

Focus Area 4.0 — Object, Orbit, and Engagement Characterization

Confidence, Pattern of Life, and Engagement Risk for Resident Space Object 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.....	5
4. Stakeholders.....	5
5. Context & Assumptions.....	5
6. Problem Statement 1.1 — Orbit and Object Characterization from Data Sets and Historical Records	7
6.1. Problem Statement.....	7
6.2. Operational Context.....	7
6.3. Requirement.....	8
6.4. Data Sources.....	8
6.5. Operational Value.....	9
6.6. Desired Capability Outcomes / Acceptance Criteria.....	9
7. Problem Statement 1.2 — Pattern of Life Baselines and Out-of-Family Deviation Detection.....	10

7.1. Problem Statement.....	10
7.2. Operational Context.....	10
7.3. Requirement.....	10
7.4. Data Sources.....	11
7.5. Operational Value.....	11
7.6. Desired Capability Outcomes / Acceptance Criteria.....	11
8. Problem Statement 1.3 — Engagement Geometry, Risk, and Threat Characterization.....	12
8.1. Problem Statement.....	12
8.2. Operational Context.....	12
8.3. Requirement.....	13
8.4. Data Sources.....	14
8.5. Operational Value.....	14
8.6. Desired Capability Outcomes / Acceptance Criteria.....	14

1. Focus Area Overview

Bottom Line Up Front (BLUF): Focus Area 4.0 — Object, Orbit, and Engagement Characterization develops an analytic capability that carries a Resident Space Object (RSO) from a set of orbital elements to a characterized engagement: what kind of object it is, what its Pattern of Life is (e.g., what orbit it occupies and how it normally behaves), which other orbits it can intersect or interact with, and whether that interaction presents risk to space operations. Space Domain Awareness (SDA) operators and intelligence analysts manually assemble this picture today across separate tools, which costs hours against events that develop within a single orbital revolution. Focus Area 4.0 will be addressed by an integrated team of two to five companies that delivers a three-stage characterization chain (Problem Statements 1.1 through 1.3) by ingesting orbit-determination data sets, public information, and historical catalog records. This effort will establish per-object behavioral baselines (also known as an object Pattern of Life/Pattern of Behavior (PoL/PoB)), detect out-of-family deviation, and characterize engagements involving one or two objects acting on a single object of interest, with a per-dimension confidence value and a traceable, operator-readable explanation. Components enter Technology Readiness Level (TRL) 4 and exit at TRL 7, transitioning to the Kronos Family of Systems (FoS).

The capability is built as a three-stage analytic chain in which each stage consumes the output of the one before it. Problem Statement 1.1 establishes the object and its orbit: it ingests orbit-determination data sets, General Perturbations (GP) element history, and other relevant information to estimate the orbit and its uncertainty, classify orbit regime and orbit family, infer object class and mission class, and record a confidence value and full source lineage for each of those outputs. Problem Statement 1.2 consumes the characterized object and builds its PoL (PoL), the behavioral envelope the object has historically held, then detects, classifies, and scores deviations in orbital parameters that fall outside that envelope. Problem Statement 1.3 consumes both, compares the object's orbit against the orbits it can intersect, approach, or otherwise interact with, characterizes the resulting engagement as one-on-one or two-on-one, and scores the risk that engagement presents to space operations with the driving factors exposed to the operator.

An analyst assembling this picture today works across a conjunction screening service, a catalog data set service, a separate maneuver-detection tool, and open-source mission records, then reconciles the result by hand into a written assessment. Orbit regime and object mission class are read from human-maintained reference material rather than derived from observation with stated confidence, so a new, uncorrelated, recently maneuvered, or deliberately ambiguous object is characterized late or not at all. PoL, where it exists, is held in the working memory of individual analysts rather than in a query able baseline, so behavior is recognized as out of family only by the analyst who happens to know that object. Relationship analysis stops at pairwise conjunction screening, which answers whether two objects will pass close to one another and not whether two objects are converging on the same object of interest. The cost is measured in decision time: a characterization that should support action inside an

orbital revolution instead takes hours to days, and the confidence behind it is neither recorded nor reproducible.

Characterization of this kind is produced at the tactical level by orbital analysts and SDA crews, consumed at the operational level by the National Space Defense Center (NSDC), the Combined Space Operations Center (CSpOC), and Space Forces-Space (S4S), and carried to the strategic level by United States Space Command (USSPACECOM) and Intelligence Community (IC) partners.

The BMC3I Tools, Applications & Processing Lab (TAP) Lab seeks an unclassified technology demonstration, and every input named in this Focus Area is public, commercial, or synthetic. The classified variants of this problem differ in the sensitivity of the catalog and in the labeling of threatening behavior, not in the analytic structure, so a solution that is ontology-driven and traceable to its sources at the unclassified level moves to higher classification with limited rework.

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. Each Problem Statement consumes the output of the one before it; Problem Statement 1.1 provides the data model and ingest interface on which Problem Statements 1.2 and 1.3 depend.

2. High Level Operational Context

Space Domain Awareness supports the Space Tasking Cycle and the Tasking, Collection, Processing, Exploitation and Dissemination (TCPED) lifecycle, and Joint Publication 3-14 (Space Operations) and Space Doctrine Publication 3-100 (Space Domain Awareness) both place object characterization at the front of that lifecycle. The work described in this Focus Area sits in the Processing and Exploitation stages: the observations have been collected and correlated, and the question is what kind of object it is, how it is behaving, and what that behavior means for other objects in the domain. In Joint Targeting Cycle (JP 3-60) terms, the capability supports Find, Fix, and Track; in Observe-Orient-Decide-Act (OODA) terms it is the Orient function that feeds Decide.

The workflow today begins when an object requires characterization — a newly cataloged object, an object that has maneuvered, an object flagged by a conjunction screening product, or an object named in a request for information from an operations center. An orbital analyst retrieves the current and historical data sets for that object, reads its orbit regime and family from reference material or infers them from the elements by inspection, and looks up what is publicly known about the payload and its mission. If the object has maneuvered, the analyst compares orbit-determination data sets across epochs to estimate what changed, and judges whether that change is normal for the object from personal familiarity rather than from a recorded baseline. If the object is near others, the analyst runs or requests a conjunction screening, which returns pairwise close approaches within a screening volume.

Each of those steps produces an answer with no attached confidence and no recorded derivation. The analyst then composes them into an assessment, and the assessment inherits the weakest of them without saying which one it was. When the question is not about a single pair of objects — when two objects converge on the same object of interest, which is where coordinated behavior Initially becomes

visible — the pairwise screening tools do not express the question, and the analyst reconstructs the relationship from separate pairwise results by hand. This is the point at which characterization stops keeping pace with the domain: the analytic product arrives after the decision window it was meant to support, and Joint All-Domain Command and Control (JADC2) integration objectives cannot be met by a product that cannot state how much of itself to believe.

3. High Level Requirement

Develop capabilities that ingest publicly available orbital data sets, historical catalog records, and other relevant publicly available information for a Resident Space Object; estimate and classify its orbit and infer its object and mission class with a stated confidence and preserved source lineage; establish a historical PoL baseline for that object and detect, classify, and score deviations in orbital parameters that fall outside it; compare the object's orbit against other orbits it can intersect, approach, or interact with; characterize the resulting engagement as one-on-one or two-on-one as an analytic output rather than an operator-selected mode; determine and score the risk that engagement presents to space operations using stated, inspectable algorithms; report every output with a per-dimension confidence value, an operator-readable explanation of the factors that drove it, and traceability to originating data; and accept, through a defined ingest interface, an externally produced confidence-scored fused track in place of or alongside data-derived state, without combining the two, and without change to downstream analytic contracts.

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 (NSDC, CSpOC):** (Operational customers; engagement characterization supports watch floor assessment)
- **Combat Forces Command (CFC):** (Operational receiving organization; catalog maintenance and orbital analysis crews and primary operators of the characterization output)
- **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) Orbit-determination data sets and historical catalog records are sufficient to establish orbit regime, object class, and a behavioral baseline for most cataloged objects. They are not sufficient to establish Intelligence Community-derived identity or intent, and this Focus Area does not expressly require either. Throughout this document, characterizing an object means classifying it by orbit regime and family, object class, and mission class. Identity already present in public record is used as an input where it exists; establishing the specific identity of an object is not a required output.
- 2) Engagements within the scope of this Focus Area occur in two structural forms, and the capability must produce both rather than treating the Instead as the general case. Both forms apply whatever the ownership of the objects involved, including engagements between spacecraft of the same operator or nation:
 - a. One-on-one — a single RSO in a relationship with a single other object, which is the form current conjunction screening already expresses.
 - b. Two-on-one — two RSOs whose position, reachability, or behavior bears on the same object of interest, whether that convergence is coordinated. This is the smallest structure in which coordinated behavior is observable, and it appears in the public record as co-location, rendezvous and proximity operations, and multiple arrivals into a shared slot, so it can be evaluated against real events rather than only against synthetic ones.

Structures larger than two-on-one, including correlated behavior across a group where the significant finding is the structure of the group itself, are outside the scope of this Focus Area. Solutions should not preclude them but will not be evaluated against them.

- 3) Behavior that is anomalous for one object is routine for another. Deviation must be scored against the baseline for that specific object and its class, not against a global threshold, and a capability that applies one threshold across the catalog will be simultaneously over-sensitive and blind.
- 4) PoL and its violations are directly connected to a temporal context window. Actions that may be deemed anomalous are only classified as such until they are routine to the point of regularity. For example, a newly launched RSO may behave differently in its instead months of operation, or dual-use RSOs may behave differently depending on how they are functioning in support of a given mission. Problem Statement 1.2 is required to resolve this window explicitly through recurring re-baselining.
- 5) No single analytic discipline covers the chain. Orbit determination and classification, behavioral analytics over long time series, and engagement geometry with risk modeling are distinct specialties, and a solution built entirely within one of them will be weak in the other two.
- 6) The fused-track products available to the space enterprise will change over time as separate capabilities are fielded and replaced. The characterization chain must be insulated from that change by a defined ingest interface rather than coupled to any one fusion implementation.
- 7) Confidence is reported per output dimension, because each dimension is separately usable and separately wrong, and a single value covering several of them is not actionable. The output

dimensions are orbit state, orbit classification, object class, and mission class (Problem Statement 1.1); behavioral baseline and deviation significance (Problem Statement 1.2); and engagement structure and engagement risk (Problem Statement 1.3). Confidence dimensions reported by an externally produced fused track are inputs, carried forward to the output dimensions they bear on, and do not replace this set.

- 8) The resolved outcome is the baseline of record for engagement risk characterization where the documented record establishes what occurred, and Government analyst adjudication is the baseline where it does not. A labeled scenario set representative of routine, ambiguous, and threatening cases may be used for evaluation.
- 9) Verification and validation are conducted by the Government.
- 10) Any learned component in the chain must monitor its own performance in operation, detect when that performance falls below the level demonstrated at acceptance, and support retraining with the change recorded.
- 11) Data Sources associated with, and listed within, this Focus Area's Problem Statements are not necessarily provided to offerors as part of this solicitation.
- 12) Representative Data Sets for technology offerings may be scaled to a representative sample sufficient to demonstrate the offeror's technology.
- 13) 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 — Orbit and Object Characterization from Data Sets and Historical Records

6.1. Problem Statement

Orbit regime, orbit family, and object mission class are treated as reference attributes looked up from human-maintained sources rather than as analytic products derived from observation and reported with stated confidence. When an object is newly cataloged, uncorrelated, recently maneuvered, or deliberately ambiguous, that reference material is absent, stale, or wrong, and there is no recorded measure of how much of it to believe or of what it was derived from.

6.2. Operational Context

This is the entry stage of the characterization chain and sits in TCPED Processing. Every downstream product depends on it: a behavioral baseline cannot be built for an object whose class is unknown, and an engagement cannot be characterized without the orbit state and uncertainty of the objects involved. Orbital analysts performing catalog maintenance and operations center crews responding to a request for information both begin here, and both currently begin with a lookup rather than a derivation.

The stage is also the portfolio interface. Orbital data sets are the given for this Focus Area, but they are not the only possible source of state and separately fielded multi-source fusion capabilities produce confidence-scored tracks that are better than element sets alone for some objects and some questions. The ingest interface defined in this Problem Statement is what allows that substitution to happen later without redesigning the chain behind it.

6.3. Requirement

- **Orbit Estimation and Classification:** Develop a capability that ingests orbit-determination data, their historical record, and other relevant information for a designated RSO, estimates the orbit and its uncertainty, and classifies the orbit by regime and by family — including altitude, inclination, eccentricity, period, repeat ground track, and longitude slot or drift behavior where applicable.
- **Object and Mission Class Inference:** Infer object class (payload, rocket body, debris, unknown) and mission class from the orbital signature together with public launch, registry, mission, and payload records. Where that evidence is consistent with more than one mission class, as with dual-use objects, report each candidate class with its own confidence value rather than forcing a single class.
- **Per-Dimension Confidence and Lineage:** Report a separate confidence value for orbit state, orbit classification, object class, and mission class, and preserve source, data-set epoch, and derivation method for each, so that the age of the underlying data is reflected in the confidence reported.
- **Common Data Model:** Normalize every input and output into a defined data model in which each field carries a single documented meaning and publish that model so downstream Problem Statements and external systems bind to it rather than to implementation.
- **Fused-Track Ingest Interface:** Provide a documented ingest interface that accepts an externally produced, confidence-scored fused track in place of data-derived state or alongside it. Where both are present, they are retained side by side, and the source of state used for each output is recorded; they are not combined into a new estimate. The interface must carry forward each confidence value the incoming track reports, by dimension and with any calibration status attached, together with its source lineage, and the confidence reported for any output that uses the fused track must reflect it. Demonstrate that substitution occurs without change to the output consumed by Problem Statements 1.2 and 1.3.
- **Scope:** Independent sensor tasking, the collection of new observations, and the development of a capability that fuses observations or tracks from multiple sources are outside the scope of this Problem Statement. Fused tracks are consumed through the ingest interface, not produced by the offeror; a representative fused-track sample, which may be simulated, is sufficient for demonstration, and its quality is not evaluated. The capability works from the catalog record and from whatever state is presented at its ingest interface, and reports low confidence, and the dimension responsible, in that same published form rather than requesting collection when the record is insufficient.

6.4. Data Sources

- Space-Track.org General Perturbations (GP) and GP History element sets — TLE, JSON, CSV, and XML; approximately 28,000 on-orbit objects with multi-year daily history; public registered-account access obtained by the offeror
- CelesTrak Satellite Catalog (SATCAT), orbit-class groupings, and supplemental General Perturbations (GP) data — open access, no registration
- European Space Agency (ESA) Database and Information System Characterizing Objects in Space (DISCOS) object, launch, and physical property record; registered access
- United Nations Register of Objects Launched into Outer Space and national launch records — open access; mission and operator attribution
- Open-source mission and payload record — mission and payload descriptions from space agencies such as National Aeronautics and Space Administration (NASA), the Union of Concerned Scientists Satellite Database (a static archive last updated in 2023), Notices to Space Operators (NOTSOs), and technical and academic literature; open access, with the publication date of each source recorded against the output it informs
- Commercial SDA catalogs and observation products — commercial license; obtained by the offeror or furnished through the Lab where a cohort agreement exists
- Synthetic orbit-determination data-set histories generated for objects and behaviors not represented in public record, produced by the offeror against a Government approved scenario definition

6.5. Operational Value

Turns the initial step of characterization from a lookup into a derived, scored, and reproducible product. An analyst receives orbit classification and object mission class with a stated confidence and a visible derivation, which allows a weak characterization to be recognized as weak immediately rather than discovered downstream in an assessment that has already been briefed. The published data model and ingest interface make the rest of the chain portable across changes in the underlying source data.

6.6. Desired Capability Outcomes / Acceptance Criteria

- Classifies orbit regime and orbit family for a high rate of agreement of a designated evaluation set drawn from the public catalog, with most of the agreement against the Government reference classification.
- Infers object class and mission class with most agreement with Government analyst validation and reports a separately calibrated confidence value for each.
- Maintains source and derivation traceability for every output, including (but not limited to):
 - Originating element set, epoch, and provenance for each state estimate
 - Derivation method and parameters used for each classification
 - Explanation, in operator-readable form, of why a classification was assigned and what would change it

- Reflects data age in every reported confidence value and demonstrates the confidence response to a deliberately stale input.
- Normalizes all inputs and outputs into a defined data model with single-meaning fields, published as an Interface Control Document (ICD).
- Demonstrates substitution of an externally produced, confidence-scored fused track for data-derived state through the ingest interface, with no change to the output contract consumed by Problem Statements 1.2 and 1.3, and with the incoming track's per-dimension confidence values and source lineage preserved and traceable in the outputs.
- Delivers the data model and ICDs, test procedures, classification performance results, and recommendations for operational transition.

7. Problem Statement 1.2 — Pattern of Life Baselines and Out-of-Family Deviation Detection

7.1. Problem Statement

PoL for a Resident Space Object exists as analyst familiarity rather than as a recorded, query able baseline. Existing maneuver detection reports that a change occurred against a general threshold, but not whether that change is normal for the object in question, how far outside its historical behavior it falls, or how significant it is relative to objects of the same class. Behavior that is out of family is therefore recognized inconsistently and only where an experienced analyst is watching.

7.2. Operational Context

This stage sits in TCPED Exploitation and supplies the temporal dimension the rest of the chain depends on. PoL is in scope for this Focus Area, and it is what allows an engagement to be characterized as unusual rather than merely geometric. An engagement between two objects that have always shared a regime is a different finding from the same geometry produced by an object that departed a two-year station-keeping pattern last week.

The consumers are the orbital analyst maintaining the catalog and the operations center crew assessing whether an observed change warrants attention. Both need the deviation stated as a magnitude against a baseline with a stated window, because both must justify the assessment to an echelon that did not observe the object's history.

7.3. Requirement

- **Pattern of Life Baseline:** Develop a capability that establishes a PoL baseline for a designated RSO from its historical orbital data-set record, and a companion baseline for the object's class. The baseline must characterize the parameters that describe how the object holds its orbit — station-keeping cadence and magnitude, semi-major axis and eccentricity maintenance, inclination control, longitude slot occupancy and drift, orbital period stability, repeat ground track behavior, and

maneuver frequency and magnitude distribution — and must state the time window it covers and the confidence attached to it.

- **Deviation Detection and Scoring:** Detect deviations in those parameters that fall outside the established envelope, classify each deviation by type and direction, estimate its magnitude including a velocity change estimate where derivable from the data-set record, and score its significance against both the object's own history and its class baseline.
- **Explainability and Replay:** Expose the parameters driving each significance score, make detection sensitivity a documented, adjustable setting whose effect on detection and false alarm rates is reported, and support replay of the detector across historical windows so a scored deviation can be reproduced and audited.
- **Re-Baselining:** Re-establish the baseline on a stated, adjustable interval and whenever a deviation recurs with enough regularity to constitute a new pattern: absorb the recurring behavior into the baseline, report each absorption rather than applying it silently so that behavior repeated deliberately to become routine remains visible, and retain the prior baseline for comparison.
- **Scope:** Attribution of intent, and detection of behavior deliberately engineered to remain within an object's baseline, including dormancy, mimicry of benign or debris behavior, and concealed maneuvering, are outside the scope of this Problem Statement. The capability establishes that behavior is out of family and by how much; what that behavior means in relation to other objects is characterized in Problem Statement 1.3, and the judgment of intent remains with the operator and the analyst.

7.4. Data Sources

- Multi-year GP element-set history from Space-Track.org for the designated evaluation set, at daily or better cadence
- Object and mission classification output from Problem Statement 1.1, bound to its published data model
- Publicly documented maneuver, relocation, and proximity-operations events with dates, used as labeled positive cases for replay
- Commercial ephemeris and owner-operator published orbit data where available under license, used as a higher-accuracy comparison against data-derived detection
- Synthetic maneuver and behavior-change injections at controlled magnitudes, produced against a government approved scenario definition; the offeror produces injections for development, and the injection set used for grading is withheld from the offeror

7.5. Operational Value

Moves PoL out of individual analyst memory and into a queryable baseline that any crew position can interrogate and converts maneuver detection from a binary alert into a scored statement of how far the object has departed from its own history. This lets an operations center distinguish a routine station-

keeping burn from a change in behavior without needing the one analyst who has followed the object for years.

7.6. Desired Capability Outcomes / Acceptance Criteria

- Establishes a behavioral baseline for each object in the evaluation set from relevant history, stating the window, the parameters covered, and the confidence in the baseline.
- Detects injected out-of-family behavior changes in a withheld synthetic test set with a high detection rate at a false alarm rate per object-day no greater than a government set ceiling stated in advance of evaluation, and reports detection rate across the full range of operating thresholds.
- Classifies each detected deviation by type and direction and estimates its magnitude, including a velocity change estimate where derivable, with stated error bounds.
- Scores deviation significance against both the object's own history and its class baseline and exposes the parameters driving the score in operator-readable form.
- Reproduces at least three publicly documented historical behavior changes in replay and reports the interval between initial detectable deviation and the publicly reported event and demonstrates re-baselining on at least one case in which recurring behavior is initially scored as out of family and later absorbed into the baseline, with the absorption reported and the prior baseline retained.
- Provides detection sensitivity as a documented, adjustable setting, reports the effect of a change on detection and false alarm rates before it is applied, and retains an auditable record of the setting in force for any scored deviation.
- Delivers baseline model documentation, detector test procedures, detection and false alarm performance results, and recommendations for operational transition.

8. Problem Statement 1.3 — Engagement Geometry, Risk, and Threat Characterization

8.1. Problem Statement

Relationship analysis between orbits stops at pairwise conjunction screening, which answers whether two objects will pass within a volume and not whether two objects are converging on the same object of interest, the smallest structure in which coordinated behavior can be seen. Determination of which engagements are risky to space operations is consequently manual, inconsistent between analysts and between shifts, and produced without a recorded confidence or an inspectable rationale.

8.2. Operational Context

This is the terminal stage of the chain and the point at which characterization becomes decision relevant. It supports Find, Fix, and Track in the Joint Targeting Cycle and occupies the Orient function feeding Decide in the OODA loop, and its products feed operations center assessment during Space Tasking Cycle execution. The consumers are the NSDC and CSPOC watch floors and the S4S staff, who

must state not only that a relationship exists but what kind of relationship it is and how much of the assessment to rely upon.

The stage is also where the structure of the question changes. Conjunction screening treats the domain as a set of pairs and reports each pair independently. Operationally significant behavior is frequently not pairwise: two objects arriving at the same object of interest, from different orbits or on different timelines, is a single finding that no pairwise product expresses, and reconstructing it from separate pair results is exactly the manual work this Problem Statement removes.

8.3. Requirement

- **Engagement Identification:** Develop a capability that compares a designated RSO against the catalog to identify orbits that intersect, approach, co-orbit, share a regime or slot, or are reachable within a stated velocity-change and time budget, and that produces a time-phased engagement set describing closest approach, relative velocity, access and illumination windows, reachability envelopes, and persistence of the relationship.
- **Scale and Degraded Input:** The comparison must operate at full public catalog scale and must degrade gracefully and visibly when data sets for participating objects are stale or absent.
- **Engagement Structure:** Characterize the structure of each engagement as one-on-one or two-on-one. The structure must be an analytic output derived from the data rather than a mode the operator selects in advance, and the capability must present the two-on-one case as a single finding rather than as two unrelated pairwise results. Detection of correlated behavior across groups larger than two-on-one is deferred to future development and is not required here; the capability should not preclude it, but it will not be evaluated against it.
- **Risk Determination and Scoring:** Determine and score the risk each engagement presents to space operations using algorithms that combine engagement geometry from this Problem Statement, deviation significance from Problem Statement 1.2, and object and mission class from Problem Statement 1.1. Produce an enumerated set of risk factors together with a scored characterization that distinguishes routine conjunction, debris hazard, benign co-location or proximity, close approach and inspection behavior, and behavior threatening to space operations.
- **Confidence, Explanation, and Traceability:** Each output must carry a separate confidence value for each output dimension it draws on, including those inherited from Problem Statements 1.1 and 1.2, an operator-readable explanation naming the factors that drove the result and their relative contribution, and traceability to the originating data set, deviation event, and geometry computation.
- **Algorithm Transparency:** The algorithms themselves must be documented and inspectable, and any learned weighting must be recorded, reviewable, attributable, and reversible.
- **Retention:** Characterized engagements are retained and queryable by owner-operator, so that later tactics, techniques, and procedures (TTPs) analysis can draw on them without reprocessing.
- **Confidence Gaps:** Where confidence is insufficient to characterize an engagement, the capability states the gap, the object, and the dimension responsible as an element of the published data

model, so that a subsequent collection decision can consume it without reinterpretation; making that decision and generating the resulting tasking are outside the scope of this Focus Area.

- **Scope:** Declaration of adversary intent, capabilities employment analysis, defensive course-of-action recommendation, sensor tasking, detection of group structure beyond two-on-one, and profiling of operator TTPs are outside the scope of this Problem Statement.

8.4. Data Sources

- Characterized object, orbit classification, and confidence output from Problem Statement 1.1, bound to its published data model
- PoL baselines and scored deviation events from Problem Statement 1.2
- Full public catalog data sets from Space-Track.org for screening at scale, approximately 28,000 on-orbit objects
- Conjunction Data Messages and public conjunction screening products, including CelesTrak Satellite Orbit Conjunction Reports Assessing Threatening Encounters in Space (SOCRATES), in standard Consultative Committee for Space Data Systems (CCSDS) format
- Publicly documented rendezvous and proximity-operations, relocation, and co-location events, used as labeled cases
- Synthetic scenarios covering one-on-one and two-on-one structures at controlled difficulty, produced by the offeror against a government approved scenario definition, used to fill gaps in the documented record rather than to carry the evaluation
- Government analyst adjudication labels for a designated scenario set covering routine, ambiguous, and threatening cases, with more than one analyst labeling each case the documented record does not resolve, withheld from the offeror during development and furnished by the Government for scoring

8.5. Operational Value

Gives an operations center a characterized engagement instead of a list of close approaches, and states how much of that characterization to trust and why. It expresses the two-on-one case that current screening products cannot, which is where coordinated behavior of operational concern initially becomes visible, and it compresses an assessment that today takes hours of manual reconciliation into a product available inside the decision window it is meant to support.

8.6. Desired Capability Outcomes / Acceptance Criteria

- Screens a designated RSO against a full public catalog of at least 25,000 objects and returns the engagement set within 10 minutes on a government specified reference computing configuration stated in advance of evaluation, with the timing reported and repeatable.

- Characterizes engagement structure as one-on-one or two-on-one as an analytic output and demonstrates both structures on a government approved scenario set drawn predominantly from documented real events.
- Produces an enumerated risk-factor list and a scored risk characterization for every engagement, with a separate confidence value reported for each output dimension it draws on.
- Achieves a high rate of agreement with the resolved outcome on scenario cases where the documented record establishes what occurred, and a similarly high level of agreement with government analyst adjudication on cases the record does not resolve; reports agreement among the adjudicating analysts on those cases as the reference ceiling and reports every disagreement with the factors that produced it.
- Exposes the reasoning behind each characterization, including:
 - The factors that drove the result and their relative contribution, in operator-readable form
 - Traceability to the originating data set, deviation event, and geometry computation
 - An auditable record of any learned weighting in force, with the means to review, attribute, and reverse it
- Supports replay against historical and synthetic one-on-one and two-on-one events and demonstrates behavior under degraded input where data sets for participating objects are stale or absent.
- Delivers algorithm documentation, ICDs, scenario and scoring results, and recommendations for operational transition.