International Journal of Sustainability and Innovation in Engineering (IJSIE)
https://www.doi.org/10.56830/IJSIE2026
Vol. 4 2026

Configuration-Aware Aircraft Maintenance Program Bridging: Data Model, Decision Rules and Scheduling Optimization Illustrated with FAA Airworthiness Directives

Short title: Configuration-Aware Maintenance Program Bridging

https://www.doi.org/10.56830/IJSIE202616

Arun Kumar Reddy Kudurmalla
Independent Researcher, Aviation Maintenance Programs, Columbus, Ohio, USA
Email: kudurumallaarunreddy55@gmail.com

Received:21 July, 2026, Accepted:30h October, 2026, Published: 11th November, 2026

Abstract

This design-science paper specifies the IMPB-TDO framework: a relational data model with explicit limit anchors and evidence links; seven ordered decision rules that assign each requirement–aircraft pair to one of seven disposition classes, including an engineering exception that is never defaulted to "not applicable"; flight-hour, flight-cycle and calendar forecasting with "whichever occurs first/later" logic; and a network-flow mixed-integer program that schedules only approved tasks. Four Federal Aviation Administration (FAA) Airworthiness Directives (ADs) from 2024–2025 serve as illustrative cases, and the framework is demonstrated on a synthetic 14-aircraft fleet and on randomly generated capacity-constrained fleets.

The ADs differ in what determines applicability (all serials, in-service embodiment, repair method, certificate date) and in costing unit; two published fleet cost totals are ceilings, not expected costs. In the worked example on AD 2025-16-06, production embodiment, a planned modification and an undocumented configuration each changed the disposition. On a fleet built to contain these situations, an SB-aware tracking comparator still differed on 3 of 14 configuration-sensitive pairs, all traceable to provenance or coupled actions; a model-level filter differed on 10. On 18 generated instances with fleet-wide inspector capacity, the optimization lowered the common objective by a mean of 15% (up to 22%) relative to the best of two planner heuristics wherever a heuristic found a plan, found compliant plans on 5 instances where both heuristics failed, and proved infeasibility on 1; no run was proven optimal within 180 s, so gaps are reported.

The paper contributes a specified, reproducible artefact; the synthetic results demonstrate mechanisms rather than operational effect sizes, which require the validation protocol provided.

Keywords: Aircraft maintenance program; Configuration management; Airworthiness directives; Maintenance program bridging; Mixed-integer programming

1. Introduction

Commercial aircraft are maintained under an operator's approved maintenance program, built from but not identical to manufacturer recommendations (Federal Aviation Administration, 2024); (European Union Aviation Safety Agency, 2025) and derived through MSG-3 logic (Airlines for America, 2022). The program changes continually: Maintenance Planning Document (MPD) revisions, Service Bulletins (SBs), Airworthiness Limitations (ALS) and Airworthiness Directives (ADs), the last being enforceable rules under 14 CFR Part 39 (Code of Federal Regulations), alter what must be done, to which aircraft and by when. **Maintenance program bridging** is the engineering process that carries each change into the aircraft-specific tracking record: deciding whether a requirement applies to each aircraft, how it relates to what is already tracked, when it falls due and which maintenance event it should be packaged into.

Applicability often depends on facts that the model designation does not carry. Under FAA AD 2025-16-06, which adopts EASA AD 2024-0204, the inspection applies to aircraft that had one of 33 listed modification SBs embodied *in service* on the effective date, while other aircraft must have rivets installed before the next flight after any later listed modification (Federal Aviation Administration, 2025), (Agency, 2024). Under FAA AD 2024-18-07, the new post-repair inspection applies only to repairs installed per the referenced Boeing SB; the FAA's response to comments states that repairs installed by other means are not subject to it (Federal Aviation Administration, 2024); (The Boeing Company, 2022). A tracking record that knows only the model, or a binary “SB embodied” flag without its provenance, will assign the wrong requirement or the wrong limit.

Scheduling optimization for check scheduling and task allocation is well developed (Deng, Santos, & Curran, 2020); (Witteman, Deng, & Santos, 2021), but it takes tasks, applicability and limits as inputs. Data standards define how publications and records are structured (Airlines for America, 2025); (S1000D Steering Committee (ASD, AIA & ATA), 2024), configuration-management standards define status accounting (SAE International, 2019), and commercial maintenance-and-engineering systems record effectivity and modification status (Swiss AviationSoftware, 2025). Text-analytics and knowledge-graph research extracts maintenance knowledge from documents (Kang, Zhou, Guo, Zhou, & Wu, 2024). What we did not find in the literature reviewed is a published, auditable specification that links these elements: how a revised requirement is dispositioned against an aircraft's as-maintained configuration, how unresolved evidence is carried, and how only the resulting approved tasks reach an optimizer.

This design paper contributes (C1) a relational data model linking revisions, requirements, limits, aircraft, configuration status, accomplishments, dispositions and maintenance windows; (C2) seven ordered decision rules with an explicit engineering-exception class; (C3) compliance forecasting across flight-hour (FH), flight-cycle (FC) and calendar limits, including “whichever occurs first/later” and terminating actions; and (C4) a network-flow mixed-integer linear program (MILP) for scheduling approved task instances. It demonstrates them (C5) on four published ADs, a step-by-step worked example on AD 2025-16-06 with mixed SB embodiment, and computational experiments. No operator records were available, so the paper makes no claim about delays, discrepancy rates or cost savings at any airline.

2. Background and Related Work

2.1 Program governance and sources of change

The Maintenance Review Board Report (MRBR) and MPD are recommendations for a type; AC 121-22D describes how operators use them to build their programs (Federal Aviation Administration, 2024), and point M.A.302 governs how programs are established and amended (European Union Aviation Safety Agency, 2025). Fig. 1 shows this pathway, with ADs and airworthiness limitations as separate mandatory inputs. Records must make program status auditable: 14 CFR 121.380 requires, among other items, "the current status of applicable airworthiness directives, including the date and methods of compliance" (Code of Federal Regulations). Changes differ in authority: ADs are mandatory (Code of Federal Regulations); post-certification limitation revisions are commonly mandated through ADs, as AD 2025-02-04 does for Airbus single-aisle aircraft (Federal Aviation Administration, 2025); SBs and MPD revisions are recommendations unless mandated or adopted. Fig. 2 summarises the configuration checks this requires.

2.2 Data standards, configuration management and records

ATA iSpec 2200 specifies the content, structure and exchange of maintenance information, including maintenance requirements and configuration control data (Airlines for America, 2025). S1000D organises technical publications as data modules in a common source database and carries applicability as metadata evaluated against product and condition attributes (S1000D Steering Committee (ASD, AIA & ATA), 2024). ATA Spec 2000 (for example, Chapter 11 on reliability data exchange) and Spec 2500 on aircraft transfer records address records exchange (Airlines for America, 2023); (Airlines for America, 2023), and AC 120-78B covers electronic recordkeeping (Federal Aviation Administration, 2024). SAE EIA-649C defines configuration management through planning, identification, change management, status accounting, and verification and audit (SAE International, 2019). These standards fix how requirements and records are represented; they do not prescribe how a revised requirement is classified as new, amended, retained, superseded, terminated, not applicable or unresolved for a given aircraft. IMPB-TDO is positioned as a decision layer that consumes these representations.

2.3 Maintenance-and-engineering systems and document understanding

Commercial maintenance-and-engineering (M&E) systems such as AMOS and TRAX eMRO include engineering modules. AMOS documentation describes tracking compliance with ADs and SBs, revision control and configuration control, setting aircraft and component effectivities for modification documents, and reporting modification status by aircraft and fleet (Swiss AviationSoftware, 2025). A vendor has recently described agent-based extraction of SB and AD effectivity data (serial-number ranges, manufacturing dates, retrofit conditions) into review workspaces, with engineers approving the result (Ramco Systems, 2026). These systems are proprietary, and their decision logic and its treatment of missing evidence are not published in a form that can be evaluated or reproduced. Research on document understanding offers related methods: technical language processing adapts natural-language processing to maintenance text (Brundage, Sexton, Hodkiewicz, Dima, & Lukens, 2021); knowledge graphs have been built from maintenance manuals and records (Yue, Xiao, Li, & Wang, 2022) and an ontology with rule-based inference has been proposed for aircraft maintenance processes (Kang, Zhou, Guo, Zhou, & Wu, 2024). Manual expert classification of a sample of ADs has also been reported (Lercel, Patankar, & Steckel, 2024). This work

addresses extraction and representation of knowledge; IMPB-TDO addresses the subsequent decision step, and could take its inputs from such extraction.

2.4 Program development and scheduling optimization

Program-development research decides which tasks and intervals should exist: Ahmadi et al. examined scheduled-task development within the MRB process (Ahmadi, Söderholm, & Kumar, 2010), Gonçalves and Trabasso used fleet data and reliability analysis to support a 12,000-flight-hour interval for a fire-extinguisher functional test on a new system against 6,000 flight hours in service (Gonçalves & Trabasso, 2018), and operators escalate intervals using reliability data (Kinnison & Siddiqui, 2013). Scheduling research decides when tasks are performed: from aircraft maintenance scheduling and re-assignment (Sriram & Haghani, 2003) to dynamic-programming check scheduling for a fleet of about 40 aircraft (Witteman, Deng, & Santos, 2021), integrated check scheduling and task allocation (Airlines for America, 2025); (Deng, Q.; Santos, B. F.; Verhagen, W. J. C., 2021), hangar and staff planning (Qin, Zhang, Chan, Chung, Niu, & Qu, 2020), line maintenance in layover windows (Shaukat, Katscher, Wu, Delgado, & Larrain, 2020) and dynamic maintenance operations (Lagos, Delgado, & Klapp, 2020). In these models the task list and limits are given. The MILP here reuses the standard assignment structure; what is specific is what it receives: only approved dispositions, provisional tasks for unresolved exceptions and coupled "before next flight after" actions. Maintenance information-systems research has addressed data capture and traceability (Lampe, Strassner, & Fleisch, 2024); (Ho, Tang, Tsang, Tang, & Chau, 2024), and the FAA DRS provides authoritative publication of ADs (Federal Aviation Administration, 2025). Table 1 summarises the positioning.

Table 1. Positioning of IMPB-TDO against related bodies of work

Body of work	Representative sources	What it provides	What it leaves open	IMPB-TDO element
Program development and governance	[1–3, 25–27]	Tasks and intervals for a type; approval of operator programs	Correct application of revisions to each diverged aircraft	Revision register; lineage (R3)
Data and configuration-management standards	[(Airlines for America, 2025)–(SAE International, 2019), 18–20]	Structure, exchange and applicability metadata; CM functions	Disposition logic against as-maintained configuration	Applicability evaluated against CONFIG_STATUS (R1–R2); DISPOSITION record
Commercial M&E systems	AMOS, TRAX eMRO; vendor AI extraction [14, 21]	Effectivity, modification status, compliance tracking	Decision logic and treatment of missing evidence not published or reproducible	Specified rules; explicit exception class with records action

Body of work	Representative sources	What it provides	What it leaves open	IMPB-TDO element
Document understanding	Technical language processing, knowledge graphs, ontologies, expert AD review (Kang, Zhou, Guo, Zhou, & Wu, 2024); 22–24]	Extraction and representation of maintenance knowledge	The disposition decision after extraction	Could supply REQUIREMENT and CONFIG_STATUS inputs
Scheduling and task allocation	[9, 10, 28–32]	Check scheduling, task allocation, hangar and staff planning	Task list, applicability and limits taken as given	Network-flow MILP receiving only approved, provisional and coupled tasks

3. Research Design and Methods

3.1 Design-science approach

The study follows design-science research, in which the output is an artefact whose utility is shown through demonstration and evaluation (Hevner, March, Park, & Ram, 2004). The six activities of Peffers et al. (Peffers, Tuunanen, Rothenberger, & Chatterjee, 2007) map onto the paper as problem identification (Sections 1–2), objectives (correct disposition, traceable evidence, valid limits and feasible schedules), design (Section 4), demonstration (Section 5) and evaluation against stated criteria (Section 3.5), with field evaluation deferred to the protocol in Section 6.3. The ADs are illustrative cases that supply real rule structures and parameters, not a sample for population inference.

3.2 Case selection and extraction

Candidates were FAA final-rule ADs published in the Federal Register between 1 January 2024 and 31 December 2025 and indexed in the DRS (Federal Aviation Administration, 2025), applicable to transport-category airplanes, requiring at least one action that enters a maintenance or inspection program, and quantifying work-hours in their Costs of Compliance section; engine- and appliance-only ADs were excluded. Four cases were chosen by diverse-case selection (Seawright & Gerring, 2008) to cover different values of four dimensions that matter to bridging: applicability determinant, requirement type, compliance-time logic and costing unit (Table 2). Each AD was abstracted into a 23-field extraction matrix, with the source paragraph of every field recorded; EASA AD 2024-0204 and Boeing Alert SB 737-25A1544 Revision 4, which the ADs incorporate by reference, were abstracted for group definitions and compliance tables [6, 8]. The compliance table of Boeing Requirements Bulletin 737-53A1410 RB, referenced by AD 2025-15-01, is not reproduced in the AD or its proposed rule and was not available. The matrix is Supplementary Material S1.

Table 2. Diverse-case selection: the dimension values contributed by each AD

Case (AD)	Applicability determinant	Requirement type	Compliance-time logic	Costing unit (FAA)
C1: 2025-15-01 (Boeing 737) (Federal Aviation Administration, 2025)	All serial numbers of listed models	Repetitive inspection	In referenced requirements bulletin (not public)	Per airplane per inspection cycle
C2: 2025-16-06 (A319/A320) [5]	Serial list plus modification SB embodied in service	One-time inspection; coupled action after future modification	Later of 24 months and 17,000 FC after embodiment	Per affected airplane, upper bound
C3: 2024-18-07 (Boeing 737) [7]	Model, SB group and repair installation method	Supersession; retained and new requirements; repetitive post-repair inspection	Later of 7,000 FC after effective date and 36,000 total FC; repeat 24,000 FC [8]	Per post-repair inspection cycle
C4: 2025-02-04 (A318–A321) (Federal Aviation Administration, 2025)	Model plus certificate-of-airworthiness date	Program revision; terminating provision	Within 90 days; initial tasks at limit or 90 days, whichever later	Per operator

3.3 Synthetic fleet and benchmark instances

The worked example uses a synthetic fleet of eight Airbus single-aisle aircraft (A01–A08) and six Boeing 737NG aircraft (B01–B06) under one operator, planned from 1 October 2025 over 36 months. Published content is used wherever it exists: applicability models, group definitions and the limits of AD 2025-16-06 (later of 24 months after the 22 September 2025 effective date and 17,000 FC after embodiment); the limits of AD 2024-18-07 paragraph (j) from SB Table 2 as modified by the AD (later of 7,000 FC after the 12 November 2024 effective date and 36,000 total FC for Revision 3 repairs; repeat 24,000 FC); the repair-provenance rule; and the FAA labour estimates as task durations. Aircraft data, utilization, histories, windows and routine MPD and ALS tasks are synthetic. The only regulatory placeholders are the AD 2025-15-01 limits (hypothetical, HYP), which a sensitivity analysis varies over a 4 × 3 grid. Configurations were chosen so that every rule and class occurs, including three traps found in the cases: production versus in-service embodiment, repairs by a non-listed method and records that do not establish the configuration. Because the fleet contains each trap by construction, its comparison results show which errors the rules prevent, not how often they occur.

To test scheduling where the problem is genuinely hard, 18 instances were generated with 10, 20 or 30 aircraft (three seeds each) over 36 months. Each aircraft has eight repetitive tasks (intervals 90–720 days, 2–24 hours, two skills), two one-time tasks of up to 42 hours, regular windows every 30 days (8 h inspection, 12 h mechanical), two heavy checks and optional additional ground events every 60 days. A fleet-wide inspection (NDT) capacity per

30-day period couples the aircraft, set at 2.0 or 1.6 times the average NDT demand. Parameters are listed in Appendix 1.

3.4 Computation

Rules, forecasting and both MILP formulations were implemented in Python with PuLP (Mitchell, O'Sullivan, & Dunning, 2011) and the HiGHS solver (Huangfu & Hall, 2018) (one thread, relative gap 0.01%). FC limits are converted to dates with a 10% utilization margin. Case-study scenarios were S1 (base), S2 (regular-window NDT capacity halved) and S3 (S2 without additional ground events). Two planner heuristics were run on every instance and scored on the same objective: *latest* places each task in the latest feasible planned window, and *consolidate* prefers windows already opened. In the benchmark, the MILP was given a 180-s limit and started from the better heuristic plan where one existed; we report the optimality gap with every result. An independent routine re-checked every plan for late or missing tasks, interval violations, chains ending before the horizon, capacity overruns and coupling. Code and outputs are Supplementary Material S2.

3.5 Evaluation criteria and coding reliability

Four demonstration criteria were applied: (E1) every rule and class is exercised by a pair traceable to a published AD structure; (E2) computed limits agree with hand calculation; (E3) dispositions are compared pair by pair with two comparators: C1, a model-level filter (aircraft family, binary embodiment flag, unknown treated as not applicable, “first” logic by default), and C2, an SB-aware tracker that checks model and serial number, applies the published first/later logic and treats unknown status as applicable, but stores no embodiment mode or repair provenance and does not couple future actions; (E4) the MILP is compared with both heuristics on the common objective. Given encoded facts, dispositions are produced deterministically by the published code, so any reader can reproduce them. The encoding of AD text into rules was done by one analyst; to allow an inter-rater check, Supplementary Material S3 provides a blind coding sheet (facts without dispositions) and S2 a script computing percentage agreement and Cohen's kappa against the first coding.

4. The IMPB-TDO Framework

4.1 Architecture and data model

IMPB-TDO has five stages (Fig. 3): technical-data acquisition into a revision register, effectivity verification, reconciliation against the current program and tracking, compliance forecasting, and scheduling. Stages 1–4 are rule-based; Stage 5 is the optimization. Fig. 4 shows the bridging workflow for one revision, with engineering review resolving exceptions. Mandatory requirements (ADs, ALS) are hard constraints that no optimization step can relax.

Fig. 5 gives the data model. SOURCE_DOCUMENT holds each revision with its supersession link, mandatory flag and receipt date; REQUIREMENT holds one actionable paragraph with its applicability expression, first/later logic, terminating link and lineage; LIMIT holds each threshold or interval with its basis (FH, FC, calendar), anchor (birth, effective date, embodiment, repair, last accomplishment) and value. AIRCRAFT holds serial, model, certificate date, current usage and planning utilization; CONFIG_STATUS holds each SB, modification, repair or part with a three-valued status (embodied, not embodied, unknown), embodiment mode (in service or production) and the source record; ACCOMPLISHMENT holds the Part 121 record content (Code of Federal Regulations). DISPOSITION is the status-accounting record [13]: class, rule, evidence, computed limit and

approving engineer, one row per aircraft, requirement and program revision. MAINT_WINDOW and ASSIGNMENT form the scheduling layer, and an assignment may reference only an approved disposition.

4.2 Decision rules

Each requirement–aircraft pair passes through the ordered rules in Table 4 and Fig. 6. Applicability is three-valued: a predicate on an attribute whose status is unknown yields an engineering exception (R2), which is scheduled provisionally at the earliest limit consistent with the requirement, raises a records-research action and re-enters at R2 when resolved. Supersession and terminating actions (R6) are tested before lifecycle classification (R3), so closed requirements generate no tasks. R4 applies the limit logic, R5 couples "before next flight after" actions to their triggering event and R7 fixes AD and ALS limits as hard constraints. Evidence is mandatory for every class, so negative decisions are as auditable as positive ones.

Table 4. IMPB-TDO decision rules (evaluated in order for each requirement–aircraft pair)

Rule	Condition evaluated	Outcome	Evidence recorded	Exercised by
R1	Model, serial number and certificate date within the applicability statement	FALSE → not applicable	Type certificate data; serial list	C2 (A07 serial, A08 model); C3 (B06 -900ER); C4 (date)
R2	Configuration predicate (embodiment and mode, repair method, part number), three-valued	FALSE → not applicable; UNKNOWN → engineering exception (provisional task, records action); TRUE → continue	CONFIG_STATUS rows and source records	C2 (production embodiment; unknown); C3 (non-listed repair; undocumented repair)
R6	Source superseded, or terminating action accomplished	Superseded or terminated (closed, no task)	Supersession link; terminating accomplishment	C3 (AD 2009-01-02); C4 (AD 2022-22-10; para (g))
R3	Lineage predecessor exists; content identical?	None → new; identical → retained (credit prior compliance); changed → amended	Prior requirement and accomplishment	C1, C2 (new); C3, C4 (retained); C4 (amended)
R4	Limits per dimension with anchors; first/later logic	Limit date by Equations (1)–(2); repetitive chain	Utilization basis	C2 (later); C3 (later; repetitive); C1 (repetitive)
R5	Action required "before next flight after" an event	Coupled task fixed to the triggering window	Planned modification record	C2 (Group 2 rivet installation)
R7	Requirement is an AD or ALS limit	Hard constraint; no escalation or tolerance unless stated	Mandatory flag	All cases

Note: R6 is evaluated before R3 so that closed requirements generate no tasks.

4.3 Compliance forecasting

For requirement r on aircraft a with limit set K_r, limit k has basis b_k, anchor value α_{ak} and value v_k; with current usage U_{ak} at as-of date t_0, planning utilization $\hat{u}_{ak}$ per day and conservatism factor $\rho \geq 1$ (calendar limits used directly):

$$\Lambda_{ak} = \alpha_{ak} + v_k, \quad D_{ak} = t_0 + \lfloor (\Lambda_{ak} - U_{ak}) / (\rho\, \hat{u}_{ak}) \rfloor \tag{1}$$

$$D_{ar} = \min_{k \in K_r} D_{ak} \text{ (“whichever occurs first”)} \quad \text{or} \quad \max_{k \in K_r} D_{ak} \text{ (“whichever occurs later”)} \tag{2}$$

Compound limits are evaluated recursively as trees of min and max nodes. For repetitive requirements the next anchor is the previous accomplishment, so the interval enters the scheduler as a network constraint rather than a fixed date; a terminating action closes the chain. No tolerance is applied to AD or ALS limits unless the source provides one. Equation (2) matters: for AD 2025-16-06, “first” instead of the published “later” brings every Group 1 limit forward and shows aircraft whose 17,000-FC term has passed as overdue. The revision register timestamps (published, effective, received, dispositioned, approved, incorporated) define the incorporation-latency measure in Section 6.3.

4.4 Scheduling formulation

Let T be the approved tasks (new, amended or retained with an open limit, plus provisional exception tasks), split into one-time tasks T^O, repetitive tasks T^R and coupled tasks T^C. Task t belongs to aircraft a(t), needs h_t hours of skill s(t) and has first limit L_t (Equation 2); repetitive tasks have interval δ_t. W is the set of windows, with aircraft a(w), day d_w, capacity C_{ws}, cost c_w and calendar period p(w); H is the horizon end. Each repetitive task is a unit flow on a directed graph G_t whose nodes are the windows of a(t): source arcs (S, w) for $d_w \leq L_t$; arcs (w, w′) for $0 < d_{w'} - d_w \leq \delta_t$; and sink arcs (w, T) for $d_w + \delta_t > H$. Any source–sink path is a compliant chain, so interval limits hold by construction. Arc costs are unused interval: $\kappa_{Sw} = (L_t - d_w)/\delta_t$, $\kappa_{ww'} = (d_w + \delta_t - d_{w'})/\delta_t$, $\kappa_{wT} = 0$. With binary x_{tw} (task t in window w), y_w (window used) and continuous arc flows $f_e \in [0,1]$:

$$\min \ \Sigma_w c_w y_w + \lambda \Sigma_{t \in TR} \Sigma_{e \in Gt} \kappa_e f_e + \mu \Sigma_{t \in TR} \Sigma_{e \in Gt:\, e=(w,w')} f_e + \nu \Sigma_{t \in TO} \Sigma_w ((L_t - d_w) / (L_t + 1))\, x_{tw} \tag{3}$$

$$\Sigma_{w:\, dw \leq Lt} x_{tw} = 1, \quad t \in T^O; \qquad x_{tw^*} = 1, \quad t \in T^C \text{ (triggering window } w^*) \tag{4}$$

$$\Sigma_w f_{Sw} = 1, \quad \Sigma_{e \text{ into } w} f_e = \Sigma_{e \text{ out of } w} f_e, \quad x_{tw} = \Sigma_{e \text{ into } w} f_e, \quad t \in T^R, w \in G_t \tag{5}$$

$$\Sigma_{t:\, s(t)=s} h_t x_{tw} \leq C_{ws} y_w, \quad w \in W, s \in S \tag{6}$$

$$\Sigma_{t:\, s(t)=NDT} \Sigma_{w:\, p(w)=p} h_t x_{tw} \leq C^{fleet}_p, \quad \text{for every period p (optional)} \tag{7}$$

The objective (3) has four terms with separate meanings. Window costs distinguish regular, heavy and additional events. The unused-interval term applies only to repetitive tasks, where performing a task early shortens every later interval. A small penalty μ counts repetitions. The deferral term applies only to one-time tasks and is normalised by the time available to the limit, so it is defined without an interval; it expresses a planner preference for keeping near-term capacity free and leaving exceptions open to resolution by records research, not a safety requirement. Coupled tasks carry no timing term. Constraint (4) makes one-time and coupled tasks mandatory, (5) routes each repetitive chain, (6) enforces window capacity by skill and (7) adds an optional fleet-wide inspector capacity per period. Binary node variables make any fractional split of flow infeasible, so arc flows can be continuous. Engineering approval acts through the task set, and the model never alters a limit. We also implemented the equivalent big-M chain formulation as a check. Without (7) the model

decomposes by aircraft; with (7) it does not. Demonstration weights are c_w = 1, 5 and 50 for regular, heavy and additional windows, $\lambda = 10$, $\mu = 2$ and $\nu = 2$. Parts, tooling, precedence beyond coupling and stochastic utilization are outside the formulation.

5. Results

5.1 Cases and published cost estimates

The four ADs differ on every selection dimension (Tables 2 and 3) and together exercise all seven rules (criterion E1; Table 4). Their labour estimates use four units, and two fleet totals are upper bounds. For AD 2025-15-01, \$1,210,995 per inspection cycle (1,583 airplanes × 9 h × \$85) is a meaningful fleet figure because applicability covers all serial numbers, although on-condition repairs are excluded (Federal Aviation Administration, 2025). For AD 2025-16-06, "up to" 8 h (\$680) per product and "up to \$196,520" for 289 airplanes is a ceiling, because the inspection applies only to Group 1 [5]. For AD 2024-18-07, the published \$2,634,660 per post-repair inspection cycle equals 738 × \$3,570, every airplane in the applicability, whereas paragraph (j) applies only to airplanes with SB-installed repairs [7]. For AD 2025-02-04, 90 h (\$7,650) is per operator, and without the number of operators no fleet total follows (Federal Aviation Administration, 2025). A per-product estimate can be multiplied only by the aircraft confirmed applicable, which is the bridging decision itself.

Table 3. Published FAA cost estimates, their units and interpretation

AD (Federal Register)	U.S. airplanes	Labour estimate	Unit	Cost per product	Cost on U.S. operators (as published)	Interpretation
2025-15-01 (90 FR 34171)	1,583	9 h	Per airplane per inspection cycle	\$765 per cycle	\$1,210,995 per cycle	Valid fleet figure (all serials); excludes on-condition repairs
2025-16-06 (90 FR 40028)	289	Up to 8 h	Per affected airplane	Up to \$680	Up to \$196,520	Ceiling: Group 1 only; may be covered under warranty
2024-18-07 (89 FR 81320), new post-repair inspection	738	42 h	Per post-repair inspection cycle	\$3,570 per cycle	\$2,634,660 per cycle (= 738 × \$3,570)	Ceiling: applies only to airplanes with SB-installed repairs
2025-02-04 (90 FR 9214)	1,857	90 h (average)	Per operator	\$7,650 per operator	Not stated	Operator-level; no fleet total without number of operators

Note: Labour rate \$85 per work-hour in all four ADs. Figures as published.

5.2 Worked example: AD 2025-16-06 with mixed SB embodiment

Table 5 and Fig. 7 apply the workflow to the eight Airbus aircraft. *Receive:* the AD is registered with its 22 September 2025 effective date, the 33 listed modification SBs and the two terms, A (24 months after the effective date: 22 September 2027) and B (17,000 FC after embodiment), whichever occurs later [5, 6]. *Verify (R1–R2):* A08 (A321) is not applicable by

model and A07 by serial number. A03's listed modification was embodied in production; Group 1 requires embodiment in service, so A03 is Group 2 and, with no modification planned, not applicable. A binary embodiment flag would have placed it in Group 1. A04 is unmodified but has a listed cabin modification planned for its heavy check on 27 August 2026, so R5 couples rivet installation to that window. A05's transferred records do not establish embodiment, so R2 raises an engineering exception. A01, A02 and A06 are Group 1. *Reconcile (R3, R6):* the inspection has no predecessor, so applicable pairs are new. *Forecast (R4):* for A01 and A06, term B (May 2021; 19 January 2026) precedes term A, so the limit is 22 September 2027; for A02, embodied at 30,000 FC, B falls on 25 July 2032 and governs; for the A05 exception the earliest consistent limit is A. *Package:* the A01, A06 and provisional A05 inspections and the A04 rivet installation enter the scheduler, at most 32 hours using the FAA's 8-hour bound, against 56 hours if the per-product estimate were applied to every aircraft whose model is listed.

Table 5. Worked example: bridging AD 2025-16-06 on eight synthetic Airbus aircraft (as of 1 October 2025)

A/C	Model	Serial in App. 1	Modification SB status (mode)	FC now	FC at embodi-ment	Term B forecast	Limit (later of A, B)	Disposition
A01	A320-214	Yes	Embodied (in service)	42,000	18,000	May 2021 (passed)	22 Sep 2027	New (Group 1)
A02	A320-214	Yes	Embodied (in service)	35,500	30,000	25 Jul 2032	25 Jul 2032 (beyond horizon)	New (Group 1); tracked
A03	A319-112	Yes	Embodied (in production)	39,000	–	–	–	Not applicable (Group 2; no modification planned)
A04	A320-251N	Yes	Not embodied; modification planned 27 Aug 2026	14,000	–	–	Before next flight after modification	New (Group 2; coupled rivet installation)
A05	A320-214	Yes	Unknown (records gap)	44,500	Unknown	–	22 Sep 2027 (provisional)	Engineering exception
A06	A320-214	Yes	Embodied (in service)	40,500	24,000	19 Jan 2026	22 Sep 2027	New (Group 1)
A07	A320-214	No	Embodied (in service)	41,000	20,000	–	–	Not applicable (serial number)
A08	A321-211	–	Embodied (in service)	38,000	21,000	–	–	Not applicable (model)

Note: Term A = 24 months after the 22 September 2025 effective date. Term B = embodiment FC + 17,000 FC, converted with Equation (1) ($\rho = 1.10$). Aircraft data are synthetic; regulatory logic from [5, 6].

5.3 Other cases and comparison with tracking comparators

For AD 2024-18-07, B06 (737-900ER) is outside the applicability (R1). B01 and B02 have Revision 3 SB repairs, giving limits of 52,385 and 52,885 FC (6 June 2028); B03's repair was approved on FAA Form 8100-9 and is not subject to paragraph (j); B05 has no repair; and B04 has a recorded repair whose method is undocumented, an exception scheduled

at the earliest consistent limit (24 September 2028). Retained paragraphs (g) and (h) are retained with prior compliance credited, and the AD 2009-01-02 items superseded. For AD 2025-15-01, all six 737s receive new repetitive inspections. For AD 2025-02-04, the AD 2022-22-10 items are superseded, paragraph (g) is terminated by the paragraph (j) revision, and two illustrative ALS tasks are retained and amended; the 90-hour program revision is one operator-level engineering task, not 90 hours per aircraft. The register holds 64 dispositions covering every class (Table 6), and all limits matched hand calculation (E2; for example A06: (41,000 − 40,500)/(4.1 × 1.1) = 110 days, 19 January 2026).

Table 7 compares the 14 configuration-sensitive pairs. The model-level filter C1 differs on 10: four false inclusions (A03, A07, A08, B03), three omissions (A04 coupled rivets; A05 and B04 unknowns treated as not applicable) and three wrong limits from "first" logic. The SB-aware tracker C2, which already checks serial numbers, applies the published logic and treats unknowns conservatively, differs on 3: A03 (no embodiment mode), B03 (no repair provenance) and A04 (no coupling). C2 schedules A05 and B04 like the rules do, but without the records action that can close the exception. These differences follow from the configurations included by design; the result is which information each comparator lacks, not an error rate.

Table 6. Disposition register for the synthetic fleet across the four ADs (number of requirement–aircraft pairs)

Requirement (aircraft)	New	Amd.	Ret.	Sup.	Term.	N/A	Exc.
AD 2025-16-06 SDI / rivet installation (A01–A08)	4					3	1
AD 2022-22-10 tracking items (A01–A08)				8			
AD 2025-02-04 para (g) revision (A01–A08)					8		
ALS task HYP-T1 (unchanged) (A01–A08)			8				
ALS task HYP-T2 (interval reduced) (A01–A08)		8					
AD 2025-15-01 frame inspections (B01–B06)	6						
AD 2009-01-02 tracking items (B01–B06)				5		1	
AD 2024-18-07 (g)/(h) retained (B01–B06)			5			1	
AD 2024-18-07 (j) post-repair inspection (B01–B06)	2					3	1
Total (64)	12	8	13	13	8	8	2

Note: Amd. = amended; Ret. = retained; Sup. = superseded; Term. = terminated; N/A = not applicable; Exc. = engineering exception. The AD 2025-02-04 program revision is one operator-level engineering task (90 h), not counted per aircraft.

Table 7. Rule-based dispositions compared with two tracking comparators for 14 configuration-sensitive pairs

A/C	Requirement	C1 model-level filter	C2 SB-aware tracker	IMPB-TDO rules	Missing information
A01	AD 2025-16-06	Applicable; May 2021 (“first”): appears overdue	Applicable; 22 Sep 2027	Applicable; 22 Sep 2027	C1: logic
A02	AD 2025-16-06	Applicable; 22 Sep 2027	Applicable; 25 Jul 2032	Applicable; 25 Jul 2032	C1: logic
A03	AD 2025-16-06	Applicable (flag “embodied”)	Applicable (flag “embodied”)	Not applicable (production embodiment)	C1, C2: mode
A04	AD 2025-16-06	Not applicable	Not applicable	Coupled rivet installation	C1, C2: coupling
A05	AD 2025-16-06	Not applicable (unknown)	Applicable; 22 Sep 2027	Exception; provisional; records action	C1: unknown
A06	AD 2025-16-06	Applicable; 19 Jan 2026	Applicable; 22 Sep 2027	Applicable; 22 Sep 2027	C1: logic
A07	AD 2025-16-06	Applicable (family)	Not applicable (serial)	Not applicable (serial)	C1: serial
A08	AD 2025-16-06	Applicable (family)	Not applicable (model)	Not applicable (model)	C1: model
B01	AD 2024-18-07 (j)	Applicable	Applicable	Applicable; 6 Jun 2028	–
B02	AD 2024-18-07 (j)	Applicable	Applicable	Applicable; 6 Jun 2028	–
B03	AD 2024-18-07 (j)	Applicable (any repair)	Applicable (any repair)	Not applicable (repair not per SB)	C1, C2: provenance
B04	AD 2024-18-07 (j)	Not applicable (unknown)	Applicable	Exception; provisional; records action	C1: unknown
B05	AD 2024-18-07 (j)	Not applicable	Not applicable	Not applicable	–
B06	AD 2024-18-07 (j)	Not applicable	Not applicable	Not applicable (model)	–

Note: C1 differs from the rules on 10 pairs. C2 differs in scheduled work on 3 (A03, A04, B03); on A05 and B04 it schedules the same work but raises no records action. The synthetic configurations were constructed to contain each situation, so the counts show which information each comparator lacks, not error frequencies.

5.4 Scheduling the synthetic fleet

Table 8 reports the case-study scheduling results. The network-flow model reached the proven optimum of S1 in 1.1 s (1,919 binary variables); the big-M chain formulation gave the identical objective (521.7) in 6.2 s, which checks the two formulations against each other. S1 schedules 197 task instances (2,062 work-hours) in 152 windows without additional ground events (Fig. 8); because of the deferral term, the A01 inspection is placed in the window before its September 2027 limit; without it, one-time tasks would be placed in any window with spare capacity, including the first. On this instance the gains over the heuristics are small: 1.3% and 6.0% lower objective than *latest* and *consolidate*. With regular-window NDT capacity halved (S2), the 9-hour and 12-hour inspection tasks no longer fit regular windows, the optimum uses 17 additional ground events, and its objective is 8.3% and 2.6% below the heuristics; S3, which forbids additional events, is infeasible, so the model shows that the shortfall cannot be absorbed by the planned checks. Across the 12 combinations of hypothetical AD 2025-15-01 limits, the base case never required an additional event (window cost 164–171), while S2 required 17–29 additional events; the qualitative conclusions do not depend on the placeholder values, but the counts do.

Table 8. Scheduling results for the synthetic 14-aircraft fleet (36-month horizon, all methods scored on objective (3))

Scenario and method	Status	Objective	Window cost	Windows used	Addit. events	Instances (work-hours)	Unused interval	Time (s)
S1 base, network-flow MILP	Optimal	521.7	168	152	0	197 (2,062)	5.95	1.1
S1 base, big-M MILP (check)	Optimal	521.7	168	152	0	197 (2,062)	5.95	6.2
S1 base, heuristic latest	Heuristic	528.3	183	163	0	197 (2,062)	5.11	<1
S1 base, heuristic consolidate	Heuristic	554.9	148	132	0	200 (2,090)	10.64	<1
S2 NDT −50%, network-flow MILP	Optimal	1451.5	1047	158	17	197 (2,062)	10.96	0.8
S2 NDT −50%, heuristic latest	Heuristic	1583.0	1161	174	19	200 (2,098)	12.14	<1
S2 NDT −50%, heuristic consolidate	Heuristic	1491.0	1048	159	17	202 (2,114)	13.81	<1
S3 = S2 without additional events	Infeasible	–	–	–	–	–	–	0.2

Note: Window cost 1 per regular, 5 per heavy and 50 per additional event; $\lambda = 10$, $\mu = 2$, $\nu = 2$. All plans passed independent verification. Data synthetic (Appendix 1).

5.5 Benchmark with fleet-wide capacity

Table 9 summarises the 18 generated instances. Where at least one heuristic found a plan (12 instances, on 2 of which only one of the two did), the MILP objective was lower than that of the best heuristic plan in 10 of 12 cases, by 0.0–22.4% (mean 15.1%). On 2 of the 30-aircraft instances the solver did not improve on its heuristic warm start within 180 s; with a 600-s limit it improved them by 11.8% and 0.0% (remaining gaps 10.1%, 25.1%). On 5 instances, both heuristics failed to find a plan that met every limit within capacity, while the MILP found one. On 1 instance the MILP proved that no compliant plan exists within the stated capacity, a conclusion that a heuristic failure cannot establish. None of the 17 MILP runs that returned a plan was proven optimal within 180 s; all stopped at the time limit with a mean remaining optimality gap of 16.6% (maximum 80.2%, on the tightest instances), so the reported improvements are conservative. All MILP and heuristic plans passed independent verification. A likely source of the advantage, consistent with fewer additional events in every group except the 30-aircraft instances at capacity factor 2.0 (Table 9), is coordination that greedy rules do not attempt: moving a long-interval task to free inspector capacity in a congested period, and using windows that other tasks have already opened.

Table 9. Benchmark on generated fleets with fleet-wide NDT capacity (36-month horizon; three seeds per row)

Capacity factor	Aircraft	Tasks	Latest: plan found	Consolidate: plan found	MILP: plan found	MILP vs latest (%)	MILP vs consolidate (%)	Addit. events MILP / consolidate	MILP gap % mean (max)
2.0	10	100	2/3	2/3	3/3	22.0	23.2	18.0 / 24.5	4.8 (6.1)
2.0	20	200	3/3	3/3	3/3	18.0	22.6	37.3 / 44.3	5.1 (9.1)
2.0	30	300	3/3	3/3	3/3	4.7	11.0	72.3 / 68.7	17.0 (25.3)
1.6	10	100	0/3	0/3	2/3	–	–	–	42.6 (80.2)
1.6	20	200	2/3	1/3	3/3	17.6	24.0	44.0 / 59.0	8.7 (11.5)
1.6	30	300	2/3	1/3	3/3	17.5	23.7	51.0 / 67.0	29.9 (76.6)

Note: Capacity factor = fleet NDT capacity per 30-day period divided by average NDT demand. Improvement = reduction in objective (3) relative to that heuristic, averaged over the instances where it found a plan. MILP limited to 180 s with a warm start from the better heuristic plan; additional events compared on the instances where the consolidating heuristic found a plan; gap = remaining relative optimality gap reported by HiGHS. All data synthetic.

6. Discussion

6.1 Contribution and implications

IMPB-TDO turns a set of practices that M&E systems implement in proprietary form [14, 21] into a specification that can be examined and reproduced: explicit anchors and

provenance, three-valued applicability with a carried exception, recursive first/later limits and an optimizer that cannot receive unapproved tasks or relax limits. The cases show that each rule answers a structure in current FAA rulemaking. Three implications for practice follow, each needing confirmation with operator data. Published AD cost totals should be multiplied only by confirmed-applicable aircraft, and per-operator estimates not at all. Provenance attributes (embodiment mode, repair method) must be stored explicitly: they account for every residual difference between the SB-aware comparator and the rules. Finally, in the generated instances, once inspection capacity is shared across a fleet, planner heuristics left measurable cost on the table and sometimes failed to find a compliant plan that exists.

6.2 Limitations

(i) Fleet, configurations, utilization, windows and routine tasks are synthetic; the AD 2025-15-01 limits are placeholders whose effect was bounded by sensitivity analysis but not removed. (ii) The comparator results reflect configurations chosen to contain each trap, and the benchmark reflects generated instances; neither estimates operational frequencies or savings. (iii) Four diverse cases do not support generalisation about ADs as a population. (iv) The encoding of AD text was done by one analyst; the blind second-coding package (S3) allows an independent check, but its result is not reported here. (v) The MILP is deterministic in utilization and omits parts, tooling and access; no benchmark instance was solved to proven optimality within 180 s, gaps are reported, and at 30 aircraft longer runs or decomposition would be needed. (vi) Interpretations follow the published AD text and comment responses; operational use requires the operator's approved procedures.

6.3 Validation protocol

A prospective validation should run IMPB-TDO in parallel with an operator's existing process over several consecutive revisions and ADs, using tracking exports, configuration and repair records, accomplishment histories, the approved program and window plans, with the existing process as baseline. Table 10 defines the measures. Accuracy should be scored against an independent engineering audit of a random sample of requirement–aircraft pairs stratified by class, including not-applicable decisions, with inter-rater agreement reported (S3 provides the instrument). Ground-time, additional-event and AOG effects should be reported only where operator logs allow attribution.

Table 10. Proposed measures for prospective validation with operator data

Measure	Definition	Data required
Incorporation latency	Days from source effective or received date to approved disposition, and to tracking update, per revision and aircraft	Revision register timestamps; tracking change log
Disposition accuracy	Share of sampled pairs whose class matches an independent engineering audit, by class (including not applicable)	Stratified audit sample; configuration and repair records
Inter-rater agreement	Percentage agreement and Cohen's kappa between independent coders	Blind coding sheet (S3)
Limit-date agreement	Absolute difference between computed and audited limit dates	Accomplishment histories; FH/FC records
Exception ageing	Open exceptions, their age, and share closed as applicable vs not applicable	DISPOSITION history

Measure	Definition	Data required
Schedule feasibility	Share of approved instances scheduled within limit without unplanned ground events	Window plan; executed work packages
Engineering effort	Engineering hours per revision (operator-level unit)	Engineering time records
Operational outcomes	Ground time, AOG events and cost attributable to bridging errors, only where attributable	Operations and AOG logs

7. Conclusions

This paper specified a configuration-aware framework for maintenance program bridging: a data model with anchors, provenance and evidence; seven ordered rules with an explicit exception class; recursive first/later forecasting; and a network-flow MILP that schedules only approved tasks. Four FAA ADs show that applicability can depend on in-service embodiment, repair method and certificate date, and that published costs come in four units, two of them ceilings. Applied to AD 2025-16-06 and extended to the other ADs, the rules exercised every class and matched hand calculation, and the remaining differences from an SB-aware tracker were all due to provenance and coupling. With fleet-wide capacity on generated fleets, the optimization outperformed planner heuristics on most instances and found compliant plans where they failed, although at 30 aircraft a 180-s limit was not always enough to improve on the heuristic start. Whether the framework reduces incorporation latency, reconciliation errors or ground time at an operator remains to be tested with the protocol provided.

Data and Code Availability

Supplementary Material S1 (spreadsheet) contains the extraction matrix with source locations, the synthetic fleet, disposition register, comparator outputs, task list, schedule, scenario, sensitivity and benchmark results. S2 contains the Python code (rules, forecasting, both MILP formulations, heuristics, verification, benchmark generator, figures and the agreement script) with a requirements file and instructions. S3 is the blind second-coder package. All regulatory inputs are public; no proprietary or personal data were used.

Appendix

Appendix 1. Synthetic data, published parameters and hypothetical (HYP) placeholders

Item	Value used	Status
Case-study fleet	8 Airbus single-aisle (A01–A08), 6 Boeing 737NG (B01–B06); one operator; as-of date 1 Oct 2025; horizon 36 months	Synthetic
Utilization	3.8–5.2 FC per day by aircraft; conversion margin $\rho = 1.10$	Synthetic
Configuration and repair histories	Table 5; S1 sheet Synthetic_fleet	Synthetic
AD 2025-16-06 limits and groups	Later of 24 months after 22 Sep 2025 and 17,000 FC after embodiment; Group 1/2 definitions	Published [5, 6]

Item	Value used	Status
AD 2024-18-07 (j) limits	Revision 3 repair: later of 7,000 FC after 12 Nov 2024 and 36,000 total FC; Revision 4 Part 3 repair: later of 24,000 FC after repair and 36,000 total FC; repeat 24,000 FC	Published [7, 8]
Task durations for AD tasks	9 h (C1); 8 h (C2, upper bound); 42 h (C3); 90 h engineering per operator (C4)	FAA cost estimates
AD 2025-15-01 limits	Base: initial first of 3,000 FC and 24 months after 25 Aug 2025; repeat 3,000 FC. Sensitivity: initial 1,500–6,000 FC; repeat 1,500–4,500 FC	HYP (RB 737-53A1410 RB not public)
ALS tasks HYP-T1/T2; routine MPD tasks	ALS 6 h, 365 days; zonal 10 h, 120 days; structural 12 h NDT, 540 days	HYP
Case-study windows	Regular every 60 days (16 h NDT, 24 h mechanical, cost 1); one heavy check (120 h, 300 h, cost 5); additional events every 90 days (48 h, 48 h, cost 50)	Synthetic
Benchmark instances	10, 20, 30 aircraft × 3 seeds × 2 capacity factors; 8 repetitive tasks (90–720 days, 2–24 h) and 2 one-time tasks (6–42 h NDT) per aircraft; regular windows every 30 days (8 h NDT, 12 h mechanical); two heavy checks (60 h, 150 h); additional events every 60 days (48 h, 48 h)	Synthetic (generator in S2)

References

Agency, E. U. (2024). *Airworthiness Directive 2024-0204: ATA 53 – Fuselage – Rivet holes at Frame 64 – Inspection* . (issued 22 October 2024, effective 5 November 2024). https://ad.easa.europa.eu/ad/2024-0204.

Ahmadi, A., Söderholm, P., & Kumar, U. (2010). On aircraft scheduled maintenance program development. *Journal of Quality in Maintenance Engineering*, 16(3), 229–255. https://doi.org/10.1108/13552511011072899.

Airlines for America. (2022). *MSG-3: Operator/manufacturer scheduled maintenance development.* Volume 1 – Fixed wing aircraft (Revision 2022.1). A4A.

Airlines for America. (2023). *Spec 2000, Chapter 11: Reliability data collection/exchange* . (Revision 2023.1). A4A.

Airlines for America. (2023). *Spec 2500: Aircraft transfer records.* (Revision 2023.1). A4A.

Airlines for America. (2025). *iSpec 2200: Information standards for aviation maintenance* . (Revision 2025.1). A4A.

Brundage, M. P., Sexton, T., Hodkiewicz, M., Dima, A., & Lukens, S. (2021). Technical language processing: Unlocking maintenance knowledge. *Manufacturing Letters*, 27, 42–46. https://doi.org/10.1016/j.mfglet.2020.11.001.

Code of Federal Regulations. (n.d.). *Title 14, §121.380 – Maintenance recording requirements. U.S. Government Publishing Office.* https://www.ecfr.gov/current/title-14/chapter-I/subchapter-G/part-121/subpart-V/section-121.380.

Code of Federal Regulations. (n.d.). *Title 14, Part 39 – Airworthiness directives. U.S. Government Publishing Office.* https://www.ecfr.gov/current/title-14/chapter-I/subchapter-C/part-39.

Deng, Q., Santos, B. F., & Curran, R. (2020). A practical dynamic programming based methodology for aircraft maintenance check scheduling optimization. *European Journal of Operational Research,* , 281(2), 256–273. https://doi.org/10.1016/j.ejor.2019.08.025.

Deng, Q.; Santos, B. F.; Verhagen, W. J. C. (2021). A novel decision support system for optimizing aircraft maintenance check schedule and task allocation. *Decision Support Systems*, 146, 113545. https://doi.org/10.1016/j.dss.2021.113545.

European Union Aviation Safety Agency. (2025). *Easy Access Rules for Continuing Airworthiness (Regulation (EU) No 1321/2014).* EASA. https://www.easa.europa.eu/en/document-library/easy-access-rules/online-publications/easy-access-rules-continuing-airworthiness.

Federal Aviation Administration. (2024). Airworthiness directives; The Boeing Company airplanes (AD 2024-18-07, Amendment 39-22841, Docket FAA-2023-1886). *Federal Register*, 89(195), 81320–81324. FR Doc. 2024-23116.

Federal Aviation Administration. (2024). *Electronic signatures, electronic recordkeeping, and electronic manuals (Advisory Circular 120-78B, 11 December 2024).* U.S. : Department of Transportation.

Federal Aviation Administration. (2024). *Maintenance review boards, maintenance type boards, and original equipment manufacturer/type certificate holder recommended maintenance procedures (Advisory Circular 121-22D, 31 May 2024). U.S.* Department of Transp Transportation. https://www.faa.gov/documentLibrary/media/Advisory_Circular/AC_121-22D.pdf.

Federal Aviation Administration. (2025). *Airworthiness directives; Airbus SAS airplanes (AD 2025-02-04, Amendment 39-22937, Docket FAA-2024-2323).* Federal Register. 90(26), 9214–9216. FR Doc. 2025-02375.

Federal Aviation Administration. (2025). Airworthiness directives; Airbus SAS airplanes (AD 2025-16-06, Amendment 39-23104, Docket FAA-2025-0907). *Federal Register*, 90(157), 40028–40030. FR Doc. 2025-15685.

Federal Aviation Administration. (2025). Airworthiness directives; The Boeing Company airplanes (AD 2025-15-01, Amendment 39-23088, Docket FAA-2024-1880). *Federal Register*, 90(137), 34171–34173. FR Doc. 2025-13595.

Federal Aviation Administration. (2025). *Dynamic Regulatory System (DRS).* U.S. : Department of Transportation. https://drs.faa.gov.

Gonçalves, F. C., & Trabasso, L. G. (2018). Aircraft preventive maintenance data evaluation applied in integrated product development process. *Journal of Aerospace Technology and Management*, 10, e1718. https://doi.org/10.5028/jatm.v10.706.

Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research . *MIS Quarterly*, 28(1), 75–105. https://doi.org/10.2307/25148625.

Ho, G. T., Tang, Y. M., Tsang, K. Y., Tang, V., & Chau, K. Y. (2024). A blockchain-based system to enhance aircraft parts traceability and trackability for inventory management. *Expert Systems with Applications*, 179, 115101. https://doi.org/10.1016/j.eswa.2021.115101.

Huangfu, Q., & Hall, J. A. (2018). Parallelizing the dual revised simplex method. *Mathematical Programming Computation*, 10(1), 119–142. https://doi.org/10.1007/s12532-017-0130-5\.

Kang, Z., Zhou, D., Guo, Z., Zhou, Q., & Wu, H. (2024). An ontology-based knowledge representation framework for aircraft maintenance processes to support work optimization. *The International Journal of Advanced Manufacturing Technology*,, 134(11), 5577–5601. https://doi.org/10.1007/s00170-024-14428-4.

Kinnison, H. A., & Siddiqui, T. (2013). Aviation maintenance management (2nd ed.). . *McGraw-Hill Education*, ISBN 978-0-07-180502-5.

Lagos, C., Delgado, F., & Klapp, M. A. (2020). Dynamic optimization for airline maintenance operations. *Transportation Science*, 54(4), 998–1015. https://doi.org/10.1287/trsc.2020.0984.

Lampe, M., Strassner, M., & Fleisch, E. (2024). A ubiquitous computing environment for aircraft maintenance. *In Proceedings of the 2004 ACM Symposium on Applied Computing (SAC '04)*, (pp. 1586–1592). ACM. https://doi.org/10.1145/967900.968217.

Lercel, D., Patankar, M., & Steckel, R. (2024). Assessing past airworthiness directives and how safety management systems may benefit aviation product design and manufacturing. *Journal of Aviation/Aerospace Education & Research*, 33(2), Article 6. .

Mitchell, S., O'Sullivan, M., & Dunning, I. (2011). PuLP: A linear programming toolkit for Python. . *Department of Engineering Science, The University of Auckland*.

Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, 24(3), 45–77. https://doi.org/10.2753/MIS0742-1222240302.

Qin, Y., Zhang, J. H., Chan, F. T., Chung, S. H., Niu, B., & Qu, T. (2020). A two-stage optimization approach for aircraft hangar maintenance planning and staff assignment problems under MRO outsourcing mode. *Computers & Industrial Engineering*, 146, 106607. https://doi.org/10.1016/j.cie.2020.106607.

Ramco Systems. (2026). *How agentic AI for aviation maintenance transforms SB & AD processing [Vendor article]*. https://www.ramco.com/blog/aviation/agentic-ai-for-aviation-maintenance.

S1000D Steering Committee (ASD, AIA & ATA). (2024). *S1000D: International specification for technical publications using a common source database (Issue 6).* https://www.s-series.org/s1000d/.

SAE International. (2019). *Configuration management standard (SAE EIA-649C).* SAE International.

Seawright, J., & Gerring, J. (2008). Case selection techniques in case study research: A menu of qualitative and quantitative options. *Political Research Quarterly*, 61(2), 294–308. https://doi.org/10.1177/1065912907313077.

Shaukat, S., Katscher, M., Wu, C. L., Delgado, F., & Larrain, H. (2020). Aircraft line maintenance scheduling and optimisation. *Journal of Air Transport Management*, 89, 101914. https://doi.org/10.1016/j.jairtraman.2020.101914.

Sriram, C., & Haghani, A. (2003). An optimization model for aircraft maintenance scheduling and re-assignment. *Transportation Research Part A: Policy and Practice*, 37(1), 29–48. https://doi.org/10.1016/S0965-8564(02)00004-6.

Swiss AviationSoftware. (2025). *AMOS Engineering module; AMOS Modification Control – Part 1 (training course description).* https://swiss-as.com/modules/engineering; https://www.swiss-as.com/public/training/AMOS_Modification_Control_Part_1.pdf.

The Boeing Company. (2022). *Alert Service Bulletin 737-25A1544, Revision 4 (15 February 2022).* . Placed in FAA docket FAA-2023-1886. https://downloads.regulations.gov/FAA-2023-1886-0002/attachment_1.pdf.

Witteman, M., Deng, Q., & Santos, B. F. (2021). A bin packing approach to solve the aircraft maintenance task allocation problem. *European Journal of Operational Research*, 294(1), 365–376. https://doi.org/10.1016/j.ejor.2021.01.027.

Yue, S., Xiao, L., Li, J., & Wang, N. (2022). Research on application of knowledge graph for aircraft maintenance. *Advances in Mechanical Engineering*, 14(7), 1–13. https://doi.org/10.1177/16878132221107429.

Figure Legends

Fig. 1. Aircraft maintenance program development and regulatory control. MSG-3 analysis produces Maintenance Review Board Report (MRBR) and Maintenance Planning Document (MPD) recommendations, which the operator adapts into its approved program and tracks for each aircraft; Airworthiness Directives (ADs) and airworthiness limitations enter separately as mandatory inputs.

Fig. 2. Aircraft configuration and task effectivity verification. The aircraft serial number, Service Bulletin (SB) embodiment status and component configuration are checked before a task is classed as applicable, not applicable or requiring engineering review.

Fig. 3. IMPB-TDO framework architecture. Five stages, with their inputs (left) and outputs (right): stages 1–4 are rule-based and produce the disposition register (rules R1–R7); stage 5 is the scheduling optimization, which receives only approved task instances. Unresolved exceptions and scheduling infeasibilities return to engineering review (dashed).

Fig. 4. Maintenance program revision and bridging workflow for one source revision, showing the decision rules applied at each step. Engineering review resolves exceptions and records questions; a bridging package is approved only when every requirement–aircraft pair has a disposition with evidence.

Fig. 5. IMPB-TDO relational data model. Blue: technical-data layer; green: configuration and records layer; orange: disposition (status-accounting) record; purple: scheduling layer. PK, primary key; FK, foreign key; FH, flight hours; FC, flight cycles; CAL, calendar. Arrows show one-to-many relationships or evidence links; only approved dispositions can be assigned to windows.

Fig. 6. Ordered decision rules for classifying a requirement–aircraft pair. Applicability (R1–R2) is evaluated in three-valued logic; an UNKNOWN configuration predicate yields an engineering exception, never "not applicable". Supersession and terminating actions (R6) are tested before lifecycle classification (R3); limit logic (R4), coupling (R5) and hard mandatory limits (R7) are then applied.

Fig. 7. Worked example of AD 2025-16-06 on eight synthetic Airbus aircraft (as of 1 October 2025). For Group 1 aircraft, red ticks show term A (24 months after the 22 September 2025 effective date), blue ticks the forecast of term B (embodiment flight cycles + 17,000) and green circles the resulting limit (later of A and B); A01's term B (May 2021) lies before the plotted range. The orange circle is the provisional limit for the engineering exception, and the purple diamond the planned modification that triggers rivet installation on A04. Shading marks the 36-month planning horizon. Aircraft data are synthetic.

Fig. 8. Optimal maintenance plan for scenario S1 (synthetic 14-aircraft fleet, 1 October 2025 to 30 September 2028). Each row is an aircraft; markers show planned task instances by type, red ticks the limits of first AD instances and open circles provisional tasks raised by engineering exceptions. No additional ground events are required in S1.

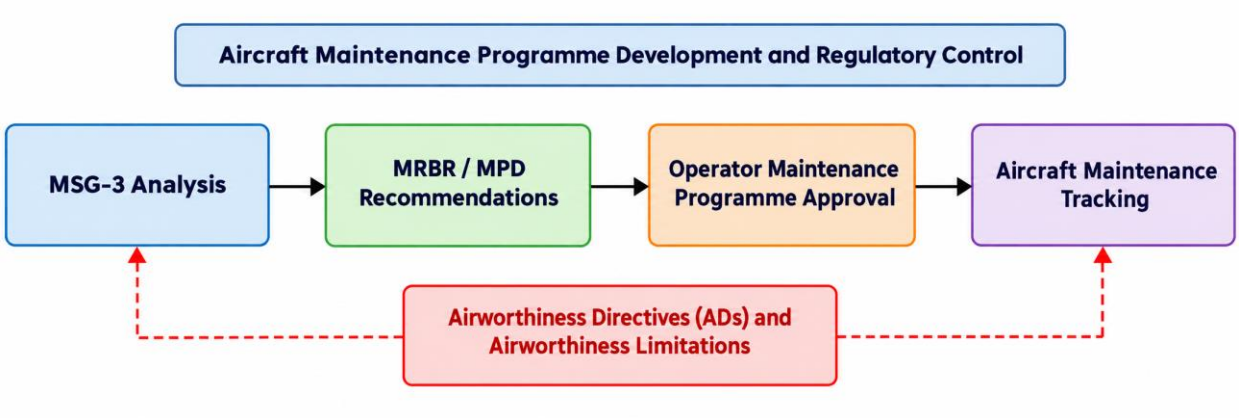


Fig. 1.

International Journal of Sustainability and Innovation in Engineering (IJSIE)
https://www.doi.org/10.56830/IJSIE2026
Vol. 4 2026

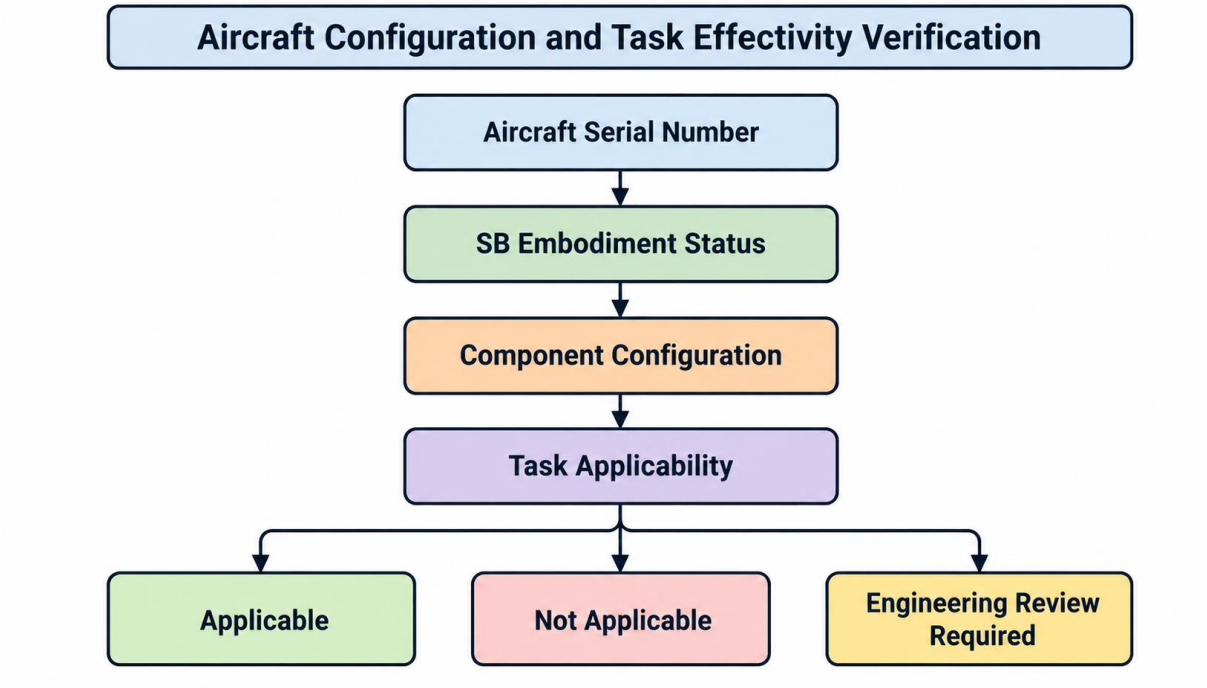


Fig. 2.

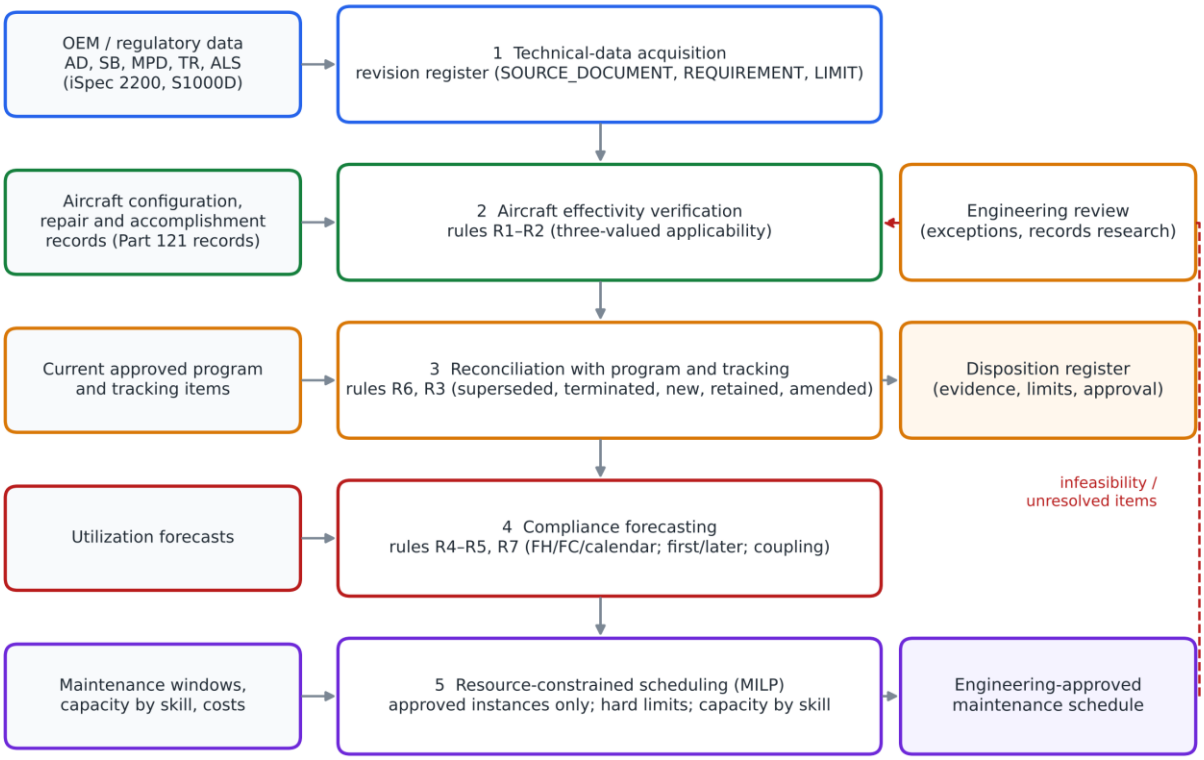


Fig. 3.

International Journal of Sustainability and Innovation in Engineering (IJSIE)
https://www.doi.org/10.56830/IJSIE2026
Vol. 4 2026

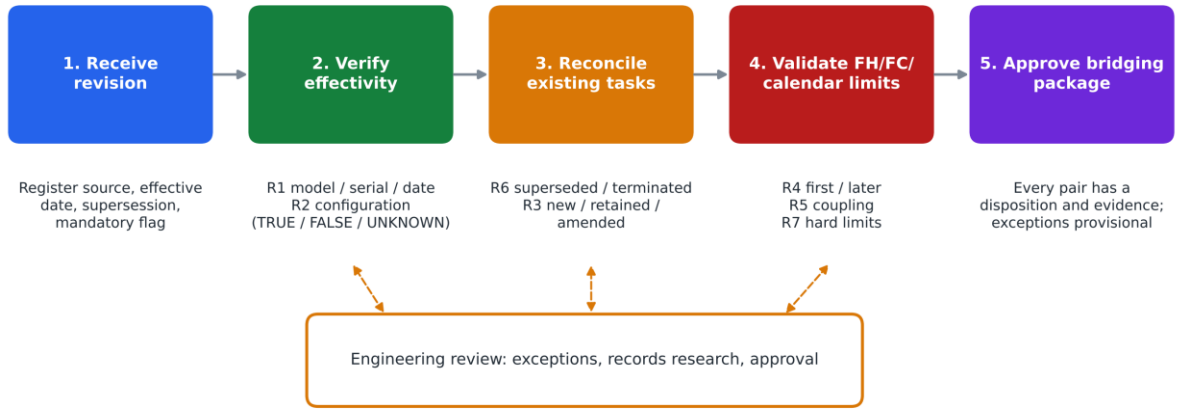


Fig. 4.

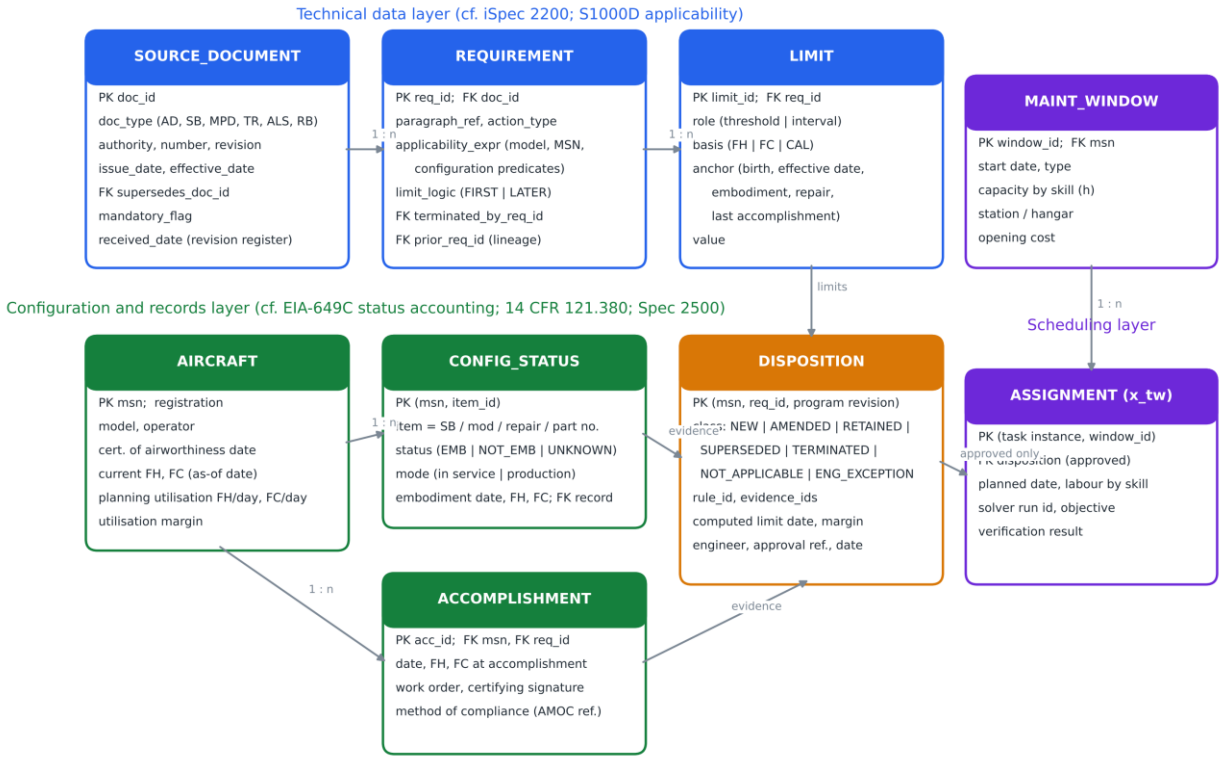


Fig. 5.

International Journal of Sustainability and Innovation in Engineering (IJSIE)
https://www.doi.org/10.56830/IJSIE2026
Vol. 4 2026

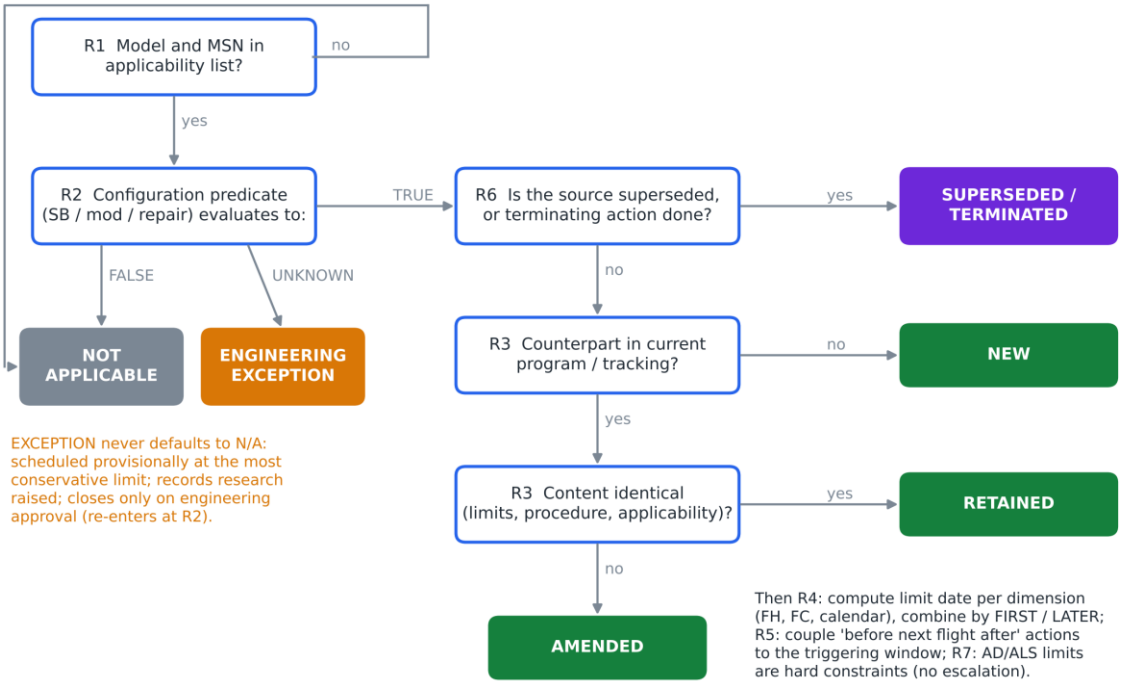


Fig. 6.

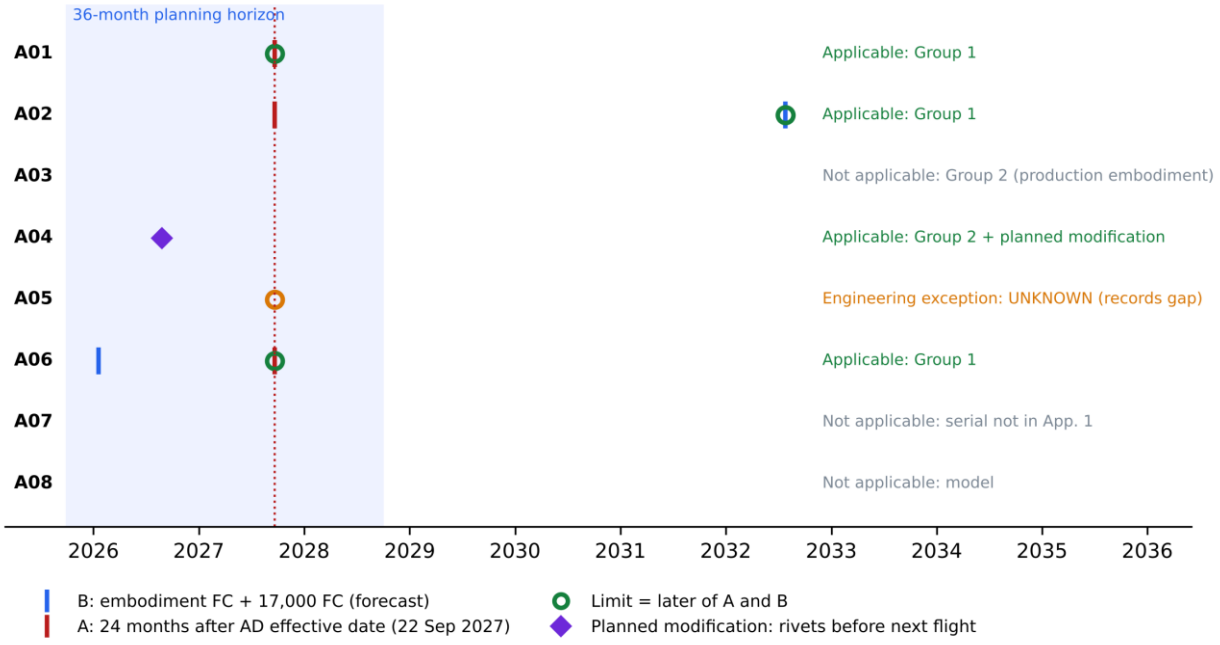


Fig. 7.

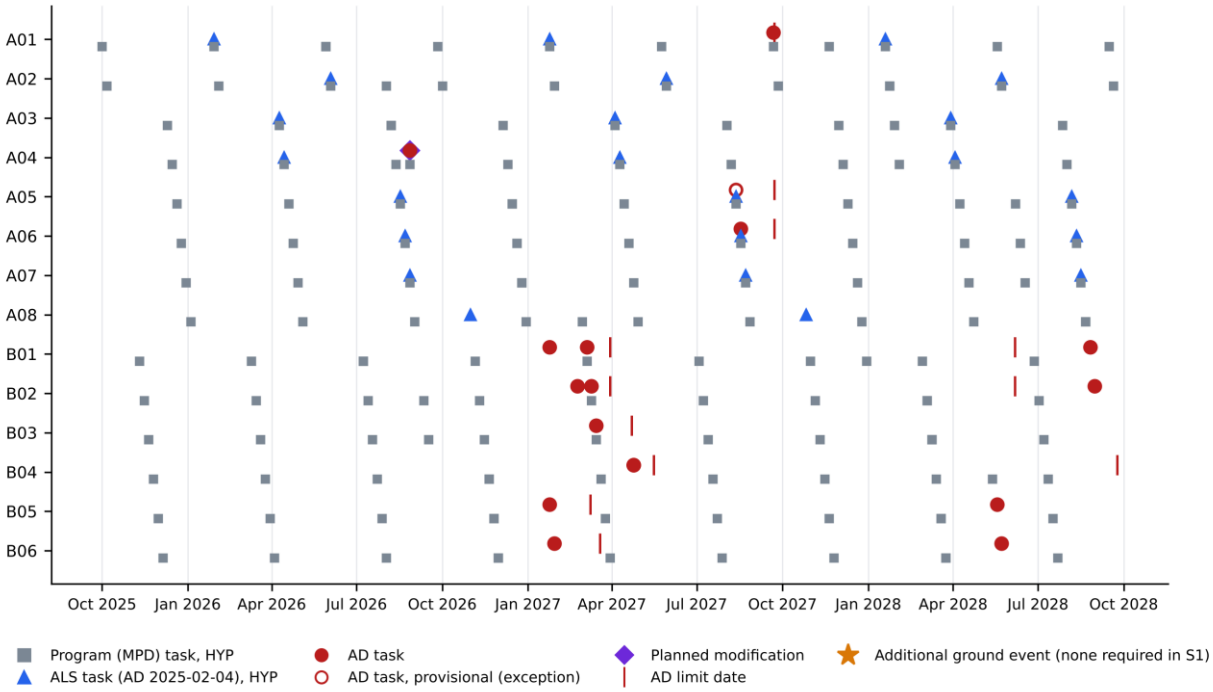

A01
A02
A03
A04
A05
A06
A07
A08
B01
B02
B03
B04
B05
B06
Oct 2025
Jan 2026
Apr 2026
Jul 2026
Oct 2026
Jan 2027
Apr 2027
Jul 2027
Oct 2027
Jan 2028
Apr 2028
Jul 2028
Oct 2028
Program (MPD) task, HYP
ALS task (AD 2025-02-04), HYP
AD task
AD task, provisional (exception)
Planned modification
AD limit date
Additional ground event (none required in S1)

Fig. 8.