How the clarity score works

The full methodology behind the clarity score: inputs, weighting, and why it changes slowly by design.

The score is designed to reflect patterns, not moments

A single distracted session should not define a person. Laivara treats clarity as a slow-moving signal built from repeated patterns across focus, recovery, reflection, and AI reliance.

That makes the score less dramatic, but more useful. It reduces false alarms and makes improvement feel like a trend rather than a daily verdict.

Inputs

What the score considers

The engine can weigh focus block duration, context switching, break regularity, reliance rhythm, task sensitivity, review behavior, and whether users keep returning to unaided thinking.

It does not need to read private content to do this. Metadata and local behavioral signals are enough to show whether the user is building or losing agency over time.

Weighting

Why richer weighting is stronger

A strong score should not treat every AI use as risky or every long session as healthy. It should look for context: the task, the user pattern, the timing, and whether verification or reflection occurred.

This is why Laivara avoids a simple prompt counter. The score is meant to represent calibration, not raw usage.

Privacy

Strength should not require surveillance

The product can be powerful without collecting sensitive text.

Privacy-preserving scoring depends on choosing signals that are meaningful enough for feedback and limited enough to protect users.

That is the balance: honest measurement without turning human attention into content inventory.