

BEYOND BRITTLE EFFICIENCY

Supply Chain Efficiency Without Resilience Is Optimised Fragility

Why the shortages of 2021 demand designed options, not simply more stock

Giovanni Leonardi

September 2021



Contents

1. Executive Summary	3
2. The evidence visible in 2021	4
3. How efficiency became fragility	5
3.1 Inventory was treated as waste without asking what service it buys	5
3.2 Purchase price displaced total exposure	5
3.3 Supplier reduction simplified management and concentrated dependency	5
3.4 Forecast accuracy became a substitute for adaptability	5
3.5 Lean was confused with removing every margin	6
4. The serious case for returning to efficiency	6
5. A better decision model: consequence, substitutability and time	6
5.1 What is the consequence of interruption?	7
5.2 How substitutable is the requirement?	7
5.3 How long does recovery take?	7
6. The recommended resilience architecture	8
6.1 Create an exposure map, not a supplier directory	8
6.2 Design options before they are needed	8
6.3 Use buffers deliberately	9
6.4 Change planning from a single forecast to bounded choices	9
6.5 Align incentives across functions	10
6.6 Govern resilience as a design principle	10
7. Data is necessary, but visibility is not resilience	11
8. Implementation without another transformation programme	11
8.1 In the first 30 days	11
8.2 By 60 days	11
8.3 By 90 days	12
9. What good looks like	12

Supply Chain Efficiency Without Resilience Is Optimised Fragility

Transformation · No. 322

Originally written September 2021 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“Resilience is not spare capacity everywhere; it is preserved choice where failure would otherwise become irreversible.”

1. Executive Summary

The supply disruptions of 2021 have exposed a mistake that has been building for years. Many organisations believed they had designed efficient supply chains when they had, in fact, designed supply chains that performed efficiently only while a narrow set of conditions held.

Low inventories, concentrated suppliers, long international routes, tightly scheduled transport and aggressive working-capital targets all reduced visible cost. They also removed options. When factories closed intermittently, containers accumulated in the wrong places, ports became congested, semiconductor availability tightened and border processes changed, the problem was not simply that individual suppliers failed. The whole system lacked room to absorb variation.

The immediate reactions have been uneven. Some organisations have placed duplicate orders, inflated forecasts and accumulated stock without knowing which constraints matter. Others have treated the disruption as temporary and protected the same efficiency measures that created their exposure. Both responses are expensive. The first converts uncertainty into uncontrolled inventory; the second leaves the operating model dependent on a rapid return to conditions that may not return soon.

This paper recommends a different position: resilience should be designed as a portfolio of options, allocated according to the consequence and recoverability of failure. That means:

- segmenting products and components by criticality, substitutability and recovery time;
- measuring end-to-end exposure rather than the performance of direct suppliers alone;
- holding buffers deliberately where they buy decision time;
- qualifying alternatives before disruption, including suppliers, specifications, routes and operating procedures;

- changing incentives so procurement, finance and operations share the cost of fragility;
- funding data quality and supplier visibility as operating capabilities, not one-off projects;
- governing resilience through explicit risk appetite and regular stress tests.

The recommendation is not to abandon efficiency or to retreat indiscriminately from global supply. A resilient network can remain lean. But it cannot be brittle by design and still be called efficient.

Resilience is not spare capacity everywhere; it is preserved choice where failure would otherwise become irreversible.

2. The evidence visible in 2021

In early 2021, a manufacturer in a composite industrial supply chain reported 96 per cent on-time delivery from its principal component suppliers. Its procurement scorecard was green. Six months later, final-product output had fallen by 28 per cent.

The apparent contradiction was easy to explain once the chain was viewed end to end. The direct suppliers were shipping almost every order they confirmed. But they were confirming fewer orders because a specialist electronic component had become scarce. That component represented less than 1 per cent of the product's purchase cost. It was present in 62 per cent of the finished range and had no qualified substitute. The manufacturer held nine days of it. Requalification of an alternative design would take at least 22 weeks.

The scorecard had measured supplier promise performance. It had not measured dependency, commonality, time to recover or the consequence of shortage. The organisation had excellent data about transactions and poor knowledge of its supply system.

This pattern recurs across sectors. Several conditions have converged:

- Demand has shifted abruptly. Forecasts built on stable channels and familiar seasonal patterns have been overwhelmed by rapid changes in what customers buy, where they buy it and how quickly they expect fulfilment.
- Production has been interrupted unevenly. Local restrictions and workforce absence have affected plants and regions at different times, creating stop-start availability rather than one clean interruption.
- Transport has lost synchronisation. Containers, vehicles and labour have not always been available where the next movement requires them. Lead times have lengthened and become more variable.
- Concentration has become visible. Specialist production often sits behind a larger number of direct suppliers. Separate vendors may depend on the same sub-tier factory, material or transport route.
- Administrative friction has increased. Changes in border, documentation and customs arrangements have added delay and uncertainty to routes previously

managed as routine.

- Forecast distortion has amplified shortage. Buyers facing uncertain supply have ordered earlier or in greater volume. Suppliers cannot easily distinguish real demand from protective ordering, so allocation and lead-time signals become less reliable.

None of these forces alone proves that globally distributed or tightly managed supply is mistaken. Together, they demonstrate that a system designed around average conditions can fail when variability becomes correlated.

3. How efficiency became fragility

The language of efficiency has hidden several different management choices.

3.1 Inventory was treated as waste without asking what service it buys

Inventory has a cost: capital, storage, obsolescence, handling and loss. Reducing it can expose process problems and improve flow. But an inventory target applied without reference to risk turns a useful principle into arithmetic detached from operations.

Nine days of a critical component and nine days of standard packaging are not equivalent. One may have a 22-week recovery time and no substitute; the other may be available from several local sources within days. A single days-of-inventory target ignores the difference.

The more important question is not “How much stock do we hold?” It is “How much decision time does this stock buy against the credible disruption?”

3.2 Purchase price displaced total exposure

Procurement teams have often been rewarded for unit-cost reduction and contract compliance. A distant or highly concentrated source can perform well against both measures. The additional exposure—long transport routes, minimum order quantities, currency, border delay, recovery time and sub-tier concentration—sits elsewhere in the organisation or remains unpriced.

A component that saves £0.40 per unit but can stop a £4,000 finished product has a risk profile that the purchase-price variance does not reveal.

3.3 Supplier reduction simplified management and concentrated dependency

Reducing the supplier base can strengthen relationships, improve quality and increase negotiating power. It can also create common-mode failure. The risk is not simply having one named supplier. It is several apparently independent suppliers sharing the same vulnerable factory, tool, material, port or skilled workforce.

Organisations frequently know their direct vendors and still cannot answer a more important question: which single events could interrupt several product lines at once?

3.4 Forecast accuracy became a substitute for adaptability

Forecasting remains necessary. Yet there is a limit to the precision available when demand patterns, transport times and supply allocations are changing together.

The practical danger is that organisations spend more effort disputing the forecast than reducing the consequence of being wrong. A forecast is an estimate; resilience is the ability to continue making acceptable decisions after the estimate fails.

3.5 Lean was confused with removing every margin

Properly understood, lean management improves flow, exposes problems and removes activity that does not create value. It does not require ignorance of risk. A network with shorter replenishment, stable processes and close supplier collaboration may carry less inventory and be more resilient than a poorly understood network with large stocks.

The failure was not lean thinking. It was using low inventory and high utilisation as visible proxies for efficiency while neglecting lead-time variation, recoverability and customer consequence.

4. The serious case for returning to efficiency

The strongest objection to the resilience agenda is financial and practical.

Buffers cost money. Dual sourcing can dilute volume leverage and increase qualification effort. Local supply may cost more and lack scale. Redesigning products for substitution consumes engineering capacity. Mapping lower tiers is difficult when suppliers consider their sources commercially sensitive. Maintaining unused capacity is expensive. If every business function protects itself against every conceivable disruption, the organisation will carry duplicated stock, underused assets and confused priorities.

There is also a danger that decisions made under the pressure of 2021 will outlive the conditions that justified them. Orders placed defensively can arrive together after demand softens. Warehouses can fill with the wrong material. Temporary premiums can become permanent contracts. A broad instruction to “increase resilience” can become a licence to spend without clear risk reduction.

This objection should be accepted, not dismissed. Resilience without economic discipline becomes precaution without limit.

The answer is not a choice between maximum efficiency and maximum protection. It is segmentation. The organisation should spend on resilience where the loss of supply would be severe, alternatives are limited and recovery is slow. It should continue to optimise cost where substitution is easy and interruption is tolerable.

The failure of the old model was indiscriminate efficiency. The failure of a reactive model would be indiscriminate redundancy.

5. A better decision model: consequence, substitutability and time

Supply decisions should be made against three questions.

5.1 What is the consequence of interruption?

Assess the operational and customer impact, not merely the value of the item. Consequences may include:

- complete stoppage of a product or service;
- safety or regulatory exposure;
- breach of customer commitments;
- damage to a strategically important relationship;
- loss of revenue or contribution;
- costly rework, expiry or restart;
- reputational harm.

A low-cost item can have a very high interruption consequence.

5.2 How substitutable is the requirement?

Substitution may exist at several levels:

- another approved supplier for the same item;
- an equivalent specification;
- a product design that accepts alternative components;
- another manufacturing site or process;
- another transport route;
- an alternative service proposition for the customer.

A supplier list alone understates these options. Engineering, quality, operations, commercial and customer teams must participate in defining them.

5.3 How long does recovery take?

Time to recover includes more than supplier lead time. It may require tooling, testing, regulatory approval, customer acceptance, system changes, transport capacity and the restart of upstream operations.

Compare time to recover with time to survive. If a critical component has 15 days of coverage and a credible recovery time of 20 weeks, the gap is not an inventory issue alone. It is a design exposure.

Segment	Consequence	Substitutability	Recovery gap	Appropriate posture
Protected constraint	High	Low	Long	Multiple options, strategic buffer, executive oversight

Segment	Consequence	Substitutability	Recovery gap	Appropriate posture
Managed dependency	High	Moderate	Manageable	Qualified alternative, trigger-based buffer, frequent review
Flexible flow	Moderate	High	Short	Responsive sourcing, modest buffer, routine monitoring
Commodity flow	Low	High	Short	Cost and service optimisation

This segmentation prevents every item from receiving the same expensive treatment. It also makes the trade-off visible to leaders.

6. The recommended resilience architecture

A credible response requires six connected changes. None is sufficient alone.

6.1 Create an exposure map, not a supplier directory

Begin with the products and services whose interruption matters most. Trace the critical components, materials, sites, tools, routes and sub-tier dependencies that enable them.

The objective is not to map every transaction in equal detail. It is to identify concentrations and recovery gaps that could defeat the operating plan.

For each protected constraint, record:

- the finished products or services affected;
- direct and known sub-tier sources;
- geography, site and shared infrastructure;
- current coverage and lead-time variability;
- qualified alternatives and their available capacity;
- time to recover and time to survive;
- contractual rights and practical access to information;
- accountable business owner.

This map must be owned as operational data. A one-off consultancy exercise that decays after three months offers false reassurance.

6.2 Design options before they are needed

An alternative supplier that has not completed technical, quality and commercial qualification is not an option. It is a name on a slide.

Options may include dual sourcing, pre-approved substitute components, portable tooling, alternate routes, postponement of final configuration, standardised product designs, reserved production slots or agreed customer substitutions.

Each option has a carrying cost. That cost should be compared with the avoided consequence and the time it buys.

For the composite manufacturer, carrying an additional 21 days of the specialist component would require £780,000 of working capital. Qualifying a second design would cost £420,000 and take six months. A 10-day plant stoppage would place roughly £14 million of output at risk and create a customer backlog taking a further month to clear.

The figures do not prove that every component deserves more stock and a second design. They make the decision economically intelligible. For this constraint, the annual carrying cost and qualification expense are not overhead; they purchase recoverability.

6.3 Use buffers deliberately

Buffers include more than inventory. They can be time, capacity, information, cash, flexible labour, transport reservations or product simplification.

The choice of buffer should match the failure mechanism. Inventory helps against a short interruption but may not solve a prolonged structural shortage. Spare capacity helps only if materials and skills are available. Extra suppliers help only if they do not share the same constraint. More information helps only if someone has authority to act on it.

Every strategic buffer should have:

- a stated risk it addresses;
- a trigger for use or replenishment;
- an owner;
- a cost;
- a review date;
- a condition under which it can be reduced.

Without these, buffers accumulate politically and are cut arithmetically.

6.4 Change planning from a single forecast to bounded choices

When uncertainty is high, planning should use ranges and scenarios rather than one agreed number presented with false precision.

A practical monthly review might examine:

1. a central demand and supply case;
2. an upside case that tests scarce components and capacity;
3. a downside case that tests inventory and cash exposure;
4. a disruption case that removes one critical node or route.

The purpose is not to predict which scenario will occur. It is to agree decisions in advance: which orders to defer, which customer segments to protect, when to activate an alternative, how much premium transport is acceptable and who can authorise a product substitution.

This converts planning from forecast defence into option management.

6.5 Align incentives across functions

Resilience cannot be owned by supply-chain management alone.

If procurement is rewarded for purchase-price reduction, finance for working-capital reduction, operations for utilisation and sales for revenue without allocation discipline, each function can optimise its measure while the network becomes more fragile.

Shared measures should include:

- service delivered during disruption;
- exposure without qualified alternatives;
- recovery gap for protected constraints;
- forecast bias and amplification, not only accuracy;
- inventory quality and critical coverage, not only total value;
- time to make and execute allocation decisions;
- concentration at site, sub-tier and route level.

The executive team must set the risk appetite. Otherwise, each resilience investment appears inefficient in the function that pays for it while its benefit sits somewhere else.

6.6 Govern resilience as a design principle

Resilience should enter product design, sourcing, capital approval, merger integration, supplier management and network change.

A new product that depends on one specialist component with a long requalification cycle should carry that exposure into its business case. A sourcing decision that creates concentration should state the risk and mitigation. A system change should preserve the data needed to see critical dependencies. A contract should provide practical routes for continuity, information and transition.

Periodic stress tests should ask concrete questions:

- What happens if this site is unavailable for four weeks?
- Which products stop if this sub-tier component is allocated at 60 per cent?
- How quickly can the alternative route carry volume?
- Which customer commitments take priority, and who decides?
- What manual processes fail if key staff remain unavailable?
- How long before cash, storage or obsolescence becomes the limiting constraint?

The value of the exercise lies in decisions made before the event, not the production of another risk document.

7. Data is necessary, but visibility is not resilience

The disruptions of 2021 have accelerated interest in control towers, supplier platforms, predictive analytics and end-to-end visibility. Better data can improve detection and coordination. It cannot substitute for options.

A dashboard may show that a shipment is delayed. It does not create an alternative component, route or capacity. A model may identify a sub-tier dependency. It does not qualify another source. Technology increases the value of a decision system only when the organisation has authority, resources and pre-designed responses.

Three data priorities matter more than an ambitious promise of total visibility:

- Identity: a consistent way to connect products, components, suppliers, sites and orders.
- Criticality: the consequence and recovery characteristics attached to those relationships.
- Timeliness: events and exceptions arriving soon enough for an available option to be exercised.

Data quality should be improved first around protected constraints. Attempting to cleanse and integrate the entire supply estate before acting will delay the decisions that the data is intended to support.

8. Implementation without another transformation programme

Organisations do not need to wait for a multi-year system replacement. A 90-day mobilisation can establish the management core.

8.1 In the first 30 days

Select a limited set of products or services representing the greatest customer, safety or financial consequence. Identify their likely protected constraints. Establish one executive owner and a cross-functional team spanning supply, procurement, operations, finance, engineering, commercial, risk and data.

Use existing records, supplier conversations and operational knowledge. Mark uncertainty openly. An incomplete exposure map with confidence levels is more useful than a polished directory that conceals unknowns.

8.2 By 60 days

Calculate time to survive and credible time to recover for the principal constraints. Define the most plausible disruption scenarios. Cost the available options: inventory,

qualification, route, specification, capacity and customer response.

Agree decision triggers and interim controls. Where exposure is acute, take proportionate action before the full analysis is complete.

8.3 By 90 days

Approve the resilience posture for each protected constraint. Fund the chosen options. Assign data ownership and review cadence. Add resilience criteria to relevant sourcing, design and investment decisions. Run one end-to-end stress test and capture the decisions it exposes.

The work should then move into ordinary operating governance. If it remains a special project owned by a temporary team, the insight will decay as the immediate shortage eases.

9. What good looks like

A resilient supply chain is not one that never suffers interruption. That promise is neither credible nor affordable.

It is one that can see its most important exposures, absorb limited variation, choose among prepared responses and recover before the consequence becomes unacceptable. It knows where efficiency is appropriate and where efficiency without options is a false economy.

The recommendation is therefore precise: retain the discipline of flow, cost and working capital, but subordinate it to a segmented view of consequence, substitutability and recovery time. Build options around the few constraints that can stop the enterprise. Measure those options as part of performance, not as insurance that disappears whenever the quarterly numbers tighten.

The events of 2021 have made fragility visible. Visibility will be wasted if organisations answer only with more stock or better dashboards.

The enduring shift is from optimising transactions to designing choices. That is how resilience becomes an operating principle rather than a reaction to the latest shortage.