

Disaster Recovery Fundamentals: RTO, RPO, and Somewhere Else to Run

Wed, Aug 26, 10:00 AM EDT · <https://scottslab.io/posts/disaster-recovery-fundamentals>



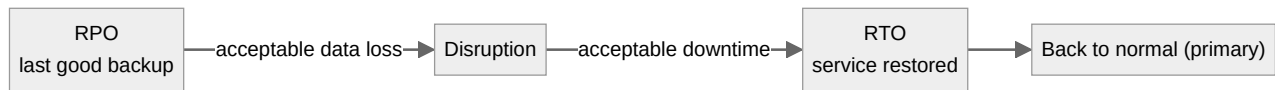
TL;DR

Disaster recovery is the subset of business continuity that gets normal operations back after a disruption: a hurricane, ransomware, or a data-center outage. It isn't "done" until you're fully back in the primary environment. Three numbers frame it: RTO (how fast you must restore a service), RPO (how much data loss you can accept), and RSL (the minimum service level to hold during the disaster). And you need somewhere to run: a hot, warm, or cold site, trading cost against activation time.

Disaster recovery is a subset of business continuity. Where continuity is keeping the business alive through a disruption, DR is the specific job of restoring normal operations after one. The trigger can be natural (a hurricane takes the building), man-made (ransomware encrypts everything), or internal (the data center just goes dark). And here's the part people declare too early: DR concludes only when you're fully back to normal in the primary environment, not when you've limped onto a backup. Running on the spare tire isn't "recovered."


Three metrics do the framing, and every DR conversation should start with them. RTO (Recovery Time Objective) is the target time to get a service back after it goes down; it answers "how long can this be offline?" RPO (Recovery Point Objective) is the maximum acceptable data-loss window; it

answers "how much recent work can we afford to lose?" and it's really a statement about how often you back up. And RSL (Recovery Service Level) is the minimum percentage of a service that must stay available during the disaster, because "degraded but up" is often the real goal, not all-or-nothing.



Disaster Recovery Fundamentals: RTO, RPO, and Somewhere Else to Run

Then you need somewhere else to run, and the three options are a straight cost-versus-speed trade. A hot site is a fully operational mirror of your primary data center. It activates immediately, and it costs the most because you're paying to run a second everything. A warm site has the hardware racked and ready but not running in parallel, so activation takes hours to days while you load data and bring it up. A cold site is an empty shell of power, cooling, cabling, and floor space: cheap to hold but weeks to months to turn into a working environment. Which one you pick falls straight out of your RTO: a two-hour RTO cannot be met from a cold site, full stop.



Disaster Recovery Fundamentals: RTO, RPO, and Somewhere Else to Run

One modern piece ties it together: Resource Capacity Agreements, contracts with providers that guarantee you'll actually get the compute, storage, and networking you need during a disaster. It's the unglamorous detail that saves you, because a disaster is exactly when everyone else is also scrambling for capacity, and "we assumed the cloud would have room" is not a recovery plan.