BLACK BOOK RESEARCH

HOSPITAL AI CYBERSECURITY READINESS

Q3 2026

What Hospitals Must Do Now Before AI Pilots, Data Loading, Agents and Third Parties Create the Next Breach Surface

A Hospital Action Companion to the Black Book State of Healthcare Cybersecurity 2026 Series

Executive Action Report

THESIS Hospitals are adding AI faster than many are extending identity, data-loss prevention, asset inventory, third-party governance and incident response around it. The resulting exposure is not theoretical: AI now accelerates attacker reconnaissance and social engineering while hospital users and vendors create new paths for regulated data to leave controlled environments.

Executive Brief: The State of Healthcare Cybersecurity Has Entered the AI Era

The hospital cyber threat model changed in 2026. This report extends the Black Book State of Global Healthcare Cybersecurity 2026 and US Hospital Cyber Resilience 2026 research series into the emerging AI control gap facing U.S. hospitals. The core attack methods remain familiar—ransomware, credential theft, vulnerability exploitation, social engineering and third-party compromise—but AI makes them faster, cheaper, more personalized and more scalable. At the same time, hospitals are creating new internal exposure through generative AI pilots, ambient and documentation tools, coding and revenue-cycle copilots, imaging algorithms, research models, private LLMs, cloud APIs, retrieval-augmented generation (RAG) stores, AI-enabled SaaS features and emerging agents.

Signal	What it means for hospitals
2026 Verizon healthcare data	1,438 confirmed healthcare breaches; System Intrusion, Miscellaneous Errors and Social Engineering accounted for 81%. External actors were 81%; 99% of actor motives were financial. Vulnerability exploitation was the leading identified initial-access vector at 20%, followed by phishing at 14% and credential abuse at 11%. [1]
AI is entering the attack chain	Verizon reports threat actors using GenAI for targeting, initial access, malware/tool development and other attack stages. The issue is speed and scale layered onto known TTPs. [1]
Shadow AI is a data-loss channel	Netskope's 2026 healthcare telemetry says regulated data represented 89% of healthcare GenAI data-policy violations. [2]
AI-enabled attacks now lead health-sector concern	Health-ISAC's 2026 threat landscape and CISO benchmarking identify AI-enabled attacks as the top emerging concern while ransomware and supply-chain compromise remain major operational threats. [3][4]
Regulators are enforcing basic cyber governance	OCR's July 2026 OSF Healthcare settlement emphasized accurate and thorough risk analysis after ransomware and ePHI exfiltration; OCR's 2026 enforcement actions repeatedly target failures in risk analysis and risk management. [5]

BLACK BOOK POSITION A hospital should treat every AI deployment as a new information system and a new trust boundary—not as a software feature. No model, agent, AI-enabled application or clinical pilot should receive production credentials, PHI, image data, notes, claims, financial data or tool access until it has passed an AI-specific security gate.

Black Book's 2026 healthcare cybersecurity research has repeatedly moved the definition of resilience beyond perimeter defense toward identity, third-party dependency, clinical continuity and recoverability. For an unprepared CIO, the likely failure is not "AI itself." It is governance debt: unknown tools, unclassified data flows, excessive service-account privileges, weak vendor controls, unmonitored model/API activity, untested downtime procedures and an inability to prove what happened after an incident. That combination converts a technical event into a patient-care, regulatory, legal and leadership crisis.

Prepared for hospital boards, CEOs, CIOs, CISOs, CMIOs, clinical engineering, privacy, compliance, legal and risk leaders. This report provides hospital leaders with a practical framework for securing AI-enabled clinical, operational and administrative environments against emerging cyber threats.

Where AI Creates New Hospital Breach Paths

Consistent with the Black Book State of Global Healthcare Cybersecurity 2026 framework, hospitals should inventory AI as an attack surface across six layers. A control that protects the EHR does not automatically protect a prompt, vector database, agent, model endpoint or third-party AI feature.

AI layer	Hospital exposure	Likely failure mode
User / prompt	Clinicians, coders, finance, researchers and vendors paste or upload data into AI.	PHI, credentials, screenshots, reports, source code or contracts leave governed repositories; non-corporate accounts defeat retention and DLP visibility.
Model / application	Clinical copilots, summarizers, decision support, chatbots, coding and workflow models.	Prompt injection, jailbreaks, insecure output handling, unsafe model behavior, leakage of system prompts or sensitive context.
RAG / knowledge stores	EHR extracts, policies, clinical documents, imaging metadata, claims, call transcripts.	Overbroad indexing, stale permissions, poisoned documents, retrieval of records a user could not otherwise access.
Agent / tools	Agents call EHR APIs, scheduling, email, ticketing, finance, pharmacy or cloud tools.	Tool misuse, privilege escalation, memory poisoning, unauthorized transactions or exfiltration through an allowed connector.
AI supply chain	Models, libraries, plugins, MCP servers, SaaS vendors, cloud services and subcontractors.	Compromised model files, dependency attacks, vendor breach, hidden subprocessors, uncontrolled data residency.
Infrastructure	GPU/cloud instances, notebooks, object stores, secrets, API keys, model gateways.	Exposed storage, hard-coded credentials, weak IAM, public endpoints, unpatched libraries, insecure development environments.

NIST's AI security work now explicitly separates three needs: securing AI system components, using AI to defend, and thwarting AI-enabled attacks. NIST's adversarial-machine-learning taxonomy includes evasion, poisoning, privacy and misuse attacks; its 2025–2026 Cyber AI work also highlights governance, AI attack surfaces and AI-enabled defense. [6][7][8]

The most underestimated hospital scenario is indirect prompt injection. A model or agent can ingest an instruction hidden in a document, webpage, email, file, image metadata or other retrieved content. If the agent has access to clinical or operational tools, the attacker may not need to compromise the user's password; the attacker instead manipulates what the agent reads. This is why agent permissions, tool allowlists, runtime inspection and transaction limits become security controls—not just AI governance features.

Immediate rule No AI agent should have broader privileges than the human role it supports. High-impact actions—medication, orders, payments, identity changes, mass export, code deployment, configuration changes or external messaging—should require deterministic policy checks and, where clinically or financially material, human approval.

The Threats Hospitals Should Expect Through 2027

Threat	AI amplification	Hospital consequence	Priority
Hyper-personalized phishing / vishing	Models generate credible language, role context, voice and pretexts at scale.	Help-desk takeover, MFA reset, EHR/cloud access, payroll and wire fraud.	Critical
Vulnerability discovery + exploitation	AI accelerates scanning, exploit research and adaptation.	Internet-facing gateways, VPNs, edge appliances, cloud services and third-party systems compromised before patch cycles catch up.	Critical
Credential and session theft	AI improves lures and automates account reconnaissance.	Privileged accounts, service accounts and vendor remote access become lateral-movement paths.	Critical
Ransomware + extortion	Automation improves target selection and operational mapping.	EHR downtime, imaging/lab disruption, ambulance diversion, delayed care, data publication.	Critical
Prompt injection / agent hijack	Malicious content steers AI to disclose data or invoke tools.	Unauthorized retrieval, transactions, messaging, workflow changes or external exfiltration.	High
Data poisoning / RAG poisoning	Attackers alter source data or trusted knowledge.	Corrupted recommendations, manipulated security analysis, clinical or operational misinformation.	High
Shadow AI / accidental disclosure	Convenient AI tools bypass formal procurement and security review.	PHI and regulated data copied to external services; unknown retention/training/residency.	High
AI supply-chain compromise	Malicious models, libraries, plugins, MCP servers or vendor updates.	One compromised dependency reaches many hospital environments.	High
Synthetic identity / deepfake fraud	Voice/video generation targets support desks, executives, clinicians and vendors.	Password resets, false approvals, vendor-payment diversion, fraudulent patient or workforce activity.	High

Verizon's 2026 DBIR reports that vulnerability exploitation became the most common initial-access vector across its broader dataset at 31%, while only 26% of CISA Known Exploited Vulnerabilities identified in the underlying remediation data were fully remediated during 2025 and the median time to full resolution rose to 43 days. Ransomware appeared in 48% of all breaches and third-party involvement reached 48% of breaches in the broader dataset. [1]

For healthcare specifically, the 2026 DBIR attributes 32% of breaches to third-party involvement and 54% to a human element. This means an AI security program that focuses only on model testing is incomplete. Hospitals must simultaneously harden identity, human verification, patching, third-party connections, data controls and recovery.

Clinical implication Cyber events must be managed as operational disasters. A JAMA Network Open study of a month-long ransomware event found significant increases in patient census, ambulance arrivals, wait times, patients leaving without being seen, length of stay and stroke-care disruption at nearby unaffected emergency departments. [9]

The Hospital AI Security Control Plane: What Must Exist Before Production

Building on Black Book's 2026 cyber-resilience and hospital-readiness work, Black Book recommends a hospital-level AI Security Control Plane built on existing cybersecurity architecture rather than a separate AI committee with no enforcement authority. The control plane should govern data, identity, models, agents, third parties and recovery through one auditable operating model.

Control	Minimum hospital requirement
1. AI asset inventory	Discover sanctioned and unsanctioned AI applications, models, APIs, agents, plugins, MCP servers, notebooks, vector stores, embedded SaaS AI and service accounts. Owner, business purpose, data classes, vendor and environment must be known.
2. AI data-flow map	Document what data enters, is retrieved, is cached, is logged, is used for training/fine-tuning, and exits the system. Include PHI, PII, payment data, credentials, images, source code and research.
3. Identity-first access	Phishing-resistant MFA for privileged/remote access; least privilege; short-lived credentials; workload identity; service-account inventory; privileged access management; no shared AI keys.
4. AI gateway / DLP	Route approved model/API interactions through policy enforcement that can classify prompts/files, block PHI or secrets where prohibited, log transactions and apply model-specific rules.
5. AI application testing	Red-team models and agents for prompt injection, data leakage, unsafe tool calls, privilege escalation, jailbreaks, poisoned retrieval and failure under malformed input before deployment and after material change.
6. Runtime protection	Inspect prompts, responses, tool calls and agent actions; enforce allowlists, rate limits, transaction boundaries, network egress controls and human approval for high-impact actions.
7. Third-party controls	BAA where applicable, data-use restrictions, model-training prohibitions as required, subprocessor inventory, breach notification, logging access, deletion/return, security evidence, resilience testing and exit plan.
8. Clinical continuity	Department-by-department downtime playbooks that assume EHR, cloud, AI vendor and network services may be simultaneously unavailable for days or weeks.
9. Immutable recovery	Segmented and tested backups for identity, EHR, imaging, interfaces, critical SaaS exports and security tooling. Recovery priorities defined by patient-safety impact.
10. Evidence and audit	Centralized logs for AI usage, admin changes, agent actions, data access, model versions and vendor events with retention adequate for forensic reconstruction and HIPAA obligations.

HHS's Healthcare and Public Health Cybersecurity Performance Goals already emphasize fundamentals such as asset inventory, third-party vulnerability disclosure, access controls, incident planning, network segmentation and recovery. AI should be mapped into those controls rather than treated as an exception. [10]

Cybersecurity Vendor Landscape: Where Innovation Is Moving and Where Hospitals Must Push Harder

The Black Book State of Global Healthcare Cybersecurity 2026 documented a market already moving away from point-solution accumulation toward measurable resilience, identity control, third-party visibility and recovery. In AI security specifically, the market is moving quickly from alert-centric security toward AI-aware discovery, data control, autonomous investigation, runtime protection and cyber-physical visibility. Hospitals should avoid assuming that a vendor's use of "AI" means it secures AI systems. The questions are different: can the platform see the AI estate, stop sensitive data movement, test models, constrain agents, correlate medical-device context, and recover clinical operations?

Firm / platform	Notable current innovation	Hospital validation question
Microsoft Security / Project Perception	Agentic security architecture intended to perceive, reason and act continuously across security signals; Microsoft describes machine-speed defense against AI-era attacks. [11]	Can it correlate hospital identity, endpoint, cloud, SaaS and AI activity without creating opaque automated actions? What human approval remains for disruptive containment?
Palo Alto Networks Cortex XSIAM / AgentiX	Unified SecOps data, AI/automation, agentic response and frontier-model support; positioned for machine-speed detection and response. [12][13]	Can response automation safely isolate clinical endpoints, shared workstations and critical servers without interrupting care?
CrowdStrike Falcon / Charlotte AI	Agentic investigation, malware analysis, workflow automation and agent orchestration in Falcon. [14]	How are AI-generated decisions evidenced, audited and constrained for high-risk healthcare assets?
SentinelOne Singularity / Purple AI	Autonomously initiated investigations with evidence chains and machine-speed response. [15]	What healthcare-specific guardrails prevent autonomous containment from disrupting clinical systems or medical-device workflows?
Cisco AI Defense	AI asset discovery, model/application validation, AI supply-chain scanning, runtime protection and agent/MCP controls. [16][17]	Can it discover every hospital model/agent/data source and enforce prompt/tool policies across on-prem, cloud and embedded clinical AI?
Zscaler AI Security	Shadow-AI discovery, AI gateway, DLP, AI asset management, agent/MCP controls and runtime security. [18][19]	Can inline controls identify PHI and prevent uploads without blocking legitimate clinical workflows or causing latency?
Netskope	Healthcare-specific GenAI data visibility and policy enforcement; 2026 telemetry highlights regulated-data leakage in GenAI use. [2]	Can the hospital identify personal AI accounts, embedded SaaS AI and API-based use—not only browser traffic?
Claroty xDome for Healthcare	Cyber-physical and medical-device visibility, clinical context, exposure management, segmentation and threat detection. [20]	Can the platform tie AI-enabled workflows to the devices, building systems and clinical infrastructure on which they depend?
Armis Centrix for Medical Device Security	Agentless asset discovery, clinical context, exposure prioritization, behavioral monitoring and medical/IoT/OT coverage. [21]	Can unknown assets and risky communications be discovered fast enough to support segmentation and incident containment?
Ordr AI Protect for Security	Behavior-based device discovery, risk prioritization and safe segmentation for healthcare environments. [22]	Can policies be validated before enforcement and integrated with NAC, firewall, SIEM and CMMS processes?

Where the market is still immature AI security remains fragmented. Hospitals still need to integrate SOC platforms, network controls, DLP/SSE, identity, cloud security, medical-device/CPS security, AI runtime controls and clinical downtime planning. “Single-pane” claims should not substitute for tested end-to-end containment and recovery.

The Failure Pattern: What Happens to Hospitals That Do Not Prepare

The Black Book 2026 cybersecurity series treats cyber failure as a clinical-continuity and enterprise-governance problem, not merely an IT outage. The consequences are sequential. The first failure is usually visibility; the last is institutional trust. Leadership should evaluate AI cybersecurity by the maximum credible operational consequence, not by whether a pilot is “small.” A pilot with production PHI, privileged APIs or broad document access can create enterprise exposure.

Stage	What happens	Leadership / enterprise outcome
1. Unknown exposure	A clinician, department or vendor connects data to AI outside the approved architecture.	CIO/CISO cannot attest to where ePHI went, what was retained, or which controls applied.
2. Initial compromise	Attacker exploits a vulnerable edge service, steals a credential, socially engineers support or compromises a third party.	Security team starts with incomplete asset and identity context; attacker dwell time expands.
3. AI-assisted acceleration	Automation maps systems, crafts lures, analyzes defenses or adapts tooling; an agent or RAG store may create additional paths.	Human-paced SOC triage cannot match machine-paced exploitation and exfiltration.
4. Clinical disruption	Systems are isolated or encrypted; EHR, lab, imaging, pharmacy, registration or interfaces degrade.	Diversion, delays, cancellations, paper workflows, safety events and regional spillover.
5. Data extortion	PHI, employee, financial or research data is stolen and threatened with publication.	Breach notification, patient harm, litigation, credit monitoring, reputational damage and negotiation costs.
6. Regulatory scrutiny	OCR asks whether the entity performed an accurate and thorough risk analysis and implemented reasonable safeguards.	Corrective action plans, monitoring, settlement costs and years of remediation. OCR’s 2026 OSF action is a direct example. [5]
7. Board / executive accountability	Post-incident reviews expose undocumented exceptions, delayed remediation, ignored audit findings or ungoverned pilots.	Loss of confidence in technology leadership; budget freezes or forced investment; leadership changes are possible where governance failures are material.
8. Strategic drag	Insurers, partners, lenders, ratings agencies and acquisition counterparties reassess cyber maturity.	Higher insurance friction, contracting constraints, capital demands and slower digital transformation.

It is important to distinguish accountability from hindsight. A CIO is not personally at fault because a sophisticated attacker succeeds. The leadership risk rises when known material gaps are not documented, prioritized, funded or escalated; when AI systems are allowed to bypass established risk processes; or when the organization cannot show that reasonable controls and recovery planning existed before the event.

OCR’s 2026 ransomware settlements reinforce the governance point: risk analysis is not a paperwork exercise performed after a breach. In the OSF Healthcare settlement, OCR cited an inadequate risk analysis, impermissible disclosure of PHI and delayed breach notifications. [5]

A 30 / 90 / 180-Day Hospital Action Plan

Time	Required actions	Board-level evidence
First 30 days	Freeze uncontrolled PHI uploads to public/unsanctioned AI; create enterprise AI inventory; identify all production AI with PHI or privileged API access; require named executive and technical owners; enforce phishing-resistant MFA for privileged/admin and remote vendor access; review exposed edge systems against CISA KEV; validate immutable backups; identify top 20 mission-critical third parties; issue interim AI acceptable-use rules.	Signed AI risk directive; inventory baseline; list of exceptions; recovery test evidence; critical-vulnerability aging; third-party dependency map.
Days 31–90	Implement AI gateway/DLP for approved GenAI; map data flows and retention; test RAG permissions; inventory service accounts/API keys; red-team highest-risk models and agents; add AI events to SIEM/SOC; establish help-desk identity verification resistant to voice/deepfake pretexting; tabletop ransomware + AI/data-exfiltration scenario; test 72-hour downtime by clinical service line.	Risk register with owners/dates; red-team findings; SOC use cases; tabletop after-action report; clinical downtime readiness score.
Days 91–180	Segment critical clinical/biomedical networks; deploy/expand IoMT/CPS visibility; implement agent runtime controls and tool allowlists; formalize AI procurement security addendum; require vendor evidence for model/data handling and subprocessors; test vendor outage and data-return scenarios; measure recovery time against clinical priorities; conduct board cyber-AI exercise; tie unresolved high risk to capital plan.	Independent control validation; recovery-time evidence; third-party test results; board exercise; documented residual-risk acceptance.

The action plan below converts the risk themes in the Black Book State of Global Healthcare Cybersecurity 2026 and US Hospital Cyber Resilience 2026 into an AI-specific operating sequence. The hospital should establish one non-negotiable gate: no production AI with PHI, clinical data or action-taking capability without a documented security architecture and recoverability plan. The gate should be jointly owned by the CIO/CISO, privacy/security officer, clinical owner and enterprise risk function—not procurement alone.

A practical approval packet for each AI use case should contain: owner and use case; data classes; architecture and data-flow diagram; vendor/subprocessors; identity and privileges; retention/training terms; logging; prompt/RAG/agent controls; test results; failure mode; clinical downtime procedure; incident notification contacts; exit/data-deletion plan; and residual risk acceptance.

Do not wait for a "perfect AI policy." Hospitals already know enough to act. The immediate objective is to stop unknown data movement, reduce identity compromise, close exploitable vulnerabilities, constrain agent permissions, map critical dependencies and prove that clinical operations can continue during a prolonged cyber outage.

8. Board, CIO and CISO Operating Model: The Questions That Must Be Answered Every Month

Metric / question	Target state
Do we know every AI system touching PHI or critical operations?	100% of sanctioned production AI registered; shadow AI discovery continuously monitored.
Can we prove where PHI enters, flows, is retained and leaves?	Documented data lineage for every high-risk AI use case; no unexplained external destinations.
Are privileged AI and vendor accounts phishing-resistant?	100% phishing-resistant MFA for privileged/admin and remote vendor access; exceptions time-limited and approved.
How many critical exploited vulnerabilities are beyond SLA?	Zero unaccepted critical KEV exposure; compensating controls documented where patching is clinically constrained.
Can an AI agent take a high-impact action by itself?	Only where explicitly approved; tool allowlists, transaction caps, runtime logging and rollback are enforced.
Can the SOC see and investigate AI events?	AI prompts/data alerts, agent actions, identity anomalies and runtime violations feed centralized monitoring with tested playbooks.
Can we operate without EHR/cloud/AI dependencies?	Department-specific procedures exercised; 72-hour minimum operational continuity tested, with longer-duration plans for critical services.
Can a third-party failure stop care?	Critical vendor concentration mapped; alternatives/manual procedures/contractual obligations tested.
Do we have immutable recovery that actually works?	Restores tested by system and dependency order; RTO/RPO tied to clinical impact, not only IT convenience.
Who accepts residual AI cyber risk?	Named executive/business owner signs; material risks escalated to board risk/audit oversight.

Across Black Book's 2026 cybersecurity research, the recurring management requirement is evidence: hospitals must be able to show what is connected, who can access it, how third parties are controlled, what can be recovered, and whether clinical operations can continue. The governing principle is straightforward: AI adoption speed cannot exceed security observability and recovery capability. When it does, the hospital is borrowing operational risk from the future.

Boards should expect the CIO and CISO to describe cyber-AI risk in clinical and financial terms: "Which services stop? For how long? What data can leave? Who can act through an agent? Which third party is a single point of failure? What is the maximum tolerable downtime? What do we know, and what remains unknown?" A dashboard of blocked attacks is not a substitute for these answers.

This also reframes budgeting. Hospitals do not necessarily need another standalone AI-security product for every use case. They need coverage across a control architecture: identity, data security, endpoint/network, cloud, AI runtime, cyber-physical/medical devices, third-party risk, SOC/response and recovery. Product consolidation is useful only if control coverage and clinical safety improve.

Black Book readiness test If a hospital cannot identify every production AI system, trace its regulated data, explain its privileges, show its runtime controls, name its vendor dependencies and demonstrate downtime/recovery procedures, the AI deployment is

not cyber-ready regardless of its clinical or operational promise.

Black Book 2026 Cybersecurity Research Continuum

This hospital AI cybersecurity action report should be read as a focused extension of Black Book's broader 2026 cybersecurity research program. The State of Global Healthcare Cybersecurity 2026 establishes the cross-market threat and resilience baseline; US Hospital Cyber Resilience 2026 concentrates on identity, third-party exposure, insurance and incident readiness; US Hospital Cyber Readiness 2026 addresses operational preparedness; and the present report isolates the additional control failures created when hospitals introduce generative AI, machine learning, RAG, agents and AI-enabled third-party services into clinical and administrative workflows.

The unifying Black Book position is that cybersecurity maturity can no longer be measured by tool ownership or policy completeness. For hospitals, maturity is the ability to prove—under real operating conditions—that critical assets are known, identities are constrained, third parties can be disconnected, AI data flows are traceable, clinical workflows can continue through disruption, and recovery can be executed within clinically tolerable timeframes.

2026 Black Book research title	Primary hospital question	Connection to this report
State of Global Healthcare Cybersecurity 2026	How is healthcare cyber risk changing across threats, vendors, governance and resilience?	Provides the overarching threat, governance and vendor-market context.
US Hospital Cyber Resilience 2026	Can a hospital maintain care, control identity and upstream dependencies, and recover under attack?	Supplies the resilience lens for AI-origin and AI-amplified events.
US Hospital Cyber Readiness 2026	Are controls, ownership, evidence and response procedures operational now?	Supports the readiness test, operating metrics and executive accountability model.
Hospital AI Cybersecurity Readiness — this report	Can hospitals safely deploy AI without creating unobserved PHI, identity, agent, API and third-party attack paths?	Adds AI inventory, RAG, agent, model, prompt, runtime and AI-supply-chain controls.

Healthcare Cybersecurity Vendor Directory for US Hospitals

The directory below is a procurement and market-navigation reference for hospitals and health systems. Vendors are organized by the primary capability a hospital buyer is likely to evaluate and are listed alphabetically within each subsection. Many vendors span multiple categories; placement is intended to help buyers build an apples-to-apples shortlist, not to imply exclusivity. Inclusion is not a Black Book ranking, award, endorsement or statement of client satisfaction. Hospitals should validate healthcare references, interoperability, implementation burden, data handling, BAAs where applicable, service-level commitments, downtime behavior, clinical-device safety and current product ownership before purchase.

Black Book's prior healthcare cybersecurity editions catalogued hundreds of vendors and dozens of functional categories. This condensed 2026 directory emphasizes providers most applicable to acute-care hospitals, health systems, connected clinical environments, AI-enabled workflows and managed cyber operations.

10.1 Healthcare Cybersecurity Advisory, Risk, MDR and Managed Security Services

Vendor	Core capability	Hospital applicability
Arctic Wolf	Managed detection and response, security operations and incident preparedness	Hospitals seeking an outsourced or co-managed SOC operating model.
Avertium	Managed security, risk, compliance, penetration testing and incident response	Health systems needing advisory plus operational security services.
Clearwater	Healthcare-focused cyber risk, HIPAA, compliance, managed security and advisory	Provider organizations needing healthcare-specific risk and compliance depth.
Deloitte Cyber	Cyber strategy, transformation, identity, cloud, resilience and incident services	Large IDNs and academic systems undertaking enterprise cyber transformation.
Fortified Health Security	Healthcare-focused managed security, advisory, incident response and third-party risk	Hospitals wanting a provider-specialist security partner and managed program.
GuidePoint Security	Security advisory, architecture, managed services and technology evaluation	Complex health systems needing independent architecture and vendor-selection support.
Kroll	Incident response, digital forensics, cyber risk, managed detection and threat intelligence	Hospitals preparing for or responding to material cyber incidents.
Lumifi	Managed detection and response, vulnerability identification and healthcare security services	Hospitals seeking MDR with healthcare-specific service packaging.
Mandiant (Google Cloud)	Threat intelligence, incident response, red teaming and managed defense	Health systems requiring high-end incident response and threat expertise.
Optiv	Security advisory, managed services, architecture, identity, cloud and risk	Large hospitals consolidating multiple security workstreams.
Pondurance	MDR, incident response, vulnerability management and compliance services	Midmarket and regional hospitals needing managed security coverage.

PwC Cybersecurity	Cyber strategy, risk, privacy, cloud, resilience and transformation	Large health systems integrating cybersecurity with enterprise risk and audit.
Unit 42 by Palo Alto Networks	Incident response, threat intelligence, proactive services and digital forensics	Hospitals needing rapid breach response tied to a broad security platform.
Verizon Business Security	Managed security, threat intelligence, network security and breach expertise	Distributed health systems with large network and communications footprints.

10.2 Identity, Privileged Access and Access Governance

Vendor	Core capability	Hospital applicability
BeyondTrust	Privileged access management, remote access and identity security	Vendor remote access, administrator privilege and service-account control.
CyberArk	Privileged access, secrets, workforce identity and machine identity	High-risk privileged users, non-human identities and application secrets.
Delinea	Privileged access management and credential governance	Hospitals reducing standing privilege across servers and applications.
Duo Security (Cisco)	Phishing-resistant MFA, device trust and zero-trust access	Clinician and workforce access with contextual device controls.
HID	Identity, authentication and credential technologies	Physical/digital identity convergence and strong authentication.
Imprivata	Healthcare identity, access, single sign-on, privileged access and clinical workflow security	Shared clinical workstations, fast clinician access and healthcare identity workflows.
Microsoft Entra	Identity, conditional access, workload identity and identity governance	Microsoft-centric hospitals standardizing workforce and workload identity.
Okta	Workforce identity, SSO, MFA and identity governance	Cloud/SaaS access and enterprise workforce authentication.
Ping Identity	Identity orchestration, authentication and access management	Complex patient, workforce and partner identity environments.
RSA	Authentication, identity governance and access assurance	Hospitals with legacy and hybrid identity requirements.
SailPoint	Identity governance and administration	Role, entitlement and lifecycle governance across large application estates.
Saviynt	Identity governance, application access and privileged access	Cloud-first identity governance and access certification.
SecureAuth	Adaptive authentication and identity security	Risk-based authentication in mixed legacy/cloud environments.
Semperis	Active Directory and Entra identity threat detection and recovery	Hospitals prioritizing identity-system resilience and ransomware recovery.

Silverfort	Identity protection and MFA extension across legacy and nonstandard resources	Protecting legacy clinical systems and service accounts that resist normal MFA.

10.3 Medical Device, IoMT, OT and Cyber-Physical Security

Vendor	Core capability	Hospital applicability
Armis	Agentless asset intelligence, exposure management and medical-device security	Enterprise visibility across IT, IoMT, OT and unmanaged connected assets.
Asimily	IoMT inventory, risk analytics and medical-device vulnerability management	Hospitals prioritizing clinical-device risk by exploitability and patient impact.
Claroty	Healthcare cyber-physical systems security, IoMT visibility and clinical-context risk	Large connected-care environments requiring IT/OT/IoMT visibility and segmentation insight.
Cynerio	Healthcare IoT security, asset visibility, risk reduction and segmentation support	Hospitals seeking healthcare-specific connected-device security.
Elisity	Identity-based microsegmentation for connected and clinical environments	Health systems enforcing least-privilege network access without endpoint agents.
Forescout	Asset intelligence, network access control and cyber-physical exposure management	Mixed IT/IoT/OT environments requiring discovery and policy enforcement.
MedCrypt	Medical-device product security, cryptography and vulnerability management	Hospitals and device manufacturers addressing device-level security and SBOM needs.
Nozomi Networks	OT/IoT visibility, anomaly detection and cyber-physical security	Hospitals with facilities, utilities and operational technology exposure.
Ordr	Connected-device discovery, behavior analytics, risk and segmentation workflows	Health systems managing large numbers of unmanaged clinical and IoT assets.
Phosphorus	Cybersecurity for the extended IoT and xIoT device estate	Credential, firmware and device-hardening needs across unmanaged assets.
TRIMEDX	Clinical asset management and medical-device cybersecurity managed services	Hospitals combining biomedical engineering, device inventory and cybersecurity operations.

10.4 Endpoint, Network, XDR, Zero Trust and Security Operations Platforms

Vendor	Core capability	Hospital applicability
Check Point Software	Network, cloud, endpoint, email and threat-prevention security	Hospitals consolidating network and threat-prevention controls.
Cisco Security	Network security, Duo identity, XDR, Secure Access and AI Defense	Cisco-heavy health systems pursuing integrated network, identity and AI controls.
CrowdStrike	Endpoint/XDR, identity security, cloud security, threat intelligence and agentic SecOps	Hospitals seeking rapid endpoint detection with broad threat visibility.

Fortinet	Firewalls, network security, zero trust, endpoint and security operations	Large distributed health systems and hospital campuses.
Microsoft Security	Defender XDR, Sentinel SIEM, Entra identity, Purview data security and cloud protection	Microsoft-centric providers seeking integrated security telemetry and response.
Netskope	Secure access service edge, cloud security, DLP and GenAI usage controls	Hospitals governing cloud, SaaS, web and AI data movement.
Palo Alto Networks	Network security, Prisma Cloud, Cortex XSIAM, AI security and Unit 42 services	Large health systems seeking platform consolidation across network, cloud and SOC.
SentinelOne	Endpoint/XDR, cloud security and AI-assisted/agentic security operations	Hospitals seeking autonomous investigation and endpoint/cloud response.
Sophos	Endpoint, firewall, MDR and incident response	Midmarket hospitals wanting integrated endpoint/network protection plus MDR.
Trend Micro	Endpoint, cloud, email and XDR security	Hospitals with mixed endpoint and cloud workloads.
Trellix	XDR, endpoint, email and data security	Enterprises maintaining broad threat detection and endpoint controls.
Zscaler	Zero-trust access, secure web, DLP, cloud and AI security	Hospitals replacing VPN-centric access and controlling internet/cloud/AI traffic.

10.5 Data Security, DLP, Privacy, Email and Human-Layer Defense

Vendor	Core capability	Hospital applicability
Abnormal Security	AI-driven email security and account-takeover detection	Hospitals facing business email compromise and socially engineered attacks.
Barracuda Networks	Email, application, network and data protection	Providers needing email and application-layer protection.
BigID	Data discovery, classification, privacy and security posture management	Hospitals locating sensitive data across structured and unstructured repositories.
Broadcom Symantec	Enterprise DLP, endpoint and information protection	Large providers with mature enterprise DLP programs.
Cyberhaven	Data lineage and data-loss prevention for endpoints and cloud workflows	Organizations needing context on how sensitive data moves into AI and SaaS.
Forcepoint	Data security, DLP, insider risk and secure access	Hospitals emphasizing data-centric policy enforcement.
IRONSCALES	Cloud email security and phishing protection	Healthcare organizations targeting mailbox-based social engineering.
Mimecast	Email security, continuity and human-risk management	Hospitals protecting email while preserving communications during disruption.
Nightfall AI	Cloud-native DLP and sensitive-data detection for SaaS and GenAI	Organizations monitoring regulated data in modern cloud collaboration and AI workflows.

Proofpoint	Email security, DLP, insider threat and human-risk management	Large health systems focused on phishing, data loss and high-risk users.
Securiti	Data intelligence, privacy, governance and AI governance	Hospitals combining data governance, privacy and AI oversight.
Varonis	Data security posture, access analytics and threat detection	Health systems controlling excessive access to files, cloud data and collaboration stores.

10.6 Cloud, Application, API and AI Security

Vendor	Core capability	Hospital applicability
Aqua Security	Cloud-native application protection, containers and software supply-chain security	Hospitals running cloud-native apps, Kubernetes and DevSecOps pipelines.
CalypsoAI	AI security testing, governance and runtime controls	Organizations evaluating model and generative-AI security controls.
Cisco AI Defense	AI discovery, validation, runtime protection and agent/MCP controls	Hospitals building an explicit AI security layer across applications and agents.
Cloudflare	Web application, API, zero-trust and network edge security	Internet-facing hospital applications, APIs and distributed access.
HiddenLayer	AI/ML model detection, monitoring and model-security controls	Health systems protecting internally developed or deployed ML assets.
Microsoft Defender for Cloud	Cloud security posture and workload protection	Azure and multi-cloud healthcare workloads.
Palo Alto Networks Prisma Cloud / AI security	Cloud-native security, application security and AI runtime controls	Large cloud estates and AI-enabled applications needing unified policy.
Snyk	Developer security, open-source, container, code and IaC scanning	Hospital software teams and digital-health development programs.
Sysdig	Cloud-native detection, posture and runtime security	Containerized healthcare workloads and Kubernetes environments.
Wiz	Cloud security posture, exposure and workload risk	Multi-cloud hospitals seeking rapid risk visualization and prioritization.
Zscaler AI Security	Shadow-AI discovery, AI gateway, DLP and agent/MCP controls	Hospitals governing enterprise generative-AI access and data leakage.

10.7 Third-Party, Supply-Chain and Vendor Cyber Risk

Vendor	Core capability	Hospital applicability
Bitsight	External cyber ratings and third-party risk monitoring	Health systems monitoring large vendor ecosystems.
Black Kite	Third-party cyber risk intelligence and ransomware susceptibility analysis	Provider supply-chain and vendor-risk prioritization.

Censinet	Healthcare-focused third-party cyber risk and risk operations	Hospitals integrating cyber risk assessment with healthcare procurement.
Exiger	Supply-chain risk, entity intelligence and software supply-chain analysis	Hospitals assessing upstream technology and supplier exposure.
Mastercard RiskRecon	Third-party cyber risk monitoring and external security assessment	Continuous supplier and partner cyber posture monitoring.
ProcessUnity	Third-party risk management workflows and assessment automation	Large health systems formalizing vendor governance and evidence.
SecurityScorecard	Cyber ratings, third-party risk and attack-surface intelligence	Enterprise supplier monitoring and board-level external risk visibility.
UpGuard	Vendor risk, attack-surface monitoring and security questionnaires	Hospitals requiring continuous vendor monitoring and workflow automation.
Whistic	Third-party risk assessment and security-profile exchange	Procurement teams accelerating vendor security reviews.

10.8 Vulnerability, Exposure, Asset and Attack-Surface Management

Vendor	Core capability	Hospital applicability
Axonius	Cyber asset management and security control validation	Hospitals reconciling fragmented asset inventories across security tools.
Qualys	Vulnerability management, policy compliance and cloud/asset security	Large heterogeneous infrastructure requiring continuous vulnerability assessment.
Rapid7	Exposure management, vulnerability management, detection and response	Hospitals integrating vulnerability and detection workflows.
runZero	Asset discovery and exposure visibility	Networks with unknown, unmanaged or difficult-to-scan devices.
Tenable	Exposure management, vulnerability management, identity and OT security	Health systems prioritizing exploitable risk across IT and OT.
XM Cyber	Continuous exposure management and attack-path analysis	Hospitals identifying paths from initial compromise to critical assets.

10.9 Backup, Recovery, Cyber Resilience and Ransomware Recovery

Vendor	Core capability	Hospital applicability
Cohesity	Data security, backup, recovery and cyber resilience	Large health systems consolidating backup and recovery operations.
Commvault	Backup, recovery, cyber resilience and cloud data protection	Hybrid hospital environments requiring broad workload recovery.
Dell Technologies	Data protection, cyber recovery vault and infrastructure resilience	Hospitals with substantial Dell infrastructure and critical recovery requirements.

Druva	SaaS-delivered data protection and cyber resilience	Cloud-oriented providers protecting SaaS, endpoint and infrastructure data.
HYCU	Multi-cloud and SaaS backup and recovery	Hospitals needing protection across diverse SaaS and cloud applications.
Rubrik	Cyber resilience, immutable recovery, identity recovery and data security	Hospitals prioritizing ransomware recovery and clean restoration.
Veeam	Backup, replication and ransomware recovery	Broad virtualized, cloud and physical workload recovery.

10.10 Incident Response, Digital Forensics and Breach Recovery

Vendor	Core capability	Hospital applicability
CrowdStrike Services	Incident response, compromise assessment and threat hunting	Material endpoint, identity and cloud incidents.
Kroll	Digital forensics, incident response, crisis management and breach support	Complex breaches requiring forensic, legal and communications coordination.
Mandiant	Incident response, threat intelligence and advanced adversary investigation	Sophisticated intrusions and ransomware events.
Microsoft Incident Response	Incident investigation and recovery across Microsoft environments	Microsoft-centric identity, cloud and endpoint incidents.
Sophos Incident Response	Ransomware and breach response tied to endpoint/MDR capabilities	Midmarket and regional provider breach containment.
Unit 42 by Palo Alto Networks	Incident response, ransomware response and threat intelligence	Large-scale attacks requiring rapid containment and forensic expertise.

Directory reference basis

The directory was assembled from Black Book's healthcare cybersecurity market taxonomy and 2026 research library, current healthcare-focused vendor and industry program information, and public product/service descriptions available in August 2026. Useful cross-checks include the American Hospital Association Preferred Cybersecurity & Risk Provider Program and current healthcare cybersecurity market directories. Product portfolios, ownership and packaging change frequently; hospital buyers should verify current capabilities directly with each vendor.

Sources and Reference Notes

[0A] Black Book Research, State of Global Healthcare Cybersecurity 2026. https://blackbookmarketresearch.com/state-of-global-healthcare-cybersecurity-2026

[0B] Black Book Research, US Hospital Cyber Resilience 2026. https://blackbookmarketresearch.com/us-hospital-cyber-resilience-2026

[0C] Black Book Research, 2026 Healthcare IT Research Library: Cybersecurity, Resilience & Downtime portfolio. https://blackbookmarketresearch.com/

[0D] Black Book Research, The Black Book of Healthcare Cybersecurity, 2025 Edition; includes functional taxonomy and alphabetical vendor directory. https://blackbookmarketresearch.com/uploads/pdf/399U5-TBB-Healthcare-Cybersecurity-v2.pdf

[0E] American Hospital Association, Preferred Cybersecurity & Risk Provider Program. https://www.aha.org/center/cybersecurity-and-risk-advisory-services/preferred-cybersecurity

[1] Verizon, 2026 Data Breach Investigations Report: Healthcare Snapshot. Healthcare dataset: 1,492 incidents, 1,438 confirmed data disclosures; System Intrusion, Miscellaneous Errors and Social Engineering = 81%; external actors 81%; financial motive 99%; third-party 32%; human element 54%. Also includes broader 2026 findings on vulnerability exploitation, ransomware, shadow AI and GenAI-assisted attacks. https://www.verizon.com/business/resources/reports/2026-dbir-healthcare-snapshot.pdf

[2] Netskope Threat Labs, Healthcare 2026. Reports regulated data accounted for 89% of healthcare GenAI data-policy violations and details shadow AI, API use and cloud-app risk. https://www.netskope.com/resources/threat-labs-reports/threat-labs-report-healthcare-2026

[3] Health-ISAC, 2026 Global Health Sector Threat Landscape. AI-enabled attacks ranked as the #1 concern for 2026; ransomware and supply-chain vulnerabilities remain central. https://health-isac.org/annual-threat-report-health-sector-2026/

[4] Health-ISAC, 2026 CISO Benchmarking Report, based on 76 health-sector CISOs. https://health-isac.org/2026-health-isac-ciso-benchmarking-report/

[5] HHS OCR, July 29, 2026, ransomware settlement with OSF Healthcare System; OCR cited risk-analysis deficiencies, PHI disclosure and breach-notification issues and required a corrective action plan. https://www.hhs.gov/press-room/hhs-ocr-settles-ransomware-investigation-with-healthcare-system.html

[6] NIST AI 600-1, Generative Artificial Intelligence Profile, updated 2026. https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence

[7] NIST AI 100-2e2025, Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations. https://www.nist.gov/news-events/news/2025/03/nist-trustworthy-and-responsible-ai-report-adversarial-machine-learning

[8] NIST IR 8596 / Cyber AI Profile work and August 2026 workshop reports: secure AI systems, AI-enabled defense and AI-enabled attacks. https://csrc.nist.gov/pubs/ir/8596/iprd

[9] JAMA Network Open, Ransomware Attack Associated With Disruptions at Adjacent Emergency Departments in the US. https://jamanetwork.com/journals/jamanetworkopen/fullarticle/2804585

[10] HHS Healthcare and Public Health Cybersecurity Performance Goals and Hospital Resiliency Landscape Analysis. https://hhscyber.hhs.gov/cybersecurity-performance-goals.html

[11] Microsoft Security, “Rethinking security for the age of AI,” July 27, 2026; Project Perception. https://blogs.microsoft.com/blog/2026/07/27/rethinking-security-for-the-age-of-ai/

[12] Palo Alto Networks, Cortex XSIAM 2026 data sheet. https://www.paloaltonetworks.com/resources/datasheets/cortex-xsiam

[13] Palo Alto Networks, “The Frontier AI SOC Has Arrived,” June 11, 2026. https://www.paloaltonetworks.com/blog/security-operations/cortex-frontier-ai/

[14] CrowdStrike, Charlotte AI Data Sheet. https://www.crowdstrike.com/en-us/resources/data-sheets/charlotte-ai/

[15] SentinelOne, Purple AI Agentic Investigation announcement, June 17, 2026. https://www.sentinelone.com/press/sentinelone-opens-purple-ai-agentic-investigation-to-all-customers-bringing-frontier-ai-directly-into-the-soc/

[16] Cisco AI Defense Data Sheet, 2026; AI cloud visibility, supply-chain scanning, validation, runtime and MCP/agent controls. https://www.cisco.com/c/en/us/products/collateral/security/ai-defense/ai-defense-ds.html

[17] Cisco AI Runtime Protection, 2026. https://www.cisco.com/site/us/en/products/security/ai-defense/ai-runtime/index.html

[18] Zscaler AI Security, 2026; AI asset management, AI gateway, runtime, agent/MCP controls. https://www.zscaler.com/products-and-solutions/ai-security

[19] Zscaler Secure the Use of Generative AI; prompt visibility, inline DLP and browser isolation. https://www.zscaler.com/products-and-solutions/securing-generative-ai

[20] Claroty xDome for Healthcare; cyber-physical systems, medical-device visibility and clinical-context risk prioritization. https://claroty.com/healthcare-cybersecurity/xdome-for-healthcare

[21] Armis Centrix for Medical Device Security; agentless discovery, clinical context and exposure management. https://www.armis.com/platform/armis-centrix-for-medical-device-security/

[22] Ordr for Healthcare, 2026; behavior-based device discovery, risk prioritization and segmentation. https://ordr.net/resources/ordr-for-healthcare

[23] HHS OCR, HIPAA Security Rule NPRM and 2026 ransomware enforcement actions. https://www.hhs.gov/hipaa/for-professionals/security/hipaa-security-rule-nprm/index.html

Scope note: Vendor descriptions and directory entries summarize publicly described capabilities and Black Book market taxonomy as of August 2026. They are not Black Book rankings, awards or endorsements. Hospitals should independently validate current ownership, product scope, healthcare references, interoperability, security architecture, contractual protections and clinical-operational fit.

About Black Book Research

Black Book Research is an independent healthcare-centric market research and competitive-intelligence organization focused on the experience of healthcare technology and managed-services users. For more than two decades, Black Book has surveyed hospitals, health systems, physician organizations, payers and other healthcare stakeholders on technology adoption, vendor performance, operational priorities, cybersecurity, digital transformation, revenue cycle, analytics, interoperability, outsourcing and emerging innovation.

Black Book's healthcare research is designed to inform buyers and industry stakeholders without vendor sponsorship, vendor control of survey responses, or pay-to-play influence over client satisfaction findings. Its vendor-performance programs use structured qualitative key performance indicators and validated user feedback to compare technologies and services on an apples-to-apples basis within defined market segments.

Cybersecurity has been a recurring Black Book research domain as healthcare organizations have moved from perimeter-oriented security toward enterprise resilience, third-party risk, identity protection, connected-device security, incident readiness and clinical continuity. The 2026 cybersecurity research portfolio includes the State of Global Healthcare Cybersecurity 2026, US Hospital Cyber Resilience 2026, US Hospital Cyber Readiness 2026, European Healthcare Cyber Resilience and Readiness 2026, and related hospital cyber-exposure and downtime research.

This report extends that body of work into AI-enabled hospital operations, where generative AI, machine learning, RAG, agents, APIs, cloud platforms and embedded vendor intelligence introduce new data paths and trust boundaries. Black Book's objective is to give healthcare leaders a practical, vendor-neutral framework for evaluating those risks while preserving innovation, clinical continuity and patient safety.

Black Book Research | Healthcare IT Market Research and Competitive Intelligence | www.blackbookmarketresearch.com | research@blackbookmarketresearch.com