CASE STUDY · ENERGY & UTILITIES · ELASTIC COMPUTE

A grid operator held peak agent demand inside a fixed compute envelope

Storm events multiplied agent demand eleven-fold in under an hour. Governed elasticity absorbed the surge without an unbounded bill or a queue of stalled directives.

Partner
Electric utility · 6.1m customers

Published
2026-09-09

Orcher components
Cost Governance, Observability, Logic Scrubber

Summary

Outage-management agents at an electric utility were fine on an ordinary Tuesday and catastrophic during a storm. Orcher's elastic compute governance let demand scale with the event while holding a hard ceiling, shedding low-stakes directives first and keeping restoration work at full priority.

11×
peak surge in agent demand absorbed
March storm event

0
restoration directives dropped or delayed beyond 90 seconds

1.0×
of the pre-agreed cost envelope at peak
ceiling held exactly

40 min → 0
dispatch queue delay at peak

The situation

— Outage-management, crew-dispatch and customer-notification agents shared one unmanaged capacity pool.

— During a March storm the pool saturated: customer notifications crowded out crew dispatch, and restoration directives queued for 40 minutes.

— Provider rate limits produced silent failures that only surfaced in customer complaints.

What Orcher enforced

— Directive classes were ranked by restoration impact, and the ceiling was expressed as a total envelope rather than a per-workflow quota.

— Under pressure Orcher sheds from the bottom of the ranking — notification retries first, dispatch never — and records every shed decision.

— Cross-provider routing spread the surge across three providers, so one provider's rate limit no longer became an outage in the utility's own operations.

— Observability exposed the live envelope to the storm-room dashboard, so the incident commander could see agent capacity next to crew capacity.

The data

Directive class	Priority	Peak volume	Shed at peak	Max latency
Crew dispatch	1 — never shed	8,400/hr	0	9s
Outage confirmation	2	21,000/hr	0	22s
Restoration ETA	3	63,000/hr	4%	71s
Customer notification retry	4 — shed first	118,000/hr	38%	deferred
Post-event summary	5 — deferrable	n/a	100%	batched after event

Peak hour of the March storm event, measured at the control plane.

Outcomes

— The utility retired its per-workflow capacity quotas, which had been the cause of the crowding-out behaviour.

— Storm-room operations now treat agent capacity as a managed resource with a named owner, alongside crews and materials.

— Post-event reporting is generated from the ledger rather than assembled by hand.

— The same envelope model was applied to wildfire-risk inspection agents in the following quarter.

> “The failure mode was never the model. It was that everything shared one bucket and nothing decided what mattered when the bucket ran out.”
>
> — Director of Grid Operations Technology, electric utility (partner declined attribution by name)

Read further

— Elastic compute governance — dippai.com/research/elastic-compute-governance

— Solutions — elastic compute — dippai.com/solutions/elastic-compute

— Full case study — dippai.com/case-studies/utility-elastic-compute-governance

Disclosure. Design-partner engagement, anonymised at the partner's request. Figures are measured by Orcher's own observability and ledger instrumentation over the stated period and have not been independently audited. Market figures carry their own source line.