The Consumer Test

Four questions for any KRI

Sep 23, 2026 · @Sophy Merrick

Why this exists

An indicator earns its place only if somebody uses it to make a decision. If nobody does, it is not a control. It is a reporting habit with a monthly cost.

Indicators accumulate. One gets added after an incident. One arrives with a regulation. One was inherited and nobody remembers who asked for it. Nothing in the normal cycle removes them, so the set only ever grows.

Run this against your top ten KRIs in about an hour. No project, no permission, no budget line. You need the indicator list, the grid at the end of this page, and nobody's agreement.

It will not tell you what to measure. It will tell you what you are measuring for no reason, and what to do about it.

The four steps

Take your indicator list. Work down it one row at a time, filling the grid at the end of this page as you go. Write the answers down rather than thinking them through, because the blanks are what you will take to the review.

1. **Name the consumer.** Write a person or a forum in column two. Not a distribution list, not "the monthly pack". If you cannot name one, leave it blank.
2. **Write the decision it changes.** Describe what someone would do differently at a different reading. If the honest answer is "we would look into it", leave it blank.
3. **Ask what happens if it stops arriving.** If nobody would notice inside a quarter, mark column two and three blank and move on.
4. **Write the risk it covers.** What was this put here to guard against, and is that still live? Write "unknown" if you do not know. Do this last, and do it before you conclude anything.

Steps one to three test whether the indicator is used. Step four tests whether it matters anyway. An indicator can be a poor instrument and still be the only thing standing where a control should be.

What each failure means

Once the grid is filled, read the blanks against this. Different failures call for different responses, so do not treat them all as removal candidates.

Blank at step	What it means	What to do
1 only	The indicator may be sound but has lost its audience	Find the owner, or stop producing it and see who asks
2 only	You are measuring something real that nobody acts on	Change the threshold so it triggers a decision, or downgrade it to management information
1 and 2	It is a reporting habit	Propose removal
3	Confirms the first two. Nothing depends on it	Propose removal
4 says the risk is live	Stop. The instrument is wrong, the exposure is not	Replace it, do not remove it

Two failures are worth naming separately, because they look like success on a dashboard.

It measures activity, not risk. Counts of things done. Volumes processed. Items raised. These move when the team gets busier and stay still when the risk changes.

It is a ratio with an uncontrolled denominator. Your score moves because something outside your control moved. Nothing about your environment changed and the number did.

Before you remove anything

The most expensive mistake this test can cause is removing something that was quietly holding a risk nobody had articulated.

A worked case: a reconciliation of back-to-back trades between two entities of the same firm. On its face, pointless. Both legs match by construction. It fails questions one to three outright.

It was not pointless. Nobody had written down what the check was actually covering against, so nobody could see what its removal would expose until it was gone.

So before you propose removal, walk the RACI:

- **Who built it, and why?** If that person has left, who inherited the reasoning?

- **Who would find out first if the thing it covers went wrong?** If nobody can answer, that is your finding, not the indicator's.
- **Is there a control elsewhere that assumes this exists?** Indicators get referenced in other functions' procedures without anyone telling you.
- **Does anything outside the firm depend on it?** An auditor, a regulator, a client contract.

If you cannot get a clear answer on what it covers, that is not permission to remove it. It is a reason to write down what you do not know and take that to the review instead.

The awkward answers

Most KRIs that fail this test are not orphans. Somebody wants them there.

An indicator with no consumer is often evidence that a function is busy. It appears in a pack, it gets a green rating, and it demonstrates coverage to whoever asks whether the area is controlled. Removing it takes that evidence away from a colleague who did nothing wrong.

That makes removal a political act, not an analytical one. Two things follow.

First, do not frame it as a cull. A list of things to delete invites defence. A list of what each indicator is for, with the gaps marked, invites a conversation about what should replace them.

Second, expect the response that nothing can be removed because a regulator or auditor expects it. Sometimes true. Often it means somebody said so once and nobody has checked since. The useful question is not whether it is required, but who required it, when, and in what words.

What you take to the review

You almost certainly cannot remove anything on your own authority. That is fine. The output of this test is not a deletion. It is a paper.

A worked example

Illustrative only. An invented non-financial risk indicator set for an operations function, filled as it would look after an hour's work. The point is the blanks and the awkward rows, not the indicators themselves.

Indicator	Consumer	Decision it changes	Risk it covers	Keep / replace / retire
Items processed per FTE per day			Unknown. Reads as a productivity measure, not a risk one	Retire
Training completion rate, %		Nothing below 100%. Chased by the team, read by nobody	Capability gaps, in theory	Retire
Breaks aged over 30 days	Ops Risk Forum, monthly	Escalation to the COO above 15	Settlement failure and client impact	Keep
Inter-entity reconciliation, exceptions			Live. Nobody could explain it, but it is the only check catching static data drift between entities	Keep, and document what it covers
Number of near misses raised	Head of Ops	Only upward pressure. A low count is read as under-reporting, a high count as a problem	Reporting culture, indirectly	Replace. Measures willingness to report, not exposure
Manual workaround count	Unknown. Appears in the quarterly pack, source unclear	Unknown	Process fragility, possibly	Unresolved. Ask who added it before deciding

Four of these six are not clean. That is a normal result, and it is the point. Row four is the one to watch: it fails the first three steps outright and is the last thing you should remove.

Your grid

Now yours. One row per indicator, filled as you work. Take it to the next reporting or governance review as it stands, blanks included.

Indicator	Consumer	Decision it changes	Risk it covers	Keep / replace / retire

The blanks are the argument. A row with three of them makes the case without you having to make it.

Two things make this land rather than sit in a folder. Put a number on it: how many indicators, how many failed, and the hours a month they consume. And bring one replacement proposal, not none, so the meeting has something to agree to rather than only something to defend.

Where else this works

The four questions were written for KRIs, because that is where the waste is easiest to see. They apply anywhere a function produces something on a cycle and calls it control.

RCSA lines are the obvious next place, with one change of emphasis. There the first question is not who consumes it but who owns it. A line covering a control the function does not own produces an attestation nobody can act on, kept because it has always been there and everyone else has one.

The same test also works on standing reports, recurring reconciliations, and any committee paper that has never changed a decision.

If you run this and find a KRI that fails all four, but you cannot get it removed, I would be interested to know why. The reason it survives is usually more instructive than the indicator itself.

Authored by Sophy Merrick. Generative AI used as an editor and structuring aid.