Per-Agent Learned Drift Models

by Nick Clark | Published April 25, 2026

What Per-Agent Learned Drift Models Specifies

Each unit's local clock drifts with temperature, age, and oscillator characteristics. A learned drift model captures the unit-specific drift pattern; predictive correction reduces the deviation between the local clock and the consensus.

The learned model is governance-credentialed. The model parameters, training history, and uncertainty bounds are recorded structurally. Other units can evaluate the contributing unit's drift quality before weighting its time observations.

Why It Matters Structurally

Static drift compensation (factory calibration, periodic re-calibration) leaves residual drift that grows with operating time. The residual drift produces consensus quality degradation across the fleet.

Learned per-agent models close the residual gap. Each unit's drift compensation improves with operating experience; the fleet-wide consensus quality improves as the contributing units' models mature.

How It Composes With Mesh Operation

The drift model trains continuously against the consensus. Each consensus update produces a residual against the unit's predicted clock; the residual feeds the drift learning.

Cross-validation operates structurally. Outlier residuals identify possible clock failures or environmental anomalies; the architecture surfaces the outliers as credentialed diagnostic events.

What This Enables for Resilient Timekeeping

Defense-grade timekeeping under contested conditions gains improving fleet quality with operating time. Civilian fleet operations gain the same.

The architecture supports unit-specific quality declaration. High-stability oscillator units contribute with high weight; lower-stability units contribute with appropriate weight; the consensus reflects the actual quality distribution of the fleet.