ZURICH

The AI resilience challenge

Keeping AI-related risks manageable



University of St.Gallen

Zurich AI Lab

Contents

Foreword	3
Why resilience matters	4
Managing uncertainty	5
Navigating AI-related risks	8
Practical steps to enhance resilience	13
Keeping AI-related risk manageable	15
Methodology, Data Collection	16
Acknowledgments	17

Making the most of AI means being prepared for the risks that come with it.

Artificial intelligence (AI) is already changing how decisions are made and how businesses operate day-to-day. As AI becomes embedded in more consequential activities, failures can reach beyond individual systems and organizations, weakening public trust and economic resilience.

In that context, managing risk is a condition for continued progress, not a brake on it. The goal is not to eliminate every failure, but to ensure that failures can be detected, contained and recovered from.

This report forms part of Zurich's collaboration with the University of St. Gallen through the Zurich AI Lab and builds on the broader conditions for safe AI adoption. We focus on how AI-related losses can arise, connect and spread, and on the conditions that make those risks more manageable for organizations and society.

Drawing on insights from 50 experts, this report examines AI resilience through the lens of insurance and insurability; how the characteristics of AI-related exposures affect their manageability and insurability, and where wider resilience measures may be needed.

The four-zone framework provides a conceptual lens for exploring these characteristics rather than a methodology for assessing AI-related risks. It examines how AI-related risk becomes loss and highlights two mechanisms that can widen its impact: concentration around common providers and infrastructure, and less visible correlations between systems that share data, assumptions or patterns of behavior.

These shifts matter to organizations managing their own exposures, to insurers looking across portfolios and to policymakers concerned with risks that cross organizational or national boundaries. The limits of insurability can help reveal where wider resilience needs to be strengthened.

Complete certainty about AI-related risk is not possible, nor is it a prerequisite for action. No institution can build AI resilience alone. It will require research, insurance expertise, business practice and public policy to work together. The goal now is to keep risk visible, responsibilities clear and disruption manageable, keeping losses within society's capacity to absorb them.



Mario Greco
Group CEO
Zurich Insurance Group



Prof. Dr. Karolin Frankenberger
Professor of Strategic Management and Innovation, University of St.Gallen (HSG)

Why resilience matters

AI is moving rapidly into the everyday operations of businesses, governments and society. Its potential to drive innovation and productivity is creating strong competitive pressure to adopt it, but as it becomes embedded in more consequential activities, the impact of failure can extend beyond an individual system or organization.

Confidence in the ability to manage AI-related risks and to limit their potential impact depends on understanding vulnerabilities, detecting problems early, containing their impact and being able to recover. Developers need to build for safety, security and reliability. Organizations need to understand the dependencies, liabilities and operational risks created as AI reshapes their operations. This starts with responsible data and AI practices that provide the governance, transparency and accountability needed to use AI safely and effectively. Governments must consider the potential consequences for economic security, essential services, labor markets and public trust.

Building on How to navigate a safe AI transformation, insurance provides a useful lens through which to examine these questions. Insurers look at how risk becomes loss: What can go wrong; how widely those consequences could spread; who bears responsibility; and what can be done to prevent or reduce the impact.

This report applies that perspective to AI resilience. Developed by Zurich in collaboration with the University of St. Gallen, it asks what determines whether AI-related risks remain contained and manageable or spread more widely across organizations and society. The answer depends on how visible and measurable the exposure is, how clearly responsibility can be established, how strongly organizations share dependencies or patterns of behavior, and whether effective controls and recovery options exist.

AI will continue to evolve, but organizations cannot wait for complete certainty. The sections that follow show how concentration and correlation can widen losses, introduce a four-zone map to guide proportionate action, and set out what organizations, technology providers, insurers and public authorities can do to strengthen resilience as adoption accelerates.

> “How can insurance act as a risk partner to mitigate either the probability or the severity of a loss if it happens?



Paul Bantick
Chief Underwriting Officer
Beazley



Managing uncertainty

AI-related risks are not entirely new. Losses are likely to emerge through familiar channels such as cyber incidents, product liability, negligence, fraud and business interruption. What is less certain is how AI may affect the frequency, severity and interconnectedness of those losses.

Insurance provides a mechanism to translate that uncertainty into risks that can be assessed, priced and pooled. Risks are easier to manage and insure when exposures can be identified and measured, potential losses are bounded, responsibility can be established and failures do not automatically affect large numbers of organizations at once.

Box 1.

Key factors shaping insurability[1]

Identifiable.
Exposure, system, use case and relevant dependencies can be described with enough precision to assess and monitor.

Measurable.
The likelihood and potential impact of failure can be estimated through data, testing, experience or informed judgment.

Diversified.
Losses are not so highly correlated or concentrated that a single cause can affect large parts of the portfolio at once.

Bounded.
The potential scale of loss can be reasonably estimated and contained.

Attributable.
The causes of loss and the responsibilities of the different parties involved can be established.

AI can place pressure on several of these conditions at the same time. An individual system may appear well understood while depending on the models, infrastructure, data or providers used by many other organizations. Separate systems may share characteristics that cause them to respond in similar ways. And while some of those connections may be visible in advance, others only emerge through testing, experience or loss.

The challenge posed by AI is not primarily that it creates entirely new categories of risk, but that it can amplify or alter existing risks in ways that make them more difficult to understand and manage. The technology is evolving so fast that the real-world data and claims experience needed to understand and assess these changing loss patterns remain limited.[2] Where exposures lack transparency, uncertainty is difficult to quantify or seemingly separate risks prove correlated, they become harder to manage and insure. As AI becomes more deeply embedded, resilience will therefore depend on how those connections are managed.

How AI-related risks connect and spread

As AI becomes more widely adopted, the connections between systems, providers and users determine how risk may spread. A failure that appears contained in one organization can have much wider consequences when other organizations share the same dependencies or respond to the same conditions in similar ways.

Two related mechanisms are particularly relevant to the growing pervasiveness of AI: concentration and correlation. The zones describe characteristics of AI-related exposures and how losses may manifest and spread; they are not intended to represent distinct categories of risk. Concentration typically arises through a shared dependency, where disruption to a common provider or infrastructure can affect many organizations at once. Correlation is broader: It can also arise when organizations use AI in similar ways, causing otherwise separate systems to behave – or fail – in similar ways. While concentrations of providers and infrastructure can often be identified in advance, these common patterns of behavior may be much harder to anticipate.

Concentration of risk: shared dependencies

Concentration of risk arises when many organizations depend on the same provider, model, infrastructure component or service.[3] An outage, compromise, material change or withdrawal of service can create losses across otherwise unrelated organizations and sectors.

A direct supplier may itself rely on several other providers for cloud capacity, compute, foundation models, data services, identity systems or widely embedded software. As a result, organizations that appear to use different vendors may still depend on the same underlying technology. Mapping dependencies beyond the first tier can reveal these common points of exposure before one failure affects many organizations at once.

A defective software update from a widely used cybersecurity provider in 2024 illustrates the potential propagation effect. It caused around 8.5 million devices running the same operating system to crash, disrupting airlines, banks, hospitals and other organizations around the world.[4] It was not an AI-related loss, but it demonstrated how disruption at one common technology layer can spread rapidly across otherwise unrelated organizations and sectors.

Concentration of risk can be identified and, at least in principle, managed by understanding how widely dependencies are shared and what alternatives exist if they become unavailable. AI may raise the stakes as shared technology layers become more deeply embedded and harder to substitute.

Correlation of risk: shared logic and behavior

Correlation of risk can extend beyond shared providers or infrastructure. Organizations using different systems may still experience similar outcomes because those systems draw on overlapping training data, architecture or decision rules – or are optimized toward similar objectives.[5] They may therefore share the same blind spots.

> If you have mistakes, you have mistakes at scale. One of the biggest AI risks is not isolated failure, as with human errors; it's correlated failure.
>
> Underwriting Expert

The connection is behavioral, arising from how systems are designed and used, rather than from a single external point of failure. Under these conditions, individually bounded errors can become a repeated pattern across otherwise separate systems. A market shift, adversarial technique or regulatory reinterpretation could reveal or trigger failures across the market without any outage to point to.

This form of correlation may be harder to anticipate than concentration of risk. There may be no clear incident, no provider failure to identify, and no rapid restoration path, with the pattern only emerging later as claims, complaints, regulatory action or litigation. Addressing the risk requires organizations to identify and change common assumptions, decision logic or practices.

How the mechanisms interact

A useful if imperfect analogy is that concentration of risk behaves like a natural catastrophe without geography: a defined trigger and footprint determine how far and how fast loss propagates. Correlation of risk behaves more like a latent liability: cumulative, multi-defendant and difficult to attribute, determining how wide and deep loss may go.

The two patterns can overlap. A widely used provider creates concentration and if disrupted can cause losses to move together. But correlation can extend beyond those identifiable dependencies: A market may have several providers and still produce similar outcomes if systems draw on similar data, assumptions and design choices.

> The AI itself brings new risks of accumulation that has been largely unexplored. Whenever we have connection points with many, essentially pulling from the same model, same database, same points, anything that goes wrong will multiply.
>
> Accumulation Underwriting Expert

Managing these risks therefore requires two separate questions: Where are organizations exposed to the same dependencies? And what could cause separate systems to behave alike?

Box 2.

Concentration and Correlation of risk

Dimension	Concentration	Correlation
Primary link	Shared provider, model, platform, data source or component	Similar data, objectives, architecture, decision logic or behavior
Typical trigger	Failure, compromise, degradation or withdrawal of a shared dependency	Common conditions revealing shared assumptions, blind spots or failure patterns
Loss pattern	Event-like, with an identifiable node and footprint	Latent or cumulative, with less obvious causation and dispersed manifestation
Key evidence	Dependency maps, provider shares, substitution paths and recovery times	Outcome distributions, scenario tests, model diversity and behavioral comparisons
Main response	Resilience, substitutability and aggregation controls	Diversity, challenge of common assumptions, tail testing and oversight

Source: Zurich Insurance Group and University of St. Gallen analysis, informed by the Financial Stability Board (2024), European Systemic Risk Board (2025).



Navigating AI-related risks

Different AI-related risks require differentiated responses. Some are already relatively well understood and likely to remain “contained” using existing controls and approaches. Others may require more evidence, stronger governance or closer cooperation between different stakeholders.

To help distinguish between these scenarios, the map brings the five insurability factors together in two practical dimensions. The first concerns how well a risk can be understood: whether the exposure can be recognized (identifiable), its likelihood and potential losses assessed (measurable), and its causes and responsibility established (attributable). Together, these factors help position a risk on the spectrum from “observable” to “opaque”.

The second concerns the potential reach of losses: whether their extent can be limited (bounded) and whether exposures are sufficiently independent for losses to be pooled rather than occur together (diversified). These factors help position a risk on the spectrum from “contained” to “system-wide spread”. Concentration and correlation explain how shared dependencies or similar behavior can cause losses across organizations, making them harder to contain and diversify.

Together, these two dimensions, which describe characteristics of AI-related exposures rather than distinct categories of risk, point to what is needed to strengthen resilience: from risk-management and risk-transfer approaches, to better evidence, stronger controls or broader coordination. Insurability is an outcome of these characteristics, rather than a fixed characteristic of the risk itself. Greater transparency, testing, stronger governance or reduced dependencies make a risk more manageable and insurable; growing concentration, hidden correlation or new patterns of use can move it in the opposite direction.

Mapping AI-related risk

The map below combines the two dimensions into four zones, distinguishing risks by how well they are understood and how widely losses could spread.

The zones describe characteristics of AI-related exposures and how losses may manifest and spread; they are not intended to represent distinct risk categories.

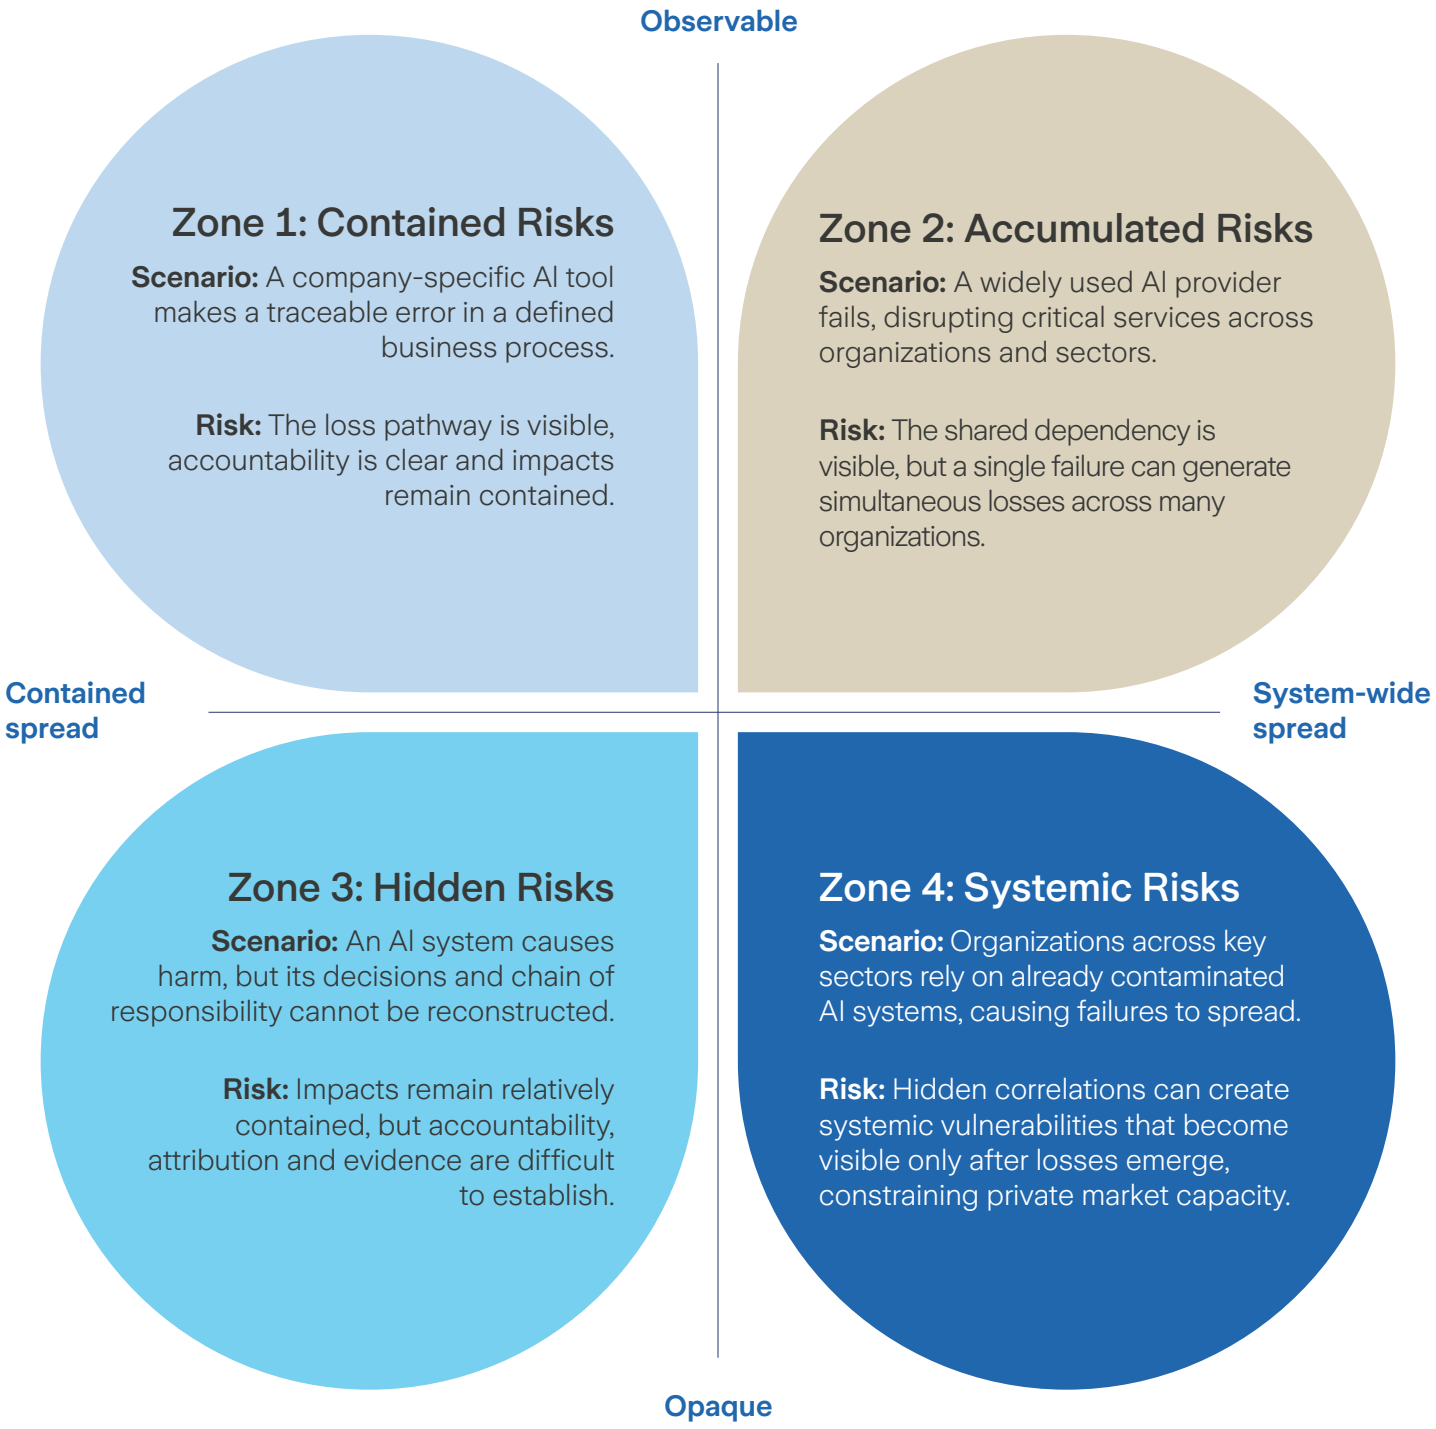

Figure 1.

Mapping AI-related risk

Zurich Insurance Group and the University of St. Gallen analysis (2026).

The map is a conceptual lens developed for this research rather than a risk-management methodology. It does not replace or modify Zurich's established risk-management frameworks or processes.

Zone 1: Contained Risks

The first zone captures risks that are relatively well understood and likely to remain contained within a defined organization, system or use case. In many cases, these risks are identifiable enough to price, reserve and settle through existing insurance products. Cyber, technology errors and omissions, professional indemnity, crime, product liability and directors' and officers' liability insurance already cover a lot of this ground.[6]

The task therefore is one of refinement, rather than the development of entirely new forms of protection. The focus is on where AI changes an existing exposure and whether existing controls and coverage are appropriate. For example:

- **Classification:** If an AI tool gives a customer incorrect advice, is the resulting loss a technology error, professional services failure or something else?
- **Liability:** If an organization is sued after an AI-powered recruitment tool systematically disadvantages certain candidates, where does responsibility sit between the organization and external provider?
- **Coverage:** If an employee authorizes a fraudulent payment in response to AI-generated communications, do existing crime and cyber policies respond as intended?
- **Terms:** If an AI-enabled machine causes physical damage, do existing terms, limits and exclusions adequately reflect the respective roles of the software, manufacturer and operator?

As claims experience develops, better loss data should allow insurers to differentiate these risks more precisely.

Zone 2: Accumulated Risks

The second zone captures risks that are relatively visible but have the potential to accumulate across many organizations. The pathway may vary through widespread reliance on a critical model or infrastructure provider; through AI-enabled techniques, such as synthetic identity or deepfake fraud, which can exploit common control weaknesses across many organizations; or through similar deployment practices and interactions between AI systems that create common patterns of exposure. In each case, understanding the extent of potential propagation is essential to managing accumulation risk.

A widely reported fraud case in 2024 illustrates this. An employee at a multinational company authorized payments after joining a video call in which other apparent participants had been synthetically generated.[7] For one organization, this is a relatively contained crime and controls failure. The accumulation risk emerges if the same technique can exploit similar processes, controls or platform weaknesses across many organizations.

Managing that potential accumulation requires action on both sides of the risk transfer. Where conventional indemnity is difficult to apply at scale, structured or parametric approaches may offer alternative ways to respond to clearly defined events or thresholds.

Zone 3: Hidden Risks

The third zone captures risks whose expected reach appears relatively contained but where legal, evidentiary or technical gaps make the likelihood or potential scale of loss difficult to assess with confidence. For example:

- **Intellectual property:** A company uses generative AI in a defined business process, but uncertainty around training data, generated content and the respective responsibilities of model provider and user makes the resulting exposure difficult to assess.
- **Autonomous agents:** A company gives an AI agent authority to negotiate contracts, purchase services or execute transactions with limited human approval. There may be no established loss history showing how often unintended actions occur or how severe they could become.
- **Multi-agent systems:** Several autonomous systems interact to complete a defined business process. Each may be understood individually, but there may be insufficient evidence about how their interaction changes the probability or nature of failure.[8]

The priority in this zone is to make these risks more observable and assessable. Better evidence, clearer operating boundaries, auditable decision trails and greater legal and contractual clarity can help establish how losses occur and where responsibility sits. As these conditions develop, risks may move toward Zone 1.

Zone 4: Systemic Risks

The fourth zone captures risks with potential systemic reach where the scale, connections or loss pathways are difficult to identify or quantify in advance. The exposure may involve correlated decision failures emerging across a sector, cross-border cascades between automated systems or widespread degradation of information integrity where impacts accumulate across organizations and society.

It may also include the deliberate weaponization of AI by state or other sophisticated actors where AI-enabled attacks could amplify disruption across critical infrastructure and economies and create impacts beyond the capacity of individual organizations or private insurance markets to absorb.

Moreover, potential losses may cross organizations, sectors or geographies without a single provider, event or cause that fully explains them.

No individual organization, insurer or regulator may have sufficient visibility to understand the full exposure. Managing these risks therefore depends on combining information and capabilities across organizations, sectors and, in some cases, borders. For a narrower group of high-impact risks, the resulting losses may also exceed the capacity that private insurance markets can provide on their own.

> Zurich's task is to look at it from a systemic perspective: what is our role, what is not our role? As an asset owner, what are our levers and what is out of our control? And where do we need support from policy makers or from asset managers.

Resilience Expert



How AI-related risks move

The map is not static. Changes in adoption, technology, evidence and controls can alter both how visible a risk is and how widely potential losses could spread. Some developments increase uncertainty or the potential reach of loss, moving risks closer to the boundaries of insurability. Conversely, better evidence, controls and resilience can move them toward a position where they are more manageable for organizations, policymakers and insurers alike.

Adoption changes reach

A model error affecting one bespoke deployment may begin in Zone 1. If the same model becomes widely adopted, the exposure may move toward Zone 2 – with the potential consequences increasing further as AI becomes embedded in more consequential processes.

Adoption can also change the scale of existing threats. AI may, for example, reduce the cost of fraud, social engineering or vulnerability discovery, increasing their frequency or potential reach.

Dependencies influence risk

A critical provider may create an identifiable source of widespread exposure in Zone 2 – but greater substitutability, recovery capacity and exposure transparency can reduce that concentration and move the risk back toward Zone 1.

The direction of travel therefore depends on how well the shared dependency is understood and whether effective and practicable alternatives are built into the system. Consolidation around models, compute and data can reduce substitutability and increase accumulation across organizations.

Evidence improves visibility

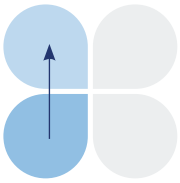

A novel use of autonomous AI may begin in Zone 3 because there is insufficient experience to assess its behavior and potential loss profile with confidence. This is particularly relevant as systems move from recommending to acting. As delegated authority increases, organizations have less established evidence about how systems will behave under limited supervision – or how effective intervention and recovery mechanisms will be.

Testing, audit trails, defined operating boundaries and clearer accountability can make the exposure better understood, moving it toward Zone 1.

Patterns shape exposure

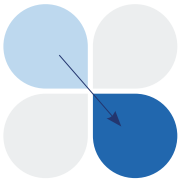

A discriminatory outcome arising from one AI-supported decision may initially appear to be a Zone 1 risk. But if similar outcomes emerge across organizations using apparently different systems, shared data, assumptions or decision logic may reveal a previously hidden correlation, moving the exposure toward Zone 4.

Movement across the map will not always be immediately visible. No single indicator proves that a risk is becoming harder to manage or insure, but a consistent pattern can show that exposure is outrunning evidence, prevention or controls:

- Rapid exposure growth without equivalent improvement in dependency data.
- Rising concentration ratios in customers' technology supply chains.
- Material gaps between stated governance and deployment practice.
- Repeated incidents or claims clustering around the same provider, component, model or control weakness.
- Claims emerging across separate lines of business from a common event.

Read alongside regulatory information and technical evidence, these signals can indicate where exposure may be moving faster than organizations can understand or control it.

Insurers can contribute two additional early-warning signals:

- Changes in terms and capacity available, including higher premiums, narrower terms, lower limits or broader exclusions.
- Changes in the information or capital required to support that capacity, including more extensive information requirements or reduced reinsurance appetite.[9]

These shifts may reveal concern before the full scale of an exposure is apparent, providing a valuable signal to policymakers, rather than simply reflecting a commercial outcome.

Practical steps to enhance resilience

Where an AI-related risk sits on the map depends on its business context, not just the technology. The quality of evidence, transparency of dependencies, strength of controls and oversight, clarity of responsibility and ability to recover from failure all shape resilience, as does the extent to which dependencies are concentrated.

From mapping risks to informing action

The same AI application may be manageable in one organization or market but pose greater risks in another. The map points to two priorities: improving visibility where risks are opaque, and limiting spread where losses could affect many organizations. Hidden Risks in Zone 3 require better evidence and clearer responsibility; Accumulated Risks in Zone 2 require stronger alternatives and recovery; Systemic Risks in Zone 4 require both, supported by coordinated action across organizations, sectors and borders. Even Contained Risks in Zone 1 need continued monitoring as adoption and dependencies evolve.

The following recommendations set out how organizations, technology providers, insurers and public authorities can contribute. Refining these to provide further practical tools will be part of the ongoing work of Zurich's AI Lab.

Managing AI within organizations

As AI becomes embedded in business processes and value chains, failures can have wider consequences. Organizations should scale governance, oversight, intervention and recovery measures to reflect each system's decision-making authority, operational importance, reliance on external providers and potential impact on other activities.

- **Map AI use and dependencies.** Identify significant AI applications, the authority granted to them, and the providers, models, data, and infrastructure they depend on, including dependencies beyond direct suppliers. This reduces opacity and reveals shared points of failure that could place exposures in Zone 2-Accumulated Risks.
- **Set limits and assign accountability.** Define what these systems may do, who is accountable for the risks, and when human review or intervention is required. Clear responsibility addresses opacity in Zone 3-Hidden Risks and Zone 4-Systemic Risks; operating boundaries help contain potential.
- **Prepare for disruption and recovery.** Establish practical options to maintain critical activities, override AI decisions, switch to alternatives, or safely pause or shut down affected systems. These measures address the potential for spread in Zone 2-Accumulated Risks and Zone 4-Systemic Risks.
- **Learn from incidents and near misses.** Record them consistently and preserve the evidence to establish what happened, identify recurring patterns and clarify responsibilities. This makes risks in Zone 3-Hidden Risks and Zone 4-Systemic Risks easier to assess and helps to detect changes that were previously contained.

Building resilience into technology

Technology providers shape the risks their customers inherit. By making critical dependencies transparent, demonstrating resilience, and providing practical recovery options, they can help customers make those risks better understood and less likely to spread.

- **Make critical dependencies visible.** Give customers enough information about the infrastructure, models, data, and services they rely on to assess concentration of risk and identify alternatives. This can reveal shared exposure that individual assessments overlook.
- **Test resilience as systems evolve.** Test how systems respond to failure, degraded performance,

and unexpected behavior, including common failure patterns across customers. Share findings to reduce uncertainty and identify safeguards against risk spreading.
- **Make responsibility clear and traceable.** Define the respective responsibilities of providers and deployers for updates, monitoring, intervention, and recovery, and preserve the records needed to establish what happened when something goes wrong. This addresses the evidentiary and contractual gaps characteristic of Zone 3-Hidden Risks and Zone 4-Systemic Risks.
- **Build in practical recovery options.** Design services and commercial arrangements so customers can maintain critical activities, restore disrupted services, or switch to alternatives promptly when a critical service fails. Test whether these options remain usable during widespread disruption, helping move accumulated exposures toward more contained conditions.

Turning insurance insights into action

Insurers and reinsurers can draw on their portfolio-wide view to identify shared dependencies and emerging loss patterns. These insights can guide action to prevent losses, limit their spread and recurrence, manage accumulation, and sustain capacity for risk transfer.

> Where are you using AI? How are you using it? Where is it coming through? This is becoming the number one question we ask at underwriting meetings.
>
> Cyber Underwriting Expert

- **Manage accumulation across portfolios.** Track exposure to critical providers and shared dependencies, using risk modeling and scenario analysis to test both common points of failure and correlated behavior in Zone 2-Accumulated Risks and Zone 4-Systemic Risks, informing exposure limits and safeguards.
- **Underwrite how AI is used.** Look beyond adoption to assess governance, delegated authority, dependencies, controls, and recovery capabilities. Identify the evidence needed to make opaque risks assessable and the safeguards needed to contain potential losses as part of routine underwriting.
- **Turn loss data into early warnings.** Collect comparable incident and near-miss data across customers and insurance lines. This can make hidden risks more observable and reveal when apparently separate exposures are accumulating. Share aggregated insights with public bodies where they can improve prevention, while protecting confidentiality and competition.
- **Clarify the path to broader coverage.** Specify the evidence and controls that support greater capacity. Where conventional insurance falls short, explore alternative risk-transfer structures tied to clearly defined events or thresholds. These solutions may provide additional capacity, but they do not reduce the underlying exposure or make potential losses more contained.

Strengthening AI safeguards through policy

Policymakers and regulators can make AI-related risks easier to identify, trace, and manage across organizational boundaries. Existing AI governance and operational resilience frameworks provide a foundation. Building on them requires a clearer view of shared dependencies and emerging loss patterns, alongside the transparency and accountability needed for effective risk transfer.

- **Establish clear accountability.** Define the obligations and liabilities of developers, providers, deployers, and users. This reduces legal opacity in Zone 3-Hidden Risks and Zone 4-Systemic Risks. Scale requirements to potential impact and systemic reach, not AI use alone.[10]
- **Expose shared vulnerabilities.** Support proportionate disclosure of material dependencies and comparable incident reporting. This makes opaque exposures more observable and reveals accumulation across organizations, while protecting genuinely sensitive information.
- **Stress-test the wider system.** Assess failures or degradation in widely used models, infrastructure, and data services, alongside correlated behavior across otherwise separate systems. Use findings to strengthen continuity and coordinated response, addressing the spread potential of Zone 2-Accumulated Risks and Zone 4-Systemic Risks.
- **Strengthen the foundations for risk transfer.** Improve legal and evidentiary clarity so exposures can be assessed and insured. For the narrow set of extreme exposures that cannot be diversified or fully absorbed by private insurance markets, consider public-private risk sharing, drawing on natural catastrophe and terrorism arrangements while preserving incentives for prevention and private-market participation.[11] Such arrangements address loss-bearing capacity, not the underlying opacity or systemic reach of the risk.

Keeping AI-related risk manageable

AI is becoming part of the infrastructure on which businesses and societies depend. As it becomes more deeply embedded, failures that appear isolated may spread through shared providers, data, infrastructure and patterns of behavior. Insurance provides a practical lens of whether these risks remain manageable by asking whether exposures are visible, responsibility can be established and potential losses can be contained and absorbed.

The four-zone framework makes clear that manageability is not fixed: adoption can widen exposure; evidence, accountability, and controls can contain it. Insurability offers a practical test. Changes in coverage, terms, or capacity can reveal gaps in understanding, governance, diversification, or loss-bearing capacity. Alongside technical and regulatory evidence, these signals help identify where stronger safeguards or collective action are needed.

Insurance contributes more than financial compensation. Coverage, accumulation analysis, and loss experience can guide prevention and provide early warnings. But risk transfer cannot replace risk reduction or absorb every systemic loss. A narrow set of extreme exposures may require public-private arrangements that preserve incentives for prevention and private-market participation.

> “Let’s not figure it out after there is a catastrophic event. Let’s try and be more thoughtful about what are these plausible pathways.
>
> Underwriting Expert

No institution sees the whole exposure. Businesses, technology providers, insurers, public authorities, and researchers must connect evidence across sectors and borders. Comparable testing, documentation, and incident reporting, backed by clear accountability, can make shared dependencies easier to manage without imposing uniform technical rules.

The test of progress is not how quickly AI is adopted, but whether the capacity to manage its consequences keeps pace. That means exposing dependencies, assigning responsibility, containing disruption, and proving recovery works before a major loss tests those arrangements. Keeping failure within society’s capacity to absorb it is not a constraint on AI’s potential. It is a condition for realizing it.



Methodology

This report forms part of Zurich's contribution to the discussion leading up to the Geneva Summit 2027. The findings presented here represent a stage in a continuing research program rather than its conclusion.

This report reflects findings from a joint research initiative between Zurich and the University of St. Gallen, undertaken through the Zurich AI Lab, publicly launched in October 2025. The analytical framework was developed for this research and external discussion and does not replace Zurich's established internal risk-management framework, methodologies or terminology. The collaboration combines Zurich's practical experience in risk assessment, loss prevention and claims management with academic research into how AI transformation is shaping business and society. The aim is to draw on expertise from both fields to better understand the conditions needed for organizations and societies to adopt AI safely and with confidence.

Data Collection

Researchers from the University of St. Gallen conducted 50 semi-structured interviews between July and September 2026. Semi-structured interviews are particularly appropriate for examining a rapidly evolving phenomenon. Participants were treated as "knowledgeable agents" capable of providing informed interpretations of the phenomena they encounter in their professional environments **(Gioia et al., 2013).**

The sample comprised 36 Zurich executives and 14 external experts selected for their complementary areas of expertise and institutional perspectives. Participants were selected through theoretical sampling to capture a range of relevant expertise and experience. The sample covered multiple functions and geographies, with participants based in Europe, Asia, North America and Africa. Within Zurich, participants represented risk, underwriting, claims, technology, cyber, legal and resilience functions.

Data collection was iterative, with external experts incorporated throughout to challenge, complement and extend perspectives emerging from within the insurance industry.

---

1 The factors in Box 1 draw on the established insurability literature, particularly Baruch Berliner, Limits of Insurability of Risks (Englewood Cliffs, NJ: Prentice Hall, 1982), and Baruch Berliner, **"Large Risks and Limits of Insurability,"** The Geneva Papers on Risk and Insurance 10, no. 37 (1985): 313–329.

2 Sasha Romanosky and Celine Robinson, **The Insurability of Artificial Intelligence**, Research Report RR-A5130-1 (Santa Monica, CA: RAND Corporation, 2026).

3 Financial Stability Board, **The Financial Stability Implications of Artificial Intelligence**, 2024; European Systemic Risk Board, Advisory Scientific Committee, **Artificial Intelligence and Systemic Risk**, 2025.

4 US Cybersecurity and Infrastructure Security Agency, **"Widespread IT Outage Due to CrowdStrike Update,"** 19 July 2024; CrowdStrike, **Channel File 291 Incident Root Cause Analysis**, 6 August 2024.

5 European Systemic Risk Board, Advisory Scientific Committee, **Artificial Intelligence and Systemic Risk**, ASC Report No. 16, December 2025; Financial Stability Board, **The Financial Stability Implications of Artificial Intelligence**, 2024.

6 Sasha Romanosky and Celine Robinson, **The Insurability of Artificial Intelligence**, Research Report RR-A5130-1 (Santa Monica, CA: RAND Corporation, 2026); Geneva Association, **Gen AI Risks for Businesses: Exploring the Role for Insurance**, 2025.

7 Professional Risk Managers' International Association, **The Arup Deepfake Fraud: A Case Study**.

8 National Institute of Standards and Technology, **Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile**, NIST AI 600-1 (Gaithersburg, MD: NIST, 2024).

9 Financial Stability Board, **Monitoring Adoption of Artificial Intelligence and Related Vulnerabilities in the Financial Sector**, 10 October 2025.

10 National Institute of Standards and Technology, **Artificial Intelligence Risk Management Framework 1.0**, NIST AI 100-1 (Gaithersburg, MD: NIST, 2023); OECD, **Towards a Common Reporting Framework for AI Incidents** (Paris: OECD Publishing, 2026); Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence, Official Journal of the European Union, 12 July 2024.

11 U.S. Department of the Treasury, **"Terrorism Risk Insurance Program"**; Pool Re, **"Official Documentation"**; Flood Re, **"Official Documentation."** These arrangements employ different forms of public-private participation for difficult-to-insure risks and differ materially in coverage, funding and government involvement. They are cited as broad institutional precedents, not as ready-made models for AI.

Acknowledgments

This report was produced through a collaboration between Zurich Group and the University of St. Gallen, bringing together insurance and risk expertise with academic research.

We are particularly grateful to Tomoko Yokoi and Caroline Mann of the University of St. Gallen for their contribution to the research, interviews and analysis underpinning this report.

We are grateful to the Zurich colleagues who contributed their time and expertise to the research, including **Paige Adams**, Group Chief Information Security Officer; **Maria Apazoglou**, Group Head of AI Engineering & Platforms; **Jay Bell**, Head of Business Transformation ZNA; **Jörg Bertogg**, Head of Commercial Insurance Switzerland; **Rene Billing**, Member of the Board, DA Direkt; **Vivien Bilquez**, Global Head Cyber Resilience; **Debra Burford**, Global Head of Financial Lines Claims; **David Campos**, former General Counsel Data, Technology, Operations & IP; **Keith Daly**, Chief Claims Officer ZNA; **Eliya Devay**, Responsible Investment Analyst; **Yvonne Dharshing**, Head of Data, Ethics, Governance & Controls, Compliance; **Antony Elliott**, Group Head of AI R&D; **Mark Fletcher**, Head of Zurich Resilience Solutions APAC; **Kara Higginbotham**, Head of Professional Liability & Cyber, ZNA; **Ross Hutchison**, Head of Eurozone Economics & Market Strategy; **Vanessa Isler**, Group Head of Regulatory Affairs; **Adrian Jenner**, Head of Financial Lines; **Louise Kerrigan**, Global Head of Casualty; **Sami Mahmal**, Head of Data & AI, APAC; **Robin Matthias**, Global Head of Zurich Resilience Solutions; **Doug Meyers**, SVP, Claims Chief Operations Officer; **Guy Miller**, Chief Market Strategist & Economist; **Clairy Moraitou**, Group Head of Customer, Market and Brand Insights; **Ha Nguyen**, Head of Market Strategy Asia; **Don Salisbury**, Chief Pricing Actuary Specialties; **Penny Seach**, former Chief Underwriting Officer CI International; **Risto Schmid**, Head Climate Resilience Switzerland; **Valeria Schmitke**, General Legal Counsel, Latin America; **Arooran Sivasubramaniam**, Head of Zurich Resilience Solutions U.S.; **Iwan Stalder**, Head of Group Accumulation Management; **Tom Thornberry**, Chief Claims Officer; **Stephan Unterberger**, Group CI Technical Director; **Susie Watt**, Head of Underwriting Transformation; **Michael White**, Head of Complex & Strategic Litigation; **Richard Winata**, Senior Compliance Expert, Data, Ethics, Governance & Controls; and **John Shane**, Head of Casualty Claims, CI International.

In addition to Zurich's internal experts, we would like to thank the external participants who contributed their perspectives and expertise to this research.

PD Dr. Sandra Andraszewicz, Senior Academic Researcher and Lecturer at the Chair of Cognitive Science, ETH Zurich; **Dr. Bastian Bergmann**, Executive Director, ETH Zurich FinsureTech Hub; **Anna Sophia Burch**, PhD Candidate & Research Associate, University of St. Gallen; **Dave Chatterjee**, Ph.D., Adjunct Associate Professor at Duke University; **Öykü Işık**, Professor of Global Digital Transformation at Thunderbird School of Global Management at Arizona State University & Co-Chair, Global Future Council on Cybersecurity at World Economic Forum; **Dr. Dag Lohmann**, Co-Founder & CEO Emeritus of KatRisk; **Vincent Martin**, Director of the FAO Office of Innovation (OIN); **Bozidar Stojadinovic**; Professor at Chair of Structural Dynamics and Earthquake Engineering at ETH Zurich; **Dr. Eva Thélisson**, Co-Founder and Executive Director of the AI Transparency Institute; **Lukas Walker**, Project Manager at Intersections & Co-Founder of Ethics Forum at VSETH; **Isabelle Wildhaber**, Professor of Private and Business Law, University of St. Gallen; **Rolf Wüstenhagen**, Professor & Director, Institute of Responsible Innovation, Sustainability and Energy (RISE-HSG), University of St. Gallen and one contributor who wished to remain anonymous.

Contact

Angel Serna
Senior Vice President of Public Affairs
angel.serna@zurich.com

Clara Löhr
Communications Manager
clara.loehr@zurich.com

About us

Zurich Insurance Group (Zurich) is a leading global multi-line insurer founded more than 150 years ago, which has grown into a business serving more than 82 million customers in more than 200 countries and territories, while delivering industry-leading total shareholder returns. Reflecting its purpose to 'create a brighter future together,' Zurich offers protection services that go beyond traditional insurance, to support its customers in building resilience. Since 2020, the Zurich Forest project has supported reforestation and biodiversity restoration in Brazil's Atlantic Forest. The Group has more than 65,000 employees and is headquartered in Zurich, Switzerland. Zurich Group Ltd (ZURN), is listed on the SIX Swiss Exchange and has a level I American Depositary Receipt (ZURVY) program, which is traded over-the-counter on OTCQX. Further information is available at **www.zurich.com.**

ZURICH

University of St.Gallen

Zurich **AI** Lab

Disclaimer:
This publication has been produced solely for informational purposes. The analysis contained and opinions expressed herein are based on numerous assumptions concerning anticipated results that are inherently subject to significant economic, competitive, and other uncertainties and contingencies. Different assumptions could result in materially different conclusions. All information contained in this publication has been compiled and obtained from sources believed to be reliable and credible but no representation or warranty, express or implied, is made by Zurich Insurance Group Ltd or any of its subsidiaries (the 'Group') as to their accuracy or completeness. Opinions expressed and analyses contained herein might differ from or be contrary to those expressed by other Group functions or contained in other documents of the Group, as a result of using different assumptions and/or criteria.

This publication is not intended to be legal, underwriting, financial investment or any other type of professional advice. No content in this publication constitutes a recommendation that any particular investment, security, transaction or investment strategy is suitable for any specific person. The content in this publication is not designed to meet any one's personal situation. The Group hereby disclaims any duty to update any information in this publication. Persons requiring advice should consult an independent adviser (the Group does not provide investment or personalized advice).

The Group disclaims any and all liability whatsoever resulting from the use of or reliance upon this publication. Certain statements in this publication are forward-looking statements, including, but not limited to, statements that are predictions of or indicate future events, trends, plans, developments or objectives. Undue reliance should not be placed on such statements because, by their nature, they are subject to known and unknown risks and uncertainties and can be affected by other factors that could cause actual results, developments and plans and objectives to differ materially from those expressed or implied in the forward-looking statements.

The subject matter of this publication is also not tied to any specific insurance product nor will it ensure coverage under any insurance policy.

This publication may not be reproduced either in whole, or in part on other communication channels, without prior written permission of Zurich Insurance Group Ltd, Mythenquai 2, 8002 Zurich, Switzerland. Neither Zurich Insurance Group Ltd nor any of its subsidiaries accept liability for any loss arising from the use or distribution of this publication. This publication is for distribution only under such circumstances as may be permitted by applicable law and regulations. This publication does not constitute an offer or an invitation for the sale or purchase of securities in any jurisdiction.