DRZEROTRUST RESEARCH DIVISION

An Analysis of the Cyber Insurance Market for 2026

A final analysis of market structure, claim uncertainty, governance risk, and Zero Trust implications

By Dr. Chase Cunningham

FINAL ANALYSIS

Executive summary

Cyber insurance has become a common executive response to cyber risk because it offers a clean, familiar narrative: acknowledge the risk, buy a policy, transfer some exposure, and gain access to incident-response support.[1] That comfort is not imaginary, but it is frequently overstated because a policy does not change the security condition of the environment, the trust relationships inside the architecture, or the organization’s ability to prevent, detect, contain, and recover from a cyber event.

That distinction is the central thesis of this paper. Cyber insurance is a risk-financing mechanism, not a risk-control mechanism. It may help fund selected losses after a covered event, but it cannot fix the governance failures, identity weaknesses, architectural flaws, visibility gaps, and resilience shortfalls that make those losses possible in the first place.

This is not an anti-insurance argument. It is a category correction. Insurance and security belong in the same enterprise-risk conversation, but they perform different functions and should never be treated as substitutes for one another. A mature organization may use security and resilience to reduce risk, then use insurance to transfer part of the financial uncertainty that remains.

The governance failure begins when leaders mistake policy ownership for security maturity. A board asks whether the organization is covered. Management answers yes. The premium has been paid, the broker has been engaged, and the policy binder exists. The conversation ends, even though the harder questions remain unanswered: why compromise would occur, how far an attacker could move, whether privileged access is constrained, whether telemetry is sufficient, whether recovery has been tested, and whether the organization can survive if payment is delayed, partial, disputed, or unavailable.

DrZeroTrust takes a clear position on that failure. Cyber insurance is not the enemy. False reliance on cyber insurance is the enemy. The issue is not whether organizations should buy cyber insurance. The issue is whether they are using it as a residual-risk tool inside a disciplined security program or as a symbolic substitute for architecture, operational accountability, and Zero Trust maturity.

The false comfort of the policy binder

Cyber insurance appeals to leadership because it fits the language of enterprise finance and governance. It offers a defined cost, the promise of assistance after an event, and the appearance of prudent risk management. In boardrooms already saturated with competing priorities, that can feel like progress.

But a policy is not a secure architecture. It is not identity governance, least privilege, segmentation, endpoint detection, immutable recovery, a tested incident-response plan, or a threat-informed defense capability. It is a contract, and contracts do not reduce attack surface or prevent lateral movement by themselves.

This is where the illusion begins. Organizations often confuse financial risk transfer with operational risk reduction. Insurance may help answer who pays for some covered losses after the incident, but it does not answer the more important questions about why the compromise happened, how quickly it was detected, whether it could be contained, and whether mission operations could be restored with confidence.

The danger is not the existence of cyber insurance. The danger is the substitution of insurance for security accountability. That substitution becomes especially dangerous in environments that already know their fundamentals are weak, including unmanaged privileged access, inconsistent multifactor authentication, poor asset inventory, flat networks, unpatched critical systems, limited logging, untested backups, weak incident response, and undefined cyber-risk ownership.

In those environments, insurance becomes a boardroom sedative. It makes risk feel handled while the conditions of compromise remain largely unchanged. A policy may change the financial conversation after failure, but it does not make an insecure system trustworthy before failure.

Scope and definitions

This paper focuses on cyber insurance as a strategic risk-management tool, particularly in organizations that use it to justify, delay, or compensate for weak security practice. The analysis depends on keeping several concepts separate because the market often blurs them together.

Cyber insurance is a contractual risk-financing mechanism. Depending on policy wording, it may cover selected first-party losses, third-party liabilities, legal expenses, forensic costs, notification expenses, business interruption, extortion-related costs, restoration expenses, privacy liability, or regulatory defense. The policy does not operate automatically. Payment depends on policy language,

exclusions, sublimits, retentions, representations, notice requirements, cooperation duties, causation analysis, documentation, and claim handling.

A cybersecurity control is different. A control changes the organization's security behavior or risk profile by reducing likelihood, reducing blast radius, increasing detection, preserving evidence, improving response speed, accelerating restoration, or strengthening governance. Cyber insurance does not directly change attacker access, trust relationships, privilege, segmentation, telemetry, or resilience.

Residual risk is the risk that remains after prevention, detection, response, and recovery measures are in place. Insurance belongs in that layer. Treating insurance as residual-risk financing is rational. Treating it as an alternative to reasonable controls is not.

One more distinction matters for executive discussions: claim denial is not the same as a claim closed without payment. A denial usually means the insurer refused payment based on coverage language, exclusions, or conditions. A claim closed without payment is broader and may include denials, claims below the retention, withdrawn claims, duplicate claims, insufficiently documented claims, or claims with no covered loss. That distinction matters because the strongest public regulator-backed signal currently available is the NAIC category "claims closed without payment," not a pure market-wide denial-rate series.

Risk transfer is not risk reduction

The central thesis of this paper is simple: cyber insurance is a risk-financing mechanism, not a risk-control mechanism. It may help fund recovery from selected covered losses, but it cannot correct the security, governance, identity, visibility, and resilience failures that create those losses.

This framing fails when insurance is sold or consumed as a solution to weak security. That mistake is temporally wrong because security matters before, during, and after compromise, while insurance primarily becomes relevant after loss. It is functionally wrong because security changes the technical and operational environment, while insurance changes the financial allocation of some consequences. It is institutionally wrong because security requires architecture, accountability, operational discipline, and continuous validation, while insurance requires underwriting, policy terms, premiums, and legal interpretation.

A policy can help finance a loss. It cannot make an insecure environment trustworthy. That single sentence captures the category error at the center of most boardroom misunderstandings about cyber insurance.

This is the core market position of the paper. Zero Trust is about reducing implicit trust, constraining access, continuously validating identity and device posture, enforcing least privilege, and limiting blast radius through architecture and policy. Cyber insurance does none of those things directly.

The evidence problem

Organizations looking for clean cyber-insurance denial statistics quickly discover that public evidence is scarce. There is no simple, market-wide, peer-reviewed public denial-rate series that reliably tells buyers what proportion of cyber claims are denied.

That scarcity is not accidental. Claims data can expose policyholder identities, incident details, forensic findings, legal strategies, insurer loss experience, and control failures, all of which are commercially sensitive. The market has strong incentives to keep those details private.

The user’s draft notes peer-reviewed support for that point through work by Nurse and others, describing the challenge of building shared cyber-insurance datasets because claims and underwriting data are proprietary and competitively sensitive. That creates a strategic problem for buyers because the market asks them to treat cyber insurance as an important risk-management instrument while giving them limited public evidence to evaluate how reliable payment may be in practice.

That does not mean insurers never pay. It means policyholders should be cautious about any executive narrative that treats payment as predictable without close attention to policy wording, exclusions, retentions, representations, documentation, and claim conditions. The best public proxy is therefore not a denial rate. It is the NAIC measure of claims closed without payment.

The NAIC nonpayment signal

The NAIC Cybersecurity Insurance Market Report is the strongest public U.S. source for cyber-insurance market data in the draft. According to the paper, the 2024 report showed approximately $9.14 billion in U.S. cyber direct written premium, down from $9.84 billion in 2023, while cyber claims rose almost 40 percent to nearly 50,000 claims reported.[2]

The more important signal for this paper is claim closure. NAIC reported an overall 2024 claim closure rate of 82.05 percent and, among U.S.-domiciled insurers excluding alien surplus lines, listed 9,941 claims closed with payment

and 28,555 claims closed without payment. That means claims closed without payment outnumbered claims closed with payment by nearly three to one.

NAIC 2024 category	Claims closed with payment	Claims closed without payment	Closed without payment share
Primary policies	8,838	16,422	65.0%
Excess policies	514	10,802	95.5%
Endorsements	589	1,331	69.3%
Total	9,941	28,555	74.2%

This data must be used carefully. It should not be paraphrased as “74.2 percent of cyber claims were denied,” because that would overstate what the NAIC data proves. Closed without payment is broader than denial and includes multiple claim outcomes.

But the governance significance remains substantial. A cyber policy is not a payout. It is a conditional instrument, and the existence of a policy does not guarantee payment after an incident. Boards and executive teams should therefore treat insurance as an uncertain recovery path governed by policy language, facts, documentation, exclusions, retentions, and claim handling rather than as guaranteed recovery capital.

Why nonpayment matters

The distinction between denial and claims closed without payment should not become an excuse to dismiss the strategic message in the data. The operational consequence for the policyholder may still be the same: the organization suffered a cyber event and did not receive payment for that claim closure.

There are several reasons this can happen, including loss below the retention, exclusions, claims outside the insuring agreement, insufficient documentation, unmet policy conditions, disputed causation, withdrawn claims, and losses that are real but not quantifiable as covered loss. Each of those reasons reinforces the same governance lesson: policy ownership does not equal cash recovery.

This matters because many organizations model insurance as if the sequence is simple: event, claim, payment, restoration. In practice, the sequence may be event, investigation, legal review, documentation problems, coverage interpretation, disagreement over cause, sublimit analysis, retention analysis, and delayed or absent payment. The farther an organization is from evidence discipline and scenario planning, the more exposed it becomes.

The defensible position is not that insurers never pay. The defensible position is that public evidence does not support treating cyber-insurance recovery as certain, and that uncertainty is enough to make insurance an inappropriate substitute for actual risk reduction.

A policy is not a security control

The simplest test is to ask what cyber insurance changes in the environment. Does it reduce exposed attack surface, eliminate standing administrative access, prevent credential theft, validate device posture, stop lateral movement, segment critical systems, detect malicious behavior, preserve logs, or restore systems? The draft's answer is no across the board.

Insurance may create incentives for some controls during underwriting, and it may connect policyholders to response vendors after an event. But it does not function as a technical, administrative, or operational control by itself.

The paper uses the NIST Cybersecurity Framework 2.0 functions as a useful contrast, noting that mature cybersecurity requires Govern, Identify, Protect, Detect, Respond, and Recover outcomes. It also points to Zero Trust as a model focused on removing implicit trust and continuously validating access before sessions are established. Cyber insurance performs none of those functions directly.

Security function	What mature security does	What cyber insurance does
Govern	Assigns accountability, risk appetite, and oversight	May influence governance discussions, but does not govern
Identify	Inventories assets, dependencies, identities, data, and risk	May ask about assets during underwriting
Protect	Enforces safeguards and access controls	May require safeguards as underwriting conditions
Detect	Finds malicious behavior and anomalies	Does not detect compromise
Respond	Contains and manages incidents	May fund or coordinate some response vendors
Recover	Restores operations and validates integrity	May reimburse selected recovery costs
Zero Trust	Removes implicit trust and enforces continuous verification	Does not alter trust relationships

This is the category correction that many executive conversations need. Insurance can support some post-incident activity, but it cannot replace the security operating model itself.

The residual-risk rule

Cyber insurance is most defensible when used for residual risk. Residual risk is what remains after an organization has made reasonable decisions about prevention, detection, response, and recovery. No organization can reduce cyber risk to zero, and even mature environments remain exposed to new vulnerabilities, third-party failures, cloud misconfigurations, insider behavior, supply-chain compromise, and human error.

In that context, insurance can make sense. It can finance selected losses that remain after risk has been reduced to a level consistent with the organization's risk appetite. That is what mature cyber insurance should be used for.

But residual risk is not the same as neglected risk. Insurance should not be used to rationalize known failures such as no enterprise MFA, flat network architecture, poor privileged-access management, unpatched critical systems, untested backups, weak logging, missing incident-response exercises, weak asset inventory, poor vendor-risk discipline, or undefined cyber-risk ownership.

The board-level rule from the draft is worth preserving exactly because it is memorable and operationally clear: insure residual uncertainty, not preventable negligence. That does not mean every security failure is negligence. It means organizations should not use insurance as a financial excuse for refusing to correct known, material, and controllable weaknesses.

The underwriting illusion

One of the most dangerous assumptions in executive conversations is that obtaining cyber insurance proves the organization is secure. It does not.

Underwriting approval means an insurer accepted the risk at a given premium, subject to policy terms, exclusions, sublimits, retentions, representations, and conditions. It does not mean the organization has achieved security maturity, that controls are effective, or that the environment was deeply validated.

Research shows that underwriting depends on complex and imperfect information and that claims decisions require substantial post-incident fact gathering. It also cites analysis of application forms showing that they may encourage some good practices but underrepresent governance, incident response, and recovery topics

while often asking only whether a control exists rather than whether it is effective.

That distinction matters because organizations can answer “yes” to having backups without testing restore, “yes” to MFA while leaving privileged or third-party pathways exposed, and “yes” to endpoint protection while lacking coverage on critical systems. The existence of a control is not the same as the effectiveness of a control.

An insurance application is not a penetration test, not a control assessment, not a resilience exercise, not a Zero Trust maturity review, and not proof of operational security. The sooner boards and buyers internalize that fact, the better their decisions become.

The questionnaire trap

Cyber-insurance questionnaires can be useful when they push organizations toward basic controls. That is the charitable interpretation, and sometimes it is true.

But questionnaires can also distort security behavior. When the goal becomes answering the insurance form correctly, organizations may optimize for insurability rather than resilience. They may implement the minimum visible control required for underwriting rather than designing around adversary behavior, mission dependency, blast-radius reduction, recovery integrity, and identity risk.

That is a compliance trap. It occurs when an organization confuses evidence of having answered a question with evidence of having reduced risk. This is strategically weak because insurer forms are designed for risk selection and pricing, not for complete security architecture. They may also lag threat behavior and often lack depth in the areas that determine whether an organization can survive an actual incident.

A questionnaire may be useful input. It should never be treated as the security program itself. Insurer requirements are commercial hygiene, not maturity.

The payment illusion

The next false comfort is the belief that the policy will pay when the organization needs it most. A cyber-insurance policy is not a reserve account, not an emergency fund, and not a guaranteed reimbursement mechanism. It is a contract that may respond if the facts of the event align with the coverage grant and avoid exclusions, limitations, and condition failures.

There are multiple layers of uncertainty, including whether the event triggers coverage, whether the loss category is covered, whether the retention is exceeded, whether sublimits apply, whether exclusions apply, whether underwriting representations were accurate, whether policy conditions were satisfied, whether the loss can be measured, and whether payment will arrive when the business actually needs operational support.

This is not theoretical uncertainty. The NAIC closure data shows a large number of claims closed without payment relative to claims closed with payment in 2024. The right conclusion is not that insurers never pay. The right conclusion is that payment is contingent.

Boards should therefore separate five categories in their risk register: gross cyber risk before controls, control-reduced risk after security and resilience measures, retained risk after acceptance, insured risk subject to policy terms, and uninsured or disputed risk that may remain after claim adjudication. Many organizations collapse the last two into a single assumption and treat insured as though it means paid. That is a governance failure.

Systemic and correlated risk

Traditional insurance works best when losses are sufficiently independent or diversifiable. Cyber risk often violates that assumption because many organizations depend on the same cloud providers, identity systems, operating systems, endpoint platforms, SaaS applications, and software components.

That creates aggregation risk for insurers and shared-failure risk for policyholders. A single vulnerability, vendor compromise, or authentication failure can affect many insureds at once. In those scenarios, the most severe events may also be the ones most likely to trigger exclusions, sublimits, litigation, or market hardening.

The paper notes that GAO has warned catastrophic cyber incidents, particularly those affecting critical infrastructure, could produce losses that private insurers and existing federal mechanisms may not cover.[4] The Lloyd’s market has also required attention to exclusionary wording for state-backed cyber exposure because of systemic-risk concerns.

From a Zero Trust standpoint, the lesson is straightforward. Dependence is not trustworthiness. Shared infrastructure and shared identity systems create shared failure modes, and insurance cannot remove that dependency, reduce the blast radius, or validate trust relationships inside the environment.

War, attribution, and exclusion uncertainty

Cyber insurance becomes especially uncertain when incidents involve destructive malware, state-linked actors, widespread disruption, or ambiguous attribution. Those are not fringe scenarios. They are among the exact events executive teams care about most because they carry the highest disruption and reputational stakes.

The paper uses the Merck and NotPetya litigation as an illustration of how a major cyber event can become years of legal dispute over whether a traditional warlike-action exclusion applies to a cyberattack.[3] The key lesson is not simply who prevailed in that case. The key lesson is that catastrophic cyber events can convert technical and geopolitical uncertainty into contractual and legal uncertainty.

For boards, that means the right question is not whether cyber insurance exists. The right question is which scenarios leadership believes are covered, which are excluded or sublimited, and what assumptions would need to be true for payment to occur. Without that mapping, policy ownership offers emotional comfort without operational clarity.

Market hardening and the access problem

Cyber insurance is often defended as a governance lever because insurers may require better controls before issuing or renewing coverage. That has some truth in limited circumstances, and many baseline controls pushed by insurers are reasonable.

But the evidence is more complicated. Research on the ransomware era argues that the cyber-insurance market hardened, raised the security standards required for purchase, and made coverage less accessible or less meaningful for some buyers.

This creates a contradiction. If insurance is supposed to improve security by forcing weak organizations to meet underwriting standards, then the very organizations that most want insurance to compensate for weak security may be least able to obtain affordable or meaningful coverage. Cyber insurance can reward maturity. It cannot manufacture maturity.

That is an important market point. The brand should speak directly to organizations that need actual maturity work rather than reassurance theater.

The moral-hazard issue

Moral hazard appears often in cyber-insurance debates because critics worry that insurance weakens the incentive to invest in security. The analysis takes a more disciplined position, noting that the empirical evidence does not support a simplistic claim that every insured organization becomes careless. Some organizations are highly mature, and some insurers require meaningful controls.

The stronger argument is governance-based. Cyber insurance creates moral-hazard risk when leaders treat the existence of coverage as evidence that cyber risk has already been managed. That misunderstanding can delay architecture modernization, identity risk reduction, recovery testing, staffing, and third-party risk work even when nobody is acting in bad faith.

This is why the paper calls the phenomenon governance theater. The policy binder becomes a symbol of action while the underlying system remains vulnerable. That is a powerful phrase because it captures the boardroom failure without requiring any accusation of dishonesty.

The value-add argument

The cyber-insurance market often presents policies as more than reimbursement. Policies may provide access to breach counsel, forensic vendors, ransomware negotiators, notification support, crisis communications, tabletop resources, risk scans, and underwriting feedback. Some of that help is real and useful.

But usefulness is not the same as security capability. Supporting services operate around the incident; they do not solve the underlying architectural, identity, telemetry, or resilience problems that made the incident possible.

Claimed value-add	Real value	Limitation
Breach counsel	Helps manage legal response	Does not prevent compromise
Forensic vendor panel	Provides post-incident investigation	Does not create prior visibility
Ransomware negotiator	May support extortion response	Does not reduce initial blast radius
Notification support	Helps meet breach obligations	Does not protect data before exposure
Underwriting questionnaire	May push baseline controls	Does not validate control effectiveness

Claimed value-add	Real value	Limitation
Security scan	May identify external weaknesses	Does not equal architecture review
Crisis communications	Helps manage messaging	Does not restore lost trust
Claim reimbursement	May finance selected losses	Conditional, delayed, and limited

The problem is not that these services are worthless. The problem is that they are sometimes marketed or mentally consumed as a substitute for strategy. Cyber insurance may add value to a mature program. It does not add enough value to compensate for a failed one.

Non-indemnifiable harm

Cyber insurance is financial by design. It is built to reimburse or indemnify defined categories of loss. But cyber harm is broader than covered cost.

Cyber incidents can damage customer trust, employee morale, operational legitimacy, regulatory credibility, market confidence, patient safety, partner relationships, brand value, and strategic momentum. Some of those harms have financial consequences, but they are not fully reducible to claim payment.

A useful taxonomy of cyber harms includes physical, digital, economic, psychological, reputational, and social harms. That framing matters because it shows why reimbursement alone cannot fully restore what a serious breach takes away.

This issue is especially important for Zero Trust. Zero Trust assumes that trust should never be implicit and that access should be continuously earned through identity, posture, context, least privilege, segmentation, telemetry, and enforcement. When a major breach occurs, the organization has not merely lost money. It has exposed a failure in its trust model. An insurance payment may help settle invoices. It does not rebuild that trust model or restore institutional confidence by itself.

The correct relationship between insurance and Zero Trust

Zero Trust and cyber insurance do not occupy the same strategic layer. Zero Trust is a security model. Cyber insurance is a financial instrument.

Zero Trust reduces implicit trust, constrains access, validates identity and device posture, enforces least privilege, segments resources, improves visibility, and limits blast radius. Cyber insurance finances some selected losses after a covered event.

The correct relationship is therefore simple: Zero Trust reduces the amount of risk that must be financed, and cyber insurance finances some of the risk that remains. That ordering is what mature organizations should preserve.

Layer	Primary question	Proper function
Governance	Who owns cyber risk and what level is acceptable?	Accountability and decision rights
Zero Trust architecture	How does the organization reduce implicit trust and constrain compromise?	Risk reduction
Security operations	How does the organization detect, investigate, contain, and respond?	Active defense
Resilience engineering	How does the organization maintain and restore mission operations?	Continuity and recovery
Cyber insurance	Which residual financial losses should be transferred?	Risk financing

In this model, insurance is not rejected. It is subordinated to security. A disciplined Zero Trust organization may still buy insurance, but it buys from a position of maturity: it understands the wording, maps coverage to scenarios, validates control representations, tests response and recovery, preserves evidence, and models uninsured loss.

The board, CFO, and CISO traps

Cyber insurance often becomes a board-level reporting shortcut. A board asks whether the organization is covered, when the more important question is whether the organization is defensible. An insured organization may still be indefensible if it cannot show credible governance, control effectiveness, monitoring, vendor oversight, incident response, and recovery testing.

Boards should therefore insist on a more rigorous cyber-insurance discussion. That discussion should include what scenarios are covered, what scenarios are excluded or sublimited, what retention applies, what representations were made, what controls are policy conditions, what evidence must be preserved, what uninsured loss remains, what recovery-time assumptions exist, and what circumstances could trigger dispute. They should also ask whether the organization can survive without payment.

Finance leaders face a different trap. A premium can create budget predictability, but it does not create operational resilience. A major cyber event can stop production, disrupt patient care, trigger regulatory scrutiny, expose data, consume executive attention, and damage customer trust even when some costs are reimbursed. The better CFO question is not only how much coverage exists, but what losses remain economically material even if the policy pays.

CISOs face their own distortion risk. When leadership pressures the security team to “do what is needed to get insured,” useful investment may occur, but the roadmap may also be redirected toward insurer-visible controls rather than identity modernization, segmentation, data-flow understanding, logging architecture, recovery isolation, and operating-model maturity. Insurer requirements should be treated as the floor of commercial hygiene, not the ceiling of Zero Trust maturity.

The market position

This paper takes an explicit position because the market is crowded with comforting half-truths. The position is not that cyber insurance is useless, illegitimate, or inherently deceptive. The position is that cyber insurance becomes dangerous when it is treated as a substitute for architecture, governance, operational discipline, and resilience.

That position can be summarized in one line: cyber insurance is not the enemy; false reliance on cyber insurance is the enemy. This creates a pragmatic posture rather than an absolutist one, which is important for credibility with executive buyers, boards, and brokers who understand the limits of coverage.

The central message of this paper rests on three pillars. First, category clarity: insurance is a financial instrument, while Zero Trust is a security model. Second, governance realism: coverage does not equal defensibility, and underwriting approval does not equal maturity. Third, operational accountability: identity, privilege, segmentation, telemetry, response, and recovery are what determine resilience, not policy ownership alone.

This framing works because it speaks plain English without sacrificing technical credibility for executive and mixed technical audiences. It also creates reusable messaging for website copy, keynote talks, briefings, podcasts, advisory offers, and LinkedIn content.

What leaders should do next

Organizations do not need to abandon cyber insurance. They need to put it in the right place. That means treating the policy as a residual-risk instrument that sits behind security architecture, security operations, and resilience engineering rather than in front of them.

Boards should require scenario-based coverage mapping tied to actual threat paths and operational dependencies. Management should validate the accuracy of underwriting representations, identify policy conditions that could complicate recovery, and build evidence-preservation practices that support both response and claims handling. Finance teams should model uninsured and disputed loss rather than assuming coverage equals cash. Security leaders should use insurer requirements as minimum hygiene while keeping the roadmap centered on Zero Trust, identity, privilege, visibility, segmentation, and clean recovery.

The organization should also test a hard question: can it survive a severe cyber event if the insurer pays late, pays partially, or does not pay at all? If the answer is no, then the organization has not transferred risk so much as created dependency on an uncertain recovery path. That is not resilience. It is wishful thinking.

Conclusion

Cyber insurance has a legitimate role in a mature cyber-risk program. It can finance some selected losses, connect organizations to useful response resources, and help manage financial volatility. But it is not a control plane, not a trust model, not a resilience architecture, and not proof that leadership has done the hard work of governing cyber risk.

The category error is what this paper rejects. Risk transfer is not risk reduction. Underwriting is not maturity. A policy is not a payout. Coverage is not defensibility.

The right role for cyber insurance is clear. Use security and Zero Trust to reduce risk. Use resilience engineering to preserve operations and recover with integrity. Use insurance to transfer some of the uncertainty that remains. Anything else is governance theater dressed up as prudence.

That is the central conclusion of this paper. You cannot insure your way into trust.

Endnotes

[1] Source draft provided by the author, section on the executive comfort narrative around cyber insurance.

[2] NAIC Cybersecurity Insurance Market Report (2024), as cited and summarized in the underlying draft.

[3] Merck & Co. insurance litigation related to the NotPetya attack, discussed in the underlying draft as a coverage-uncertainty example.

[4] U.S. Government Accountability Office reporting on catastrophic cyber risk and insurance limitations, as summarized in the underlying draft.