JOINT RESEARCH REPORT · CYERA RESEARCH & OSO · 2026

96% of Enterprise Permissions Go Unused.

AI Agents Will Inherit Them — and Turn Dormant Access Into a Security Crisis

96%
of permissions granted go unused

2.4M
workers analyzed

3.6B
permissions examined

cyera research | OSO

Executive Summary

The access surface hiding in plain sight.

A joint research effort by Cyera Research & Oso analyzing 2.4 million enterprise workers and 3.6 billion permissions has uncovered a structural vulnerability embedded in how modern organizations manage access — one that has operated silently for years, and is now on the verge of becoming critical.

The findings are unambiguous: 96% of permissions go unused. Sensitive data sits within reach of millions of workers who will never touch it. Administrative privileges are distributed far beyond any reasonable operational need.

AI agents remove the behavioral constraints that have kept this exposure theoretical. They inherit the full permission surface — and they act on it immediately.

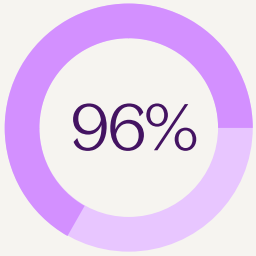

of permissions granted to employees are never used

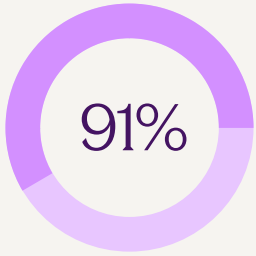

of sensitive data available to workers goes untouched

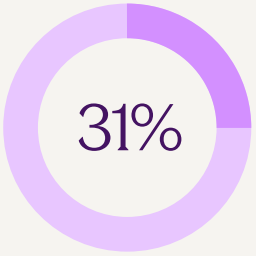

of workers can modify or delete sensitive data

4%

of workers log into most enterprise apps in any 90-day period

<20%

of available permissions used, even among active workers

When an AI agent inherits an existing employee permission set, it inherits the full scope of that access — not just the small slice the employee typically uses. Unlike humans, agents operate continuously, interact directly with APIs and data systems, and can be directed by adversaries to take actions no legitimate user would ever perform.

Key Insight

Agents Don't Behave Like Humans — They Run Without Limits

Humans are mostly understandable: they move slowly, apply judgment, and don't want to get fired. So most permissions stay untouched. Agents don't work that way. They run continuously, testing what they can do. Then they do it.

Human - 90 days of activity

sparse activity

Agent - first 30 seconds

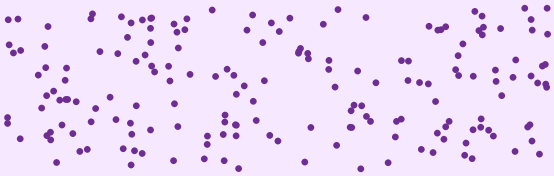

dense, immediate activity

Illustrative comparison of activity patterns — not drawn from the dataset. Real-world agent incidents confirm the pattern.

	Human Worker	AI Agent
Speed	Actions spread across hours, days, or weeks	Hundreds of actions per second
Judgment	Trusted to apply judgment; limited by hours in a day	Easily tricked, prone to hallucination; can confidently take wrong actions
Accountability	Professional consequences for misuse	No personal consequences for mistakes
Operating hours	~8 hours/day, within normal workflows	Continuous, 24/7, no natural stopping point

Finding 01

Massive Permission Exposure Exists Across the Enterprise

Cyera Research & Oso analyzed 2.4 million workers and 3.6 billion application permissions. The gap between granted access and actual usage is staggering — creating a dormant attack surface embedded directly into enterprise access models.

Organizations grant broad access to avoid blocking work. Permissions accumulate — granted to unblock a project, fix an issue, run a report — and are never revoked. Humans move slowly enough that the gap is tolerable. Agents remove that constraint.

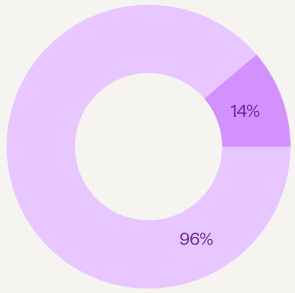

96% of permissions go untouched
All granted permissions

14% used 96% unused

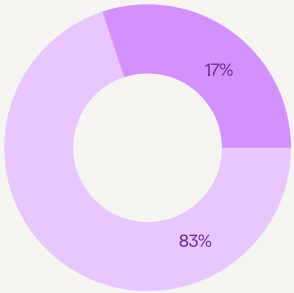

83% permissions used
by active workers

17% used 83% unused

4 in 100

Corporate workers take any action in most apps over a 90-day period

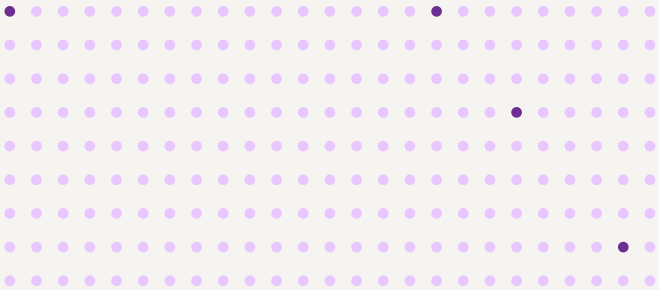

Only 4 in 100 corporate workers take any action in most applications over a 90-day period. Even among those who do, 83% of available permissions sit dormant — and agents inherit all of it.

At 1Password, we're seeing the same pattern this research highlights as teams start putting AI agents into real production workflows. Access models built for humans don't map cleanly to agents. When agents are handed broad, static permissions, the unused ones don't just sit there, they quietly expand the attack surface."- **Nancy Wang — CTO, 1Password**

Finding 02

Most Employees Retain Broad Access to Sensitive Data — Including Data They Rarely Touch

Despite most permissions going unused, many employees still retain the ability to interact with sensitive information. In many environments these permissions remain permanently available, even when they are rarely required.

Workers touch 9% of the data they can reach

Human - 90 days of activity

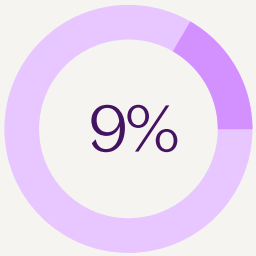

records accessed

Agent - seconds

if reachable, accessible

Illustrative — not a direct finding, but a logical consequence of the data.

13 of 100

workers can reach regulated data

100 workers. Each dot is one person Can access regulated data (13 of 100)

Regulated data includes PII, financial records, and health information. 13 in 100 corporate workers have access to regulated data. When agents inherit those accounts, they inherit that access too.

"The biggest mistake companies are making with AI agents is assuming yesterday's identity model will hold. At Brex, we're deploying agents aggressively, but we're designing for failure modes upfront, not after an incident. Speed without control is risk, and control without speed is a blocker." - **Mark Hillick — CISO, Brex**

Finding 03

Enterprise Access Models Systematically Overprovision Users

Excessive permissions are not accidental — they are often built directly into how enterprise systems are configured. Rigid, profile-heavy configurations grant broad privileges by default and are rarely revisited.

Admin privilege distribution across enterprise environments

Configuration Method

80%+

of access still managed through static profiles — despite platform guidance recommending modular permission sets

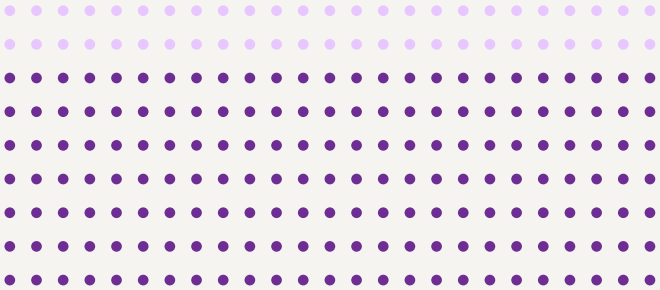

Profile-only Users

1 in 4

users rely solely on profiles, creating rigidity, overexposure, and audit challenges

Admin Access

~30%

admin assignments in some environments — up to 6× the expected level

High-risk permissions such as Modify All Data and View All Data can override normal data-sharing controls, allowing users to access or alter large datasets across an entire environment. These are exactly the kinds of broad, static permissions that AI agents will inherit by default.

Finding 04

How Enterprise Systems Quietly Accumulate Excess Permissions

Permission sprawl rarely happens through a single decision. It accumulates gradually as systems evolve, roles expand, and access bundles are reused across teams. Enterprise CRM environments offer a clear example: access is still largely structured around static profiles, even though the platform itself recommends a more granular model.

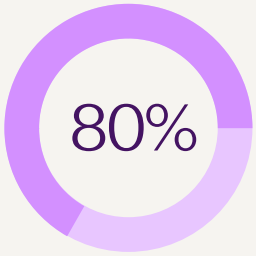

of Salesforce access managed through static profiles across environments analyzed

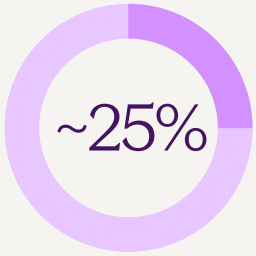

of users relying entirely on profiles for their permissions

Profiles bundle large numbers of privileges together. Over time, as organizations add new workflows, integrations, and administrative needs, these profiles accumulate permissions that exceed what most users actually require. "View All Data" and "Modify All Data" amplify this risk further — overriding normal sharing controls and allowing unrestricted access across an entire environment.

Broad, durable permissions persist long after their original operational need has passed — and AI agents will inherit every one of them.

Case Study - Cyera Research Deep Dive

Salesforce: Where Permission Sprawl Becomes Concrete Risk

The dynamics described in Finding 04 are not hypothetical. Cyera Research Labs conducted a dedicated deep-dive analysis of Salesforce environments across multiple organizations, and the findings put specific numbers to the structural patterns that permission sprawl produces in one of the most widely deployed enterprise platforms in the world.

Salesforce sits at the center of most organizations' customer data ecosystems — holding contact records, deal pipelines, support histories, and regulated customer information. Its layered permission model is both powerful and perilous: enabling precise access control in theory, while silently producing overexposure when not actively governed.

How Salesforce Access Is Supposed to Work

Salesforce's recommended model separates access into three layers: a minimal Profile per user that defines only baseline access, Permission Sets that grant role-specific privileges additively, and Permission Set Groups that bundle sets by function for easy assignment and review. This modular approach is designed to support least-privilege at scale — specific, auditable, and easy to revoke.

Profile

One per user. Base-level access only.

→

Permission Sets

Additive. Multiple per user. Rolespecific grants.

→

Permission Set Groups

Bundle sets by function. Assign once, manage once.

Salesforce best-practice access model — minimal profiles, modular permission sets

What Cyera Research Found in Practice

The reality in production environments diverges sharply from the recommended model. Cyera Research's analysis revealed that organizations are still managing access predominantly through static profiles — the very configuration pattern that Salesforce itself advises against.

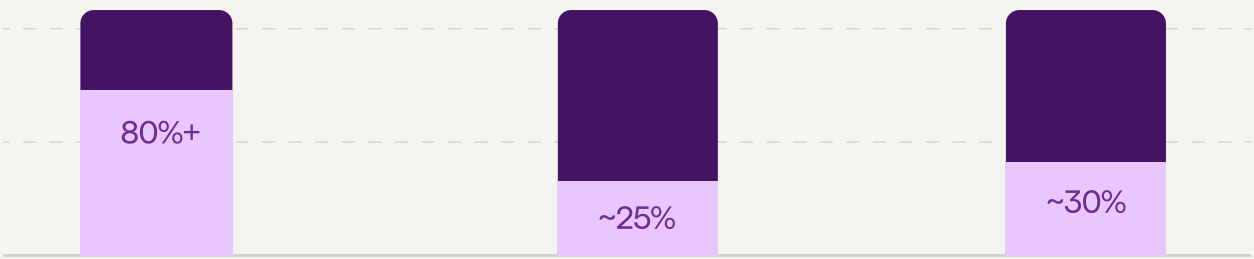

Salesforce recommends profile-minimal, permission-set-driven governance — most orgs invert this

5× higher than best-practice benchmarks — creating rigidity, overexposure, and audit complexity

Typical governance standard is ~5% — some environments showed 6× that level

The ‘Nuclear Buttons’: View All Data and Modify All Data

Within Salesforce, two permissions stand above all others in their potential for damage. View All Data grants read-only access to every record in the org, regardless of the sharing model. Modify All Data goes further — providing read, write, and delete access across the entire environment, effectively elevating the holder to super-admin status.

These permissions are frequently granted temporarily to resolve a technical issue or expedite a project — and then never revoked. Cyera Research’s analysis found these high-privilege capabilities distributed far more broadly than intended, in many cases persisting long after the original justification had passed.

From Static Configuration to Living Governance

The conclusion Cyera Research draws from this analysis applies well beyond Salesforce: access management cannot be treated as a one-time configuration. It must be a living governance process — continuously reviewed as teams change, integrations are added, and roles evolve. Profiles should be kept minimal. The majority of access should be managed through modular permission sets that are easy to assign, audit, and revoke.

This is precisely the kind of governance infrastructure that must be in place before AI agents are introduced. An agent connected to a Salesforce environment still running on over-provisioned profiles will inherit the full weight of that legacy configuration — with none of the human hesitation that has kept it dormant.

Further Reading: Are Your Salesforce Permissions Protecting You — or Exposing You? cyera.com/research

This analysis is the first in Cyera Research’s Salesforce Access Control Deep Dive series, covering profiles, permission sets, high-privilege capabilities, record-level access, and public data exposure.

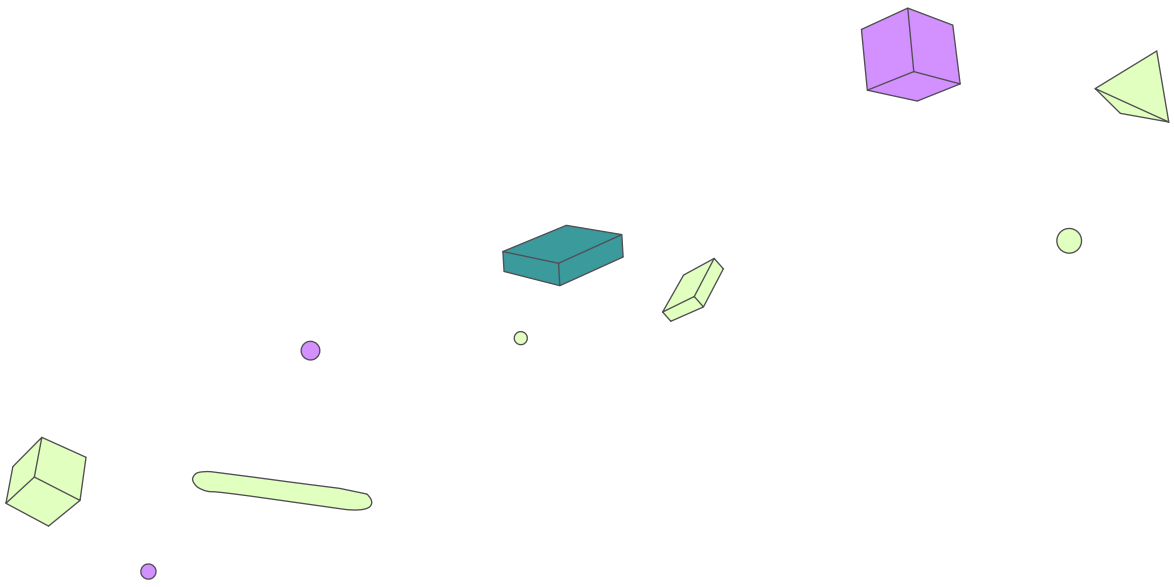

Finding 05

A Significant Minority of Workers Hold Dangerous Write and Export Permissions

Delete, modify, export — capabilities intended for selective use are held by a meaningful share of the workforce. Humans use them sparingly. Agents don't apply that restraint.

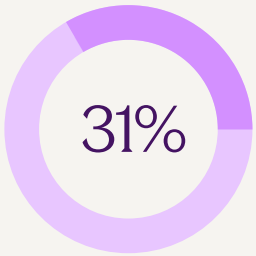

of corporate workers hold modify/delete capabilities

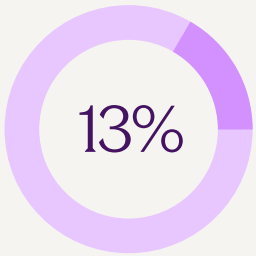

of corporate workers hold data export capabilities

Prevalence of high-risk capabilities across the workforce

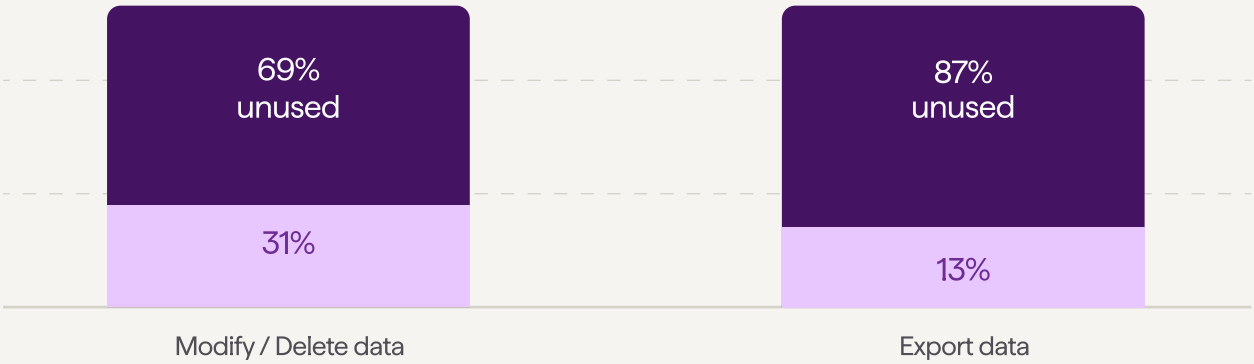

"At HashiCorp, we've spent many years helping organizations manage their sprawl of secrets. We consistently see that both humans and services are over-privileged and that bad hygiene turns into a major security threat. Now with agents, we are seeing the risks compound exponentially." - **Armon Dadgar — Co-Founder & CTO, HashiCorp**

Key Implication

AI Agents Will Inherit the Entire Permission Surface

For human users, unused permissions often remain dormant. AI agents change that dynamic completely. When agents inherit existing user permissions, they gain access to the entire permission surface — not just the subset employees normally use.

- access sensitive data the employee never viewed
- modify records the employee never opened
- export information the employee never downloaded

The enterprise permission model built for human users becomes a security crisis once AI agents inherit it. Dormant access is not safe access — it is latent exposure waiting to be activated.

Three Implications for Enterprise Security

Unused permissions represent unnecessary exposure.	When the vast majority of permissions remain unused, organizations can reduce risk significantly by identifying and eliminating dormant access without disrupting normal operations.
The real risk lies in access to sensitive data and high-impact actions.	Permissions that allow data modification, bulk export, or unrestricted visibility across datasets determine the potential blast radius of both operational mistakes and malicious activity.
Autonomous systems require purpose-built access models.	Agents should not inherit broad human permission bundles. Their access should be narrowly scoped to the specific systems, data, and actions required for their tasks — nothing more.

Recommendations

Ten Actions to Close the Gap Before Your Next Agent Deployment

Ordered from immediate to strategic. Each is linked to the research finding that motivates it.

01

Audit permission sprawl before deploying agents

If 96% of human access goes unused, that same access should not be handed to an agent. Use identity governance tools to surface dormant and excessive permissions before any agent deployment.

02

Create dedicated agent identities

Don't let agents inherit user credentials. Purpose-built identities, minimum permissions. Agent identities should be auditable, revocable, and entirely distinct from the human access model.

03

Start agents in read-only mode

Observation first, write access later. Deploy in observation mode before granting write or delete access. Use logs to right-size permissions before expanding.

04

Log every agent action from day one

You cannot govern what you cannot see. Every tool call, API request, data access event, and permission exercise must be captured from the start.

05

Configure agent-specific detection rules

SIEM alerts for out-of-scope queries, unfamiliar data access, privilege escalation attempts, and activity at anomalous times or volumes.

06

Triage permissions by blast radius

Lock down modify, delete, and export first. Read-only access is lower priority. The permissions that can cause irreversible or mass-scale damage require the most urgent attention.

07

Require continuous visibility into sensitive data access

Know where regulated and sensitive data resides, which identities — human and agent — can reach it, and how that access is exercised in real time.

08

Use only vendor-maintained integrations

Battle-tested MCP servers over community tools with untested security boundaries. Recent reports found over 40,000 exposed agent instances running malicious community-contributed integrations.

09

Expand agent access incrementally

One integration at a time. Monitor, validate, then expand. Treat each new integration as a controlled rollout. Broad permissions granted at deployment are far harder to reduce after the fact.

10

Make this a board-level conversation

Frame permissions infrastructure as what accelerates safe agent adoption. Few situations allow security to so clearly unlock business value.

Appendix

Definitions & Methodology

Permission

Can this user perform this action on this resource?

User

A person, service, or AI agent taking action.

Resource

The data or object being acted on.

Action

What the actor does: view, edit, delete, export, share.

Permission usage

Tracks whether you exercise a capability at all. If you can read documents and you read one, that's 100% usage — regardless of whether you could access 1 or 1,000 documents.

Resource access

Measures how many individual records get touched out of the total available. For example: if your CRM contains 500 opportunities and users access 10 of them in 90 days, then 2% of resources have been accessed.

Fifty Years of Principle, Almost No Measurement

To our knowledge, this is the first research examining how permissions are exercised in production. Most discussions focus on policy — how access should be structured. Far less is known about how access is actually used.

Users relying solely on profiles for all access

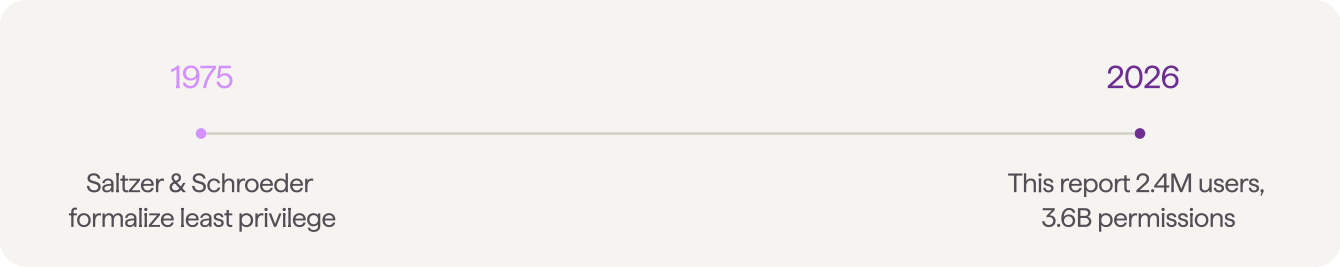

Our dataset is substantial but represents a slice of the ecosystem — companies who already take security seriously by investing in tools to manage access and data risk. Similar analyses from large infrastructure providers would deepen the industry's understanding considerably.

About Cyera Research

Cyera Research is the data-centric research arm of Cyera, dedicated to advancing vulnerability research and transforming real-world data insights into decisive security action. Led by a multidisciplinary team of researchers, scientists, security engineers, and security vulnerability researchers, they uncover critical vulnerabilities, emerging attack vectors, and AI-driven risks across modern data environments. By combining hands-on vulnerability discovery with rigorous, evidence-based research, Cyera Research delivers actionable intelligence and practical guidance that empower organizations to proactively secure, govern, and protect their data and AI assets with confidence.