Credentialed Range Observations With Lineage

by Nick Clark | Published April 25, 2026

What Credentialed Range Observations Specify

Each range observation carries: the identity of the contributing unit (cryptographic identity tied to the unit's keying material), the modality of the observation (UWB, lidar, radar, RFID, etc.), the declared uncertainty (range estimate plus uncertainty bound under the modality's error model), the timestamp, and a signature binding all of the above.

The receiving solver evaluates each observation for admissibility before integrating. Admissibility includes credential validity, modality acceptance under current operating context, uncertainty acceptance under solution requirements, and freshness against the operating window.

Why Lineage Matters for Position Audit

When a position-derived decision is later audited (post-incident analysis, regulatory review, adversarial-action investigation), the question 'what observations produced this position' has an architecturally-supported answer only when the observations carry lineage.

Without lineage, the audit reconstructs from logging that wasn't structured for the audit purpose. With lineage, the audit reads structured records that the architecture

maintained for exactly this purpose. The difference is the difference between defensible audit and best-effort reconstruction.

How Lineage Composes With Solution

The multilateration solution carries forward lineage from each contributing observation. The position estimate links to the observation set; each observation links to its contributing unit, modality, and timestamp. Downstream consumers of the position can traverse the lineage to inspect the contributing observations.

Lineage retention is governance-credentialed. The retention period, the retention authority, and the access controls are declared structurally. The architecture supports lineage retention requirements that vary by operating jurisdiction and operational class.

What This Enables for Audit-Grade Positioning

Defense operations subject to engagement-decision review gain positioning audit that survives adversarial scrutiny. Civilian operations subject to liability review (autonomous-vehicle incidents, surgical-robotics outcomes) gain the same.

The architecture also supports diagnostic differentiation. When a position is later identified as bad, the lineage permits identification of which observation contributed the error and whether the error pattern indicates sensor failure, environmental condition, or adversarial interference.