

HEALTHCARE SECURITY

A GUIDE TO OFFENSIVE TESTING



BISHOPFOX

Application Penetration Testing (CPT) builds on Cosmos Attack Surface Management to provide the highest level

of attack surface protection with post-exploitation activities.

WHO IS THIS GUIDE FOR?

- Healthcare CISOs and security program directors
- Clinical engineering and biomed leaders responsible for medical device and OT security decisions
- Compliance and risk officers
- Third-party and vendor risk leaders overseeing software and hardware embedded in clinical workflows

INTRODUCTION

The next stage of healthcare security maturity is evidence: proof that controls can interrupt the attack paths most likely to affect care delivery.

Healthcare security leaders need confidence that the controls protecting patient data and clinical operations will hold under real-world conditions. Compliance establishes an important foundation, but documented controls alone cannot show how an environment will respond when an attacker combines weaknesses across systems.

Offensive security testing helps answer that question. It examines where an attacker could gain access, how far that access could take them, and whether the controls intended to contain them work as expected. The resulting evidence helps leaders decide where remediation, deeper testing, and security investment are most needed.

BISHOP FOX ANALYSIS

Healthcare Security by the Numbers

Offensive security work completed over two years.

40+

HEALTHCARE CLIENTS

160+

OFFENSIVE SECURITY
ASSESSMENTS

825

INDIVIDUAL FINDINGS

IN THIS GUIDE

Drawing on 160+ assessments across 40+ healthcare clients, alongside case studies and documented incidents, this guide examines how controls hold up under realistic attacks. Readers will learn:

- Where risk develops across connected care environments.
- What testing reveals about defenses protecting patient data and care delivery.
- How testing methods and scope shape findings and inform what to test next.



HEALTHCARE HAS CHANGED

Two decades of compliance investment have not slowed how often attackers breach healthcare organizations.

CARE DELIVERY HAS OUTGROWN TRADITIONAL BOUNDARIES

For decades, healthcare security programs have been organized around protecting patient data and demonstrating compliance with requirements such as HIPAA and frameworks such as HITRUST. That work established an essential foundation for privacy, governance, and control design. Much of it developed around systems and organizational boundaries that were easier to define than those supporting care today.

A single patient journey may now touch an electronic health record, a digital front-door application, a connected infusion pump, a third-party laboratory API, and cloud infrastructure before a clinician sees the patient. Each connection creates a dependency, and the access supporting those connections can extend beyond the healthcare organization's direct control.

Medical devices introduce operational constraints that conventional IT security practices may not accommodate. Vendors can hold access to systems embedded in clinical workflows. AI-enabled applications can introduce additional data access and actions through existing integrations. Evaluating the security of each component individually leaves questions about how an attacker could move between them.

Protecting patient information therefore requires protecting the systems and relationships that make care possible. Confidentiality remains essential, while availability and integrity determine whether clinicians can access information, trust it, and continue delivering care.

CYBERATTACKS PUT CONTINUITY OF CARE AT RISK

Security architecture, staffing, and testing cadence do not always evolve at the same pace as that environment. The gap widens with each integration, connected device, and externally managed service. An organization can execute well against its regulatory obligations while still lacking evidence about whether its controls can interrupt a realistic attack path.

The [Department of Health and Human Services' Office for Civil Rights](#) logged 772 large healthcare data breaches in 2025, a record high, and hacking has driven more than 80% of those large breaches every year since 2023. [IBM's Cost of a Data Breach Report](#) has ranked healthcare the most expensive industry to breach for 13 consecutive years, at \$6.64 million per incident in its 2026 report. Those figures describe an industry absorbing record breach volume despite two decades of HIPAA-compliance investment, which is exactly the gap this guide sets out to explain.

For healthcare leaders, the security question also extends to what happens during a compromise: which clinical services become unavailable, whether affected systems remain trustworthy, and how long care can continue under disrupted conditions.

THE MODERN HEALTHCARE ATTACK SURFACE

Most attack-surface conversations are asset inventories, but healthcare's attack surface is a web of trust relationships that an asset inventory cannot see.

Servers, applications, and endpoints: the traditional attack-surface conversation is an inventory problem, and it made perfect sense when trust boundaries were mostly physical, like a hospital network, or organizational, like a single IT department. In modern healthcare organizations, security leaders operate in a fundamentally different environment, one built less on boundaries and more on relationships.

Today's healthcare attack surface is better understood as a lattice of trust relationships: clinical device to hospital network, hospital to cloud identity provider, hospital to third-party vendor, vendor to vendor. An asset inventory cannot see that web, and four vectors in particular are expanding it faster than the rest.



Figure 1. The healthcare attack surface as trust relationships.

IDENTITY IS THE NEW PERIMETER

The network perimeter used to be what a healthcare security program defended, back when defending it meant defending a place. In an environment built on cloud EHR platforms, remote clinical access, and third-party integrations, a single perimeter rarely exists to defend, and what exists instead is a lattice of identity relationships that collectively functions as the perimeter.

That shift carries a practical consequence: identity governance (meaning who can authenticate as what, to which system, and under which conditions) now performs the job the firewall used to do, and any program still budgeted and staffed as though network segmentation is the primary control is defending the wrong boundary.

In practice, identity exposure shows up as over-provisioned service accounts, shared clinical-system credentials that often reflect a defensible tradeoff for speed of care, and third-party vendor accounts that quietly outlive the relationship that created them.

Credential and secrets-management issues of exactly this kind appear in roughly 25% of assessments in Bishop Fox’s dataset and half of all categorized Critical and High findings, further evidence that identity is where exposure actually concentrates.

Identity-Related Weaknesses Carry Disproportionate Risk




Figure 2. Identity-related categories and severity represented.

Inside a typical hospital system, a clinician authenticates once through single sign-on, but that single identity fans out into dozens of downstream application-to-application trust relationships: EHR to lab system, lab system to billing, billing to a third-party clearinghouse, each inheriting some portion of the original access without a fresh authentication event.

A single compromised credential at the top of that chain can cascade through every system that trusted it, and no firewall rule sits anywhere in that chain to stop it.




Figure 3. Illustrative overprovisioning of modern hospital environment.

MEDICAL DEVICES ARE OPERATIONAL TECHNOLOGY

Folding medical devices into general IT asset management misreads what they actually are. A connected infusion pump or imaging system shares little with a laptop beyond a network connection. It is operational technology, purpose-built and long-lived, with security and lifecycle constraints IT was never designed around. Since March 2023, the FDA has [required newly authorized devices](#) to include a cybersecurity management plan with a defined patching process as a condition of market clearance, but that requirement does not retroactively apply to devices already in service, which means most hospital device fleets still run a mix of newly compliant devices and older ones that don't have a formal update path.

None of this makes devices unfixable, but it does call for the discipline that operational technology deserves: segmentation appropriate to device criticality, compensating controls where native hardening is not possible, and validation that treats "can we reach this device from somewhere we shouldn't be able to" as a first-class question rather than an afterthought bolted onto a general network test.

A connected infusion pump illustrates the tension well. A device the FDA cleared before that requirement took effect can still trigger a recertification process that runs months, not days, for even a minor patch, because its original clearance never accounted for a formal update path. Security teams typically inherit a fleet spanning both eras, which means the discipline has to plan for legacy devices without a defined patch process alongside newer ones that have one, rather than assuming the current regulatory standard already covers everything already in the building.

The [FBI's Cyber Division](#) has found that 53% of networked medical devices carry at least one known critical vulnerability, and [separate research](#) puts the average at 6.2 distinct software flaws per device. As of 2025, roughly 1.2 million internet-connected healthcare devices and systems sit publicly accessible online, nearly 3x the exposure recorded in 2021. The stakes go beyond data breaches; [published research](#) has linked hospital cyberattacks rooted in medical device vulnerabilities to a 29% increase in inpatient mortality, and neighboring hospitals absorbing diverted patients saw an 81% surge in cardiac arrest cases during disruption. Bishop Fox's own client base reflects that scale directly: medical device environments account for roughly one in eight of the healthcare organizations in this dataset, a mainstream category of exposure rather than a rounding error.

MARKET SNAPSHOT

Medical Devices by the Numbers

1.2M

internet-connected healthcare
devices publicly accessible

53%

of networked medical devices contain
at least 1 known critical vulnerability

6.2 avg.

software vulnerabilities per medical
device

3x

increase in publicly exposed
healthcare devices since 2021

THIRD PARTIES EXPAND THE ATTACK SURFACE

A hospital's attack surface no longer ends at its own infrastructure. Lab-results integrations, billing clearinghouses, digital health platforms, staffing and scheduling vendors, and cloud infrastructure providers are each a trust relationship the organization depends on without fully controlling, and that dependency is precisely the point.

The healthcare-specific wrinkle is that these third parties often sit inside the clinical workflow rather than merely adjacent to it, so a compromised scheduling vendor can affect care delivery directly. That is a materially different risk than third-party exposure in most other sectors, where a vendor breach is primarily a data or reputational event.

Third-party risk management built around annual questionnaires and point-in-time attestations validates a moment rather than an ongoing relationship, which is exactly why offensive validation (i.e., testing where a third-party connection could be abused) answers a different question than a compliance questionnaire ever will.

A staffing platform is a useful illustration. It rarely appears on a healthcare organization's own asset inventory, since the organization neither owns nor operates it, yet it routinely holds credentials, scheduling data, and sometimes direct system access for contract clinicians moving through multiple facilities. A compromise there does not register as a healthcare breach on most risk registers, right up until it happens.

The [American Hospital Association's](#) analysis of federal breach data found that attackers stole more than 80% of the patient records exposed in recent years from the third-party vendors, business associates, and software platforms that hospitals depend on. The [2023 breach at Perry Johnson & Associates](#), a medical transcription services vendor, rippled out to more than a dozen of its healthcare clients and 13 million individuals combined.

AI IS EXPANDING THE SURFACE AGAIN, ON TOP OF THE REST

AI adoption in healthcare adds a new layer to existing identity, application, and third-party risk. Clinical documentation assistants, prior-authorization tools, and diagnostic-support systems often depend on broad data access and vendor integrations. Their security, therefore, depends on the same authorization boundaries, secrets, APIs, and trust relationships already carrying the rest of the environment.

For security leaders, AI is best treated as an accelerant and an overlay rather than an isolated attack-surface category:

- A model with broad access can magnify an identity weakness.
- An agent connected through a third-party API can extend the reach of that vendor relationship.
- A workflow allowed to act can convert inappropriate access into operational consequence more quickly.

An AI-powered clinical documentation assistant illustrates the intersection. It may read patient records through an enterprise identity, integrate through a vendor API, and draft or initiate actions within a clinical workflow. Thus, testing must examine both the AI-specific behavior and the software, authorization, and integration boundaries surrounding it.

As adoption moves faster than formal governance, security teams need a consistent way to bring AI-enabled workflows into existing architecture review, application testing, identity testing, and third-party risk processes. Treating AI as a lighter productivity review can leave the underlying access relationships untested.

These forces rarely operate independently. The most consequential exposure often appears where identity, software, devices, and third parties intersect. The next section follows how an attacker can move through those intersections and turn an initial weakness into operational leverage.

HOW MODERN ATTACKS UNFOLD

Attack chains in healthcare don't start where defenders expect, and they don't stop where a compliance checklist ends.

Most healthcare security narratives still describe attacks as ending in a data breach, a records dump, or a HIPAA notification letter, and that description is not wrong as much as incomplete. This section and the two that follow present what emerged from analyzing 825 findings across the 160+ assessments in Bishop Fox's dataset from the last two years. This includes how attacks actually unfold, why traditional programs miss them, and what closes the gap.

The most useful way to think about a modern healthcare attack is as a chain that runs straight through the four vectors from the previous section. The chain typically opens with an identity weakness, an exposed or over-permissioned credential often tied to a third-party or vendor account, and from there, lateral movement frequently exploits the lack of segmentation between IT and clinical OT or medical devices. The device itself may not be the target, and what follows is escalation into a position with real operational leverage like the ability to affect a system that care delivery depends on. That escalation does not distribute evenly across the four vectors, and the unevenness is the real finding.

Severity Concentrates at Identity & Lateral Movement




Figure 4. Severity conversion rate by finding category, across all Bishop Fox findings.

Let’s look at how this could play out in a real environment. An assessment team obtains a valid but over-permissioned service account credential, exposed through a misconfigured integration between a scheduling platform and the hospital network. That credential does not grant administrative access on its own, but it does grant enough lateral reach to discover a segment where the organization never fully separated clinical and corporate IT systems.

From there, escalation follows a familiar pattern: enumerate service accounts with excessive privilege, identify one with domain-level reach, and pivot into a position to affect systems well beyond the original point of entry. None of the individual steps requires a novel technique. The chain works because each weakness, taken alone, looks survivable.

RANSOMWARE & DATA EXTORTION AS AN OPERATIONAL PROBLEM

Ransomware shows why healthcare attacks must be understood through both operational and data impact. Attackers may encrypt systems, steal highly sensitive records, threaten disclosure, interrupt care, or combine these tactics. The leverage comes from healthcare’s dependence on available systems and the sensitivity of the information those systems hold.

In 2024, ransomware attackers stole an employee’s credentials at [Ascension](#), one of the largest nonprofit health systems in the country, through a phishing message, then moved laterally through the network, escalated privileges, and deployed ransomware across systems supporting more than 140 hospitals. Electronic health records and the patient portal went offline nationwide, hospitals diverted ambulances, and clinical staff reverted to paper-based processes for weeks before systems were fully restored, even as the attackers separately stole records belonging to more than 5.6 million patients and employees.

In 2025, encryption occurred in [34% of healthcare ransomware attacks](#), a five-year low, while data extortion attacks continue to rise. Attackers do not need to encrypt systems when the threat of disrupting essential services creates enough pressure to pay. In healthcare, even limited disruption can delay care, interrupt clinical workflows, and erode confidence in the systems clinicians depend on. Organizations must therefore be prepared to contain data theft and sustain clinical operations when attackers use disruption itself as leverage.

HOW FAR AN ATTACKER CAN PROGRESS

The red team and social engineering engagements included in our research documented 16 attack scenarios against healthcare organizations.

Each scenario tested whether a realistic technique could achieve a defined objective.

How Far Could Attackers Go?




Figure 5. Outcomes across 16 attack scenarios documented in red team and social engineering reports.

Direct external network breach attempts achieved partial success once in three scenarios and no full successes. Social engineering achieved full success in four of nine scenarios, with one additional partial success. All four assumed-breach scenarios made at least partial progress after beginning with a foothold, and two reached full success. The samples are small and directional, but the pattern reinforces the importance of testing containment after initial access.

The perimeter resisted direct technical compromise, but other paths produced access more often. Once testers began inside the environment, each scenario demonstrated some ability to progress. The leadership question is therefore broader than whether the perimeter holds: if access occurs through a user, vendor, or already-compromised system, how far can an attacker move before another control stops them? Next, we turn to why programs built around the traditional model consistently miss this chain.

WHY TRADITIONAL SECURITY PROGRAMS MISS RISK

Prevalence and severity are different axes. Most programs target the wrong one, and a passed audit often reports on the depth of the test rather than the depth of the risk.

Compliance reviews, scanners, and dashboards provide necessary visibility into control gaps and known weaknesses. Their results become misleading when finding frequency is treated as a proxy for consequence.

Across the 19 categories in Bishop Fox's data, prevalence and severity conversion were essentially unrelated. Figure 6 illustrates this gap:

- **Data exposure and storage:** 44% of assessments, but only 6% of categorized findings were Critical or High.
- **Access control and authorization:** Less common, but 43% of findings were Critical or High.
- **Credential and secrets management:** 27% of findings were Critical or High.

Frequency Does Not Predict Severity





Figure 6. Finding prevalence v. severity conversion across 19 categories. Bubble size represents categorized finding volume.

A program that prioritizes only what dominates the dashboard can direct attention toward volume while missing the weaknesses most capable of changing an attacker’s position.

SCOPE THEATER

Scope creates a second source of false confidence. A clean result can describe the boundaries of the assessment more accurately than the broader environment. If identities, vendor connections, clinical devices, or adjacent systems sit outside scope, the assessment cannot show whether an attacker could use them to continue an attack path.

Multi-scope assessments surfaced 9.8 findings on average, compared with 4.6 in single-scope assessments, and uncovered at least one Critical or High severity issue in 55% of engagements v. 34%. Broader coverage increases the chance that testing will reach the part of the environment where consequential risk is present.



THE ARCHITECTURE BEHIND THE BLIND SPOT

The blind spot comes from testing individual control domains without examining the relationships between them. Network and endpoint controls may work as configured while an over-permissioned vendor identity still reaches a clinical system. An application may pass a focused test while a connected service account creates a path beyond the tested boundary. No single tool or assessment type owns the complete question.

Programs often accumulate capable tools that each examine a different slice of the environment. The missing discipline is a way to ask cross-domain questions directly: Can a compromised scheduling account reach clinical infrastructure? Can a corporate foothold cross into a medical-device segment? Can a newly deployed AI workflow access more data or take more action than its intended role requires?

A well-run traditional program can therefore remain blind to a path that only becomes visible when the relationships are tested together. Offensive security provides the means to test those paths deliberately.

VALIDATION THROUGH OFFENSIVE SECURITY TESTING

Offensive security testing replaces assumed protection with evidence that security controls hold against the ways attacks realistically unfold.

Healthcare security programs have no shortage of controls. The challenge is knowing whether those controls work together when an attacker moves across identity systems, third-party connections, clinical technology, and other interdependent parts of the environment.

Traditional assessments often examine those components individually. Offensive security tests the relationships between them. It asks whether a compromised vendor credential can reach a clinical system, whether weak segmentation allows movement from corporate IT into a medical device environment, or whether an over-permissioned identity can turn limited access into operational control.

That distinction matters because the greatest risks in our dataset rarely appeared as isolated vulnerabilities. They emerged when an attacker could connect multiple weaknesses into a viable path through the environment.

WHAT TESTING ACTUALLY PROVES

Across 160+ assessments, only 3.8% of assessments in our dataset came back with no findings at all. The remaining 96% surfaced at least one issue, and more than a third, 36.7%, surfaced at least one Critical- or High-severity finding.

Those results do more than confirm that vulnerabilities exist. They demonstrate how frequently healthcare environments contain weaknesses capable of producing meaningful technical or business impact, even when the organization has established security controls and successfully completed required audits.

Testing Consistently Finds Real Things




Figure 8. Finding outcomes across all assessments

The value of offensive security is that it shows which issues an attacker can use, how far those issues can take them, and whether the controls intended to stop that progression actually hold. Each outcome gives the security team something a scanner, questionnaire, or compliance review cannot: direct evidence of how the environment responds under adversarial pressure.

SCOPE DETERMINES WHAT TESTING CAN PROVE

Meaningful offensive security testing must reflect how the healthcare environment actually operates. That means examining not only individual applications and network ranges, but also the identities connecting systems, third parties embedded in clinical workflows, medical devices with distinct security constraints, and AI-enabled tools gaining access to sensitive data and operational processes.

A narrowly scoped assessment can thoroughly test everything included and still miss the path an attacker would be most likely to use. This does not mean every assessment must cover the entire enterprise; each test should be designed around a security question the organization needs answered.

- Can an external attacker gain an initial foothold?
- Can a compromised identity reach a clinical system?
- Does segmentation contain movement between corporate and medical device environments?
- Can a trusted third-party connection become an entry point?

The scope should follow the attack path being examined, rather than the assets that are easiest to place within the boundaries of a test. Attackers aren't bound by scope, and limiting testing artificially can be counterproductive.


FROM FINDINGS TO SECURITY DECISIONS

A finding is most valuable when it changes a security decision.

Offensive security testing gives teams evidence they can use to prioritize remediation, adjust architecture, strengthen detection, or reconsider where existing controls provide less protection than expected. It connects technical weaknesses to plausible attack paths and operational consequences, helping leaders distinguish issues that are merely present from those that could materially affect care delivery.

That evidence has a shelf life. It reflects the systems, relationships, and controls present during the engagement. A new vendor connection, identity role, clinical application release, connected-device class, or AI-enabled workflow can create a path the previous assessment never examined.

Offensive security testing should therefore operate as an ongoing part of the security program. Deep assessments answer consequential questions at a moment in time, continuous visibility identifies changes between them, and targeted retesting confirms whether remediation continues to hold.



BUILDING TESTING INTO THE SECURITY PROGRAM

Continuous security testing means maintaining enough visibility and testing coverage to evaluate risk as the environment changes.

This matters in healthcare because every new application, vendor connection, identity relationship, medical device, or AI-enabled workflow can create an attack path that did not exist during the previous assessment. An annual penetration test still provides valuable evidence about a defined environment at a particular moment. Building offensive security into the broader security program allows that evidence to remain current as the environment evolves.

1. BUILD TESTING AROUND THE RISKS THAT MATTER MOST

Testing priorities should begin with the outcomes the organization most needs to prevent, such as unauthorized access to clinical systems, movement into medical device environments, disruption of care delivery, or compromise through a trusted third party.

Teams can then work backward to identify the systems and trust relationships an attacker would need to traverse. This organizes testing around realistic attack paths instead of treating applications, networks, identities, and vendors as disconnected assets.

A healthcare organization concerned about ransomware-driven disruption, for example, should examine whether a compromised user or vendor identity could move laterally, reach privileged systems, cross into clinical environments, and obtain enough access to interrupt operations.

The central question becomes: Have we tested whether an attacker can reach the outcome we are trying to prevent?

2. MATCH THE TESTING METHOD TO THE QUESTION

Different testing methods answer different security questions.

Application and network penetration testing identify exploitable weaknesses within a defined scope. Red team exercises evaluate whether weaknesses across systems can be chained toward a larger objective. Social engineering tests whether attackers can gain access through people and processes. Continuous attack-surface testing identifies new exposures and verifies whether they create exploitable risk.

Healthcare environments also require methods appropriate to operational constraints. Medical devices may not tolerate the same techniques as conventional IT assets. Third-party connections can carry contractual limitations, while production testing of clinical systems must account for patient safety and continuity of care. Clinical engineering, product teams, legal, procurement, and vendor management may all need a role in defining safe and meaningful scope.

The assessment mix in this dataset provides a directional illustration. Network penetration tests uncovered at least one Critical or High issue in 56% of assessments, compared with 29.6% of application tests. Differences in scope, target selection, and sample size prevent a causal conclusion, but the result reinforces why one testing method cannot answer every security question.

3. COMBINE POINT-IN-TIME AND CONTINUOUS TESTING

Point-in-time assessments provide depth, human judgment, and the ability to evaluate complex attack paths. Continuous security testing provides visibility between those engagements by identifying new assets, services, and externally observable changes as they appear.

Targeted testing can then investigate whether those changes introduce exploitable risk, while periodic penetration tests and red team exercises provide deeper evaluation of the systems and attack paths with the greatest potential impact.

The cadence should account for both time and change. An annual schedule can establish a baseline, while material changes trigger additional testing when they affect access, create new trust relationships, or involve systems critical to care delivery.

Common triggers include:

- Deploying a patient-facing or clinical application
- Connecting a new vendor to internal systems or data
- Changing identity or privileged-access models
- Introducing AI into a clinical or operational workflow
- Adding a new class of connected medical devices
- Making a significant cloud or network architecture change
- Remediating a Critical- or High-severity attack path

This creates a continuous cycle in which monitoring identifies change, targeted testing evaluates the resulting exposure, and deeper assessments determine how far an attacker could progress.

4. TURN TESTING RESULTS INTO RISK REDUCTION

Offensive security testing creates value when its results improve the broader security program. Findings should inform remediation priorities, architecture decisions, detection logic, incident-response planning, and future testing scopes.

Retesting is essential to that process. Closing a finding may confirm that a technical change was made, but it does not necessarily prove that the associated attack path has been eliminated. Retesting should determine whether the fix prevents further attacker progress and whether compensating controls perform as intended.

Program measurement should therefore extend beyond the number of findings produced. Useful measures include:

- Coverage of critical systems and trust relationships
- Time between a material change and its first test
- Critical and High-severity findings successfully remediated and retested
- Recurrence of previously identified weaknesses
- Progress an attacker can make toward priority objectives
- New exposures identified between scheduled assessments

A PRACTICAL STARTING PATH

Organizations can strengthen security incrementally by adding testing disciplines over time:

1. **Map the outcomes** that would most affect care delivery and the systems, identities, devices, and third parties they depend on.
2. **Test the highest-consequence paths** with the method appropriate to the question and the operational constraints of the environment.
3. **Define the changes that trigger targeted testing** between scheduled assessments, including new vendor access, major releases, identity changes, and new connected technology.
4. **Retest significant findings** and feed the results into architecture, detection, response planning, and the next testing scope.

FROM A PERIODIC SNAPSHOT TO AN ONGOING PRACTICE

Building offensive security into the security program creates a coordinated testing model that operates throughout the year. Scheduled assessments provide depth, continuous monitoring identifies emerging exposure, change-triggered testing keeps coverage aligned with the environment, and retesting confirms that remediation interrupted the attack path.

The result is a continuously updated understanding of where attackers can gain access, how far they can move, and whether controls protecting care delivery remain effective. That evidence helps leaders prioritize testing, remediation, and investment.

17

ASSESSMENT BRIEF

Scaling a Healthcare Platform Without Scaling Risk

A healthcare technology provider was transforming its business from on-premises deployments to a shared SaaS platform supporting cloud-connected medical devices. That shift fundamentally changed the security questions. Instead of securing individual customer deployments, the organization needed confidence that a shared cloud architecture could protect customer environments while supporting connected devices and an expanding external ecosystem.

THE CHALLENGE



Shared multi-tenant cloud architecture



Cloud-connected medical devices



Tenant isolation and cloud identities



A growing external attack surface

[READ THE FULL CASE STUDY →](#)

THE OUTCOME

An architecture security assessment confirmed key design decisions while identifying a cloud identity issue that could have exposed database credentials across tenant environments.

The organization strengthened tenant isolation and paired the assessment with continuous security testing to maintain visibility as new products, domains, and services were introduced.



KEY TAKEAWAY

As healthcare platforms evolve, so do the assumptions that protect them.

APPENDIX

The insights in this guide are informed by more than 160 offensive security assessments conducted between January 2024 and August 2026. The dataset represents over 40 healthcare organizations across eight industry segments and includes 825 individual security findings.

DISTRIBUTION OF ORGANIZATIONS BY HEALTHCARE SUB-INDUSTRY




Figure 9. Share of client organizations by healthcare sub-industry.

ENGAGEMENTS BY ASSESSMENT TYPE




Figure 10. Share of engagements by assessment type.

SEVERITY DISTRIBUTION BY FINDING CATEGORY

Data Exposure & Storage	4%	3%	14%	68%	11%
Other / Uncategorized	3%	9%	32%	38%	18%
Vulnerable / Outdated Software	0%	7%	20%	61%	12%
Network & Transmission Security	2%	2%	41%	52%	3%
Configuration & Hardening	0%	5%	23%	49%	23%
AuthN & Session Mgmt	0%	12%	20%	55%	12%
Credential & Secrets Mgmt	4%	24%	35%	36%	2%
Cryptography & Encryption	0%	8%	12%	48%	32%
Access Control & AuthZ	5%	39%	34%	23%	0%
Injection	0%	12%	35%	38%	15%
Client-Side / Web (XSS, CSRF)	0%	35%	43%	13%	9%
Business Logic & Anti-Automation	0%	0%	0%	100%	0%
Process / SDLC Gaps	0%	0%	0%	91%	9%
Device / Filesystem / Physical	22%	11%	56%	0%	11%
SSRF	0%	0%	50%	50%	0%
API & Service Endpoint Exposure	0%	0%	12%	88%	0%
Memory Safety / Exploitation Mitigation	0%	12%	38%	12%	38%
Denial of Service	0%	0%	71%	29%	0%
Cloud & Identity Misconfiguration	0%	50%	0%	25%	25%
	Critical	High	Medium	Low	Informational

Figure 11. Severity distribution within each finding category. Each category is calculated independently and totals 100%.

FINDING PREVALENCE BY ENGAGEMENT




Figure 12. Finding prevalence in single-scope and multi-scope assessments. Prevalence is the share of assessments in each group containing at least one finding in the category.

SEVERITY DISTRIBUTION BY ENGAGEMENT SCOPE




Figure 13. Severity distribution within single-scope and multi-scope assessments. Percentages are calculated within each scope group.

SCOPE DEFINITIONS

A single-scope assessment contains one recorded testing type. A multi-scope assessment contains more than one, such as application and network testing within the same report.

SEVERITY DEFINITIONS

Bishop Fox determines severity ratings using in-house expertise and industry-standard rating methodologies such as the Open Worldwide Application Security Project (OWASP) and the Common Vulnerability Scoring System (CVSS).

The severity of each finding in this report was determined independently of the severity of other findings. Vulnerabilities assigned a higher severity have more significant technical and business impact and achieve that impact through fewer dependencies on other flaws.

- Critical** Vulnerability is an otherwise high-severity issue with additional security implications that could lead to exceptional business impact. Findings are marked as critical severity to communicate an exigent need for immediate remediation. Examples include threats to human safety, permanent loss or compromise of business-critical data, and evidence of prior compromise.
- High** Vulnerability introduces significant technical risk to the system that is not contingent on other issues being present to exploit. Examples include creating a breach in the confidentiality or integrity of sensitive business data, customer information, or administrative and user accounts.
- Medium** Vulnerability does not in isolation lead directly to the exposure of sensitive business data. However, it can be leveraged in conjunction with another issue to expose business risk. Examples include insecurely storing user credentials, transmitting sensitive data unencrypted, and improper network segmentation.
- Low** Vulnerability may result in limited risk or require the presence of multiple additional vulnerabilities to become exploitable. Examples include overly verbose error messages, insecure TLS configurations, and detailed banner information disclosure.
- Informational** Finding does not have a direct security impact but represents an opportunity to add an additional layer of security, is a deviation from best practices, or is a security-relevant observation that may lead to exploitable vulnerabilities in the future. Examples include vulnerable yet unused source code and missing HTTP security headers.

OFFENSIVE SECURITY FOR HEALTHCARE ORGANIZATIONS

GIVE YOUR CYBER DEFENSES A HEALTH CHECK.

Bishop Fox provides a comprehensive portfolio of offensive security solutions to harden defenses against the latest cyber threats, helping to protect sensitive data and patient safety across the healthcare sector.

The healthcare industry's vast network of interconnected systems and devices, coupled with the highly sensitive nature of its data, makes it a prime target for cyberattacks. Despite increasing investments in security tools and employee training, healthcare remains the most exploited industry by hackers.

As attackers become more sophisticated in their tactics, the potential for exposing personal information, disrupting medical services, and endangering patient safety grows. To safeguard patient well-being and data integrity, healthcare organizations must adopt an offensive approach to cybersecurity.

With two decades of expertise in offensive security, Bishop Fox is the trusted partner for leading healthcare organizations seeking to assess the security of their applications, devices, networks, and cloud infrastructure, as well as their physical locations and employees.

Our solutions leverage cutting-edge tools, proven methodologies, and the best minds in offensive security to simulate industry-specific threats across your environment and rigorously test your defenses against ransomware, phishing, IoMT attacks, supply chain compromises, and more.

Contact Bishop Fox today to start planning your offensive security testing.

ABOUT BISHOP FOX

Bishop Fox is the leading authority in offensive security, providing solutions ranging from continuous penetration testing, red teaming, and attack surface management to product, cloud, and application security assessments.

LEARN MORE AT [BISHOPFOX.COM](https://bishopfox.com)

