CHAPTER 19 / 4 MINUTE BRIEF

Managing Risks and Anomalies

Read the full chapter → / Read this brief online →

The main idea

Earlier in my career, I had a small business building a search engine. We had a primary system and a backup, which was considerably less reassuring after lightning struck the building and took out both. That was effectively the end of the business. I went back to school.[1]

The story takes a few sentences now. What mattered then was that we had prepared for a narrower problem than the one that occurred: having a backup had not given us a recovery we could perform after losing the place that contained it.

That experience is behind the way I use the word **anomaly**: an event outside the ordinary operating assumptions of the business, one that changes what the capacity plan can promise. No capacity team can predict these events, and I would not ask one to try. I would ask it to know, before one arrives, which decisions will be its own, what each would cost, and how long each would take.

Explore this section in the full chapter →

[1] Tim O'Brien, *Leaving the Cloud Isn't for Everyone*, October 31, 2025, "The Risk Premium You're Paying Cloud to Carry." The passage is the author's recollection.

Force: the kinds of events that hit the business

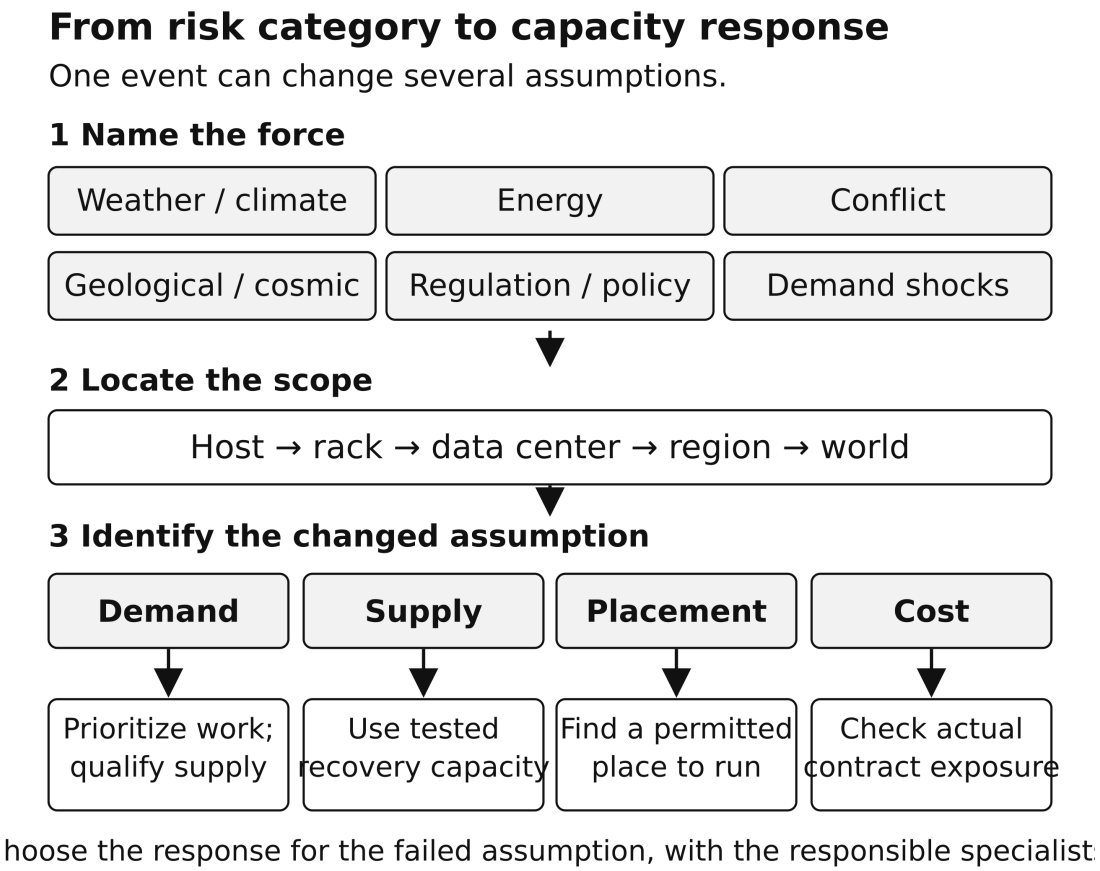


Diagram 19.1. Classify an event by force and scope, then identify which capacity assumptions it changes. The same event can affect demand, supply, permitted placement, and cost; each calls for a different response with the responsible specialists.

Explore this section in the full chapter →

Weather, climate, and the map you inherit

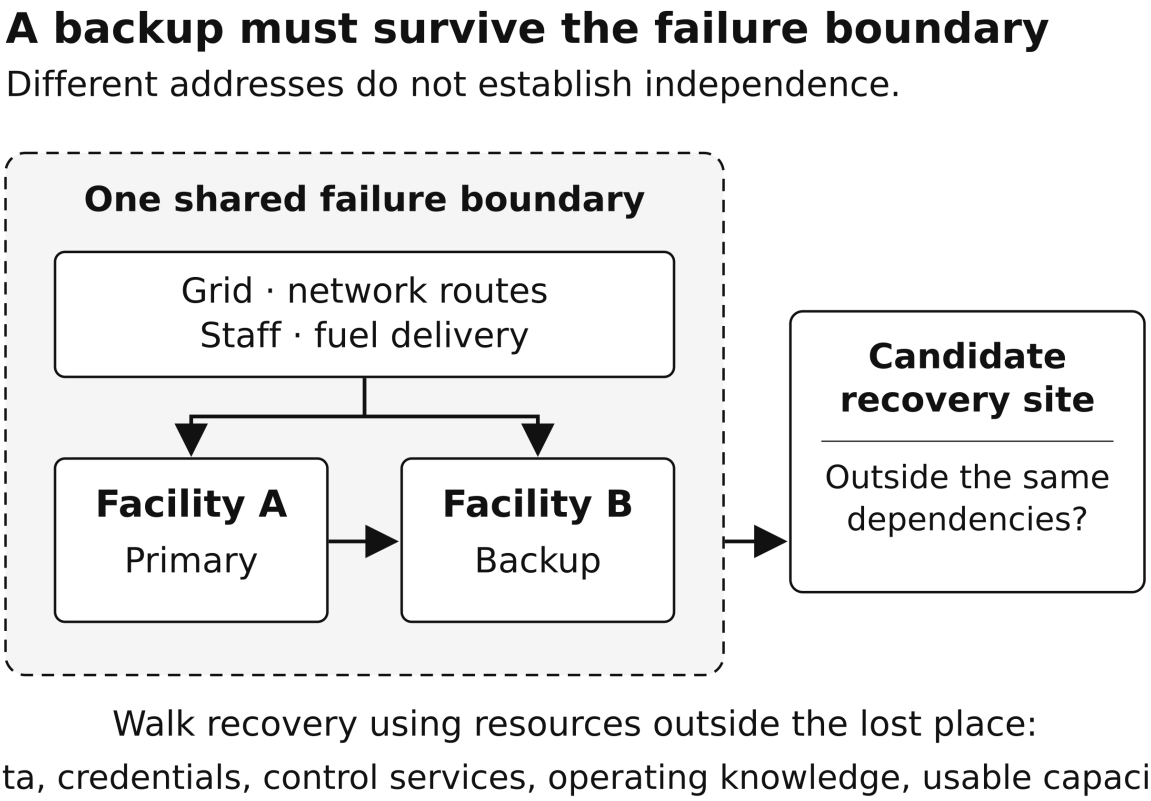


Diagram 19.2. Different addresses do not establish independence. This schematic places two facilities inside one shared failure boundary; a candidate recovery site still needs its dependencies checked and its recovery resources qualified.

Explore this section in the full chapter →

Solar storms: warning and preparation

Solar flares and coronal mass ejections are different phenomena; for the power grid, the one that matters is an Earth-directed coronal mass ejection, or CME. Forecasters can watch a CME leave the Sun. The fastest reach Earth in 15 to 18 hours, and slower ones take several days. Spacecraft near the L1 point, about 1.5 million kilometers toward the Sun, measure the disturbance as it passes, which often gives 15 to 60 minutes of warning.[2] That last interval is far too short to provision and qualify replacement infrastructure. Compare the warning with the actual recovery time, and hold in advance whatever the promised response needs.

Explore this section in the full chapter →

[2]NOAA SWPC, Coronal Mass Ejections, checked October 4, 2026: the fastest Earth-directed CMEs arrive "in as little as 15-18 hours," slower ones in several days; DSCOVR at L1 "can often provide 15 to 60 minutes advanced warning." Typical intervals, not guarantees.

Meteors and the loss of a place

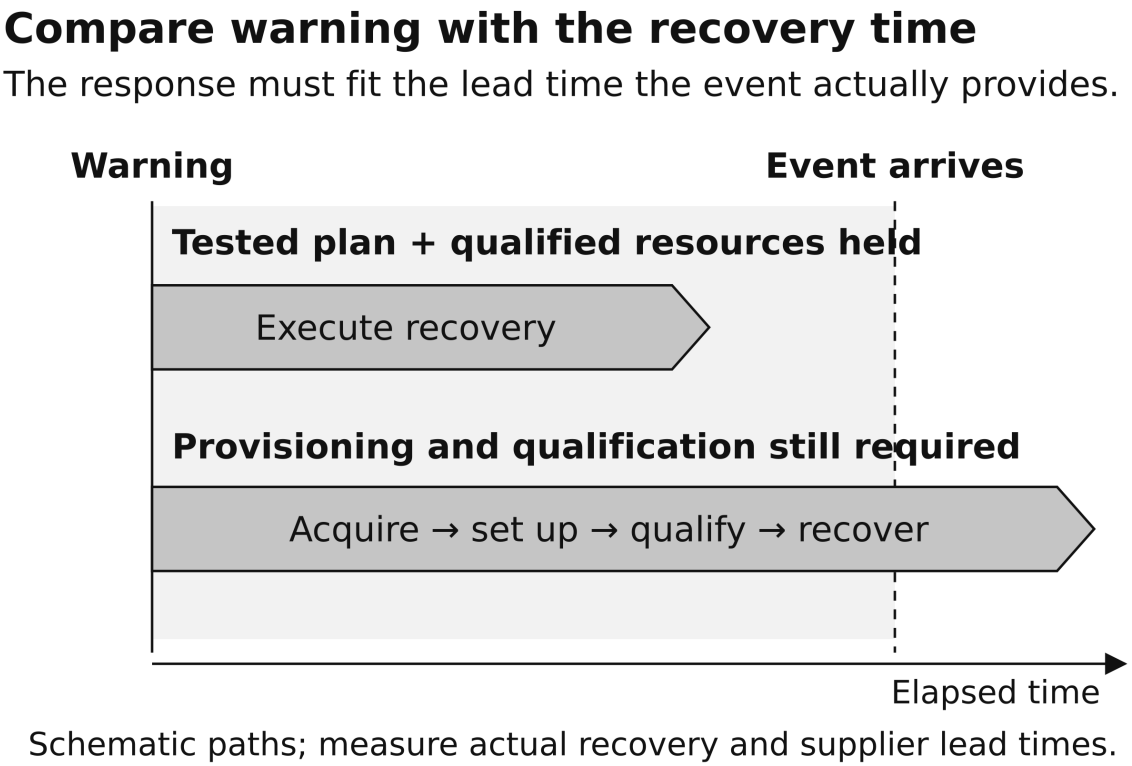


Diagram 19.3. Compare the warning available with the work still needed to recover. These schematic paths illustrate why resources that still need acquisition and qualification cannot be treated as ready capacity; actual response and supplier lead times must be established for the promised recovery.

Explore this section in the full chapter →

COVID, and the anomaly that changes demand

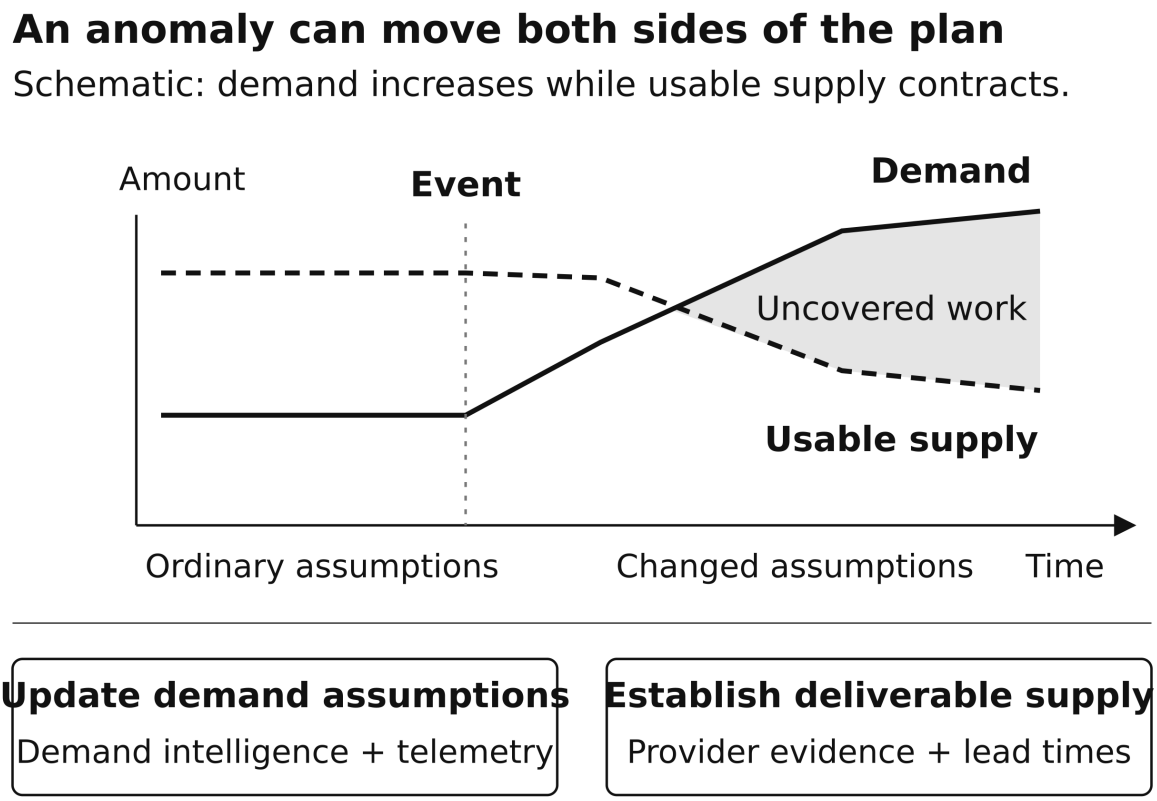


Diagram 19.4. An anomaly can change demand and supply together. This schematic, rather than a measured historical series, shows why updated demand assumptions and evidence of deliverable supply are separate parts of the response.

Explore this section in the full chapter →

The headroom you hold is the failover you can do

This is Chapter 18's failure-group arithmetic with regions as the groups. Take N equal regions sharing the work evenly, each able to sustain a qualified rate C, and let $u_{\max}$ be the highest fraction of C a survivor may carry after one region is lost. Normal utilization u must then satisfy

$$u \le \frac{N-1}{N}\, u_{\max}.$$

With two regions and $u_{\max} = 0.8$, the ceiling is 40 percent: two regions at $0.4C$ send $0.8C$ to the survivor, leaving 20 percent of its qualified capacity free. Letting the survivor run to its full tested limit raises the ceiling to 50 percent. Three regions with the same margin can each run at about 53.3 percent ($3 \times 0.533C \div 2 \approx 0.8C$); at two-thirds each, a loss leaves the survivors no headroom. These are upper bounds under stated assumptions, not utilization targets. Uneven routing, recovery work, a changed request mix, and shared failures each need their own calculation, and no percentage guarantees zero downtime.

Explore this section in the full chapter →

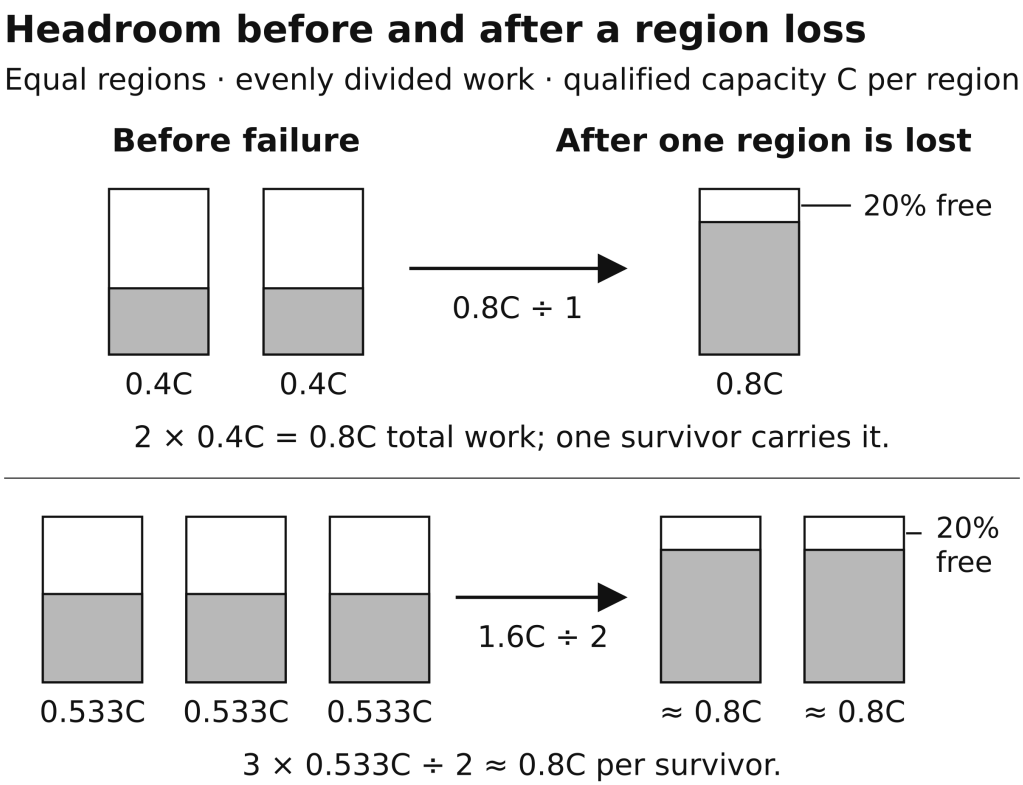


Diagram 19.5. Each outlined bar is a region's qualified capacity C; its gray portion is the work it carries. Under the stated equal-region assumptions, both arrangements leave about 20 percent free after one region is lost. Routing, state, recovery work, and shared failures still need qualification.

Explore this section in the full chapter →

What the capacity team is for, when it happens

The organization should assign ownership of business continuity and disaster recovery explicitly. In the arrangement described here, the capacity team contributes without owning the whole plan: when someone asks “can we handle this?”, we answer alongside the service operators, risk specialists, and business owners whose requirements define what handling it means.

That answer cannot be assembled during an event. Keep a short record for each scenario the organization has decided to plan for:

- **Event and scale**, with the date of the evidence.
- **Affected workloads**, their criticality, and the dependencies they share (Volume I, Chapter 1).
- **Surviving capacity**, and what it is qualified to do (Chapter 18).
- **Recovery deadline**, and the lead time for anything not already held.
- **Permitted destinations**, including legal eligibility (Chapter 21).
- **Provider commitments**, and what they entitle us to in a shortage (Chapter 20).
- **Decision owner and trigger**, and the evidence that would change the response.
- **Accepted gaps**: who accepted them, and when they will be reviewed.

I would write the first record for the scenario the organization is most confident it has covered, because confidence is what stops people asking what two sites share. Its findings go to the next monthly review (Chapter 22) and the next provider conversation.

Explore this section in the full chapter →