

Tidewell Brain

One agentic brain, two modes: Solo on the robot, Crew across the site.

What it does

- Solo: full local execution on one robot, no network dependency (Both)
- Crew: shared core, short-term and long-term memory across every robot on the site (Both)
- Task planning, verification gate and human approval (Both)
- Cross-embodiment handoff between wheeled and legged robots (Both)
- Learning from corrections, promoted to the fleet through a human gate (Both)

Specifications

SPECIFICATION	DESIGN TARGET	MEASURED
On-robot manipulation policy	3B-class VLA at 10 to 25 Hz on Jetson Thor	—
On-robot planner	7 to 8B model, 40 to 80 tokens per second	—
On-robot episodic buffer	30 to 60 minutes, captions with pose and time	—
Core memory replica on robot	Under 50 MB, pinned summary in planner context	—
Link-loss detection	Under 2 s, automatic downgrade to Solo	—
Re-sync after reconnect	Append-only upload, robot never blocks	—
Site memory read latency	Under 1 s on the site LAN	—
Site planner and judge	30 to 70B model, 0.5 to 3 s per call at task boundaries	—
Fleet arbitration latency	100 ms to 1 s through Open-RMF	—
Robots per site server	5 to 20 on one RTX Pro 6000 class appliance	—
Site network assumption	Private 5G or Wi-Fi 6E, 10 to 30 ms typical, tails to 150 ms	—

SPECIFICATION	DESIGN TARGET	MEASURED
Offline mode	Documented, audited by a third party	—
Firmware signing keys	Held by the customer	—
Vendor telemetry	None	—
Measured values are added as each test completes and carry the test date.		

Safety and certification status

IEC 61508 · Link-loss detection, automatic downgrade to Solo and the verification gate as safety functions · in-design

Deployment model

Per robot licence · Crew site server as a Tidewell appliance · Crew software on a qualified customer server

114 Lavender Street, #07-51, CT Hub2, Singapore 338729 · 5 September 2026