

11 Microsoft Fabric Cost-Waste Checks

The recurring ways a Fabric capacity bleeds money — worst-first. Rows 1–3 recover the most with the least risk; row 2 is the one that *creates* cost if you get it wrong. Figures as of June 2026; estimates are labeled.

- 1. Idle PAYG capacity running 24/7 for a windowed workload**
Impact: ~\$12/idle hr on F64 → ~\$4,200/mo at 12 idle hrs/day
Fix: Automate scheduled pause/resume (REST API + Azure Automation/Logic Apps).
- 2. The pause trap — pausing to clear throttling debt**
Impact: Settles all smoothed overage at full PAYG at once; on reserved you pay reservation + overage
Fix: Never pause to clear throttling. Size up or fix the query instead.
- 3. Oversized SKU never validated against smoothed load**
Impact: One tier down ≈ 50% of that line (F32 → F16 saves ~\$2,102/mo)
Fix: Read 14 days of Capacity Metrics; right-size to smoothed-under-100%.
- 4. Smoothing/throttling debt misread as "high utilization"**
Impact: Mis-sizing both ways: over-buy or chronic throttling
Fix: Size to 24-h smoothed background <100%, not to peak.
- 5. Pipeline-vs-ADF premium on short/frequent copies**
Impact: "Up to ~10x" — only vs self-hosted-IR ADF, specific profiles
Fix: Use Copy job/incremental load; consolidate runs; scope before migrating.
- 6. Minute-rounding on sub-minute copy activities**
Impact: A 14s → 60s run is ~328% over actual; compounds across 250+ table loads
Fix: Batch short copies; cut run frequency; avoid 1-table-per-run patterns.
- 7. Dataflows Gen2 for heavy ETL**
Impact: 1.5 CU baseline per item plus engine meters; CU-hungry
Fix: Move heavy transforms to Warehouse stored procs or Spark notebooks.
- 8. OneLake duplication / soft-delete retention**
Impact: ~\$0.023/GB-mo per duplicated copy; orphaned + soft-deleted data accrues silently
Fix: Dedupe via shortcuts; prune retention; delete orphaned items.
- 9. Mirroring storage billed on pause**
Impact: Free TB-per-CU allowance flips to billable OneLake storage when paused
Fix: Don't pause capacities hosting active mirrors; budget storage separately.
- 10. Copilot / AI meters with no native kill switch**
Impact: ~0.11 CU-hr per Copilot request (est.); runaway loops compound
Fix: Monitor the Copilot-and-AI meter; gate AI features; watch Data Activator loops.
- 11. Workspace sprawl / Direct Lake fallback**
Impact: Destroys attribution; silent fallback spikes CU with no warning
Fix: 4-workspace topology; monitor fallback telemetry; fix schema breaks.

WANT YOUR NUMBER, NOT THE GENERIC ONE?

- Drop one Capacity Metrics CSV into the free SpendWeave audit and get your recoverable spend in ~60 seconds — it runs entirely in your browser; your raw data never leaves it. spendweave.com/audit