dynamic quest

MID-MARKET LEADER'S GUIDE

2026 CYBERSECURITY SURVIVAL GUIDE

What Mid-Market Leaders Need to Know About AI, Insurance, and Compliance in 2026

BY MARTIN CAPURRO

Chief Technology Officer & SVP of Operations
Dynamic Quest

FOREWORD

A note from the CTO

Martin Capurro
Chief Technology Officer & SVP of Operations, Dynamic Quest

The cybersecurity conversation changed in 2026, and most mid-market companies haven't caught up.

A year ago, the dominant question I heard from prospective clients was *"Are we doing enough?"* Today the question has shifted to *"How do we know what enough even looks like anymore — and at a price we can defend?"* That's a meaningful change, and it reflects something real. The threat landscape is genuinely different than it was twelve months ago. AI has rewritten how attackers move, how defenders detect, and how regulators and insurers measure whether you're a credible counterparty. The cost of getting it right has gone up, lead times on certain hardware and licensing have stretched, and the cost of getting it wrong has gone up faster. The companies I talk to most often aren't underfunded on security. They're under-oriented. They have tools, but no framework. They have policies, but no playbook. They have a vendor, but no advisor.

This guide is written for the leaders running mid-market companies — manufacturers, CPA firms, healthcare practices, credit unions, professional services firms — that have outgrown the small-business cybersecurity playbook but haven't yet built the program a 200-person company actually needs. You are the segment most aggressively targeted right now, and you are the segment with the biggest gap between what your insurance carrier expects, what your customers ask about in security questionnaires, and what your IT setup can actually demonstrate. That gap is where breaches live.

We've structured this guide around the questions our clients ask us when they're three months out from an insurance renewal, two weeks out from a customer audit, or one ransomware story away from a board conversation. It's not a 101 on what phishing is. It's a working framework for what mid-market cybersecurity should look like in 2026, what's changed because of AI, and how to evaluate whether the partner managing your environment is actually positioned to defend it.

The good news embedded in this year's data is that defenses are working — when they're configured right and monitored well. The harder news is that "configured right and monitored well" is doing a lot of work in that sentence. My goal with this document is to help you tell the difference, and to help you make the investment case for the right level of protection at a cost your business can absorb.

— Martin

EXECUTIVE SNAPSHOT

Six numbers your leadership team needs to know

The data below is drawn from the most credible cybersecurity research published in the last twelve months — Gartner, IBM, Sophos, Marsh McLennan, Coalition, Verizon, and the IRS Security Summit. Each of these stats reflects something we are seeing daily in client environments across our footprint.

$10.22M

Average U.S. breach cost — all-time high

U.S. breach costs jumped 9% to a record in 2025, even as the global average fell to $4.44M.

IBM Cost of a Data Breach Report 2025

$244B

Global security spending projected for 2026

Up 13.3% YoY. Cloud security growing fastest at 28.8%. Investment is accelerating.

Gartner Forecast 2023–2029

99%

Of cyber insurance applications ask about MFA

82% of denied claims involved organizations without MFA. The questionnaire is the audit.

Marsh McLennan 2025 Market Report

88%

Of SMB breaches involved ransomware

Compared to 39% at large enterprises. Attackers are targeting mid-market companies more often, not just harder.

Verizon 2025 DBIR

97%

Of AI-related incidents lacked access controls

63% of organizations have no AI governance policy at all. The largest controllable risk most companies have.

IBM Cost of a Data Breach Report 2025

$1.9M

Average breach cost savings using AI in security

Organizations using AI extensively contained breaches 80 days faster than those that did not.

IBM Cost of a Data Breach Report 2025

SECTION 1

The 2026 Mid-Market Cybersecurity Reality

What's actually different this year

If you read only one section of this guide, read this one.

The cybersecurity landscape has shifted in four ways that matter specifically to mid-market companies in 2026. None of these are theoretical. All of them are showing up in the way our clients are renewing insurance, fielding customer security questionnaires, and absorbing the second-order effects of attacks on their suppliers. We will spend the rest of this document going deeper into each, but here is the orientation.

The first shift: attackers found the middle of the market.

For most of the last decade, cybersecurity narratives were written by enterprise CISOs for enterprise CISOs. The data, the tools, the frameworks all assumed a Fortune 500 reader with a 30-person security team. Mid-market companies — meaning organizations roughly between $3 million and $150 million in revenue, and 25 to 1,000 employees — were treated as either an afterthought or a smaller version of the same problem.

That is no longer accurate. The 2025 Verizon DBIR found that 43–46% of all cyberattacks now target businesses with fewer than 1,000 employees, with SMBs experiencing nearly four times the incident volume of large organizations. The reason is straightforward: ransomware groups have figured out that mid-market companies typically allocate 6–9% of their IT budget to cybersecurity versus 12–15% at enterprises, and that 43% of SMBs have no dedicated cybersecurity staff member at all. The math from the attacker's perspective is simple. Smaller targets mean smaller ransoms per hit, but they also mean lower defenses, faster compromise, and many more available targets. The aggregate revenue model favors volume.

For Dynamic Quest's clients, this shift is most visible in three patterns we have seen accelerate in 2026: more attempts against companies that have never been targeted before, more attacks that exploit known but unpatched vulnerabilities (which now account for 32% of ransomware incidents in manufacturing, the leading root cause, per Sophos), and more attacks that succeed because of a missing basic control rather than a sophisticated technique.

The second shift: AI changed both sides of the equation.

We will spend a full section on this later, but it deserves to be framed up front. AI is not a 2027 problem or a 2028 problem. It is a 2026 operational reality, and it cuts both ways.

On the attacker side, AI has dramatically lowered the cost and raised the quality of phishing, deepfake impersonation, and social engineering. IBM's 2025 research found that 16% of breaches now involve attackers using AI tools, most often for phishing and deepfake impersonation. The voicemail your CFO got

last week from "the CEO" asking for an urgent wire transfer no longer has the obvious tells it had two years ago. Generative AI removed them.

On the defender side, AI has materially shortened breach lifecycles and reduced costs for organizations that have deployed it well. Organizations using AI extensively in security operations saved an average of $1.9 million in breach costs and contained breaches 80 days faster than those that did not.

The gap that matters most, and the one Gartner explicitly flagged in its 2026 trends report, is the governance gap in between. Gartner's 2026 Top Trends in Cybersecurity report identifies "agentic AI demanding cybersecurity oversight" as the leading trend reshaping how security leaders must operate, citing the rapid proliferation of unmanaged AI agents through no-code platforms and "vibe coding" as a structural new attack surface. 63% of breached organizations either have no AI governance policy or are still developing one. That is the exposure we see most often in mid-market environments today, and it is almost always invisible to the company until something goes wrong.

The third shift: insurance, regulation, and customer due diligence converged into a single de facto audit.

Five years ago, your cyber insurance renewal was a form. Today it is the most rigorous security audit most mid-market companies will ever face — more granular, in many cases, than the audits required by HIPAA, SOC 2, or PCI. According to Marsh McLennan's 2025 Cyber Insurance Market Report, 99% of cyber insurance applications in 2025 included specific questions about MFA implementation, and Coalition's data shows that 82% of denied claims involved organizations without MFA in place. Carriers are not only asking whether you have controls. They are asking for screenshots, policy exports, restore-test logs, and incident response plans with documented test dates.

At the same time, your customers are asking the same questions. Mid-market manufacturers selling into automotive, aerospace, or healthcare supply chains are now routinely fielding 80-question security assessments from buyers who would not have asked anything five years ago. Mid-market accounting firms are subject to the AICPA's updated Statements on Standards for Tax Services Section 1.3, effective January 1, 2024, which explicitly requires safeguarding taxpayer data and documenting the protections. And every mid-market company is subject to a regulatory environment that, per Gartner, is volatile enough to be its own 2026 trend: *"Global Regulatory Volatility Drives Cyber Resilience Efforts."*

The practical effect is that the bar for being a credible counterparty has moved. It is no longer "we have antivirus and we back things up." It is "we can produce, on 48 hours' notice, evidence of MFA enforcement on every privileged account, EDR coverage on every endpoint, restore-tested backups, and a current incident response plan with documented tabletop exercises." Companies that can produce that evidence get renewed, get awarded, and get trusted. Companies that cannot, increasingly do not.

The fourth shift: the economics of "enough" got harder.

Cybersecurity got more expensive in 2026, on multiple fronts at once. Software licensing tied to AI-enhanced features is up. Cyber insurance premiums are reflecting tighter underwriting, with carriers raising premiums 30% to 50% year over year for organizations that cannot demonstrate the required controls. Gartner projects information security spending will grow 13.3% in 2026 alone — meaningfully faster than overall IT spend. Hardware lead times for certain firewalls, switches, and on-prem infrastructure have stretched as supply

chains absorb both AI-driven demand and ongoing component constraints. The same security stack that cost X eighteen months ago costs noticeably more today, and takes longer to deploy.

For a mid-market CFO trying to defend a security budget, this is a genuinely difficult moment. The temptation is to defer — to push the firewall refresh another quarter, to delay the EDR rollout to next fiscal year, to take a smaller cyber insurance limit because the premium jumped. We understand the pressure. We see it in every budget conversation we have with mid-market leadership.

The honest framing is this: a defensible security program at mid-market scale typically runs 1.5% to 3% of revenue, depending on industry and regulatory exposure. That number has crept up about 30 to 50 basis points over the last two years. Compared to the cost of a breach — which now averages $10.22 million for U.S. companies, an all-time high — it remains overwhelmingly the right side of the math. Compared to the cost of a denied insurance renewal, a lost contract because you failed a customer security questionnaire, or a regulatory fine, it is also the right side of the math. The companies we see making the best decisions in 2026 are the ones treating cybersecurity as an operating cost line tied to their license to operate, rather than a project budget tied to a CapEx cycle.

We will return to the economics in Section 5 when we walk through what a mid-market program actually looks like and what the right line items are.

SECTION 2

Where AI Changed the Equation

If the four shifts in Section 1 are the macro picture, AI is the thread running through all of them. It is the single largest change to the cybersecurity environment between the 2025 version of this guide and this one. Mid-market companies that understand this shift will navigate the next twenty-four months meaningfully better than those that do not.

How attackers are using AI

The effective bar for a successful phishing or social engineering attack dropped sharply in 2025. The reason is that generative AI removed the production constraints attackers used to operate under. Writing a convincing spear-phishing email used to take a skilled human attacker thirty to sixty minutes per target, with most of that time spent on grammar, voice-matching, and finding the right pretext. Today, an AI agent does the same work in under a minute, at higher quality, in any language, with personalization drawn from the target's LinkedIn profile and recent public activity. The economics of attack volume changed.

IBM's 2025 research found that 16% of breaches now involve attackers using AI tools, most often for phishing or deepfake impersonation. The deepfake angle is particularly important for mid-market companies. We are seeing voice-cloned calls to CFOs requesting urgent wire transfers, video deepfakes of executives in fraudulent Teams or Zoom meetings, and AI-generated text messages from "the CEO" that match the actual CEO's writing patterns closely enough to fool people who work with that CEO every day. The tells we used to teach employees to look for — bad grammar, generic greetings, off-brand voice — are no longer reliable. The attack got better. Employee training has to evolve to match it.

Two specific AI-enabled attack patterns are showing up most frequently in mid-market environments right now:

AI-augmented phishing at scale. Attackers are running thousands of personalized phishing variants against employee email lists, with copy generated and tailored by an AI model. The volume is high and the quality is high, which means traditional spam filters miss more of them and employees fall for more of them. The downstream effect is that the credential theft rate is climbing — 46% of compromised business credentials in 2025 came from non-managed BYOD devices, and a meaningful share of those compromises started with an AI-crafted phishing message that didn't trigger the user's instinct to be suspicious.

Deepfake-enabled business email compromise. The classic BEC scam — "the CEO is in a meeting, please process this wire urgently" — got a major upgrade. Voice cloning now requires only a few minutes of audio, easily harvested from a podcast appearance, a webinar recording, or a publicly posted earnings call. The result is that a CFO who would have caught a written BEC attempt may not catch a thirty-second voicemail in the CEO's actual voice instructing the same action.

How defenders are using AI

The defender side of the AI story is the part that doesn't get enough airtime, and it's the part most relevant to whether a mid-market company actually emerges from 2026 better off than they started.

Organizations using AI extensively in security operations saved an average of $1.9 million in breach costs and contained breaches 80 days faster than those that did not. That's not marginal. That's a generational gap in defensive capability between organizations whose security partner has integrated AI into their operations and those whose partner is still running 2023's playbook.

The specific places AI is making the biggest defensive difference today:

Threat detection at human-impossible speed. Modern Managed Detection and Response platforms use AI to process billions of log events per day, surfacing the few hundred that warrant human attention. A mid-market company generates more security telemetry than any in-house team can review manually. AI-driven correlation is the only way to find the needle, and it is meaningfully better than it was even twelve months ago.

Faster investigation and response. When an alert does fire, AI agents can perform the initial enrichment and triage work that used to take a tier-one analyst thirty to sixty minutes. The human analyst inherits a mostly-investigated incident with context attached, which lets them make a containment decision faster.

Phishing detection that learns. Email security platforms increasingly use AI to evaluate not just the content of an email but the behavioral pattern of the sender — including subtle anomalies that a rule-based filter would miss. This is one of the most direct counters to the AI-augmented phishing problem above.

How mid-market companies should be using AI — safely

The defensive benefits of AI are one side of the conversation. The other side, and the one most mid-market leaders are actually wrestling with, is how to put AI to work productively across the business — in finance, in operations, in customer service, in engineering — without creating the exposure we just described. This is not a theoretical question. Your employees are already adopting AI tools, your competitors are already using them, and the operational case for getting good at AI is real. The question is how to do it in a way your security posture can defend.

The mid-market companies we see doing this well are following a small set of consistent principles, and we increasingly find ourselves advising clients on this directly:

Sanction the tools, don't suppress the demand. When IT bans AI tools without offering a sanctioned alternative, employees move to unsanctioned tools and the data exposure gets worse, not better. The better pattern is to evaluate, license, and deploy enterprise-grade AI tools with the right access controls, then make those the easy path. Microsoft Copilot, ChatGPT Enterprise, Claude for Work, Gemini for Workspace — each has a meaningful security and governance story that the consumer versions do not.

Classify your data before you connect it to AI. The single biggest mistake we see is companies enabling AI features across their entire data estate without first deciding which data should be reachable by which models. Customer financial records, HR data, legal documents, and regulated information typically need different treatment than internal product documentation or marketing copy. The classification work is unglamorous, but it is the foundation everything else sits on.

Treat AI agents as machine identities. When you wire an AI agent into a workflow — to summarize support tickets, to draft proposals, to route invoices — that agent has access to systems and data on behalf of the humans who deployed it. Gartner's 2026 framing of this as an identity and access management problem is the right one. The agent needs an identity, permissions scoped to the smallest necessary footprint, and audit

logging of what it did and why.

Run a real pilot before a real rollout. The mid-market companies that get the most value from AI are the ones that pilot small, measure actual outcomes (time saved, error rates, employee adoption), and roll out the use cases that demonstrably work. The ones that struggle are the ones that mandate AI use across the company without a meaningful learning loop.

For Dynamic Quest's clients, this work increasingly sits at the intersection of cybersecurity and what we used to call IT strategy. It is not a project. It is an ongoing program — and one that benefits substantially from a partner who is doing it across many client environments rather than figuring it out alone.

The governance gap is the real risk

The most important AI security finding of the year is also the least discussed. Gartner named it the leading 2026 cybersecurity trend: *"Agentic AI Demands Cybersecurity Oversight" — driven by no-code and low-code platforms, "vibe coding," and the rapid proliferation of unmanaged AI agents creating new attack surfaces and unsecured code.* The corresponding IBM finding: 97% of organizations that experienced an AI-related security incident lacked proper AI access controls, and 63% have no AI governance policy at all.

Translated for a mid-market context: your employees are using AI tools right now, with or without your knowledge or sanction. They are pasting client data into ChatGPT to summarize it. They are using Microsoft Copilot, Claude, Gemini, or a dozen other tools to draft documents that contain confidential information. They are wiring AI agents into business workflows through low-code platforms that no one in IT has reviewed. Each of these creates a data path that doesn't exist on any inventory and isn't covered by any policy. Organizations with high levels of shadow AI experienced an average $670,000 in higher breach costs than those with low or no shadow AI usage.

The fix is not to ban AI usage. It is to govern it. That means a documented inventory of which AI tools are sanctioned for which kinds of data, with the same rigor a mid-market company would apply to its CRM or ERP. An access control model that recognizes AI agents as machine identities — a category traditional identity and access management systems were not built to handle, and one Gartner explicitly calls out as a 2026 priority. A data classification scheme that tells employees which categories of company data are appropriate to put into a public AI tool and which are not, in language they can actually apply at their desks. And a monitoring layer that can detect unsanctioned AI usage and flag it for review.

We are working through some version of this conversation with most of our clients right now. The companies that get ahead of it in 2026 will be in materially better shape than the ones that are still treating AI as someone else's problem.

SECTION 3

The Five Threat Patterns You Will Actually Encounter

The 2025 version of this guide listed eight categories of cybersecurity threat. That was useful as taxonomy. It was less useful as a working framework. In practice, mid-market companies are most likely to encounter five threat patterns, and the difference between them is operational rather than theoretical. Each one suggests a different set of controls and a different set of questions to ask the partner managing your environment.

PATTERN 01

Ransomware that starts with an unpatched vulnerability

This is the most common pattern in 2026, and it is most often missed by mid-market companies because it doesn't start with a phishing email or an obvious mistake. It starts with a known vulnerability in an internet-facing system — a firewall, a VPN, a web application, an unpatched server — that the attacker exploited because the patch wasn't applied in time.

Sophos's 2025 manufacturing research found that exploited vulnerabilities are the leading root cause of ransomware attacks on manufacturing organizations, responsible for 32% of incidents. The Verizon DBIR found that exploitation of vulnerabilities as an initial attack vector grew 34% year over year, driven heavily by zero-day exploits targeting perimeter devices and VPNs. The attackers in this pattern are not particularly skilled. They are running automated scans against the public internet looking for the specific software versions that have known flaws, and they are exploiting them faster than mid-market IT teams can patch.

The defenses that matter: a vulnerability management program with documented patch SLAs, externally-validated vulnerability scanning on a recurring schedule, and an engineering partner who treats edge devices (firewalls, VPN concentrators, remote access) as the highest-priority patching targets. Most mid-market companies do not have this in place. They have antivirus and they have backups, and they assume those cover ransomware. They do not.

PATTERN 02

Credential theft from an unmanaged device

The second-most-common pattern, and the one growing fastest. 46% of compromised business credentials in 2025 came from non-managed BYOD devices. The mechanism is usually some flavor of infostealer malware on an employee's personal device — their home laptop, their phone, a contractor's machine — that captures saved business credentials when the employee logs into work systems from that device. Those credentials end up on a credential marketplace, and an attacker buys them and walks in through the front door of your environment with valid login information.

The defenses that matter: enforced multi-factor authentication everywhere, with phishing-resistant methods (hardware keys, FIDO2, number-matching) for privileged accounts. Strict policies about which devices can

access which resources, supported by conditional access enforcement. And visibility into where business credentials are being entered, so an unmanaged device logging into a work resource is a flag rather than a default behavior.

PATTERN 03

AI-enhanced social engineering

Covered in Section 2. The new generation of phishing, voice cloning, and deepfake-enabled BEC. **The defenses** are a combination of technical controls (modern email security platforms with behavioral analysis, mandatory MFA on financial transactions, callback verification protocols for any wire request over a defined threshold) and human controls (regular tabletop exercises that include AI-generated scenarios, training that explicitly addresses voice and video deepfakes, a culture in which employees feel safe verifying suspicious requests with the actual executive).

PATTERN 04

Third-party and supply chain compromise

The fastest-growing category in the 2025 DBIR. The percentage of breaches involving a third party doubled year over year, going from 15% to 30%. Supply chain compromise was the second-most prevalent attack vector in 2025 at nearly 15%, costing $4.91 million on average and taking 267 days to resolve — the longest of any attack pattern.

For mid-market companies, this is a two-sided problem. You are exposed both because your vendors and partners can be a route into your environment, and because if you sell into larger customers, you are part of *their* supply chain risk and they are increasingly evaluating you accordingly. The 80-question security assessments mid-market companies are receiving from enterprise customers are direct downstream effects of this.

The defenses: a vendor risk management program that goes beyond the annual questionnaire, including ongoing monitoring of critical vendors. Network segmentation that limits the blast radius of a compromised vendor connection. And — increasingly — the ability to provide your own customers with the kind of attestation they are asking for, which is its own program of work.

PATTERN 05

Insider risk, intentional or accidental

The most uncomfortable category, and the one mid-market companies are most reluctant to invest in. Human involvement in breaches remains around 60% of all incidents, with significant overlap between social engineering and credential abuse. Most of this is not malicious. It is a well-meaning employee making a well-meaning mistake — sending the wrong file to the wrong distribution list, clicking the wrong link on a bad day, using a personal cloud storage account for client data because the official channel is too slow.

The defenses: data loss prevention controls that detect when sensitive data is moving in unexpected ways, role-based access that limits any one employee's blast radius, regular training that emphasizes

psychological safety in reporting near-misses, and a security culture that does not punish employees for raising concerns. Mid-market companies that build this culture genuinely well end up with a defensive advantage that is hard to buy.

SECTION 4

Your Cyber Insurance Renewal Is the New Annual Audit

This is the section every mid-market COO and CFO should read first. The shift in cyber insurance underwriting between 2023 and 2026 is the single biggest practical change in how mid-market cybersecurity actually gets bought, configured, and proven.

What changed

Five years ago, applying for cyber insurance was a checkbox exercise. Today it is a security audit with financial consequences. According to Marsh McLennan's 2025 Cyber Insurance Market Report, 99% of cyber insurance applications now include specific questions about MFA implementation. Carriers are no longer accepting attestation. They are asking for screenshots of your MFA enforcement settings, EDR coverage reports, backup test logs, incident response plan documents with documented test dates, and patch compliance evidence. The questionnaire is the audit.

The downstream effects are substantial. If you cannot demonstrate the required controls, your insurer may increase your premiums significantly — sometimes 30% to 50% over the previous year. They may exclude ransomware coverage entirely, removing the single most common and most expensive type of claim. In some cases, they will deny renewal altogether. There is also claim denial risk: a January 2026 case involved a mid-size accounting firm whose ransomware claim was denied because the controls in place did not match what was represented on the application. Coalition's data shows 82% of denied claims involved organizations without MFA.

The eight controls carriers ask about

The list has stabilized. If your environment can credibly demonstrate the following, your renewal will go meaningfully better than if it cannot. If your environment cannot demonstrate them, you have your roadmap.

01 **Multi-factor authentication, enforced and documented.**
Not "available" — enforced. Phishing-resistant where possible (hardware keys or FIDO2 for privileged access). Universal across email, VPN, remote access, and administrative accounts.

02 **Endpoint Detection and Response on every endpoint.**
Traditional antivirus is no longer sufficient. EDR is the baseline expectation, with documented coverage reports and an incident the carrier can verify was investigated and resolved.

03 **Email security with behavioral analysis.**
Spam filtering alone is insufficient. Carriers are increasingly asking about mailbox-level email security that can detect business email compromise, social engineering, and AI-augmented phishing.

04 **Tested, immutable backups.**
The standard now is immutable or offline backup copies, with documented restore tests. Coalition's 2026 Cyber Claims Report found that 86% of ransomware victims refused to pay ransoms in 2025 — a record high, up from 50% two years ago — with improved backups and tested incident response plans cited as the primary reason.

05 **A written incident response plan, tested at least annually.**
Not a document on a shared drive. A document that has been exercised in a tabletop scenario, with documented findings and remediation actions closed.

06 **Patch management with documented SLAs.**
Critical patches within 30 days is the standard most underwriters expect, with vulnerability scan reports demonstrating the age of any open findings.

07 **Privileged Access Management.**
Especially for organizations above 100 employees. Carriers want to see that admin accounts are managed differently from standard user accounts, with stronger authentication and audit logging.

08 **Security awareness training with measurable outcomes.**
Annual training is no longer sufficient. Carriers want ongoing training with simulated phishing campaigns and trend reports showing employee click-through rates declining over time.

What the audit costs you if you fail it

The math is straightforward and worth running. A mid-market manufacturer with $50M revenue carrying a $5M cyber liability policy might have paid $30,000–$50,000 annually two years ago. If they cannot demonstrate the controls above today, they are likely looking at a 30% to 50% premium increase, a higher deductible, ransomware sublimits that may cap recovery at $250K, or non-renewal. Compounding effects: the same controls the insurance carrier is asking for are increasingly required to win business from larger customers. The same controls are required to maintain regulatory compliance. The investment to satisfy the insurance audit is the same investment that satisfies the customer audit, the regulatory audit, and the actual security need.

This is the strategic reframe that matters. Your cyber insurance renewal is not a paperwork exercise. It is a forcing function for the security program you already needed. The mid-market companies we work with that approach the renewal as a strategic event rather than a procurement event almost always end up with a better security posture and a better-priced policy. Those that approach it as paperwork are the ones writing checks for premium increases or shopping for new carriers under deadline pressure.

> *"The framing on cybersecurity has changed. It used to be insurance against an unlikely event. Today it's table stakes for being a credible counterparty — to your customers, to your insurance carrier, to your regulators, and to your investors. The companies treating it as the cost of doing business in 2026 are the ones positioned to keep doing business in 2027."*
>
> **John Guillaume**
> CEO, Dynamic Quest

SECTION 5

What a Real Cybersecurity Program Looks Like at Mid-Market Scale

There is a tendency in cybersecurity content to either go too generic ("you need defense in depth") or too granular ("here are 247 controls to implement"). Neither is useful to a mid-market leader trying to make a real decision about a real budget against a real threat picture. This section is the working framework we use with clients.

A defensible mid-market cybersecurity program in 2026 has three components, in this order of priority: the people accountable for it, the platform that supports it, and the plan that ties them together. Most mid-market companies have a partial version of all three. Very few have all three running in coordination.

People

The first question is not "what tools do we need" but "who is accountable for security in our organization, and what are they accountable for." For most mid-market companies, the answer is some combination of an internal IT generalist, an external partner, and an executive sponsor — typically a COO, CFO, or VP of Operations.

What matters is that the accountabilities are clear. Who is responsible for patch management on edge devices? Who reviews EDR alerts overnight? Who decides when an incident becomes a security event that triggers the response plan? Who owns the relationship with the cyber insurance broker? Who responds when a customer sends a 60-question security assessment? In environments where these answers are clear and documented, response times are fast and renewals are smooth. In environments where they are not, the answers default to whoever picks up the phone first, which is rarely the right person.

For most of our clients, the right people structure includes an executive sponsor at the leadership level (typically a COO or CFO) who owns budget, vendor relationships, and the strategic conversation; a technical lead — internal or virtual — who owns day-to-day security operations and coordinates with the MSP partner; a trained Managed Detection and Response team handling 24/7 monitoring and threat response, because no internal team at mid-market scale can credibly cover overnight and weekends; and a documented escalation path to senior incident response specialists when something serious surfaces.

The mistake mid-market companies make most often here is assuming their existing IT generalist can also be the security lead. They cannot. The skillsets are different, the time demands are different, and the tools are different. Cybersecurity at mid-market scale needs dedicated attention, even if that attention is sourced from a partner rather than a hire.

Platform

The platform layer is where the eight controls from Section 4 actually live. The cyber insurance carriers are not asking for these controls because they invented them. They are asking because their actuarial data shows that organizations with these controls in place have meaningfully lower claim frequency and severity.

The same controls satisfy customer audits, regulatory requirements, and the actual security need.

A defensible mid-market platform stack in 2026 includes multi-factor authentication, enforced and phishing-resistant where possible; Endpoint Detection and Response on every endpoint, with 24/7 monitoring and human response; modern email security with AI-driven behavioral analysis; immutable backup with documented restore testing; a hardened identity layer with conditional access and privileged access management; network segmentation that limits the blast radius of any single compromise; vulnerability management with documented patch SLAs; and data loss prevention and AI governance tooling appropriate to the data classification work in Section 2.

The right answer is rarely a single vendor's full stack. It is usually a curated set of best-in-class tools integrated and managed by a partner who knows how to make them work together. The integration work is where most mid-market companies struggle — they end up with three tools that overlap in some places and leave gaps in others, and no one is accountable for the gaps.

Plan

The plan is what turns the people and the platform into a program. This is the layer most mid-market companies have least developed.

The minimum viable plan includes a documented incident response playbook with named roles and an escalation tree, tested annually in a tabletop exercise with leadership in the room; a vulnerability and patching SLA committed to in writing, with monthly reporting against it; a defined process for security event declaration — when does an alert become an incident, who decides, and what happens next; a vendor risk management process that goes beyond the annual questionnaire; a data classification scheme that translates into who can access what, where, and through which AI tools; and a quarterly business review with the security partner that covers threats observed, remediations completed, controls maturity progress, and changes for the next quarter.

The plan is also where the cybersecurity program connects to the rest of the business. The CFO needs to know how the security investment maps to the insurance renewal. The COO needs to know how it maps to customer audits. The CEO needs to know how it maps to the board's risk register. A program that runs in isolation from those conversations is a program that gets cut in the next budget cycle.

What this costs at mid-market scale

A defensible cybersecurity program for a mid-market company typically runs 1.5% to 3% of revenue annually, with the higher end of the range applicable to companies in regulated industries (defense supply chain, healthcare, financial services, accounting) and companies with elevated supply chain or contractual exposure. That investment has crept up roughly 30 to 50 basis points over the last two years, driven by AI-enhanced platform pricing, tighter insurance underwriting, and longer hardware lead times.

For a $50 million manufacturer, that benchmark range is roughly $750,000 to $1.5 million annually across all security spending — internal labor, external partner fees, software licensing, hardware refresh, and insurance premiums. For a $25 million accounting firm in the FTC Safeguards Rule scope, the same benchmark applies and likely indexes toward the high end given the regulatory exposure and the sensitivity of client data.

Compared against an average U.S. breach cost of $10.22 million, a 30 to 50 percent insurance premium increase, or the loss of a major customer relationship after a failed security questionnaire, this remains overwhelmingly the right side of the math. The mid-market companies making the most defensible budget decisions in 2026 are the ones treating cybersecurity as an operating cost line tied to their license to operate, not a CapEx project tied to a budget cycle.

SECTION 6

Three Questions to Ask Any MSP You're Evaluating

The 2025 version of this guide included a long checklist of MSP traits. The checklist was accurate. It was also the kind of document that a buyer reads, nods at, and then doesn't know how to actually use in a vendor conversation.

Mid-market companies evaluating a cybersecurity partner — whether that's a new MSP relationship or a renewed conversation with the current one — would do better to ask three specific questions and watch carefully how the partner answers. The answers tell you almost everything you need to know about whether the partner is operationally serious or selling a brochure.

> *These are the questions we ask ourselves about our own operations, and the questions we expect prospective clients to ask us. A serious cybersecurity partner welcomes them.*

QUESTION 1

"Walk me through how you'd respond to a ransomware event in our environment in the next 90 minutes."

This question reveals whether the partner has actually done this work. A serious answer includes specifics: the detection mechanism that would surface the event, the on-call team that would respond, the containment steps in the first thirty minutes, the communication protocol with leadership, the backup verification process, and the threshold at which legal counsel and the cyber insurance carrier are notified. A serious partner can describe the last time they actually executed this playbook and what they learned from it.

A weak answer is generic. It uses words like "we have a process" and "our team is highly trained" without naming specific actions, specific timelines, or specific people. If the partner cannot describe a real response in real time terms, the partner has not done this real work.

QUESTION 2

"Show me what your last quarter's threat detection and response data looked like for a client our size."

This is the operational evidence question. Any partner managing security for mid-market clients should be able to share — with appropriate anonymization — what their detection volume, response time, and remediation outcome data look like. Specifically: how many alerts were investigated, how many were

escalated, how many turned into actual incidents, what the median time-to-containment was, and what the trend lines look like quarter over quarter.

A weak answer dodges the specifics or pivots to marketing material. A serious answer includes real numbers, real time periods, and a willingness to discuss what didn't go well alongside what did. The partner that can show you their actual operational data is the partner whose operational data is worth showing.

QUESTION 3

"What's your roadmap for AI-powered security operations, and how does it map to my environment?"

The third question is the future-readiness question. Given the AI shifts described in Section 2, the partner you choose in 2026 needs to have a credible answer for how their security operations are evolving — not as a marketing claim, but as an operational reality.

A serious answer addresses both halves of the AI question: how the partner is using AI to improve their own detection and response capabilities (the $1.9 million breach savings opportunity from Section 2), and how the partner is helping clients govern AI usage in their own environments. A serious partner is doing this work today, can name specific platforms and capabilities they have deployed, and can describe how those capabilities would apply to your specific environment.

A weak answer talks about AI in the future tense — "we're evaluating," "we're planning to," "AI is on our roadmap." Eighteen months ago, that was a defensible answer. In 2026, it is the answer of a partner who is behind.

What you should be testing for, behind the questions

The three questions are designed to test for three things that matter operationally: actual incident response capability, real operational discipline, and forward-looking technical investment. A partner that answers all three credibly is a partner you can build a program around. A partner that struggles with one or more is telling you something important about where the gaps will be when you need them not to be.

SECTION 7

Vertical Considerations

Cybersecurity is universal. The threats, controls, and program elements covered in the preceding sections apply across every mid-market industry. But the regulatory environment, the customer expectations, and the specific attack patterns vary meaningfully by vertical, and a few targeted notes are worth attaching to the general framework.

Manufacturing (general)

Mid-market manufacturers face the highest concentration of ransomware activity in 2026, and the data on root cause is unusually clear: exploited vulnerabilities are the leading root cause at 32% of incidents, with malicious email at 23% and credential-based attacks at 20%. The single most important investment for a mid-market manufacturer is a vulnerability management program with credible patch SLAs on internet-facing infrastructure. Beyond that, the operational reality of manufacturing — interconnected systems, low tolerance for downtime, expensive supply chain ripple effects — makes incident response speed disproportionately valuable. Sophos found that 58% of manufacturers fully recovered within one week in 2025, up from 44% the prior year, suggesting that organizations with mature response capabilities are pulling away from those without.

Trigger event watch list: cyber insurance renewal, ERP modernization, customer security questionnaire, departure of the IT generalist who has been holding things together.

Defense Supply Chain (CMMC)

If your manufacturing or services business sells into Department of Defense contracts and handles Federal Contract Information or Controlled Unclassified Information, you are now in a regulatory environment that has fundamentally changed. The CMMC final rule took effect November 10, 2025, and the phased rollout extends through 2028. The cybersecurity controls required for CMMC compliance are substantial and the timelines for achieving them are not generous.

This guide is the right starting point for the broader cybersecurity conversation, but CMMC deserves its own deeper treatment. Dynamic Quest publishes a separate **CMMC Readiness Guide** that covers the certification levels, the phased timelines, the gap analysis process, and what a typical mid-market defense supplier needs to do to be eligible to bid in 2026 and beyond. If CMMC applies to you, request the CMMC Readiness Guide from your Dynamic Quest contact.

Large Accounting Firms

Mid-market accounting firms — particularly those above 100 staff — face a uniquely concentrated threat environment. The IRS Security Summit consistently reports that cyberattack attempts against tax

professionals spike by roughly 50% between January and April. Effective January 1, 2024, the AICPA updated its Statements on Standards for Tax Services with the addition of Section 1.3, which explicitly requires CPAs to take reasonable steps to safeguard taxpayer data and document those protections. The FTC Safeguards Rule applies. Cyber insurance carriers are denying claims when firms cannot produce documented WISPs and current risk assessments.

For firms still hosted on Rightworks or other legacy hosted-desktop platforms, performance during tax season has become a meaningful competitive concern, and the security posture of these platforms is a legitimate question for firm leadership to investigate. Firms exploring Azure Virtual Desktop or modern alternatives are doing so for both performance and security reasons, and the conversation is increasingly happening at the managing partner level rather than buried in IT.

Healthcare and Medical Practices

Multi-specialty mid-market healthcare practices operate at the intersection of HIPAA compliance, patient safety, and the highest-cost breach environment of any U.S. industry. Healthcare remained the highest-cost industry for data breaches in 2025, at $7.42 million on average — even after a 24% year-over-year reduction. The combination of patient data sensitivity, vendor ecosystem complexity (EHR vendors, billing services, imaging services, lab integrations), and regulatory exposure means healthcare programs need particular attention to vendor risk management and data flow mapping.

Two regulatory developments make this vertical's 2026 picture meaningfully different than 2025. The first is the proposed update to the HIPAA Security Rule. On January 6, 2025, HHS published a Notice of Proposed Rulemaking representing the most significant update to the rule in over a decade. The proposed changes would remove the distinction between "required" and "addressable" implementation specifications and make all controls mandatory, with new express requirements including encryption of ePHI at rest and in transit, multi-factor authentication, vulnerability scanning every six months, annual penetration testing, network segmentation, and separate technical controls for backup and recovery. The final rule has been delayed by the regulatory freeze and may be issued in modified form in 2026, but the direction of travel is unambiguous. Healthcare practices that wait for the final rule before acting will be behind. Those that align to the proposed framework now will be ready when the rule lands.

The second is the Cyber Incident Reporting for Critical Infrastructure Act, or CIRCIA. CIRCIA, signed into law in 2022, establishes the first comprehensive federal cyber incident reporting mandate spanning 16 critical infrastructure sectors including healthcare. Once the final rule takes effect — currently anticipated in 2026 — covered entities will be required to report substantial cyber incidents to CISA within 72 hours and ransomware payments within 24 hours. For healthcare organizations accustomed to HIPAA's 60-day breach notification window, this is a fundamentally different operational cadence. CIRCIA's clock starts when your team suspects something significant happened, not when forensics complete. Importantly, CIRCIA does not replace HIPAA reporting — it adds to it. Healthcare organizations will be operating under both reporting frameworks simultaneously.

The most overlooked control area in mid-market healthcare is third-party access. Vendors with privileged access to patient data and billing systems are a frequent route into healthcare environments, and many practices do not have an inventory of who has access to what. The combination of new HIPAA requirements, CIRCIA reporting timelines, and ongoing vendor risk creates a meaningful operational burden that

mid-market healthcare practices typically cannot absorb without a partner.

Financial Services and Credit Unions

Mid-market credit unions and community financial institutions face a regulatory environment that is more demanding than most other verticals — NCUA examinations, FFIEC guidance, state-level requirements, and increasingly stringent member-data protection expectations. The threat environment is also concentrated: financial services attackers are predominantly financially motivated, with extortion-only and double-extortion ransomware patterns becoming more common.

The operational reality for most mid-market credit unions is that their existing IT staff is fully consumed with member-facing systems, leaving cybersecurity as either an under-resourced internal function or a function fully outsourced to a partner. The partner relationship in this vertical is unusually high-stakes because the regulatory examiners will ask for evidence that the partner is performing at the level required, and that evidence is the credit union's responsibility to produce.

State and Local Government / Municipal

Mid-market government entities — counties, municipalities, school districts, public utilities — operate under unique constraints: legacy systems, restricted budgets, public-sector procurement timelines, and an attack environment that has grown sharply in the last 24 months. Municipalities are now among the most targeted ransomware victim categories, and the public-sector procurement cycle frequently does not match the speed at which security investments need to be made.

The most pragmatic path for mid-market government is an outsourced partner relationship that brings the controls, the monitoring, and the documented response capability inside an operational structure that procurement can defend. The technical answer is the same as for private-sector mid-market. The procurement and governance answer is different and needs careful attention.

ABOUT DYNAMIC QUEST

A Technology Partner, Not a Vendor

Dynamic Quest is a Managed Services Provider headquartered in the Southeast U.S. with deep market presence across eight cities — Greensboro, Charlotte, Raleigh, Huntsville, Birmingham, Tampa, Orlando, and Jacksonville — and clients across the country. We serve mid-market organizations nationally and are actively expanding our footprint, supported by U.S.-based service teams and a Philippines operations center providing 24/7 follow-the-sun coverage. Our blended onshore-offshore staffing model is particularly well-suited to clients with after-hours monitoring requirements, government and defense supply chain compliance needs, and round-the-clock operations.

We work with manufacturers, defense suppliers, accounting firms, healthcare practices, credit unions, and professional services organizations across the verticals described in Section 7. Our practice areas include managed IT, cybersecurity, cloud and infrastructure, compliance support, and virtual CIO advisory services.

Our cybersecurity practice is led by certified security professionals operating against frameworks including NIST 800-171, CMMC, HIPAA, and FTC Safeguards. We deliver Managed Detection and Response, vulnerability management, identity and access management, security awareness training, and incident response — supported by quarterly business reviews and ongoing advisory engagement at the leadership level.

We are backed by Spire Capital and structured for sustained growth across our footprint. We are large enough to deliver enterprise-grade security operations, and small enough that our clients know our team by name.

What sets Dynamic Quest apart is the way we engage. We are a technology partner, not a vendor. We build long-term relationships with the leadership teams we serve, we are honest about what is working and what is not, and we measure ourselves on the operational outcomes our clients can actually demonstrate to their insurance carriers, their customers, and their boards.

NEXT STEP

Let's Start the Conversation

If you read this guide and recognized your environment in any of the four shifts, the five threat patterns, or the cyber insurance audit framing, the right next step depends on where you are.

Strategic Conversation

A discovery call with a senior member of our team to discuss your current setup, the compliance or insurance pressure you're navigating, and whether Dynamic Quest is the right fit for the work ahead. We'll talk through your situation, share our perspective, and if there's mutual fit, scope an engagement together. **The right starting point for most conversations.**

Tailored Assessment

A scoped engagement that delivers a documented gap analysis against your specific environment, prioritized findings, and a remediation roadmap. **Custom-priced based on environment complexity.** Appropriate when you have a specific audit, renewal, customer requirement, or contract triggering urgency.

CMMC Compliance

If you handle Federal Contract Information or Controlled Unclassified Information in DoD contracts and know you need CMMC Level 2 certification, our compliance practice offers a policy and process gap analysis to identify what's required for certification readiness. **Custom-priced based on certification level and environment scope.**

We work best with leadership teams who are ready to invest in a defensible technology program — cybersecurity, cloud, AI governance, and the operational systems that run your business — and who want a proactive technology partner, not a vendor.

We are honest with prospects about whether we are the right partner for their situation. If we're not, we'll tell you. If we are, we'll move quickly to scope the work and get to value.

To start the conversation

dynamicquest.com/contact

or call **(833) 437-8378**

— The Dynamic Quest Team

SOURCES CITED

Research & References

This document draws from the most recent and credible cybersecurity research available as of Q2 2026:

Gartner, Inc.	Forecast: Information Security, Worldwide, 2023–2029 (2025 updates); Top Trends in Cybersecurity for 2026 (February 5, 2026)
IBM	Cost of a Data Breach Report 2025 (Ponemon Institute research, 20th annual edition)
Verizon	2025 Data Breach Investigations Report
Sophos	State of Ransomware in Manufacturing and Production 2025
Marsh McLennan	2025 Cyber Insurance Market Report
Coalition	2026 Cyber Claims Report
AICPA	Statements on Standards for Tax Services, Section 1.3 (effective January 1, 2024)
IRS Security Summit	Tax professional cyberattack reporting
U.S. Department of Health and Human Services	HIPAA Security Rule Notice of Proposed Rulemaking (January 6, 2025)
CISA	Cyber Incident Reporting for Critical Infrastructure Act (CIRCIA) framework

For additional cybersecurity guidance, visit the Cybersecurity and Infrastructure Security Agency (CISA) at **cisa.gov**.