Utilisation Reads Best Right Before It Fails

Auditing utilisation as a fitness claim

By Mario Aiello

Nobody has ever been congratulated for an empty calendar. That is how utilisation got its reputation.

The figure is simple. Of the hours people are available, how many are spent on work. A high number looks like efficiency and a low one looks like waste, so organizations push it up and report it with pride. The audit asks what the number is for.

What the number answers

On fit for what, utilisation says people are busy. It does not say busy with what, or whether anything that matters is moving. A team can be fully utilised on work nobody will use.

On fit for context, the same percentage means different things in different systems. In a factory with steady, predictable demand, high utilisation can be a reasonable goal. In knowledge work, where demand arrives unevenly and the next request is unpredictable, the same percentage describes something else entirely. The number does not know which system it is in.

On fit for how, it measures presence of work, not flow of work. A person can be fully occupied while the thing they are working on waits three weeks for the next person.

On fit until when, it is a target someone set once. The conditions that made 85 percent sensible may be gone, and no one asked.

So utilisation answers how much of the day was occupied, and presumes the other three.

Why it reads best just before it fails

Where demand varies, waiting time in a queue rises steeply as utilisation approaches full. The last stretch of utilisation costs far more in delay than the first. A team at 95 percent has nothing left to absorb an urgent request, an absence, or a piece of work that turned out larger than expected, so everything queues behind everything else.

This is where the indicator hides the failure. As utilisation climbs, the number improves and nobody is idle, so friction appears to fall. Meanwhile delay is growing in every direction. That is Hidden Misfit: everything reads fine while the system is losing the capacity to cope.

Slack is also not a trade-off. It is a boundary. You cannot offset a lack of it with strength elsewhere, because it is what lets the system absorb variation at all. Averaging it away against efficiency, which is what a utilisation target does, removes the condition the rest depends on. The argument for treating it that way is in Boundaries Aren't Trade-offs.

Where this stops

Some work is predictable and some capacity is expensive enough to keep busy, and there high utilisation is a sensible aim. I am not arguing for idle people. The point is narrower: the number

© Mario Aiello | Adaptive Ways

cannot tell you which kind of system you are in, so it cannot be used as proof that anything is working. This is desk work, and I have not measured a real queue.

The question to ask before setting a target is what the work does when something unexpected arrives. If the answer is that it waits, the utilisation figure is already telling you the system has no room, and reading it as good news is the mistake.

More AFS: AFS Library | Latest Essays | Indicators Under Audit

LinkedIn hashtags: #AdaptiveFitness #Utilisation #Metrics #Flow #Measurement

Substack tags: Metrics, Flow, Measurement, Organizational Change, Adaptive Fitness System

Licensed CC BY 4.0 — share freely with attribution