

When should an agent leave the device?

Abstract. The local-first compute thesis treats placement as a human decision as well as a scheduler decision: stay local when latency or privacy makes that right; propose a temporary burst only when the task earns it. It is a design target, not a shipped control plane in this repository.

THE PROBLEM

Can a placement policy jointly optimize latency, energy, memory pressure, cost, and consent while giving the owner a reason they can understand and override?

HARD CONSTRAINTS

- 01 The user’s instruction and data boundary outrank a throughput-only optimum.
- 02 A proposal must distinguish a local failure from a cloud-worthy escalation.
- 03 Bring-your-own-cloud credentials and job lifecycle need a defensible trust boundary.
- 04 A burst must have an accountable teardown and outcome, not an orphaned resource.

A FIRST CONTRIBUTION

Build a simulator with reproducible device/resource traces, counterfactual placement decisions, and an owner-visible explanation for every proposed escalation.

BUILD IN THE OPEN

github.com/hushh-labs/hushh-research
hushh.ai/discord

