Resilience Never Activated, Deters No One

By
Burak Oktenli

Published: September 28, 2026

Deterrence by denial rests on a practical promise. An attack on a state's infrastructure will not produce the effect an adversary expects. Resilience contributes to that promise only when continuity measures can reduce losses under the disruptions an adversary can impose.

A *continuity plan* is the plan an agency or operator uses to preserve or restore functions that must continue during a disruption. For federal agencies, The Federal Emergency Management Agency (FEMA) document Federal Continuity Directive 2 formalizes the identification and business-process analysis of mission essential functions. An *interdependence model* describes how the loss of capability in one sector propagates to sectors that depend on it. The National Institute of Standards and Technology (NIST) document NIST SP 800-82 describes U.S. critical infrastructure as highly interconnected and mutually dependent, warning that failures can cascade across systems. This is referred to as a *continuity target,* which is one that must be kept above a specified service level or restored within a specified time. It is not a formal FEMA term. The distinction matters because identifying an essential function verses proof that a target reduces systemic damage, otherwise known as a continuity target, are different tasks.

An illustrative Example: The U.S. Economy

Consider an illustrative cascade model using the U.S. economy. The model receives a service-integrity failure: a sector remains physically online, but the output other sectors depend upon can no longer be trusted. The disturbance runs for six time increments using stylized dependency coefficients. Inoperability represents the fraction of useful service loss, while total systemic loss is a cumulative inoperability across all modeled sectors.

If telecommunications is the origin, the struck sector reaches 63 percent inoperability. Yet 71 percent of cumulative systemic loss occurs outside telecommunications, in sectors that depend on it. Much of the damage from a sector failure may therefore emerge elsewhere. Now add a continuity target. Suppose each sector can cap its own inoperability at 30 percent. Applying that target to telecommunications reduces total systemic loss by 44 percent. Applying the same target to any downstream sector does not reduce the total systemic loss.

This is the point. Downstream sectors never lose more than 6 to 23 percent of their service, so their 30 percent continuity targets are never reached. In operations-research terms, the constraint never binds. An organization can spend substantial resources maintaining a target that appears prudent on paper but changes nothing in the scenario being tested.

Now, lower the modeled target to 5 percent and some downstream measures begin to matter. The best produces a 14 percent reduction in systemic loss. The specific percentages are illustrative. The structural result matters more because a continuity target can be demanding in isolation, yet irrelevant once cross-sector propagation is considered.

Upon further analysis, if one changes the origin of the failure, the result changes again. When the same disturbance begins in electric power rather than telecommunications, total systemic loss rises by 15

percent even though the originating sector peaks lower, because damage spreads more evenly. The ranking of which continuity measures provide the greatest reduction also changes.

A continuity target therefore has no deterrent meaning in the abstract. Its value depends on the disruption's origin, dependency structure, severity, and duration.

From Resilience to Deterrence

Natalie Treloar and Jean-Claude Meledje recently argued that dynamic resilience should replace the pursuit of brittle stability. The argument is compelling, but resilience becomes deterrence by denial only when it measurably reduces the damage an attacker expects to inflict. A declared capacity to endure is not the same as a tested capacity that changes an attacker's expected payoff.

An adversary also sees infrastructure differently from an individual operator. An operator plans from inside its own organization; an adversary selecting an origin point can evaluate the network as a whole. Public data already make part of that assessment possible. The Bureau of Economic Analysis (BEA) publishes input-output accounts showing direct and indirect production relationships among U.S. industries. These accounts do not reveal operational vulnerabilities by themselves, but they provide a public starting map of economic interdependence.

An adversary does not need access to every continuity plan to exploit that asymmetry. It needs to identify nodes with disproportionate downstream effects and disruptions likely to be long-lived. Sophisticated adversaries can estimate dependency structures and test alternative origins even when detailed performance targets remain protected.

Also, duration matters as much as depth. In the illustrative model, the health sector ranks only fifth by peak inoperability but recovers last because it combines heavy dependence on other sectors with slow restoration. A triage scheme based only on immediate loss would therefore underrate a function whose prolonged degradation could create greater coercive leverage.

What Should Change

Continuity targets should be tested against cross-sector scenarios before being treated as evidence of denial. A useful target should answer four questions:

1. From which disruption origins was it tested?
2. Does it bind to relevant sectors?
3. How much cumulative system loss does it avert?
4. How much does it shorten degradation or recovery?

The Critical Infrastructure Security and Resilience Agency (CISA) already provides a cross-sector framework through its National Critical Functions work, while NIST guidance recognizes physical and logical interdependencies among operational systems. CISA could sponsor a protected pilot combining public BEA input-output structures with nonpublic operational data on recovery rates, service floors, and continuity costs.

The purpose would not be to publish exploitable thresholds. Detailed targets can remain protected. The public result should instead be an assurance claim that “continuity measures were tested against multiple failure origins, became binding where expected, and produced measurable reductions in cumulative loss and recovery time.”

None of the numerical results above describes the real U.S. economy. The coefficients are illustrative, the disturbance is simplified, and detection delay sits outside the model. That is precisely why the next step should be empirical rather than rhetorical. The public economic accounts exist. Sector operators hold the missing operational parameters, and the question is testable.

Dynamic resilience is a stronger deterrent concept than a promise to return to yesterday's equilibrium. But resilience that never activates to counter the damage an adversary is positioned to inflict does not deny anything. It documents preparation without demonstrating effect.

The standard should therefore be stricter and more useful. Do not just ask whether an essential function has a continuity plan. Ask whether its performance target becomes binding in the scenarios that matter and how much adversary-imposed loss it prevents.

Burak Oktenli is an independent researcher conducting work on the governance of authority in autonomous and AI-enabled systems. His writing has appeared in Dark Reading, RealClearDefense, the Modern War Institute at West Point, and the Lieber Institute's Articles of War. Views expressed in this article are the author's own.