

Layer 1 - Missed-Risk Suppression

Foundation Effect / Missed-Risk Suppression

Do not silently miss boundary-near risks.

Problem

AI workflows may adopt or discard boundary-near cases too early. Low-frequency signals, unresolved risks, and conditionally useful recommendations can disappear before commitment.

Trust Kernel action

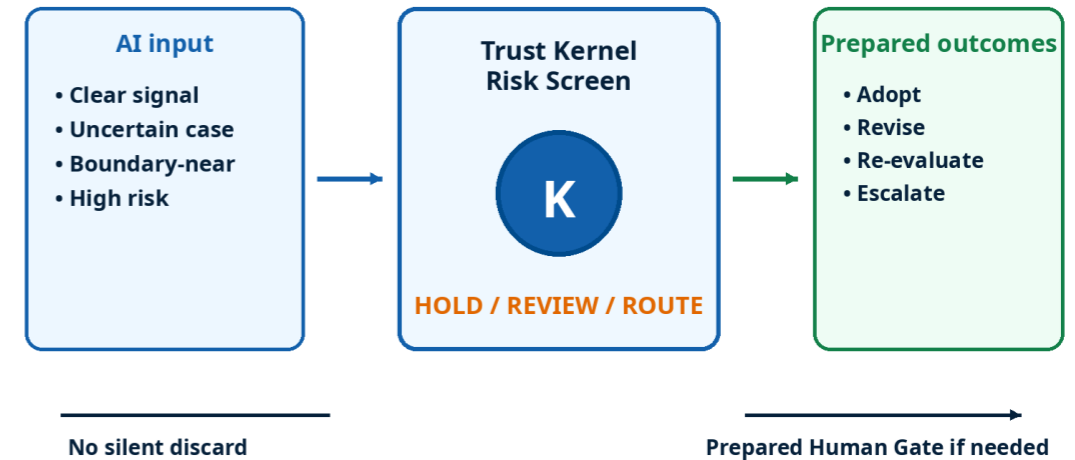
The Trust Kernel preserves uncertain or boundary-near states before commitment. It supports HOLD, re-evaluation, evidence requests, counter-evidence, repair options, and route-state selection.

Effect

Uncertain candidates become prepared candidates. They can be adopted, revised, held, re-evaluated, returned to an earlier workflow, or escalated with a prepared Human Gate package.

Public boundary

This is a control structure for reducing missed-risk exposure. It is not a public claim of zero misses or universal reliability.



Who this helps

AI workflows where near-boundary misses are costly.

Discussion scope

Workflow fit, risk classes, POC framing, NDA boundary.

Decision preparation vocabulary

Adopt Revise HOLD Re-evaluate Repair Escalate Return