

The Triad of Exposure

Impact of organizations undermining themselves in the eyes of AI

Contents

1. Executive summary	3
2. AI readiness starts outside-in	4
3. The unknown estate is now active exposure	5
4. The Triad of Exposure	6
Misrepresentation and misinformation	6
Future visibility and competitiveness	7
The cyber problem	8
5. Insurance impact: exposure, mitigation and evidence	9
6. Financial exposure: the cost of maintaining the unknown	10
7. PDFs, documents, and hidden exposure	11
8. Why the exposure compounds over time	12
9. Discovery, Mapping, and the three-stage foundation	13
Stage one	14
Stage two	14
Stage three	14
10. What leadership should establish first	15
11. Conclusion	16
12. Bibliography	17

1. Executive summary

Organizations are exposed when AI can read digital material they no longer fully know, control, or stand behind. The exposure starts in the estate they have already built and left online. AI reads that estate whether or not the organization still treats it as current.

Gallagher's 2026 AI Adoption and Risk Benchmarking reports that 1 in 5 insurance industry respondents said a client experienced loss or claims from AI-related risks in the past year, with just over half covered by insurance.^[14] That evidence does not make this an insurance paper. It makes the business exposure harder to dismiss.

Analysis from P&C / Sitemorse risk profiling, 2017 to 2023, covering more than 100 million websites, indicates that 41% of an organization's digital footprint is unknown to digital teams.^[1] This is not a value score or a maturity rating. It is evidence of a visibility gap. Where material is unknown, leadership cannot confidently judge whether it is current, accurate, necessary, secure, or suitable to remain public.

The impact sits across the Triad of Exposure: misrepresentation and misinformation, future visibility and competitiveness, and the cyber problem. They are connected because the same unknown estate produces several effects at once.

The cyber leg is now sharper. In June 2026 a US export control directive removed two frontier AI models from the market over their cybersecurity capability.^[18] The capability that drew that response is the capability now reading the estate an organization has lost sight of.

This is why AI readiness must start outside-in. The first question is not what the organization says internally. It is what AI, customers, advisers, insurers, competitors and attackers can already read from the outside. The answer set out in this report is a discovery and mapping foundation, executive-sponsored and low-impact, that gives leadership a position it can hold within 90 days.

An organisation's AI readiness is measured by what AI reads across the footprint it has built online over years, much of it forgotten, not by what it believes it presents today.

Lloyd Golley, CxO Bureau

2. AI readiness starts outside-in

AI readiness is usually discussed as an inside-out program. Organizations focus on staff tools, governance committees, approved systems, and policy statements.

Those matters have their place, but they do not address the first public test of AI readiness. The outside world sees the organization through what is already published, discoverable, readable, and reusable.

Outside-in means assessing the organization by the evidence available from outside it. It asks what AI systems, search engines, retrieval tools, agents, customers, regulators, insurers and competitors find when they meet the digital estate. That includes the current website, and it includes material built up over years that may no longer sit with a current team.


The practical question is not whether the organization has an AI policy. It is whether the public footprint is accurate enough, structured enough, and current enough to be read correctly. A footprint can carry multiple versions of the same message: legacy product material, outdated forms, withdrawn policies, duplicated PDFs, forgotten recruitment documents, old campaign pages, and third-party copies that still carry historic information.

AI systems do not respect the internal view of what is current. They retrieve, parse, and interpret what is available. If old material remains visible, it stays part of the public evidence base from which answers are formed. If clearer competitor information is easier to retrieve, the organization may appear less accurately in comparison sets.

This is why outside-in is a foundation issue. It is about confirming what exists, understanding what remains visible, and closing the gap between what the organization believes it says today and what the public estate still says on its behalf.

3. The unknown estate is now active exposure

The scale of the unknown matters because unmanaged assets cannot be reliably assessed. AAAnow / Sitemorse analysis indicates that 41% of an organization's digital footprint is unknown to digital teams.^[1] The point is not that an unknown asset is automatically harmful. The point is that unknown assets sit outside confident ownership, review, retirement, and assurance.



The unknown estate forms over time rather than through a single decision. A campaign site is launched and left online. A product range is withdrawn, but supporting documents remain accessible. A PDF is copied to a partner portal. A form is replaced, but the older version still appears through search. A subdomain is created for a short-term need and stays live after the team has moved on.

These assets are not always prominent. Many are below the surface of normal digital operations. They may not appear in executive dashboards, content plans, brand reviews, security schedules, or agency scopes. They remain externally readable. That is the exposure. What the organization no longer manages can still be the material from which others understand it.

External research shows the wider trend moving the same way. Gartner found that 41% of employees were business technologists in 2022, creating or modifying technology outside the IT department, and predicts that figure rises to 75% by 2027.^[17] That figure measures people and tools, not the footprint share above. It points in the same direction: more technology is created and changed outside central visibility, which is how unmanaged estates form and grow.

In the eyes of AI, the unknown estate is not dormant. It is a live evidence layer. It can supply old facts, outdated terminology, inaccessible structure, and contradictory statements. That turns content neglect into organizational exposure.

4. The Triad of Exposure

The Triad of Exposure explains how unmanaged digital landscapes create connected forms of organizational exposure. They should be understood together, because the same unknown estate can produce several effects at once.

Exposure area	What creates the exposure	Executive impact
Misrepresentation and misinformation	Outdated, duplicated, inaccurate, inaccessible, or poorly structured material remains visible and may be retrieved by AI systems.	The organization may be described, compared, or summarized using material it no longer considers current.
Future visibility and competitiveness	Current information is buried, contradicted, poorly structured, or harder to retrieve than competitor information.	The organization may be misread, overlooked, or included in weaker comparison sets.
The cyber problem	Forgotten subdomains, old microsites, legacy platforms, forms, scripts, and campaign assets remain outside normal oversight.	Unknown assets expand the external surface requiring monitoring, patching, access control, and retirement.

The triad is one problem with three consequences. The same unmanaged footprint feeds each column.

Misrepresentation and misinformation

Misrepresentation and misinformation begin with a practical content problem. The organization may hold accurate current information, and it may also have older, weaker material still available online. AI systems can meet both. They may retrieve a superseded policy, a product page that no longer reflects the offer, a form that is no longer correct, or a document carrying historic terminology.

This creates exposure because AI systems assemble answers from what they can retrieve and interpret. Academic work on retrieval augmented generation and document analysis shows that parsing quality and document structure affect retrieval quality.^{[8] [9] [10]} Poorly structured documents are more than an accessibility issue. They weaken the material available for AI systems to ground an answer.

The problem is not limited to one web page. Many organizations hold the same information across PDFs, HTML pages, campaign assets, portals, partner sites, and archived areas. When those sources conflict, the organization has created the conditions for confused external interpretation. AAAnow / Sitemorse analysis indicates that 37% of PDFs become untracked once they leave the producer's website, including through circulation in reports and product catalogs.^[1] A correction at source does not automatically correct copies elsewhere.

Future visibility and competitiveness

Future visibility is no longer limited to search ranking. It includes whether AI systems understand the organization correctly and place it in the right comparison set. When people ask AI tools to compare providers, summarize options, explain products, or answer buyer questions, the result depends on the public evidence base the system can read.

McKinsey reports that brand-owned sites make up 5 to 10% of the sources AI search references in many categories, with the rest coming from publishers, affiliates, user-generated content, and third-party material.^[2] AirOps analysis of more than 21,000 brands reports that 85% of brand mentions in AI search come from third-party content rather than the brand's own site.^[3]

Brand-owned sites are 5 to 10% of the sources AI search references



Source: McKinsey, New front door to the internet, 2025

An organization cannot assume its main website controls the whole picture. The AI-visible picture is assembled across many sources. If historic material remains available, if product pages are inconsistent, if PDFs cannot be parsed, or if current information is weaker than competitor information, the organization may be represented less accurately. The cost is not only reduced attention.

It is the risk of being understood through material the organization would not choose to present today.

The cyber problem

The cyber problem sits in the same landscape. An unmanaged digital estate is not only a content issue. Forgotten subdomains, old microsites, legacy platforms, unmanaged forms, redundant scripts, and abandoned campaign assets can sit outside normal monitoring and ownership. Security teams cannot confidently protect what the organization has not mapped.

The exposure is straightforward. Unknown externally visible assets may sit outside patching schedules, access review, certificate renewal, vendor management, monitoring, and retirement decisions. Unmaintained software is unpatched software. Old domains, forgotten subdomains, and legacy applications are the openings, and they multiply as the estate grows beyond what anyone tracks.

The capability reading that estate is advancing. On 12 June 2026 a US government export control directive suspended access to two frontier AI models, Anthropic's Fable 5 and Mythos 5.^[18] The concern the government cited, as Anthropic described it, was a method of bypassing the models' safeguards to reach their cybersecurity capability, with the reported technique being the model reading a codebase to find software flaws.^[19] Anthropic disputes that a narrow jailbreak warranted removal and notes the same capability exists in other public models.

The point for this report holds either way. A government treated a frontier model's capacity to read code and find flaws as serious enough to act on. That same capacity reads an organization's oldest, least-watched, externally visible assets. The cyber problem is connected to the other two exposures, because the same forgotten asset can carry outdated content, mislead AI systems, weaken visibility, and enlarge the surface that must be monitored.

5. Insurance impact: exposure, mitigation and evidence

Insurance is not separate from the Triad of Exposure. It is one of the places where exposure becomes financially visible. Misrepresentation, misinformation and cyber exposure can each lead to questions about loss, liability, policy wording, exclusions, limits, controls, and the evidence an organization can produce to show mitigation.

Gallagher's 2026 AI Adoption and Risk Benchmarking reports that AI errors, misinformation and hallucinations are the leading perceived threat from AI adoption, cited by 57% of respondents.^[14] The same report says one in five insurance industry respondents had a client experience loss or claims from AI-related risks in the past year, with just over half fully covered. Gallagher also notes that AI-related exposures are driving exclusions, endorsements, bolt-on covers, and bespoke AI policies.^[14]

Exposure	Insurance relevance	Evidence leadership should be able to show
Misrepresentation	AI may present the organization in a way leadership would not approve, creating customer, adviser, regulatory, or reputational questions.	A mapped public estate, current authoritative content, removed or corrected outdated material, and a record of decisions.
Misinformation	AI may use old or conflicting material to produce inaccurate answers that create loss, complaints, or dispute over what the organization represented.	Evidence that high-risk content has been identified, prioritized, corrected, consolidated, or withdrawn.
Cyber exposure	Unknown assets may sit outside normal security control, which can affect risk quality and underwriting confidence.	A view of externally visible assets, ownership, retirement decisions, and reduction of unmanaged surfaces.

Insurance is built on risk transfer after mitigation, not as a replacement for it. RUSI describes cyber insurance as a mechanism for transferring residual risk after other risk management practices have been applied.^[16] OECD analysis notes that insurers increasingly use data, analytical tools, and engagement platforms to support risk assessment and policyholder risk reduction.^[15] Evidence of action therefore matters.

Organizations that can demonstrate they are identifying, assessing, and reducing exposure are better placed in underwriting and renewal discussions than organizations relying on assertion. No outcome is automatic. The value lies in having evidence of mitigation, which gives executives a stronger basis to discuss risk quality, coverage terms, limits, and exclusions.

6. Financial exposure: the cost of maintaining the unknown

The cost of the unknown estate is hidden because it is absorbed into normal budgets. There is no single line item called unmanaged digital exposure. The cost is distributed across hosting, storage, CMS licensing, technical support, content management, agency maintenance, search work, accessibility review, security monitoring, customer service, legal input, and repeated clean-up.

This makes the exposure easy to tolerate. A redundant site may not look expensive on its own. A duplicated PDF may not appear material in isolation. Multiplied across hundreds or thousands of assets, the organization can pay to maintain a landscape that works against accuracy, visibility, security, and efficiency.

Cost area	How the cost appears	Why it matters
Hosting and storage	Sites, files, documents, and media continue to consume infrastructure resources.	The organization pays to maintain material that may no longer support its objectives.
CMS and platform support	Old templates, domains, microsites, and content systems remain supported.	Digital teams carry complexity that slows change and raises maintenance effort.
Search and content management	Teams manage competing versions, duplicate pages, and outdated assets.	The public evidence base becomes harder to control and correct.
Security monitoring	Unknown assets increase the surface requiring oversight and review.	Security effort is spent defending assets that may no longer need to exist.
Accessibility remediation	PDFs and legacy documents require repeated review and correction.	Manual remediation becomes expensive when scale and duplication are unmanaged.
Customer and brand effort	Wrong or old information causes confusion, complaints, and clarification work.	The organization pays again when public material creates avoidable friction.

External cost figures should be used carefully here, because costs vary by sector, estate size, and governance model.^[13] The strongest conclusion is that unmanaged estates create recurring cost centers that stay invisible until the organization maps the assets and understands the work attached to them. The organization pays to keep material online that may no longer support its objectives, while that same material can damage accuracy, weaken visibility, and raise the future cost of clean-up.

7. PDFs, documents, and hidden exposure

Documents are a major part of the hidden exposure layer. They accumulate over years because they are easy to publish, easy to copy, and easy to leave behind. Annual reports, brochures, forms, policy documents, product sheets, recruitment material, catalogs, and guidance notes can remain available long after the organization's current position has changed.

The PDF problem matters because it joins human accessibility and machine readability. Research on PDF parsing and document structure shows that AI retrieval quality depends on recognizing structure, extracting passages, and interpreting content reliably.^{[6] [7]} A PDF that looks correct to a human reader may still be weakly structured for assistive technology and machine interpretation.

AAAnow / Sitemorse analysis indicates that 19% of PDFs on organizational websites are duplicates.^[1] Duplicate documents are not only a storage issue. They create inconsistent signals about which version is current. When duplicates spread across sites and external locations, the organization loses confidence that the public evidence base reflects its current position.

Guidance from government and accessibility bodies has pushed toward structured, machine-readable formats over design-led documents that machines cannot parse.^{[11] [12]} If the organization cannot find, classify, and understand its document estate, it cannot know where outdated, duplicated, inaccessible, or poorly structured material is creating exposure.

8. Why the exposure compounds over time

The exposure compounds because the digital estate keeps moving. New content is published. New documents are produced. Older assets become outdated. Teams change. Campaigns end. Partner copies remain. Search engines continue to surface historic material. AI systems retrieve and reuse what is available. Third-party sites add their own interpretation of the organization.

This is not a static archive issue. The estate grows while the organization's confidence in it can fall. A page that was correct three years ago may be wrong today. A form that was replaced may still be in use. A document suitable for a historic product may still be discoverable. A campaign page with no current owner may still be indexed.

The cost rises because the organization must first find and understand the estate before it can decide what should remain, change, or be removed. Discovery, ownership, review, correction, consolidation, and retirement all become larger tasks as the estate grows. The longer the unknown estate stays unmanaged, the more work is needed to establish the current position.

AI adds pressure because it creates new routes through which old material is reused. A customer may not find an outdated PDF directly, but an AI system may retrieve or summarize it. A competitor may not see a forgotten microsite, but a comparison engine may use signals from it. The organization's own view of what is current is no longer enough.

9. Discovery, Mapping, and the three-stage foundation

The first leadership task is not a broad technology deployment. It is the establishment of evidence. Leaders need to know what exists, where it sits, who owns it, whether it is current, and whether it should remain online.


Without that evidence, decisions about AI readiness rest on assumption.

AAAnow provides this through a discovery and mapping service that starts where the exposure starts, outside-in, with the estate AI can already read. It is executive-sponsored but low-impact for the organization. It begins with 2 days of client resource, then continues remotely and continuously.

Within 90 days, leadership can see what exists, where exposure sits, and what should be kept, corrected, or removed. The work is independent of any single vendor's AI products and is sector-neutral.

The AAAnow Discovery and Mapping service

Executive-sponsored, low-impact. 2 days of client resource to start, then continuous and remote.



Stage one

Discovery and mapping. It identifies the full digital landscape visible from the outside, across domains, subdomains, documents, and third-party copies. It locates sites the organization has lost track of, identifies third-party properties carrying its content, and surfaces the PDFs spread across owned sites, partner sites, and supplier portals. Discovery runs continuously, because the footprint is not static.

Stage two

Establishes value and risk. Each mapped asset is assessed against a human-centric scale: an A to F value ranking and a 1 to 5 importance baseline. Value is measured through Experience, Integrity, and SEO. Risk is measured through Privacy, Accessibility, and Carbon. The assessment gives leadership a position scored against a clear framework, rather than a list of assets without judgment.

Stage three

Decision and remediation. High-value assets are kept and structured properly. Necessary assets are remediated where privacy, security, accessibility, or structure issues are found. Redundant, obsolete, or trivial assets are removed, a defensible disposition that reduces cost and exposure while leaving a clean, high-integrity foundation for AI readiness.

The service does not require each asset to be treated as equally important. It requires the organization to know what exists and to apply consistent judgment. That is the practical move from unmanaged exposure to managed evidence, and it is the evidence base the board can hold.

10. What leadership should establish first

The leadership questions are direct. They are operational, and they carry board-level consequences. Each one is answerable once the estate has been discovered, mapped, and scored.

Leadership question	Why it matters
What do we have online?	The organization cannot manage an estate it has not mapped.
Who owns it?	Unowned assets fall outside review, correction, and retirement.
Is it current and accurate?	Old material can still shape how AI systems explain the organization.
Is it accessible and structured?	Weak structure affects people, retrieval systems, and AI interpretation.
Is it necessary?	Unnecessary material creates cost without strategic benefit.
Is it creating unnecessary risk?	Unknown assets can create brand, cyber, accessibility, insurance, and operational exposure.

11. Conclusion

The Triad of Exposure is the practical consequence of unmanaged digital growth. Misrepresentation and misinformation, future visibility and competitiveness, and the cyber problem are connected by the same underlying issue. Organizations do not fully know what they have online, yet the outside world can still read it, retrieve it, compare it, and act on it.

This is why the impact of organizations undermining themselves in the eyes of AI is now a senior leadership issue. The exposure is not created by AI alone. It comes from the public material that remains online without confident ownership, accuracy, structure, and purpose. AI makes that exposure more visible because it can turn old material into current answers.

The insurance and financial points strengthen the case. Organizations that can show practical mitigation across misrepresentation, misinformation, and cyber exposure are better placed to discuss risk quality, coverage cost, limits, and indemnity terms. They also stop paying twice, once to maintain the unknown, and again to understand, correct, defend, or remove it.

AI readiness outside-in begins with the digital landscape that already exists. The organization cannot manage what it has not mapped. It cannot reduce exposure it has not identified. The first AI readiness decision is not what to build next. It is what exposure the organization is already carrying.

12. Bibliography

The following sources support the external evidence base used in this report.

- [1] P&C . Sitemorse risk profiling, 2017 to 2023, covering more than 100 million websites.
<https://www.aaanow.ai/>
- [2] **McKinsey & Company.** New front door to the internet: Winning in the age of AI search.
<https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/new-front-door-to-the-internet-winning-in-the-age-of-ai-search>
- [3] **AirOps.** The Influence of Offsite Signals in AI Search.
<https://www.airops.com/report/the-influence-of-offsite-signals-in-ai-search>
- [6] **arXiv.** UDA: A Benchmark Suite for Retrieval Augmented Generation in Real-world Document Analysis.
<https://arxiv.org/abs/2406.15187>
- [7] **arXiv.** Revolutionizing Retrieval-Augmented Generation with Enhanced PDF Structure Recognition.
<https://arxiv.org/abs/2401.12599>
- [8] **arXiv.** Capturing Logical Structure of Visually Structured Documents with Multimodal Transition Parser.
<https://arxiv.org/abs/2105.00150>
- [9] **arXiv.** Understanding the Logical and Semantic Structure of Large Documents.
<https://arxiv.org/abs/1709.00770>
- [10] **arXiv.** Structured Linked Data as a Memory Layer for Agent-Orchestrated Retrieval.
<https://arxiv.org/abs/2603.10700>
- [11] **Section508.gov.** Create accessible PDFs.
<https://www.section508.gov/create/pdfs/>
- [12] **Government Digital Service.** Why GOV.UK content should be published in HTML and not PDF.
<https://gds.blog.gov.uk/2018/07/16/why-gov-uk-content-should-be-published-in-html-and-not-pdf/>
- [13] **CACI.** Website sprawl: the cost of CMS fragmentation.
<https://www.caci.co.uk/blog/website-sprawl-cost-cms-fragmentation/>
- [14] **Gallagher.** The 2026 AI Adoption and Risk Benchmarking survey.
<https://www.ajg.com/news-and-insights/features/ai-adoption-and-risk-benchmarking-2026/>
- [15] **OECD.** Leveraging technology in insurance to enhance risk assessment and policyholder risk reduction.
<https://www.oecd.org/finance/leveraging-technology-in-insurance-to-enhance-risk-assessment-and-policyholder-risk-reduction.htm>
- [16] **RUSI.** Cyber Insurance and the Cyber Security Challenge.
<https://static.rusi.org/247-op-cyber-insurance-fwv.pdf>
- [17] **Gartner.** Rise in business technologists: 41% of employees in 2022, predicted to reach 75% by 2027.
<https://www.gartner.com/en/newsroom/press-releases/2022-03-13-gartner-survey-finds-rise-in-business-technologists-is-driving-funding-for-tech-purchases-outside-of-it>
- [18] **Anthropic.** Statement on the US government directive to suspend access to Fable 5 and Mythos 5, 12 June 2026.
<https://www.anthropic.com/news/fable-mythos-access>
- [19] **Fortune.** Anthropic disables Fable and Mythos AI models following US government export ban, 13 June 2026.
<https://fortune.com/2026/06/13/anth>