

How Much Resilience Is Worth Paying For?

Resilience can protect revenue, service, customers, and strategic options - but redundancy, inventory, excess capacity, alternate suppliers, and recovery capability all carry a price. The executive task is not to maximize resilience. It is to decide which risks the enterprise should pay to reduce and which it should deliberately carry.

Executive Operations Insights | Saša Cibula

Resilience is easy to support in principle and difficult to fund in practice. A second supplier reduces dependency but may charge more. Additional inventory protects service but ties up cash. Spare capacity creates recovery options but depresses utilization. Backup systems, alternate transportation, cross-trained labor, and redundant infrastructure all improve the ability to absorb disruption, yet every one of them competes with growth, productivity, and margin.

That tension has become harder to ignore. Trade fragmentation, geopolitical risk, cyber exposure, climate events, supplier concentration, infrastructure constraints, and volatile demand have made operational disruption a recurring executive concern. At the same time, structural cost pressure makes it unrealistic to build redundancy everywhere. BCG argued in 2026 that leaders need rigorous analysis to determine when the extra cost of resilience is worth the 'insurance policy' it provides rather than assuming that reshoring, duplication, or additional inventory is automatically the right answer.

The executive question, therefore, is not whether resilience is valuable. It is **how much resilience is worth paying for, where, and against which failure modes**. That changes the discussion from continuity planning to risk economics.

Resilience Is a Risk-Appetite Decision

COSO's Enterprise Risk Management framework provides a useful foundation because it treats risk as inseparable from strategy and performance. Its risk-appetite guidance makes the central point explicit: senior leadership and the board should understand how much risk the organization is willing to accept while pursuing its objectives. Risk is not something to eliminate after strategy is chosen; it is part of choosing and executing the strategy.

Applied to operations, that means resilience should begin with the objective at risk. A network designed to guarantee life-critical product availability should carry a different tolerance for disruption than one serving a highly substitutable, low-margin product. A plant that is the sole source of a strategic technology deserves a different protection level than a facility whose output can be shifted across the network in days.

The sequence matters. Leadership first identifies the gross exposure: what can fail, how the failure reaches the enterprise, and what the consequences could be. It then compares that exposure with its appetite and tolerance for loss. Only after that should management choose treatments - avoidance, reduction, sharing, transfer, redundancy, inventory, alternative capacity, or deliberate acceptance - and evaluate whether the residual exposure is now inside acceptable bounds.

RISK APPETITE: TURNING RESILIENCE INTO AN ECONOMIC DECISION

Executive adaptation of COSO enterprise risk management and risk-appetite logic




Figure 1. COSO risk-appetite logic applied to operational resilience. Executive adaptation based on COSO ERM and risk-appetite guidance.

This logic avoids two common extremes. One is efficiency absolutism: eliminate every buffer because it looks expensive in a stable period. The other is resilience absolutism: duplicate everything because disruption is possible. Neither is a strategy. The first silently accepts risks that may be existential; the second can permanently burden the cost structure.

The Cost of Protection Is Easier to See Than the Cost of Failure

Resilience investments face a structural disadvantage in capital discussions: their costs are visible and recurring, while the losses they prevent are uncertain. A second source may add a measurable premium every quarter. Safety stock increases working capital immediately. A backup facility may sit underutilized for years. If the disruption never occurs, the protection can look wasteful in hindsight.

The avoided loss is harder to price because disruption is rarely limited to the direct cost of stopped production. A critical failure can create lost contribution margin, premium freight, contractual penalties, customer defections, emergency procurement, overtime, quality risk, reputational damage, management distraction, and delayed strategic initiatives. Recovery time can matter as much as event probability: a rare event that can be restored in hours is fundamentally different from a rare event that disables a business for three months.

This is why expected-loss calculations are useful but insufficient. Multiplying probability by financial impact can create a false sense of precision, particularly for low-frequency, high-severity events where the probability estimate itself is weak. Executives should also consider **time to recover, strategic irreversibility, concentration, and optionality**. A disruption that permanently shifts a key customer to a competitor or compromises a regulated obligation may deserve more protection than its modeled annual probability suggests.

Harvard Business Review has highlighted the same financial difficulty in supply-chain resilience: companies often struggle to justify resilience programs because future disruption cash flows are uncertain, even when the strategic case is real. The answer is not to abandon financial discipline. It is to recognize that the economics include the value of preserving choices and avoiding losses that conventional project models may understate.

Not Every Exposure Deserves the Same Protection

The strongest resilience portfolios are selective. They spend disproportionately where failure would materially damage the enterprise and avoid paying premium protection for exposures that are cheap or fast to recover from. This is where Peter Kraljic's classic purchasing portfolio remains useful as a supporting lens. Kraljic distinguished purchases by their strategic importance and the complexity or risk of the supply market, arguing that high-impact, high-risk items require fundamentally different management than routine inputs.

The same principle extends beyond procurement. Operations leaders should segment resilience needs rather than apply one policy across every product, asset, site, supplier, system, and customer. A critical control system with a twelve-month replacement lead time should not receive the same treatment as a commodity component available from dozens of suppliers. A network node that serves one geography may require less redundancy than the only facility capable of producing a proprietary product.

Four questions sharpen that segmentation. **How consequential is the failure? How quickly can the operation recover? How concentrated is the exposure? What options exist once disruption begins?** The more severe, slow-recovering, concentrated, and option-poor the exposure is, the stronger the case for paying in advance to change the risk profile.

This also helps executives distinguish resilience from simple redundancy. Redundancy is one way to create resilience, but not always the best one. Faster detection, better supplier visibility, cross-qualification, modular equipment, flexible labor, transferable tooling, pre-negotiated capacity, alternate routes, and decision protocols can sometimes reduce recovery time without carrying the full recurring cost of duplication.

A Practical Example: The Supplier That Looks Cheap Until It Stops

Consider an illustrative manufacturer with a specialized component sourced from one supplier. The item represents only a small percentage of product cost, so procurement has historically focused on price. The supplier has performed well for years. A second qualified source would add approximately 7 percent to component cost, require engineering work, and complicate planning. Carrying additional inventory would consume several million dollars of working capital. On a conventional annual cost comparison, the single-source model wins easily.

The enterprise view is different. The component is required across the company's highest-margin product family, and no practical substitute exists. A supplier failure could stop several plants within three weeks. Qualification of a replacement source after a disruption would take months, not days. The apparent purchasing decision is therefore carrying a concentrated enterprise exposure far larger than the component's unit cost suggests.

Management now has several choices. It can continue single sourcing and deliberately accept the risk. It can qualify a second supplier but give that supplier only enough volume to keep the source viable. It can hold targeted inventory to extend the recovery window. It can redesign the component for substitutability. Or it can combine smaller measures rather than purchase maximum redundancy.

The right answer is not automatically dual sourcing. If the primary supplier is exceptionally stable, the component will be redesigned next year, or alternate production can be qualified quickly, paying a permanent premium may destroy more value than it protects. Conversely, if a six-week interruption would threaten a major customer relationship or an annual launch, the premium may be inexpensive insurance. The decision should follow the risk appetite and recovery economics, not a universal rule that single sourcing is bad or redundancy is good.

Resilience Is a Portfolio, Not a Collection of Buffers

Resilience decisions rarely remain isolated. A plant adds safety stock, procurement qualifies another supplier, IT builds redundancy, logistics contracts backup capacity, and a business unit reserves additional labor. Each decision can be individually rational while the combined portfolio becomes expensive, duplicative, or still vulnerable to the same underlying event. The enterprise therefore needs a portfolio view of resilience, not a collection of functional contingency plans.

Correlation is the critical issue. Two suppliers located in different countries may depend on the same sub-tier source. Two distribution centers may share the same port, cloud platform, energy grid, labor market, or transportation corridor. A company can spend heavily on apparent diversification and discover during a disruption that the alternatives fail together. Conversely, one flexible asset, transferable process, or shared reserve may protect several exposures at once. The economics improve when resilience investments are evaluated for the risks they jointly reduce rather than only for the asset or function that owns the budget.

This is also where COO-CFO alignment matters. Operations understands recovery paths, technical dependencies, and service consequences; finance forces comparability across competing uses of cash and capital. Together they should ask whether the current resilience portfolio is concentrated in yesterday's risks, whether protections still match the company's strategy, and whether the same reduction in exposure could be purchased more efficiently. COSO's guidance is especially useful here because it treats risk appetite as an enterprise concept tied to objectives and a portfolio view, not a collection of independent risk limits.

Resilience should therefore be periodically re-underwritten. A buffer justified when a supplier was unstable may no longer be needed after redesign. A backup site that once protected a critical product may become redundant after network expansion. A cyber recovery capability may need more investment as operational technology becomes more connected. The objective is not to preserve every protection once purchased. It is to keep the enterprise's residual risk inside its chosen appetite at the lowest reasonable total cost.

Where Resilience Spending Goes Wrong

Protecting everything equally

Uniform resilience policies feel fair and are often economically weak. Requiring the same inventory days, backup arrangements, or redundancy standards across all exposures ignores differences in criticality and recoverability. Resilience capital should concentrate where the business consequence of failure is greatest.

Pricing the protection but not the disruption

Managers naturally compare supplier premiums, carrying costs, and spare capacity against today's operating budget because those numbers are concrete. If the alternative is described only as 'risk,' the resilience investment usually loses. The comparison should include the full enterprise consequence of disruption and the value of reducing recovery time.

Assuming diversification always reduces risk

Two suppliers can still depend on the same sub-tier source, port, energy system, technology platform, or geography. Two facilities can share a single network, labor market, or critical utility. Resilience requires understanding correlated failure, not merely counting alternatives. Apparent redundancy can disappear precisely when it is needed.

The Counterargument: Resilience Can Become Structural Waste

There is a legitimate danger in overlearning the lessons of disruption. Companies that experienced shortages can respond by adding inventory, suppliers, capacity, and controls faster than they remove them. Temporary protection becomes permanent overhead. Networks become more complex. Purchasing leverage falls. Utilization deteriorates. More suppliers create more quality and compliance work. Extra inventory masks process problems and eventually becomes obsolescence.

That is why resilience requires the same discipline as any other investment. Every protective measure should have an explicit risk hypothesis: what exposure is being reduced, by how much, at what cost, and for how long. As the environment changes, the protection should be challenged again. Risk appetite is not a license to build buffers indefinitely; it is a basis for deciding which buffers continue to earn their cost.

The opposite error is equally dangerous. A resilience measure that is never 'used' is not necessarily wasted. Fire suppression, cyber recovery, emergency power, cross-qualified suppliers, and spare capacity can create value precisely because they are available when normal assumptions fail. Executives must resist judging protection solely by utilization. The relevant question is whether the capability changes the enterprise's loss profile enough to justify its total cost.

The Decision

So how much resilience is worth paying for? Enough to bring the exposures that can materially damage the strategy inside the organization's chosen risk appetite - but not so much that the cost and complexity of protection weaken the business it is supposed to protect.

That requires explicit choices. Some risks should be engineered out. Some should be diversified. Some should be insured or contractually shared. Some deserve inventory or idle capacity. Some require faster detection and recovery rather than expensive duplication. And some should simply be carried because the cost of reducing them is greater than the value of the protection.

The difficult part is not building resilience. It is deciding **where not to build it**. Every dollar committed to redundancy, buffer capacity, inventory, alternate sourcing, or recovery capability is a dollar unavailable for growth, productivity, or innovation. But every dollar saved by eliminating protection also increases an exposure somewhere in the enterprise.

A resilient operation is not one that eliminates uncertainty. It is one whose leaders know which risks they are carrying, why they are carrying them, and what it would cost to carry less.

Selected Sources

Committee of Sponsoring Organizations of the Treadway Commission (COSO). Enterprise Risk Management - Integrating with Strategy and Performance. 2017. [Source](#)

COSO. Risk Appetite - Critical to Success: Using Risk Appetite to Thrive. 2020. [Source](#)

International Organization for Standardization. ISO 31000:2018 - Risk Management - Guidelines. [Source](#)

BCG. As Trade Gets Tougher, Companies Turn to Cost Discipline. February 23, 2026. [Source](#)

BCG. Real-World Supply Chain Resilience. July 29, 2021. [Source](#)

Rice, J. B., Klibi, W., & Trepte, K. Overcoming the Financial Barriers to Building Resilient Supply Chains. Harvard Business Review, November 1, 2022. [Source](#)

Editorial standard: established frameworks, executive application, practical tradeoffs, selective sourcing.