

# Focus Area 3.0 — Trusted and Extensible Multi-Source Fusion

Confidence, Extensibility, and Traceability for Target Custody and Characterization

|                   |                                                               |
|-------------------|---------------------------------------------------------------|
| Sponsoring Agency | Space Systems Command (SSC), System Delta 85 (SYD 85), Kronos |
| Focus Area Type   | Capability (Software)                                         |
| Priority Tier     | Tier 2 — Programmed                                           |
| Document Version  | Version 1                                                     |
| Date              | 21 September 2026                                             |

## Table of Contents

|                                                                                            |    |
|--------------------------------------------------------------------------------------------|----|
| 1. Focus Area Overview.....                                                                | 3  |
| 2. High Level Operational Context.....                                                     | 4  |
| 3. High Level Requirement.....                                                             | 4  |
| 4. Stakeholders.....                                                                       | 5  |
| 5. Context & Assumptions.....                                                              | 5  |
| 6. Problem Statement 1.1 — Multi-Source Fusion, Common Data Model, and Source Lineage..... | 7  |
| 6.1. Problem Statement.....                                                                | 7  |
| 6.2. Operational Context.....                                                              | 8  |
| 6.3. Requirement.....                                                                      | 8  |
| 6.4. Data Sources.....                                                                     | 9  |
| 6.5. Operational Value.....                                                                | 9  |
| 6.6. Desired Capability Outcomes / Acceptance Criteria.....                                | 9  |
| 7. Problem Statement 1.2 — Rapidly Extensible Data Model for New Data Source Fusion.....   | 10 |
| 7.1. Problem Statement.....                                                                | 10 |

|                                                                                          |    |
|------------------------------------------------------------------------------------------|----|
| 7.2. Operational Context.....                                                            | 10 |
| 7.3. Requirement.....                                                                    | 11 |
| 7.4. Data Sources.....                                                                   | 12 |
| 7.5. Operational Value.....                                                              | 12 |
| 7.6. Desired Capability Outcomes / Acceptance Criteria.....                              | 13 |
| 8. Problem Statement 1.3 — Independence-Weighted Confidence Scoring and Calibration..... | 13 |
| 8.1. Problem Statement.....                                                              | 13 |
| 8.2. Operational Context.....                                                            | 14 |
| 8.3. Requirement.....                                                                    | 14 |
| 8.4. Data Sources.....                                                                   | 15 |
| 8.5. Operational Value.....                                                              | 16 |
| 8.6. Desired Capability Outcomes / Acceptance Criteria.....                              | 16 |

## 1. Focus Area Overview

**Bottom Line Up Front (BLUF):** Focus Area 3.0 develops a processing and exploitation capability that collects, collates, and fuses observations from many phenomenologically diverse sensors into a single coherent track for a Resident Space Object, scored separately for whether it is real, where it is and will be, whether its observations belong together, which cataloged object it is, and what kind of object it is. Each score is reduced by the age of the data behind it, carries the sources that produced it, and is empirically calibrated, so an operator can act on it, and an evaluator can check it. Contributing data products are collated rather than dissolved: each is retained in its original form and remains individually discoverable behind the coherent track. New and unfamiliar sources are admitted into the fused product without a software release, and at an influence proportionate to what has been demonstrated about them. Aligned under the System Delta 85 (SYD 85) Program Management Office (PMO), this Focus Area matures the capability from a prototype baseline (TRL 4) to TRL 7 for transition to the Kronos Family of Systems (FoS).

Three Problem Statements (PSs) form a single processing chain. PS 1.1 ingests heterogeneous observations, resolves them to objects, expresses them in a documented and versioned data model without dissolving them into it, and preserves the lineage of every input. PS 1.2 admits new and unfamiliar sources into that model without a software release, establishing from evidence what each source contributes and granting it influence in proportion to what has been demonstrated. "Without a software release" refers to technology versions and human-engineered data model and/or knowledge graph releases that would cause significant delay. It does not restrict human-machine teaming solutions for mapping fields or confirming suggestions from proposed smart tools. PS 1.3 converts the fused product and its lineage into separate, explainable confidence values for existence, state, association, identification, and characterization, weighted by the independent information each source contributes, and demonstrates empirically that those values are calibrated against observed outcomes.

PS 1.2 sits between the other two rather than after them. It extends the data model PS 1.1 publishes, and it hands PS 1.3 a provisional independence declaration for every source it admits. That relationship is deliberate: an extensibility capability built without reference to how sources are weighted will admit sources that corrupt the confidence product, and a scoring capability built without reference to how sources arrive will be calibrated only against the sources present on the day it was built.

Operators today assemble target-quality custody and identification from fragmented government, commercial, and academic sources that differ in sensor type, quality, timing, and format. The number of available sources is growing faster than the capacity to adjudicate them, and faster than the software release cycle that governs which of them can be used at all. Fielded automated event-detection and command-and-control systems present conclusions without exposing the sources, method, or confidence behind them; shared data libraries distribute source data without adjudicating it; commercial providers maintain independent catalogs using provider-specific methods. Published work exists on uncorrelated track resolution, on schema mapping, and on data model versioning, and offerors are

expected to cite and differentiate against it. What is absent is a confidence value that is traceable to its sources, does not overstate itself when those sources are not independent, has been demonstrated to be calibrated against real outcomes, and remains calibrated as new sources are admitted.

Development and demonstrations occur at UNCLASSIFIED on public, commercial, and academic data, and the product must remain valid and traceable at Secret Internet Protocol Router (SIPR) and Top Secret (TS) / Secret Compartmented Information (SCI)/ Special Access Program (SAP). Execution depends on sustained collaboration across operational squadrons, space operations centers, and commercial and academic sensor providers. Outcomes resolved by the squadrons and operations centers form the basis of the Government furnished evaluation sets against which calibration is measured.

Proposals must include an explanation of how they resolve Problem Statements 1.1, 1.2, and 1.3 in the context of an integrated solution. Problem Statements 1.2 and 1.3 address capabilities that do not exist today; Problem Statement 1.1 provides the data model on which both depend.

## 2. High Level Operational Context

The capability operates in the Processing and Exploitation stages of the Tasking, Collection, Processing, Exploitation and Dissemination (TCPED) cycle. Processing fuses heterogeneous feeds into a single track per object and admits new feeds to that process; Exploitation scores each dimension and records the reasoning. In the Joint Targeting Cycle (JP 3-60) the capability supports Find, Fix, and Track in service of positive identification and custody; Space Doctrine Publication 3-100 establishes the requirement to combine terrestrial, radio frequency, and orbital sources into a coherent domain picture.

The operational deficiency is not a shortage of data. Custody and characterization decisions are made at the crew position under time pressure, using products that assert an answer without exposing how sure that answer is. An operator who cannot see the basis of a value must either accept it without justification or reconstruct it manually, and both consume decision time. As proliferated commercial and academic sensing expands the observation pool, the cost of that manual reconciliation grows.

A second deficiency compounds the first. The rate at which the fused product can improve is currently set by the software release cycle, because every new source requires schema work, code, and a release before any of its data is usable. The operational community, which is best placed to recognize that a source would be valuable, has no route to nominate one and no evidence with which to argue the case, since nothing can be measured until it has already been integrated. After transition, this becomes the governing constraint on the capability's useful life, because the sources available to the enterprise will not be the sources available during development.

## 3. High Level Requirement

Develop capabilities that collect, collate, and fuse observations from heterogeneous, phenomenologically diverse sources into a single coherent track per Resident Space Object (RSO); express all inputs in a documented, versioned common data model while retaining each contributing

data product in its original form and keeping it individually discoverable behind the track it informs; preserve source lineage from the point of ingest; admit new and unfamiliar sources into that model and into the fused product without a software release, on demonstrated evidence and at an influence proportionate to that evidence; produce separate, explainable confidence values for existence, state, association, identification, and characterization, weighted by independent information rather than source count and reduced by the age of the data behind them; and empirically demonstrate that reported confidence is calibrated against observed outcomes and remains calibrated as sources are admitted, withdrawn, and changed.

## 4. Stakeholders

- **PRIMARY: Space Systems Command (SSC) System Delta 85 (SYD 85):** (Direct Program Management Office (PMO) and acquisition sponsor)
- **PRIMARY: Battle Management, Command, Control, Communications, and Space Intelligence (BMC3I) Portfolio Acquisition Executive (PAE) and SSC:** (Portfolio and acquisition authority)
- **PRIMARY: Kronos Family of Systems (FoS):** (Sponsoring Program of Record and receiving program)
- **Space Operations Centers (National Space Defense Center (NSDC), Combined Space Operations Center (CSpOC)):** (Operational customers; source of resolved outcomes for calibration)
- **Combat Forces Command (CFC):** (Operational receiving units; nominating authority for candidate data sources; source of resolved outcomes for calibration)
- **Intelligence Community Subject Matter Experts:** (Consumers of characterization products at higher classification)
- **Air Force Research Laboratory's (AFRL) SpaceWERX:** (Commercial transition partner)
- **Commercial and Academic Sensor Providers:** (Observation and catalog data providers)
- **Participating Cohort Companies:** (Organized as Focus Area Cohort Teams (FACTs))

## 5. Context & Assumptions

- 1) No single collection source possesses all required data, and no single analytic tool performs all required functions. Fusion across phenomenologies increases confidence more than reliance on single-phenomenology data sets.
- 2) Sources are not independent by default. Overlapping observations, shared derivation from a common public catalog, a shared timing reference, and shared observing conditions all create the opportunity for distinct sources to carry the same information. Combining them without accounting for it understates uncertainty and produces overconfidence in fusion output.
  - a. Correlation between sources must be decomposed into the portion explained by known shared derivation and the unexplained residual. A confidence penalty applies to all

- correlation, explained and unexplained; only the unexplained residual warrants a reviewable flag.
- b. Source count and phenomenology count are proxies for independence, not measures of it. Acceptance is scored on independence of information contributed, not on feeds connected.
- 3) Source lineage must be preserved from the point of ingest. Once information has propagated through multiple processing paths, separating shared from independent contributions becomes impractical at scale.
  - 4) Fusion here means collation, not amalgamation. The operational product is a single coherent track that lets a crew see every facet of an object at once, on demand; it is not a single amalgamated record that has absorbed and replaced the products that informed it.
    - a. No input is replaced by its transformed form. Expressing an observation in the common data model produces an additional representation of it, never a replacement for it.
    - b. The coherent track is attached to its contributing data products rather than substituted for them, and every contributing product remains individually discoverable and retrievable behind it.
    - c. A downstream consumer may therefore take the coherent track, an individual contributing product, or both. A consumer whose question differs from the one the fusion was built to answer is not forced to re-collect data the enterprise already holds.
  - 5) Every track is assessed in multiple distinct dimensions, each answering a different question and strengthened by different evidence. A blended score, or a dimension that bundles two questions, hides which part of a track is weak. The evidence named is illustrative, and the set is expected to grow.
    - a. Existence — whether there is a real object or a false detection. Strengthened by independent detections across sensors and phenomenologies.
    - b. State — where the object is and will be: position and velocity, with uncertainty that grows with time since observation. Strengthened by longer observation arcs, diverse geometry, radar and laser ranging, and owner-operator ephemerides.
    - c. Association — whether observations belong to the same object, including across a maneuver, breakup, or deployment. Strengthened by revisit, historical custody, and distinguishing brightness and emitter signatures.
    - d. Identification — which cataloged object this is, if any, and so it's owner of record. An uncatalogued object can carry high association and no identification confidence. Strengthened by launch history, catalog correlation, and emitter characteristics.
    - e. Characterization — what kind of object it is and its condition, including type, physical attributes, and status such as tumbling or inactive. Strengthened by signature analysis, radio frequency collection, and orbital behavior over time.

- 6) New sources are not uniform in the difficulty they present, and the capability must distinguish three classes because each demands a different mechanism:
  - a. Class A — the source expresses data the published model already represents, and admission is a mapping problem.
  - b. Class B — the source carries information bearing on a dimension the model already holds, but expresses it in field names, labels, or identifiers that do not match anything in the model. The relationship to existing fields must be discovered from co-observation of the same objects rather than declared.
  - c. Class C — the source carries information the model does not represent at all and must enter the fusion as its own independently assessed variable. Class C admission may require human authorship of the model extension; Class A and Class B admission must not.
- 7) A new source earns influence rather than receiving it. The weight a source carries is a function of two factors that are distinct and must be reported separately:
  - a. Assessed independence — the independent information the source contributes relative to the sources already present, which is the subject of Problem Statement 1.3.
  - b. Evidence maturity — how much has been demonstrated about the source against resolved outcomes. A source with high potential independence and no accumulated evidence is not yet worth its independence. Maturity is established by accumulated resolved outcomes rather than by elapsed time alone, so sources mature on independent timelines, and many will be maturing concurrently at different states of influence.
- 8) A confidence value is only useful if calibrated. Calibration is distinct from accuracy and is empirically testable. Because the source population changes, calibration is a property that must be maintained rather than established once.
- 9) All development, demonstration, and evaluation occur at UNCLASS on public, commercial, and academic data, and the product must remain valid and traceable at higher classification without rework.
- 10) Data Sources associated with, and listed within, this Focus Area's Problem Statements are not necessarily provided to offerors as part of this solicitation.
- 11) Representative Data Sets for technology offerings may be scaled to a representative sample sufficient to demonstrate the offeror's technology.
- 12) Each cohort team is a Focus Area Cohort Team (FACT) of one Prime and one to four Subcontractors, for a recommended team size of two to five companies. Technologies enter at TRL 4, scored on the lowest TRL of any component to be integrated.

## 6. Problem Statement 1.1 — Multi-Source Fusion, Common Data Model, and Source Lineage

## 6.1. Problem Statement

Information for a single RSO is spread across sources that differ in sensor type, accuracy, timing, and format, with no common data model governing how inputs are expressed. Analyst reconciliation consumes decision time, introduces error, and does not scale as proliferated sensing expands the observation pool. Products that do exist tend to resolve the disagreement by replacing their inputs with a single merged record, which settles the question at the cost of the evidence: an operator cannot examine what the answer was built from, and a consumer with a different question cannot reuse the underlying collection. Without a coherent product that retains both the provenance and the original form of every input, no downstream capability can determine how much verifiably independent information stands behind any element of the track.

## 6.2. Operational Context

There is a need for a fusion and normalization capability that aggregates, associates, and resolves heterogeneous observations into a single track per object while recording where each contribution came from. This occupies the Processing stage of TCPED and supports the Find and Fix functions of JP 3-60.

The data model published here is the portfolio interface. Problem Statement 1.2 extends it as new sources arrive and Problem Statement 1.3 scores against it, so it must be documented and versioned from the outset rather than treated as an internal artifact.

## 6.3. Requirement

- **Aggregation and Association:** Ingest observations from multiple heterogeneous sources over distinct phenomenologies, associate them with objects, and resolve duplicate or conflicting reports into a single track per object.
- **Common Data Model:** Express every input in a documented data model in which each field carries one precise meaning. Publish that model and its interface control documentation under version control and structure it so that elements can be added without invalidating a consumer bound to an earlier version.
- **Non-Destructive Collation:** Retain every contributing data product in its original form, fields, and meaning. The coherent track is attached to its contributing products, not substituted for them, and any contributing product must remain individually discoverable and retrievable from every track it informs. A consumer must be able to take the coherent track, an individual contributing product, or both, and a contributing product must remain usable by a consumer that does not use the common data model at all.
- **Source Lineage:** Record source identity, sensor type and collection geometry, collection time, and ingest time for every observation, and make that lineage retrievable for any value in the fused track. Problem Statement 1.3 cannot assess independence without it, and Problem Statement 1.2 cannot establish what a new source contributes without it.

- **Object Identity:** Maintain an internal object identity independent of an external catalog, support uncatalogued objects, and permit a catalog association to be absent, assigned later, or corrected without rebuilding the track.
- **Merge and Split:** Support multiple tracks merging into one object and one-track splitting into several, preserving the prior state so a superseded association can still be examined.
- **Method Independence:** Document the model, the interfaces, and the lineage record in sufficient detail that the Government can evaluate a structurally different implementation by the same method, without access to proprietary internals. This is what allows offerors to be compared during the cohort and components to be replaced after transition.
- **Timeliness:** Meet stated ingest-to-availability targets tiered by source class, set at source onboarding as a term of the data agreement, under representative observation volume.
- **DevSecOps and Security:** Align with BMC3I Tools, Applications & Processing Lab (TAP) Lab DevSecOps pipeline standards, including containerized deployment, Single Sign-On, and Cybersecurity Maturity Model Certification (CMMC) compliance.

#### 6.4. Data Sources

- Commercial space domain awareness (SDA) catalog and observation providers
- Wide-field ground-based optical sensors
- On-orbit SDA sensors; deployable multi-modal sensor kits
- Passive and bistatic radar using illuminators of opportunity
- Event-based and neuromorphic camera systems; university and academic telescope networks
- Government optical and radar feeds, and the Unified Data Library (UDL)
- Formats, access process, volumes, and delivery cadence for each class are set at source onboarding as a term of the data agreement.

#### 6.5. Operational Value

Delivers higher-confidence custody and characterization by combining many varied sources, including data collected today but rarely fused. Because contributing products survive collation rather than being consumed by it, the same collection serves the fused picture and any consumer that needs the underlying observation for a different purpose. Establishes the common data model and lineage record every downstream capability depends on, reducing the integration cost of each additional source and consumer.

#### 6.6. Desired Capability Outcomes / Acceptance Criteria

- Collates and fuses observations from multiple distinct sources across distinct phenomenologies into a single coherent track per object, using public, unclassified, or commercial data.

- Produces existence, state, association, identification, and characterization estimates as separately addressable values on every track.
- Retrieves, for any value in a fused track, the complete lineage of the inputs that produced it — source identity, sensor type, collection geometry, band, timing reference, collection and ingest time, and derivation path.
- Retrieves any contributing data product from any coherent track it informs in its original as-collected form, fields, and meaning, and demonstrates that a consumer can take the coherent track, an individual contributing product, or both.
- Maintains tracks for uncatalogued objects, and demonstrates assignment, later correction, and absence of a catalog association without rebuilding the track.
- Demonstrates track merge and split against a defined scenario set, with the superseded prior state recoverable.
- Publishes the data model under version control and demonstrates that an element can be added without invalidating a consumer bound to the prior version.
- Meets stated ingest-to-availability targets by source class under representative volume, measured and reported per tier, and reflects data age in the fused product without operator action.
- Delivers the common data model specification, interface control documentation, test procedures, performance results, and transition recommendations.

## 7. Problem Statement 1.2 — Rapidly Extensible Data Model for New Data Source Fusion

### 7.1. Problem Statement

The common data model is a contract, and a contract is what makes each additional source expensive. A new provider, an unfamiliar field name, or an unrecognized identifier requires schema work, code, and software release before any of its data reaches the fused product, so the observation pool grows faster than the capability to admit it. The operational community, which is best placed to recognize that a source would be valuable, has no route to nominate one and no evidence with which to argue the case, because nothing can be measured until it has already been integrated. There is no means of determining whether an unfamiliar data element corroborates something the model already carries or represents something novel, and no means of admitting it at an influence proportionate to what has been demonstrated about it.

### 7.2. Operational Context

There is a need for a capability that admits new and unfamiliar sources into the fused product without a software release, on evidence rather than on assertion. It occupies the Processing stage of TCPED alongside Problem Statement 1.1, whose data model and lineage record it extends, and it supplies Problem Statement 1.3 with a provisional assessment of what each newly admitted source contributes.

Today, adding a source is an engineering action, not an operational one. A unit that recognizes a useful feed can ask for it and then wait: the request becomes a schema change, a code change, and a place in a release queue, and nothing about the source can be measured until that work is already complete. The judgment that matters most — whether the source improves custody or characterization of the objects this crew cares about — is therefore made last, on intuition, by the people least able to test it. After transition, this becomes the governing constraint on the capability's useful life, because the sources available to the enterprise will not be the sources that were available when it was built.

A candidate source arrives as a feed or a file with its own field names — the labels a provider attaches to the values it delivers — its own identifiers for the objects it reports, its own units, and its own timing reference. The same physical quantity may arrive under one label from one provider and a different label from another, with neither matching the published model, and a source may carry a quantity the model does not represent at all. Sorting out which of those situations applies is the primary thing this capability must do, because each requires a different mechanism and carries a different level of risk.

Two properties make the problem tractable. First, the difference between those situations is characterizable in advance, so a nominated source can be classified and routed to the right mechanism rather than treated as a bespoke integration. Second, a source does not have to be trusted or rejected on arrival. It can be admitted at no influence, observed against resolved outcomes, and granted weight only as it earns it, which converts an integration decision into a measurement.

### 7.3. Requirement

- **Nomination and Profiling:** Accept a nominated candidate source with sample data and available metadata, profile its field names, identifiers, units, timing reference, and coordinate frame, and classify it as Class A, B, or C. Fields that cannot be mapped are listed explicitly and retained for later assessment, not silently discarded.
- **Admission Without a Software Release:** Admit Class A and Class B sources to the lowest state of influence with no code change and no software release. Class C admission may require a human to author the model extension, but the capability must produce the evidence on which that extension is based, and the extension must be committed without rebuilding the fused product.
- **Evidence-Based Relationship Discovery:** For a Class B source, establish the relationship between its unmapped elements and the dimensions the model already carries from co-observation of the same objects over time, rather than from declared metadata. Record the evidence supporting each discovered relationship, and make each relationship individually reviewable, attributable, and reversible.
- **Additive Model Extension:** For a Class C source, extend the published data model to carry the new element additively and under version control, such that a downstream consumer bound to the prior version continues to operate unchanged, and such that the new element is separately addressable rather than folded into an existing field.
- **Graduated Admission:** Admit a source through no fewer than three states of increasing influence. Initially, the source runs in a parallel path and has no effect on the live fused product. In the

second, it affects only the single dimension in which it has demonstrated efficacy, at a weight discounted for evidence maturity, and cannot affect any other. In the third, it contributes to every dimension in which efficacy has been demonstrated, at the weight its assessed independence warrants. The number of states, the thresholds between them, and the discount applied are the offerors to propose, must be documented, and must be derived from evidence rather than set by an operator.

- **Evidence Sufficiency Rather Than Elapsed Time:** State, for each candidate source, the evidence required to advance to the next state, and report progress against it. Advancement is governed by accumulated resolved outcomes, not by a fixed evaluation window. Sources will therefore mature on independent timelines, and the capability must carry many concurrently at different states without one blocking another.
- **Calibration Checkpoint at Each Advance:** Demonstrate at each advance that the fused product remains calibrated at the source's new weight, reported for each affected dimension. A source that cannot advance without degrading calibration does not advance.
- **Demotion and Withdrawal:** Reduce an admitted source's state automatically and immediately when its agreement with resolved outcomes degrades, notify the operator of the demotion and the evidence behind it, and allow the operator to reverse it. Support complete withdrawal of an admitted source with the fused product returning to its prior state.
- **Provisional Independence Declaration:** Declare the independent information a candidate source is expected to contribute, validated against collection geometry and modality rather than accepted as asserted, and pass that declaration to Problem Statement 1.3 marked as provisional. Evidence of maturity and assessed independence are separate factors and must be reported separately at every state.
- **Scope:** Negotiating data agreements, obtaining authority to access a source, and accrediting a source for operational use are outside the scope of this Problem Statement. The capability states what access it requires and hands off.

## 7.4. Data Sources

- Fused tracks, lineage records, and the published common data model from Problem Statement 1.1
- Candidate source sample sets withheld from the offeror during development and furnished by the Government for onboarding demonstration, covering all three classes
- Commercial, academic, and public catalog and observation feeds not used during development
- Resolved outcome sets for each dimension a candidate source claims to inform
- Source and sensor metadata for candidate sources: phenomenology, aperture location, band, timing reference, illumination geometry, and declared derivation
- Operationally nominated candidate sources and data sets originating from the receiving units

## 7.5. Operational Value

Decouples the growth of the observation pool from the software release cycle, which is the constraint that currently governs how quickly the fused product can improve. Gives the operational community a route to nominate a source and an evidentiary basis on which to argue for it, rather than requiring the case to be made before anything can be measured.

Limits the risk of admitting an unproven source by making influence proportional to demonstrated efficacy, so a source that turns out to be poor degrades nothing while it is being assessed. Preserves the value of the capability after transition, when the sources available to the enterprise will differ from those available during development, and when the alternative is a return to the release cycle for every addition.

## 7.6. Desired Capability Outcomes / Acceptance Criteria

- Onboards a Class A source withheld from the offeror during development to the lowest state of influence with no code change and no software release and reports elapsed time from nomination to initial shadow scoring.
- Onboards a Class B source withheld during development with no code change, discovering the relationship between its unmapped elements and the dimensions the model already carries from co-observation, and producing the supporting evidence for each discovered relationship.
- Produces, for a Class C source, a proposed additive model extension with its supporting evidence, and demonstrates after commitment that a downstream consumer bound to the prior model version operates unchanged and that the new element is separately addressable.
- Lists every field of a nominated source that could not be mapped and retains those fields for later assessment.
- Advances a source through no fewer than three states of increasing influence, reporting at each advance the evidence that justified it, the weight applied, and the measured effect on the fused product.
- Demonstrates that the fused product remains calibrated at each advance, reported separately for each affected dimension, and that a source that would degrade calibration is held at its current state.
- Demonstrates automatic demotion of an admitted source when its agreement with resolved outcomes degrades, with operator notification and the ability to reverse it, and complete withdrawal of an admitted source with the fused product returning to its prior state.
- Reports evidence maturity and assessed independence as separate values for every admitted source at every state.
- Delivers the classification and admission methodology, the state-advancement criteria and discount function with their assumptions, model versioning and interface control documentation, onboarding and withdrawal test results, and transition recommendations.

## 8. Problem Statement 1.3 — Independence-Weighted Confidence Scoring and Calibration

### 8.1. Problem Statement

Overconfidence in fused output increases the risk of error in decision-making, and current fused products give the operator no means of detecting it. A fused track without a trustworthy measure of its own reliability transfers the adjudication burden back to the operator. Confidence values that rise with the number of contributing sources, rather than with the independent information those sources carry, systematically overstate certainty, because commercial and academic sources routinely share underlying observations, derive from the same public catalog, share a timing reference, and observe under shared conditions.

A confidence value is also an assertion until it is shown to be calibrated. Calibration is distinct from accuracy and is empirically testable, and it is currently unaddressed: no method, evaluation set, or outcome definition exists for establishing that reported confidence in any assessed dimension means what it claims.

### 8.2. Operational Context

There is a need for a scoring capability that consumes the fused track and lineage record from Problem Statement 1.1 and the provisional independence declarations from Problem Statement 1.2, produces separate and explainable confidence values weighted by independent information rather than source count, and demonstrates empirically that those values are calibrated. This occupies the Exploitation stage of TCPED and corresponds to the Orient function of the Observe-Orient-Decide-Act (OODA) loop.

Calibration is treated here as a property of the scoring capability rather than as a separate activity, because the two cannot be developed apart: a scoring method that cannot be shown to be calibrated has not been finished, and a calibration method developed independently of the scoring it grades tends to measure what is convenient rather than what matters. The outcome definitions and labeled evaluation sets against which calibration is measured are furnished by the Government so that results remain comparable across the cohort and cannot be tuned by the party being graded.

### 8.3. Requirement

- **Per-Dimension Scoring:** Produce separate confidence values for existence, state, association, identification, and characterization on every track, with state covering the predicted state over a stated horizon as well as the current estimate. A single blended value is insufficient, because it does not identify which dimension is weak or what would strengthen it. Version the dimension set so a dimension can be added without invalidating a consumer bound to the existing set. Assessment of mission, function, or intent is outside the scope of this Problem Statement.
- **Independence Assessment:** Assess the independent information each source contributes, using declared source metadata validated against collection geometry and modality and refined by observed residual correlation as history accumulates. Declared metadata must not be trusted on its

own; disagreement between declared and geometrically supportable independence must surface for review at onboarding.

- **Common-Mode Decomposition:** Decompose correlation between sources into the portion explained by known shared derivation, shared observations, or a shared reference, and the unexplained residual. A confidence penalty applies automatically to all correlation, explained and unexplained; an attributable, reviewable flag is raised only on the unexplained residual.
- **Marginal Contribution:** Express the confidence gain from an added source as a function of its independence from the existing set, with diminishing returns. Fixed or assumed gain values are not acceptable.
- **Evidence Maturity as a Separate Factor:** Apply the maturity state assigned by Problem Statement 1.2 as a factor distinct from assessed independence and report the two separately wherever a source's contribution is exposed. A source's potential independence and what has been demonstrated about it are different quantities and collapsing them into a single weight makes both unauditale.
- **Data Age and Staleness:** Treat the age of contributing data as a factor in the score rather than as an attribute displayed beside it. Reported confidence must decay as the data behind it becomes stale, without operator action, at a rate that is documented, defensible, and stated separately for each dimension, because a predicted state degrades with age far faster than an object's type does. The capability must also expose how a score is expected to change over time, so an operator can distinguish a value that is weak now from one that is about to become weak.
- **Explainability:** Produce on request a plain-language explanation naming the contributing sources, the independent information each provided, its maturity state, and the age of the data behind it, with each contributing data product retrievable in its original form.
- **Change Control:** Where the capability learns or adjusts relationships between sources over time, record those changes so each is individually reviewable, attributable, and reversible, with the supporting evidence available to a human reviewer.
- **Calibration Measurement:** Measure calibration empirically against the Government furnished outcome definitions and labeled evaluation sets, comparing reported confidence against observed outcome rates on held-out data, and report results in a form an accreditation audience can evaluate, including reliability curves and a stated calibration error metric. Report calibration separately for each dimension; an aggregate figure is insufficient because the dimensions are informed by different evidence and may be calibrated to different degrees. Where the Government furnished sets cannot support calibration of a dimension, its value is still reported, marked as uncalibrated.
- **Independence Model Validation:** Validate the independence model against observed outcomes rather than against inter-source agreement, and state the method used to avoid crediting agreement that arises from shared derivation. Scoring independence by consensus among sources is circular and is not acceptable as the sole method.

- **Robustness and Operating Envelope:** Measure calibration under conditions absent from development data — degraded or denied collection, reduced coverage, and shifts in the observed object population — and state where reported confidence remains valid and where it weakens.

## 8.4. Data Sources

- Fused tracks and associated lineage records from Problem Statement 1.1
- Provisional independence declarations and maturity states from Problem Statement 1.2
- Source and sensor metadata: phenomenology, aperture location, band, timing reference, illumination geometry
- Historical source performance and reliability records; sensor uncertainty characterization records
- Public catalog derivation records for commercial and academic sources
- Government furnished outcome definitions and labeled evaluation sets for each scored dimension, with held-out partitions reserved for calibration measurement. The sets are drawn from outcomes resolved by operational squadrons and space operations centers and corroborated by independent confirmation, such as owner-operator ephemerides, later high-quality observations, and confirmed catalog updates, rather than from operator adjudication alone
- Historical event repositories supporting replay against known outcomes; synthetic events for controlled comparison under conditions absent from the record

## 8.5. Operational Value

Converts a fused track into a product an operator can act on and an evaluation system can validate, and ensures operators receive validated products from independently assessed sources rather than a large sensor population producing the appearance of corroboration without the substance of it.

Because the dimensions are scored separately, the weakest dimension identifies which collection would close the gap. That identification is the input a subsequent collection decision requires; making that decision and generating the resulting tasking are outside the scope of this Focus Area.

Calibration converts the confidence product from plausible to defensible. It provides the evidence an accreditation authority requires, gives the operator a documented basis for reliance, and gives the Government a method-agnostic standard for comparing offerors and re-evaluating the capability as sources and components change after transition.

## 8.6. Desired Capability Outcomes / Acceptance Criteria

- Produces separate confidence values for existence, state, association, identification, and characterization on every track, with state confidence reported for the current estimate and for the predicted state over a stated horizon.
- Separates real objects from injected false detections on a Government furnished set, with existence confidence rising as independent detections accumulate and staying low for an uncorroborated detection.

- Reports association and identification confidence independently, showing an uncatalogued object with high association and no identification confidence as a valid track state, and that correcting a catalog association leaves association confidence unchanged.
- Demonstrates that adding a source that shares underlying observations, catalog derivation, or timing reference with an already-contributing source does not inflate reported confidence, across a test set built for that purpose.
- Demonstrates diminishing marginal gain as sources of decreasing independence are added, with the gain function documented and defensible.
- Reports assessed independence and evidence maturity as separate values wherever a source's contribution is exposed.
- Identifies outlying data from sensor observation and track sources and provides proper disposition, including exclusion from track determination, uncorrelated track identification, and calibration of data source integrity.
- Produces on request, for any score, a plain-language explanation naming the contributing sources, the independent information each provided, its maturity state, and the age of the data behind it.
- Demonstrates that reported confidence decays with the age of contributing data without operator action, at a documented rate stated separately for each dimension, and that the decayed value is what calibration is measured against rather than being exempted from it.
- Demonstrates in replay of a government furnished event containing an unannounced maneuver that predicted-state and association confidence fall without operator action and recover as post-maneuver observations accumulate.
- Demonstrates calibration empirically on held-out data: outcomes at each reported confidence level occur at the reported rate within a stated tolerance, reported via reliability curves and a stated calibration error metric, separately for each dimension the Government furnished sets support, with any remaining dimension reported and marked as uncalibrated.
- Validates the independence model against observed outcomes rather than inter-source agreement, documenting the method used to avoid circular credit from shared derivation.
- Measures calibration under at least degraded or reduced-coverage collection and a shift in the observed object population and states the validated operating envelope.
- Delivers the scoring methodology, the independence model and its assumptions, calibration methodology and reliability results, interface control documentation for the score and explanation payload, and a statement of validated operating conditions suitable for an accreditation package.