



A MEASUREMENT GUIDE

The Service-to-Outcome Gap

Your returns number tells you people came back. It can't tell you whether the support matched the need.

Every system tracks **Returns to Homelessness** — the HUD measure of how many people, after exiting to permanent housing, end up back in the system. A return is a signal that something didn't hold. But the number is silent on the more useful question: **was the support this person received matched to what they actually needed?** The answer to that lives in your data too — just not anywhere the returns number can reach.

SAME SUPPORT, DIFFERENT NEED — ONE COMES BACK

Household A

Brief crisis · few underlying barriers

SUPPORT RECEIVED

A short rental subsidy and a light-touch check-in. Roughly what their situation called for.

AT 12 MONTHS

Still stably housed

Household B

Long homelessness history · high acuity

SUPPORT RECEIVED

The *same* short subsidy and light-touch check-in — far less than their situation required.

AT 12 MONTHS

Returned to homelessness

The problem wasn't the service — it was the same service given to two very different needs. Household B didn't need more of everything; it needed support matched to a harder situation.

THE POINT ISN'T TO STANDARDIZE — IT'S TO MATCH

Good case management is supposed to look different for different people. A brief-crisis household and a high-acuity one *should* receive different support — that's not a flaw, it's the work done well. The goal is never to give everyone the same package. The goal is to make sure the support each person receives is **matched to the need in front of them** — and to be able to see, across a whole caseload, whether that matching is actually happening.

This is where the returns number and the service record fail to meet. Your **Returns to Homelessness** measure can tell you that Household B came back. Your case files can tell you what Household B received. But the two aren't connected in any form you can analyze at scale — so the system can see *that* people return without being able to ask whether the ones returning were **under-matched to their level of need**. HUD itself reads a high return rate as a sign that people may have needed more support before or after exit. That's a matching question — and it's the one the data can't currently answer.

- 1 **Need and support live in different places.** Acuity, history, and barriers sit in assessments and notes; support delivered sits in service records; returns sit in exit data. Nothing was designed to connect them at the person level.
- 2 **More support isn't automatically better.** The highest-need households receive the most intensive help — so a raw link between support and returns can mislead as easily as it informs. The real question is fit, not volume. Correlation is not causation.
- 3 **The numbers get small fast.** Split a caseload by need level and support pattern and each group shrinks — and small groups swing wildly, making it easy to over-read a mismatch that won't hold up.

A return tells you the support didn't hold. Whether it didn't hold because it wasn't matched to the need is the question worth answering — and the one your data can't answer on its own.

Connecting need, support, and returns — carefully enough to trust the answer — is exactly the work we do. To close the gap in your own data, start at homelessresponselab.com.