Space Traffic Management Performance for Large Constellations – A Self-reporting Aggregated Realized Risk Approach from the OneWeb Constellation

Pablo Minguijon Pallas[1], **Calum Hervieu**[1], **and Ignacio Dominguez-Adame Palomo**[1]
[1]*Flight Dynamics Engineering, Eutelsat Group*

29 August 2025

ABSTRACT

As space traffic congestion increases, academia, industry, and regulatory bodies are working from different perspectives to ensure the safety and sustainability of space operations. Broadly, safety and sustainability assessments rely on metrics that can be categorized as either ”predictive” (i.e., events that may occur) or ”realized” (i.e., events that have already occurred).

One often overlooked ”realized” behavior is the quantification of a space operator’s effective and practical ability to manage collision risk while in orbit. This type of risk can be looked at from two angles: (a) risk that has been actively or passively reduced or eliminated (i.e., abated), and (b) risk that has been intentionally or unintentionally ”accepted” (i.e., residual). In given circumstances, and for an increasingly congested space environment, accepted residual risk could become a non-negligible contributor to both the overall risk that deployed systems undergo and the impact that those systems pose to the space environment.

This paper presents methods for Owner/Operators (O/Os) to monitor their operational residual risk and use it as a benchmark with a dual purpose: (a) to explicitly expose and quantify the ability of an O/O’s system (ground- or space-based) to achieve safe navigation in their orbital neighborhood, and (b) to support a two-fold feedback loop, an internal one for the business to tweak their risk posture and an external one for regulatory agencies to verify that systems are safely operated. This discussion opens the question of how to expose relevant information to the appropriate stakeholders. Finally, this paper present and propose some reporting methods as operationally used and in development by the Eutelsat’s LEO Constellation (formerly OneWeb).

ACRONYMS

Δv Delta-v
19th SDS 19th Space Defence Squadron
CA Conjunction Assessment
CARA Conjunction Assessment Risk Analysis
CDM Conjunction Data Message
CSpOC Combined Space Operations Center
FDS Flight Dynamics System
Gen1 1st Generation
HBR Hard Body Radius
ICC Intra-Constellation Conjunction
LEO Low Earth Orbit
NASA National Aeronautics and Space Administration
O/O Owner/Operator
Pc Probability of Collision
rCDM Representative CDM
RMM Risk Mitigation Maneuver

SOC Satellite Operations Centre

SSA Space Situational Awareness

STM Space Traffic Management

t2TCA Time to Time of Closest Approach

TCA Time of Closest Approach

1. INTRODUCTION

The rapid growth of activity in low Earth orbit has turned (Space Traffic Management (STM)) into a daily operational challenge rather than a licensing formality. Most frameworks used for design review and authorization emphasize *predictive* metrics for non-maneuverable objects or non-trackable debris —e.g., model-heavy collision probabilities and debris-generation analyses—produced before launch. These are necessary, but they do not fully capture how systems actually behave once deployed: a design can be "safe on paper" yet operated in ways that accumulate nontrivial collision risk. To successfully foster space safety, operational STM practices and performance should get a more standardized and explicit exposure.

This paper focuses on *realized, operational* metrics, with particular attention to the risk an operator *accepts* at a concrete decision point. We distinguish three related concepts: (i) *abated risk*, the reduction achieved through active or passive measures; (ii) *accepted risk*, the level of risk intentionally or unavoidably carried at the chosen decision horizon (relative to Time of Closest Approach (TCA)) based on information available at that time; and (iii) *residual risk*, a retrospective view using the best information available after the event. In congested regimes, the aggregation of accepted risk across a large, maneuverable fleet is not negligible when compared to industry standards for single events, and can meaningfully inform both operator practice and ecosystem sustainability.

Why high-level metrics matter for innovation. Elevating accepted risk to a first-class, outcome-based metric enables a shift from prescriptive "how-to" rules toward *performance-based oversight*. Rather than mandating specific thresholds, maneuver strategies, or software architectures, operators can pursue diverse designs—higher autonomy, novel guidance and control, differentiated ground segments, and alternative concepts of operations—so long as their *observed* accepted risk remains within declared bounds. This expands the design space for innovation while preserving accountability. It also creates a direct business incentive to invest in their STM systems, estimation quality (e.g., covariance realism, multi-source Space Situational Awareness (SSA) fusion) and decision latency (e.g., faster screening and approval loops), because improvements are immediately reflected in the metric. Finally, outcome metrics support safer experimentation: operators can trial new thresholds or maneuver policies under guardrails (risk budgets, escalation triggers), tightening or relaxing constraints as empirical performance warrants.

Within the Eutelsat Group operational context for the OneWeb Low Earth Orbit (LEO) constellation, we introduce an empirical method to monitor accepted risk using historical conjunction data and operator decision points. The approach consolidates multi-provider Conjunction Data Messages (CDMs) into a representative Representative CDM (rCDM), recomputes Probability of Collision (Pc) internally under consistent assumptions, and evaluates risk at a fixed relative time-to-TCA to reflect what was actionable in operations. The method is designed to be (a) *operator-facing*, as a benchmark and feedback loop for tuning thresholds, timelines, and maneuver policies; and (b) *stakeholder-facing*, as a basis for transparent reporting of how well a system performs in its environment.

Contributions. Specifically, this paper:

1. Formalizes an *accepted-risk* metric at a chosen decision horizon and shows how to aggregate it across time and fleet scale for an actively maneuverable constellation. It includes risk from both mitigated and unmitigated events.
2. Describes an operational pipeline that selects an rCDM from multiple SSA sources, recomputes Pc with internal models, and flags late-notice "pop-up" events that challenge decision timelines.
3. Applies the method to historical data from the OneWeb fleet (service orbit) to characterize empirical patterns: tail-driven daily risk, the nonzero contribution of maneuverable systems, and the utility of aggregated views (daily/monthly/yearly) for management.
4. Proposes reporting formats that expose accepted risk to internal and external stakeholders while acknowledging current limitations.

Scope and limitations. The metric depends on assumptions and data quality: hard-body radius choices, covariance realism, secondary-object tracking quality, provider latencies, and geometric criteria all influence Pc estimates and their timely availability and, consequently, any aggregation at a given point. Known fundamental Pc definition limitations impact the absolute meaning behind the aggregation. We therefore treat accepted risk primarily as a *trending and*

benchmarking indicator—useful for guiding operations and transparency—not as a frequentist predictor of collision counts nor as an immediately comparable cross-operator league table. Future work should aim to address some of these limitations.

Paper organization. Section 2 summarizes the Eutelsat Group STM concept of operations and reaction thresholds. Section 3 defines aggregation across abated, accepted, and residual risk and clarifies the data availability and decision-horizon framing. Section 4 presents empirical results and discusses drivers, including late-notice events. Section 5 concludes with implications and next steps.

2. SPACE TRAFFIC MANAGEMENT AT EUTELSAT GROUP

The Eutelsat Group maintains two highly-automated Satellite Operations Centres (SOCs) in Virginia, U.S, and London, U.K., responsible for operating the OneWeb LEO fleet.

2.1 Overview

At the highest level, the operational model of Eutelsat Group heavily adheres to a set of orbital safety best practices[1, 2], whereby the system shall at all times:

- Maintain quality Operator/Operator predicted ephemerides and spacecraft status information and update this information regularly with the relevant Conjunction Assessment (CA) screening authority;
- Perform CA risk assessment to identify high-risk conjunctions that require mitigation;
- Pursue adequate mitigation actions to avoid such high-risk conjunctions.

A key tenet for robust STM is transparency between the involved actors and the appropriate screening authority. This includes sharing information, where relevant, regarding applicable risk thresholds, claiming maneuver responsibility, delivering updated ephemeris with realistic covariance and modeled mitigation maneuvers, typical reaction timelines, satellite state of health, operator contact information, and any necessary mission constraints.

2.2 Conjunction Assessment

Eutelsat Group utilitzes multiple SSA providers, including the 19th Space Defence Squadron (19th SDS) and LeoLabs, to provide regular ephemerides screening. With the current full 1st Generation (Gen1) fleet at service altitude, these providers typically generate 100,000 CDMs per day relevant to the OneWeb LEO fleet. Upon publication, these CDMs are automatically ingested and associated to conjunction events before being rapidly recomputed with the most up to date ephemeris available in our system. Alongside external CDMs, the OneWeb Flight Dynamics System (FDS) triggers an internal *one vs all* screening to identify Intra-Constellation Conjunction (ICC) events the moment a new ephemeris is published for any vehicle. These internal CDMs are grouped and processed into the same CA flow as external events.

In order to consolidate event data from multiple sources, Eutelsat's OneWeb fleet utilizes the concept of a rCDM - a dataset that is considered the most trustworthy representation of a conjunction event at a given instant in time.

The rCDM selection process is flexible and modular, allowing for highly configurable, preferential selection based on various CDM criteria. These include CDM creation age, provider, secondary observation/ephemeris age and quality, as well as preferences tuned for selected operators or even individual secondary objects. It is the rCDM that ultimately drives the reaction logic for a given conjunction. In practice, the choice of rCDM is a proxy for the choice of secondary information to be used for a given conjunction.

2.3 Reaction Thresholds

There are three primary thresholds that will trigger mitigation events in the Eutelsat FDS that rely on the industry standard, internally-recomputed Pc metric. These thresholds determine actionability based on severity & orbital geometry, but are only mitigated within certain Time to Time of Closest Approach (t2TCA) criteria; typically 48 hours.

In cases where Pc is not defined, such as those where an Owner/Operator (O/O) has published an ephemeris without covariance, there is a fail-over criterion that was derived in coordination with National Aeronautics and Space Administration (NASA) Conjunction Assessment Risk Analysis (CARA) that considers event geometry.

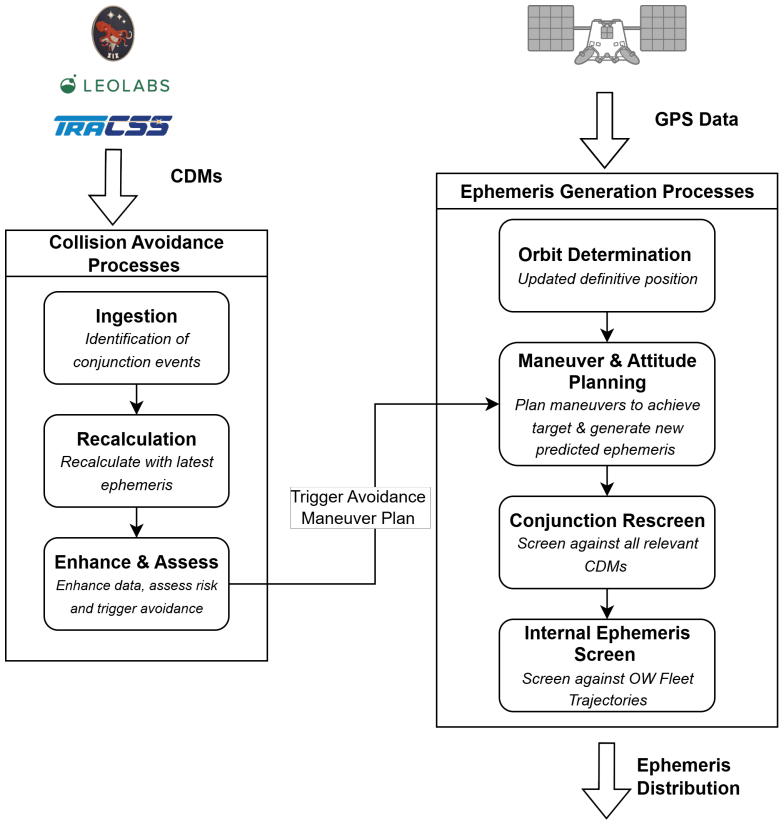

Fig. 1: Overview of automated Collision Avoidance Process for the OneWeb fleet.

Table 1: Reaction thresholds based on recomputed Pc.

Pc		Miss Distance [m]		Radial [m]		Intrack [m]		Crosstrack [m]
$>$ 2e-5	OR	$<$ 100						
$>$ 1e-6	AND	$<$ 2000	AND	$<$ 100				
				$<$ 50	AND	$<$ 500	AND	$<$ 500

Due to the high levels of automation involved, the FDS also respects *reaction lockouts*, whereby the system is able to understand and verify the current state of a conjunction event based on agreements in place concerning maneuver responsibility with another O/O.

2.4 Conjunction Mitigation

If a threshold is breached, the system will begin automatic planning of mitigation maneuvers. When changing trajectory at short notice, Eutelsat Group utilizes candidate maneuver screenings, which are submitted as *special* ephemeris to the 19th SDS. This allows the system to plan multiple responses to a concerning conjunction, and choose which to execute preferentially e.g. the lowest aggregate Pc, the least Delta-v (Δv) required, the minimally invasive to mission goals, or a combination therein. To avoid confusion, these candidate maneuvers are not visible to other O/O, but will be published as operational the moment a maneuver plan is chosen.

When generating candidate plans, the FDS will trigger an internal recalculation of all known conjunction events for that given satellite, and another *one vs all* to internally validate mitigation maneuvers, before requesting final clearance from the appropriate SSA providers. All candidate plans are derived to minimize Pc within a set of mission boundaries; such as station-keeping requirements, Δv budget, operational restrictions, etc.

The number of events requiring mitigation varies with each mission phase, due to various factors, such as higher debris density and number of O/O at lower altitudes, as well as physical effects on prediction accuracy. At the OneWeb service altitude ($\approx 1,200$ km), Eutelsat Group executes approximately 150 Risk Mitigation Maneuvers (RMMs) per month

across the fleet.

3. RISK FOR A LARGE MANEUVERING FLEET

There are two main dimensions to the problem of realized risk for a large constellation: fleet scale and time. This section will try to cover how the authors have approached these two dimensions within this work.

Note on purpose. Throughout this paper, "risk" refers to a specific, operationally grounded construct used for benchmarking and trending within an actively maneuvered LEO constellation. The intent is to make observed operator behavior measurable; it is *not* to provide a frequentist prediction of collision counts.

Note on *relevant* risk. All risk metrics mentioned above are implicitly derived from relevant risk notification by SSA providers. It is an obvious caveat to this work, that O/Os can only meaningfully track risk of which they are aware, and that is notified within their satellite & system's operational envelope. As the SSA landscape across LEO rapidly evolves, the authors expect this to substantially impact the observed Accepted Risk metric.

3.1 Aggregated Risk - Abated vs Accepted vs Residual Risks

3.1.1 Aggregated Risk

For a large fleet of satellites of a few orders of magnitude over the unit, the relation between the per-satellite risk management and the overall risk that the constellation faces stops being an easily inferred parameter from the choice of per-conjunction reaction thresholds. The RMM thresholds, which have historically been used as a proxy for system-wide collision risk management, fall short, and O/O require higher level of abstractions to enable effective system-wide risk management.

As an operator, one is typically not only interested in avoiding a collision with any one satellite individually, but on not having a single collision among one's many satellites. Some form of aggregation is required to create that metric and expose the right level of information to the relevant internal and external stakeholders.

Assuming a constellation faces a series of N Conjunctions, C_i, between two given times, TCA_n and TCA_{n+1}, each with a Pc as observed at t_j of $pc_i^{t_j}$, and assuming total independence between these events, then the total aggregated risk as of information available at time t_j can be computed as:

$$PC_{constellation}^{t_j} = 1 - \prod_{i=1}^{N} (1 - pc_i^{t_j})$$

Where $PC_{constellation}^{t_j}$ can be read as the probability that one or more events as observed at t_j results in a collision. The addition of the time of observation t_j is relevant because of the empirical nature of the proposed metric. Note that C_i aims to describe all conjunctions encountered by the constellation between TCA_n and TCA_{n+1}, regardless of their perceived risk or whether this risk is pre or post mitigation.

Despite all the shortcomings of the industry-standard Pc, and discussed in previous works [3, 4] and the results of this work not necessarily representing expected frequencies, we believe there is physical meaning behind the magnitudes presented and that they can be used, if not on an strict frequentist sense, at least as a trending/benchmarking metric. Future work will focus on refining the underlying metric so that aggregation can lead to more meaningful results in an absolute sense.

3.1.2 Abated vs Accepted vs Residual Risk

Assessing the realized safety of flight can be seen as the combination of two performance indicators that address the temporal component in different ways:

- **Abated risk:** This is the risk that has evolved (hopefully reduced) as the conjunction evolved over time. Exactly how this metric is measured is contentious as one can take into account very different contributors to the reduction of risk. Depending on what contributors are considered, this metric can have a very different significance.

In the most general sense and taking into account the full risk variation as a function of time as best known for the conjunction, the metric highlights the ability of a system to reduce risk within its environment. Because of its very definition, aggregating abated risk and the meaning of that aggregated metric is also not trivial. Because of the complexity behind this metric, abated risk has been left out of this work, although the team continues working in this direction and recognizes its synergy with the overall goal behind the proposal.

- **Accepted/Residual risk:** This risk represents the state of a conjunction at a given point in time. It proves useful to understand, in absolute terms, the risk that particular conjunction represented as of information available at that epoch. This type of risk is the main focus of this work and a more lengthy discussion on it, and the differences between Residual and Accepted flavors will be expanded below.

Previous work has focused its attention on the concept of Residual Risk [5, 6, 7], which focuses on the ”final” risk of a conjunction after any applicable RMM has been performed. This is to say, it assumes knowledge on the actual risk that a conjunction event actually posed.

Accepted risk acknowledges that the ”final/real risk” of a past conjunction is a controversial topic, and instead focuses on the operationally ”witnessed” risk at a given decision point. If you align this decision point with the last point where your system can reliably react to high risk events, it becomes an explicit and direct measure of how safe a system operates on its orbital environment - with all the nuances that come with real life conditions embedded in the assessment. Your overall system design, concept of operations and environment in which they operate are making you ”accept” that risk.

Even if we accept the shortcomings of the currently discussed metric and agree that there is no absolute or frequentist meaning behind it, we argue that it can still become a useful trend/predictor and can help operators and other stakeholders monitor their risk-exposure emanating from their fleet’s STM given their current spacecraft, ground system, environment and operational practices. And what is most important - enable them to potentially do something about it and benchmark the improvements.

The team acknowledges the interest and potential behind a ”realized” metric like Residual risk and is working on blending in the three risk metrics described in this section together for a more comprehensive characterisation of a system’s performance in its environment. However, for the remainder of the work, the discussion will focus on Accepted risk as it is admittedly a more humble first take on the matter.

3.2 Risk Evaluation Point - Analysis Data Availability

For a given risk analysis window between two arbitrary absolute times TCA_n to TCA_{n+1}, Figure 2 tries to depict three different data availability strategies (i.e., possible definitions for t_j). The observer simulates a relative or absolute time with respect to the conjunction C_i TCA and the information that was available for that event at that observation time (t_j) is the one considered in the analysis.

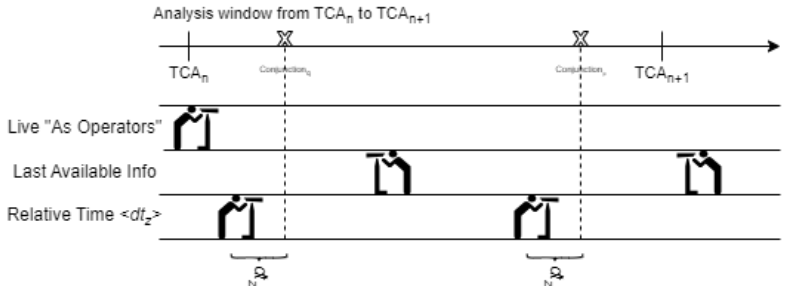

Fig. 2: Risk Evaluation Point. Different ”points of view” to Historic Conjunction Analysis. Observer in relation to existing information made available to the analysis.

- **Live / ”Operational”:** This view tries to mimic what is observed operationally at a given absolute time. This is particularly useful when trying to recreate how events were observed by the operational team and how they evolved over time.

- **Last Available Information:** This view tries to make use of the information that typically is considered to best characterize the event. This includes information delivered after TCA is passed. This is particularly useful when trying to determine ”real” risk of a conjunction as accurately as possible. The concept of ”residual risk” makes use of this approach.
- **Relative Time dt_z:** This view is a combination of the previous two. The risk is assessed from a relative distance to the event's TCA which simulates an ”information cut off” on the decision making. This is useful to analyze risk as typically addressed operationally in an ”accepted risk” assessment.

In a ”relative time” analysis, risk that is only notified later than your assessment epoch is completely disregarded. Even though this is consistent with the definition of the ”accepted risk” concept that is the main focus of this work, it is a short coming when the intention becomes to try to characterize absolute risk endured by a maneuvering constellation under lethal and trackable debris. A residual risk approach does a better job at including this risk where information is considered even if it was available after the event's TCA. However, known inconsistencies in post-event data delivery by SSA providers puts at risk the ability to confidently characterize this risk.

3.2.1 Note on last minute risk events - Pop Up conjunctions

A sensitivity analysis between different risk assessment relative times and the ”last available” strategy highlighted that last-minute/pop-up events play a substantial role in the total risk the constellation undergoes and has to operationally manage. This sensitivity analysis is excluded from this work, but will be explored further in future.

Due to these findings, and their expected higher importance for constellations flying at lower altitudes where state knowledge uncertainties of primary and secondary are much larger, these events have been loosely identified and highlighted in the results presented in the next section.

In this work, this potential shadowing of ”last-minute/pop-up” risk from the ”relative time” approach has been partially addressed by choosing an assessment relative time of 0 hours ahead of TCA. This is relatively consistent with OneWeb's platform and ground segment high agility, where the system can successfully execute remediation of high risk events made available to us less than 2 hours from TCA. Future work should focus on a more consistent translation between operational capabilities and analysis relative time.

4. ACCEPTED RISK MONITORING FOR A LARGE CONSTELLATION

In this section, we apply the empirical approach described previously to historical CDM data used operationally by the Eutelsat's OneWeb constellation over the period from July 2024 to June 2025, both included.

4.1 Scope limits and threats to validity

- **Probability of Collision inherent limitations.** As argued in literature [3, 4], Pc suffers from some conceptual issues that affect its typically assumed frequentist sense, and these propagate upwards in the paper as aggregated metrics build upon it.
- **Probability of Collision computation.** All Pc values shown throughout the paper have been internally computed by the OneWeb ground segment algorithms and do not represent originalCombined Space Operations Center (CSpOC) CDM values.
- **Aggregation meaning.** Summing Pc across encounters, vehicles, and time is a useful *indicator* of exposure but is *not* a statement of expected collision frequency. Events are not guaranteed independent, and Pc is not a calibrated rate [2].
- **Decision-horizon dependence.** Accepted risk excludes information that arrives after the chosen dt_z; late-notice (“pop-up”) events therefore contribute only when detected by dt_z. Where identifiable, such events are flagged but not retroactively added.
- **Modeling sensitivity.** Results are sensitive to Hard Body Radius (HBR) choice, covariance realism, force modeling, process-noise tuning, secondary ephemeris age/quality, provider latencies, and internal recomputation

Table 2: Filter and parameters applied in the analysis

Parameter	Value	Description
t2TCA	0 hrs	Data displayed corresponds to the closest CDM available before TCA
CDM Source	CSpOC	All CDMs used come from CSpOC and were retrieved via Space-Track.org
Primary HBR	6.7 m	All Pcs computed with a OneWeb HBR of 6.7 m, which is ≈ 3 times larger than reality, as a passive safety margin. For most conjunctions this can be rougly considered as an inflation of Pc by up to 1 order of magnitude in all Pc formulations (Pc $\propto$ combined HBR^2)[8].
Pc	Internally recalculated with 2D-Pc method	All Pcs used throughout the analysis have been internally recomputed using OneWeb's software libraries, latest available as of dt_z OneWeb satellite ephemerides and CDM secondary information (state, covariance, HBR.
Tracking Type	Operational & HAC	CDMs include both ephemeris-based and radar-based tracking of both primary and secondaries
Data age policy	12 hours	Conjunction is considered clear if the rCDM (HAC or Operational) are older than 12 hrs from t2TCA assessment point.
Fleet subset	Service altitude fleet (630-650 vehicles)	Only Eutelsat's OneWeb satellites at service altitude (≈ 1200 km) have been considered
Pop Up t2TCA	< 12 hrs	Conjunctions that become high risk events within 12 hours of TCA are classified as "Pop Up" / last minute

settings. Consequently, absolute magnitudes should be interpreted with caution. In particular the HBR choice is intentionally ≈ 3 times larger than physically real, which substantially affects recomputed Pc results.

- **Metric choice.** This paper emphasizes *accepted* risk. *Residual* risk (using last-available/post-TCA information) and *abated* risk (risk reduction over an event's evolution) are discussed conceptually but not used for the main aggregation unless stated.
- **Comparability.** As of current maturity level, values are intended for within-operator benchmarking and transparency. Cross-operator comparisons require harmonization of dt_z, HBR, altitude/regime, CDM/SSA source policies, covariance-quality criteria, and filtering rules.

Practical interpretation. The reported risk should be treated as an *outcome metric* for operational performance—useful for trend detection, risk budgeting, and evaluating the effect of thresholds, timelines, and maneuver policies—rather than as a direct predictor of collision counts or a league table across operators. Future work should aim to address some of these limitations.

4.2 Results

The data and analysis presented has been pre-processed and filtered according to the parameters discussed in previous chapters and are summarized in Table 2. After filtering, the analysis processed around 12 million unique CDM for the entire July 2024 - June 2025 span.

As an example, Figure 3 shows the daily Accepted Risk for the 30 days in June 2025. This is a dashboard similar to what we are currently using internally to monitor the metric.

As a reminder, all the information presented reflects the available data as of the analysis time (i.e., 0 hours to TCA).

- Each cross represents an individual conjunction with its corresponding risk and classification. Thick blue crosses represent pop up/last minute high-risk events.
- Each grid square displays by the color shade an aggregated amount of risk for given 24 hours and risk bin. This is helpful to understand what areas of the plot drive daily aggregation.

- The red continuous line represents the daily aggregated accepted risk.
- Total monthly risks are provided for the aggregated daily risks in the legend. These values represent risks for the month of June 2025 (as plotted).

Some interesting notes for this particular month in display:

- Number of low risk events does not drive the total risk. This is an interesting take away that could look very different for a constellation flying in a different environment or with a different ability to react to high risk events. For example, a constellation flying in a busier environment, at lower altitudes (i.e. larger covariances) and with less reaction agility (i.e. accepted risk further out from TCA) could find its accepted risk dominated by a majority of "medium" risk events. The operator might then want to consider adopting a finite risk mitigation strategy as discussed in previous works [9].
- Daily risk is anchored and somewhat driven by a few worst-case conjunctions. Where the overall risk is low, it can be seen as driven by the aggregation of multiple medium-risk events.
- For this month, pop-up events contribute to $\approx 5\%$ of the risk accepted. This value does vary substantially from month to month ($\approx 5\% - 50\%$). This is expected to be a major contributor to constellations flying in more congested and uncertain environments. Considering how much of this is driven by secondary tracking quality, tracking the incidence of these events and their relative risk contribution could be used as a decision factor for O/Os to invest in improved SSA services.
- Reported Pop Ups numbers or classification does not match with displayed values because:
  - Pop Ups can be remediated ahead of analysis time, lowering their risk below the 1e-10 plot lower bound.
  - Pop Ups can be triggered by thresholds that do not rely on risk, since Eutelsat's OneWeb fleet has several non-Pc driven RMM thresholds.

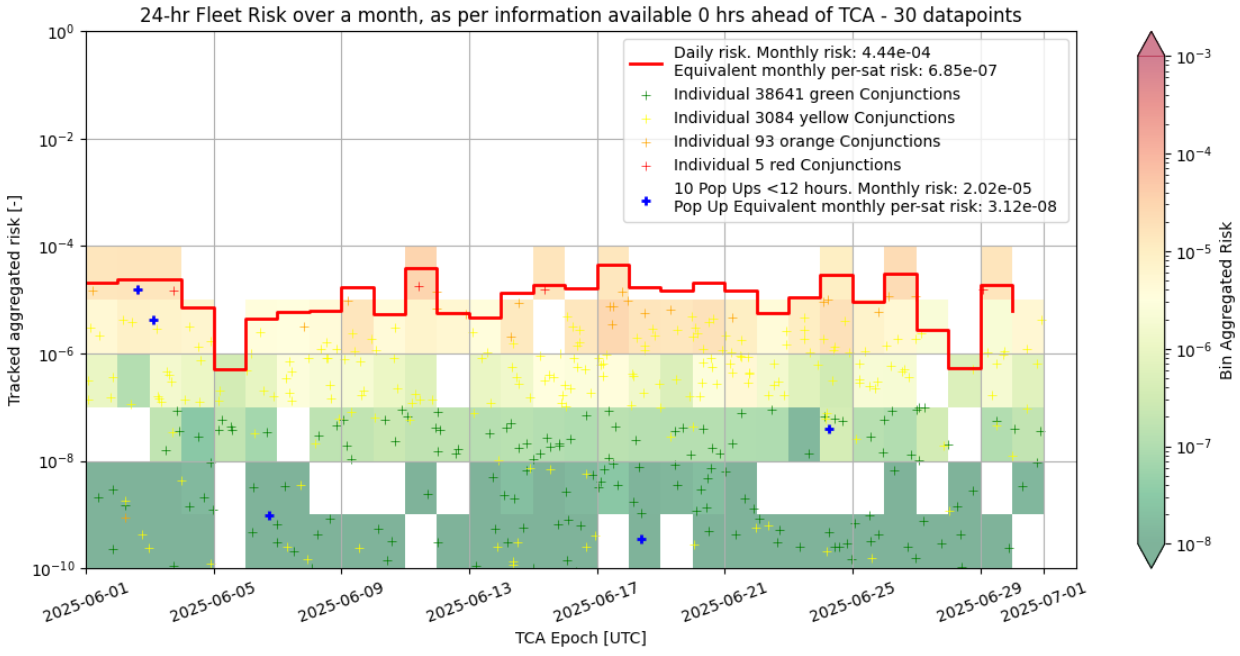

Fig. 3: Daily Accepted Risk overview for the 30 days spanning June 2025 - July 2025 for the service OneWeb fleet. Crosses represent individual conjunction events.

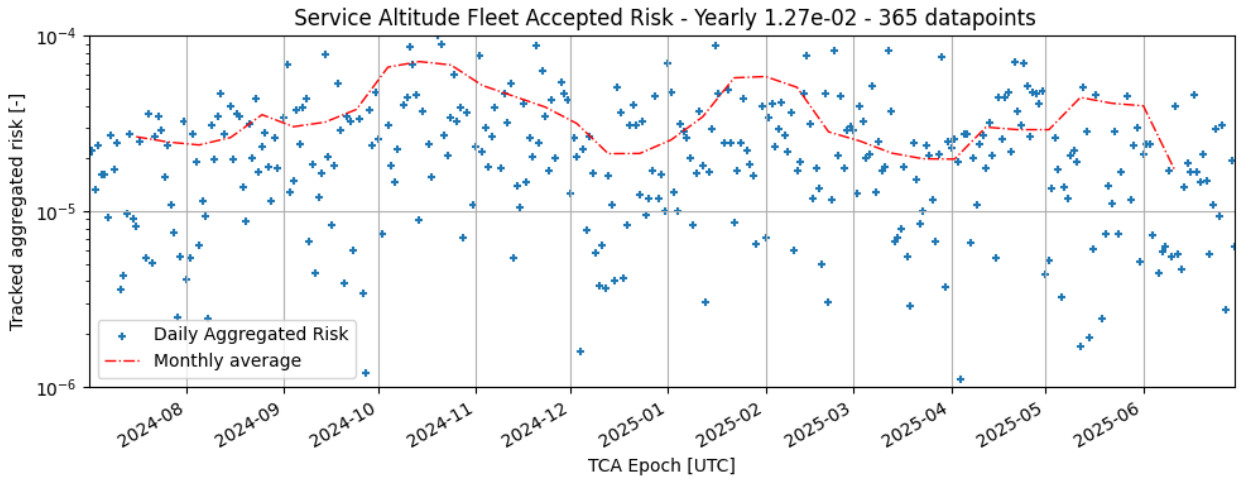

Fig. 4: Yearly view of Daily Accepted Risk experienced by Eutelsat's OneWeb Fleet. Note likely over-inflation of Pc by ˜1 order of magnitude due to HBR used. Yearly aggregated risk amounts to 1.27e-02

Figure 4 shows a high level overview of the yearly risk faced by Eutelsat's OneWeb constellation from an "Accepted Risk" sense. This plot has been created from the sequence of same data that was shown in Figure 3 for the entire July 2024 - June 2025 span.

Some observations can be made:

- The fleet-wide, yearly *aggregated accepted risk* is high relative to typical single-event risk tolerances used across the industry.
  - Even with the conservative HBR and the resulting ∼one–order-of-magnitude oversizing of Pc, the aggregated accepted risk is of the same order as mission-level risk thresholds commonly used for design and sizing.
- The yearly accepted risk *per satellite* is roughly aligned with the single-event RMM criterion (1.27e-2/640 ˜2e-5). This is likely contingent on this fleet's size, RMM thresholds, reaction timelines, and operating environment rather than a generalizable property.
- The time series is relatively smooth and bounded. For sufficiently large fleets, an aggregated accepted-risk metric can serve as a practical indicator of higher-level STM performance drivers (e.g., space weather, neighborhood changes, special maneuvering activities, system upgrades).
  - A rigorous attribution of fluctuations to specific external or internal drivers is out of scope here. Future work should quantify sensitivities, enabling data-driven decisions on investments in SSA providers or STM technologies, space/ground segment upgrades, and the appetite for specific maneuver activities.
- In data-rich operations where high-level metric trending and monitoring is key, large deviations in the monitored risk could act as symptoms that help triage attention to underlying causes (e.g., new co-locations, clusters of conjunctions, new space neighbours, space-weather events).

Until current limitations are addressed, accepted risk should be treated as a *within-operator trending* metric rather than an absolute cross-constellation comparator. Normalizing by exposure (e.g., per-satellite-day, per 1,000 CDMs, conjunction load, others) might be a practical first step toward harmonization; fair comparisons additionally require alignment on decision horizon, HBR, altitude/regime, CDM/SSA source policies, and covariance-quality criteria. Sensitivity of the metric vs O/O available controls in operational scenarios should be explored further [6].

We argue than even if we accept limited statistical meaning in the aggregate presented, O/Os of large constellations would be leaving a large risk management gap in their operational concepts by not looking at some form of aggregation and sticking to a standard per-event management.

A consequence-weighted variant (e.g., mass or kinetic-energy proxy) could be a natural extension to better align accepted risk with accepted environmental impact [10, 7].

5. CONCLUSION

This paper introduced an operationally grounded method to monitor *accepted collision risk* for a large, maneuverable LEO fleet. By evaluating risk at a defined decision horizon and aggregating it across vehicles and time, we provided a practical lens on how a system actually performs in its environment. The empirical results show that (i) the contribution of maneuverable fleets to overall collision risk is not zero; (ii) daily exposure is often dominated by a small number of high-severity encounters; and (iii) aggregated, time-series views can reveal meaningful trends relevant to operations, business risk management, resource allocation, and transparency.

The approach is intentionally modest about interpretation: accepted risk, as constructed here, is best used for benchmarking within an operator, not as a universal or frequentist estimate. Its utility rests on data availability and quality (e.g., covariance realism, secondary-state quality), provider latencies, and modeling choices such as hard-body radius. Late-notice events remain a salient challenge, underscoring the coupling between screening performance, reaction timelines, and achievable risk posture.

Future work will extend the analysis across mission phases and altitude regimes; explore integration of accepted, abated, and residual risk into a unified performance framework; incorporate consequence weighting (e.g., mass) and small-debris background; and run operational sensitivity studies on control variables such as reaction time, clearance criteria, and maneuver thresholds. Advancing standardized reporting and definitions would enable more meaningful comparisons over time and, ultimately, support performance-based oversight that complements traditional predictive assessments.

ACKNOWLEDGMENTS

We thank the entire OneWeb team—past and present—for their dedication and craft. In particular, we acknowledge the Fleet Management Segment team and the Flight Dynamics team for building, operating, and continuously improving the systems and analyses that underpin this work.

We are grateful to the 18th Space Defense Squadron, the 19th Space Defense Squadron, and Space-Track for sustained SSA/STM support and for permitting the use of data referenced in this paper.

Finally, we recognize the broader O/O community, who go above and beyond every day to keep the space environment safe through transparency, coordination, and professionalism.

Any opinions and errors are the authors' alone.

REFERENCES

[1] SpaceX Iridium, OneWeb. Satellite orbital safety best practices. *Facilitated by AIAA, September*, 2022.

[2] Nasa spacecraft conjunction assessment and collision avoidance best practices handbook. Technical Report NASA/SP-20230002470 Rev 1, National Aeronautics and Space Administration, Office of the Chief Engineer, Washington, DC, February 2023. Revision of the 2020 handbook; supersedes earlier versions.

[3] M. D. Hejduk and D. E. Snow. Satellite conjunction probability, possibility, and plausibility: A categorization of competing conjunction assessment risk assessment paradigms. *AAS/AIAA Astrodynamics Specialist Conference*, 2019.

[4] Soumaya Elkantassi and Anthony C. Davison. Space oddity? a statistical formulation of conjunction assessment. *JOURNAL OF GUIDANCE, CONTROL, AND DYNAMICS*, 45(12), 2022.

[5] Hugh G. Lewis and Georgia Skelton. Safety considerations for large constellations of satellites. In *2nd NASA Orbital Debris Conference*, Houston, TX, USA, 2023.

[6] Paul Diaz, Pol Mesalles Ripoll, Matthew Duncan, Mike Lindsay, Toby Harris, and High G. Lewis. Data-driven lifetime risk assessment and mitigation planning for large-scale satellite constellations. *The Journal of the Astronautical Sciences*, 70(21), 2023.

[7] Darren McKnight, Erin Dale, Rachit Bhatia, Chris Kunstadter, Matthew Stevenson, and Mohin Patel. A map of the statistical collision risk in LEO. In *73rd International Astronautical Congress (IAC-22-A6.2, x67913)*, Paris, France, September 2022. International Astronautical Federation.

[8] Doyle T. Hall, Luis G. Baars, and Stephen J. Casali. A multistep probability of collision computational algorithm. In *AAS/AIAA Astrodynamics Specialist Conference*, Big Sky, MT, USA, August 2023. American Astronautical Society.

[9] Ryan C. Frigm, Matthew D. Hejduk, Lauren C. Johnson, and Dragan Plakalovic. Total probability of collision as a metric for finite conjunction assessment and collision risk management. *Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference*, 2015.

[10] Mark Sturza and Robert Janka. Realized collision risk: A tool for comparing the safety of satellite operators. In *10th IAASS Conference: Safety is Not an Option*, Ostend, Belgium, October 2024.