



Buying Resilience

How Risk Managers Can Close the Gap Between the CISO and the Boardroom

James Tuplin, SVP and Head of International Cyber, Mosaic;
Jay Vinda, Global CISO and Cyber Risk Engineering Lead, Mosaic;
Sanith Wijesinghe, Chief Engineer, Public Sector, MITRE;
Chris Folk, Director, Cyber Solutions Cell, MITRE

“

The emergence of frontier AI models has triggered a debate across the cyber security and risk management communities – what happens when attackers gain access to systems capable of discovering vulnerabilities, generating exploits and reasoning through attack paths at unprecedented speeds. While this is a reasonable debate to be having, this framing has the risk of overstating what is actually changing.

This Thought Leadership Paper Series is a joint project from Mosaic Insurance and MITRE.

Introduction | The New Economics of Cyber Risk

How CISOs, Insurers, and Risk Managers Respond to the Age of Frontier Artificial Intelligence

Paper 1 | Beyond Patch Velocity

Rethinking Cybersecurity Strategy in a Post-Frontier AI World

Paper 2 | Beyond Vulnerability Counts

Reimagining Cyber Insurance Underwriting in a Post-Frontier AI World

Paper 3 | Buying Resilience

How Risk Managers Can Close the Gap Between the CISO and the Boardroom

Buying Resilience

How Risk Managers Can Close the Gap Between the CISO and the Boardroom

Executive Summary

Your CISO has a risk assessment your insurer has never seen. Your insurer has policy exclusions your CISO has never read. Your broker knows both exist but does not necessarily connect them. Nobody in that chain is doing anything wrong, which is exactly why the disconnect goes unnoticed until a claim is already in dispute.

That arrangement got considerably more expensive in April 2026. Three public demonstrations within eight weeks, Anthropic's Project Glasswing, the ExploitGym benchmark, and Microsoft's MDASH system, showed that frontier AI can find high-severity vulnerabilities faster than traditional human review processes and turn what it finds into working exploits in hours. When Glasswing was announced, fewer than 1% of the vulnerabilities it found were patched. Most cyber programs are still priced on three assumptions: that exposure is reasonably visible, that patches close risk on a manageable timeline, and that the questionnaire gives an adequate picture. None of the three survived the spring.

There is a window, and it is closing on a published schedule. Anthropic's own interim Glasswing update, released in May 2026, reported more than ten thousand high- or critical-severity findings across roughly fifty partners, of which 530 had been disclosed to maintainers, 75 patched, and 65 carried public advisories. The disclosure embargoes behind the rest expire through the second half of 2026. As that reporting lands, every insurer, broker, and risk manager in the market will be having this conversation at the same time. Some organizations will have restructured already. Most will not. The market has repriced three times in the past decade, and the pattern is well documented. What is unusual this time is that the event is visible before it closes the window.

None of what follows is a threat briefing. It is four changes to how the risk function operates, one for each of the four places the current model breaks, and none of them requires a budget line, a new hire, or a reorganization. Any risk manager can start any of them on a Monday morning without asking permission.

The risk manager is also the only person in the building with standing on both sides of the gap. A CISO cannot walk into the CFO's office and reopen the insurance budget. A CFO cannot interrogate the CISO's risk assessment without stepping outside the CFO's lane. The broker has relationships with both and internal context in neither. That leaves one person to do the translating.

Buying Resilience

How Risk Managers Can Close the Gap Between the CISO and the Boardroom

Where Paper 2 Left Off

“Beyond Vulnerability Counts,” the second paper in this series, made the case that the environment against which insurance is priced has shifted on every axis at once. Organizations are rethinking what software they are willing to run at all.

Security teams are rebuilding operating models around AI-augmented defense. Legal is treating discovery findings as a new category of material information, and supply chain is reopening vendor risk. The insurance conversation sits at the end of that chain, which makes it the one most likely to still be priced on assumptions the spring called into question.

That paper reached a specific underwriting conclusion: remediation velocity is the signal worth pricing, not KEV presence. An insurer that limits coverage because a KEV-listed vulnerability happened to be present at the time of a breach is charging the insured for knowing. That distinction is where this paper picks up.

A Familiar Market Pattern

The cyber insurance market has repriced three times in a decade. NotPetya in 2017, when premiums rose as the market absorbed nation-state capability deployed at scale; the ransomware surge of 2020 and 2021, when premiums rose again, coverage narrowed, and exclusions multiplied; and Log4Shell at the end of 2021, when supply chain exposure stopped being theoretical. Every time, the repricing came after the losses. Every time, the buyers who had structured their programs in advance came out of it better than those reacting to conditions that had already changed.

A CFO watching premiums fall since 2022 has a fair objection, and it is grounded in real data: Howden puts global cyber rates roughly 22% below the 2022 peak, and roughly 6% lower in 2025 than in 2024. That objection describes where the market has been. The forecast describes where it is going. S&P Global Ratings expects premium increases of 15% to 20% in 2026, citing rising claim severity, growing data theft, and the cost of AI-driven attacks. The soft market is ending.

What makes this cycle different is that the evidence arrived before the losses. First: Anthropic’s Project Glasswing found thousands of high-severity vulnerabilities across every major operating system and browser in a matter of weeks, including a flaw in FreeBSD that had survived 17 years of expert human review. Fewer than 1% were patched when it was announced.¹

Second: the

ExploitGym benchmark from Berkeley, Anthropic, OpenAI, and Google showed frontier agents taking a known vulnerability

to a working exploit in under two hours.²

Third: Microsoft’s MDASH system, running in production, found 16 new vulnerabilities in the Windows networking and authentication stack in one Patch Tuesday cycle, four of them Critical remote code execution flaws. In a separate retrospective test, it rediscovered 96% of 28 Microsoft Security Response Center (MSRC) cases in clfs.sys and 100% of 7 in tcpip.sys, spanning five years of confirmed findings in components Microsoft reviews more heavily than almost any other.³

The cyber insurance market has repriced three times in a decade.

Buying Resilience

How Risk Managers Can Close the Gap Between the CISO and the Boardroom

It has not slowed down. On June 24, curl shipped the largest security patch in its history, 18 CVEs in one release, set off by Mythos finding a single vulnerability in May. The oldest of them had been sitting in the code since 2001. Curl runs on something north of 30 billion devices, and most of the organizations running those devices have no way to know the library is there, let alone control when it gets patched.⁴ IBM's 2026 X-Force Threat Intelligence Index found vulnerability exploitation to be the leading cause of attacks, at 40% of observed incidents.⁵ M-Trends 2026 puts the mean time to exploit at negative seven days, which is a polite way of saying the average vulnerability is being exploited before the patch exists.⁶

So the organizations inside the Glasswing cohort now know more about their own exposure than the people insuring them do. The questionnaire that went out last year captured none of it. That is not a pricing nuance. It is the setup for a coverage dispute at scale.

That gap belongs to the risk manager. Nobody else can carry a moving technical picture into policy terms, financial protection, and a board conversation. The job is to make current risk visible before the claim, the renewal, or the repricing makes it visible by force.

The Failure: A Missing Conversation

Consider the following composite claim scenario:

The Claim

A ransomware group encrypts 40% of an organization's production environment. The CISO had flagged the specific

attack vector, a chained exploit across two unpatched edge devices, in a risk assessment eight months earlier. The risk manager had not seen that assessment. The broker had not asked for it. The insurance policy contained a policy-limiting trigger for vulnerabilities that appeared on the Cybersecurity and Infrastructure Security Agency's (CISA's) Known Exploited Vulnerabilities (KEV) catalog and had not been remediated within 90 days. One of the two devices had a KEV-listed vulnerability. The claim is disputed. The organization is in litigation with its insurer while simultaneously managing the breach. This is precisely the failure mode described above: treating KEV presence as grounds for a policy-limiting trigger rather than a prioritization signal and discovering the distinction at the worst possible moment.

The CISO in that scenario did the right thing. The assessment was accurate, and it was written down. The insurer arguably did the right thing too: the terms were disclosed, and the KEV endorsement was not unreasonable on its face. The broker did what brokers do: matched a client to available coverage at a competitive price. Every decision was defensible on its own. The outcome was still bad, because nobody connected them.

That is the gap not technical failure, not bad faith, but an organizational architecture that keeps the people who understand cyber risk and the people who buy financial protection against it in separate rooms until the moment they should have been in the same one.

Buying Resilience

How Risk Managers Can Close the Gap Between the CISO and the Boardroom

CISO

Knows the actual threat. Understands the exploitability of the software the organization runs. Loses sleep over what a motivated attacker could chain together from the known vulnerability inventory.

The Missing Conversation

No structured channel exists between the CISO's actual risk picture and the insurance program. They do not talk until a loss event forces them to. By then, the coverage question has already been answered by the claim.

Risk Manager/CFO

Buys what the broker recommends. Optimizes for premium. Has no reliable way to know whether the coverage reflects the CISO's actual concerns or the broker's benchmark.

This gap has three structural causes.

- **Language:** The CISO talks in exploit chains and vulnerability identifiers, the broker in limits and exclusions, the CFO in ratios and expected loss. There is no translation layer anyone uses consistently.
- **Incentives:** The CISO is measured on zero material incidents. The CFO is measured on cost-effective coverage. The objects are related, but not the same, and the difference surfaces the moment a tradeoff has to be made.
- **Timing:** Renewals are annual, the CISO's risk picture is continuous, and almost nothing triggers an update in between. Therefore the translation work goes undone. Not because anyone refused it, but because nothing on the calendar asks for it.

From Risk Transfer to Resilience Partnership

The traditional model is transactional. The insured works out its exposure, buys a policy, and hands a defined slice of the financial consequence to the insurer. The insurer prices what it can see: questionnaires, external scans, industry benchmarks and claims history. Both sides proceed on the assumption that coverage is roughly proportionate to risk. That assumption is doing a great deal of work.

Glasswing, ExploitGym, and MDASH each put pressure on it. Coverage that was adequate six months ago may have been priced against a risk picture that has since moved, and the person responsible for saying so out loud, before a claim proves it, is the risk manager.

Resilience partnership is the alternative organizing principle. Under risk transfer, success is a policy that pays out. Under resilience, success is an organization that is harder to breach, quicker to recover, and able to walk into the next renewal with evidence instead of attestations. Insurers want fewer claims. Insureds want fewer claims. That is more common ground than either side usually trades on.

The risk manager who approaches renewal asking, "What is the cheapest coverage that meets the minimum requirement?" is buying risk transfer. The one who asks, "What does this insurer offer that makes us more resilient, and how do we demonstrate that improvement to get pricing that reflects it?" is buying a partnership. The difference shows up at renewal, at claim, and in the boardroom conversation about whether the cyber budget is being spent on the right things.

Buying Resilience

How Risk Managers Can Close the Gap Between the CISO and the Boardroom

What Buying Resilience Actually Means

Three parts of the relationship have to be structured differently from the way most organizations handle them now.

The Financial Safety Net

The financial safety net is what the risk manager already buys, and it is not going anywhere. First-party coverage handles breach response, business interruption, extortion, and recovery. Third-party coverage handles regulatory fines and liability to customers and partners. The foundation is sound. How it gets structured is the problem.

Under risk transfer, you buy the broadest coverage the budget allows and hope it stays in the drawer. Under resilience, the coverage is built around a specific understanding of what the loss scenarios actually look like, grounded in the CISO's read of exposure rather than the broker's benchmark. That means having the conversation that mostly does not happen: which components do we run, which of them carry the highest exploitation risk right now, and do the policy terms speak to those scenarios or to generic ones.

The claim scenario above went wrong partly because nobody connected the KEV limiter in the policy to the CISO's knowledge that a KEV-listed vulnerability was sitting on an exposed device.

Pre-Negotiated Incident Response Access

This is the one that shows up in post-breach reviews as the thing people wish they had done differently. Pre-negotiated access to an incident response panel through the insurance relationship is different from finding a firm on the spot market while a breach is underway.

The organizations that recover fastest are rarely the ones with the largest limits. They are the ones with vendors already under contract, protocols already tested, and a playbook they can run instead of invent. Fourteen days to containment versus forty-five usually comes down to preparation, not budget.

The new environment raises the stakes. When an agent can go from a single known vulnerability to full system compromise in under two hours, whether the incident response firm answers in hours or days changes the outcome materially. Treat the vendor panel as a primary purchase criterion rather than a footnote in the quote comparison.

Access to Services That Build Actual Resilience

Some insurers and brokers now bundle security services: vulnerability scanning, tabletop exercises, awareness training, and threat intelligence feeds. These are worth having. They are not resilience. For many organizations, they duplicate something the security function already runs, which may cut costs or improve efficiency without touching the gap that actually matters. The question to stop asking is, "What cyber services are included in our policy?" The question to start asking is, "Which of these will help us execute our security strategy faster?"

That means going to the CISO first. Where is the program constrained? Which capabilities is the CISO trying to mature? What got deferred when the budget got tight? Where is outside expertise the only way forward? Which improvements would actually move business risk?

The best insurance relationship stops being a catalog and becomes a match between what the

Buying Resilience

How Risk Managers Can Close the Gap Between the CISO and the Boardroom

insurer can reach and what the CISO is already trying to do. The CISO gets capability that was never going to be funded. The risk manager gets measurable movement toward risk appetite. The insurer lowers its odds of writing a large check. Nobody is doing anybody a favor.

Under risk transfer, these services are a marketing differentiator. Under resilience, they are chosen to close a gap the CISO has already named.

The services that matter now are the ones that close what the adversary has opened. Has anyone run AI-scale discovery against our own repositories? Can we tell whether the vulnerabilities we know about chain into a viable path, or do we only know they exist? Do we have continuous visibility into the software we run, or are we working from snapshots? Those questions decide whether the risk posture reflects this year's threat environment or last year's.

Four Failure Points and the Fix for Each

The organizations that are easiest to work with at renewal are the ones where the risk manager has clearly had a detailed conversation with their CISO before they walked in the door. They know what their organization runs, they know where the exposure is concentrated, and they can tell us specifically what has changed since last year. When someone walks in with that, we can price it effectively. We can build a program around actual risk rather than benchmarks. That is the kind of relationship where both sides get a better outcome."

—Jay Vinda, Mosaic Insurance

That conversation is not happening in most organizations. The sections above cover what to say in it. This one covers why it keeps failing to happen.

The analogy that follows comes from banking, where the function this paper calls the risk manager sits at the executive level as a chief risk officer. The title differs. The structural position, translating a moving risk picture into decisions the business acts on, is the same. In April 2022, the chief risk officer (CRO) at Silicon Valley Bank (SVB) stepped down. The seat stayed empty for eight months. Interest rate risk and liquidity pressure were building across the balance sheet the entire time. Federal Reserve supervisors had been citing risk management deficiencies since 2021. The board's risk committee met 18 times in 2022, more than double the seven meetings it held in 2021, and was widely assessed afterward as thin on large-bank risk expertise. The bank failed in March 2023. The Fed's own postmortem said supervisors had the authority to cite the vacancy as a violation of enhanced prudential standards but waited instead, allowing the hole in oversight to remain open through the most consequential eight months in the bank's history.

JPMorgan Chase operated in the same rate environment and came through it. The difference was not balance sheet size. It was a risk function that was in the room continuously, with standing authority over strategic and operational decisions, rather than a reporting exercise that woke up at audit time.

Buying Resilience

How Risk Managers Can Close the Gap Between the CISO and the Boardroom

SVB is not a cyber story. The structure transfers exactly. When a risk category starts moving faster than the cadence built to manage it, a risk function that is passive, periodic, or cut off from the people generating the picture does not fail for lack of information. It fails because it sits at the wrong point in the process to use the information in time. The reforms after 2008 did more than add rules. They moved the CRO from an end-of-process reviewer to a continuous participant in capital allocation, portfolio decisions, and planning. The banks that made that move handled what came next, including the 2023 regional crisis, visibly better than those that did not.

Cyber risk is at that inflection point now. The conversations that close the gap with the CISO, the insurer, and the CFO are the right conversations. But conversations are episodic, and frontier AI has broken the cadence assumption underneath episodic risk management. The threat environment moves faster than the renewal cycle. The space between those two clocks is where coverage disputes are born. More conversations will not close it. Structure will.

None of what follows is a reorganization. No budget line, no headcount, no reporting change. Each one names a specific place the function breaks, and each one can be started by a single risk manager who decides to. Taken together, they move the function from procurement to intelligence.

Failure Point One: The CISO Conversation Has No Standing Instrument

The reason it does not happen is not unwillingness. There is no standing instrument to make it productive. The risk manager does not know what to ask in language the CISO will take seriously. The CISO does not know what the risk manager needs in language an underwriter will act on. With no shared artifact, the conversation either never happens, or it happens but leaves no trace anywhere in the insurance program.

The fix: a standing quarterly brief. One page, written by the CISO, built around five questions that map onto the dimensions an underwriter actually prices. It should exist ahead of renewal rather than during it, and its job is to give both sides the same picture of where exposure sits: what the CISO is investing in, what residual risk the organization is carrying in systems it cannot patch quickly, whether coalition participation belongs in the operating model, and what has changed since the previous brief.

Conversations are episodic, and frontier AI has broken the cadence assumption underneath episodic risk management.

Buying Resilience

How Risk Managers Can Close the Gap Between the CISO and the Boardroom

Question to the CISO	Why it matters	What a good answer looks like
Has the organization conducted AI-scale vulnerability discovery against its own code repositories in the past 12 months, or identified if their critical software providers have?	The Glasswing cohort now has materially better visibility into its own exposure than the market does. This question establishes which side of that information asymmetry the organization is on.	Yes, with specific findings, severity profile, and remediation outcomes. If the answer is no, the follow-up is what the plan is and on what timeline.
What is our KEV remediation track record over the past 24 months?	Demonstrated ability to patch under actual operational constraint is the most observable indicator of how the organization performs under pressure. The policy terms that depend on KEV remediation should reflect what the CISO knows about this track record.	A specific percentage of KEV-listed vulnerabilities remediated within CISA timelines. If that data does not exist, it should. The absence is itself a signal.
Do we have continuous visibility into our software inventory, or are we running on periodic snapshots?	AI-scale discovery, whether by defenders or adversaries, operates on continuous time. Periodic visibility means the organization is always working from a picture that is already out of date.	Continuous, automated inventory with known gaps documented. Point-in-time visibility is a starting point, not an answer.
Can we assess whether our known vulnerabilities can be chained into a credible exploit pathway, or do we just know they exist?	This is the gap between the first two layers of the underwriting stack and the third layer that currently does not exist in most programs. ExploitGym documents AI agents chaining low- and medium- findings into full compromise routinely. The organization that cannot reason about chain viability cannot accurately represent its exposure.	An honest "no" is valuable. It identifies the specific investment the CISO needs to make and gives the risk manager a specific question to ask the insurer about how that gap affects pricing.
What third-party components do we run that we did not build and cannot directly patch?	Organizations with limited visibility into their software supply chain may be accepting the risk without full knowledge of their exposure. .	Software Bill of Materials data where available. A component inventory at minimum. If the organization does not know what software it runs, that is the first problem to solve.

Buying Resilience

How Risk Managers Can Close the Gap Between the CISO and the Boardroom

Once the brief exists, the risk manager knows where exposure is concentrated, where the distance between what the insurer sees and what the CISO knows is widest, and how the CISO's spending will change the picture before the next renewal.

None of this asks the risk manager to become a security expert. Ask for the brief, read it, and carry what it says into the insurer conversation. The CISO answers five questions each quarter. The broker carries the output into underwriting. Without it, every renewal starts from scratch, and the same ground gets covered again.

Failure Point Two: Legal Does Not Know It Has a New Exposure

Legal has picked up a new exposure, and nobody has told the legal department. Disclosure obligations written for vulnerabilities that somebody else found and reported on a human timeline now attach to findings arriving continuously and in volume. That lands on the General Counsel (GC), who has no framework for it.

Not every organization has this problem in the same shape, and the shape matters. A pure consumer of software, running curl and commercial operating systems and third-party software as a service without maintaining or distributing code, has the narrow version. A participant in an AI-discovery program, whether running discovery against its own environment or receiving findings through a coalition arrangement, has the wide one. Know which you are before starting this conversation.

For the consumer, the GC's question is narrow and specific. When a critical vulnerability is disclosed in software we run and it lands on the CISA KEV list, does that trigger notification under our current policy terms? Most policies

carry material change language. Almost nobody has tested whether an unpatched KEV-listed vulnerability in a core system clears that bar.

For the AI-discovery participant, the exposure is wider. When a program the organization runs or contributes to surfaces something material, three obligations can be triggered at once: Securities and Exchange Commission (SEC) material cybersecurity incident disclosure under the 2023 rules, vendor notification under whatever contract requires disclosure of known vulnerabilities in supplied software, and policy notification for a material change in risk posture. These involve different clocks, different audiences, and different legal tests for what counts as a trigger. None of them map to an existing workflow, because every existing workflow assumes the vulnerability came from outside.

The fix: the same either way. Get in front of the GC before a discovery event forces the conversation. For the consumer, ask one question: What does our policy language actually require when a KEV-listed vulnerability sits unpatched in a critical system, and who decides whether that triggers notification? For the AI-discovery participant, ask four: Is this finding in production software we run? Is it exploitable in our environment? Does it meet the SEC materiality threshold? Does it trigger policy notification? Walking through either version in the abstract, before anything has been found, takes ninety minutes and leaves the organization with a framework it can use when something is found.

The GC will want this conversation. Most GCs have not built the framework because nobody has walked in with the agenda. The risk manager is the only person touching all three pieces: the policy terms, the SEC obligations, and the CISO's discovery program. That makes it the risk manager's responsibility to start the conversation.

Buying Resilience

How Risk Managers Can Close the Gap Between the CISO and the Boardroom

Failure Point Three: The CFO Argument Gets Made Too Late and in the Wrong Frame

The CFO conversation means turning a technical and actuarial argument into a business case, usually inside a budget meeting where the only live question is whether the cyber line item is the right number. The argument is fine. The timing and the frame are wrong. It gets made at budget time, against last year's premium and whatever the broker is recommending this year. By then, the CFO is optimizing cost on a number they already approved, while somebody asks them to raise limits based on a threat environment they have no way to evaluate independently. That argument does not get won in a budget meeting.

The fix: make the argument 90 days before budget season instead, in market timing language rather than security language, anchored to an external event the CFO can verify without taking anyone's word for it. The market history set out at the start of this paper is the whole argument: three repricings in a decade, with preparation beating reaction every time. That is a market history conversation. CFOs have those all day.

It also changes the ask. Not more money, but an earlier decision on money already committed. Lay out the market history, name the frontier model reporting as the event that will put every buyer in the market into the same conversation at once, and ask for one decision: Do we want to be having this before the market adjusts, or after? Framed that way, the CFO is not being asked to evaluate a security claim. They are making a timing call in a domain they know cold.

The CFO who buys cheaper insurance because the premium difference is not explained in business terms is not making a bad decision. They are making the only decision available to them given the information they have. The risk manager's job is to give them better information, not to win an argument about security spending"

Failure Point Four: The Insurer Relationship Resets to Zero Every Renewal

The annual questionnaire is a snapshot. It says what the posture looked like on the day somebody filled out the form. It says nothing about what the CISO has built since the last renewal, what the KEV remediation record looks like, whether anyone has run AI-scale exposure analysis against the known vulnerability inventory, or whether the profile is improving or sliding. So the insurer prices every renewal from a standing start, and an organization that has genuinely improved has no way to show it until the next form goes out, by which point the improvement has already been averaged into an undifferentiated market rate.

The fix: a standing posture update provision in the broker agreement. A structured one-page annual summary, pulled straight out of the CISO brief from Failure Point One, sent by the broker to the underwriter and put in the file before renewal. Nobody is asking the insurer to change how it underwrites. It just gets a posture history to price against instead of a questionnaire to price from. Do that for three years, and the record itself becomes leverage. An underwriter who can see improving remediation velocity, real investment in resilience maturity, and an operating model that has moved is pricing a different risk than one working from a blank form every twelve months.

Buying Resilience
How Risk Managers Can Close the Gap Between the CISO and the Boardroom

This is the dental model. The patient who turns up every year with x-rays, hygiene visits, and a documented history of flossing gets priced differently from the one who arrives with nothing. Organizations that start building posture history now, while the market is still hardening, will negotiate on evidence rather than attestation. The document already exists. It is the output of the quarterly brief.

The renewal conversation itself changes once the risk manager walks in having had the CISO conversation. The point is not to talk the premium down by claiming good posture. It is to make sure the coverage matches the actual risk profile and to find out how the insurer is thinking about the parts of this environment that matter most to the organization.

Three questions are worth asking directly.

Question to the insurer	What the answer reveals
How does your underwriting methodology account for the difference between an organization that has conducted AI-scale vulnerability discovery against its own environment and one that has not?	Whether the insurer has engaged with the post-Glasswing risk environment at all. An insurer that cannot answer this question is pricing on assumptions that no longer hold.
What is your approach to KEV-listed vulnerabilities in policy terms, and specifically, do you treat KEV presence as a prioritization signal or as a trigger to limit policy coverage?	Whether the insurer has absorbed the distinction set out at the front of this paper. An insurer that changes coverage when a KEV is present at time of breach is conflating awareness of risk with cause of loss. That approach will become more costly as AI-scale discovery populates the KEV catalog rapidly. The time to understand the insurer’s framework is before a breach, not during a claim.
Does your services panel include access to cyber resilience services aligned to our priorities?	Whether the insurer’s resilience partnership offer is substantive or marketing.

Buying Resilience

How Risk Managers Can Close the Gap Between the CISO and the Boardroom

What It Looks Like When It Works

The risk manager who has done all four has reorganized nothing. No budget, no headcount, no reporting change. What moved is where the function sits in the information chain: from the end of it to the middle. The CISO's real risk picture now reaches the insurance program. The insurer prices a track record instead of a form. The GC has a decision framework in hand when a finding arrives. The CFO gets a timing argument before budget season instead of a security briefing during it. The conversation that used to depend on somebody remembering to have it now happens by default.

That is what the post-2008 reforms did to the CRO role, but not by mandate alone. The continuous version of the job turned out to work better in fast-moving risk environments than the periodic one. The risk managers making these changes now are not doing anything extraordinary. They are doing early what the market will eventually require of everybody. The ones who wait will do the same work under worse conditions, after the loss that takes the choice away.

The Larger Picture

So far, this has been an organizational argument: close the gap between the CISO's world and the boardroom, use the renewal as the forcing function, and buy resilience instead of transfer. It sits inside something larger.

Frontier AI lowers the cost of sophisticated attack and speeds up discovery, exploit development, and execution all at once. The practical consequence is that controls that used to hold residual risk inside appetite may no longer do so.

That is not a verdict on the controls. The environment moved around them. Risk velocity changed, and residual risk can drift away from appetite while spending holds steady and the security program

does nothing wrong. For a risk manager, the useful question is not whether cyber risk went up, but whether the appetite, the investment priorities, and the treatment strategies still make sense in an environment that has materially changed.

None of this is confined to individual organizations. Every node in the defense ecosystem is adapting at once. CISOs are rebuilding operating models. The disclosure infrastructure is carrying a load it was never designed for: the Common Vulnerabilities and Exposures (CVE) Program, the CVE Numbering Authority (CNA) framework, and the National Vulnerability Database (NVD) enrichment pipeline are all absorbing volume at a rate nobody built for, and how fast they adapt determines how much of the window between discovery and exploitation defenders actually get. The Glasswing coalition is showing that collective discovery and coordinated disclosure can work at AI scale. The insurance market is starting to build the signals it needs to tell those two populations apart.

There is a collective dimension that no individual program reaches. The organizations that gain most directly from coalition discovery are the ones that build software. They can point frontier AI at their own codebases, find what is sitting there before it becomes public, and coordinate remediation before anyone gets exploited.

The Glasswing coalition is showing that collective discovery and coordinated disclosure can work at AI scale.

Buying Resilience

How Risk Managers Can Close the Gap Between the CISO and the Boardroom

An enterprise that buys commercial software rather than writing it cannot do any of that, because it does not have the source code. What it gets instead depends on whether its critical suppliers are taking part in arrangements like Glasswing. Most of the resilience an enterprise ends up with will be inherited rather than built.

That makes supplier participation an upstream indicator, visible from outside the organization, and it belongs in the underwriting conversation for the same reason remediation velocity does. “Beyond Vulnerability Counts” sets out the signals the market is building to price that difference, and coalition participation sits underneath several of them.

Underneath all of this is a governance story rather than a technology story, continuous with what this community has been building for twenty-five years through the CVE Program, the CNA framework, and every sector Information Sharing and Analysis Center that ever argued for sharing threat intelligence across organizational lines. Coordinated defense at ecosystem scale produces outcomes no single organization reaches alone. The principle has not changed. What has changed is the scale at which it now has to work.

The governance question stopped being academic on June 2, 2026, when Executive Order 14409, “Promoting Advanced Artificial Intelligence Innovation and Security,” directed the Secretary of the Treasury, in consultation with the National Cyber Director, NSA, and CISA, to form an AI cybersecurity clearinghouse in voluntary collaboration with the AI industry and critical infrastructure operators, coordinating and deconflicting vulnerability scanning, validating findings, and prioritizing remediation and patch distribution. Federal policy now has a governance response to AI-scale discovery built into it. The argument here is the same one at organizational scale. They are the same problem at different altitudes. The risk manager sits where they meet, and the renewal is not a procurement exercise. It is the moment the organization’s position in that ecosystem becomes legible to the people setting its price.

That conversation is worth having. The window for having it on your own terms is shorter than it looks

Buying Resilience

How Risk Managers Can Close the Gap Between the CISO and the Boardroom

About the Authors

James Tuplin, is the Head of International Cyber underwriting at Mosaic Syndicate where he is responsible for the cyber underwriting teams based in London, Cologne, and Dubai. James has 23 years of underwriting experience specializing in cyber and technology E&O insurance for the last 20 years.

Jay Vinda, is the CISO and Cyber Risk Engineering Lead at Mosaic Insurance. He bridges the gap between technical security concepts, business objectives and risk management strategies. Through his work, he aims to contribute towards establishing cyber insurance as a powerful risk management partner, proactively helping organizations to protect systems and become resilient to emerging threats.

Sanith Wijesinghe, Dr. Sanith Wijesinghe is the Chief Engineer for MITRE's public sector. Over the past 10 years he has overseen MITRE's R&D initiatives in blockchain, including development of the AADAPT cyber threat framework for digital assets. Most recently he led a NIST cybersecurity workshop panel discussion on "Systemic Threats, Vulnerabilities, Interoperability and Other Shared Challenges" for digital ledger based technologies.

Chris Folk, is a Senior Cybersecurity Leader at MITRE, with 25 years in cybersecurity, national critical infrastructure protection, including eight years leading the CVE Program through its 2016 federation redesign. His work focuses on the governance layer of cyber risk: how disclosure, coordination, and interpretation reduce exposure risks.

Buying Resilience

How Risk Managers Can Close the Gap Between the CISO and the Boardroom

Key Sources

Source	Finding Cited
ExploitGym, Berkeley RDI / Anthropic / OpenAI / Google (May 2026)	157 of 898 exploited by Claude Mythos Preview and 120 by GPT-5.5, mitigations disabled; separate GPT-5.4 case study produced a 229-line V8 exploit chain in 71 minutes; frontier/non-frontier performance cliff
Microsoft MDASH, Security Blog, Taesoo Kim (May 12, 2026)	16 CVEs on single Patch Tuesday; 4 Critical RCEs; retrospective recall of 96% on 28 MSRC cases in clfs.sys and 100% on 7 in tcpip.sys; 2 UDP packets to LocalSystem
Project Glasswing, Anthropic (April 7, 2026)	Thousands of zero-days; fewer than 1% patched at announcement; 135-day coordinated disclosure window
Verizon DBIR 2026 (2025 edition where noted)	43-day median to fully remediate a KEV, up from 32; only 26% of KEV vulnerabilities fully remediated in 2025, down from 38%; edge devices and VPNs 22% of exploitation breaches, up from 3%. 2025 edition: 0-day median time to mass exploitation of edge devices
IBM X-Force Threat Intelligence Index 2026 (February 2026)	Vulnerability exploitation the leading cause of attacks at 40% of observed incidents
Beyond Patch Velocity, Paper 1 in this series (2026)	CISO operating model; cost imposition; AI-augmented defense; resilience over patch velocity
Mandiant M-Trends 2026 (March 2026)	Exploitation leading initial infection vector for sixth consecutive year at 32%; mean time to exploit negative seven days; access-broker handoff to ransomware affiliate now 22 seconds
curl 8.21.0 security release (June 24, 2026)	18 CVEs in single release (record); 25-year-old authentication bypass present since 2001; triggered by Mythos finding one CVE in May 2026; AI cascade demonstrating embedded component exposure at 30B device scale

Buying Resilience

How Risk Managers Can Close the Gap Between the CISO and the Boardroom

References

- 1 Anthropic, Project Glasswing announcement, April 7, 2026. <https://www.anthropic.com/glasswing>. Fewer than 1% of identified vulnerabilities were patched at the time of the announcement.
- 2 ExploitGym: Benchmarking AI Systems on Offensive Security Tasks, Berkeley RDI / Anthropic / OpenAI / Google, May 2026. <https://rdi.berkeley.edu/blog/exploitegym/>. 157 of 898 real-world vulnerabilities exploited by Claude Mythos Preview; and 120 by GPT-5.5, with mitigations disabled. The 71-minute, 229-line V8 exploit chain was a separate GPT-5.4 case study, produced with ASLR and the V8 heap sandbox disabled.
- 3 Microsoft Security Blog, Taesoo Kim, "Defense at AI Speed," May 12, 2026. <https://www.microsoft.com/en-us/security/blog/2026/05/12/defense-at-ai-speed/>
- 4 curl 8.21.0, released June 24, 2026. CVE-2026-8932 introduced in curl version 7.7, March 22, 2001.
- 5 IBM X-Force Threat Intelligence Index 2026, published February 2026: vulnerability exploitation was the leading cause of attacks, at 40% of observed incidents.
- 6 Mandiant M-Trends 2026: exploitation of vulnerabilities was the leading initial infection vector for the sixth consecutive year, accounting for 32% of intrusions. Mean time to exploit is now negative seven days.

