infrastructure, SaaS applications, endpoints, and on-premises environments, DSPM gives security teams the foundational visibility that DLP policy development requires. Without it, policy writing is guesswork. With it, you're building coverage around an accurate picture of your actual data landscape.

When DSPM and DLP operate on the same platform, the benefit compounds. Discovery feeds directly into enforcement, with no translation layer between the posture insights that reveal risk and the policy engine that acts on it.

CRITERIA 4

Combine Data Classification with Behavior Analysis to Better Detect Insider Threats

For years, DLP and insider risk management (IRM) have approached the same problem from opposite ends. DLP tools focused on classifying data and controlling its movement, but had no way to understand the person behind the action. Insider risk management tools analyzed behavior, flagging unusual upload volumes or off-hours access, but had no insight into what data was actually being handled.

The result was a coverage gap that neither category could fill on its own.

That gap has real consequences. A user uploading a large volume of files might be backing up a project or walking out the door with your source code. The action looks the same to a behavior-only tool. Conversely, a DLP alert on a file containing sensitive data tells you nothing about whether the user is a new employee following an unfamiliar workflow or someone who resigned last week and has been quietly staging data for weeks.

The most significant insider risk scenarios, particularly those involving departing employees, rarely manifest as a single obvious event. They unfold over time and across systems: downloading files from SaaS applications, syncing data to personal cloud storage, copying to removable media, interacting with job platforms. Each action evaluated in isolation can appear entirely normal. Evaluated together, they tell a different story. But connecting those signals requires a platform that understands both what data is moving and who is moving it.

This is why DLP and insider risk management are converging. Merging signals from separate tools in a SIEM improves alert quality, but it doesn't enable action. To intervene when data is actually at risk, the detection logic, the policy engine, and the response mechanism need to operate from the same unified understanding of data and behavior. That means risk scores that account for the sensitivity of the data being handled, not just the volume of actions taken. It means policies that can respond differently to a flagged user than to a first-time offender. And it means visibility that can correlate events happening weeks apart into a coherent picture of intent.

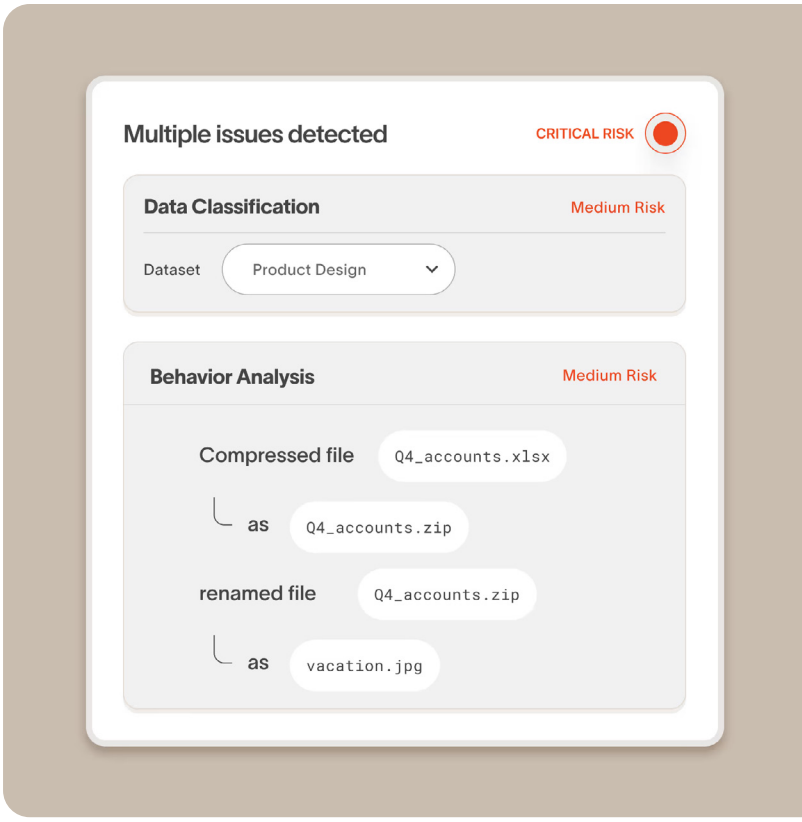

"By enriching DLP events with the context around the user's behavior, it will be far easier to distinguish between malicious and negligent acts and apply controls." — Gartner, Inc*

AI plays an increasingly important role here. The most advanced platforms apply AI-driven analysis to continuously surface meaningful patterns across billions of data events, connecting early signals, such as job search activity or unusual access patterns, to later actions in a way that static rules cannot. That kind of intelligence is what separates identifying a risk in progress from discovering a breach after the fact.

Features to look for:

- The ability to combine data classification and user behavior signals within a single policy and policy engine
- User risk scores that factor in the sensitivity of data being handled, not just behavioral volume or frequency
- A user-centric view of risk over time, not just per-incident alerting
- AI-driven analysis that can correlate signals across days or weeks to identify threats that unfold slowly

How to leverage this functionality:

Insider risk is most dangerous when it goes undetected long enough for data to leave the organization. Use combined data and behavior signals to get ahead of it:

- Review user risk scores regularly to identify elevated or trending risk before an incident occurs
- Place high-risk users, including those flagged by departure signals or performance activity, on watchlists with stepped-up policy enforcement
- Investigate users whose behavior shows a pattern over time rather than waiting for a single high-confidence alert
- Use AI-generated summaries and data lineage context to move quickly from detection to decision without manual log correlation

* Gartner Market Guide For DLP

CRITERIA 5

Distinguish Between Personal and Corporate Accounts Across All Applications

As cloud adoption expanded, most organizations addressed the risk of unsanctioned applications through a combination of CASBs, web proxies, and firewall rules. Blocking access to unknown or unapproved applications is a tractable problem. The harder problem, and the one that has grown significantly more complex, is employees using personal accounts of applications the organization has explicitly sanctioned.

The challenge is no longer limited to cloud storage. Today it spans every major category of business application. An employee working in corporate Google Drive or OneDrive should not be able to move that data into a personal account for the same service. An engineer committing code to a corporate GitHub repository should not be able to push that same code to a personal one. A user working in a corporate genAI environment operates under a completely different set of data governance, retention, and training policies than someone using a personal account. And one of the most consistent exfiltration paths for departing employees remains forwarding corporate email to a personal Gmail or Outlook account before their access is revoked.

In each of these cases, the application is approved. The data flow looks routine at the network level. What changes the risk profile is the account the user is authenticated into, and whether that account is under organizational control.

DLP solutions need to resolve this distinction at the account level, not the application level. That means detecting whether a user is logged into a corporate or personal instance of a given service, and enforcing policies that treat those two states differently. Without that capability, sanctioned application lists provide a false sense of coverage, because the same tool that is approved for corporate use becomes an exfiltration channel the moment an employee switches to a personal account.

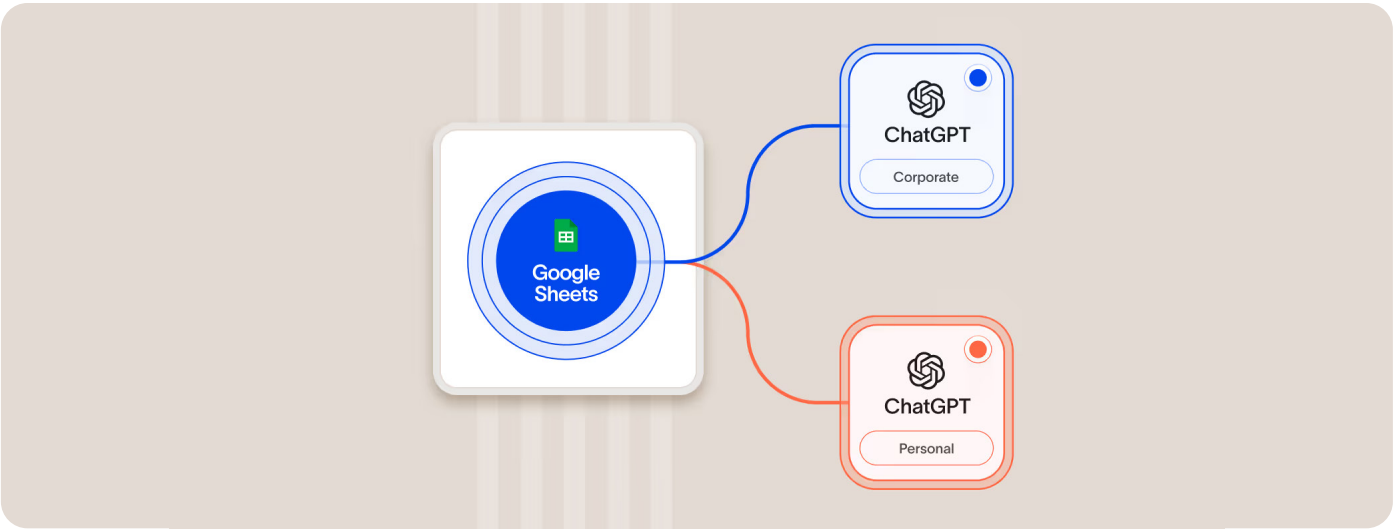

Features to look for:

- The ability to distinguish between corporate and personal account instances across cloud storage, email, collaboration tools, AI platforms, and code repositories
- Policy enforcement that applies different rules based on account type for the same application
- Coverage across the full range of account-boundary risk, including AI tools such as ChatGPT and Copilot, where the governance implications of personal vs. corporate accounts are significant

How to leverage this functionality:

The account boundary is only enforceable if you know where it applies. Start by mapping the risk surface:

- Compile a list of sanctioned corporate applications that also have consumer versions, including cloud storage, email, collaboration tools, code repositories, and AI platforms
- Define policies that restrict movement of corporate data to personal accounts of those same services
- Pay particular attention to AI tools, where the distinction between a managed enterprise instance and a personal account carries significant implications for data governance and compliance
- Apply heightened scrutiny to departing employees, for whom personal account forwarding across email and cloud storage is one of the most common pre-departure data collection patterns

What is Shadow AI?

Shadow AI refers to the use of AI tools and services by employees outside the visibility or control of IT and security teams. This includes personal accounts of sanctioned AI platforms as well as entirely unapproved AI tools employees adopt independently. When sensitive data enters a shadow AI environment, it may be processed, retained, or used to train models under terms the organization never agreed to. DLP must be able to detect when data flows into AI tools, distinguish between managed and unmanaged AI instances, and enforce policies that follow the data regardless of which AI surface it reaches.

CRITERIA 6

An Endpoint-First Platform Built to Cover Every Environment

The most important architectural question in a modern DLP evaluation is not how the software is delivered. It's where the platform starts.

The answer should be the endpoint. The endpoint is where data risk is highest. It is where files are created, copied, renamed, compressed, and moved to external destinations. It is where employees interact with generative AI tools through browser windows, where sensitive content is pasted into prompts, where agentic AI operates on local data with limited oversight, and where departing employees copy files to personal USB drives minutes before their last day. Network controls and cloud-based visibility cannot see any of this. They observe traffic at the perimeter or inventory data at rest in cloud repositories. But the actions that put data at risk happen on the device, and without an agent built to capture those actions with full fidelity, critical context is lost before an investigation even begins.

The foundational questions that matter most in a real investigation cannot be answered from the cloud alone: where did this data originate, was it modified, was it combined with other sensitive information, and was this action actually risky or routine behavior taken out of context? Answering those questions requires visibility into what happened on the device, at the moment it happened.

That visibility breaks down the moment data crosses an environment boundary, and that's exactly where most DLP architectures fail. CASB solutions provide visibility into cloud application activity but lose the thread the instant data is downloaded to a local device. Standalone endpoint DLP picks up where cloud tools leave off, but treats downloaded data as though it appeared from nowhere, stripped of the context that defined its sensitivity at the source. Consider a data scientist pulling an export from Snowflake to a local machine for analysis. To a cloud-focused tool, that data disappears at the point of download. To an endpoint tool, that file arrives without any inherited context from Snowflake. Neither tool can answer the question that actually matters: is this a sanctioned workflow, and is this data being handled appropriately given where it came from?

Data lineage is what makes that question answerable. By tracking data from its origin through every transformation, copy, and transfer it undergoes, a lineage-aware platform maintains a continuous record that spans cloud and on-premises environments, endpoints, browsers, and SaaS applications. The sensitivity context established at the source travels with the data as it moves. A file exported from a customer database in Snowflake carries that classification whether it's sitting on a local drive, attached to an email, or uploaded to a cloud storage service.

Endpoint coverage is the foundation, not the ceiling. A modern DLP platform needs to extend that same depth of visibility across every environment where data lives and moves: cloud infrastructure, SaaS applications, browsers, collaboration tools, email, and on-premises systems. The goal is not endpoint DLP, cloud DLP, or browser DLP in isolation. It is a unified platform built on an endpoint foundation that extends coherent, lineage-aware visibility across the entire environment without losing context at the boundaries that matter most.

Features to look for:

- A purpose-built endpoint agent with the performance and stability to capture data events at the device level without degrading user experience
- Endpoint visibility that covers all relevant channels: Local file operations, browser activity, copy/paste, removable media, AI tool interactions, and agentic AI behavior
- The ability to follow data as it moves between cloud platforms, on-premises systems, endpoints, browsers, and SaaS applications without losing lineage context
- Tracking of derivatives and exports, so that a copy or transformed version of sensitive data inherits the classification of its source
- Consistent policy enforcement across all environments, regardless of where data currently resides or how it arrived there
- Cloud-based management and policy infrastructure, so security teams benefit from modern delivery without sacrificing endpoint depth

How to leverage this functionality:

A platform built endpoint-first enables visibility that cloud-only tools cannot provide. Put that foundation to work:

- Ensure your endpoint agent is deployed across all managed devices, including both Windows and macOS, before evaluating coverage gaps elsewhere
- Use endpoint visibility to close the AI tool blind spot: Verify that your platform can observe data flowing into generative AI prompts and agentic AI actions at the device level, not just at the network edge
- Validate that lineage context is preserved as data moves from endpoint to cloud and back, so that a file exported from a cloud system and opened locally carries its origin and classification with it
- Identify access mismatches: Users who lack permission to access data in one system should not have uncontrolled copies of that data on a local device or in a personal cloud account
- Configure forensic evidence collection, including screen recordings and file captures, to be stored in your own cloud infrastructure rather than the vendor's, maintaining control over sensitive evidentiary data

CRITERIA 7

Real-Time User Coaching While Allowing Employees to Override

The majority of data exposure incidents are not the result of malicious intent. Most happen because an employee is trying to get something done, doesn't know the approved way to do it, and reaches for whatever tool is in front of them. An employee uploading files to a free online PDF converter isn't staging an exfiltration. They have a task, a deadline, and no obvious alternative. The risk is real, but the response needs to reflect the actual situation.

This is where traditional DLP approaches create as much friction as they prevent. A hard block with no explanation leaves the employee unable to complete their work and with no understanding of what they did wrong or what they should do instead. Annual security training doesn't bridge that gap. Instruction delivered months before or after the moment of risk rarely changes behavior in the moment that matters.

Real-time user coaching addresses this directly. When a risky action triggers a policy, a just-in-time notification delivered on the device, at the moment the action occurs, does something that no training session can: it connects the policy to the specific behavior in the specific context where it happens. That immediacy is what makes the lesson stick. Research consistently shows that organizations using real-time notifications reduce incident volume significantly over time, not because the tool is blocking more, but because employees are learning what acceptable behavior looks like through direct, contextual feedback.

The notification itself should be calibrated to the risk. Not every policy trigger warrants a hard block. A tiered response model, where low-severity events surface an informational notification, medium-severity events block the action but allow the employee to provide a business justification and proceed, and high-severity events enforce a hard block without override, gives security teams proportionate control without treating every user as a threat. It also creates a legitimate channel for exceptions, so that an employee asked by a client to upload documents to an unsanctioned platform has a way to document the business reason rather than finding a workaround.

The override and justification mechanism serves a secondary purpose that is easy to overlook. When employees consistently request exceptions for the same tool or workflow, that pattern is a signal. It often indicates a gap in the organization's approved technology stack, an opportunity to evaluate whether a tool employees are reaching for independently should be formally sanctioned and supported.

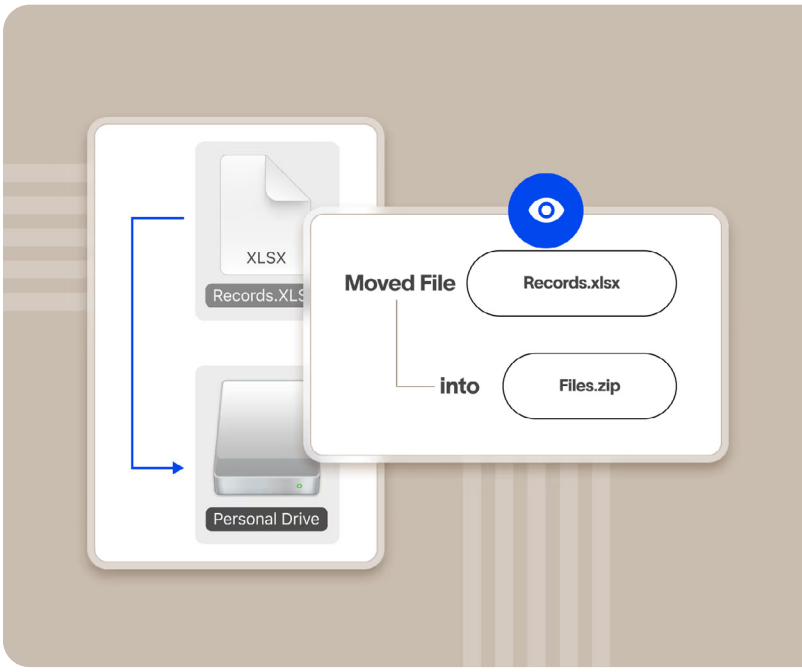

Features to look for:

- Just-in-time on-device notifications that fire across all exfiltration channels, not only within specific applications or browsers
- The ability to configure notifications with or without blocking, so that coaching can occur regardless of whether the action is prevented
- User override with business justification, allowing employees to document legitimate exceptions without circumventing the policy framework
- Tiered response options that can be calibrated by risk level and data sensitivity

How to leverage this functionality:

Match the response to the risk, and use the data your policies generate to improve both security and employee experience over time:

- Configure low-severity events to display an informational notification without blocking, coaching the employee without interrupting their workflow
- Configure medium-severity events to block the action but offer the employee an override option requiring a business justification
- Configure high-severity events as hard blocks with no override, reserved for actions involving your most sensitive data or highest-risk users
- Track changes in user behavior over time: real-time coaching measurably reduces incident rates, and that reduction is reportable progress
- Review override justifications regularly to identify tools and workflows that employees need but the organization hasn't provisioned, and use that data to close technology gaps proactively

CRITERIA 8

Full Picture of Incidents to Accelerate Investigations

An alert without context is not an investigation. It's a starting point that forces an analyst to go hunting, pulling logs from multiple systems, reconstructing a timeline manually, and ultimately spending more time on the mechanics of the investigation than on understanding what actually happened. Legacy DLP tools surfaced alerts in exactly this form: A user did something, here is the event, good luck.

The cost of that approach compounds quickly. Incident response time is one of the most significant drivers of breach impact. The longer it takes to determine whether an alert represents a genuine threat, a careless mistake, or a sanctioned exception, the more exposure accumulates. And in cases involving sophisticated insiders, delay is not just inefficient. It is the window in which data exits the organization entirely.

Modern DLP should eliminate that reconstruction work by assembling the full context of an incident automatically. That means showing analysts not just the triggering event, but everything that happened to the data before and after it. An employee who attempted to exfiltrate sensitive data through one channel, was blocked, then tried again through a different channel, then compressed the data into a ZIP file and renamed it before a third attempt is demonstrating a clear pattern of intent. That pattern is only visible if the system connects those events into a single coherent view. Presented as three separate alerts, each looks ambiguous. Presented as a sequence, the picture is unambiguous.

The same context that reveals intent also reveals access. If a user is attempting to exfiltrate data they were never supposed to have, the incident view should be able to answer how they obtained it. Perhaps a colleague uploaded a Workday export to a shared Google Drive folder with permissions that were too broad. Perhaps the data passed through a third party who shouldn't have had access. Data lineage traces that path from origin to incident, surfacing both the immediate risk and the underlying access control failure that enabled it.

Investigation context also needs to extend to collusion. Insider threats don't always operate alone. When the same sensitive data is being passed between employees in a pattern that doesn't reflect normal workflow, and one of those employees subsequently attempts exfiltration, that connection should be visible without requiring a manual cross-reference across multiple tools. The ability to identify that pattern automatically, and to surface it as part of the incident view, is what separates a DLP platform that accelerates investigations from one that merely logs events.

> "By 2027, organizations incorporating intent detection and real-time remediation capabilities into DLP programs will realize a one-third reduction in insider risks." — Gartner, Inc*

* Gartner Market Guide For DLP

Features to look for:

- An incident view that automatically surfaces the full history of events for a piece of data before, during, and after the triggering alert, without requiring manual log correlation
- The ability to trace data all the way back to its origin, revealing how a user obtained it and whether access was legitimate
- Visibility into obfuscation attempts: file extension changes, compression, renaming, and repeated exfiltration attempts across different channels
- Pattern detection for data transfers between users that may indicate collusion
- Remote forensic capabilities, including screen recordings and forensic file capture, stored securely in the cloud without requiring physical access to a device

How to leverage this functionality:

An alert should be the beginning of an answer, not the beginning of a search. Use full incident context to work faster and reach better conclusions:

- Review events before and after the triggering alert to assess intent: look for obfuscation attempts, repeated exfiltration through different channels, and file manipulation that suggests deliberate concealment
- If the user shouldn't have had access to the data in question, trace how they obtained it: overly broad sharing permissions and misconfigured cloud storage are common culprits
- When data was passed to the user by a colleague before the incident, examine the transfer pattern for signs of coordination between users
- Use screen recordings and forensic file captures to build evidentiary context for cases that may require HR, legal, or law enforcement involvement
- Generate records of all data an employee accessed or copied before offboarding, to support severance agreements and competitive departure investigations

Section 03

DLP Evaluation Checklist

BUSINESS CONSIDERATIONS

Strategic Fit and Risk Alignment

☐ **Beyond Compliance**	The solution addresses not just regulatory requirements (PCI DSS, HIPAA, GDPR) but also unregulated high-value data like source code, financial models, product designs, and trade secrets.
☐ **Intellectual Property Protection**	The platform can classify and protect IP that contains no recognizable content pattern, including unstructured documents, design files, and proprietary formats.
☐ **Business Continuity**	The solution supports continuous data monitoring so that incidents are contained before they escalate into operational disruptions or brand damage.
☐ **Trust and Reputation**	The vendor demonstrates practices that support stakeholder confidence, including transparent data handling, strong access controls, and a documented security posture.

Deployment and Time to Value

☐ **Deployment Speed**	The solution does not require months of tuning before delivering coverage. Historical data testing and DSPM integration can compress time to value significantly.
☐ **No On-Prem Infrastructure**	Management and policy infrastructure are delivered from the cloud, eliminating the need to manage application servers or databases on-premises.
☐ **Out-of-the-Box Coverage**	The platform includes pre-built policy templates for common use cases and compliance frameworks so teams can start from working coverage, not a blank slate.
☐ **Vendor Viability**	The vendor has a credible roadmap, a stable customer base, and demonstrated investment in AI-driven capabilities. DLP is core to their product, not a bundled add-on.
☐ **Onboarding**	The vendor must provide expert-led onboarding that deploys and integrates the platform and delivers meaningful coverage within weeks

Organizational and HR Risk

☐ **Departing Employee Risk**	The platform provides specific capabilities for monitoring and responding to pre-departure data staging activity, including personal account forwarding, bulk downloads, and removable media transfers.
☐ **HR and Legal Readiness**	Incident records, screen captures, and forensic file collections are stored in the organization's own cloud infrastructure to support HR, legal, or law enforcement involvement.
☐ **Security Culture Support**	The solution includes real-time user coaching that educates employees at the moment of a risky action, reducing incident volume over time rather than relying on blocking alone.

OPERATIONAL CONSIDERATIONS

Policy Management and Maintenance

☐ **Single Policy Engine**	All exfiltration channels are governed by one policy engine and managed through one interface. No channel requires a separate tool, console, or policy framework.
☐ **Historical Policy Testing**	Policy changes can be previewed against a complete record of historical events before going live. No need to deploy and wait weeks to evaluate whether a policy is working.
☐ **Policy Simplicity**	The policy editor supports Boolean logic, auto-complete, and graph-to-policy conversion. Using data lineage, the solution enables simple, high-confidence policies that outperform complex content-only rules.
☐ **Exception Management**	Users can submit business justifications when overriding a policy. Override patterns are surfaced to administrators to identify gaps in sanctioned tooling.

Alert Quality and Analyst Experience

☐ **Low False Positive Rate**	The solution combines content analysis with data lineage dramatically reducing false positives from common patterns like phone numbers, emails, and generic numeric strings.
☐ **Alert Context**	Each incident surfaces the full history of events before and after the triggering action, without requiring manual log correlation across separate tools.
☐ **Obfuscation Detection**	The platform detects file renaming, extension changes, compression, and encryption attempts, and flags repeated exfiltration attempts across different channels as a connected pattern.
☐ **Collusion Visibility**	The platform can identify when the same sensitive data is passed between multiple users in patterns that suggest coordination, and can surface that connection as part of the incident view.

User Impact and Adoption

☐ **Endpoint Performance**	The endpoint agent is purpose-built to capture data events at the device level without degrading system performance or breaking cloud applications.
☐ **Tiered Response Model**	The platform supports informational notifications to security, educational or warning-level pop-up to an end user, and hard block responses calibrated to data sensitivity and user risk level.
☐ **Employee Education**	The platform supports Just-in-time coaching delivered at the moment of a risky action, which has been shown to reduce incident volume by up to 80% over time. Coaching operates across all channels, not just within specific apps.
☐ **Minimal Workflow Disruption**	Prevention features are actually enabled in production. High false positive rates from legacy tools mean prevention is often turned off. Accuracy here is an operational requirement, not a feature.

Investigation and Response

☐ **Incident View Completeness**	Analysts can view the full data lifecycle from origin through the triggering event without needing to pivot to a SIEM or external log system.
☐ **Access Path Tracing**	If a user is exfiltrating data they were not authorized to access, the platform can trace how they obtained it, including overly broad sharing permissions or misconfigured cloud storage.
☐ **Remote Forensics**	Screen recordings and forensic file captures are available and stored in the organization's own cloud, not the vendor's, for evidentiary integrity.
☐ **SIEM and API Integration**	Incidents are exportable to SIEM tools and accessible via API so DLP data can feed existing security workflows without requiring tool replacement.

TECHNICAL CONSIDERATIONS

Data Classification

☐ **Data Lineage as Core Capability**	Lineage is embedded in the classification engine, not bolted on. It tracks data from point of origin through every copy, transformation, and transfer across the organization.
☐ **Content + Lineage Combined**	Policies can combine content detection (regex, EDM, OCR, NLP) with lineage attributes (origin, movement history, who handled it) for higher accuracy and fewer false positives.
☐ **Lineage-Only Classification**	Sensitive data with no recognizable content pattern (strategy documents, financial models, recorded meetings, design files) can be classified on lineage attributes alone.
☐ **File Type Breadth**	Classification and policy enforcement covers the full range of relevant file types, including CAD files, design files, source code, video, and proprietary formats, not only office documents.
☐ **Persistence Across Transformations**	Sensitivity context survives copy/paste actions, format changes, application transitions, compression, encryption, and movement between file systems.

Coverage and Architecture

☐ **Endpoint-First Architecture**	The platform is built on a purpose-built endpoint agent that captures file operations, browser activity, copy/paste, removable media, AI tool interactions, and agentic AI behavior at the device level.
☐ **Unified Environment Coverage**	The same platform covers cloud infrastructure, SaaS applications, endpoints, browsers, email, collaboration tools, and on-premises systems without losing lineage context at environment boundaries.
☐ **Cross-Environment Lineage**	The solution must preserve data classification and end-to-end handling history as data moves across environments and systems, including when it is exported from cloud services to local devices.
☐ **Derivative and Export Tracking**	Copies, exports, and transformed versions of sensitive data must automatically inherit the classification and policy treatment of their source, with lineage preserved even when the data changes format or location.

Channel and Application Coverage

☐ **All Exfiltration Channels**	Policy enforcement covers personal and unsanctioned SaaS apps, copy/paste, USB and removable media, encrypted messaging (Signal, WhatsApp), generative AI tools, and AI agents.
☐ **Personal vs. Corporate Account Distinction**	The platform can distinguish between a corporate and personal instance of the same application (Google Drive, GitHub, AI tools, email) and enforce different policies for each.
☐ **AI Tool Coverage**	The platform monitors data entering and leaving generative AI tools and AI agents at the device level, including data pasted into prompts, and distinguishes between managed enterprise instances and personal accounts.
☐ **Copy/Paste as First-Class Event**	Copy/paste is treated as a data movement event in the policy model. Sensitive content copied out of a source application retains its classification in the destination.

Insider Risk and Behavioral Detection

☐ **Unified DLP + IRM**	Data classification and user behavior signals are combined within a single policy engine. Risk scores account for the sensitivity of data being handled, not just behavioral volume or frequency.
☐ **User Risk Scoring**	The platform maintains a user-centric view of risk over time, correlating signals across days or weeks to identify threats that unfold slowly, not just per-incident alerting.
☐ **AI-Driven Pattern Detection**	AI analysis continuously surfaces meaningful patterns across billions of data events, connecting early signals (job search activity, unusual access) to later actions in a way that static rules cannot.
☐ **Watchlist and Stepped-Up Enforcement**	High-risk users can be placed on watchlists with elevated policy enforcement. Policies can respond differently to a flagged user than to a first-time offender.

DSPM Integration

☐ **Discovery Feeds Enforcement**	DSPM capabilities allow the platform to discover and classify sensitive data at rest across cloud, SaaS, endpoints, and on-premises environments, providing the data map that DLP policy development requires.
☐ **Single Platform, Single Source of Truth**	DSPM and DLP operate on the same platform with the same lineage engine. Discovery context is available to the enforcement layer without a translation layer between tools.
☐ **Posture to Policy**	DSPM findings can be used to identify coverage gaps, prioritize policy development, and validate that policy coverage matches actual data risk, compressing time to meaningful DLP coverage.

Section 04

The Cyberhaven Difference

What makes AI-native DLP fundamentally different

The criteria above describes what a modern DLP program requires. Each criteria is demanding by design, because the problem is demanding. Most legacy DLP tools fail on several of them structurally, not because of implementation decisions, but because of architectural ones made a decade ago for an environment that no longer exists.

Cyberhaven was built to meet them. Here is where the architecture differs.

01 **Data lineage as the classification engine, not a bolt-on.** Legacy DLP classifies data by scanning what a file contains. Cyberhaven tracks where data came from, how it moved, who handled it, and what systems it touched. That lineage context is embedded in the policy engine and available as a classification condition in its own right. It covers sensitive data that contains no recognizable pattern and data that contains no text at all.

02 **One product, one policy engine, every exfiltration channel.** Cloud, endpoint, browser, email, SaaS, encrypted messaging, generative AI tools, AI agents, and removable media are all governed by the same policy engine and managed through one interface. No siloed consoles, no inconsistent coverage, no channel that falls between tools.

03 **Endpoint-first architecture with unified cross-environment visibility.** The platform is built on a purpose-built endpoint agent that captures every meaningful data action at the device level. That depth of visibility extends across cloud, SaaS, and on-premises environments without losing lineage context at environment boundaries. Sensitivity travels with the data, not with the tool that last touched it.

04 **DLP and insider risk management from one platform.** Data classification and user behavior signals are combined within a single policy engine. Risk scores account for what data is being handled, not just how much activity a user generates. AI-driven analysis connects signals across days or weeks to surface slow-moving threats that static rules miss entirely.

05 **Policy calibration against historical data before going live.** Cyberhaven maintains a complete record of every data event across the organization. When you update a policy, you can preview what it would have triggered against real historical activity before it touches a live environment. What previously required weeks of observation can be validated in minutes.

06 **DSPM and DLP on one platform, from discovery to enforcement.** Cyberhaven's DSPM capabilities discover and classify data at rest across cloud, SaaS, endpoints, and on-premises infrastructure. That visibility feeds directly into the DLP enforcement layer. There is no translation layer between posture insights and policy action. Discovery and enforcement are part of one continuous workflow.

DLP has been deployed as a compliance checkbox for long enough that most security teams have learned to expect it to underperform. Real protection requires a platform that understands where data comes from, not just what it contains. One that follows data across every environment, every account type, and every exfiltration channel. One that connects the behavior of the person moving the data to the sensitivity of the data itself.

That is what Cyberhaven is built to do.

Request a demo