GUIDE

5 Steps to Secure the Data Behind Every AI Initiative



cyera

A practical guide to building the data foundation AI can trust

Every assistant, copilot, workflow, and agent depends on the data already inside the enterprise.

That changes the consequence of weak data security.

Before AI, an overexposed file might sit untouched for months. A user had to know it existed, know where to find it, open it, understand it, and decide what to do next. AI compresses those steps into a prompt. Agents go further, repeating the same sequence across systems and actions without waiting for a person to drive each step.

AI does not create your data security gaps. It operationalizes them.

Sensitive data that is unknown, mislabeled, overexposed, or reachable by the wrong identities becomes easier to find, easier to use, and easier to move. The question is no longer simply whether sensitive data exists somewhere it should not. The question is whether a user, assistant, workflow, or agent can turn that exposure into action.

Before scaling AI, enterprises need a data foundation they can trust.

This guide breaks that foundation into five steps:

Discover

what AI could reach

Classify

what actually matters

Govern

what AI can inherit

Protect

data when it is used or moved

Validate

that the foundation still holds as AI expands

AI Changes the Consequence of Data Risk

Security teams have been trying to answer the same questions for years: Where does sensitive data live? Who can access it? How is it used? Where does it move? Are the right controls working?

Those questions were already difficult in the cloud era. Data spread across SaaS, cloud infrastructure, databases, file shares, collaboration platforms, endpoints, and on-premises systems. Unstructured data grew faster than teams could label it. Ownership blurred. Permissions accumulated. DLP queues filled up. Compliance evidence was assembled after the fact.

Most organizations responded by adding tools. Each solved a real part of the problem, but the answers remained fragmented. One system knows where some sensitive data lives. Another knows who can access it. Another sees file movement. Another manages labels. Another tracks incidents.

AI turns that fragmentation into a control problem.

A user can ask broadly and let an assistant locate information that was technically accessible but practically hidden. An agent can inherit an identity's permissions and act on them repeatedly. A model can summarize sensitive content from multiple sources in seconds. A workflow can move data before anyone reviews the full context.

Every prompt, retrieval, file read, summary, upload, share, and agent action ultimately depends on one decision:

Can this identity use this data, in this way, for this purpose, right now?

If the data foundation is incomplete, that decision is made with too little information.

That is why securing AI starts before the model, prompt, copilot, or agent. It starts with the data environment they inherit.

The shift security teams have to make

Move from	Move to
Discovery projects that produce periodic snapshots	Discovery that starts fast and stays current
Classification based only on known patterns	Classification that understands business meaning
Permission reviews disconnected from sensitivity	Access decisions informed by data, identity, usage, and exposure
Noisy protection based on partial signals	Data-aware protection across use and movement
Evidence assembled after the fact	Current proof that controls are working

The objective is continuous data trust: a current, connected understanding of the data, identities, access paths, activity, movement, policies, and evidence behind every AI initiative.

This five-step roadmap is designed as a practical operating model. Apply it to one real AI initiative at a time, such as a Copilot rollout, an agentic workflow, or a single sensitive data domain. Each step builds on the last: first find the data AI could reach, then understand what matters, reduce unnecessary access, protect data when it moves, and continuously prove the foundation is working as AI expands.

Step	What it helps your team answer
Discover	What data could AI reach, and where does it live?
Classify	Which data is sensitive, proprietary, business-critical, or regulated?
Govern	Who and what can access that data, how is it being used, and what access should be reduced?

Step	What it helps your team answer
Protect	How is sensitive data moving or being used, and which activity needs action?
Validate	Can your team prove controls are working and keep them current as AI expands?

Step 1: Discover What AI Could Reach

Find the data before AI does

You can't secure data you can't see.

Enterprise environments are full of blind spots: abandoned folders, departmental SaaS applications, M&A systems, file shares, developer environments, stale copies, and repositories that were never brought into security scope.

If an AI assistant, workflow, or agent can reach a repository, the data inside it becomes part of the AI risk surface whether or not the security team has looked at it recently.

The problem is that traditional discovery often operates like a project. Teams deploy infrastructure, define scope, tune scans, wait for results, and repeat the process every time the environment changes. In a world where data is constantly being created, copied, moved, and connected to new AI workflows, point-in-time visibility ages quickly.

Discovery has to become a continuous capability.

A team should be able to understand:

- Where enterprise data lives across SaaS, cloud, databases, file shares, endpoints, and on-premises systems
- Which repositories are unknown or only partially covered
- Which data is duplicated, stale, abandoned, or shadow
- Which data stores are connected to AI initiatives
- Which users, systems, workflows, and agents may be able to reach them

The goal is not another inventory. It is a current map of the data AI could put into play.

Proof at enterprise scale

At Genuine Parts Company, Cyera scanned more than 82 million on-premises files in 40 days across an environment spanning roughly 40 petabytes of data.

The important point is not that every environment should match that timeline. It is that discovery at enterprise scale should no longer require months of setup before useful visibility appears.

Cyera in practice

Cyera provides agentless discovery across SaaS, IaaS, DBaaS, and on-premises environments, giving teams a living view of enterprise data without requiring them to deploy and maintain new scanning infrastructure for every environment.

Pressure-test one AI initiative

Can you answer:

- What data could this AI initiative reach?
- Which repositories are missing or only partially covered?
- Which data is duplicated, stale, or shadow?
- Who owns the highest-risk data in scope?

If those answers require multiple teams and several tools to assemble, you do not yet have a trusted foundation.

Step 2: Classify What Actually Matters

Your most important data may not match a pattern

Finding a file does not tell you what it means to the business.

Traditional classification works best when sensitive information has a recognizable structure: a credit card number, Social Security number, or other known pattern.

But much of the data enterprises care most about does not come with a convenient signature.

An M&A plan does not look like a credit card number. Neither does a product roadmap, pricing model, proprietary formula, customer strategy, executive communication, source code repository, or confidential research project.

AI makes that gap much more consequential because AI understands and retrieves information based on meaning.

A generic classifier might see an ordinary spreadsheet. The business may see next quarter's pricing strategy.

Classification therefore has to answer more than **"Does this data match a known pattern?"**

It has to answer **"What does this data mean to this organization?"**

A modern classification program should combine deterministic techniques for known structured data with AI-native approaches that can identify unstructured, proprietary, and business-specific information. Teams should also be able to describe new sensitive concepts in the language the business already uses rather than building endless rules around keywords and regex.

That lets security understand:

- What the data is
- Why it matters
- Who owns it
- Which policies apply
- Whether it is exposed
- Whether it is labeled correctly
- How it should be protected

In one documented deployment, **82% of the sensitive data classes Cyera identified had no equivalent in the customer's previous classification coverage.** The data did not lack importance. The existing tools simply had never been taught how to recognize it.

Cyera in practice

Cyera's Enriched Classification combines deterministic techniques and LLM validation for structured data with Learned Classification for information that has no predefined pattern. Classification Topics organize those findings around business concepts such as M&A Planning, IP Strategy, Customer Records, and Executive Communications.

Cyera reports **95%+ classification precision** across structured and unstructured data.

Pressure-test your classification

For one important data domain, ask:

- Which data is regulated?
- Which is proprietary or business-critical?
- Which is sensitive because of its meaning rather than its format?
- Which sensitive data is unlabeled or mislabeled?
- What would your existing classifiers never know to look for?

Finding data is visibility. Understanding it is intelligence.

Step 3: Govern What AI Can Inherit

AI inherits your access reality

Once you know what data exists and what matters, the next question is not simply:

Who has access?

It is:

Who and what should still have access, given what this data actually is?

That distinction matters because AI turns overly broad permission into a retrieval path.

If an employee can reach sensitive information they no longer need, an AI assistant may be able to retrieve it for them. If a service account retains broad permissions, an automated workflow can use them. If an agent is connected to too many systems, it can act beyond the scope anyone originally intended.

Access accumulates because organizations change faster than permissions do. Employees move roles. Contractors leave. Groups expand. Projects end. Service accounts survive long after their original purpose. Shared folders become permanent. New systems arrive through M&A.

Traditional access reviews tend to begin with identities and permissions. A data-first approach begins with the asset at risk, then connects sensitivity, identity, permissions, ownership, usage, and exposure.

That changes prioritization.

Instead of treating every entitlement equally, teams can focus on:

- Sensitive data with broad or public access
- Risky identities that can reach critical information
- Dormant or departing users with lingering permissions
- Service accounts with excessive reach
- External sharing paths
- AI-connected workflows with access beyond their purpose
- Permissions that exist but are rarely or never used

Permissions tell you what someone can access. Usage tells you what they do access. Data sensitivity tells you what that access could cost.

Together, those signals make least privilege far more actionable.

Reduce the data footprint too

Access is only half the problem.

Stale, duplicated, abandoned, and ownerless data is additional attack surface and additional AI retrieval surface. If no one needs it, there is little value in continuing to expose and govern it forever.

For every AI initiative, teams should ask not only what access can be removed, but what data can be archived, deleted, or otherwise taken out of reach.

The objective is simple:

Make the blast radius smaller before AI makes the environment faster.

Cyera in practice

Cyera connects data sensitivity with identity, permissions, ownership, access activity, and exposure. Access Trail helps teams understand who can reach sensitive data and how that access is actually being used. Remediation workflows help reduce unnecessary access and exposure, while discovery and classification identify stale, duplicate, and shadow data that can be prioritized for cleanup.

Pressure-test one critical dataset

Can you prove:

- It has an accountable owner
- Everyone with access still needs it
- Rarely used or excessive access has been reduced
- Stale and duplicate copies have a disposition plan
- The AI initiative cannot reach more than it requires

If governance ends with a spreadsheet of findings but nothing actually gets smaller, you have measured the problem without changing it.

Step 4: Protect Data in Use and in Motion

Legitimate access does not make every action safe

A user may legitimately need access to sensitive data.

That does not mean every way they use or move it is appropriate.

Viewing a customer file for work is different from sending it to a personal email account. Opening source code is different from pasting it into an unmanaged AI tool. Accessing a financial model is different from syncing it to an external drive or passing it into an automated workflow.

Protection has to understand the difference.

That requires more than a pattern match or label. A protection decision needs to know:

- What the data is
- Why it matters
- Who or what is using it
- Whether that identity should have access
- Where the data is going
- Whether the destination is approved
- Which policy applies
- Whether the behavior fits its business context

Without those signals, DLP is forced to make decisions with partial information.

AI makes noisy DLP more expensive

DLP programs already struggle with duplicate alerts, false positives, and fragmented policies. AI introduces even more movement paths: prompts, outputs, summaries, generated files, browser sessions, cloud shares, SaaS applications, endpoints, and agentic workflows.

More events do not create better security if analysts cannot tell which ones matter.

The goal is not maximum alert volume.

The goal is a queue the team can trust.

A useful protection decision should explain what moved, why it was sensitive, who moved it, where it went, whether the action made sense, and why the event was blocked, escalated, or deprioritized.

That same context should improve the system over time. Repetitive low-value detections should become easier to tune out. High-risk events buried in low-severity queues should become easier to surface.

In documented deployments, Cyera Omni DLP has driven up to 98% fewer false positives and 80% fewer alerts.

"We went from 28K false positives to 0 with Cyera." - Insurance customer, IDC Business Value Study sponsored by Cyera

Close the loop

Detection without action is just another queue.

Every meaningful finding needs an owner, a target response time, a documented action, and verification that the risk was actually addressed.

There is an important difference between **closing the ticket** and **closing the risk.**

A mature program can show what happened, who owned it, how it was resolved, whether the fix worked, and whether the root cause was addressed so the same issue does not simply return.

Cyera in practice

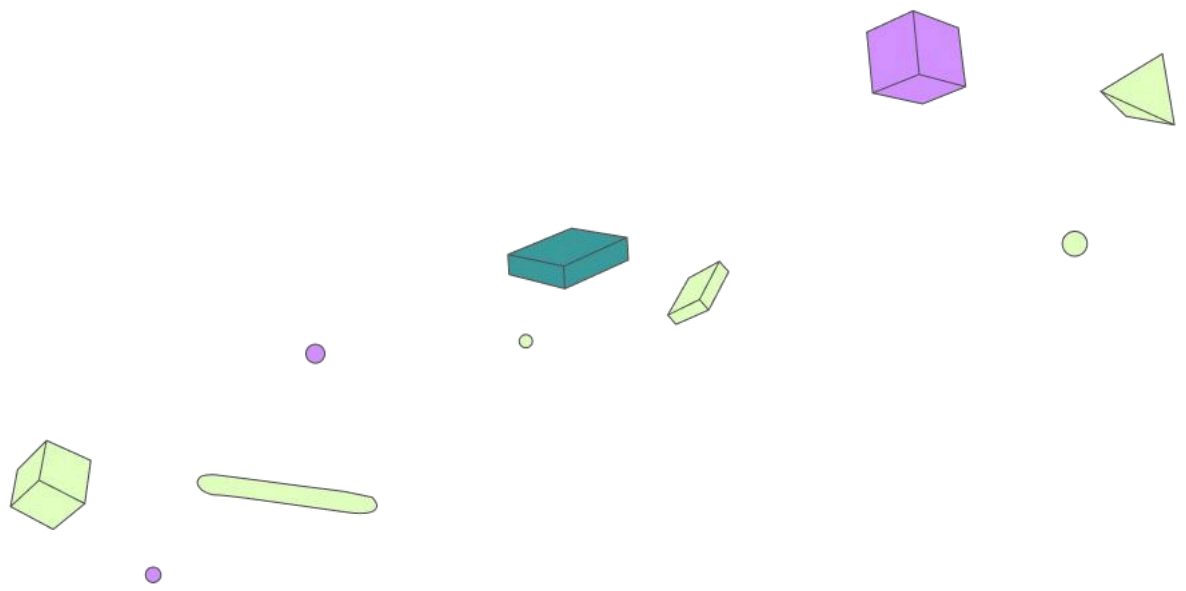
Cyera carries data intelligence from discovery and classification into protection and remediation. Omni DLP enriches DLP decisions with data, identity, usage, and destination context, helping teams reduce noise and prioritize meaningful activity. Findings can then move into workflows that revoke unnecessary access, remediate exposure, and support incident response when needed.

Pressure-test one event

Pick one recent protection alert and ask:

- What data moved, and why did it matter?
- Who or what moved it?
- Should that identity have been able to access it?
- Where did it go?
- Why was the event prioritized?
- Who owned the response?
- Was the risk actually closed?

If those answers live across separate tools and teams, protection still does not have the full story.



Step 5: Validate That the Foundation Still Holds

Trust is not a deployment state

A data foundation can be secure at launch and unreliable three months later.

New data appears. Files are copied. Labels drift. Employees change roles. Service accounts remain active. New AI tools connect. Agents gain capabilities. DLP policies need tuning. Incidents happen. Auditors request evidence. Leadership asks whether risk is actually going down.

That is why AI readiness cannot be a one-time assessment.

Validation means proving that the controls behind an AI initiative are working now, with evidence that stays current as the environment changes.

Many organizations can produce a point-in-time readiness report. Far fewer can recreate that proof a month later without manually rebuilding it.

The evidence is typically fragmented: a discovery report, an access-review spreadsheet, a DLP dashboard, a ticketing system, an audit folder.

A trusted foundation makes proof a by-product of the program itself.

A trusted foundation makes proof a by-product of the program itself.

At any point, security should be able to show:

- Discovery coverage is current
- Sensitive and proprietary data is being identified
- Unnecessary access and exposure are decreasing
- Protection controls are surfacing meaningful events
- Remediation is assigned, tracked, and verified
- Evidence is available for audits, incidents, and AI go/no-go decisions

The leadership conversation then changes from:

“We designed the controls correctly.” to: **“Here is the evidence that they are working.”**

The operating model matters too

Technology does not keep a program current by itself.

A common failure pattern is predictable: discovery creates a baseline, the obvious access issues get fixed, policies are tuned, evidence is collected, and then the environment continues moving.

Findings accumulate. Ownership blurs. Tuning becomes a backlog. Evidence quietly goes stale.

For lean security organizations, the answer cannot be to hire enough analysts to maintain every control manually.

Continuous posture, policy tuning, remediation, materiality assessment, and evidence all have to become sustainable operating motions.

Cyera's Data Security Approach at a Glance

Step	What teams need	Common gap	Cyera's approach	Why it matters for AI
Discover	A current view of what data AI could reach	Discovery requires infrastructure, manual scoping, and scans that take too long or go stale too quickly	Agentless discovery across SaaS, IaaS, DBaaS, and on-premises environments, with no scanning infrastructure to maintain	AI can retrieve data quickly once it has access. Blind spots become exposure paths.
Classify	Classification that automatically understands business context	Generic and manual classification only finds what it was told to find	Classification automatically identifies sensitive, proprietary, shadow, and business-specific data with no manual tuning	AI can expose data that is sensitive because of meaning, not just format.
Govern	Data-aware governance of inherited access	Access reviews are disconnected from sensitivity, usage, and exposure context	Data sensitivity is connected to identity, permissions, ownership, usage, and exposure	AI-connected tools can inherit overbroad access. Access risk becomes retrieval risk.
Protect	Context-rich protection across data in use and in motion	DLP alerts are noisy because they lack enough data, identity, and usage context	Shared data intelligence informs alert analysis, policy tuning, and protection decisions	AI increases data movement. Teams need better signal, not more alerts.
Validate	Continuous proof that controls are working	Evidence, policy tuning, materiality assessment, and response are often manual or reactive	Ongoing posture support, policy tuning, materiality assessments, and 24×7 response through managed services	AI adoption keeps changing. The foundation must stay current.

Cyera in practice

Cyera keeps discovery, classification, identity, access, protection, and remediation signals connected so teams do not have to repeatedly rebuild the story across separate systems. Cyera's managed services can also help maintain posture, tune policy, conduct materiality assessments, and support 24×7 breach response where internal capacity is limited.

Pressure-test your evidence

If leadership asked today whether an AI rollout is safe to expand, could you show:

- What sensitive data the initiative can reach
- Whether discovery and classification are current
- Whether unnecessary access has decreased
- Whether meaningful findings are being fixed
- Whether protection is becoming more effective
- Whether the evidence can be produced in minutes rather than days

If the answer depends on rebuilding the story by hand, validation is not continuous yet.

Why a Connected Data Security System Matters

The five steps are not independent controls.

Each depends on intelligence created by the others.

Discovery needs classification to understand which assets matter. Governance needs sensitivity to prioritize access. Protection needs data and identity context to distinguish legitimate use from risk. Validation needs all of those signals to prove whether the program is improving.

That is why fragmented tooling becomes harder to sustain as AI expands.

The issue is not simply the number of consoles. It is what gets lost in the handoffs between them.

A discovery tool can tell you a file exists. An IAM system can tell you someone has access. A DLP product can tell you the file moved. A ticketing platform can tell you an issue was closed.

None of those facts alone answer the decision AI forces security teams to make:

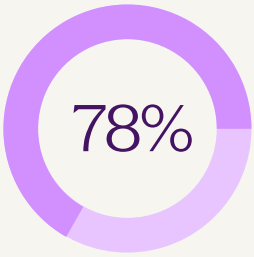
Can this identity use this data, in this way, for this purpose, right now?

Cyera's Data Security solution connects those facts through shared data intelligence. The same understanding of sensitive, proprietary, and shadow data can inform access, minimization, DLP, remediation, and evidence.

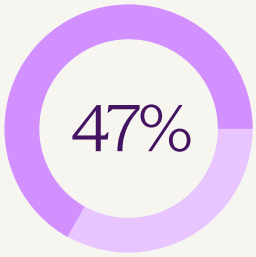
For practitioners, that means fewer handoffs and less manual correlation.

For security leaders, it means a defensible way to show how data risk is changing as AI adoption grows.

In an IDC Business Value Study sponsored by Cyera, customers reported:



more data threats identified and remediated



less time spent resolving alerts

Proof in Practice: Closing a 40-Petabyte Blind Spot

Genuine Parts Company operates more than 10,800 locations across 17 countries, with decades of data distributed across AWS, Azure, GCP, SaaS, and on-premises systems.

Before Cyera, the team could not confidently answer a basic question:

Where does our sensitive data live?

Legacy approaches forced tradeoffs between structured and unstructured data, cloud and on-premises systems, and one environment versus another.

During evaluation, a scan of just 2% of the environment surfaced roughly **$145 million in exposure,** about **500% more** than the team had previously estimated.

Cyera then scanned the company's global Office 365 tenant in six weeks and **82 million on-premises NetApp files in 40 days.**

Protection changed too.

With Cyera Omni DLP layered onto the company's existing Microsoft Purview deployment, Cyera analyzed **27,472 alerts** across four weeks and found that **1,043,** about 3%, required a person's attention.

Six months in:

532	98%	$45	12th of 13
billion sensitive records had been mapped	of high-risk cloud issues had been resolved	The team quantified $45 million in risk reduction	Data protection moved from a top-five risk to 12th of 13 on GPC's risk register

> "If you really want to be sincere and understand where your risk and exposure is, there's only one solution that will give you that, and that's Cyera." - Damian Apone, Director of Global Retail Cybersecurity, Genuine Parts Company

The GPC story shows the five steps working together: find the data at enterprise scale, understand what matters, reduce exposure, make protection more actionable, and produce evidence leadership can use.

The Next Layer: Secure the Agent

The five steps answer a data-security question:

Can we trust the environment AI is entering?

Once the answer is yes, the next question is different:

What should agents be allowed to see and do with it?

An agent's risk is bounded by what it can reach and what it can do. If sensitive data is unknown, overexposed, mislabeled, or moving without the right controls, agent security inherits those gaps.

Cyera Agent Guardian builds on the same platform context to help teams discover agents, understand which data and tools they can access, monitor their actions, and reduce risk when agents are overexposed, misconfigured, or behaving unexpectedly.

Data security helps you trust the foundation AI depends on. Agent security helps you trust the agents acting on top of it.

You need both.

The data foundation comes first.

The 5-Minute AI Data Foundation Check

Apply these questions to one real AI initiative, not the enterprise in the abstract.

1. Discover

Can we locate the data this AI initiative could reach?

☐ Yes ☐ Partially ☐ No

2. Classify

Can we identify the sensitive, proprietary, regulated, and business-specific data in scope?

☐ Yes ☐ Partially ☐ No

3. Govern

Can we see who and what has access, understand how that access is being used, and reduce what AI does not need?

☐ Yes ☐ Partially ☐ No

4. Protect

Can we identify and prioritize risky data use or movement with enough context to take action?

☐ Yes ☐ Partially ☐ No

5. Validate

Can we prove the foundation is working now and keep that evidence current as AI expands?

☐ Yes ☐ Partially ☐ No

If you answered "Partially" or "No"

Do not turn this into a generic maturity exercise.

Find the gap with the largest potential blast radius:

- Sensitive data you cannot see
- Business-critical data you cannot classify
- Overbroad access AI could inherit
- Unnecessary data that remains exposed
- Protection events your team cannot prioritize
- Evidence you cannot produce quickly

Start there.

The objective is not perfect data hygiene before AI can move.

It is knowing which risks are real, which ones matter, and whether the foundation is strong enough for the next step.

The Bottom Line

AI security starts before the model, prompt, copilot, or agent.

It starts with the data those systems inherit.

When sensitive data is unknown, misunderstood, overexposed, or reachable by identities that do not need it, AI turns passive weakness into active use.

When teams can discover what AI could reach, classify what matters, govern what AI can inherit, protect data when it moves, and continuously validate that the controls still work, security gets a defensible path to say yes to AI.

The standard is not perfect data hygiene.

The standard is trust.

Sources

IDC Perspective, Preparing Data for Safe AI Enablement, Jennifer Glenn and Ryan O'Leary, #US54310426, 2026.

IDC InfoBrief, Enabling Your AI Initiatives with Safe and Protected Data, #US54819626-IS, August 2026, sponsored by Cyera.

IDC Business Value White Paper, sponsored by Cyera, The Business Value of Cyera, #US54797226, August 2026.

Cyera whitepaper, Legacy Classification is Losing the AI Race. Here's How to Win It, 2026.

Cyera case study, Genuine Parts Company Closes a 40-Petabyte Data Blind Spot to Safely Enable AI With Cyera, 2026.

Cyera campaign messaging and approved customer results.

SACR, The Agentic Shift: Solving the Elusive Equation of Data Maturity, 2026.

cyera

See Cyera in action

Ready to build your AI roadmap
on a trusted data foundation?

Book a demo →