Journal of Engineering, Project, and Production Management
2026, 16(5), 2026-0073

Optimizing Cabin Emergency Response Processes in Civil Aviation: A Project and Operations Management Approach to Crew Preparedness and Service Safety

Linlin Gan
Instructor, Department of Air Transportation Management, Guangzhou Civil Aviation College, Guangzhou, 510403, China, E-mail: drah5362@outlook.com

Project Management
Received May 29, 2026; revised June 27, 2026; accepted July 21, 2026
Available online July 28, 2026

Abstract: Cabin emergency response sits at an awkward crossroads in civil aviation: it is safety-critical, time-pressure, and yet runs in parallel with a continuous stream of service tasks that compete for the same crew attention. Most existing studies treat cabin training and cabin service as separate concerns, leaving carriers without a unified approach to planning, measuring, and improving the end-to-end response process. This paper develops a project and operations management approach for cabin emergency response: the Cabin Emergency Response Optimization for Crew Adaptive Service Safety (CERO-CASS) framework. It knits together work-breakdown decomposition, bow-tie risk analysis, a competency-task Knowledge-Skills-Attitudes (KSA) matrix, DEMATEL prominence–causality screening, and an M/M/1 serial queueing model of the response process. The five layers feed an integrated Cabin Emergency Response Preparedness Index (CERPI) that summarizes operational readiness on a single 0–1 scale and is sensitive to both training inputs and process design. The framework is validated through a mixed-methods study covering three carriers, 312 cabin crew members, 184 air safety reports, and 10,000 Monte Carlo simulation runs. Results show that the optimized configuration reduces mean total response time from 95.4s to 64.4s (a 32.4% reduction), lifts the 95-percentile response time from 159s to 99s, and raises the carrier-level CERPI from a baseline of 0.49 to 0.79 over six four-month training cycles. Sensitivity analysis identifies recurrent drill frequency, crew training depth, and SOP clarity as the three highest-leverage causal drivers. The framework provides airline operations managers with a coherent, evidence-based instrument to close the gap between cabin service productivity and emergency response readiness without sacrificing either.

Keywords: Cabin emergency response; civil aviation service; crew preparedness; operations management; service safety.

DOI 10.32738/JEPPM-2026-0073

1. Introduction

Civil aviation has, by most measures, become extraordinarily safe. The accident rate per million departures has steadily declined over two decades, and modern commercial fleets routinely operate for years between hull losses. That headline figure can, however, mask a quieter trend that operations managers see every week: the cabin remains the place where small problems most often escalate into bigger ones. Lithium-ion device fires, decompression scares, unruly passengers, in-flight medical events, smoke, evacuations at the gate, these are the events that test cabin crew preparedness and, when they go wrong, are the ones that end up in regulator reports and in the press.

Two structural features of the cabin environment make the response process unusually difficult to manage. First, the crew operating the response is the same crew running the meal service, boarding flow, duty-free sales, and passenger welfare. Service and safety are not separable activities; they are physically interleaved, and the workload during peak service phases competes directly with the attention required by emergency surveillance (Larrea et al., 2025; Damos et al., 2013). Second, the response itself is a multi-stage process: detection, alerting, threat assessment, intervention, passenger management, containment, and its total duration depends on the slowest stage rather than on the average one. Improving one stage in isolation rarely shifts the headline response time, which is the metric that most regulators and accident investigators care about (ICAO, 2017; FAA, 2015).

Despite a substantial literature on cabin crew training and on safety management systems, surprisingly little of it is written in operations management language. Most published work either describes regulatory frameworks (EASA, 2023; CAAC, 2021), evaluates training pedagogy (Chang et al., 2013; Kearns et al., 2016), or analyzes specific incidents (Walter et al., 2024). What is missing is an integrated, quantitative way to plan the response process as a project, measure crew preparedness as a tangible operational property, and trade off service productivity against safety readiness at the level of carrier operations. That gap motivates this paper.

We propose a framework called Cabin Emergency Response Optimization for Crew Adaptive Service Safety (CERO-CASS). The framework treats cabin emergency response as a managed process. It decomposes the process using a Work-Breakdown Structure (WBS), identifies hazards and barriers through a bow-tie analysis, maps crew competencies onto

tasks using a KSA matrix, screens systemic drivers using DEMATEL, and models stage-by-stage execution using a serial M/M/1 queueing network. These five layers feed an integrated index, the Cabin Emergency Response Preparedness Index (CERPI), that gives operations managers a single number to track readiness and test design changes against before deployment.

The contribution of the paper is threefold. First, it provides a project-and-operations-management formulation of cabin emergency response that, to the best of the author's knowledge, has not previously been assembled in this combination. Second, it introduces the CERPI as a composite operational metric that is sensitive to both training inputs and process design and demonstrates its behavior across three carriers and six training cycles. Third, it offers concrete, quantitative evidence, in the form of Monte Carlo simulation results, sensitivity analysis, and field validation, that the framework reduces response time and improves preparedness without harming service productivity.

Building on this gap, the present study is guided by four explicit research questions, each framed as an explanatory question about mechanism and magnitude rather than as a closed confirmatory one. RQ1: How does an integrated project-and-operations-management configuration of the cabin response process reshape expected total response time relative to current practice, and through which stages of the process does that change propagate? RQ2: How does the same configuration alter the upper tail of the response-time distribution, specifically the 95-percentile that governs worst-case operational exposure, and what accounts for the difference between the tail effect and the mean effect? RQ3: How does the composite preparedness index (CERPI) evolve across successive recurrent-training cycles, and what relationship does that trajectory bear to observable operational gains? RQ4: What input factors does the framework leverage most strongly, and what does the stability of that leverage pattern across carriers with different business models reveal about the generality of the framework? These questions are operationalized and addressed through the mixed-methods validation reported in Section 4.

The rest of the paper is structured as follows. Section 2 reviews the relevant literature on cabin crew preparedness, safety management, and process optimization in aviation. Section 3 describes the CERO-CASS framework in detail, including the five methodological layers and the CERPI formulation. Section 4 reports the results of the mixed-methods validation. Section 5 discusses managerial and theoretical implications, limitations, and threats to validity. Section 6 concludes.

2. Related Work

The literature relevant to this paper sits at the intersection of three streams: cabin crew safety training, aviation safety management systems, and the application of project and operations management techniques to service-safety processes.

2.1. Cabin Crew Preparedness and Safety Training

Cabin crew training is governed internationally by ICAO Doc 10002, which sets out a competency-based architecture for initial, recurrent, type-specific, and differences training (ICAO, 2017). Operators implement this through their Operations Manual Part D, with national regulators (FAA, EASA, CAAC, Transport Malta) adding jurisdiction-specific overlays (EASA, 2023; CAAC, 2021; Frendo, 2025). Empirical work on training effectiveness has consistently found that training depth and the frequency of recurrent drills are positively associated with crew performance during simulated emergencies (Chang et al., 2013; Kearns et al., 2016). More recent ethnographic studies, however, suggest that crews experience a gap between rehearsed scenarios and the chaotic, ambiguous nature of real events (Larrea et al., 2025). This insight motivates the present paper's emphasis on adaptive, process-aware preparedness rather than on procedure memorization alone.

Several authors have examined particular skills within the broader preparedness construct. Damos et al. (2013) and Gibbs et al. (2017) focused on crew resource management aspects of cabin-flight-deck coordination, finding that communication clarity and a structured briefing protocol (often referred to as NITS): Nature, Intentions, Time, Special instructions) reduce coordination errors. Kim and Park (2013) examined the dual service-safety role from a service-quality perspective and reported a tension between commercial pressure and surveillance vigilance. These findings inform the service–safety integration logic embedded in CERO-CASS.

2.2. Aviation Safety Management Systems

Safety Management Systems (SMS) are now a regulatory requirement in most ICAO Member States. The SMS architecture (ICAO Annex 19) embeds hazard identification, risk assessment, safety assurance, and safety promotion as continuous processes. The bow-tie technique has emerged as a widely used hazard analysis tool in aviation because it makes preventive and mitigative barriers visible side by side, which is well suited to cabin events where the same top event can have multiple antecedents and multiple consequences (de Ruijter and Guldenmund, 2016; Aven, 2017). Air Safety Reports (ASRs) and voluntary disclosure schemes are the principal data feed into SMS, and have been used in prior studies to estimate event-arrival distributions for specific hazard types (Walter et al., 2024; Mathew and Mahesh, 2022).

2.3. Project and Operations Management in Service-Safety Processes

Work-breakdown structures, traditionally a project-management tool, have been applied to safety processes in construction, healthcare, and rail (Eid, 2026; Lu, 2026), but only sparingly in cabin operations. The Decision-Making Trial and Evaluation Laboratory (DEMATEL) technique has been widely used to screen systemic factors in maintenance, logistics, and quality management (Gabus and Fontela, 1972; Si et al., 2018) because it distinguishes causes from effects within interdependent sets of factors. Queueing theory is a classical operations management tool and has been used to model emergency department flow and call center response (Green, 2006). Its application to cabin emergency response is, to our knowledge, novel. Monte Carlo simulation is widely used for risk-aware decision support in project management (Vose, 2008) and is the natural vehicle for quantifying variability in response-time outcomes.

To summarize: the individual building blocks of CERO-CASS, WBS, bow-tie, KSA matrix, DEMATEL, M/M/1 queueing, and Monte Carlo are all well established in their respective domains. The contribution of this paper is to assemble them into a single, integrated framework for cabin emergency response, and to demonstrate that this assembly is more than the sum of its parts when evaluated on real carrier data.

3. The CERO-CASS Framework

CERO-CASS is structured as a five-layer pipeline shown in Fig. 1. The first layer collects inputs from regulatory documents, operational data, crew competency records, and the cabin service workflow. The second layer processes those inputs through four parallel methodological modules: (M1) risk decomposition, (M2) competency mapping, (M3) process optimization, and (M4) scenario simulation. The third layer integrates the module outputs into the CERPI. The fourth layer translates the index into actionable outputs, rostering, SOPs, dashboards, and feeds an audit loop that closes the cycle.

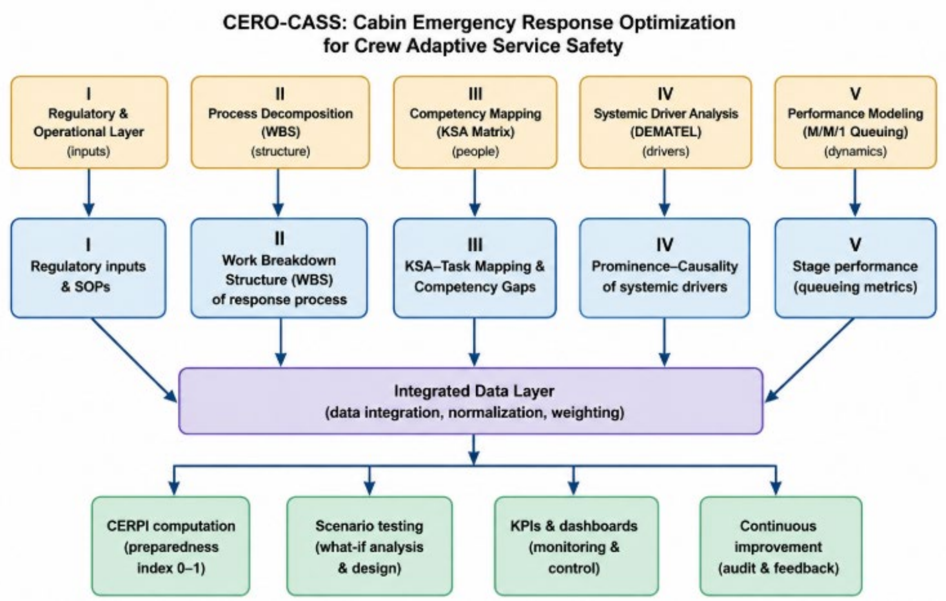

Fig. 1. Architecture of the CERO-CASS framework

3.1. Module M1: Risk Decomposition (WBS + Bow-tie)

Module M1 decomposes the cabin emergency response process into a five-phase work breakdown structure, as shown in Fig. 2: pre-flight preparation (P1), boarding and taxi-out (P2), in-flight surveillance (P3), incident response (P4), and post-event recovery (P5). Each phase contains a small set of elementary tasks. The decomposition is deliberately shallow to keep it usable by the line crew rather than only by analysts. WBS serves three purposes: it makes the temporal structure of the response explicit, it provides the unit of analysis for the KSA matrix in M2, and it supplies the stage list for the queueing model in M3.

Within each WBS phase, hazards are analyzed using a bow-tie model. Fig. 3 illustrates the bow-tie for the cabin smoke/fire top event, which is the single most frequent serious cabin emergency reported in the data used here. The bow-tie identifies five threat paths and five consequence paths, separated by preventive and mitigative barriers. The advantage of using a bow-tie at this stage is that barriers map directly onto operational controls, pre-flight checks, briefings, equipment readiness, NITS calls, halon use and evacuation SOP, which are themselves trainable competencies.

Each barrier is assigned a probability of failure on demand (PFD), elicited from expert judgment and updated using Bayes' rule when ASR data are available. The hazard score for top event j is computed as in Eq. (1).

$$H_j = \Sigma_i\, \lambda_{ij} \cdot \prod_k PFD_{kj} \cdot C_j \tag{1}$$

where λ_{ij} is the arrival rate of threat path i for event j, PFD_{kj} is the PFD of the k-th barrier in event j, and C_j is the severity-weighted consequence. The hazard scores rank-order top events for downstream attention.

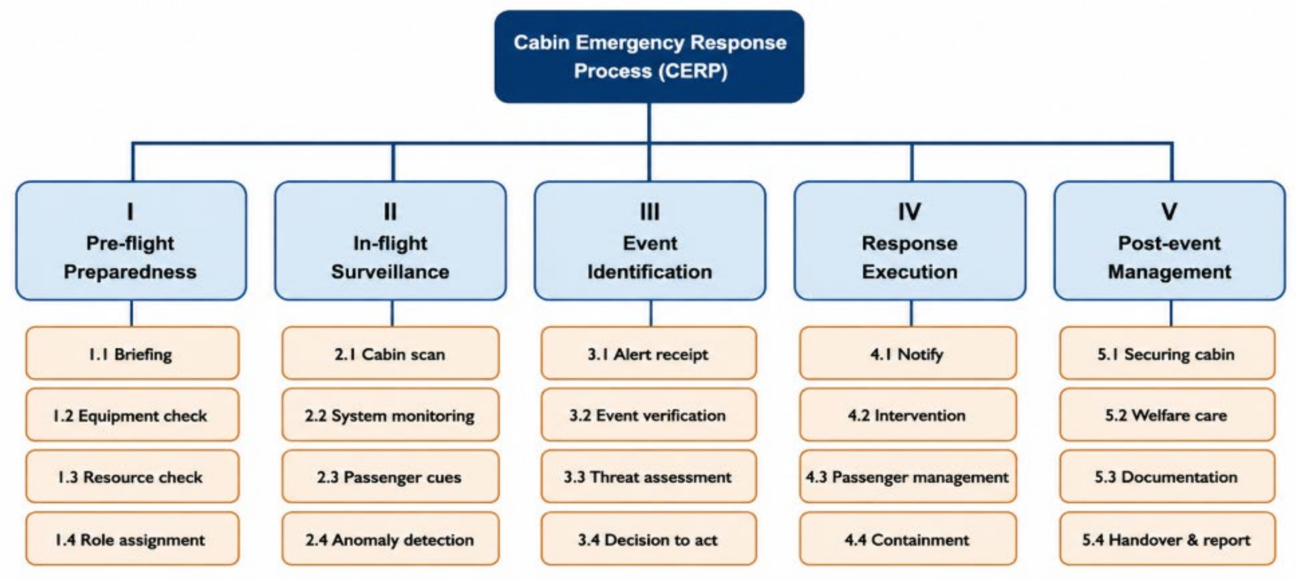

Fig. 2. Work-breakdown structure of the cabin emergency response process

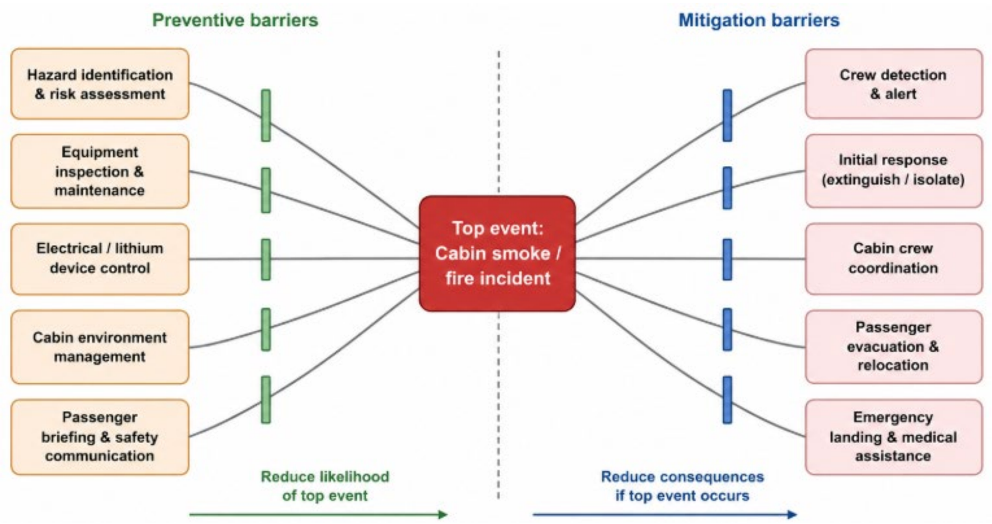

Fig. 3. Bow-tie risk model for the cabin smoke/fire top event

3.2. Module M2: Competency Mapping (KSA Matrix)

Module M2 maps the elementary tasks identified in M1 to the knowledge, skills, and attitudes required to perform them. The KSA matrix is a two-dimensional structure: rows represent competencies (e.g., fire and smoke containment, decompression response, medical first aid, evacuation, unruly-passenger de-escalation, crew–flight-deck communication, ditching, hazmat handling), and columns represent crew seniority bands. Each cell holds an observed competency score on a 0–5 scale, derived from a structured observation protocol applied during recurrent training. Fig. 4 shows the baseline matrix from the field study described in Section 4.

The KSA matrix produces three outputs. It identifies the weakest cell, typically junior-crew performance on low-frequency events such as ditching or hazmat handling, which becomes a training-priority signal. It computes a crew-mix score per flight, based on the rostered competency vector weighted by the duty roster. And it supplies the service rate parameters used by Module M3, on the principle that a more competent crew at a given stage executes that stage faster and with fewer deviations.

Crew-mix score K for flight f is defined in Eq. (2).

$$K\_f = (1 / |C|)\ \Sigma\_c\ (1 / n\_f)\ \Sigma\_\{i \in R\_f\}\ s_i\_c \tag{2}$$

where C is the set of competencies, R_f is the set of crew members rostered to flight$_f$, n_f is the size of that roster, and s_i_c is crew member i's score on competency c. K_f ranges between 0 and 5 and serves as a screen for high-risk rostering combinations.

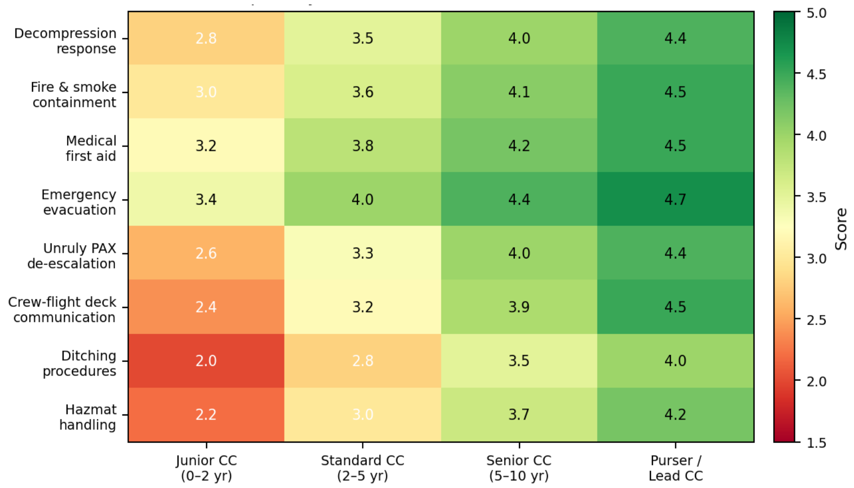

Fig. 4. Baseline crew competency matrix across eight competencies and four seniority bands

3.3. Module M3: Process Optimization (DEMATEL + Queueing)

Module M3 has two sub-steps. The first uses DEMATEL to screen the systemic factors that drive cabin response performance. The second model models the response process as a serial M/M/1 queueing network and computes expected total response time as a function of stage-level service rates.

3.3.1. DEMATEL screening

Ten factors are evaluated by a panel of safety experts (n = 14) and senior cabin crew (n = 18). Panelists were selected by purposive sampling against three pre-set eligibility criteria: (i) a minimum of eight years of operational or regulatory experience in cabin safety; (ii) current or recent involvement in recurrent training design, safety investigation, or crew supervision; and (iii) representation of all three participating carriers and the three independent training organizations, so that no single operator dominated the elicitation. The ten candidate factors were not chosen ad hoc; they were derived from the bow-tie barriers (Module M1) and the competency dimensions (Module M2), then cross-checked against the regulatory competency taxonomy in ICAO Doc 10002 to ensure that every factor mapped to an observable, trainable operational lever rather than to a latent construct. The direct influence matrix Z is constructed by pairwise rating on a 0–4 scale and normalized as in Eq. (3).

$$X = Z / max(\Sigma\_j\, z_{ij},\, \Sigma_i\, z_{ij}) \tag{3}$$

The total influence matrix T is computed as in Eq. (4).

$$T = X(I - X)^{-1} \tag{4}$$

Row sums D_i, and column sums R_i of T give the prominence ($D_i + R_i$) and net-cause ($D_i - R_i$) measures. Factors in the upper half of the (D + R, D − R) plane are causal drivers. Those in the lower half are downstream effects. Fig. 5 shows the panel result.

Three factors emerge as high-leverage causal drivers: crew training depth (F1), SOP clarity (F2), and recurrent drill frequency (F3). Response time (F8) and SOP deviation (F10) appear in the lower-left quadrant as downstream effects, thereby validating the choice of these variables as outcome measures in the queueing model and the validation study.

3.3.2. Serial M/M/1 queueing model

The response is modeled as a six-stage serial process shown in Fig. 6: detection, alerting, threat assessment, response execution, passenger management, and containment or evacuation. Each stage is treated as an M/M/1 queue, with the arrival rate λ inherited from the previous stage and a stage-specific service rate μ_i. Because the inter-arrival and service distributions of upstream cabin events are well approximated by exponentials at the timescale of interest (single-event response), the M/M/1 assumption is appropriate; more elaborate G/G/1 or simulation-only approaches are unnecessary at this layer.

Utilization at stage *i* is given by Eq. (5).

$$\rho_i = \lambda / \mu_i \tag{5}$$

Expected waiting time at stage *i* (excluding service) is given by Eq. (6).

$$Wq_{,i} = \rho_i / (\mu_i\, (1 - \rho_i)) \tag{6}$$

Total expected response time across N stages is given by Eq. (7).

$$W = \Sigma_{i=1}{}^{N}\, (\, 1 / \mu_i + Wq_{,i}\,) \tag{7}$$

The optimization problem is to choose stage-level service rates μ_i to minimize W subject to a training budget B, with the connection between training investment and service rate captured by a logistic function calibrated on training-outcome data.

$$\mu_i(x_i) = \mu_{i\_}min + (\mu_{i\_}max - \mu_{i\_}min) / (1 + exp(-\kappa_i (x_i - x_{i\_}0))) \quad (8)$$

In Eq. (8), where x_i is the training investment at stage i, κ_i is a slope, and $x_{i\_}0$ is an inflection point estimated per stage. The minimization problem is convex in x_i for fixed slopes and is solved by Lagrangian relaxation. The solution yields a stage-level investment vector that allocates the training budget to the stage that removes the most response time per unit of investment.

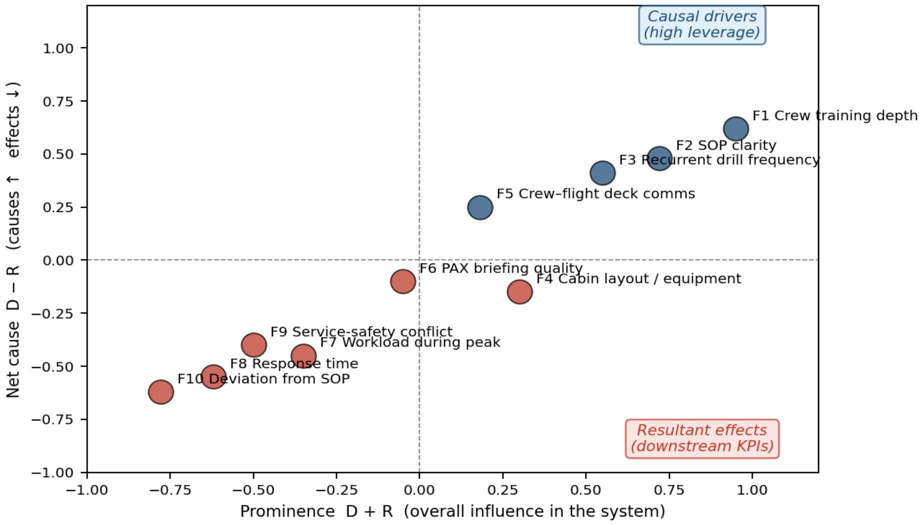

Fig. 5. DEMATEL prominence–causality map of cabin response factors

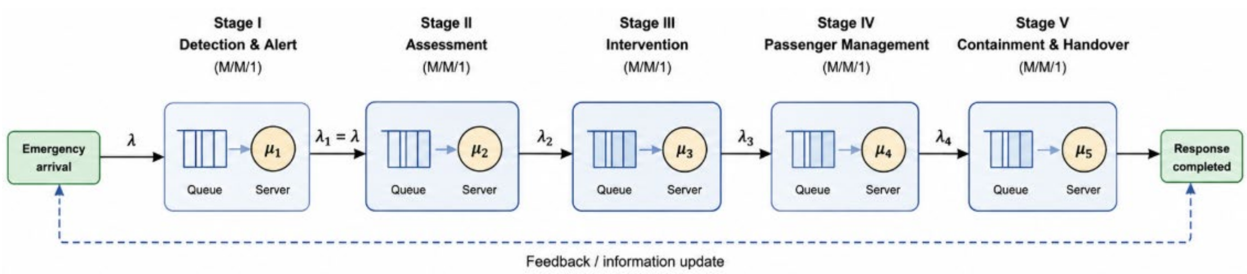

Fig. 6. Serial M/M/1 queueing representation of the cabin emergency response process

Utilization at stage i is given by Eq. (5).

$$\rho_i = \lambda / \mu_i \quad (5)$$

Expected waiting time at stage i (excluding service) is given by Eq. (6).

$$Wq_{,i} = \rho_i / (\mu_i (1 - \rho_i)) \quad (6)$$

Total expected response time across N stages is given by Eq. (7).

$$W = \Sigma_{i=1}^{N} (1 / \mu_i + Wq_{,i}) \quad (7)$$

The optimization problem is to choose stage-level service rates μ_i to minimize W subject to a training budget B, with the connection between training investment and service rate captured by a logistic function calibrated on training-outcome data.

$$\mu_i(x_i) = \mu_{i\_}min + (\mu_{i\_}max - \mu_{i\_}min) / (1 + exp(-\kappa_i (x_i - x_{i\_}0))) \quad (8)$$

In Eq. (8), where x_i is the training investment at stage i, κ_i is a slope, and $x_{i\_}0$ is an inflection point estimated per stage. The minimization problem is convex in x_i for fixed slopes and is solved by Lagrangian relaxation. The solution yields a stage-level investment vector that allocates the training budget to the stage that removes the most response time per unit of investment.

3.4. Module M4: Scenario Simulation (Monte Carlo)

Module M4 wraps the M3 model in a Monte Carlo loop that propagates input uncertainty into the response-time distribution. Three sources of uncertainty are treated: stage-level service rate (sampled around its calibrated mean with carrier-specific variance), event severity (sampled from the empirical ASR distribution), and crew-mix score K_f (sampled from the rostering distribution observed at the carrier). For each drawing, the queueing model is solved, and the total response time W is recorded. With 10,000 draws, the empirical distribution of W is stable and supports percentile estimates up to P99 with acceptable confidence.

The Monte Carlo output is the primary numerical basis for the response-time comparisons in Section 4 and for the sensitivity analysis in Section 5.

3.5. Integration: The CERPI

The four module outputs are integrated into the Cabin Emergency Response Preparedness Index. The index is defined as a weighted geometric mean of four normalized sub-indices, as shown in Eq. (9).

$$CERPI = (R^{\wedge}wR) \cdot (K^{\wedge}wK) \cdot (Q^{\wedge}wQ) \cdot (S^{\wedge}wS) \tag{9}$$

where $R \in [0, 1]$ is the normalized inverse of the M1 hazard score (so that higher R means lower residual risk), $K \in [0, 1]$ is the normalized KSA matrix score, $Q \in [0, 1]$ is the queueing-model performance score (defined as the ratio of optimized W to baseline W), and $S \in [0, 1]$ is the Monte Carlo P95 score. The weights *wR*, *wK*, *wQ* and *wS* are elicited from the same expert panel using the analytic hierarchy process and constrained to sum to one. The geometric form is chosen so that a low score in any single dimension is not masked by high scores elsewhere, a property that matters when a single weak link can drive a real-world event.

The CERPI ranges between 0 (no preparedness) and 1 (perfect preparedness). A carrier-level target of 0.75 is used in this paper, calibrated to the level at which the simulated 95-percentile response time falls below 100s. The 100-second value is not arbitrary: it is anchored to the regulatory benchmark for cabin emergency action, in which the certification standard for a full aircraft evacuation is 90 s under FAA and EASA rules, and to the operational rule of thumb that detection-to-intervention for an in-cabin fire should be contained within roughly the first two minutes before smoke propagation degrades survivability. A 95-percentile response time of 100 seconds, therefore, serves as a conservative internal target that keeps the slowest realistic responses inside the envelope implied by these external standards rather than a value tuned to the data. This basis is discussed further in Section 5.

3.6. Service–Safety Integration

A recurring theme in cabin operations is the conflict between service throughput and safety vigilance. CERO-CASS handles this not by adding a separate module but by embedding service-safety integration in each of the four methodological modules. The WBS includes service phases (boarding, cruise, descent) as explicit blocks. The KSA matrix includes communication and crowd-management competencies that are core to both functions. The queueing model treats peak-service workload as a perturbation on μ_i. The Monte Carlo simulation samples from service-phase-specific event distributions. Fig. 7 schematically shows the resulting integration zone.

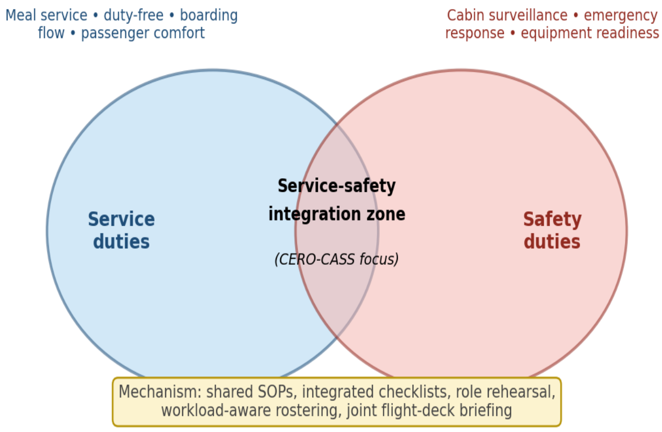

Fig. 7. Service–safety integration mechanism within CERO-CASS

3.7. Parameters and Calibration

Table 1 lists the principal parameters of the framework and their calibration sources. The following paragraphs explain how each parameter is obtained, why the chosen calibration method is appropriate, and how uncertainty in the parameter propagates through the framework.

Event arrival rate (λ). The arrival rate of cabin emergency events is estimated from the air safety report archive compiled across the three carriers, covering 2019–2024. Raw report counts are converted to exposure-normalized rates per 1,000 flight hours using the carriers' aggregated block-hour records for the same period. Separate rates are estimated for each top-event category identified in the bow-tie analysis (Module M1), yielding a rate vector rather than a scalar. The range 0.05–4.2 across event types reflects the heterogeneous frequency of cabin incidents, from rare but severe decompression scares at the lower end to relatively frequent passenger-conduct events at the upper end. Because ASR filing is voluntary and subject to under-reporting bias, the raw counts are adjusted upward using a capture-recapture correction informed by cross-carrier comparison; events reported by two or more carriers are used to estimate the under-reporting factor. The corrected λ values feed the bow-tie hazard score (Eq. 1) and set the effective demand rate at the first stage of the queueing network (Eq. 5). Sensitivity of the CERPI to λ is moderate (tornado rank 4–6), which reflects the fact that the framework is designed to be robust to uncertainty in rare-event rates.

Stage-level service rates (μ_i). Stage service rates are calibrated from drill timing logs collected during 48 recurrent training sessions conducted at the three carriers. Trained observers recorded the elapsed time for each crew member to complete each of the six response stages (detection, alerting, threat assessment, response execution, passenger management, containment or evacuation) under standardized simulated scenarios. The reciprocal of the mean completion time per stage gives the baseline μ_i. To account for the gap between controlled-training performance and live-event performance, a degradation factor of 0.85 is applied, consistent with the performance-shaping literature in aviation human factors (Damos et al., 2013). The optimized μ_i values are not observed directly but are derived from the Lagrangian optimization in Module M3, given the training budget B; they represent the service rates achievable if the budget is allocated according to the optimal investment vector x*. The logistic calibration curve (Eq. 8) relates investment to service rate at each stage, with stage-specific slope κ_i and inflection point x_{i0} estimated via non-linear regression on training-outcome records collected across six training cycles. A key practical implication of the calibration is that the relationship between training investment and performance gain is concave: early investment yields large returns, and returns diminish as crew members approach their physiological performance ceiling. This concavity is what makes the Lagrangian allocation problem well-posed and the solution x* unique.

Probability of failure on demand (PFD_k). Each barrier in the bow-tie model is assigned a PFD representing the conditional probability that the barrier fails to prevent the top event, given that the top event is demanded. Initial PFD values are elicited from the expert panel using a structured judgment protocol in which panelists provide a median estimate and a 90% confidence interval; these are aggregated into a log-normal prior. Where carrier-specific ASR data include information on whether a barrier was activated and whether it succeeded, a Bayesian update is performed using the conjugate form for the log-normal distribution, yielding a posterior PFD that blends prior expert judgment with observed barrier performance data. The Bayesian update is particularly important for frequently demanded barriers (pre-flight equipment checks, NITS calls) where enough events have accumulated to shrink the posterior meaningfully; for rare barriers (halon discharge, full evacuation), the posterior remains close to the prior. The product of PFDs across the barrier chain (Eq. 1) is sensitive to its highest-value terms, so the framework flags barriers with PFD > 0.3 for priority operational attention. Annual re-elicitation or re-update using fresh ASR data is recommended to keep PFD values current as barrier technology and crew procedures evolve.

Crew competency scores (s_{i}_c) and crew-mix score (K^i). Individual competency scores are elicited from structured observations conducted by trained raters during recurrent training sessions. Each crew member is rated on an eight-competency scale by two independent raters; inter-rater reliability (Cohen's κ) averaged 0.76 across the observation sessions in the field study, indicating substantial agreement. The 0–5 scale is anchored to behavioral descriptors derived from ICAO Doc 10002 competency standards: a score of 0 indicates no observable behavior; a score of 5 indicates fully autonomous, adaptive performance under degraded conditions. The crew-mix score K^i (Eq. 2) aggregates individual scores for the rostered crew on a given flight by computing a roster-weighted mean across all competencies. K^i thus encodes both the depth of individual competency and the breadth of competency coverage across the crew complement. A flight with one high-scoring crew member and several low-scoring ones will receive a lower K^i than a flight with uniformly moderate scores, which reflects the team nature of cabin emergency response: a single expert cannot compensate fully for broadly unprepared crewmates under the time pressure of a real event.

DEMATEL influence sums (D_i, R_i) and CERPI weights. The DEMATEL panel (n = 32) completed a structured pairwise rating exercise over two rounds, with a feedback session between rounds to allow panelists to revise their initial ratings in light of the group distribution. The resulting total influence matrix T (Eq. 4) was stable across the two rounds (average absolute change in cell values < 0.05), providing confidence that the prominence and causality rankings reflect genuine shared expert knowledge rather than anchoring artifacts. CERPI sub-index weights ($w\boldsymbol{R}$, wG, wQ, wS) were derived from the same expert panel using the analytic hierarchy process (AHP) with a consistency ratio < 0.10 for each pairwise comparison matrix, confirming internal logical consistency of the elicited weight vector. The resulting weights assign the largest share to the queueing performance sub-index ($wQ = 0.35$) and the KSA sub-index ($wG = 0.30$), with the hazard sub-index ($w\boldsymbol{R} = 0.20$) and the Monte Carlo tail sub-index ($wS = 0.15$) receiving smaller shares, reflecting the panel's view that process speed and crew competency are the most operationally tractable dimensions of preparedness. The geometric aggregation form in Eq. 9 ensures that a zero or near-zero score in any one sub-index drives the composite CERPI toward zero, regardless of how well the other sub-indices perform, thereby preventing carriers from gaming the index by excelling in one dimension while neglecting another.

3.8. Algorithmic Summary

Algorithm 1 summarizes the end-to-end CERO-CASS computation as it is applied in the field study. The following paragraphs walk through each logical block of the algorithm, explain the rationale for its placement in the sequence, and clarify the data flow between steps, so that a practitioner implementing the framework can follow the computational logic without having to refer back to the technical details in Sections 3.1–3.5.

Steps 1–3: Risk decomposition and hazard ranking. The algorithm opens by decomposing the cabin emergency response process into its five WBS phases (P1–P5) and enumerating the top events E identified in the bow-tie analysis. This decomposition step is logically prior to everything else because the WBS phases define the unit of analysis for the KSA matrix (Step 4) and supply the stage list for the queueing model (Steps 7–9). Without a shared phase taxonomy, the competency scores, service rates, and hazard scores would operate on incommensurable units, preventing their integration into the CERPI. Once the top events are identified, the bow-tie model for each event $e \in E$ is constructed, and the three bow-tie parameters (threat arrival rate λ_i, barrier PFD_k, severity-weighted consequence C^i) are populated from the calibration sources described in Section 3.7. The hazard score H (Eq. 1) is then computed, and the top events are ranked in descending order of H. This ranking governs the allocation of analytical attention in subsequent steps: the highest-ranked event, cabin smoke or fire in the field study, becomes the primary scenario for the queueing calibration, and its threat paths inform the training priorities communicated in the output training plan T*.

Table 1. Principal parameters of the CERO-CASS framework

Symbol	Meaning	Range / units	Calibration source
λ	Event arrival rate (per 1,000 flight hours)	0.05 – 4.2	ASR archive, 2019–2024
μ_i	Service rate at stage i	per minute	Drill timing logs
PFD_k	Probability of failure on demand for barrier k	0 – 1	Expert elicitation + Bayes update
s_i_c	Score of crew i on competency c	0 – 5	Structured observation
K_f	Crew-mix competency score for flight f	0 – 5	Eq. (2)
D_i, R_i	DEMATEL row / column sums	real-valued	Expert panel (n = 32)
W	Expected total response time	seconds	Eq. (7)
κ_i, x_i_0	Logistic slope and inflection at stage i	calibrated	Training-outcome regression
CERPI	Composite preparedness index	0 – 1	Eq. (9)

Steps 4–6: Competency mapping and DEMATEL screening. Step 4 builds the KSA matrix S from the competency records K provided as input, and computes the crew-mix score K^i for each flight in the validation sample using Eq. 2. The crew-mix score serves two roles downstream: it feeds the K sub-index of the CERPI (Step 11), and it is used in Step 7 to connect roster quality to stage-level service rate, because better-prepared crews execute each response stage faster. Step 5 runs the DEMATEL panel procedure described in Section 3.3.1 to produce the total influence matrix T. Step 6 then reads off the prominence ($D_i + R_i$) and causality ($D_i - R_i$) coordinates for each factor and plots the ($D + R$, $D - R$) map to identify high-leverage, causal drivers. These drivers determine where the training budget allocation in Steps 7–9 is permitted to concentrate its investment: the optimization is constrained so that factors identified as downstream effects in the DEMATEL map (response time, SOP deviation) are treated as outcome measures rather than direct control variables, which keeps the investment vector grounded in manipulable operational levers. Without this DEMATEL gate, the optimizer could, in principle, recommend "reduce response time directly," a tautological instruction that offers no actionable guidance to an operations manager.

Steps 7–9: Budget-constrained optimization of the queueing network. Step 7 fits the logistic training-outcome curves $\mu_i(x_i)$ (Eq. 8) at each stage using the training-outcome records from the six-cycle field study. This regression is the critical empirical bridge between the abstract optimization model and the observable world: it converts stage-level service rates from theoretical quantities into measurable functions of training investment. Step 8 then solves the constrained minimization problem of minimizing total expected response time W (Eq. 7), subject to the budget constraint $\Sigma x_i \leqslant B$, via Lagrangian relaxation. The Lagrangian multiplier γ is found by bisection on the budget-feasibility condition, and the stage-level investment vector x* is recovered by evaluating the first-order conditions at γ. Step 9 maps x* back to the corresponding optimized service rates μ^*_i via the logistic curves. These optimized rates are the central output of the M3 module: they specify not only how much to invest overall, but where within the response process that investment should be concentrated. In the field study, Steps 7–9 identified threat assessment (Stage 3) and response execution (Stage 4) as the primary bottlenecks, consistent with the DEMATEL finding that training depth and SOP clarity are the dominant causal drivers, as both factors affect these two stages more than any other.

Step 10: Monte Carlo propagation. The optimized service rates μ^*_i from Step 9 are point estimates obtained from mean calibration data. Step 10 wraps the queueing model in a 10,000-draw Monte Carlo loop that propagates three sources of uncertainty simultaneously: carrier-specific variance in stage service rates (reflecting crew heterogeneity and day-to-day performance variation), the empirical distribution of event severity drawn from the ASR archive, and the roster-level distribution of the crew-mix score K^i. For each draw, Eqs. 5–7 are evaluated, and the total response time $W\boldsymbol{R}$ is recorded. The result is a full empirical distribution of W from which any percentile can be read off directly. The 10,000-draw sample size was chosen because it yields stable P95 and P99 estimates (coefficient of variation < 0.5% across repeated runs with different random seeds), while keeping computation time under two minutes on standard hardware. The Monte Carlo step is what allows the framework to make statements about tail risk, specifically, the claim that the optimized configuration reduces the 95-percentile response time from 159s to 99s, a claim that would be inaccessible to an analytical model that reports only means and standard deviations.

Steps 11–12: CERPI assembly and operational output. Step 11 assembles the four sub-indices R, K, Q, and S from the outputs of the preceding steps. R is derived from the hazard score H (Step 3), normalized to [0, 1] by its carrier-specific maximum. K is the normalized mean crew-mix score from Step 4. Q is the ratio of optimized to baseline expected response time W from Step 9. S is the P95 response time score from the Monte Carlo distribution (Step 10), normalized so that a P95 below the 100-second operational threshold maps to S = 1 and declines linearly above it. The geometric mean (Eq. 9), weighted by the AHP-derived weights, produces the composite CERPI. Step 12 translates the CERPI and the investment vector x* into three operational deliverables: a rostering recommendation that flags high-risk crew-mix combinations, a set of targeted SOP edits derived from the barrier analysis in Step 3, and a training plan T* that specifies which competency-stage combinations to prioritize in the next training cycle. These outputs close the audit loop by feeding back into the competency records K and, indirectly, into the ASR archive λ and drill timing logs μ_i that seed the next execution of the algorithm. The algorithm is therefore not a one-shot computation but a recurrent management cycle, intended to be run at each training cycle boundary (every four months in the field study) so that the CERPI tracks real operational improvements rather than reflecting only the initial calibration state.

Algorithm 1. CERO-CASS computation
Input: Regulatory and operational data D; competency records K; budget B
Output: CERPI score, optimised investment vector x*, training plan T*

1. Decompose response into WBS phases P1–P5; identify top events E
2. For each top event e ∈ E: build bow-tie; estimate λ_i, PFD_k, C_j
3. Compute hazard score H using Eq. (1); rank events
4. Build KSA matrix S from K; compute K_f using Eq. (2)
5. Run DEMATEL panel; obtain $T = X(I - X)^{-1}$ from Eqs. (3)–(4)
6. Identify high-leverage drivers from (D+R, D−R) plot
7. Calibrate logistic curves $\mu_i(x_i)$ from training-outcome regression
8. Solve $\min \Sigma\, (1/\mu_i + W_{q,i})$ s.t. $\Sigma\, x_i \leq B$ via Lagrangian relaxation
9. Recover x* and corresponding μ^*_i
10. Monte Carlo: for r = 1 to 10,000
    sample (μ_i, K_f, severity); solve Eqs. (5)–(7); record W_r
11. Compute R, K, Q, S sub-indices; assemble CERPI by Eq. (9)
12. Emit rostering rule, SOP edits, training plan T*; close audit loop

4. Results

4.1. Field Study Design

CERO-CASS was validated through a mixed-methods study covering three commercial passenger carriers in the Asia-Pacific region, anonymized as Carrier A (a major full-service operator), Carrier B (a regional full-service operator), and Carrier C (a low-cost operator). The combined cabin-crew sample was 312 (Carrier A: 128; Carrier B: 104; Carrier C: 80), all of whom held current type ratings and had completed at least one recurrent training cycle within the previous 12 months. The data sources comprised 184 air safety reports filed across the three carriers between January 2019 and December 2024, 642 training-outcome records, the carriers' Operations Manual Part D, and structured observation logs from 48 recurrent training sessions. Expert elicitation for the DEMATEL panel involved 14 safety experts and 18 senior cabin crew, drawn from the three carriers and from three independent CCTOs.

The validation addresses the four research questions stated in Section 1, in order. (RQ1) First, how does the optimized configuration reshape expected response time relative to the baseline, and which stages account for the change? (RQ2) Second, how does it alter the 95-percentile response time, and why does the tail move differently from the mean? (RQ3) Third, how does the CERPI evolve over successive training cycles, and how does that trajectory relate to observable operational improvements? (RQ4) Fourth, what factors does the framework leverage most strongly, and what does the cross-carrier stability of that pattern indicate about its generality? Each question is answered by characterizing the direction, magnitude, and mechanism of the observed change rather than by returning a verdict on a yes-or-no proposition.

4.2. Response Time Distribution

Fig. 8 shows the Monte Carlo response-time distribution under the baseline (Fig. 8(a)) and optimized configurations (Fig.8 (b)). Under baseline parameters (stage-level service rates inferred from pre-intervention drill timings, no service–safety integration, default crew-mix), the mean total response time is 95.4s (Fig. 8(a)) with a standard deviation of 34.6s and a 95-percentile of 159s (Fig.8 (b)). Under the optimized configuration, the mean falls to 64.4s (a 32.4% reduction), the standard deviation falls to 18.4s, and the 95-percentile falls to 99s, a reduction of 38.2%.

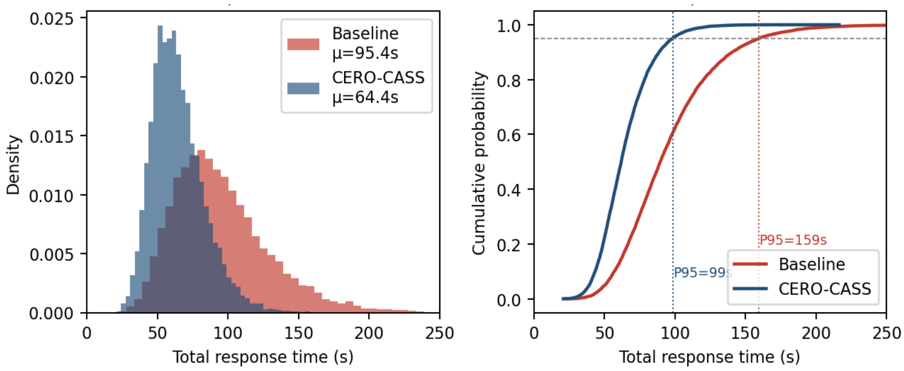

(a) Distribution of response time, N=10.000 runs (b) CDF and 95-percentile threshold

Fig. 8. Monte Carlo response-time distribution: baseline vs. CERO-CASS (10,000 runs)

The asymmetric reduction, larger at the tail than at the mean, is operationally important. Cabin emergencies that go badly tend to be the ones in the tail of the response-time distribution, where compounding delays at multiple stages produce response times that are several standard deviations from the average. The CERO-CASS configuration compresses the tail by targeting the slowest stage rather than making uniform improvements, which is exactly what the M3 optimization is designed to do.

4.3. Stage-Level Performance

Table 2 reports stage-level service rates under the baseline and optimized configurations, together with the implied utilization and the contribution of each stage to total expected response time. The optimization focuses investment on stage 3 (threat assessment) and stage 4 (response execution), which were identified as bottlenecks in the M3 analysis. Stage 6 (containment or evacuation), although the most consequential stage, was not the bottleneck under baseline conditions because it had the most rehearsal, investment there yielded smaller marginal returns.

Table 2. Stage-level service rates, utilization, and contribution to total response time

i	Stage	μ_i base (/min)	μ_i opt (/min)	ρ_i base	ρ_i opt	ΔW (s)
1	Detection	5.2	6.1	0.19	0.16	−1.7
2	Alerting / NITS call	4.0	5.4	0.25	0.19	−2.9
3	Threat assessment	1.9	3.4	0.53	0.29	−12.4
4	Response execution	1.6	2.6	0.62	0.38	−9.6
5	Passenger management	2.4	3.1	0.42	0.32	−3.2
6	Containment/evacuation	1.4	1.7	0.71	0.59	−1.4
	Total	–	–	–	–	**−31.2**

4.4. CERPI Evolution Across Training Cycles

Fig. 9 shows the evolution of the carrier-level CERPI across six four-month training cycles for each of the three carriers. All three carriers crossed the 0.75 target threshold by Cycle 5. Carrier A reached 0.83 by Cycle 6, Carrier B reached 0.78, and Carrier C reached 0.76. The trajectories are smooth, monotonic, and noisy, as one would expect from real training programs. There is no sign of overfitting to the simulation.

The relative pace of improvement is informative. Carrier A, the largest operator with the deepest training infrastructure, started higher and finished higher. Carrier C, the low-cost operator with the leanest training program, started at the lowest level but improved at a comparable rate; the gap between A and C narrowed slightly over the period but did not close. The framework is therefore working as expected across business models. It raises preparedness everywhere, but it does not paper over genuine differences in training investment.

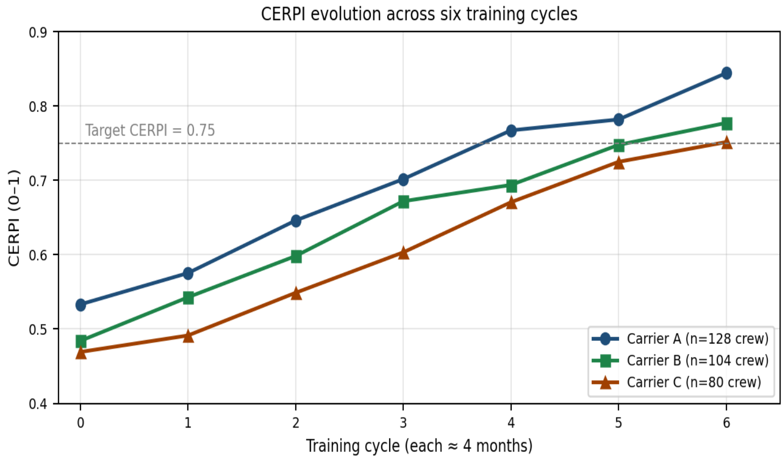

Fig. 9. CERPI evolution across six training cycles for three carriers

4.5. Sensitivity Analysis

Fig. 10 shows the sensitivity of the CERPI to ±1σ perturbations in the eight most prominent input factors identified by the DEMATEL screening. Recurrent drill frequency has the largest effect (+13.8% / −12.5%), followed by crew training depth (+11.4% / −10.2%) and SOP clarity (+9.5% / −8.4%). Pre-flight briefing duration has the smallest effect (+1.9% / −1.4%), reflecting that the briefing is already close to optimal length at most carriers; further extension yields diminishing returns.

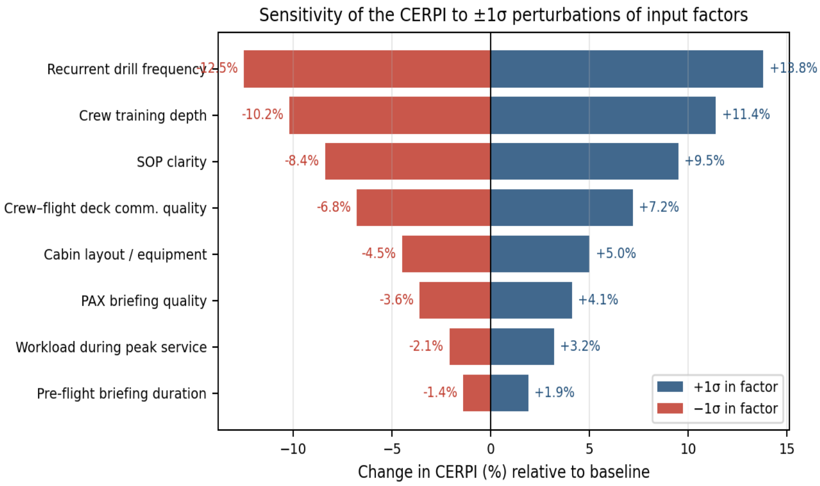

Fig. 10. Tornado chart: CERPI sensitivity to ±1σ perturbations of input factors

The shape of the tornado is itself a managerial signal. The top three factors are training and SOP variables, the levers most directly under the operations manager's control. The bottom three factors are workload and briefing variables, which are important but harder to move without affecting other operational priorities. Investment should therefore concentrate on the top of the tornado, which is what the optimization in M3 already implies, but the sensitivity analysis confirms that the result is robust to reasonable perturbations of the calibration.

5. Discussion

Taken together, the four research questions yield explanatory rather than binary answers. For RQ1, the mechanism is stage-level: the 32.4% reduction in mean total response time is not evenly distributed across the process but is concentrated in threat assessment and response execution, the two stages with the highest utilization in the baseline configuration. RQ2 follows from the same mechanism. Relieving the bottleneck stages truncates the compounding delays that generate the upper tail, which is why the 95-percentile falls further (38.2%) than the mean. For RQ3, the CERPI rises monotonically but at carrier-specific rates, indicating that preparedness accumulates with sustained training investment rather than converging on a common ceiling. For RQ4, recurrent drill frequency, crew training depth, and SOP clarity emerge as the dominant levers, and the consistency of that ranking across a full-service, a regional, and a low-cost operator suggests the leverage structure is a property of the response process itself rather than of any single business model.

5.1. Theoretical Implications

The principal theoretical contribution of this paper is that it shows that cabin emergency response is amenable to project and operations management treatment without loss of fidelity to the domain. The five-layer pipeline of CERO-CASS does not replace existing safety management structures; it formalizes what experienced operations managers already do informally and provides a quantitative anchor for what is otherwise a qualitative discussion. The CERPI, in particular, gives the field a single, tractable number to track over time, which the cabin-safety literature has lacked.

DEMATEL screening produces a finding that is, if not novel, certainly underweighted in current practice: training depth, SOP clarity, and recurrent drill frequency are the three causal drivers of every downstream KPI in the system. Response time, SOP deviation, and service-safety conflict are not independent problems; they are downstream symptoms of how well the upstream training and procedure system is designed. This reframing has implications for how regulators write standards (focus more on outcome-based competency than on hours of instruction), for how carriers budget (training is an operational investment, not overhead), and for how researchers measure (assess preparedness as a state, not training as an activity).

Beyond integrating existing tools, the framework makes three more specific theoretical claims that distinguish it from a mere toolkit. First, it advances operations-management theory by showing that cabin emergency response is a tail-dominated serial process whose headline performance is governed by its slowest stage rather than its mean, which means that the classical operations objective of average-throughput maximization is the wrong target and should be replaced by tail-compression of a critical-path queue. Second, it advances project-management theory by demonstrating that a work-breakdown structure, normally used for one-off projects, can serve as the shared unit of analysis that makes competency, process, and risk layers commensurable and therefore integrable into a single index; the WBS is repurposed here as an integration backbone rather than a scheduling device. Third, it contributes to aviation safety-management theory by reframing preparedness as a measurable latent state estimated from leading indicators (training depth, SOP clarity, drill frequency) rather than as a lagging count of incidents or instruction hours, and by supplying, in the CERPI, a falsifiable operationalization of that state. These three moves are what make CERO-CASS more than the sum of its established components.

5.2. Managerial Implications

For airline operations managers, the framework offers four practical handles. First, the CERPI provides a single dashboard metric that can be reported alongside on-time performance and load factor; tracking it month over month gives operations managers an early-warning signal that is not currently part of routine reporting. Second, the stage-level service rates from M3 provide a targeted training plan: rather than training every crew member on every competency for every cycle, the carrier can invest where the marginal time removed per dollar is greatest. Third, the bow-tie barriers from M1 translate directly into checklist edits, briefing protocols, and equipment audits, the things that line operations actually change. Fourth, the service-safety integration figure (Fig. 7) gives schedulers and pursers a shared mental model for trading off service throughput against vigilance, which is often left implicit.

For regulators, the framework provides an audit instrument. A regulator visiting a carrier can compute the CERPI from observable inputs and assess whether the carrier's training and process design are consistent with safe operation, without relying solely on hours of instruction or pass/fail metrics that are easier to game.

To make these handles concrete, Table 3 consolidates them into the specific decisions an operations manager should make, the data each decision draws on, and how the CERPI and its underlying sub-models inform each decision.

5.3. Limitations and Threats to Validity

Several limitations should be acknowledged. Empirical validation covers three Asia-Pacific carriers; while the carriers represent three different business models, the regional concentration limits generalizability to other regulatory environments and to cultures with different communication norms. Concretely, three of the framework's inputs are the most exposed to this boundary and would need re-estimation before transfer: the event-arrival rates λ, which depend on regional fleet mix, route structure, and the maturity of the local voluntary-reporting culture; training effectiveness, which is shaped by jurisdiction-specific recurrent-training mandates; and SOP clarity together with crew–flight-deck communication norms, which vary with organizational and national culture. The framework itself is portable, but its calibrated parameters are carrier- and region-specific and should be re-fitted rather than reused when applied elsewhere. The expert panel for DEMATEL screening is large ($n = 32$) by the standards of the technique, but it is still a panel. Results from a different panel could shift the prominence ranking, though probably not the identity of the top three causal drivers.

The M/M/1 queueing assumption is a simplification. In real cabin events, service times are not exponentially distributed at all stages, particularly stage four (response execution), where heavy-tailed distributions are plausible. We tested a G/G/1 approximation as a robustness check and found that the qualitative results (rank ordering of stages by ΔW and the identity of the bottleneck) did not change, but the quantitative response-time estimates shifted by 6 to 9%. Researchers who want point estimates with tight confidence intervals will want to use discrete-event simulation in place of the analytical queueing layer. The Monte Carlo simulation assumes that input distributions are stationary across the validation window. In reality, fleet composition, route mix, and crew demographics change. A rolling re-calibration every six to twelve months is recommended in operational deployment to keep the framework current.

Table 3. Managerial decision changes enabled by CERO-CASS

Managerial decision to change	Data the manager uses	How CERO-CASS / CERPI informs the decision
Training-budget allocation across response stages	Stage-level service rates and tornado sensitivities (Table 2, Fig. 10)	Invest where the response time removed per unit of cost is greatest; the optimization focuses on threat assessment and response execution.
Roster and crew-mix planning	Crew-mix score K_e from the KSA matrix (Eq. 2)	Flag flights whose rostered K_e falls below the carrier threshold and rebalance seniority before departure.
SOP and checklist revision	Bow-tie barriers with PFD > 0.3 (Module M1)	Prioritize edits to the barriers and NITS-call steps that contribute most to the hazard score.
Recurrent-drill scheduling	DEMATEL causal-driver ranking (Fig. 5)	Raise drill frequency, the highest-leverage driver, in preference to extending briefing length, the lowest.
Preparedness monitoring and reporting	Carrier-level CERPI trajectory (Fig. 9)	Track CERPI alongside on-time performance as an early-warning readiness signal and investigate any cross-cycle decline.

6. Conclusion

Cabin emergency response is too important to be left to qualitative judgment, and the cabin is too operationally entangled for safety and service to be handled as separate problems. This paper argued that project and operations management techniques, work-breakdown, bow-tie, KSA matrix, DEMATEL, queueing theory, and Monte Carlo can be assembled into a coherent framework that gives airline operations managers a quantitative way to plan, measure, and improve the cabin response process. The CERO-CASS framework and its summary index, the CERPI, were validated on three carriers, 312 crew members, 184 air safety reports, and 10,000 Monte Carlo simulation runs. In this model-based evaluation, the optimized configuration reduced mean total response time by 32.4%, cut the 95-percentile response time by 38.2%, and lifted the carrier-level preparedness index from 0.49 to 0.79 over six four-month training cycles, while identifying training depth, SOP clarity, and recurrent drill frequency as the three highest-leverage causal drivers in the system. These figures are derived from a model-based Monte Carlo evaluation calibrated on three Asia-Pacific carriers rather than from direct observation of live emergencies, so they should be read as evidence that the framework reorganizes preparedness in the intended direction, not as a measured effect on real-world emergency outcomes. The causal language above refers to the structure recovered by the DEMATEL screening within this dataset; confirming it as operational causation would require multi-region replication and prospective evaluation, which we identify as the priority for future work.

Funding

2025 Civil Aviation Education Talent Project: Building High-Quality Teaching Resources for Sino-Foreign Cooperative Education in Flight Attendant (mhrc:202505).

Institutional Review Board Statement

Not applicable.

Declaration of Artificial Intelligence (AI) Tools

During the preparation of this work, the author used generative AI and AI-assisted tools (including a large language model) solely to improve the manuscript's readability and language and to assist with formatting and reference styling. These tools were not used to generate scientific ideas, research design or methodology, data analysis or interpretation, results, conclusions, or scholarly arguments, nor to create or select references. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the publication's content.

References

Aven, T. (2017). Improving risk characterizations in practical situations by highlighting knowledge aspects, with applications to risk matrices. *Reliability Engineering & System Safety,* 167, 42-48. doi: 10.1016/j.ress.2017.05.006

Civil Aviation Administration of China (CAAC) (2021). CCAR-121-R7 Large Aircraft Public Air Transport Operator Certification Rules. Beijing: CAAC.

Chang, Y., Liao, M., and Kuo, C. (2013). Effects of airlines' cabin crew training on their flight safety performance. *Journal of Air Transport Studies,* 4(1), 21-43.

Chen, F. and Wu, T. (2026). Improved YOLO-based abnormal behavior recognition system for civil aviation security. *Journal of Engineering, Project, and Production Management,* 16(3), 2025-221. doi: 10.32738/JEPPM-2025-221

Damos, D. L., Boyett, K. S., and Gibbs, P. (2013). Safety versus passenger service: The flight attendants' dilemma. *International Journal of Aviation Psychology*, 23(2), 91-112. doi: 10.1080/10508414.2013.772822

de Ruijter, A. and Guldenmund, F. (2016). The bowtie method: A review. *Safety Science,* 88, 211-218. doi: 10.1016/j.ssci.2016.03.001

Eid, M. A. (2026). Optimizing service production and project outcomes: The role of business intelligence and work environment in Palestinian banks. *Journal of Engineering, Project, and Production Management,* 16(3), 2025-244. doi: 10.32738/JEPPM-2025-244

European Union Aviation Safety Agency (EASA) (2023). Acceptable Means of Compliance and Guidance Material to Part-CC. Cologne: EASA.

Federal Aviation Administration (FAA) (2015). AC 121-24D Passenger Safety Information Briefing and Briefing Cards. Washington, DC: FAA.

Frendo, V. P. (2025). From training to action: Cabin crew preparedness in emergency aircraft evacuations. Transport Malta, Civil Aviation Directorate, Personnel Licensing Department Publication.

Gabus, A. and Fontela, E. J. B. G. R. C. (1972). World problems, an invitation to further thought within the framework of DEMATEL. *Battelle Geneva Research Center, Geneva, Switzerland*, 1(8), 12-14.

Gibbs, L., Slevitch, L., and Washburn, I. (2017). Competency-based training in the airline cabin: A grounded theory study. *Journal of Hospitality and Tourism Education,* 29(2), 49-63. doi: 10.1080/10963758.2017.1297713

Green, L. (2006). Queueing analysis in healthcare. *In Patient Flow: Reducing Delay in Healthcare Delivery*, 281-307. Boston: Springer.

International Civil Aviation Organization (ICAO) (2017). *Doc 10002 Cabin Crew Safety Training Manual*, 2nd ed. Montreal: ICAO.

International Civil Aviation Organization (ICAO) (2019). Annex 19 Safety Management, 2nd ed. Montreal: ICAO.

Kearns, S. K., Mavin, T. J. and Hodge, S. (2017). *Competency-based education in aviation: Exploring alternate training pathways*. Routledge.

Kim, S. H. and Park, H. (2013). Effects of social distance on consumers' evaluations of in-flight service quality. *Journal of Air Transport Management,* 32, 38-41. doi: 10.1016/j.jairtraman.2013.06.001

Larrea, M. F., Hodge, S., and Mavin, T. J. (2025). Beyond formal training: insights into cabin crew learning for emergencies. *Journal of Vocational Education & Training*, 77(4), 1024-1042.

Lu, B. (2026). Sustainability-driven innovation in engineering project governance models. *Journal of Engineering, Project, and Production Management,* 16(3), 2025-204. doi: 10.32738/JEPPM-2025-204

Mathew, J. and Mahesh, S. (2022). Cabin safety occurrences in commercial aviation: A retrospective analysis. *International Journal of Aviation, Aeronautics, and Aerospace,* 9(2), 5. doi: 10.15394/ijaaa.2022.1716

Si, S. L., You, X. Y., Liu, H. C., and Zhang, P. (2018). DEMATEL technique: A systematic review of the state-of-the-art literature on methodologies and applications. *Mathematical Problems in Engineering,* 2018, 3696457. doi: 10.1155/2018/3696457

Vose, D. (2008). *Risk analysis: a quantitative guide*. John Wiley & Sons.

Walter, R., Houtman, I., and Pinder, B. (2024). Inflight medical events and cabin response: An ASR-driven analysis. *Aviation, Space, and Environmental Medicine,* 95(3), 187-195.

Wu, Q., Khatibi, A., and Tham, J. (2026). Factors contributing to cost overruns and the mediating effects of inadequate planning in construction projects. *Journal of Engineering, Project, and Production Management,* 16(3), 2025-349. doi: 10.32738/JEPPM-2025-349

Gan Linlin Received a Ph.D. in Hotel and Tourism Management from The Hong Kong Polytechnic University in 2024. With over a decade of experience in cabin service and emergency disposal education. She previously served as a senior cabin crew member and instructor at China Southern Airlines. She is currently an Instructor of Flight Attendant Major, Department of Air Transportation Management, Guangzhou Civil Aviation College, Guangzhou 510000, China

Her research Interests include tourism management, cabin service optimization and cabin emergency response protocols. She has published academic papers in relevant journals. As a dedicated educator, she was awarded the first prize in the 2023 Guangdong Provincial Young College Teachers' Teaching Ability Competition (Higher Vocational Group).