




CASE STUDY · STANDBY EQUIPMENT

Standby equipment: ready on paper, not on demand

How a national energy company turned 32% readiness into 86% for its standby critical equipment — closing the window where a single failure could have stopped production.

BEFORE



AFTER



standby equipment readiness

SWIPE →



THE SITUATION

Where standby readiness slowly erodes

- Standby status was recorded but rarely checked. A unit could sit marked “ready” for months with no real proof it would start when needed.
- When a unit wasn’t ready, no one had one clear picture of why. A decision waiting on approval. A part on order. A repair that needed a rare shutdown window — wherever this lived, it wasn’t in one place, and often it wasn’t tracked at all.
- Risk was judged one unit at a time, if at all. No one could see where it was building up across the whole fleet.
- Repair timing was worked out case by case. With no shared way to weigh the trade-offs, the most urgent voice in the room usually won. That was often production — and fairly so.

WHAT GOOD STANDBY MANAGEMENT REQUIRES

A clear, current status for every standby unit — not a guess · one place to see what’s blocking a unit, and when it clears · a full view of risk across the fleet, not just one unit at a time.

~6 weeks

average time a standby unit sat unconfirmed before first evaluation

~1 in 2

open standby issues with no agreed date for a decision



THE RESULT

~86%

of standby equipment now
confirmed ready on demand


BEFORE

~32%

AFTER

~86%

Heat map of standby readiness across the fleet



The fleet's critical rotating equipment — pumps, compressors,
and similar assets per production Unit.

Ready Partially ready Not ready

TARGET 80%



THE APPROACH

Three of the changes that mattered most

01

A clear view of what's blocking every unit

For any standby unit not fully ready, there is now **one place to see why** — a decision waiting on approval, a part on order, a repair needing a shutdown window — and a **realistic date** for when it clears.

02

A live risk map across the whole fleet

Standby status is no longer judged one unit at a time. A single view now shows **where risk is concentrated across the fleet**, and where it's slowly building before it becomes urgent.

03

A shared basis for prioritising repair

Every open issue is now weighed the same way — safety, production impact, cost of waiting — so when production needs more time, that's a visible, **deliberate trade-off**, not something that just happens by default.



WHAT IT TAUGHT

A backup you can't see isn't a backup at all

Having a backup was never the hard part — installing a second unit is a decision made once. The harder work never stops — knowing whether **that unit would actually start when needed.**

*That's not a maintenance question, and it's not a production question either. It's a question of whether the **organisation can see its own risk clearly enough to manage it** — before it's too late to do anything about it.*

IF THIS RESONATED

- A gap between what looks covered and what's actually confirmed? **Let's talk.**
- Visit annaplekhanova.com for more of my thinking and work.
- Follow **Anna Plekhanova** for more case studies.
- **Save** for your next CI conversation.