

You added redundancy.
You hit your SLA.
Your incident playbook
is current.

So why does this keep happening?

Primary circuit	↑UP 99.97%
Secondary circuit	↑UP 99.94%
Failover tested	PASS last run 14d ago
SLA compliance	99.9% within target
Incident tickets	0 open last closed 3d ago

Your monitoring sees healthy components.
Your users are still struggling to complete the task.

THE RESILIENCE GAP

Three common moves that can still leave the workload exposed.


FIX ATTEMPTED	EXPECTED RESULT	✗ ACTUAL RESULT
Added a second circuit from a second carrier.	Two independent paths. Real redundancy.	A second carrier can reduce risk, but independence has to be verified. Two paths may still share a conduit, building entry, facility or upstream network dependency. If that shared dependency degrades, both paths can be impaired while both circuits remain available.
Upgraded monitoring thresholds.	Earlier detection. Faster response.	Brownouts may stay below outage or failover thresholds. Dashboards can stay green while the critical task becomes slow, intermittent or unusable.
Ran post-incident RFOs.	Root cause identified. Problem solved.	The immediate symptom gets addressed, but a shared dependency, failover behavior or task-level monitoring gap may remain. A similar incident can return under different symptoms.

THE STRUCTURAL GAP

The safeguards can work as designed
and still leave the workload exposed.

A second carrier can reduce risk, but path independence has to be verified. Circuits may still share a conduit, building entry or facility, and they may also depend on the same upstream network service. A shared dependency can impair both paths even while both circuits remain up.

WHAT CORRELATED FAILURE LOOKS LIKE



WHAT THE CYCLE COSTS

THE MEASUREMENT GAP

The cost of a hard outage is usually easier to see. Degradation is harder to quantify because it can show up as slow transactions, retries, dropped sessions, support calls and operational work without a clean outage ever being declared. When those effects are not measured, their business impact can stay hidden.

“We don’t measure the cost of slow.
We measure the cost of downtime.”

— F5

What outage costs alone look like:

\$14K/min

The average cost of enterprise downtime continues to climb, reaching as high as \$23,750/min for large enterprises. That figure measures downtime. Degradation can create additional impact that downtime-only reporting does not capture.

EMA Research/BigPanda, 2024

And it’s getting more frequent:

51%

Most enterprises are reporting major financial impact from recurring network incidents, with over half now citing losses of more than \$1M per month, up from 43%. And those figures include degradation events, not just full outages.

Catchpoint Internet Resilience Report, 2025

OUTAGE COSTS ARE EASIER TO SEE. BROWNOUT COSTS ARE EASIER TO MISS.
BOTH ARE REAL.

Most organizations have a clearer view of what downtime costs
than what degradation costs.

Brownouts can show up as slow transactions, retries, dropped sessions, customer
calls and operational work without ever becoming a clean outage.

That impact is real even when downtime-only
reporting never captures it.