unosecur

WHITEPAPER

Identity security in the pharmaceutical industry

Why identity is the common thread through pharma cyber risk, from R&D and clinical trials to manufacturing, patient support, and AI agents.

Table of contents

Executive summary

The pharmaceutical enterprise runs on shared access. Drug discovery flows through cloud-connected research platforms and AI tools. Clinical trials depend on CROs, trial-site networks, and third-party data systems. Manufacturing runs on validated execution systems that were never designed for a connected threat environment. Patient support programs sit at the intersection of pharma commercial operations, specialty distributors, and protected health information.

Every stage of the molecule-to-market chain depends on identities that are difficult to govern. Scientists, vendors, trial-site users, CRO and CDMO staff, service accounts, lab instrument credentials, API tokens, cloud roles, manufacturing system accounts, data integrations, and AI agents all touch regulated and proprietary systems.

The threat landscape reflects this exposure. CybelAngel analyzed 172 pharmaceutical-sector incidents between January and late September 2025, with the pattern driven by research data, regulated IP, and data-centric cybercrime. Atos Eviden recorded around 50 ransomware attacks against pharmaceutical companies in 2026 alone. Comparitech reported that 130 ransomware attacks targeted healthcare businesses, including pharmaceutical manufacturers, in the first nine months of 2025, alongside 293 attacks on hospitals and direct care providers.

The pattern across every major incident is consistent. Identity is the attack surface. Compromised service accounts, misconfigured access controls, overprivileged third-party credentials, ungoverned data integrations, and AI agents with unscoped access form the path from initial compromise to operational disruption, IP exposure, or a patient data breach.

Metric	Figure
Pharmaceutical-sector incidents analyzed, Jan to Sep 2025	172
Ransomware attacks on pharma companies in 2026	-50
FDA submissions with AI components reviewed, 2016 to 2023	500+
Records exposed in a single clinical-trial database misconfiguration	1.67M

This whitepaper maps the identity threat landscape specific to the pharmaceutical industry, examines the incidents that define it, and shows what an identity-first security posture looks like across R&D, clinical operations, manufacturing, patient support, and AI-enabled drug development.

Section 1: The pharma identity landscape

Most enterprises have complex identity environments. Pharmaceutical enterprises have a uniquely high-stakes one, with three identity types coexisting across every major organization and each carrying a different risk profile.

Human identities

Human identities in pharma span research scientists, lab technicians, clinical operations teams, quality assurance, manufacturing operators, regulatory affairs, pharmacovigilance, medical affairs, commercial teams, contractors, CRO staff, CDMO users, trial-site personnel, auditors, and supplier representatives.
The defining challenge runs beyond internal workforce governance. A single clinical program may involve dozens of organizations sharing access to sensitive research, patient, and safety data, with the extended identity perimeter created by federated access across sponsors, CROs, CDMOs, trial sites, distributors, and patient-support vendors.

Non-human identities (NHIs)

Non-human identities are embedded throughout the pharma technology stack. They include LIMS integrations, electronic lab notebook service accounts, EDC and CTMS APIs, eTMF connectors, manufacturing execution system accounts, ERP integrations, QMS workflows, lab instrument service accounts, robotic lab automation, batch-record platforms, cloud data pipelines, pharmacovigilance feeds, and API tokens shared with vendors and partners.
These identities are created continuously as workflows automate and systems integrate. They rarely receive the governance that human accounts do. Ownership disappears as teams change. Rotation lags. The scope is rarely reviewed after the initial deployment. Many of these systems operate within validated or GxP-sensitive environments, where changes carry compliance implications that further slow remediation.

AI Agents

AI agents are the newest and least-governed identity class in pharma. They now operate across drug discovery, clinical trial intelligence, regulatory writing, safety signal detection, manufacturing deviation analysis, QA review, knowledge retrieval, and supply-chain planning.

An AI agent interprets instructions, selects tools, chains actions, and operates across multiple systems at the same time. Unosecur describes agentic AI as "an identity expansion event," and that framing carries regulatory and operational weight in pharma. An AI agent with access to proprietary compound libraries, clinical trial data, or regulatory submission documents holds credentials, access paths, and behavioral patterns that demand identity-grade governance.

Pharma carries an identity governance problem across R&D, trials, manufacturing, patient support, and AI-assisted operations. The convergence of human identities, non-human identities, and AI agents across cloud infrastructure, SaaS platforms, validated on-premise systems, and operational technology creates an identity perimeter that most pharmaceutical organizations have only partially mapped.

Section 2: How identity is being attacked

The following incidents represent recurring patterns of identity exploitation across the pharmaceutical and life sciences sector.

1. **2025 | Inotiv ransomware-style cyberattack**
   Pattern: CRO operational disruption. Unauthorized access and system encryption disrupted business applications and internal data storage at a publicly traded preclinical contract research organization.

   Inotiv supports drug discovery and development for pharmaceutical clients across preclinical research, safety assessment, and laboratory testing services. Its August 2025 SEC filing confirmed that a threat actor gained unauthorized access and encrypted certain systems, disrupting business operations and temporarily affecting access to internal data storage and business applications. Inotiv's subsequent notice confirmed unauthorized access between approximately August 5 and August 8, 2025, with possible acquisition of certain data.
   The public filings do not confirm the initial access vector. The identity lesson centers on blast radius. A threat actor operating inside a CRO environment can reach research data, client program information, study records, and the business applications that support ongoing drug development. Governing privileged human and non-human access across research systems carries continuity weight far beyond compliance documentation.

2. **2025 | DM Clinical Research database exposure**
   Pattern: Clinical-trial data exposure. A misconfigured database left approximately 1.67 million records and 2 TB of sensitive trial participant data accessible online without authentication.

   In early 2025, security researcher Jeremiah Fowler discovered a publicly accessible database holding approximately 1.6 million records belonging to DM Clinical Research, a clinical trials network operating sites across the United States. The exposed records included names, contact details, medical histories, medication information, and other sensitive personal and medical data drawn from clinical research participants. The database was not password-protected, and the exposure was confirmed and reported by the HIPAA Journal.
   The access-control reality here reflects the realities of distributed clinical research. Trial data flows across sponsors, CROs, site networks, and third-party data systems. Responsibility for securing that data is distributed across organizational boundaries that identity governance rarely crosses cleanly. The failure was the absence of enforced access controls on a system that held regulated, sensitive data about real trial participants.

3. **2024 - 2026 | Cencora and Lash Group breach**
   Pattern: Pharma supply-chain concentration. A patient-support and drug-distribution platform became a shared data-exposure point for more than two dozen pharmaceutical manufacturers.

   Cencora, formerly AmerisourceBergen and one of the world's largest pharmaceutical distributors, disclosed in February 2024 that data had been stolen from its information systems. Reuters subsequently reported that the stolen data included personally identifiable information and protected health information, primarily from its Lash Group subsidiary, which provides patient-support services. Lash Group's notice confirmed exfiltration from its systems and described its role across pharmaceutical partnerships, patient support, drug distribution, business analytics, and technology services. The breach ultimately triggered notifications from a long list of pharmaceutical manufacturers whose patient data flowed through Lash Group platforms, with downstream notifications extending into 2025.
   This incident defines a structural risk in pharma commercial operations. Patient-support programs, specialty distributors, and drug-access services are shared platforms. A breach at one partner becomes a breach across every pharma company that shares data with that platform. The supply-chain identity backdoor is equally present in pharmaceutical distribution and patient-support ecosystems.

4. **2025 | Healthcare ransomware surge across the broader ecosystem**
   Pattern: Attackers shifting focus to vendors, service partners, and healthcare businesses, including pharmaceutical manufacturers.

   Comparitech data published by Industrial Cyber in 2025 shows that in the first nine months of 2025, 293 ransomware attacks targeted hospitals and direct care providers, while another 130 targeted healthcare businesses, including pharmaceutical manufacturers, medical billing providers, and healthcare technology companies. Overall, healthcare ransomware volume grew approximately 30 percent year over year, with vendors and service partners absorbing a disproportionate share of the increase.
   Pharmaceutical organizations operate as nodes in a shared ecosystem with providers, CROs, payers, distributors, patient-support vendors, and technology partners. Attackers move through the ecosystem rather than against individual organizations. The identity perimeter in pharma extends far beyond the enterprise boundary, and the access paths that span it are where the risk is concentrated.

The Five Identity Attack Patterns

Across these incidents, five patterns recur:

- **Ransomware disruption through privileged access:** threat actors reaching research systems, CRO operations, manufacturing applications, or business platforms through compromised or overprivileged accounts.
- **Third-party and supply-chain data exposure:** CROs, patient-support programs, trial-site networks, and specialty distributors serving as shared entry points for pharma data.
- **Weak access control on sensitive data stores:** misconfigured trial databases, cloud storage, and data integration endpoints left accessible beyond their intended scope.
- **Non-human identity sprawl:** LIMS, ELN, MES, EDC, lab instruments, and automation credentials accumulating without ownership, rotation, or lifecycle management.
- **AI-agent access drift:** agents connected to research data, regulatory documents, safety data, and internal tools operating without identity-grade governance.

In every case, the identity access was the target.

Section 3: The emerging frontier: AI agents as pharma identities

AI agents are already operating inside pharmaceutical environments. They query

proprietary compound libraries and external research databases. They access trial schedules, patient cohorts, site performance data, and protocol deviations. They summarize adverse events and case narratives for pharmacovigilance teams. They draft regulatory submission modules from internal source documents. They read batch records, deviations, CAPA notes, and quality events to support manufacturing QA. They optimize API sourcing, cold-chain logistics, and demand planning across supply chains.

In most deployments, these agents hold broad access because restrictions slow adoption. That tradeoff carries weight in pharma that it carries in few other industries. The FDA's 2025 draft guidance on AI used to support regulatory decision-making places AI systems inside workflows that may produce or support regulated evidence, safety data, clinical assessments, and submission documents. Governing the identities of those systems is part of the regulatory record.

The specific risk scenarios for pharma are distinct from other industries.

- **Drug-discovery agent credential exposure:** an agent with access to proprietary compound libraries, research databases, and molecular modeling systems is manipulated or misconfigured to expose IP. The agent's existing access provides the path without any breach of the research platform itself.
- **Clinical operations agent overreach:** an agent responsible for trial schedules, patient cohort data, and protocol deviation records has broader access than their function requires. A compromised instruction or tool causes it to retrieve information outside its intended scope.
- **Regulatory writing agent data exposure:** an agent drafting submission modules has access to clinical study reports, statistical analyses, and proprietary trial data. Without scoped access controls, the agent becomes a retrieval path for sensitive regulatory content.
- **Pharmacovigilance agent manipulation:** an agent that summarizes adverse events and case narratives is manipulated via crafted input to retrieve or alter safety records, threatening the integrity of safety evidence that supports regulatory obligations.
- **Confused deputy:** the agent holds legitimate permission to access a regulated system, and the attacker tricks the agent into using that permission on their behalf.
- **Cross-agent propagation:** in multi-agent architectures common in R&D and clinical operations platforms, a compromised agent passes credentials, files, or instructions to downstream agents, spreading access across the research or manufacturing environment.
- **Orphaned agent credentials:** agents created for discrete tasks in drug discovery, trial operations, or supply-chain planning retain tokens and data access long after the original use case ends.

- **Audit gap:** When an action is taken by an agent acting on behalf of a user acting through a platform, attribution across regulated workflows becomes genuinely difficult.

AI adoption in pharma is outpacing identity governance. The FDA reports experience with more than 500 submissions containing AI components from 2016 to 2023, with the trajectory continuing to accelerate. Governance frameworks to match that scale are still being built across most organizations.

Section 4: Why this is hard to solve

Understanding the attack patterns is necessary. Acting on them is harder because security leaders responsible for pharmaceutical environments work against structural conditions that complicate identity security,, independent of investment and intent. The environment is federated by design. Sponsors, CROs, CDMOs, trial sites, labs, distributors, patient-support vendors, and regulators all need controlled access to shared data and systems. Every partner boundary is a potential identity control gap with uncertain enforcement on the other side, and the federation that enables drug development at scale also creates the access paths attackers can chain.

The data is high-value and multi-context. A single identity in a pharmaceutical organization may access proprietary compound data, patient records, safety data, trial data, batch records, and evidence for regulatory submissions. The value of that data to competitors and threat actors makes pharma a priority target, and the regulatory sensitivity of that data makes a breach consequential far beyond financial loss. Legacy and validated systems slow security change. LIMS, electronic lab notebooks, manufacturing execution systems, electronic batch record platforms, QMS workflows, and lab instruments often cannot be modified quickly without triggering revalidation. That creates a lag between the identity risk on these systems and the governance controls that security teams can apply to them.

Non-human identities are everywhere and rarely governed. Lab instruments, workflow automation, APIs, service accounts, robotic systems, and cloud pipelines operate continuously across research and manufacturing. They are created for specific integrations and rarely reviewed afterward. Credentials persist, ownership dissolves, and scope rarely shrinks. Compliance produces documentation without runtime control. Audit evidence may show approved access and completed access reviews, while exercised access, dormant tokens, orphaned accounts, and service credentials remain poorly governed. Periodic review cycles do not reflect how identity risk actually moves through a pharmaceutical environment.

AI adoption is entering regulated workflows. The FDA's 2025 guidance on AI in regulatory decision-making creates a new category of accountability. When an AI agent contributes to a regulatory submission, safety assessment, or clinical evidence package, the governance of that agent's identity, access, and behavior becomes part of the regulatory record.

Section 5: What an identity-first defense looks like

Attack pattern	Identity-first control required
CRO and research-system disruption	Runtime behavioral monitoring, Identity Timeline, effective permission analysis
Patient-support and supply-chain data exposure	Unified Identity Fabric across human, third-party, and non-human identities
Clinical-trial database misconfiguration	Access-path visibility, unused privilege reduction, third-party identity governance
Lab and manufacturing service account sprawl	Non-human identity discovery, ownership, lifecycle, and activity history
AI agents in R&D and regulatory workflows	AI agent discovery, MCP governance, agent permissions, tool access, auditability
Compliance evidence gaps	Identity-level evidence across access, behavior, ownership, and remediation history

An effective response covers all three identity types: human, non-human, and AI agent. Coverage extends across cloud, SaaS, validated on-premises systems, and operational technology simultaneously. Runtime behavior analysis is required because configuration alone cannot show whether access is being used safely or whether it has drifted outside its intended scope.

Unosecur's Unified Identity Fabric is built for this identity reality. The platform helps pharmaceutical enterprises discover human, non-human, and AI agent identities, understand what they can access, monitor what they actually do, and act when identity risk appears across cloud, on-prem, SaaS, identity providers, manufacturing systems, and AI environments. Integration is read-only and agentless, which matters where validated systems cannot tolerate disruption. Native API integrations reach roughly six times as deep as standardized protocols such as SCIM, SAML, and OIDC. Tenant provisioning completes in 15 minutes.

Unified Identity Fabric

Pharmaceutical identity security fails when identity data stays fragmented across systems. Research platforms know who can access compound libraries. Clinical systems know who can access trial data. Manufacturing systems know which service accounts touch batch records. Cloud providers know which roles are active. Patient-support platforms know which partner integrations are connected. Attackers move across all of them.
The Unified Identity Fabric connects those signals into one operating view. It surfaces the relationships, access paths, unused privileges, risky behaviors, and identity dependencies that fragmented governance leaves invisible across the molecule-to-market chain.

Non-human identity governance

The pharmaceutical environment is dense with non-human identities. LIMS service accounts, ELN tokens, EDC and CTMS API credentials, lab instrument accounts, robotic automation identities, MES connectors, cloud pipeline credentials, and vendor integration tokens all operate continuously inside regulated environments.
Unosecur discovers these identities, connects them to ownership, baselines their behavior, and reduces standing privileges through activity-based right-sizing. Every non-human identity should carry an owner, a purpose, a scope, a lifecycle, and a behavior history. Findings include locally managed accounts within validated and GxP systems, partially offboarded CRO and vendor credentials that remain active after a study or contract closes, and orphaned lab integration tokens.

Runtime behavioral monitoring

Pharmaceutical identity risk cannot be measured solely from permissions. A service account accessing batch records outside its normal pattern matters. A lab integration credential used from an unexpected location matters. A vendor API token suddenly exercising privileges it has never used matters. Whether a research system account remains active after a contractor's departure matters.
Unosecur monitors identity behavior at runtime across connected environments, so security teams can distinguish expected access from anomalous access before it becomes lateral movement or data exfiltration. Findings include Shadow Admins inside research, clinical, and manufacturing platforms, identity drift across vendor integration tokens, and MFA blind spots in partner connections.

Activity-based right-sizing

Pharmaceutical environments accumulate privilege quickly and reduce it rarely. A cloud role receives broad access for a system migration. A service account is

provisioned for a vendor integration with more scope than required. A CRO user account retains access to a sponsor's trial platform after the study closes. An AI agent receives access to accelerate a project and keeps it permanently.
Unosecur aligns permissions with actual usage and surfaces JEP recommendations through activity-based right-sizing. Standing privileges shrink, JIT access decisions become measurable, and toxic privilege combinations across research, clinical, and regulatory systems are flagged before exploitation. Cross-cloud identity chains spanning research-to-manufacturing data flows are fully reconstructed.

Identity Timeline

When an incident crosses research platforms, cloud infrastructure, manufacturing systems, partner integrations, and patient-support platforms, logs alone are not enough. Security teams need to know which identity acted, what access it held at the time, what systems it touched, whether its behavior had changed, and which other identities or data stores were connected to it.
The Identity Timeline provides security teams with behavioral history for each identity across connected systems. In pharma, that means reconstructing what a researcher, vendor account, service credential, lab automation identity, CRO user, or AI agent actually did, in sequence, across every connected platform.

Security for AI

AI agents must be governed as identities across the pharmaceutical environment. The AI Agent Dashboard provides security teams with a dedicated way to view agents, assess risk, review data access, inspect permissions, and trace execution paths. For pharma, that means visibility into which agents are operating across drug discovery, clinical operations, regulatory workflows, pharmacovigilance, and manufacturing; which data sources they can access; and whether their behavior aligns with their intended purpose.
The MCP Auth Gateway extends governance into agent-tool interactions. It controls which tools agents can call, what data they can reach, and how agent activity is audited across environments where regulated evidence and proprietary IP are at stake.

AI for Security

ARK AI is Unosecur's AI copilot that plans and executes actions across the Identity Fabric. A security engineer makes a request in natural language. ARK builds an execution plan, waits for approval or modification, executes the approved steps, and automatically writes an audit trail for each step.
ARK reaches across permissions, NHI inventory, AI agent inventory, the identity

graph, compliance evidence, ticketing, and threat response. In a pharmaceutical context, ARK can surface dormant vendor accounts with active access to trial data, overprivileged AI agents operating within regulatory workflows, service accounts without an owner across manufacturing systems, or identity-level evidence for compliance reporting. Response time compresses while approval and traceability remain in place.

Identity-first defense across the pharmaceutical chain

Pharmaceutical identity security cannot be solved one system at a time. A validated system access review will not govern a cloud pipeline credential. A third-party risk assessment will not indicate whether a CRO service account remains active after a study closes. A cloud posture tool will not explain why an AI agent can reach proprietary compound data. A SIEM will not automatically reconstruct the identity chain across research platforms, clinical systems, manufacturing, and patient-support integrations.
The Unified Identity Fabric connects these surfaces, shows which identities can move across them, and helps reduce the access paths attackers chain before they become a breach, a regulatory event, or a manufacturing disruption.

Conclusion

Drug discovery, trial management, manufacturing, and patient support have become interconnected identity environments spanning dozens of organizations. Every CRO relationship, CDMO integration, trial-site network, patient-support platform, and AI tool deployed in a drug development workflow forms part of the identity surface. Most of these surfaces are governed by separate tools, separate teams, and separate policies, without a unified view of who has access to what.

The incidents of recent years show a consistent pattern. The Inotiv attack showed that a CRO environment, when compromised, can expose research operations and client data across the drug development ecosystem. The DM Clinical Research exposure showed that a single misconfigured access control can expose 1.67 million clinical trial records. The Cencora and Lash Group breach showed that patient-support and distribution platforms, when compromised, become breach vectors across the pharma companies they serve.

These patterns share a common root. The identity perimeter has expanded faster than the governance frameworks designed to protect it. AI agents make this more urgent. As pharmaceutical organizations deploy autonomous agents across drug discovery, regulatory, clinical operations, pharmacovigilance, and manufacturing, they

create identities that operate with broad access, make autonomous decisions, and produce evidence that may sit inside regulated workflows. The FDA's 2025 AI guidance makes the stakes explicit, with identity governance for AI systems becoming a regulatory concern alongside a security one.

Three things every pharmaceutical security leader should do now.

- **Inventory every identity type across the organization:** human, non-human, and AI agent. Without an answer to which identities exist across research, clinical, manufacturing, and commercial environments, and what each one can access, identity risk cannot be governed.
- **Extend identity governance across every partner boundary:** CROs, CDMOs, trial-site networks, patient-support platforms, and specialty distributors all hold or access data and systems that belong to your organization. Third-party identity governance is a baseline requirement in a federated drug development model.
- **Establish AI agent governance before adoption scales further into regulated workflows:** every agent operating in a drug discovery, regulatory, pharmacovigilance, or manufacturing context should carry ownership, permission scoping, tool visibility, behavioral history, and auditability. The window between deployment and governance failure is short, and the consequences in a regulated industry are significant.

The organizations hit by the next wave of pharmaceutical cyberattacks will be those that allowed identities to accumulate access across research, clinical, manufacturing, and AI systems without unified governance across the molecule-to-market chain.