Fortified HEALTH SECURITY

MID-YEAR

2026 HORIZON REPORT

The State of Cybersecurity in Healthcare

Table of Contents

CEO MESSAGE

What's Next.

That's not a question.
It's what we are always ready for.

In healthcare cybersecurity, no organization moves forward alone. The pace of change, the scale of risk, and the realities of healthcare operations demand partnership. Staying ahead isn't about reacting faster in the moment. It's about working together before the moment arrives. That kind of readiness is built through shared awareness, open communication, and trust.

What's next is shaped by listening closely to healthcare organizations and learning from real-world experience. Every incident, assessment, and conversation adds clarity. Collaboration becomes resilience. Security becomes something designed with healthcare, not imposed on it. Feedback becomes action.

That feedback continues to guide how we evolve alongside our partners. You asked for a clearer way to manage work across Central Command, so we built it by launching Workbench in Central Command Pro to provide a single, clear view of every service. You asked for third-party risk management that actually works for healthcare, so we launched TPRM with VendorIQ.

Proactive security doesn't mean doing more work. It means doing the right work earlier, together. It means aligning people, processes, and expectations so risk is addressed before it disrupts care.

This mindset guides how we support our partners every day. The focus isn't on handing off responsibility or leaving teams to navigate complexity alone. It's on standing beside healthcare organizations, helping them anticipate challenges, make informed decisions, and build programs that hold up under pressure.

As we move into the second half of 2026, "what's next" isn't about predictions. It's about preparation through partnership. It's about momentum built on communication and trust. And moving forward together, always ready for what's next.

Warm regards,

Dan L. Dodson

2026 Mid-Year Cybersecurity Check-In

Providers are finding more risk than ever under NIST CSF 2.0, and closing less of it.

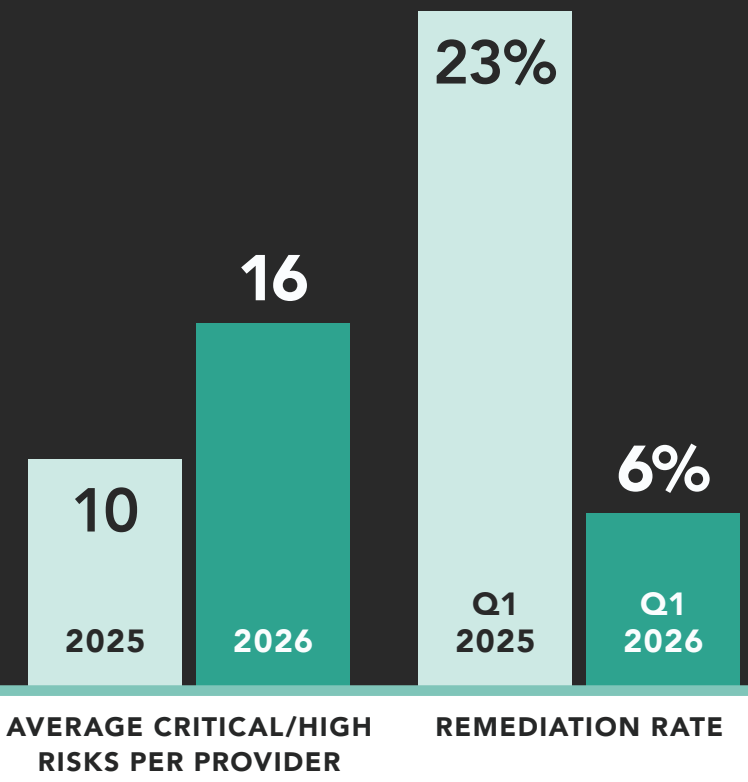

Healthcare spent the past year absorbing cyberattacks the way it absorbs most pressure. We keep the lights on while the strain builds underneath.

That strain always has tells, and the numbers from Fortified's rolling risk assessments give us part of the story: *organizations are identifying more risk than ever, but remediating less of it.*

Based on Fortified's analysis of NIST 2.0 risk assessment scoring, the overall risk remediation, as a percentage, has dropped to 6.4%. This is down from 23.3% from the same period last year.[1] Over that same period, we also observed that the average healthcare organization saw an increase of 60% in critical and high risk findings.

This isn't the story of a single catastrophic breach. It's the story of visibility outpacing capacity.

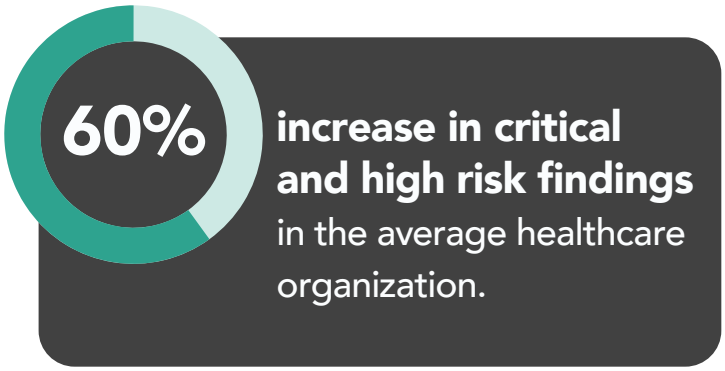

1 *Source: Fortified Health Security rolling NIST Cybersecurity Framework (CSF) 2.0 client assessment data.*

What's Driving the Increase in Identified Risk

Two forces are pushing the numbers in the same direction. One is good news wearing a ski mask. The other is the oldest problem in healthcare security.

6X
INCREASE
in Cybersecurity Supply Chain Risk Management Findings

4X
INCREASE
in Identity Management, Authentication, and Access Control Findings

345
DAYS
Average Time Identified Risk is Open

HIGH RISK CATEGORIES WHERE REMEDIATION HAS STALLED

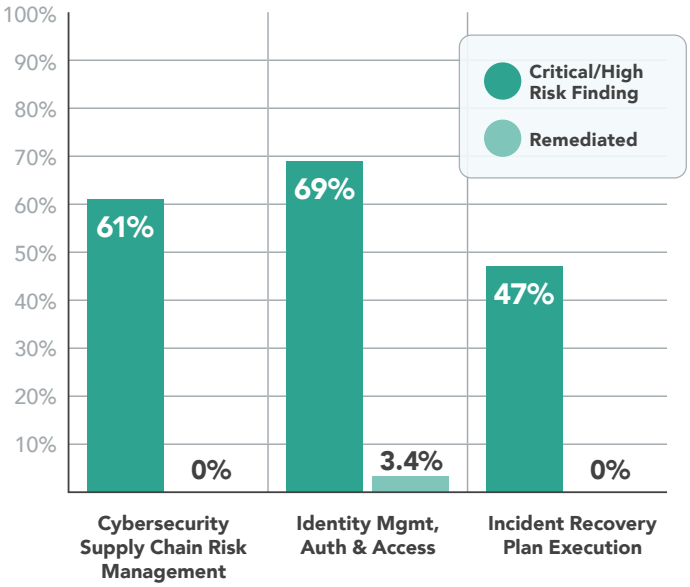

NIST CSF 2.0 Adoption Is Surfacing Maturity Gaps

Framework adoption is doing exactly what it should: shining a light on the places where maturity hasn't caught up.

Cybersecurity Supply Chain Risk Management findings are tracking toward a 6x increase in 2026 over 2025, with 63% of those findings rated critical or high. Assessments are exposing third-party risk gaps that many healthcare organizations have no current program to address, and they are struggling to remediate them once they are identified.

Identity Management, Authentication, and Access Control findings are tracking toward a 4x increase, with 64% rated critical or high. We're seeing real improvement in the organizational processes around Identity and Access Management (IAM), but the underlying work, authenticating users, services, and hardware, and protecting, conveying, and verifying identity assertions, continues to challenge teams.

Roles, Responsibilities, and Authorities controls are exposing previously unscoped gaps. Still, the progress on controls here is encouraging. With recent focus from executive teams, organizations are successfully establishing who owns what. Knowing who is accountable is the first step. Funding and staffing the work is the next one.

Human Capital & the Remediation Bottleneck

The human capital challenge is not new to healthcare cybersecurity. Modern tooling and security frameworks are just making the gaps it leaves more visible. Even where teams have prioritized well, remediation has stalled because you simply can't close a finding if you don't have the people to do the work. No framework fixes a headcount problem.

When it comes to protecting patients, the faster we surface risk, the more it matters what we do next.

Where That Leaves Us Now:

A Snapshot of Healthcare thru NIST CSF 2.0

Organizations will continue to mature against the NIST CSF 2.0 framework, even as the HIPAA update aims to compel compliance-motivated organizations.

Here is the snapshot of where the healthcare industry stands as relates to NIST CSF 2.0.

GOVERN

↓ **Oversight**

↑ **Roles, Responsibilities, & Authorities**

Oversight is slipping, while Roles, Responsibilities, and Authorities are improving. Now that we know who owns things, we can start to improve them.

PROTECT

↓ **Awareness & Training**

= **Platform Security**

Awareness and Training is slipping, while Platform Security is holding. Taken together, this seems to reflect a focus on technology, even though humans remain the leading attack vector.

IDENTIFY

↓ **Risk Assessment**

↑ **Improvement**

The Risk Assessment category is losing ground, with notable challenges in assessing critical suppliers prior to acquisition. Improvement is gaining ground, largely in incident response plans and other cybersecurity plans, indicating organizations are getting better at writing down what they will do even as the raw assessment workload grows.

DETECT

↓ **Continuous Monitoring**

= **Adverse Event Analysis**

Continuous Monitoring has lost ground due to the challenges organizations face in monitoring both the physical and technological environments for adverse events. Meanwhile, Adverse Event Analysis is largely stable as organizations continue to work towards the correlation of information from multiple sources.

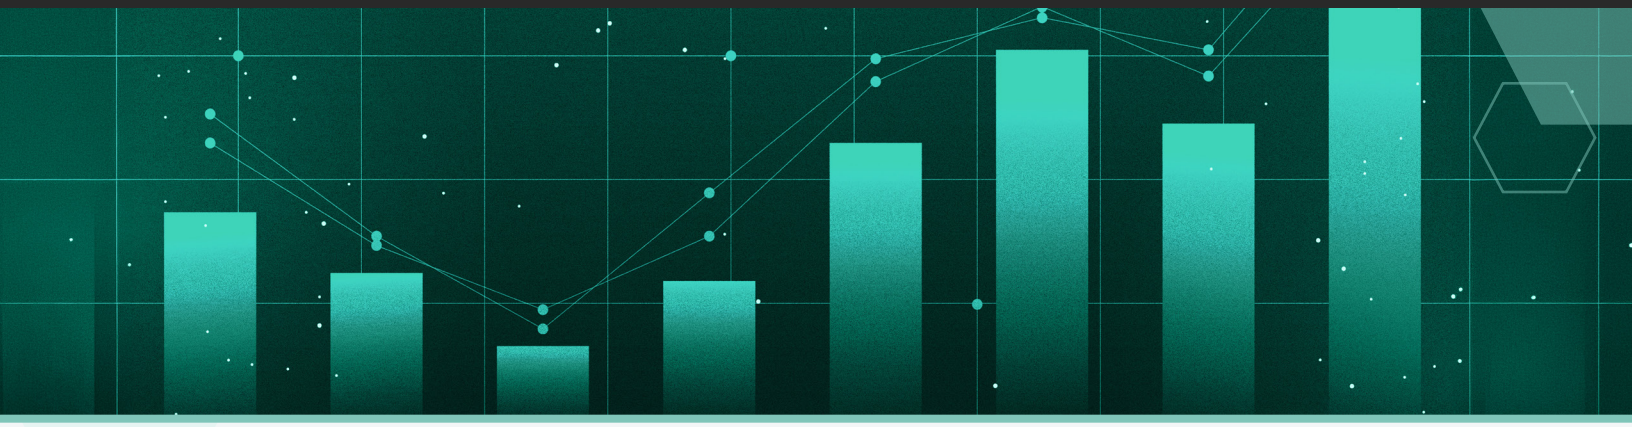

RESPOND

Incident Analysis

Incident Mitigation

Incident Analysis took a step back, but Incident Mitigation improved. Specifically, incident eradication is trending positive. The biggest hit landed on documentation during an investigation. This is indicative of the push and pull of continuity of care and business recovery against the hand of compliance and regulatory considerations.

RECOVER

Incident Recovery Communications

Incident Recovery Plan Execution

A fundamental truth about healthcare cybersecurity shows up in the essentially flat scores related to Recover. The biggest gap in control is RC.RP-04 – critical mission functions and cybersecurity risk management are considered to establish post-incident operational norms. Critical mission functions, a.k.a. patient care, will and should always take precedence.

Our challenge is to position cybersecurity as a component of critical mission functions, not as a separate consideration.

Healthcare is getting better at seeing its risk. The organizations that use that visibility to drive prioritization, a strategic roadmap, and the resourcing towards remediation will be the ones that turn a growing list of findings into risk reduction rather than the next headline.

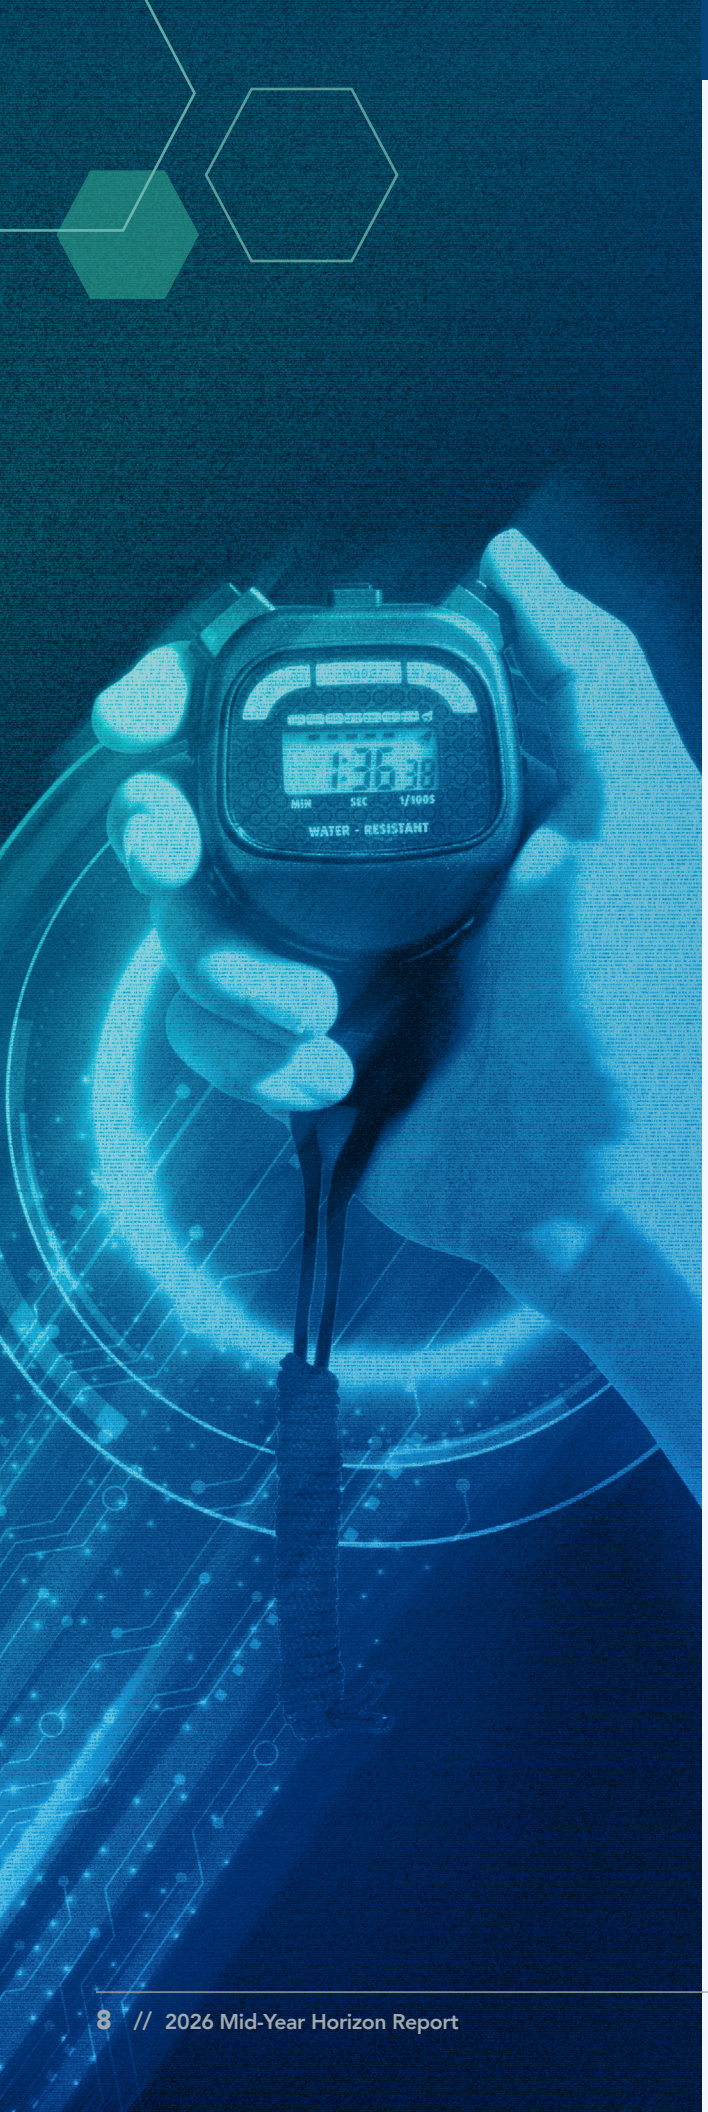

Intelligent Remediation:

Why MTTR Matters More Than Ever in Healthcare

Preston Duren // Vice President, Threat Services

Mean Time to Remediate (MTTR) is not a new cybersecurity metric. Cybersecurity teams have been measuring remediation timelines for years, but the role they play inside healthcare cybersecurity programs is changing.

Recently, healthcare organizations have begun to prioritize their cybersecurity investments. This allows for improved visibility, with many cybersecurity programs putting new budget towards tools that can identify vulnerabilities, suspicious activity, and emerging threats.

They can ***see*** the risk now.

So, the challenge changes from how to see the risk to how to reduce it quickly enough to stay ahead of attackers without disrupting patient care.

In other words, the bottleneck has shifted from detection to remediation, making MTTR one of healthcare's most important cybersecurity metrics and a new battle strategy.

Detection Is Only the Start

One of the most common misconceptions in cybersecurity is that identifying a threat is the same thing as reducing risk. It's not.

Detection creates awareness. Remediation reduces risk.

A critical vulnerability, suspicious activity, or security alert remains a risk until someone evaluates it, makes a decision, and takes action. That gap between awareness and action is where MTTR becomes important.

> **Detection creates awareness. Remediation reduces risk.**

While organizations have become more effective at detection, this improved visibility alone does not reduce risk. And the value of detection is limited by the organization's ability to respond.

If a security platform identifies malicious activity at 2 a.m., but nobody reviews the alert until 9 a.m., that is seven hours for the attacker to work on moving through your systems to launch an attack. In the past, that may not have been much time, but as threat actors leverage AI to move faster and "smash-and-grab" attacks become more attainable, how quickly the threat is understood, prioritized, and addressed becomes much more critical.

That is ultimately what MTTR measures.

Why Healthcare Struggles with Remediation

The issue with remediation is operational. Security teams face the same reality as every other department in healthcare: there are more priorities than there is time. Analysts are investigating alerts while supporting users, participating in projects, meeting with stakeholders, and helping clinical teams. Context switching is constant. Resources are finite.

The challenge with remediation is rarely identifying a problem. The challenge is determining whether that problem deserves immediate action. Is the vulnerability internet-facing? Is it associated with active ransomware campaigns? Is the affected system critical to patient care? Are compensating controls already in place?

The more time teams spend gathering answers to those questions, the longer the risk remains unresolved.

False positives, incomplete information, and uncertainty can all slow the response process. Before taking action, analysts need confidence that what they are seeing really matters and warrants attention.

Until those questions are answered, the risk remains unresolved, making quick access to complete information, context, and, of course, minimized false positives the keys to accelerating remediation.

Not All MTTR Is Created Equal

One of the biggest mistakes organizations make is treating MTTR as a single universal metric because not all risks deserve the same response.

A low-severity vulnerability on an internal system should not receive the same urgency as a critical vulnerability on an internet-facing system that is actively being exploited by ransomware groups. This may sound obvious, yet many organizations unintentionally treat them similarly.

When leaders ask, "What's our MTTR?" they may be focusing on the wrong question.

A better question is: What's our MTTR on the risks that matter most?

- Critical vulnerabilities.
- Internet-facing assets.
- Actively exploited weaknesses.
- High-priority security alerts.

These are the areas where remediation speed has the greatest impact.

> **What's our MTTR on the risks that matter most?**

Organizations that focus exclusively on lowering overall MTTR can create unintended consequences. Teams may spend significant time reducing remediation timelines for low-priority findings simply because they are easier to close, driving down overall MTTR while higher-risk issues remain unresolved.

The goal is to remediate the *right* things *faster*. That requires prioritization, business context, and a clear understanding of which risks pose the greatest threat to operations and patient care.

> **The goal is to remediate the *right* things faster.**

Why Healthcare Can't Prioritize Speed Alone

MTTR is largely an efficiency problem for many industries. But in healthcare, it's also a patient care problem.

Healthcare organizations operate in environments where availability matters:

- Clinical applications support care delivery.
- Medical devices often rely on legacy operating systems.
- Vendor-controlled technologies can limit flexibility.
- Maintenance windows may be limited or difficult to schedule.

Even well-intentioned remediation efforts can have unintended operational consequences. That's why healthcare organizations have traditionally been more cautious about change. And for good reason.

The purpose of cybersecurity is to support the mission of the organization. In healthcare, that mission is patient care, and attackers know this. They know healthcare organizations can't shut down systems or aggressively implement changes without first considering clinical impact. As a result, healthcare organizations face a balancing act that many other industries don't. They must move faster, but safely.

The Rise of Intelligent Remediation

One of the biggest obstacles to remediation is gathering enough information to determine the correct course of action. Analysts need to understand whether a system is internet-facing, whether a vulnerability is actively being exploited, whether compensating controls exist, and whether patient care could be affected by remediation efforts.

This is where I believe artificial intelligence can have the greatest impact on healthcare cybersecurity: helping analysts make decisions faster.

Historically, gathering that information required time. Analysts often had to search across multiple tools and manually piece together context before making a decision.

Instead of replacing human expertise, AI can surface that relevant information, correlate data across systems, and provide analysts with the context needed to make informed decisions more quickly.

> **The value of AI isn't that it makes decisions. The value is that it helps analysts reach the right decision *faster.***

What once took several minutes of investigation can now take a matter of seconds, and that distinction is important. The value of AI isn't that it makes decisions. The value is that it helps analysts reach the right decision faster.

Healthcare still requires human judgment, business context, and an understanding of patient care implications. Those responsibilities remain with the analyst. But better context delivered faster can significantly improve an organization's ability to reduce risk.

The Future of MTTR

As attacks continue to accelerate, MTTR will become an increasingly important measure of cyber resilience because it reflects something fundamental: How quickly can an organization reduce risk?

The organizations that will be most successful over the next several years will be the ones that can identify their highest-priority risks, provide analysts with the context needed to make informed decisions, and respond effectively without disrupting patient care.

Ultimately, that's what cyber resilience is about.

The Real Risk of AI Attacks in Healthcare Isn't the Technology

T.J. Ramsey // Senior Director, Threat Operations

Healthcare hasn't experienced its defining AI-powered cyberattack yet.

At least not publicly.

We haven't seen the massive AI-enabled ransomware event or fully orchestrated deepfake attack that suddenly changes the healthcare threat landscape overnight. No singular *"this changes everything"* moment.

But healthcare organizations are beginning to see early indicators of how AI could accelerate existing cyber threats through more believable social engineering, increasingly scalable attack tactics, and faster operational execution.

That trend deserves more of our attention.

Not because AI is some futuristic concept looming off in the distance. It's already here.

And if there's one thing cybersecurity has taught us over the years, it's that threat actors are exceptionally good at weaponizing emerging technologies faster than most organizations are prepared to defend against them.

We're already starting to see some of that happen.

Social Engineering Is Already Evolving

One of the clearest examples of how AI is already influencing cyberattacks is social engineering.

Historically, phishing emails often contained obvious indicators that something was wrong. The grammar felt awkward. The tone sounded unnatural. Sometimes you could practically read the accent of the person who wrote it.

People learned to spot those cues over time. AI removes many of them.

Today, threat actors can use AI tools to clean up language instantly, rewrite messages, improve tone, and create communication that feels significantly more believable. Ironically, legitimate users are already doing the exact same thing in their own workplaces every day.

We've all seen it:

"Make this sound more natural."
"Make this more concise."
"Rewrite this professionally."

That means some of the most recognizable phishing indicators employees were trained to identify are slowly disappearing.

Attackers don't need Hollywood-level deepfakes

Years ago, long before generative AI became mainstream, I conducted a social engineering exercise targeting a physician's account. I found publicly available videos of the physician speaking online and studied the cadence, tone, inflection, and speech patterns as closely as I could. Then I called the organization pretending to be him.

And it worked.

Now imagine those same tactics enhanced by AI tools that can replicate voices, camera presence, automate messaging, and scale deception faster than ever before.

That possibility should concern healthcare organizations.

Not because we've already seen it become widespread, but because the operational conditions that would make it effective already exist today. The technology is readily available, and the environment is primed.

Healthcare environments operate on urgency, collaboration, and trust. People move quickly because they have to. Clinicians need access. Finance teams process urgent requests. HR departments handle sensitive employee actions. Vendors constantly interact with internal systems.

That means the deepfake doesn't have to be perfect. Just good enough and at the right moment. Attackers understand that extremely well.

> **If I had to predict where AI is likely to influence healthcare cyberattacks first, it's going to be the speed of execution.**

AI May Not Change the Goal. But It Could Change the Speed.

I don't think AI fundamentally changes the core lifecycle of a healthcare breach, at least not yet.

The methodology behind most attacks are still largely the same:

- Gain access.
- Escalate privileges.
- Move laterally.
- Identify valuable data.
- Exfiltrate data.
- Encrypt systems.
- Extort the victim.

What has changed over the years is the speed at which attackers execute it.

Years ago, threat actors might quietly remain inside an environment for six months or more before launching an attack. Over time, we started seeing that timeline shrink into weeks or days. More recently, we've seen "smash-and-grab" style attacks happen incredibly quickly.

If I had to predict where AI is likely to influence healthcare cyberattacks first, it's going to be the speed of execution.

- Faster *reconnaissance.*
- Faster *lateral movement.*
- Faster *payload modification.*
- Faster *automation.*

And that's concerning in healthcare environments because there's already so much operational noise happening constantly. Thousands of legitimate processes are running simultaneously. If malicious activity starts happening fast enough, portions of it may simply blend into the background.

That's where organizations can get into trouble very quickly.

The Fundamentals Still Matter

As concerning as AI-enabled threats are, the fundamentals are still the fundamentals for a reason.

Organizations can spend all day worrying about hypothetical AI-generated exploits while still struggling to patch vulnerabilities that already exist inside their environment today.

That's backward.

Before organizations panic about AI discovering some unknown vulnerability tomorrow, they should first ask whether they are effectively addressing the known vulnerabilities sitting directly in front of them right now.

Because attackers absolutely will exploit those first.

I've compared this before to worrying about some future plague outbreak while ignoring the flu that's actively killing people today. There's still a lot we don't know about where AI-enabled threats ultimately evolve. But there are plenty of risks we already fully understand and still struggle to manage consistently.

> **The fundamentals are still the fundamentals for a reason.**

> **There are plenty of risks we already fully understand and still struggle to manage consistently.**

Strengthening fundamentals still matters:

- Patch management.
- Role-based access controls.
- Identity and access management.
- Tightened network architecture.
- Reduced external exposure.
- Security awareness training.
- Strong operational procedures.

Reducing unnecessary exposure, tightening Active Directory environments, limiting external access, and reviewing permissions remain some of the most impactful defensive measures healthcare organizations can take.

Nobody Wants to Call Their Baby Ugly

One of the biggest challenges in healthcare cybersecurity is that difficult conversations are often avoided.

Nobody wants to call their own baby ugly.

There is a reluctance to admit that your architecture has weaknesses, or that technical debt exists. Nobody wants to acknowledge that processes may be outdated or that operational shortcuts have introduced risk over time.

But avoiding uncomfortable conversations does not eliminate risk. If anything, it increases it.

Unfortunately, healthcare organizations often don't see urgency fully materialize until after an incident occurs.

Healthcare cybersecurity teams are filled with hardworking people trying to protect their organizations with the resources they have available. But the speed of modern threats, especially AI-enabled threats, is going to continue increasing.

Organizations will need to become more operationally agile in how they respond, prioritize, and defend themselves moving forward.

The Window to Prepare Is Still Open

The good news is healthcare organizations still have time to prepare for AI-enabled threats. That preparation starts with honesty.

The organizations that will be best positioned moving forward are not necessarily the ones chasing solutions to every futuristic doomsday AI scenario. They're the organizations strengthening fundamentals, reducing complexity, improving operational discipline, and taking an honest look at where their vulnerabilities, and limitations that slow the mitigation of those vulnerabilities, still exist.

Because attackers don't need organizations to be perfect.

They just need them to believe they already are.

> **Avoiding uncomfortable conversations does not eliminate risk.**
>
> **If anything, it increases it.**

Not Final, But Not Optional: A Guide to the Proposed HIPAA Security Rule

By Russell Teague // Chief Strategy and Security Officer

The Office for Civil Rights (OCR) has delayed the Notice of Proposed Rulemaking (NPRM) intended to overhaul the HIPAA Security Rule until at least July 2027, moving it from a near-term final rule posture into "Long-Term Actions" on the federal regulatory agenda. The proposed rule itself remains the same NPRM published in the Federal Register on January 6, 2025.

The NPRM would be the largest overhaul of HIPAA in over a decade and prompted more than 4,700 public comments that the OCR is still working through, including a coalition of more than 100 hospital and provider organizations that asked the agency to withdraw the proposal.

It would be easy to read all of that as a reason to wait and see where things land before taking any action. In my experience, that reading is a mistake. The direction is settled. Only the timing is delayed. This delay is giving healthcare a runway, not a reprieve, and the organizations that treat this window as preparation time will implement on their own schedule and budget rather than someone else's.

> **The organizations that treat this window as preparation time will implement on their own schedule and budget rather than someone else's.**

Why Now, After So Long?

Skepticism we saw on this proposed update, especially from executives and boards, wasn't without merit. But this isn't a death notice for HIPAA Security Rule modernization, as the NPRM isn't abstract policy. It's a direct response to some of the most notable, and highly visible, breaches of 2024.

Change Healthcare, February 2024. With the initial foothold being a Citrix remote-access portal without MFA, Senator Ron Wyden called it an incident that "could have been stopped with cybersecurity 101."

Ascension, 2024. An inadequately segmented network spanning 142 hospitals in 19 states allowed a single employee's mistake of clicking a malicious link to cause weeks of EHR downtime and ambulance diversions across the region.

These were the notorious events, but they were not isolated. The FBI's IC3 reporting counted 460 healthcare ransomware incidents, and the MOVEit campaign exposed 41 million records across 42 healthcare breaches.

OCR is closing the exact vectors that produced the largest healthcare breaches in U.S. history. The universal MFA mandate is the Change Healthcare answer. The segmentation mandate is the Ascension answer. One-hour access termination is the answer to insider-access monitoring failures like Montefiore, which closed at a $4.75 million settlement. Twenty-four-hour cross-entity notification is the response to MOVEit.

> **For thirteen years, "addressable" was the language that let organizations not encrypt, not segment, and not deploy MFA.**

The Changes

The NPRM changes included a proposal to eliminate the "addressable" versus "required" distinction at 45 CFR 164.306(d). Every implementation specification becomes required, with only narrow, enumerated exceptions for legacy technology under documented migration plans, genuine emergencies, and certain FDA-approved devices cleared before March 2023.

For thirteen years, "addressable" was the language that let organizations not encrypt, not segment, and not deploy MFA , as long as they documented why. When that word goes away, the partial rollouts and compensating-control arguments that were once defensible become gaps.

In total, the proposal carries roughly 40 new or enhanced controls. The list below captures some of the most consequential controls, not all the changes.

THE TECHNICAL SAFEGUARDS BECOMING MANDATORY:

MFA, universal (45 CFR 164.312)

Required for all access to relevant electronic information systems, with narrow exceptions.

Any Citrix portal, VPN, RDP session, or admin console without MFA is a Change Healthcare–shaped hole.

Encryption as a standalone standard (45 CFR 164.312(a)(2)(iv))(iv))

ePHI encrypted at rest and in transit, no longer buried inside access control.

TLS 1.2 at minimum, full-disk encryption, no cleartext HL7 feeds.

Network segmentation, now required (45 CFR 164.312)

The proposal calls for segmenting ePHI to "limit access and prevent lateral movement by intruders."

That is the Ascension incident written into regulation.

THE DOCUMENTATION BECOMING ENFORCEABLE:

Written asset inventory and network map (45 CFR 164.308(a)(1))

Every asset that touches ePHI, with version, accountable owner, and location, plus a network map, reviewed annually.

Change Healthcare could not confirm its scope to regulators for months.

Enhanced risk analysis with threat-vulnerability pairs (45 CFR 164.308(a)(1)(ii)(A))

The analysis must draw on the inventory and map and assign likelihood, impact, and risk level for each pair.

Risk-analysis failures appeared in every major 2025 enforcement action. It is the most weaponized provision in HIPAA enforcement.

THE OPERATIONAL CADENCE AND SERVICE LEVELS:

Vulnerability scanning every six months and penetration testing at least every twelve.

Critical patches within 15 days and high-severity within 30, with documented compensating controls when patching is not possible.

Workforce access termination within one hour, cross-entity access-change notification within 24 hours, and business-associate contingency activation within 24 hours.

A 72-hour recovery time objective and 48-hour recovery point objective for critical systems, and every business associate agreement renegotiated within twelve months.

What This Means For Healthcare Leaders

The "good enough on the crown jewels" era is over. Programs scoped to critical systems will have to be re-scoped to the full estate, and phased multi-year roadmaps will have to be compressed.

Once the clock starts, it'll run out fast. As it's currently written, the structure is:

- 60 days to the effective date.
- Then 180 more days to mandatory compliance (for 240 days total).
- Business associate agreements due within a year.

The timeline details may still shift — much of the public comment volume targets exactly this 240-day window as unrealistic against the medical-device and legacy-server realities above, so a longer runway is plausible — but the budget cycle for this work begins before the rule is final, not after.

Where the Challenges Are

Medical devices will be one of the most challenging areas for compliance. Per HIMSS, roughly 62 percent of organizations cannot adequately protect unpatchable devices, and 50 to 70 percent cannot host security agents on them.

Legacy servers that cannot run modern MFA or EDR, or meet a 15-day patch cadence, sit close behind, along with business associate renegotiation at scale and the one-hour access termination that lands squarely on identity and HR workflows.

> **The budget cycle for this work begins before the rule is final, not after.**

What To Do Today

A compliance-first program asks "when must I comply?" A security-first program asks "how fast can I implement controls that protect patients and care delivery?" My recommendation is to create the latter and lead with NIST CSF 2.0.

The HIPAA Security Rule is, by design, general. NIST CSF 2.0 is specific, prescriptive, and operationally testable. Aligning to it now insulates you from however the final rule's wording lands, because the NIST framework already includes most of what is coming.

It is also one of those rare investments that pays off in every scenario. If OCR narrows the rule or withdraws it entirely under the current deregulatory posture, a NIST-aligned program still strengthens your breach defensibility and your insurance position. There is no version of the future in which this work is wasted.

The upside reaches beyond protection. OCR has signaled a willingness to extend concessions, the modern echo of Safe Harbor, to organizations operating against a recognized framework during a breach investigation. NIST alignment also improves your cyber insurance posture at every renewal.

What to start before the clock does:

01. **Inventory honestly.** Map current state against the proposed rule, with special attention to anything classified today as "addressable." Build the asset inventory and network map as living artifacts, not stale diagrams.
02. **Reassess against the right framework.** Move from a HIPAA-only risk assessment to a NIST CSF 2.0 assessment, aligned to NIST SP 800-30, for a true readiness picture.
03. **Build a time-phased roadmap inside the normal budget cycle.** Prioritize MFA gap closure across remote access, VPN, Citrix, RDP, and admin consoles; segmentation decisions for high-risk clinical areas; and the business associate renegotiation list.
04. **Tighten the operational basics now.** Run incident-response readiness and tabletop exercises, complete a business impact analysis, build BCDR toward the 72-hour RTO, and set a corrective-action governance cadence.

Waiting compounds. Reactive compliance is the most expensive kind: premium-rate external resources, rushed deployments that need rework, and controls dropped in without change management that risk destabilizing clinical operations.

Every extra dollar pulled into an emergency scramble for compliance will be an extra dollar away from the bedside. HHS's own Regulatory Impact Analysis estimates first-year industry cost at roughly $9 billion, rising to about $34 billion over five years.

> **Every extra dollar pulled into an emergency scramble for compliance** *will be an extra dollar away from the bedside.*

Managing the Conversation

The wrong frame for the proposed rule is fear, uncertainty, and doubt. The right frame is opportunity.

I see a 5 point narrative:

01. The rule is coming and only the timing is open.
02. We have a rare window of regulatory breathing room.
03. Aligning to NIST CSF 2.0 now protects patients regardless of the final wording.
04. It improves our position with regulators, attorneys generals, and insurers if something goes wrong.
05. We can fund it inside the normal operating cadence rather than an emergency one.

That reframing shifts the boardroom question from "what will this cost us?" to "what do we get for moving now?" In my experience, that is a far easier question to win.

Stay safe, healthcare.

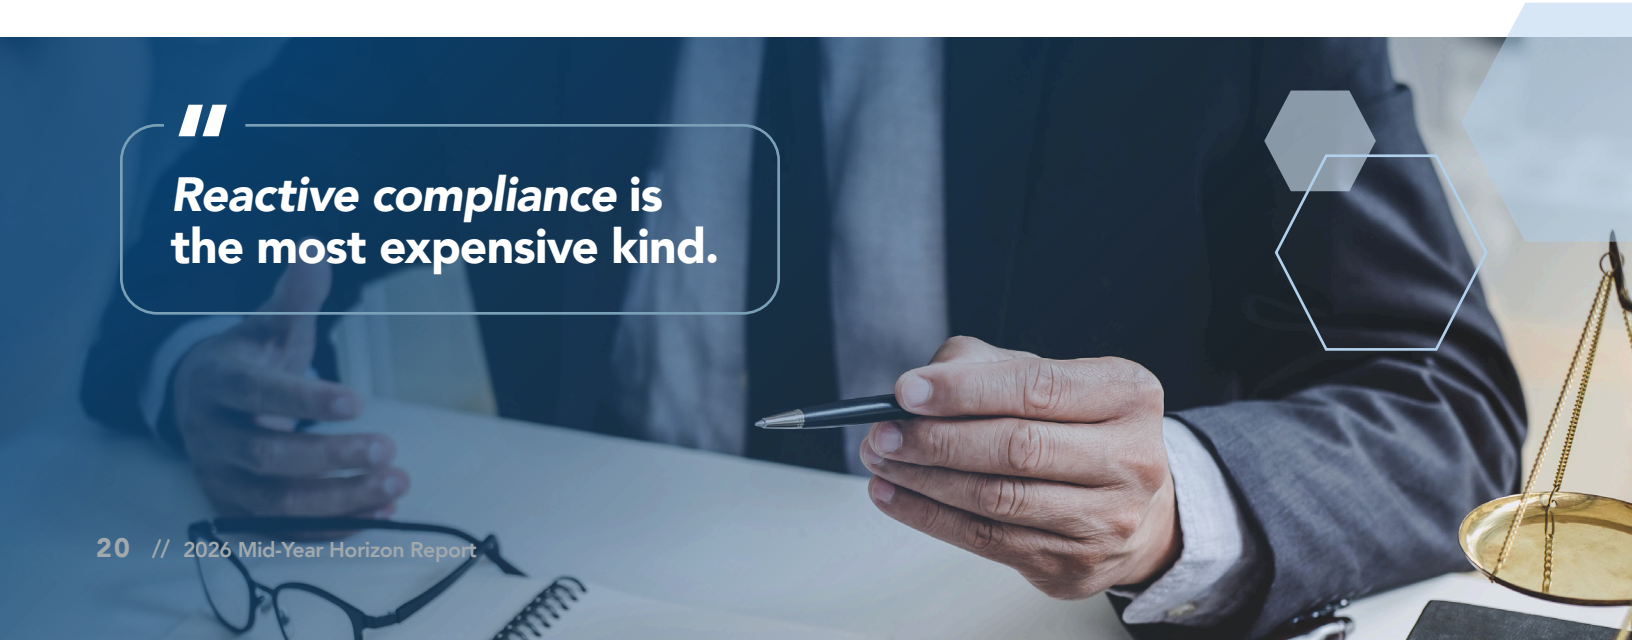

Change Your Oil: Why Identity Management Is Routine Care, Not a Project

Scott Milne // *Security Engineer*

"Change their oil."

It was a surprising answer to a question about what a healthcare organization should consider doing today to move their security program forward.

It wasn't really about cars, of course. It was about identity and access management (IAM). Stale accounts. Old passwords. Legacy Operating Systems and Authentication Methods. The quiet maintenance that gets pushed down the to-do list because the day is already full.

A car that skips its oil change doesn't fail that afternoon. It fails later: silently, expensively, and at the worst possible moment.

Nearly ⅓ of computers were inactive in over ½ of Active Directory Domains reviewed.

92% OF DOMAINS had an admin with a password 3+ years old.

IAM works the same way. Most breaches in healthcare don't stem from some brilliant, never-before-seen exploit or a masterful, digital lockpick. They come from a decades-old account nobody remembered, an over-provisioned service account, a password that hasn't changed since the server was racked, or from a computer that should have been retired with the Pontiac Aztek.

They start with the keys that are left in the doors.

As reported healthcare data breaches climbed from 39 in 2014 to 598 in 2024 and continue to be more frequent, they haven't necessarily all grown cleverer. Most were preventable with basic upkeep.

> **Role-based access does multiple things at once.**

THE FOUNDATION:
What Good Identity Looks Like

Strong identity rests on three questions. Who are you? What are you allowed to touch? Can we prove what you did? Authentication, authorization, and accountability. HIPAA's access control requirements are built on exactly this: unique user identification, controlled access to protected health information, and logging you can stand behind in an audit.

The most practical way to deliver all three is role-based access control. Put simply, assign people/accounts to roles, assign roles to permission groups, then assign permissions to resources to the permission groups – not to individual persons or accounts. It may sound like excessive layers at first glance, but in practice, it's the difference between a program you can audit and one you can only apologize for.

Role-based access does multiple things at once:

- Roles are cleanly separated from permissions.
- It enforces least privilege, so people carry only the access their job requires.
- It scales, so onboarding fifty new techs next quarter is one action, not fifty.
- It survives inspection, because when a regulator asks why someone can reach imaging, the role answers for you.

When designing role-based access control strategies, don't re-invent the wheel. Leveraging the classic strategy of AGDLP (Accounts, Global groups, Domain Local groups, Permissions) is a great way to standardize your environment with a repeatable group nesting model.

It's not really about picking the right acronym to follow. It's about using a repeatable process that keeps you rolling down the road.

LEGACY DEVICES:
Contain What You Can't Replace

You cannot retire a multimillion-dollar imaging system because it happens to run an old operating system with weak authentication. The budget for a new system would be hard to justify, and the patients still need their scans.

That's technical debt, and it's everywhere in this healthcare.

Wall it off. Segment it behind firewall policy. Where it makes sense, stand up a bastion domain so the device authenticates entirely outside your primary environment. Layer Group Policy and recognized security baselines on top to harden what you can't patch. And never, never, let a legacy box run on domain administrator credentials. One forgotten, over-privileged system is all an attacker needs to move laterally across everything else.

> **Treat AI as an "account" & wall it away from ePHI and other sensitive information.**

> Every clinician who walks in has to be credentialed before they can legally work. **That credentialing event is your trigger.**

TRAVELING NURSES & CONTRACT PHYSICIANS:
Tie Identity to Credentialing

Travel nurses and contract physicians are a real identity challenge: high turnover, outside credentialing, and a clock that's always running. A travel assignment often lasts thirteen weeks. Then the person is gone, and the account is still there.

There's a better way, and it's already sitting in your building. Every clinician who walks in has to be credentialed before they can legally work. That credentialing event is your trigger. Tie provisioning to it.

When a provider is credentialed, the account turns on. When the contract ends or the credential lapses, the account turns off. Disable it; don't delete it. A contractor who left in spring may be back in fall, and it's faster to re-enable an archived account than it is to recreate one.

Know Your Account Types

Many identity problems trace back to one mistake: treating every account the same. There are three primary kinds of accounts.

- **User accounts** belong to people. One per clinician, daily login, standard policy backed by MFA.
- **Service accounts** are non-human: the connections that sync your EHR, your PACS, your lab systems.
- **Privileged accounts** run administration, and they should never be someone's everyday login.

These account types are not interchangeable, but a single person can have more than one. A common reason to do so is to ensure that privileged accounts are not email-enabled.

For example, an IT user will often need a standard user account to log in to their laptop and access email, Teams, etc., an "ADM" (privileged account) for on-prem administration (no email on that account), and a cloud-only/Entra account for Azure/M365 administration. In some cases, it is also recommended to split IT users further into separate accounts that manage workstations and servers.

A nurse, on the other hand, is likely to only need a single account. There is a growing trend towards standard users not being email-enabled except for managers and other leadership roles. So, in the case of the nurse, that user account may or may not be email-enabled.

> **There is a growing trend towards standard users not being email-enabled.**

There is, of course, grey area. A specific physician who needs additional access. A contractor who needs a privileged account but only for a limited window on specific days. This is where Just-In-Time (JIT) access and Privileged Identity Management (PIM) come into play. Whether managed through PowerShell scripts or a third-party tool, taking the time to layer these controls over specific users in your account is a critical component to truly maintaining least privilege.

Service accounts are where it gets dangerous. Attackers love them because of a technique called Kerberoasting: a privileged service account with an old, weak password can have its ticket pulled and cracked offline, no special access required. The 2024 attack on Ascension showed how far down that path an intruder can travel. The preventative is straightforward: managed service accounts with long, regularly rotated passwords, and delegated, fine-grained permissions instead of standing administrator rights.

Which brings us to the most common finding of all. Teams often make a service account a domain administrator because it's faster than delegating the exact permissions it needs. Faster today. Far more expensive later.

HYGIENE:
Where Most Programs Fail

When our team looked across more than a hundred healthcare domains, the pattern was almost universal. Ninety-eight percent had at least ten percent of their accounts sitting stale, unused for six months or more. In some environments, the majority of objects were inactive. Every one of those is a working key, and each one of them looks like legitimate activity if an attacker manages to pick it up.

The solutions require more discipline than budget:

- Run a monthly report of accounts inactive for six months, and route it to the people who own those roles.
- Rotate the Kerberos account password on a schedule.

It's the single most important credential in your domain, and, in too many organizations, it hasn't changed in decades.

- Tie onboarding and offboarding to HR systems so accounts disable automatically when someone leaves.
- Set password and rotation policy by account type, not by habit.

None of this is glamorous. All of it is the work.

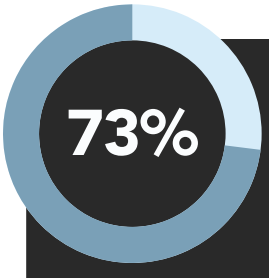

Kerberos password not changed in over a year in 73% of reviewed domains

Putting It Together

Let's look at a use case. A radiology technician sits in a single role group. That role carries exactly the permissions the job needs — read imaging, write to PACS — and nothing else. Move that person to the lab, and you change one membership, the role. The access follows, maintaining least privilege and providing a clean audit trail.

ROLE-BASED PERMISSION NESTING (AGDLP MODEL)

A radiology technician inherits only the access the role requires—nothing more

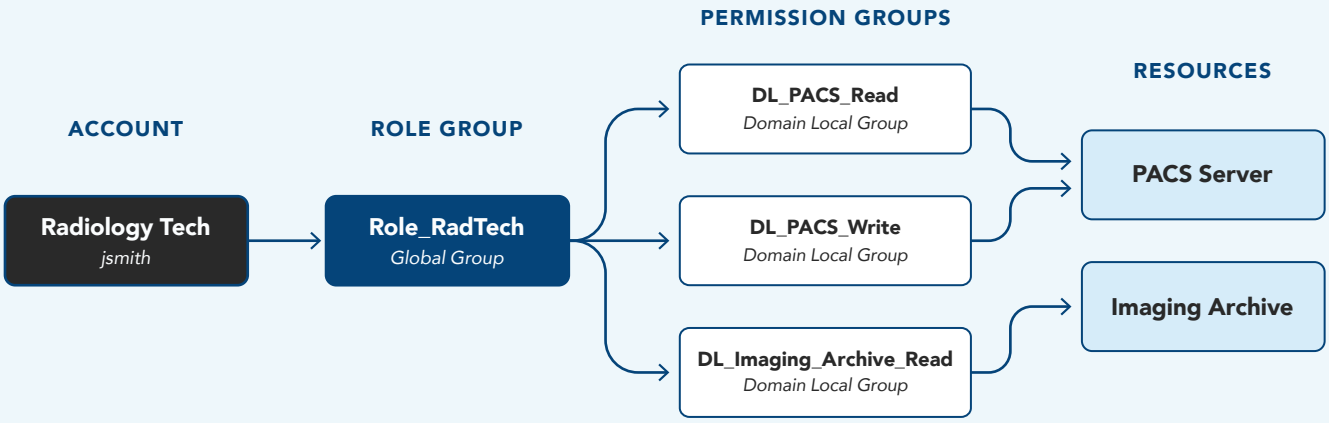

Move the technician to the lab? Change one membership. The access follows the role.

Change the Oil

Identity maintenance will never be the exciting part of cybersecurity. It is a quiet, extremely critical part of it though: it closes the doors most attackers walk through.

It also isn't a project you finish and file away. IAM is a practice you keep and the cost of doing it can be measured in hours. The cost of skipping it is measured in millions, in downtime, and in patient trust. Because, in healthcare, those doors don't just protect data. They protect care.

So, run the reports. Flush out where your keys are sitting, drain out the stale accounts, swap out your role permission filters for fresh ones, and tighten access to your service accounts so they don't leak any information.

That's your oil change. Don't skip it.

Healthcare Cybersecurity **Needs to Train Like the Fire Service**

Mark Ferrari // Vice President, Advisory Services

> Healthcare needs to approach incident response programs & cyber response exercises with the **mindset of a first responder.**

In addition to marking 30 years in Healthcare IT and Cybersecurity, 2026 marks my 30th year in the volunteer Fire and EMS service.

Over that time, one thing has always been clear: communities don't ask firefighters to prepare for the fires they expect. They prepare for the fires that will eventually come. Not as an optional investment or a compliance exercise, but as a constant operational necessity.

It's a matter of public safety.

Cybersecurity is patient safety, and healthcare needs to approach incident response programs and cyber response exercises with the mindset of a first responder.

It's Never *If.* It's Always *When.*

One of the most dangerous assumptions healthcare organizations can make is believing they're unlikely to become a target.

Smaller hospitals and healthcare providers may assume attackers will focus on larger systems. More mature organizations may believe their defenses are sophisticated enough to significantly reduce their risk. In either case, preparedness often quietly becomes more performative than operational, existing mostly on paper rather than in practice.

But emergency responders understand something deeply important: *probability does not change responsibility.*

Firefighters can't say: "We have not had a structure fire in years. Let's stop running drills and checking the equipment."

Even a rural fire department that only sees one major fire every year maintains engines, staffing, equipment, and training standards following many of the same operational expectations that a major metropolitan department responding to fires daily does.

But emergency responders understand something deeply important: *probability does not change responsibility.*

The frequency may differ, but the standard of readiness does not.

Whether organizations experience attempted breaches consistently or have never had a major cyber event, the expectation must remain the same: eventually, a serious incident *will* occur.

The only question is whether the organization is operationally prepared when it happens.

Preparedness is not measured by the absence of disaster. It's measured by the ability to respond effectively when a disaster occurs.

Preparedness Requires More Than Technology

Cybersecurity conversations often center around technology. Detection tools. Endpoint protection. Monitoring platforms. AI-driven analytics.

Yes, those tools matter, but in emergency response, tools are only one piece of readiness.

Fire departments operate under defined standards because communities know that minimum capabilities need to exist before an emergency ever happens. The National Fire Protection Association (NFPA), for example, outlines baseline equipment and operational expectations for fire departments regardless of how often they respond to major incidents, or even what steps the community may have taken to prevent them.

A fire engine is expected to be fully operational and immediately deployable at all times. Staffing levels are defined. Equipment standards are defined. Response expectations are defined.

No one argues that a community should lower those standards simply because a major fire is statistically unlikely.

Healthcare cybersecurity needs to be viewed through the same lens.

Organizations can't assume that because an attack has not yet caused operational disruption, their current capabilities are sufficient. Preparedness is not measured by the absence of disaster. It's measured by the ability to respond effectively when a disaster occurs.

Healthcare organizations need more than security tools. They need operational readiness.

Cybersecurity Teams Need Operational Muscle Memory

One of the most important traditions in the fire service is drill night.

In many departments, one evening every week is dedicated entirely to training. Sometimes it involves equipment maintenance. Sometimes, classroom instruction. But most often, it includes simulated emergency scenarios to improve coordination, sharpen decision-making, and build response speed under pressure.

These exercises are a collaborative event involving neighboring departments, because large emergencies require coordinated response across multiple teams. Training together builds familiarity, communication, and trust before an actual crisis occurs.

Drills like these should be the same for cybersecurity in healthcare.

Too often, cyber incident response plans exist primarily as documents. Teams may review them annually or complete a tabletop exercise simply to satisfy compliance requirements. But real incidents rarely unfold according to documentation. They unfold in confusion, pressure, and uncertainty.

> But real incidents rarely unfold according to documentation.
> **They unfold in confusion, pressure, and uncertainty.**

That is why operational repetition matters.

Healthcare organizations should regularly simulate realistic attack scenarios involving more than IT and security teams. Executive leadership, legal counsel, communications, clinical operations, third-party vendors, and external partners must be included as well.

The goal is not only to validate a policy, but also to build organizational muscle memory. Because during an actual cyberattack, speed, and coordination matter as much as technical capability.

> **Training together builds familiarity, communication, and trust before an actual crisis occurs.**

Staffing Matters More Than Many Organizations Want to Admit

Another lesson cybersecurity can learn from emergency response is that readiness depends *heavily* on people.

The fire service doesn't assume a small number of responders can compensate indefinitely through effort alone. Staffing minimums exist because emergency response requires enough trained personnel to execute effectively and safely.

Healthcare cybersecurity faces a similar challenge. Many organizations operate with understaffed teams carrying enormous operational responsibility. The assumption often becomes that talented individuals can simply work harder during a crisis. But burnout is not a response strategy.

Requirements for effective cyber preparedness include sufficient staffing, clearly defined responsibilities, escalation paths, leadership coordination, and external partnerships who can support the organization and the community when incidents occur.

That's especially important in healthcare, where cyberattacks can directly affect clinical operations, patient access, and care delivery.

> In both firefighting and cybersecurity, the **outcome is often determined long before the alarm ever sounds.**

> **Attackers don't need organizations to be careless. They only need them to be unprepared for the moment that matters most.**

Healthcare Cannot Afford *Good Enough*

Firefighters are expected to maintain readiness even when the emergency may never come. Communities accept that standard because the consequences of failure are too high.

Healthcare cybersecurity is patient safety. It deserves the same mindset.

Organizations can no longer view preparedness as a checklist exercise or assume maturity alone guarantees resilience. Attackers don't need organizations to be careless. They only need them to be unprepared for the moment that matters most.

The healthcare organizations that can handle future cyber threats effectively may not be the ones with the biggest budgets or the most advanced tools. They'll be the ones who operationalize readiness, train consistently, prepare collaboratively, and treat cyber response as a practiced discipline rather than a "possible" event.

Because in both firefighting and cybersecurity, the outcome is often determined long before the alarm ever sounds.

About the Contributors

Dan L. Dodson

CEO

As the CEO of Fortified Health Security, Dan Dodson brings over 17 years of experience leading healthcare and insurance organizations. As a recognized thought leader in healthcare cybersecurity, Dan is a frequent speaker at industry events and conferences including CHIME, HIMSS, and HIT Summits. His insights and data-driven expertise in cybersecurity, data privacy, risk management, and threat mitigation are regularly featured in popular media and trade publications such as Forbes, Becker's Hospital Review, and Healthcare Business.

Russell Teague

Chief Strategy and Security Officer

Russell Teague is a healthcare cybersecurity strategist with nearly three decades of experience advising healthcare organizations across complex environments. A U.S. Army Intelligence veteran, he brings a mission-driven, risk-informed approach to cybersecurity leadership, has consulted with the White House on national healthcare cybersecurity efforts, and is a frequent speaker at HIMSS, VIVE, HSCC, Health Connect Partners, and executive leadership events.

Mark Ferrari

Vice President, Advisory Services

Mark has a proven record of guiding cybersecurity strategy for healthcare, joining Fortified through the acquisition of his healthcare-focused cybersecurity firm, Latitude. Prior to that, he brought executive insight and hands-on expertise to his roles as EVP at a cybersecurity consultancy and CISO at a software development and consulting company.

Preston Duren

Vice President, Threat Services

Preston Duren brings more than 16 years of IT/security expertise to his role as VP of Threat Services at Fortified. His experience spans threat and vulnerability management, security engineering, security program development, digital forensics, and SOC. Previous roles include engineering/architecture at Community Health Systems & Information Security Officer at RCCH Health.

T.J. Ramsey
Senior Director, Threat Operations

T.J. Ramsey is a seasoned IT security professional with nearly 20 years of experience focused on healthcare and defense intelligence. He served as a U.S. Army Military Intelligence Analyst for the Department of Defense and held security roles at Obsidian Solution Group and SAIC/Leidos.

Jim Kirk
Senior Director, Consulting Services

Jim is a Security Compliance Consultant with over 35 years of experience in information security and infrastructure covering healthcare, third-party administration, technology, financial, government, legal, manufacturing, and process control. Jim has held roles as Infrastructure Architect, Chief Information Officer, and Information Security Officer. Jim has a deep understanding of technical and privacy security controls, threat management, and mitigation.

Scott Milne
Security Engineer

Scott Milne is an IT leader with more than 25 years of hands-on experience across healthcare and enterprise environments. He has served successfully as a Hospital Director of Information Systems and led large-scale desktop standardization initiatives and complex HIS, RIS, PACS, HL7, and revenue cycle implementations. Scott excels at managing multi-site deployments, and delivering practical, secure technology solutions using modern Microsoft cloud and endpoint platforms.

About Fortified HEALTH SECURITY

Fortified Health Security is healthcare's cybersecurity partner, trusted by healthcare organizations nationwide to deliver tailored, high-touch programs built exclusively for healthcare that reduce risk, simplify complexity, and protect what matters most: their patients.

As a five-time consecutive Best in KLAS winner, Fortified understands the full spectrum of healthcare cybersecurity, from rural providers to enterprise networks, third-party vendors, and connected medical devices. The company's award-winning Central Command platform, featuring innovations that translate client feedback into action, enabling smarter, faster security decisions that strengthen every layer of defense.

Fortified doesn't just guard the perimeter. The team embeds with clients, bringing context to every alert and helping healthcare leaders move from reactive to resilient, 24/7, 365.

Because in healthcare, cybersecurity isn't just an IT issue. **It's a patient safety issue.**

Learn more at fortifiedhealthsecurity.com.

Fortified
HEALTH SECURITY

www.FortifiedHealthSecurity.com

connect@fortifiedhealthsecurity.com

120 Brentwood Commons Way
Building 4, Suite 500
Brentwood, TN 37027