SUPPLY CHAIN RESEARCH · EDITORIAL FEATURE

When the Software Is the Disruption

Vendor Concentration and the Risk Your Resilience Programme Does Not Price

Supply chain organizations diversify suppliers, sites, and routes, then run the whole estate on software that thousands of other firms also run. When that software fails, it fails for all of them at once. The register that prices supplier risk does not price this one.

Independent, vendor-neutral research for supply chain and technology leaders

Supply Chain Research | supplychainresearch.com | 2026

Contents

Vendor Concentration and the Risk Your Resilience Programme Does Not Price

01 Executive summary

Supply chain organizations are disciplined about dependency. They dual-source critical components, qualify alternate suppliers, map their tiers, diversify their carriers, and rehearse what happens when a plant, a port, or a supplier goes down. Then they run the entire apparatus on software that thousands of other organizations also run, hosted by a handful of infrastructure providers that almost everyone uses, protected by endpoint agents deployed on millions of machines. The diversification stops precisely where the leverage is greatest. When the shared software fails, it does not fail for one firm; it fails for every firm that shares it, on the same morning, and the diversification that was supposed to protect the operation offers no protection at all, because there is no unaffected counterparty left to absorb the volume.

This article examines that structural exposure through the incidents that demonstrated it: the Blue Yonder ransomware event of November 2024, which sent Starbucks managers back to manual scheduling and disrupted warehouse operations at UK grocers; the CrowdStrike content update of July 2024, which crashed more than eight and a half million Windows devices and cost one airline over half a billion dollars by its own account; and the file-transfer, clearinghouse, and build-pipeline compromises that preceded them. We say honestly that concentration is not simply a mistake, and that the major providers frequently raise the security floor for customers who could not match their investment. The argument is narrower and harder to dismiss: concentration converts idiosyncratic risk into correlated risk, correlated risk is what breaks markets rather than firms, and almost no enterprise risk register prices it. The remedy is not to abandon the major vendors, which is rarely possible and often unwise. It is to map the shared components, measure how long the business can run without each of them, rehearse the fallback, and contract for continuity rather than for service credits.

8.5M
Windows devices crashed by a single endpoint-agent content update in July 2024

3,000+
customers of one supply chain software vendor; the number affected by its 2024 outage was never disclosed

63%
of cloud infrastructure spending captured by three providers, up from 61 percent two years earlier

The argument in brief

- **Diversification stops at the software.** Firms dual-source components and single-source their planning suite, endpoint agent, and cloud region. The most concentrated dependency in the operation is the one the resilience programme does not examine.
- **Concentration creates correlated failure.** The claim is not that vendors fail more often. It is that shared vendors fail for everyone at once, and diversification offers no protection against simultaneous loss of the same capability across a whole market.
- **The shared component is usually below the application.** Endpoint agents, file-transfer products, clearinghouses, identity providers, and build pipelines have each become everyone's outage. A vendor list that names only enterprise applications misses them.

- **Disclosure is not resilience.** The SEC rules improve transparency and NIS2 extends security duties, but only DORA treats concentration in critical third parties as a supervised risk. Reporting an outage does not shorten it.
- **Measure recovery, not uptime.** The number that matters is how long the operation can run manually and whether that fallback has ever been rehearsed. Service credits do not fund a week of manual operation.

02 The estate that manages disruption

There is an irony at the centre of modern supply chain technology that deserves to be stated plainly before the evidence is assembled. The software estate that an enterprise builds to manage disruption, the planning systems that model scenarios, the visibility platforms that flag delays, the warehouse systems that keep the goods moving, the transport systems that reroute freight, exists precisely so that the organization can absorb shocks. That estate has itself become one of the largest concentrated shocks the organization faces, and it is the shock least likely to appear on the risk register, because the register was designed to track suppliers of things rather than suppliers of software. Figure 1 states the structure.

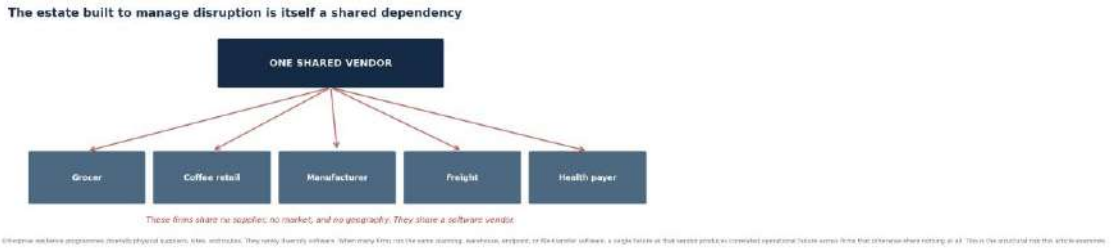

Figure 1. The estate built to manage disruption is itself a shared dependency. Firms that share no supplier, no market, and no geography nonetheless share a software vendor, and a failure there reaches all of them simultaneously.

Consider how thoroughly the discipline of supplier diversification has been absorbed everywhere else. A manufacturer that depends on a single supplier for a critical component knows it has a problem and works to qualify a second source, accepting the cost and complexity of doing so because the alternative is an unhedged dependency. A retailer that ships through a single port develops alternates. A distributor that relies on one carrier negotiates with others. The logic is well understood and the practice is mature, and the whole apparatus of supplier risk management exists to implement it. None of that apparatus is typically pointed at the software estate, where the same firm may run one planning suite, one warehouse system, one endpoint security agent, and one cloud region, with no alternate qualified, no second source, and no rehearsed fallback.

Part of the reason is categorical. Software vendors are procured as technology, through a technology function, on technology criteria, and they are managed afterward as technology relationships with service-level agreements and support tickets. They are not usually treated as suppliers in the sense that supplier risk management understands the term, which means they do not enter the supplier risk process, do not receive the diligence a component supplier receives, and do not appear in the analyses that map dependency and single points of failure. The organization has two parallel systems for managing dependency, one rigorous and applied to physical inputs, the other administrative and applied to software, and the software system is not designed to ask whether the operation could survive the vendor's failure.

Part of the reason is that the exposure was, for a long time, less acute. When enterprise software ran on the customer's own hardware in the customer's own data centre, a vendor's failure was a support problem rather than an operational one, because the software kept running whether or not the vendor did. The shift to hosted and managed services changed that fundamentally. When the planning system

runs in the vendor's environment, the vendor's outage is the customer's outage, immediately and completely, and the customer has no ability to keep the system running independently. The dependency became operational rather than commercial, and it did so gradually enough that the risk frameworks did not adapt. The estate that manages disruption came to depend on vendors whose own failures the estate cannot absorb, and the incidents in the following sections are what that looks like in practice.

There is a further reason the software estate escapes supplier scrutiny, which is that its dependencies are largely invisible to the people who would otherwise catch them. A procurement function that maps component suppliers can walk a bill of materials and see what feeds what. The equivalent map for software rarely exists, because no single function owns it: the application owners know their applications, the infrastructure team knows the platforms, the security team knows the agents, and nobody holds the composite view that would reveal which components several critical processes silently share. The dependency is not hidden by anyone; it is simply distributed across enough organizational boundaries that no one is positioned to see the whole of it, and what cannot be seen is not assessed.

The consequence is a category of exposure that surfaces only during the incident. Organizations that have lived through one of the events described later in this article report a common experience: the outage revealed dependencies that no document had recorded, as processes failed that nobody had connected to the affected component. This is the defining characteristic of unmapped concentration, and it is why the mapping exercise recommended later is worth doing before rather than after. An organization that discovers its dependency structure during a multi-day outage is discovering it at the least useful moment and at the highest possible cost, when the people who would do the analysis are occupied with the recovery.

03 Blue Yonder, November 2024

The clearest recent demonstration in the supply chain domain specifically, rather than in general enterprise technology, is the ransomware incident at Blue Yonder in November 2024. It is worth setting out in detail, because the pattern it displays recurs. Figure 2 gives the sequence.

Figure 2. Blue Yonder, November 2024. Disclosed on 21 November in the hosted managed-services environment, it reached named customers within days. The vendor has over 3,000 customers; how many were affected was never publicly disclosed.

Blue Yonder, formerly JDA Software and now a Panasonic subsidiary, is one of the largest suppliers of supply chain planning and warehouse management software in the world, with a customer base of more than three thousand organizations. On 21 November 2024 it disclosed that its hosted managed-services environment had been disrupted by ransomware. The Termite group, a ransomware operation derived from earlier Babuk code that had surfaced during 2024, subsequently claimed responsibility and listed the company on its leak site in December, asserting that it had exfiltrated a large volume of documents and data. The company worked through restoration over the following weeks.

What makes the incident instructive is not the intrusion itself, which followed a familiar pattern, but where the consequences landed. Starbucks reported that the outage affected the system its store managers use for scheduling and time tracking, and managers reverted to calculating schedules and pay manually across an extensive store estate. In the United Kingdom, the grocer Morrisons reported disruption to the warehouse management system serving its fresh produce operation across roughly five hundred stores, with service restored over about a week, and Sainsbury's also reported effects. A French manufacturer reported shipping delays. These organizations share no supplier base, no market, and no geography. They shared a software vendor, and that was sufficient to give them a common operational failure in the same week.

The most telling detail is what was never disclosed. Blue Yonder has more than three thousand customers, and the company did not publicly state how many of them were affected, for how long, or in what functions. The customers whose disruption became public did so because journalists asked them or because their own operations were visible enough that the effects could not be concealed, which means the public record captures the tip of the incident and the base remains unknown. An organization attempting to assess its own exposure to a similar event, whether at this vendor or another, has no way to calibrate from this case, because the denominator was never published. This opacity is itself part of the structural problem: concentration risk cannot be priced by the market if the scale of each realization is undisclosed, and there is no obligation on a private vendor to disclose it.

04 One faulty update, 8.5 million devices

If the Blue Yonder incident shows the exposure at the application layer, the CrowdStrike outage of July 2024 shows it at a layer beneath, and at a scale that made the structural point impossible to ignore. Figure 3 sets out the shape of it.

Figure 3. The endpoint agent as a single point of failure. A faulty sensor content update crashed more than 8.5 million Windows devices on 19 July 2024, with one airline alone attributing over half a billion dollars of cost to the event.

On 19 July 2024, a faulty content update to the CrowdStrike Falcon endpoint sensor caused more than eight and a half million Windows devices to crash worldwide. The mechanism deserves attention because it explains the blast radius. An endpoint security agent runs at a privileged level on the machine it protects, and it updates automatically and continuously, because the whole value proposition of such an agent is that it receives new threat intelligence faster than an administrator could deploy it. That design is correct from a security standpoint and it means that a defective update propagates to the entire installed base at machine speed, with no staged rollout to catch it and no administrator in the path to stop it. The property that makes the agent effective is the same property that made the failure global within hours.

The operational consequences fell hardest on organizations whose processes could not run without the affected machines. Delta Air Lines cancelled roughly seven thousand flights over five days, affecting about one and a third million passengers, and had to reset approximately forty thousand servers, many of them manually. The airline stated publicly that the disruption cost it more than five hundred million dollars, comprising lost revenue and direct costs partly offset by fuel savings, and its subsequent litigation claim is in the region of five hundred and fifty million. Logistics and freight operations experienced parallel disruption wherever the affected machines sat in the operational path.

The litigation that followed is instructive about where liability actually lands, and it is examined more fully in the section on contracts. Delta sued CrowdStrike in Georgia state court in October 2024; CrowdStrike counter-sued and publicly blamed Delta's own systems, characterising its infrastructure as antiquated and its claims as disproven. A court subsequently allowed Delta's gross-negligence and computer-trespass claims to proceed while dismissing most of the fraud allegations, and a separate passenger class action was dismissed on preemption grounds. The important point for a buyer is that the recovery of a five-hundred-million-dollar operational loss from a software vendor is a matter of multi-year litigation with a contested outcome, not a contractual entitlement, and that the vendor's first move was to argue that the customer's own environment was at fault.

The two incidents together establish the range. Blue Yonder was a supply chain application used by thousands of organizations; CrowdStrike was a security agent used by millions of machines. Neither

vendor was careless in any way that distinguishes it from its peers, and neither incident depended on an exotic failure. In one case an intrusion reached a hosted environment; in the other a routine update contained a defect. These are ordinary failure modes, and the reason they produced extraordinary consequences is entirely a function of how many organizations depended on the same component at the same moment.

05 The concentration arithmetic

Behind the individual incidents lies a market structure that has been consolidating rather than dispersing, and the arithmetic of that structure determines how large the correlated exposure is. Figure 4 shows the pattern in the layer beneath almost everything else.

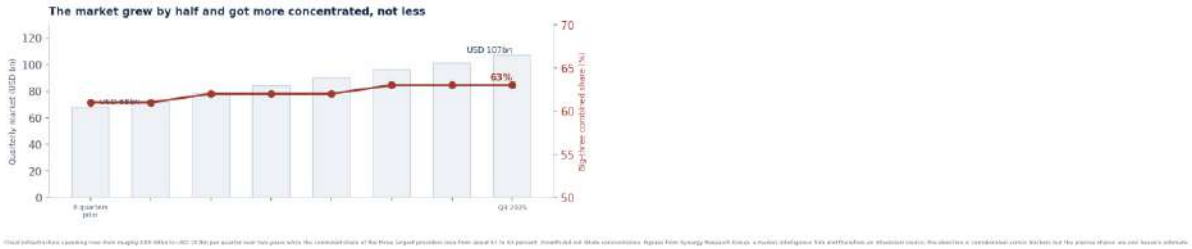

Figure 4. Cloud infrastructure spending rose from roughly USD 68 billion to USD 107 billion per quarter over two years while the combined share of the three largest providers rose from about 61 to 63 percent. Growth did not dilute concentration.

The intuition that a growing market becomes less concentrated is understandable and, in this case, incorrect. Over the two-year period shown, quarterly cloud infrastructure spending rose by more than half, from roughly sixty-eight billion dollars to roughly one hundred and seven billion, and across that expansion the three largest providers did not lose share; they gained it, moving from about sixty-one percent combined to about sixty-three. The market grew and concentrated simultaneously, which means that the substrate on which the modern software estate runs became a larger dependency in absolute terms and a slightly more concentrated one in relative terms over the same period. These figures come from a market-intelligence firm and are therefore an interested estimate rather than an audited count, and they are flagged as such, but the direction is corroborated across trackers and is not seriously disputed.

The same pattern of consolidation operates at the application layer of supply chain software, though it is harder to quantify because the market is fragmented across categories and few of the vendors are public. What can be said with confidence is that the major planning, warehouse, and transport management vendors each serve customer bases numbering in the thousands, that a handful of vendors account for a substantial share of large-enterprise deployments in each category, and that the consolidation of the sector through private equity and strategic acquisition has reduced rather than increased the number of independent suppliers over the past decade. A supply chain organization choosing a planning system today is choosing from a shorter list than its predecessor faced, and the systems on that list have more customers each.

The arithmetic that follows from this structure is what matters for risk. If a capability is supplied by twenty independent vendors each holding five percent of the market, the failure of any one of them removes five percent of the market's capability, and the remaining ninety-five percent continues to function and can, in principle, absorb some of the displaced demand. If the same capability is supplied by three vendors holding sixty-three percent between them, the failure of the largest removes a quarter or more of the market's capability at once, and there is no meaningful spare capacity to absorb it because the survivors are running their own customers. The consequence of concentration is not a

higher probability of failure but a much larger conditional loss when failure occurs, and it is the conditional loss, not the probability, that the risk register systematically fails to capture.

It is worth being precise about what this argument does not claim. It does not claim that the concentrated providers are less reliable than a dispersed field would be; the opposite is frequently true, since scale funds engineering discipline that smaller vendors cannot match. It does not claim that any given firm is more likely to suffer an outage as a result of concentration. It claims only that the losses, when they occur, arrive together rather than separately, and that a market whose participants all lose the same capability simultaneously behaves differently from one in which losses are spread. That distinction is the entire argument, and the next section develops it.

Consolidation in the supply chain software market has a further characteristic that compounds the arithmetic, which is that it has proceeded partly through acquisition rather than through competition. When a large vendor acquires a smaller one, the customers of the acquired product do not immediately change systems, so the visible diversity of products in the market overstates the diversity of the parties behind them. Two organizations running what they believe are independent systems from different suppliers may in fact be running two products owned by the same parent, hosted on the same infrastructure, supported by the same operations team, and updated through the same pipelines. The product names differ and the dependency does not, which means a concentration assessment based on product names will understate the true position.

This matters for how a buyer verifies independence. The question worth asking of any critical vendor is not only who owns it but where the product actually runs, who operates the hosting environment, which shared services it consumes, and whether the vendor's own disaster recovery depends on the same infrastructure as its production environment. Vendors will generally answer these questions when asked directly during procurement, and the answers frequently reveal that two ostensibly independent suppliers rest on the same substrate. An organization that has qualified an alternate supplier without checking this has purchased the appearance of a second source rather than the substance of one, which is precisely the failure mode that correlated risk exploits.

06 Correlated risk, not average risk

The conceptual core of this article is a distinction that risk registers routinely collapse and that finance has understood for decades: the difference between idiosyncratic risk, which diversification manages well, and correlated risk, which it does not manage at all. Figure 5 states the distinction.

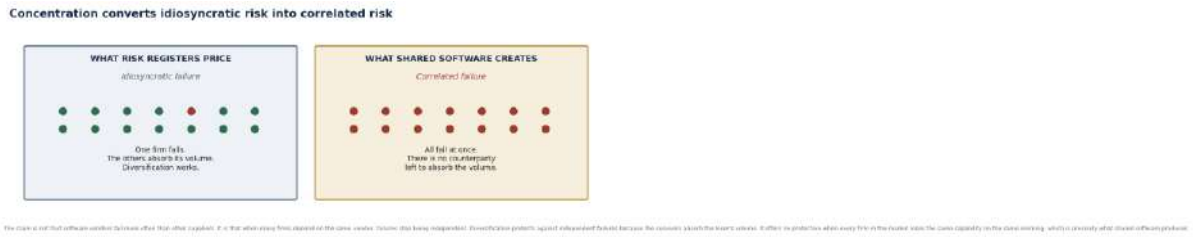

Figure 5. Concentration converts idiosyncratic risk into correlated risk. Diversification protects against independent failures because survivors absorb the loser's volume. It offers no protection when every firm loses the same capability at once.

Supplier risk management, as practised, is built for idiosyncratic failure. The premise is that a supplier may fail for reasons specific to that supplier, its finances, its plant, its management, its region, and that the correct response is to have an alternative ready so that the failure of one supplier does not stop the operation. This premise is sound and the practice built on it works. It works precisely because supplier failures are largely independent: the probability that a firm's second-source supplier fails in the same week as its primary, for unrelated reasons, is low, so holding two sources materially reduces exposure.

Shared software violates the independence assumption completely. When a firm and its alternate supplier both run the same planning system, and that system goes down, both are impaired at once, and the alternate provides no protection because the failure mode is common to both. When every freight forwarder in a lane runs the same transport management platform, an outage at that platform does not shift volume from one forwarder to another; it degrades all of them together. The diversification that the firm believes it has purchased is diversification of the entities but not of the dependencies, and the dependency is where the failure originates. A supply chain can be perfectly diversified at the level of legal counterparties and completely undiversified at the level of the software those counterparties run.

This is why the standard defence of concentration, that the major providers are more reliable than the alternatives, is true and beside the point. Suppose it is entirely correct that a large provider suffers fewer incidents per year than a smaller one would. The expected number of outages falls, which is a real benefit, and the size of the loss conditional on an outage rises sharply, because the outage now affects the whole market rather than one firm's share of it. Whether that trade is favourable depends on how much the organization cares about tail outcomes relative to average ones, and most supply chain organizations, when the question is put to them directly, care a great deal about the tail, because a simultaneous market-wide loss of capability is exactly the scenario their continuity planning is supposed to address.

The practical implication is that the risk register needs a category it usually lacks. Most registers list risks as discrete events with a probability and an impact, and they treat those events as independent when aggregating. A concentration exposure does not fit that template, because its defining property is that it correlates other entries: the same shared component failure simultaneously realises the supplier-

disruption entry, the warehouse-availability entry, the customer-service entry, and the revenue entry. Recording it as one more independent line item understates it substantially. Recording it correctly requires modelling the scenario in which several apparently separate risks all fire at once because they share a cause, which is a modelling exercise most enterprise risk functions have simply never been asked to perform for software.

The insurance market provides an instructive external check on this argument, because insurers price correlation for a living and have reached conclusions that enterprise risk registers have not. Cyber insurers have become progressively more attentive to what they term systemic or aggregation risk, the possibility that a single event triggers claims across a large portion of their book simultaneously, and they have responded with the tools available to them: sub-limits, exclusions for widespread events, and careful attention to which shared providers their insureds depend upon. An insurer asking a prospective insured which cloud regions, endpoint agents, and managed service providers it uses is asking the concentration question directly, because the insurer's own exposure depends on how correlated its insureds are with one another.

That the insurance market prices this and the typical corporate risk register does not is a useful diagnostic. Insurers cannot afford to treat correlated exposures as independent, because doing so would understate their capital requirements and eventually destroy them, so they have built the analytical apparatus to identify shared dependencies across their portfolios. A corporation faces the same analytical problem within its own operations, in that a single shared component failure can realise several apparently separate risk entries at once, but it faces no external discipline forcing it to model the correlation. The absence of that discipline, rather than any conceptual difficulty, is the reason the exposure remains unpriced, and an organization can supply the discipline for itself by adopting the insurer's question: what single failure would trigger the largest number of our entries simultaneously.

07 The shared component below the application

An organization that decides to take this exposure seriously will typically begin by listing its major software vendors, and it will typically miss most of the concentration, because the shared component that produces correlated failure is frequently not an application anybody thinks of as a supplier. Figure 6 sets out the layers.

Figure 6. Five layers at which one component became everyone's outage: the planning suite, the endpoint agent, the payments clearinghouse, the managed file-transfer product, and the software build pipeline itself.

The application layer is the visible one, and the Blue Yonder case sits there. Below it is the endpoint agent, where CrowdStrike sits, a component that most business owners have never heard of and that runs privileged code on essentially every machine in the enterprise, updating itself continuously. Below that again is the file-transfer layer, where the MOVEit compromise of 2023 illustrated the pattern: a single vulnerability in one managed file-transfer product was exploited to reach a substantial number of organizations, because that one product sat in the data-exchange path of thousands of enterprises that had no idea they shared it. Estimates of the reach run to more than two thousand five hundred organizations and tens of millions of individuals, though those tallies come from a cybersecurity vendor and the widely-quoted cost figures compound that vendor's count with another vendor's per-record estimate, so they should be treated as directional rather than measured.

Two further layers deserve attention because they are the least visible and the most consequential. The clearinghouse layer is illustrated by the Change Healthcare compromise of February 2024, which is not a supply chain case but is the clearest available demonstration of what happens when a single processing intermediary sits between an entire sector and its cash flow. The parent company disclosed costs of approximately two and nine tenths billion dollars for the year and advanced more than nine billion dollars in funding to providers who could not bill or collect while the intermediary was down. Supply chain has its own equivalents in payments, customs filing, and carrier settlement, and the lesson transfers directly.

The build-pipeline layer is illustrated by the SolarWinds compromise of 2020, in which the delivery mechanism was the vendor's own signed software update, which is the one channel that every customer is specifically instructed to trust and to apply promptly. That case also produced the clearest regulatory test of the period, when the securities regulator brought an action against the company and its chief information security officer personally; a federal court dismissed most of the claims in July 2024, including the novel theory that internal accounting controls provisions reach cybersecurity controls, and the remaining case was subsequently dismissed by stipulation with no admission of wrongdoing. The dismissal is a useful corrective for anyone inclined to assume that regulators will

convert vendor security failures into buyer remedies. For the purposes of a concentration assessment, the lesson from this layer is that the inventory must extend past the applications the business recognises to the agents, transfer products, identity providers, intermediaries, and update channels that the business has never named but on which everything else silently depends.

08 What the filings actually say

Estimating the cost of these events is where the evidence base becomes treacherous, because the field is saturated with figures produced by parties that sell remedies. This article's approach is to use what the affected organizations disclosed about themselves in regulatory filings and formal statements, and to treat everything else as commentary. Figure 7 assembles those disclosures.

Figure 7. Costs drawn from the affected companies' own filings and statements rather than third-party estimates. Using primary disclosures avoids the compounded vendor estimates that dominate this field.

The largest disclosed figure belongs to the clearinghouse case, where the parent company's quarterly filing put the total cost for the year at approximately two and nine tenths billion dollars, an estimate the company revised upward across successive filings as the scope became clear. The NotPetya cases of 2017 remain the best-documented supply chain examples: Maersk's chairman stated publicly that the company reinstalled four thousand servers, forty-five thousand personal computers, and two and a half thousand applications over ten days, and estimated damages at two hundred and fifty to three hundred million dollars, while a pharmaceutical company disclosed a substantially larger figure. Delta's claim against its security vendor is in the region of five hundred and fifty million. The freight forwarder Expeditors, whose systems were down for roughly three weeks in 2022, disclosed in its annual report approximately forty-seven million dollars of incremental demurrage charges net of recoveries plus approximately eighteen million dollars for shipment-related claims and remediation, and stated that it carried no cyber insurance.

Set against these primary disclosures are the figures that circulate in the trade press, and the contrast is instructive. A widely-quoted total cost for the file-transfer campaign is derived by taking one cybersecurity vendor's count of affected individuals and multiplying it by another vendor's average cost-per-record estimate, which produces a headline number that is the product of two interested estimates and has no independent verification. A frequently-cited loss figure for the 2022 automotive supplier shutdown appears only in vendor blog posts and cannot be traced to any statement by the manufacturer. Neither figure is necessarily wrong, but neither is evidence in the sense that a filing is evidence, and a buyer building a business case on such numbers is building on sand.

The methodological point generalises well beyond this topic and is worth stating as a rule. When a statistic quantifies the size of a problem, and the party publishing it sells the solution to that problem, the statistic requires corroboration before it can carry weight. This is not an accusation of bad faith; interested parties frequently produce careful work, and the incentive to find a large number is not the same as fabricating one. But the incentive exists, the definitions are usually the publisher's own, and the figures propagate through repetition until their provenance is invisible and a marketing estimate has

acquired the standing of a measurement. The disclosures in Figure 7 are the ones an organization can actually rely on, because the companies making them were reporting against their own accounts under a legal obligation to be accurate.

For an organization attempting to size its own exposure, the practical use of these figures is as an order-of-magnitude anchor rather than a benchmark. The disclosed range runs from tens of millions of dollars for a mid-sized freight forwarder whose systems were unavailable for weeks, to hundreds of millions for a large airline whose operations depended on the affected machines, to billions for an intermediary sitting in the cash-flow path of an entire sector. What determines position in that range is not the sophistication of the attack but the operational dependency of the affected organization: how central the failed component was, and how long the business could function without it. That is a question each organization can answer about itself in advance, and Figure 7 is chiefly useful as evidence that the answer matters enormously.

One further methodological caution applies specifically to incident cost figures, which is that the disclosed number and the true economic cost are rarely the same thing and diverge in both directions. Disclosed figures typically capture direct response costs, lost revenue in the affected period, and remediation, because those are the amounts an accounting process can attribute. They usually omit the management attention consumed, the projects displaced while the organization recovered, the customer relationships damaged, and the longer-term commercial consequences, all of which are real and none of which appear in a quarterly filing. In that direction the disclosed figure understates.

In the other direction, disclosed figures are sometimes inflated by costs the organization would have incurred anyway, particularly where an incident accelerates modernisation work that was already planned and the whole programme is charged against the event. Both effects are present in the public record and neither is usually separable from outside. The practical response is to treat the disclosed figures as a range and an order of magnitude rather than as precise measurements, which is how they are used in this article, and to recognise that the most useful figure for any given organization is not another company's disclosed loss but its own estimate of what a week without the affected capability would cost it, which it can calculate directly from its own throughput and margin.

09 NIS2, DORA, and the disclosure regime

Policymakers have responded to these events, and a buyer should understand what the resulting regimes do and, more importantly, what they do not do. Figure 8 compares the three that matter most.

Figure 8. Three regimes, only one of which prices concentration. Disclosure rules improve transparency, supply chain security duties extend obligations, and only the financial-sector regime treats third-party concentration as a supervised risk in its own right.

The securities disclosure rules adopted in 2023 and effective from December of that year require registrants to disclose material cybersecurity incidents promptly on a current report and to describe their risk management and governance annually. The rules have improved the public record substantially, and several of the figures used in this article exist because of them. What they do not do is reduce anyone's dependency. A rule that requires a company to tell the market that it has suffered a material incident is a transparency instrument, and transparency is valuable, but the disclosure arrives after the outage and does nothing to shorten it. The regime also applies to the affected registrant rather than to the vendor whose product failed, so a privately-held software vendor with thousands of customers has no equivalent obligation to disclose the scope of its own incident, which is precisely why the denominator in the Blue Yonder case remains unknown.

The European network and information security regime extends security obligations to a much wider set of entities than its predecessor and, importantly, reaches into supply chain security by requiring in-scope entities to address risks arising from their suppliers and service providers. This moves closer to the problem, because it makes an entity's own security posture partly a function of how it manages its supplier dependencies. It remains, however, primarily a security-practice regime rather than a concentration regime: it obliges entities to manage supplier risk without limiting how many entities may depend on the same supplier, which is the variable that generates correlated failure.

The financial-sector operational resilience regime is the one that addresses the structural question directly, and for that reason it is the most useful template for other sectors even though it does not apply to them. It establishes a framework for managing information and communications technology third-party risk, and it treats concentration in critical third-party providers as a supervised concern in its own right, with designated critical providers subject to oversight. The regulator, in other words, has accepted that the aggregate dependency of many regulated firms on the same handful of providers is a systemic issue that individual firms cannot solve through their own contracting, and has taken a supervisory interest in it. Supply chain has no equivalent, and it is worth asking why an interruption to the software that moves food, medicine, and industrial components is treated as a purely private matter when an equivalent dependency in payments is treated as systemic. A supply chain organization cannot wait for that question to be answered, which is why the constructive sections of this article are addressed to what a buyer can do unilaterally.

10 The contract will not make you whole

Buyers frequently assume that the commercial agreement transfers this risk to the vendor, and it is worth being precise about how little of it the agreement actually transfers. The gap between what a service-level agreement provides and what an operational outage costs is one of the widest in enterprise procurement.

A typical software service-level agreement specifies an availability target and provides service credits when the target is missed. The credits are calculated as a percentage of fees paid for the affected period, and they are almost always capped, frequently at one month of fees or a similar figure. Set that against the disclosed costs in the previous section. A freight forwarder losing tens of millions of dollars to demurrage while its systems are down does not have those losses restored by a credit against a monthly subscription, and an airline attributing half a billion dollars to an outage is not made whole by any credit regime that any vendor offers. The service credit is a partial refund of the fee, not compensation for the consequences of the failure, and the two differ by orders of magnitude.

Beyond the credit regime sit the limitation-of-liability provisions, which in most enterprise software agreements cap total liability at a multiple of fees paid, commonly twelve months, and exclude consequential, indirect, and special damages entirely. Operational losses from an outage are consequential damages almost by definition, so the exclusion removes the category of loss that actually matters to the buyer, and the cap limits whatever remains to a figure calibrated to the contract value rather than to the exposure. This structure is standard, it is not hidden, and buyers accept it routinely because it is presented as market practice and because negotiating it is difficult. But its effect is that the party best placed to prevent the failure bears little of its cost, and the party least able to prevent it bears nearly all.

What is left when the contract does not deliver is litigation, and the airline case shows what that path looks like in practice. The customer sued in October 2024; the vendor counter-sued and publicly argued that the customer's own infrastructure was responsible; a court permitted the gross-negligence claim to proceed while dismissing most of the fraud allegations; and the matter continued for years after the outage with no certain outcome. This is a well-resourced plaintiff with a large and well-documented loss, and it is still a multi-year contested proceeding whose result cannot be assumed. For a mid-sized organization with a smaller loss, the practical answer is that litigation is not a remedy at all, because the cost and duration exceed what the claim is worth. The realistic planning assumption is that the operational cost of a vendor outage stays with the buyer, and that assumption should be made explicitly at the point of purchase rather than discovered during the incident.

There is a narrow set of contractual provisions that does meaningfully improve a buyer's position, and it is worth distinguishing them from the credit regime that does not. The most valuable are the ones that shorten the outage or preserve the buyer's options during it: a firm notification obligation with a defined timeframe, so the buyer learns of an incident from the vendor rather than from the trade press; a restoration commitment with a stated target and a defined order of priority for restoring customers; data portability and escrow provisions ensuring the buyer can extract its own data in a usable form even

while the vendor's systems are impaired; and audit or attestation rights that let the buyer verify security and continuity claims rather than accepting them. None of these transfers the financial loss, and all of them reduce either the duration of the outage or the buyer's dependence on the vendor's goodwill during it.

The reason these provisions are worth pursuing even though they do not transfer risk is that duration is the variable a buyer can still influence after the failure has occurred. The financial loss from an outage is roughly the product of the operational impact per day and the number of days, and the buyer has no control over whether the incident happens but considerable interest in how long it lasts and how well it can operate meanwhile. Provisions that improve notification, restoration priority, and data access all act on duration and on the quality of the degraded operating mode, which is where the recoverable value is. Buyers who spend their limited negotiating capital arguing for a higher liability cap, which the vendor will resist strenuously and which would still not cover the loss, would frequently do better spending it on notification and portability terms, which vendors concede more readily and which measurably improve the outcome.

11 The fairness case: the cloud is safer

This article has argued hard about concentration, and the counterargument is strong enough that any serious treatment must engage it properly rather than gesture at it. A reader who concludes from the preceding sections that enterprises should retreat from major providers to smaller vendors or to their own data centres would be drawing precisely the wrong lesson.

The strongest point is that the major providers are, on the evidence, better at security and reliability than most of the alternatives. A large cloud provider employs security engineering teams larger than the entire information-technology function of most of its customers, patches continuously, operates across multiple availability zones by default, and subjects itself to audit regimes that few private data centres could pass. The counterfactual to running on a hyperscaler is frequently not a more resilient bespoke arrangement but a less resilient one, staffed by fewer people with less specialist expertise and a slower patch cycle. The Maersk case makes the point sharply and against the article's own thesis: the company was reportedly running outdated server software at the time of the 2017 incident, which is an in-house patching failure of a kind that a managed cloud service would have prevented, and the company's own stated lesson was to move toward cloud infrastructure with automatic patching and distributed backup.

A second point is that concentration also concentrates defensive investment, and that this is a real benefit rather than a rhetorical one. When thousands of organizations depend on one provider, that provider can justify security spending that no individual customer could, and the resulting protection is shared across the whole customer base. Dispersing the market across many small vendors would raise the aggregate number of poorly-defended targets even as it reduced the size of any single failure, and it is not obvious that the resulting total loss would be lower. There is a genuine trade between fewer, larger, better-defended failure domains and more, smaller, weaker ones, and reasonable analysts can disagree about where the optimum sits.

A third point is that catastrophic vendor incidents remain rare relative to the installed base. The cases in this article are drawn from a period of several years across the entire global software estate, and they are notable partly because they are unusual. An organization that reorganised its technology strategy around the assumption that its major vendors will fail regularly would be misallocating attention away from the far more common causes of operational disruption, including its own configuration errors, its own change management, and the ordinary supplier and demand shocks that supply chain risk management already addresses. The fair synthesis is that concentration is a tail risk rather than a routine one, that it is materially underpriced rather than universally catastrophic, and that the correct response is proportionate: map it, measure recovery, rehearse the fallback, and price it candidly in the register, without pretending that the alternative of dispersal is free or obviously superior.

A final point in fairness concerns the direction of travel, which is not uniformly toward greater exposure. Several developments cut the other way. Multi-region and multi-zone architectures, now standard on major platforms, protect against a broad class of infrastructure failure that would previously have been fatal. Staged rollout practices for agent updates, which the July 2024 outage prompted several vendors to strengthen, directly address the propagation mechanism that made that incident global. Portability

standards and regulatory pressure on data egress have made extraction from a hosted environment somewhat less difficult than it was. And buyers have become measurably more willing to ask concentration questions during procurement, which changes vendor behaviour over time. None of this eliminates the structural exposure, and this article's argument stands, but a reader should not conclude that the position is static or that it is deteriorating in every respect.

12 Price the correlation

The constructive principle that follows from the diagnosis is narrow and actionable: stop assessing software vendors one at a time on their individual reliability, and start assessing the estate on how much of it fails together. Correlation, not average reliability, is the property that determines the exposure, and it is the property that current practice does not measure.

In practice this means changing the unit of analysis. Vendor risk assessment as commonly practised asks, for each vendor in turn, whether that vendor is financially sound, appropriately certified, and adequately secured, and it produces a per-vendor rating. That assessment is worth doing and it does not answer the question this article poses, because a portfolio of individually excellent vendors can still contain a single component on which everything depends. The correlation question is asked of the estate rather than the vendor: which components, if they failed, would take down more than one critical process at once, and how many of the organization's supposedly independent capabilities actually rest on the same substrate. This is a mapping exercise, not a rating exercise, and it produces a fundamentally different artefact.

The mapping has to go deeper than the vendor list, for the reasons set out earlier. It should trace each critical business process to the applications that support it, then to the infrastructure and platform services those applications run on, then to the cross-cutting components that touch everything, the identity provider, the endpoint agent, the network and content-delivery layer, the file-transfer and integration middleware, and the update channels through which all of these change. The output is a dependency map in which the shared components are visible as the nodes with the highest fan-out, and those nodes are the concentration exposures, whether or not anyone previously thought of them as suppliers. In most organizations that perform this exercise for the first time, the highest fan-out nodes are surprises, because they are infrastructure rather than applications.

Once the shared components are identified, the correlation can be priced by scenario rather than by line item. Instead of asking what the probability and impact of a vendor outage are in the abstract, the organization asks what happens if this specific high-fan-out component is unavailable for a day, a week, and a month: which processes stop, which can be run manually and at what throughput, what the cost accrues to over each duration, and at what point the losses become existential rather than merely painful. That scenario is a coherent object that management can reason about and budget against, and it captures the correlation directly, because it starts from the shared cause and traces all of the simultaneous consequences rather than treating each consequence as an independent risk. The output is a small number of named scenarios with costed durations, which is a far more useful artefact than a long register of independently-scored vendor risks that collectively understate the exposure.

A useful discipline when mapping is to distinguish between components that are shared across the organization and components that are shared across the market, because they create different problems and admit different remedies. An internally shared component, such as an identity provider on which every application depends, concentrates risk within the firm: its failure stops many internal processes at once, and the remedy is architectural, whether that means a secondary path, a break-glass

procedure, or a cached credential mode. A market-shared component, such as a planning platform that the firm and its logistics partners all use, concentrates risk across counterparties, and no architectural change inside the firm addresses it, because the firm's partners fail simultaneously regardless of what the firm does internally.

The market-shared case is the harder one and it is routinely missed, because it requires looking outward at what the counterparties run rather than inward at what the firm runs. A logistics operation that has carefully arranged alternate carriers has diversified its counterparties, and if all of those carriers run the same transport management platform then the diversification does not survive an outage at that platform. Establishing this requires asking counterparties what they depend on, which is an unusual question in a commercial relationship and a reasonable one to raise during contracting. The answer determines whether the alternate arrangements the firm has paid for provide real protection or only the appearance of it, and it is worth knowing before the question is settled by events.

13 A concentration protocol, and a scoring rubric

The principles above combine into a protocol an organization can adopt, and a rubric a board or risk committee can use to judge whether the exposure is being managed or merely documented. Figure 9 sets out the discipline.

Figure 9. Owning the dependency you cannot diversify away: map the shared components, measure recovery rather than uptime, rehearse the fallback, price correlation in the register, and contract for continuity rather than credits.

The protocol runs as follows. Map the shared components below the application layer, tracing critical processes through applications to infrastructure, identity, endpoint, transfer, and update channels, and identify the nodes with the highest fan-out. For each such node, establish the recovery position: how the business operates without it, at what throughput, for how long, and at what accruing cost. Rehearse that fallback rather than documenting it, because an untested runbook is an assumption rather than a control. Price the correlation as a named scenario in the risk register rather than as a set of independent line items. And negotiate for continuity provisions, including notification, data portability, restoration commitments, and audit rights, rather than accepting a service-credit regime that cannot fund a week of manual operation.

A scoring rubric

The dimensions below distinguish an organization that manages concentration from one that has merely recorded it.

Dimension	Managed	Documented only
Unit of analysis	The estate and its shared components	One vendor at a time
Depth of map	Below the application: agent, transfer, identity	The named enterprise applications
Metric	Time the business can run without it	Contracted uptime percentage
Fallback	Rehearsed under realistic conditions	Written down, never tested
Register treatment	Named correlated scenario, costed by duration	Independent line items, aggregated wrongly
Contract	Notification, portability, restoration, audit	Service credits capped at fees
Evidence base	Primary disclosures and own testing	Vendor-published incident statistics

An organization scoring in the left column knows which components it cannot lose, how long it survives without each, and what it has actually rehearsed. An organization scoring in the right column has a vendor register, a set of contracted availability percentages, and an untested assumption that the fallback works. The rubric does not eliminate the exposure, which cannot be eliminated while the market is structured as it is. It ensures the exposure is understood and provided for rather than discovered on the morning it is realised.

14 Recovery time is the number that matters

Of everything in the protocol, one measure deserves separate treatment because it is the single most useful number an organization can hold and the one most frequently absent: how long the business can operate without each shared component, and how well. Uptime percentages, which are what contracts specify and what vendor marketing emphasises, are close to useless for this purpose.

Consider what an availability figure actually tells a buyer. A commitment of ninety-nine point nine percent availability permits roughly eight and three quarter hours of downtime per year, and a commitment of ninety-nine point ninety-five permits about four and a half. These figures sound reassuring and they say nothing about the distribution, which is the property that determines operational consequence. Eight hours of downtime distributed as a few minutes per month is an inconvenience that most operations absorb without noticing. The same eight hours arriving as a single continuous outage during a peak shipping window is a different event entirely, and the identical availability number covers both. Worse, the incidents that matter most, the multi-day restorations described earlier, breach any such commitment so comprehensively that the commitment ceases to be informative; the contract simply pays out its capped credit and the operation deals with the consequences.

The measure that carries information is the recovery position, expressed as the answer to a specific question: if this component is unavailable, what does the operation do, at what proportion of normal throughput, and for how long can that be sustained before the position becomes untenable. The Starbucks response to the Blue Yonder outage is a good illustration of a fallback that existed and functioned, in that managers reverted to calculating schedules manually, which is slower and more error-prone than the system but permits the stores to keep operating. The Morrisons response, working through a warehouse management outage across hundreds of stores for about a week, is another. In both cases the operation continued in a degraded mode, and the question worth asking in advance is exactly what that degraded mode looks like, whether the people who would have to execute it know how, and how long it holds.

There is a further reason to prefer recovery time over availability, which is that it is a number the organization can establish for itself without depending on vendor cooperation. Availability is asserted by the vendor and measured by the vendor against the vendor's own definitions of what constitutes an outage, and buyers rarely audit it. Recovery time is a property of the buyer's own operation and processes, discoverable by the buyer through its own testing, and not subject to anyone else's definitions. An organization that has rehearsed operating without its planning system for three days knows something concrete and verified about its own resilience, and that knowledge does not depend on trusting a vendor's uptime report. Where the two conflict, and they frequently do, the rehearsed recovery position is the one that will govern what actually happens.

Establishing recovery time requires a specific kind of exercise that many organizations have never run for software dependencies, though most run its equivalent for physical ones. The test is not a tabletop discussion of what people believe would happen, which reliably produces optimistic answers, but a

controlled exercise in which the component is treated as unavailable and the operation is asked to proceed without it for a defined period. What such exercises consistently reveal is that the documented fallback assumes access to information that lives inside the unavailable system, that the people expected to execute it have never done so, and that the throughput achievable manually is a fraction of what the plan assumed. These are useful discoveries and they are available only through the exercise.

The finding about information dependency is the one that most often surprises. A manual fallback for a warehouse system typically assumes that staff know what is in the warehouse and where, but that knowledge may exist only in the system that is unavailable, in which case the fallback is not a degraded mode but an impossibility until the data is reconstructed from another source. The same applies to schedules, routing, customer commitments, and inventory positions. An organization that intends to rely on manual operation during an outage needs to establish in advance where the information required to operate manually will come from, and to ensure that it is captured somewhere that survives the loss of the primary system. A daily export held outside the vendor's environment is unglamorous and is frequently the difference between a degraded operation and a stopped one.

15 Conclusion: the dependency you cannot diversify

The supply chain profession has spent two decades building sophisticated machinery for managing dependency on suppliers of physical things, and that machinery works. Firms map their tiers, qualify alternates, model disruptions, hold buffers, and rehearse responses, and the discipline has repeatedly proved its worth through shocks that would otherwise have been far more damaging. The argument of this article is that the same profession has left one dependency almost entirely outside that machinery, and that it is the dependency with the highest fan-out and the least diversification: the software estate on which everything else runs.

The incidents assembled here demonstrate the shape of the exposure rather than its frequency. A supply chain planning and warehouse vendor with more than three thousand customers is compromised, and a coffee retailer, two grocers, and a manufacturer discover in the same week that they share a dependency none of them had listed. A security agent installed on millions of machines receives a defective update, and eight and a half million devices stop, taking an airline's operation with them at a cost the airline puts above half a billion dollars. A file-transfer product used by thousands of organizations is exploited through a single vulnerability. A payments intermediary sitting between an entire sector and its cash flow goes down, and the parent company discloses costs approaching three billion dollars and advances more than nine billion to the providers it could no longer pay. None of these required an exotic failure. Each required only that many organizations depended on the same component at the same moment.

What follows for a supply chain leader is not a recommendation to leave the major providers, which is usually neither possible nor wise, and this article has set out the serious case that the concentrated providers raise the security floor rather than lowering it. What follows is a change in what gets measured and managed. The unit of analysis moves from the individual vendor to the estate and the components it shares. The metric moves from contracted uptime to rehearsed recovery time. The register entry moves from a set of independent line items to a named correlated scenario, costed by duration. The contract negotiation moves from service credits, which cannot fund a week of manual operation, to notification, portability, restoration, and audit rights, which at least shorten the outage and preserve the option of leaving. And the diligence extends below the application layer to the agents, transfer products, identity providers, and update channels that no business owner has ever thought of as suppliers and on which every business process silently depends.

The organizations that came through the incidents described here in reasonable shape were not the ones with the best contracts or the highest contracted availability. They were the ones that could still operate, in a degraded and uncomfortable mode, without the system that had failed, because someone had thought about it in advance and, in the better cases, had practised it. That capability is unglamorous, it is difficult to justify in a budget cycle when nothing has gone wrong, and it is the only thing that reliably works when the shared component fails. The software estate built to manage disruption will eventually deliver one. The question worth answering before that happens is a simple one, and most organizations cannot currently answer it: if this system were unavailable tomorrow morning, for a week, what exactly would we do.

16 Methodology, caveats, and sources

Methodology

- This article draws on regulatory filings and formal company statements, court records, regulator and legislative publications, and contemporaneous reporting, current to mid-2026. Supply Chain Research is independent and accepts no payment from the vendors, cybersecurity firms, or insurers discussed.
- Cost figures are taken wherever possible from the affected organizations' own disclosures rather than from third-party estimates, because this field is dominated by figures produced by parties that sell remedies. Where an interested estimate is used, the interest is identified.

Caveats

- The scope of the November 2024 supply chain software incident was never publicly disclosed by the vendor. The named affected customers are those identified in public reporting; the total number affected is unknown, and this article does not estimate it.
- Cloud market-share figures come from a market-intelligence firm and are an interested estimate rather than an audited count. The direction is corroborated across trackers; the precise percentages are one house's method.
- Widely-quoted aggregate cost figures for the 2023 file-transfer campaign are the product of one cybersecurity vendor's count multiplied by another vendor's per-record estimate. They are reported here only as an illustration of compounded interested estimates, not as measurements.
- A frequently-cited loss figure for the 2022 automotive supplier shutdown appears only in secondary cybersecurity commentary and cannot be traced to a statement by the manufacturer. It is therefore not used.
- Litigation described here was continuing at the time of writing. Rulings on preliminary motions are not determinations of liability, and outcomes may have changed since publication.
- Figures 1, 5, 6, and 9 are conceptual illustrations of structure rather than measured data, and are labelled as such.

Sources

[1] BleepingComputer. Blue Yonder breached by the Termite ransomware group.

[2] Cybersecurity Dive. Blue Yonder data-leak claim and customer impact.

[3] UnitedHealth Group. Quarterly report on Form 10-Q, Q3 2024 (cyberattack costs).

[4] Expeditors International. Annual report on Form 10-K, FY2022 (cyberattack costs, no cyber insurance).

[5] SolarWinds Corporation. Annual report on Form 10-K, FY2024 (loss contingency).

[6] BleepingComputer. Maersk reinstalled 45,000 PCs and 4,000 servers after NotPetya.

[7] Synergy Research Group. Cloud market share trends: the big three hold 63 percent (interested source).

[8] Harvard Law School Forum on Corporate Governance. The dismissal of the SEC action against SolarWinds and its CISO.

Additional context drawn from contemporaneous reporting on the July 2024 endpoint-agent outage and the related litigation, from the 2023 managed file-transfer campaign, and from published materials on the European network security and financial operational resilience regimes. Figures originating with parties that sell security products or market research are identified as such. This article is analysis, not security, legal, or procurement advice, and its conclusions should be validated against your own circumstances before any decision.
