RESILIENT SUPPLY CHAIN

EXECUTIVE BRIEF

Resilience without excess

How to protect service without turning resilience into permanent cost

Based on a structured review of the Resilient Supply Chain interview archive

RESILIENT SUPPLY CHAIN

Contents

© 2026 Tom Raftery. All rights reserved.

Introduction

Supply-chain resilience has become a capital-allocation question. Disruption can interrupt revenue and service, but visible forms of protection - more inventory, suppliers, capacity and facilities, tie up cash and add cost when disruption does not occur. This creates an unhelpful choice between a supply chain that is efficient but fragile and one that is protected but permanently expensive.

This brief examines an alternative: resilience as a portfolio of usable options. The relevant test is not how much redundancy the business owns, but whether it can activate dependable protection before an identified failure becomes material loss. That distinction matters because nominal alternatives often fail under stress. A second supplier may share the same upstream dependency; spare capacity may be unavailable when an entire market is disrupted; extra stock may conceal inaccurate data or a recurring process failure rather than insure against genuine uncertainty.

The analysis draws on practitioner interviews from the *Resilient Supply Chain* archive to identify the management decisions that determine whether resilience earns its cost. It does not seek a universal prescription for safety-stock days, supplier numbers or spare-capacity percentages. Instead, it asks how leaders can identify the exposures that warrant protection, verify that alternatives will work, choose between fixed and flexible capacity, remove avoidable sources of failure and govern the remaining options on a common economic basis.

Resilience and efficiency are not opposites when protection is specific, executable and reviewable. The task for leadership is to make each resilience investment answerable to a defined exposure, recovery window, activation mechanism and lifecycle cost.

© 2026 Tom Raftery. All rights reserved.

RESILIENT SUPPLY CHAIN

Executive brief

Resilience is the value of the options a business can activate before disruption becomes material loss—not the amount of redundancy it owns. Extra inventory, suppliers, capacity and routes create resilience only when they protect a defined exposure, work within the required recovery window and justify their full lifecycle cost. Otherwise, they consume cash and can conceal the failures that created the apparent need for protection.

Five decisions matter most: segment exposures by consequence; test whether backups are genuinely usable; compare permanent redundancy with flexible capacity; repair avoidable process and data failures before buffering them; and price resilience as a portfolio of options using scenarios, avoided-loss assumptions and end-to-end metrics.

The evidence exposes a central organisational tension. Finance is right to challenge assets that tie up working capital without measurable purpose; operations is right that a system engineered only for normal conditions can fail expensively under stress. Abe Eshkenazi, CEO of the Association for Supply Chain Management, describes this conflict between inventory pressure and operational protection. The resolution is to make the exposure, economics and activation conditions explicit for each critical segment. [1]

The archive does not support universal safety-stock, supplier-count or spare-capacity targets. It supports a better governance standard: approve no resilience investment without identifying the failure being insured against, the response window, the dependable protection being purchased, its full cost and the conditions for reducing, redesigning or renewing it.

Decision	The question management must be able to answer
Segment the exposure	What specific failure are we protecting against, and what would it cost?
Verify the alternative	Can the option deliver the required capacity inside the recovery window?
Choose the operating model	Is fixed redundancy necessary, or can flexible capacity provide dependable protection at lower cost?
Remove avoidable failure	Is the buffer absorbing genuine uncertainty or compensating for a recurring defect?
Price and govern the option	Does the avoided loss justify the full cost, and who owns activation, review and expiry?

© 2026 Tom Raftery. All rights reserved.

The capital-allocation problem behind resilience

Extra stock, duplicate suppliers and spare assets are not intrinsically resilient. A pallet of safety stock can prevent a shutdown or become obsolete; a second supplier can provide recovery or share the first supplier's hidden dependency; additional equipment can increase throughput or add complexity without removing the bottleneck. Their value depends on what they protect and whether the operating model can use them.

The relevant question is not “How much resilience can we afford?” but “Which exposures justify protection, what form should it take, and how will we know whether it remains economically useful?” The interviews identify the decisions and enabling conditions that make that question answerable.

© 2026 Tom Raftery. All rights reserved.

The five decisions leaders need to make

1. Segment the exposure before buying protection

The first discipline is to stop treating the supply chain as if every item, supplier, lane and customer carried the same consequence of failure. Protection should be strongest where an interruption threatens significant revenue, service, safety, market access or recovery—and deliberately lighter where substitution is easy or the consequence is limited.

This changes the buffer discussion from a portfolio-wide percentage to a consequence-based policy. Component criticality, usage, lead time and confidence in demand planning determine whether protection is warranted and how much is defensible. Andy Kohm, founder and CEO of supply-chain intelligence platform SCIP, provides the operating evidence for that distinction: exposure differs materially by component, so the policy must differ too. [2]

The same logic determines how much analytical effort to spend. The Bristlecone discussion distinguishes deep, well-governed modelling for the minority of SKUs, lanes or suppliers driving most value or risk from coarser rules for the tail. This is not simply an analytics efficiency. It prevents low-impact complexity from consuming the attention needed for genuinely consequential exposures. [3]

Segmentation must also reflect the outcome being protected. Cost, service and emissions do not carry the same weight in every market or lane; for some raw-material movements, on-time delivery is the binding priority. Constantine Komodromos, CEO of logistics emissions platform VesselBot, shows why a resilience decision needs an explicit hierarchy of outcomes rather than an average corporate score that obscures the operating requirement. That does not make sustainability irrelevant; it makes the trade-off visible. [4]

Segmentation does not eliminate judgement. Data on bills of material, supplier dependencies, demand and lane performance can reveal concentration, but leadership still has to determine which consequences are intolerable. That decision should be visible to finance and operations alike.

> **Executive decision:** Require every resilience proposal to identify the specific segment, the exposure being reduced and the business consequence of failure. Reject requests for undifferentiated inventory, capacity or supplier buffers.

© 2026 Tom Raftery. All rights reserved.

2. An alternative has value only when it can be activated

Supplier and route counts are poor proxies for resilience. A nominal backup may be commercially unavailable, technically unqualified, dependent on the same upstream producer, unable to accept sufficient volume or disconnected from the systems required to transact. The relevant measure is usable optionality: how much substitute capability can be activated, with what confidence and within what time.

Hidden dependencies make this harder than it appears. Wayne Scott, Global Governance, Risk and Compliance Lead at software escrow company Escode, describes critical supplier ecosystems containing fourth parties, proprietary components and source-code dependencies. Adding another direct supplier does not remove risk if both alternatives rely on the same invisible node. His evidence also highlights contractual readiness: organisations need rights and levers that enable them to obtain resilience controls, test exit arrangements and act when a supplier fails. [5]

Geographic diversification can be useful, particularly where climate, geopolitical or tariff exposure is concentrated. But it should be evaluated as a portfolio of direct and indirect costs. The SourceReady discussion connects alternative suppliers with lead-time, tariff and concentration information rather than unit price alone. A cheaper nominal source may remain the more expensive resilience choice once delayed and second-order exposures are included. [6]

Operational readiness is equally physical. In yard and warehouse operations, Adam Newsome, CEO of Lazer Logistics, describes backup plans that include labour, drivers, trucks and other deployable resources—not simply software or a documented intention. A rapid-response labour pool can provide surge capability only if people are available, trained and able to reach the affected operation. [7]

This suggests a better resilience inventory. Instead of counting alternatives, leaders should record qualification status, available volume, shared dependencies, contractual rights, onboarding requirements and time to activation. An alternative that cannot meet the recovery objective is not protection; it is reassurance.

> **Executive decision:** Measure dependable substitute capacity and time-to-activate. Test critical alternatives under stressed conditions and remove nominal backups that cannot meet the required recovery window.

© 2026 Tom Raftery. All rights reserved.

3. Flexibility can protect service with less fixed redundancy

Permanent spare assets are only one way to absorb disruption. Flexible labour, carrier pools, temporary warehousing, configurable automation and rapid network reallocation can provide protection without carrying the full normal-period cost of idle capacity. The trade-off is that flexibility requires orchestration and may be less certain than dedicated resources.

The choice should begin with the disruption horizon. A short seasonal peak, an unexpected carrier failure and a prolonged regional shutdown require different responses. Some justify a contracted option or temporary resource; others need physical inventory or structurally diversified supply. The archive repeatedly supports combining resilience—the ability to withstand disruption—with agility, the ability to reconfigure after conditions change. [1]

Digital platforms can increase the set of available carriers or fulfilment channels, but more options also make selection harder. Nishith Rastogi, founder and CEO of logistics technology company Locus, describes connectivity enabling faster onboarding and network reconfiguration while simultaneously increasing the decision complexity involved in choosing the right option. Flexible capacity therefore depends on the ability to evaluate and orchestrate choices, not just access to a marketplace. [8]

Physical limits remain decisive. Better sequencing can increase utilisation and defer capital expenditure, but software cannot create infinite throughput from a saturated resource—a distinction made by Keith Moore, CEO of AutoScheduler. The same warehouse discussion carries the opposite warning: adding equipment reflexively can introduce maintenance and operational complexity without resolving the constraint that triggered the investment. Together, the two mechanisms show why orchestration and physical capacity must be diagnosed separately. [9]

The comparison between fixed and flexible protection should therefore include four dimensions: activation speed, dependable capacity, normal-period cost and operating complexity. Flexible options are attractive when they can be called reliably and the disruption is shorter than the useful life of permanent capacity. Dedicated resources may be justified where consequences are severe, activation cannot be delayed or shared capacity is likely to disappear during a market-wide shock.

> **Executive decision:** Compare fixed redundancy with pooled or contracted flexibility under the same disruption scenarios. Treat orchestration capability and physical capacity as separate constraints.

© 2026 Tom Raftery. All rights reserved.

4. Buffers should absorb uncertainty, not recurring process failure

Some variation cannot be removed: weather changes, vehicles break down, suppliers fail and demand surprises occur. Buffers exist to absorb that residual uncertainty. They become wasteful when they compensate for inaccurate records, uncontrolled exceptions, weak ownership or workflows that repeatedly create the same failure.

The financial symptoms appear downstream. Forecasting errors, crisis buying and delayed decisions create excess inventory and idle assets; once those assets exist, organisations face further costs and constrained disposal choices, particularly for perishable, controlled or specialised categories. The Evergreen Trading discussion illustrates the consequence: holding the asset longer does not preserve resilience automatically and may merely postpone recognition of the loss. [10]

Operational exceptions can create the same pattern. Weak inbound setup, inventory practice and tolerated non-compliance generate recurring fulfilment problems. Additional stock or labour can keep service moving temporarily, but it also permits the faulty process to continue and creates labelling, counting and rework later. The River Plate episode shows why accountability and standardised workflows must remove avoidable variation before leaders decide what protection remains necessary. [11]

The food sector provides an unusually visible example of excess becoming the business model. Olaf van der Veen, CEO of food-waste technology company Orbisk, describes overproduction and overstocking that require operators to buy well above expected demand, converting an apparent service buffer into discarded product. Better measurement and contextual demand information can reduce that waste, although the scale of vendor-reported improvements should not be treated as an independent cross-industry result. [12]

The proper sequence is remediation first, residual protection second. This does not mean waiting for perfect data or flawless processes. It means identifying which variation is structurally unavoidable and which is being produced internally. Funding a buffer without that distinction risks paying indefinitely for a correctable cause.

> **Executive decision:** Require resilience proposals to separate unavoidable uncertainty from recurring process failure and document the remediation already attempted. Do not automate or buffer a defect without naming it.

© 2026 Tom Raftery. All rights reserved.

5. Price resilience as a portfolio of options

Once exposures are segmented and the alternatives are understood, resilience becomes an economic comparison. The company can hold inventory, reserve capacity, qualify another supplier, buy contractual options, improve a process, increase sensing or accept the risk. Each choice has a normal-period cost and a different effect under disruption.

Scenario planning makes those comparisons explicit. Robbert de Looff, Industry Commercial Lead for Chemicals at supply-chain planning company OMP, describes testing upside and downside conditions and comparing the cost of pre-stocking or external warehousing with the loss expected from a shutdown. The value lies less in predicting one future correctly than in understanding which levers are available and how the economics change across plausible conditions. [13]

The model must be concrete enough for finance. Bezrukov argues that resilience initiatives gain organisational support when they express recurring savings or revenue protection in monetary terms. Bottom-up simulation can estimate throughput and test buffers, but a technically sophisticated model is not a business case until its output connects to cash, service, margin or avoided loss. [3]

Network design illustrates the danger of evaluating a component in isolation. Richard Savoie, CEO of routing technology company Adiona, argues that transport investments should right-size the network rather than replace assets one for one. Simulation can compare vehicle types, routing and fleet size while preserving service. The broader principle is that an intervention designed at system level may achieve the objective with fewer assets than a local replacement policy. [14]

The portfolio should also recognise externalities and useful life. John Beath, founder and CEO of sustainability consultancy JBE, cautions that analyses can omit cost or optimise a visible attribute while shortening product life and worsening the total outcome. Resilience decisions need the same discipline: the protected service outcome, working-capital effect, operating cost, sustainability consequence and useful life of the intervention belong in one comparison. [15]

No model eliminates judgement. Probability estimates will be uncertain and some consequences—safety, regulatory access or strategic customer failure—may not be reducible to an average expected value. But scenario ranges are still more accountable than unexplained buffer targets. They also make review possible: if exposure, demand, lead time or activation cost changes, the portfolio can change with it.

> **Executive decision:** Approve resilience investments against explicit scenarios, avoided-loss assumptions and full lifecycle costs. Set renewal, reduction or redesign triggers when the protection is approved.

© 2026 Tom Raftery. All rights reserved.

What leaders should do differently

1. **Define and segment the intolerable failures.** Identify the products, customers, suppliers, lanes and regulatory obligations whose interruption would create material harm, then quantify criticality, time-to-impact, recovery time, demand uncertainty, lead time and concentration by segment.
2. **Create an option register.** Record inventory buffers, alternate suppliers, carriers, sites, flexible labour, temporary facilities and process interventions together with activation conditions and dependable capacity.
3. **Test executability.** Verify contracts, technical qualification, shared dependencies, system connectivity, onboarding effort and time-to-activate.
4. **Remove avoidable variation.** Correct data, ownership, workflow and KPI failures before treating their symptoms as structural uncertainty.
5. **Compare options on one economic basis.** Use common disruption scenarios and include working capital, operating cost, maintenance, disposal, sustainability effects, coordination overhead and expected avoided loss. Measure end-to-end service and recovery—not local utilisation alone.
6. **Assign governance at approval.** Name the owner, activation threshold, review date and expiry condition so protection can be deployed quickly and temporary crisis measures do not become permanent cost by inertia.

© 2026 Tom Raftery. All rights reserved.

Where the evidence is less conclusive

The archive provides a consistent decision logic but not a universal calibration. It does not establish generally applicable safety-stock days, capacity-headroom percentages, supplier counts or activation-time thresholds. Those values depend on economics and operating conditions not standardised across the interviews.

Nor does the material provide controlled comparisons of redundancy versus process remediation. Many accounts explain plausible mechanisms and practical examples, but they do not isolate causal effects or calculate marginal returns across comparable companies. Regionalisation is similarly conditional: it may reduce distance or concentration while increasing cost, emissions or exposure to another regional dependency.

Public interviews may underrepresent failed investments, internal disagreement and protection that was quietly removed after proving unnecessary. The archive also lacks a verified organisation-type classification for every guest, so this brief does not claim systematic differences between technology vendors, advisers and operating companies.

The most important research gap is quantified marginal value: how much additional service or avoided loss is produced by the next unit of inventory, capacity or supplier optionality in a defined segment? Companies should answer that question through their own scenario models, disruption reviews and controlled operating experiments rather than import a generic benchmark.

© 2026 Tom Raftery. All rights reserved.

About this research

This brief draws on a preregistered, purposive semantic retrieval from the *Resilient Supply Chain* podcast archive. Before retrieval, the research question, archive snapshot, inclusion and exclusion rules, analytical lens, challenge requirements, 50-episode ceiling, two-section-per-episode limit and 0.40 minimum similarity threshold were recorded and checksummed. No publication-date restriction was imposed.

The resulting evidence set contained 41 interviews and 79 transcript sections; the 50-episode ceiling did not truncate the result. No publication-date restriction was imposed: the selected interviews were published between 3 November 2025 and 24 August 2026. All 41 episodes were reviewed; 38 contributed 147 validated claims, while three yielded no validated claim. Every claim was assigned to an eight-theme canonical evidence pack before editorial consolidation into the five findings used here. Fifteen interviews are cited directly in this concise publication.

The interviews provide informed practitioner narratives, not a representative industry survey. Retrieval rank, support counts and co-occurrence do not establish prevalence or causation. This brief uses paraphrases. Any direct quotation intended for publication requires a separate human check against the transcript and preferably the audio.

© 2026 Tom Raftery. All rights reserved.

Source interviews

The interviews below are those cited directly in this brief; the wider canonical analysis included additional selected episodes.

7. Abe Eshkenazi, CEO of the Association for Supply Chain Management, *Why Supply Chain Agility Is Now a Competitive Advantage*.
8. Andy Kohm, founder and CEO of SCIP, *Bad Data Is Slowing Supply Chain Decisions*.
9. Simon Bezrukov, Chief AI Officer at Bristlecone, *AI in Supply Chain: Automation Is Not Autonomy*.
10. Constantine Komodromos, CEO of VesselBot, *The Logistics Blind Spot Hurting Cost, Service, and Emissions*.
11. Wayne Scott, Global Governance, Risk and Compliance Lead at Escode, *When Critical Software Becomes a Supply Chain Risk*.
12. Ricky Ho, founder and CEO of SourceReady, *AI Can Find Suppliers. Humans Still Own the Risk*.
13. Adam Newsome, CEO of Lazer Logistics; Blaine Dirker, CTO of Lazer Logistics; and Pini Usha, CEO of Buffers AI, *Why Yard Automation Is Harder Than Autonomous Trucking*.
14. Nishith Rastogi, founder and CEO of Locus, *Miss the Delivery Slot, Lose the Customer*.
15. Keith Moore, CEO of AutoScheduler; Mor Peretz, CEO of CaPow; and Gonzalo Benedit, CRO of Aera Technology, *Why More Robots Don't Always Fix Warehouse Performance*.
16. Gordon Zellner, CEO of Evergreen Trading, *Surplus Inventory and Idle Assets Are Draining Margin*.
17. Leo Rodriguez, Vice President at River Plate, *Why Small Fulfilment Exceptions Become Big Supply Chain Problems*.
18. Olaf van der Veen, CEO of Orbisk, *Buying 120% to Serve 100%: The Supply Chain Cost of Food Waste*.
19. Robbert de Looff, Industry Commercial Lead for Chemicals at OMP, *Supply Chain Resilience Fails When Decisions Move Too Slowly*.
20. Richard Savoie, CEO of Adiona, *The Maths of Supply Chain Resilience: What Happens When Algorithms Drive the Last Mile*.
21. John Beath, founder and CEO of JBE, *Measure First: Stop Spending on the Wrong Carbon Fixes*.

© 2026 Tom Raftery. All rights reserved.

RESILIENT SUPPLY CHAIN

About Tom Raftery

Tom Raftery is a sustainability and technology executive, analyst and broadcaster whose work focuses on climate, energy, digital transformation and resilient supply chains. He hosts the **Resilient Supply Chain** and **Climate Confident** podcasts, where he interviews business leaders, policymakers and innovators about the practical changes required to decarbonise industry and build more adaptable operating systems.

Previously a Global Vice President at SAP, Tom worked on strategy at the intersection of sustainability, energy and emerging technology. He has also spent more than a decade as an independent industry analyst and futurist, including seven years leading sustainability research at RedMonk and work with Gerd Leonhard's Futures Agency.

Earlier in his career, Tom co-founded an Irish software-development company, a social-media consultancy and Cork Internet eXchange, an energy-efficient data centre. That combination of technical, entrepreneurial and advisory experience shapes his approach: he connects new technology to the commercial, organisational and environmental realities that determine whether it creates value.

Tom is an international keynote speaker, a guest lecturer at Instituto Internacional San Telmo and an advisor to startups working in climate and supply-chain technology. Through Podcast Archive AI, he is also developing an evidence-led publishing practice that turns years of expert interviews into traceable, decision-focused research for senior business leaders.

© 2026 Tom Raftery. All rights reserved.