Closing the Quantification Gap in Crowd Safety Reporting

Logic Before Data:
A Side-by-Side Reference Implementation

Traditional Season Reporting versus the DIM ICE + P-index Stack

Methodology White Paper, Version 1

PneumaTheorem, Inc.

Delaware C-Corporation, incorporated 2026

Kevin Tao, Founder & CEO

Flagship publication. This document accompanies the open-access manuscript *"A Physical Quantification Framework for Crowd Compression Risk: The P-index System and Its Contact-Mechanical Foundation"* (Zenodo DOI: `pending`). The Zenodo record contains the complete physical derivation, validation methodology, and bibliography. See page 3 for the full publication series context.

PneumaTheorem, Inc.
Legalinc Corporate Services Inc., 131 Continental Dr, Suite 305, Newark, DE 19713, United States

Inquiries: kevin@pneumatheorem.org   Web: pneumatheorem.org
ORCID (Founder): 0009-0005-8136-3886

Publication Series Context

The PneumaTheorem framework is documented across a multi-layer publication series. Each document occupies a distinct methodological or applied layer; readers conducting due diligence are encouraged to consult the layer appropriate to their purpose.

Zenodo *"A Physical Quantification Framework for Crowd Compression Risk: The P-index System and Its Contact-Mechanical Foundation."*
Flagship academic manuscript. Physical derivation, four-route validation, and complete bibliography. Open-access reference for independent replication and peer review.
Zenodo DOI: pending — to be inserted upon release.

W01 *"The P-index Framework: A Physics-Based System for Crowd Compression Risk Quantification."*
Commercial white paper. Framework capabilities, application domains (insurance, regulation, venue design), and engagement architecture.

W02 *"Closing the Quantification Gap in Crowd Safety Reporting — Logic Before Data: A Side-by-Side Reference Implementation."*
Methodology white paper (*present document*). Addresses the methodological rationale under which the framework is offered: why logic-first physical grounding must precede AI-driven data augmentation, and how DIM ICE and the P-index compose as a two-layer stack rather than as competing frameworks. Includes fictionalized reference-implementation reports R01 and R02 as appendices.

E01 *"The Bipartite Nine-Cell Framework: Venue P-Report and Event DIM ICE."*
Litigation-posture white paper. Evidentiary structure of the framework under adversarial review, including the four evidentiary configurations under which a P-report may be present or absent in dispute.

Reading paths. Readers primarily interested in physical grounding should begin with the flagship Zenodo manuscript. Readers concerned with adversarial or litigation-posture questions should read E01 alongside W02.

Companion demonstration reports (appended to this document). This white paper is delivered with two reference-implementation reports as appendices:

- **Appendix A — R01 (Traditional format).** Fictional Riverside FC 2024/25 season safety report, structured under the current Green Guide / SAG reporting tradition.
- **Appendix B — R02 (DIM ICE + P-index format).** The same underlying case restructured under the two-layer stack, incorporating P_{baseline}, $P_{\max}$, P_{Ω} heatmaps, and DIM ICE-structured qualitative risk landscape.

Publication venue. All documents in the series are released through Zenodo (`zenodo.org`) under Creative Commons licensing appropriate to each document's classification. Each release receives a permanent DOI. Peer-reviewed journal or conference publication for individual documents will be pursued as appropriate; such downstream publication does not supersede the Zenodo record, which remains the canonical citation reference.

The two reports share fictionalized venue data. Neither report constitutes a safety assessment for any real venue.

Reference Materials

This white paper is grounded in three publicly available reference materials, each representing a different layer of the current crowd safety reporting ecosystem:

Reference A — Administrative Certification Layer. Notts County FC 2017 Annual Review of General Safety Certificate, issued by Nottingham City Council as Certifying Authority under the Safety of Sports Grounds Act 1975. A 9-page committee report supporting the annual General Safety Certificate for Meadow Lane Stadium (capacity 19,841). Uses the Green Guide P/S factor system (0.0–1.0 rating) as its quantification framework.

Reference B — Methodological Layer. G. Keith Still, *DIM-ICE Risk Analysis*, official framework definition. The DIM ICE meta-model organizes crowd risk into a 3×3×2 matrix (Design/Information/Management × Ingress/Circulation/Egress × Normal/Emergency), with Red/Amber/Green severity coloring. Explicitly defined by its originator as "*numerical techniques not involving computer simulations*."

Reference C — Post-Incident Review Layer. Texas Task Force on Concert Safety Report, presented to Governor Greg Abbott on April 19, 2022, following the Astroworld Festival tragedy of November 5, 2021 (10 fatalities). A state-government review whose expert panel included GK Still International, Events Safety Alliance, IAVM, TEEX, and other leading industry authorities. Produced five recurring-theme findings, all qualitative.

Throughout this document, these three references are used as anchor points against which the DIM ICE + P-index stack is compared. The three references were selected specifically because they exhibit the range of quantification maturity currently available: from administrative rating (Reference A), through structured qualitative meta-modeling (Reference B), to post-incident thematic recommendation (Reference C). No single existing framework simultaneously provides physical grounding, individual-level exposure quantification, and integration with insurance actuarial pipelines. The gap between these references and the requirements of the emerging AI-augmented risk assessment era motivates the present document.

Positioning statement. This white paper does not compete with DIM ICE. It does not replace the Green Guide P/S factor system. It does not diminish the value of expert-panel review of the type documented in the Texas Task Force report. It presents a methodological argument: *that a physically grounded quantitative substrate must be developed and validated before AI-driven data augmentation can be responsibly applied to crowd safety assessment*, and that DIM ICE together with the P-index compose a natural two-layer stack for that substrate.

Executive Summary

Crowd safety reporting currently operates across three layers, each with a mature form of expert judgment but each lacking a common quantitative substrate. The administrative certification layer (exemplified by the annual General Safety Certificate process in the United Kingdom) rates venue safety on a 0.0–1.0 P/S factor scale, but the scale measures administrative compliance, not physical risk. The methodological layer (DIM ICE, developed by Still, 2014) organizes risk into a 3×3×2 matrix with Red/Amber/Green coloring, and is deliberately positioned by its originator as a non-computational framework. The post-incident review layer (exemplified by the Texas Task Force on Concert Safety report following Astroworld, 2021) produces qualitative recommendations even when leading industry experts participate.

None of these three layers provides an individual-level, physically grounded, cumulative-exposure measurement of crowd compression risk. This quantification gap is not a failing of the three layers — each performs its function well within its intended scope. It is a structural absence in the existing methodology stack.

The gap has become critical for a reason unrelated to any of the three layers: the entry of AI-driven data augmentation into safety-critical decision workflows. Machine learning models trained on incident statistics, computer-vision-based density estimation, and generative simulation of crowd scenarios are being incorporated into venue operations, insurance underwriting, and regulatory review at increasing rates. Data availability has outpaced methodological grounding. When strong data is applied without a strong logical framework beneath it, the data dominates — and this domination is not corrected by more data.

This white paper argues that logic must precede data. The methodological substrate must be established, physically grounded, and validated *before* AI-driven augmentation is applied. The P-index framework, developed by PneumaTheorem, Inc., provides one such substrate: an individual-level, physically grounded, cumulative-exposure metric built on Hertzian contact theory, Coulomb friction, granular force-chain physics, and biomechanical lethality envelopes. When composed with DIM ICE as the structured judgment layer above it, the two form a natural two-layer stack: DIM ICE structures the qualitative risk landscape, the P-index provides the quantitative physical grounding within that landscape.

This document presents that composition through side-by-side comparison. Section 1 defines the quantification gap using the three reference materials. Section 2 develops the Logic-First argument in the context of AI-augmented decision workflows. Section 3 positions DIM ICE as the structured judgment layer. Section 4 presents a fictionalized season safety report rewritten under both the traditional format and the DIM ICE + P-index stack, allowing direct visual comparison. Section 5 provides a reader's guide to interpreting the differences. Section 6 addresses adoption paths from reference implementation to operational deployment. Section 7 positions the framework with respect to the traditions of Still, Fruin, SGSA, and the broader crowd safety literature.

The framework is available for engagement under the structures described in the companion commercial white paper W01. The present document establishes the methodological rationale under which such engagement is offered.

Contents

1. The Quantification Problem in Current Crowd Safety Assessment

1.1 Three layers, three quantification behaviors

Current crowd safety reporting occupies three functional layers. Each layer exhibits a characteristic quantification behavior. Understanding these behaviors is prerequisite to identifying where a new methodological substrate can be introduced without conflicting with existing practice.

Layer 1: Administrative Certification (Green Guide P/S factor)

The United Kingdom's Safety of Sports Grounds Act 1975 designates Certifying Authorities (local governments) to issue General Safety Certificates for sports grounds with capacity above 10,000 (or 5,000 for Premier League and Football League grounds). The Guide to Safety at Sports Grounds ("Green Guide," 5th edition, published by DCMS in 2008) specifies the assessment framework: two factors, P (Physical Condition) and S (Safety Management), each rated on $[0.0, 1.0]$; the lower factor multiplies the venue's maximum structural capacity to yield the permitted operational capacity.

In the Notts County 2017 report (Reference A), both factors were rated 1.0, yielding a permitted capacity of 19,841. The rating rationale was documented as follows: the Police, Ambulance Service, Community Protection Food/Health & Safety Service, Building Control, and Sports Grounds Safety Authority "continue to have general confidence in the safety arrangements at the Club and have not expressed objections." No physical measurement of compression, density, or contact force appears anywhere in the 9-page document. The P/S factor is a rating of *administrative compliance*, not a measurement of physical crowd risk.

Layer 2: Methodological Structuring (DIM ICE)

DIM ICE, developed by Still (2014, 2020) and now taught across the crowd safety community through GCMA, ESI, and university-level courses, organizes the qualitative structure of crowd risk into a matrix:

- **Three phases of behavior**: Ingress, Circulation, Egress.
- **Three influence dimensions**: Design, Information, Management.
- **Two operational modes**: Normal, Emergency.

Each cell of the matrix is populated with the risks associated with that phase-influence combination, and each is colored Red (needs improvement), Amber (needs monitoring), or Green (currently well managed). The output is a structured qualitative summary that supports crowd management plan development, safety officer training, and cross-scenario knowledge transfer.

Still (2020) explicitly frames DIM ICE as "numerical techniques not involving computer simulations," positioning the framework at the layer above computational implementation. The framework's strength lies in this layer: it structures human judgment reproducibly. It does not, and by design does not attempt to, provide individual-level physical measurement.

Layer 3: Post-Incident Review (Task Force reporting)

When crowd incidents occur, the standard institutional response is a task-force review producing recommendations for regulatory or industry action. The Texas Task Force on Concert Safety Report (Reference C), commissioned by Governor Abbott following the Astroworld Festival tragedy of November 2021, exemplifies this layer. The task force included GK Still International, Events Safety Alliance, the International Association of Venue Managers (IAVM), the Texas A&M Engineering Extension Service (TEEX), and other leading industry authorities.

The report's five findings — Unified Command and Control, Permitting, Training, Planning with Risk Assessment, Centralized Resources — are qualitative recommendations. No quantitative threshold, no physical measurement, no reproducible metric appears in the document. This is not a failure of expertise. The task force assembled the most experienced practitioners in the field. The absence of quantification reflects the state of the available methodological tools, not the state of the participating experts.

1.2 The gap: individual-level physical measurement

Across these three layers, one class of measurement is structurally absent: *individual-level cumulative physical exposure*. The Green Guide rates venue-level compliance. DIM ICE structures venue-level and phase-level qualitative risk. Post-incident reviews produce venue-class-level recommendations. None measures, for an individual person in a specific location at a specific time, the cumulative seconds during which that person is subject to above-threshold contact pressure.

Yet this quantity — individual cumulative exposure — is the physical basis of compressive asphyxia, the fatality mechanism in the Hillsborough (1989), Love Parade (2010), Mecca (2015), Astroworld (2021), and Itaewon (2022) incidents. The gap between what is currently measured and what physically causes fatality is the quantification gap this document addresses.

1.3 Why the gap has not been closed

The physical inputs required — Hertzian contact theory (1882), Coulomb friction (older), granular force chains (Cates et al., 1998; Majmudar & Behringer, 2005), biomechanical compression lethality (Kroll et al., 2017) — have existed in their respective literatures for decades. The obstacle has not been physical knowledge. It has been methodological assembly: no framework has integrated these physical components into an individual-level, venue-scale, insurance-compatible metric.

The framework presented in the companion flagship manuscript and W01 white paper is one such assembly. The present document addresses the methodological question that must be answered before any such assembly is adopted: *why should the physical grounding be established before, rather than after, data-driven augmentation?*

2. Logic Before Data: The Methodological Rationale

2.1 The AI-augmentation trend and its structural risk

Between approximately 2022 and 2025, three classes of AI-driven tooling have entered crowd safety practice:

1. **Computer-vision density estimation.** Real-time crowd density inference from CCTV feeds, typically producing per-frame density heatmaps at 0.5–5 Hz refresh rates.
2. **Machine-learning incident prediction.** Supervised models trained on historical incident data, producing incident probability scores over short forecast horizons.
3. **Generative crowd simulation.** Large-model-driven generation of crowd behavior scenarios for pre-event stress testing.

Each of these tools produces quantitative output. Each is increasingly incorporated into venue operations, insurance underwriting, and regulatory review workflows. The rate of adoption has outpaced the rate at which methodological grounding for the output has been established. In this configuration, a structural risk emerges that is not corrected by more data.

2.2 The structural risk stated formally

When a decision workflow lacks a strong logical framework, incoming data dominates the decision. This domination is a property of the workflow structure, not of the data quality: it holds regardless of whether the incoming data is well-founded. Three consequences follow.

Consequence 1: Weak-logic workflows launder unvalidated data into decisions.

If a venue operator's decision process is "if the CCTV density model outputs > 6 persons/m^2, close the entry gate," the decision has been outsourced to the CCTV model. The model's output enters the decision without passing through any physical validation of whether 6 persons/m^2 is the right threshold, whether the model's density estimate is accurate at the crowd configuration in question, or whether density is the correct quantity to be measuring at all (Section 1.4 of the companion W01 white paper establishes that density is neither necessary nor sufficient for compression risk).

The output of the ML model has become the logic of the decision. But the ML model contains no logic: it contains a fitted mapping from input features to output labels. Under this configuration, adding more training data to the ML model does not add logic to the decision workflow.

Consequence 2: The competitive dynamic favors early quantifiers.

In the current market, an insurer that quotes a premium supported by a specific number outperforms an insurer that quotes a premium supported by qualitative judgment, even when the number is less well-founded than the judgment. A venue operator that presents an ML-based risk score to its regulator outperforms an operator presenting a qualitative safety plan, even when the score has weaker physical basis than the plan. This is a market dynamic, not a methodological one: numbers convert to competitive positioning faster than argumentation.

The consequence is that within an estimated two to three years, the crowd safety assessment market will be dominated by quantitative outputs whose physical basis is not established. This is the "bad money drives out good" dynamic applied to methodology. Traditional qualitative frameworks — DIM ICE included — will lose institutional ground not because they are wrong but because they cannot produce numbers on demand.

Consequence 3: Point-estimate outputs cannot enter actuarial pipelines.

The first two consequences describe how weak-logic quantification captures market share. A third consequence describes why that capture is nonetheless brittle at the layer where it most needs to persist: the actuarial pipeline.

Insurance and reinsurance underwriting under Solvency II, and comparable frameworks under NAIC RBC and the IAIS Insurance Capital Standard, do not accept opaque point estimates as inputs to capital calculation. What these frameworks accept are *frequency-calibrated ordinal grades*: risk positions expressed on a discrete scale whose grades are attached to historically observed loss frequencies. Credit ratings (AAA, AA, …, D) are the mature exemplar: the letter is ordinal, but each letter carries a published default-frequency band derived from decades of default data, and it is that frequency band — not the letter itself — that enters the capital calculation.

An ML-based crowd risk model that outputs "incident probability = 0.032 for this fixture" produces a number that no actuarial pipeline can use. The pipeline requires evidence that the 0.032 was calibrated against a defined observation base, that the calibration is stable across the population of fixtures to which the pipeline will apply the number, and that the number's confidence interval is characterised. An ML model trained on historical incident reports typically provides none of these; the number is a model output, not a calibrated frequency.

The methodological substrate that Section 2.3 will introduce is designed to close this gap not by producing better point estimates, but by producing structured ordinal outputs whose boundaries are physically anchored and whose interior is deliberately left uncalibrated. Section 2.5 develops this structure.

2.3 The Logic-First response

The Logic-First position responds to this dynamic not by resisting quantification but by preceding it with adequate physical grounding. Three components are required:

- **A physically defined quantity.** The metric must derive from first principles, not from fitted correlations. Its calibration must reduce to independently established physical measurements (biomechanical thresholds, granular contact mechanics, etc.), not to historical incident statistics.

- **An individual-level output.** The metric must be defined per person, not per venue-average. Averaging over a venue permits arbitrarily many individuals to be subject to lethal exposure while the venue-average remains acceptable. The Itaewon 2022 incident (Section 1 of W01) documents this failure mode directly.
- **A specification-based reporting format.** The metric must be defined by its physical requirements, not by a specific software implementation. This is consistent with NFPA 101, ISO 16730, and Solvency II performance-based regulatory tradition, and permits validation independent of the vendor producing the number.

The P-index framework, presented in the flagship manuscript and W01, satisfies these three requirements. It provides the physical substrate onto which AI-driven data augmentation can be responsibly attached. AI is not excluded from the resulting workflow; it is positioned below the logical framework rather than above it. The internal shape of that substrate — the specific way in which it expresses safe boundaries, danger boundaries, and the interior between them — is described in Section 2.5.

2.4 Still's position and its role in the Logic-First argument

Still (2020) writes: "…the principles and applications of numerical techniques (not involving computer simulations)…[remain] the most critical element to improving crowd safety." This is not a rejection of computational quantification. It is a demarcation: Still positions DIM ICE deliberately above the computational layer, at the level of structured human judgment, because the computational layer had not, at the time of his writing, been developed to a standard of physical rigor that would support DIM ICE's evidentiary requirements.

The Logic-First position agrees with Still's demarcation and offers a specific computational substrate that meets the rigor requirement. The P-index framework's physical grounding (Hertzian contact, Coulomb friction, biomechanical lethality) is designed exactly to answer the concern Still articulates. In this sense, the P-index does not challenge DIM ICE; it addresses the reason Still restricted DIM ICE to the non-computational layer.

The composition that results — DIM ICE above, P-index below — is the natural continuation of Still's methodological program, not a departure from it.

2.5 Structured Ambiguity: the internal shape of Logic-First

Section 2.3 stated the three external requirements the substrate must meet. This subsection describes the substrate's internal shape: how it expresses risk positions in a way that satisfies actuarial pipelines without overclaiming physical precision. The shape has a name — *Structured Ambiguity* — and it consists of two components: *bipolar anchoring* at the boundaries, and *ordinal ranking* in the interior.

Bipolar anchoring: boundaries derived from independent physical measurement.

The P-index substrate is anchored at two boundaries drawn from independent physical literatures, neither of which was fitted to crowd-incident data.

The *upper boundary* is drawn from acute compression biomechanics. Kroll et al. (Kroll et al., 2017) established, using post-mortem thoracic loading data, an acute-force threshold in the region of $F_{\text{acute}} \approx 6000$ N at short duration ($\tau < 15$ s), corresponding to flail-chest structural failure. This is a short-duration, high-force anchor. The reference is co-authored by Still, and its authoritative status within the crowd safety community is not in dispute.

The *lower boundary* is drawn from restraint physiology and static compression asphyxia literature. Live-subject restraint studies (Chan et al., 1997; Michalewicz et al., 2007; Hall et al., 2012) document that healthy adults tolerate sustained back-pressure loading in the range of 25–120 kg (approximately 250–1180 N) without clinically relevant hypoxia. Static thoracic compression asphyxia in adult males occurs at sustained loads in the region of $F \approx 1100$ N for durations $\tau > 60$ s, with forensic case series (Byard et al., 2006; Gill and Landi, 2004) confirming this range in real-world crowd-crush fatalities. Sand-burial fatalities (Maron et al., 2007), in which static overburden pressure is the sole loading mechanism, provide a further independent anchor at the long-duration end. Both anchors are drawn from measurable, peer-reviewed physical data, not from fitted crowd-incident correlations.

Ordinal ranking: interior deliberately left uncalibrated.

Between the two anchored boundaries, the P-index does not claim a calibrated point-estimate probability of fatality. It expresses risk positions on a discrete ordinal scale — Safe, Grey, Danger — whose boundaries are the anchored quantities above, and whose interior is deliberately treated as an engineered extrapolation zone rather than a fitted lethality curve. The internal ordering within Grey (which venue is closer to Danger than another Grey venue) is meaningful; the internal absolute value (what specific probability of fatality corresponds to a Grey position) is not claimed.

This is the same shape that credit ratings use. A BBB and a BB rating differ ordinally; each is attached to a published default-frequency band drawn from historical data; but no rating agency claims that BBB corresponds to a specific point-estimate probability of default. The rating is a frequency-calibrated ordinal grade. The P-index substrate is the same construction applied to crowd compression risk.

Three consequences of this shape.

First, boundary judgements carry evidentiary weight, interior positions do not. A P-Report cell reading of Danger is a claim that the venue has crossed a physically anchored threshold; that claim is defensible in expert testimony and admissible under standard evidentiary rules. A P-Report cell reading of "Grey, closer to Danger than to Safe" is an ordinal comparison useful for prioritisation and insurance pricing but not offered as a point-estimate probability. The two types of output serve different institutional users and should not be conflated.

Second, the uncalibrated interior is a feature, not a defect. Insurance underwriting does not require a fatality probability for each venue; it requires zone identification for premium banding. A venue in Safe Zone qualifies for standard coverage; a venue in Grey Zone requires contract-negotiated coverage with documented mitigations; a venue in Danger Zone is uninsurable at design load until physical modification is undertaken. The market function is zone-boundary determination, not point-estimate accuracy. The P-index shape matches this function directly.

Third, AI-driven data augmentation attaches beneath the ordinal layer, not above it. Once the physical anchors and the ordinal zones are fixed, AI-driven density estimation, incident forecasting, and generative simulation can improve the temporal and spatial resolution of the readings within a zone without displacing the zone boundaries. This is the same pattern by which internal credit models, calibrated against external rating anchors, improve resolution within a rating grade without displacing the letter grade itself. The Solvency II Internal Model pathway (Section 6.5) permits precisely this composition.

Positioning of the substrate.

Structured Ambiguity is the internal shape that the Logic-First position takes when it is required to survive actuarial and evidentiary review simultaneously. It is not a retreat from quantification. It is a specification of what quantification means when the underlying physical reality — the flail-chest / long-duration-asphyxia distinction, the individual variability documented across restraint studies — does not support a single continuous lethality curve. The substrate produces the outputs the market requires (frequency-calibrated ordinal grades), attached to the anchors the physics supports (independent boundary measurements), with the interior left honest about what is not calibrated.

The composition of this substrate with the DIM ICE layer above it is developed in Section 3. The reference-implementation demonstration of the substrate in operation is presented in Section 4.

3. The DIM ICE Layer and Its Structural Role

3.1 DIM ICE as structured judgment

The DIM ICE framework, in its full form as defined by Still (2020) and taught through the GCMA, ESI, and university curricula, provides three functions that a physical metric such as the P-index cannot replace:

1. **Phase decomposition.** The Ingress–Circulation–Egress axis structures risk analysis according to the temporal sequence of the event. This decomposition organizes attention: risks specific to the Ingress phase (queue formation, entry surge, ticket validation delays) are analytically separated from Circulation risks (density fluctuation, cross-flow interference) and Egress risks (bottleneck formation at exits, dispersal into surrounding infrastructure).
2. **Influence decomposition.** The Design–Information–Management axis separates *what can be changed by whom.* Design factors (venue geometry, barrier placement, exit width)

are typically fixed at the architectural level. Information factors (signage, PA announcements, social media) can be adjusted operationally. Management factors (stewarding, decision protocols, emergency response) are subject to procedural control. Assigning risks to these categories directly maps to intervention pathways.

3. **Mode decomposition.** The Normal–Emergency axis distinguishes routine-state assessment from crisis-response assessment. A venue may perform well in Normal mode and poorly in Emergency mode, or vice versa; a single mode-blended assessment obscures this distinction.

None of these three decomposition functions is replicated by the P-index. The P-index measures a physical quantity; DIM ICE structures the analytical space within which that quantity is interpreted. They operate at different methodological layers.

3.2 The Red/Amber/Green coloring and its cognitive function

DIM ICE's Red/Amber/Green cell coloring performs a cognitive function that has been validated across multiple domains (aviation safety, medical triage, financial risk dashboards): it directs limited attention to the highest-priority items. A risk manager reviewing a fully-populated DIM ICE matrix identifies the Red cells first, addresses them, then moves to Amber, then Green. This attention allocation is more effective than either a flat qualitative list or a pure numerical ranking, both of which have been shown to overwhelm working memory when the item count exceeds approximately seven.

The P-index, by contrast, produces continuous numerical output. Continuous numbers are essential for insurance actuarial workflows and for regulatory threshold definition, but they do not directly support the attention-allocation function that operational safety managers require. The two outputs — DIM ICE's discrete colored matrix and the P-index's continuous physical metric — serve different cognitive users and different decision timescales.

3.3 The two-layer stack

The composition of DIM ICE with the P-index defines a two-layer methodology stack:

- **Upper layer (DIM ICE).** Structures the qualitative risk landscape. Answers "where in the event, in which dimension of influence, under which operational mode, do risks concentrate?" Output: a colored $3 \times 3 \times 2$ matrix, human-readable at a glance.
- **Lower layer (P-index).** Provides physical grounding for the risk levels appearing in the upper layer. Answers "for each Red or Amber cell identified above, what is the individual-level cumulative-exposure quantity that justifies the coloring?" Output: P_Ω heatmap over venue geometry, integrated P_{baseline} and $P_{\max}$ scalars for the venue as a whole, EP curve for actuarial use.

The upper layer directs attention; the lower layer supplies evidence. A Red cell in the DIM ICE matrix — for example, "Egress–Design–Emergency: pre-exit compression risk at the Jimmy Sirrel Stand end" — becomes actionable when the corresponding P_Ω heatmap identifies the specific spatial region and time window during which cumulative exposure crosses threshold. A Green cell — “Ingress–Management–Normal: turnstile allocation adequate” — is supported when P_Ω remains at zero across the corresponding region and window.

This composition is not a hierarchy in which one framework replaces the other. It is a stack in which each framework performs a function the other cannot perform.

3.4 Compatibility with existing DIM ICE practice

A crowd safety consultant currently using DIM ICE can adopt the P-index as an evidentiary layer beneath the existing DIM ICE workflow without changing DIM ICE practice. The colored matrix continues to be populated by the consultant’s structured judgment. The P-index runs in parallel, producing physical measurements that either support the consultant’s cell colorings (evidentiary confirmation) or challenge them (methodological prompt for review).

In the reverse direction, a venue operator adopting the P-index as an insurance or regulatory reporting tool retains the ability to integrate DIM ICE-based operational planning above the P-index outputs. The P-index does not require abandonment of DIM ICE; it does not require adoption of DIM ICE; it composes cleanly with either presence or absence of the upper layer.

The bipartite nine-cell framework presented in the companion litigation-posture white paper E01 formalizes this compatibility at the diagrammatic level: the venue P-Report matrix (P-index outputs) and the event DIM ICE matrix (Still framework outputs) share the Ingress–Circulation–Egress horizontal axis while differing in vertical structure, permitting side-by-side reading without merger or conflict.

4. Case Comparison: A Season Report Rewritten

4.1 Reference case and fictionalization protocol

To demonstrate the two-layer stack in operation, this section presents a fictionalized season safety report structured after Reference A (Notts County FC 2017 Annual Review of General Safety Certificate). The fictionalization protocol is as follows:

- Venue name, club name, personnel names, and match dates are replaced with fictional equivalents.
- The overall report structure (Recommendation, Reasons, Background, Key Issues, Financial, Risk Management, Equality) is preserved.
- The physical layout characteristics (four-stand configuration, capacity range, concourse geometry, road proximity) are preserved to permit realistic P-index application.
- The P/S factor ratings, the SAG process, and the Green Guide compliance framework are preserved: the traditional version reproduces the current-practice format faithfully.

The fictional venue is designated *Riverside FC at Meadowfield Stadium*, capacity 19,841, located in a fictional city. The reporting period is a fictional 2024/25 season. Two versions of the annual report are presented in parallel:

- **Version R01 (Traditional format).** Follows the Reference A template, using current best-practice content that matches or exceeds typical Certifying Authority documentation. This version does *not* represent a weakened baseline: it represents current-practice-at-its-best.
- **Version R02 (DIM ICE + P-index format).** Same underlying event data, restructured under the two-layer stack, with quantitative P_{baseline}, $P_{\max}$, P_{Ω} outputs and DIM ICE-structured qualitative risk landscape.

Both versions are attached as appendices to this document. The present section presents the comparative reading of the two.

4.2 Structural comparison

Table 1 maps the section structure of R01 (traditional Committee Report format) onto R02 (bipartite framework format). Most sections are preserved; the substantive change is that R02 folds the traditional "Recommendations" block into the titlepage and adds three annexes providing the framework overview, the full P-Report readings, and an Event DIM ICE worked example.

The intent of this structural design is that R02 remains readable as a General Safety Certificate Committee Report to a Committee member familiar only with the traditional format. The bipartite framework additions occupy annexes and per-topic readings; they do not disrupt the statutory recitation.

4.3 Quantification comparison

Table 2 lists the quantitative outputs available in R02 that have no counterpart in R01. R01 does not fail to provide these outputs through any deficiency of drafting; the traditional format simply has no place for them, because the Green Guide P/S factor system is an administrative-compliance aggregate rather than a physical-measurement framework.

The quantitative outputs above are not offered as replacements for the P/S factor. Section 1(h)(i) of R02 explicitly states that "both layers concur" in this reporting period, and that the physical readings and the administrative aggregate operate at complementary layers. The framework value is in the divergence-detection capability: should a future reporting period reveal a P-Report cell entering Danger despite an administratively compliant P/S rating, the physical reading would take precedence.

4.4 Decision-support comparison

Table 3 lists the four key decisions documented in R01 and, for each, the additional evidence R02 makes available to the Certifying Authority and the Safety Advisory Group. The additional evidence does not alter the decisions taken in this reporting period; it changes what the decisions are defensible against under external review.

R01 (Traditional)	**R02 (Bipartite framework)**	**Change**
Recommendation	Titlepage (Recommendation block)	Merged into titlepage; four recommendation items (i)–(iv) retained
Reasons for Recommendations	(merged with Background)	Absorbed into §1 Background
Background (incl. consultation outcomes)	§1 Background (a)–(h)	Preserved; §1(d) augmented with per-topic P-Report readings; §1(h) adds cross-reference to P-Report layer
Key Issues on the proposed Certificate	§1(h) (i)–(vi)	Preserved; new item (vi) introduces the bipartite framework as supplementary analytical layer
Other Options Considered	§2 Other Options	Preserved (None)
Financial Implications	§3 Financial	Preserved
Risk Management (incl. Legal)	§4 Risk Management	Preserved
Equality Impact Assessment	§5 Equality	Preserved
Background Papers	§6 Background Papers	Preserved
Published Documents	§7 Published Documents	Preserved
Annex 1 — Draft Certificate	Annex 1 — Draft Certificate	Preserved verbatim
(none)	**Annex 2 — Bipartite framework overview**	New: framework definition, composition rules
(none)	**Annex 3 — Complete nine-cell P-Report readings**	New: full reading table with threshold sources and recommended actions
(none)	**Annex 4 — Event DIM ICE worked example**	New: Coastport United fixture demonstration

Table 1: Section-map of R01 (traditional) and R02 (bipartite framework). The R02 body preserves the statutory content of R01; the framework contribution is concentrated in Annexes 2–4 and in the per-topic P-Report readings embedded within §1(d).

Quantity	R01 availability	R02 availability
P_{baseline} (venue physical floor)	Not measured; implicit in "P factor = 1.0"	Explicit per cell (Annex 3); nine cells classified Safe / Grey / Danger
$P_{\max}$ (peak sustained pressure per location)	Not measured	0.73×10^{-3} at Riverside Stand concourse, Coastport United fixture (§1(d), Annex 3 row vHC)
$P_{\Omega}(x, y, t)$ heatmap per stand	Not available (recommended for 2025/26; see Annex 3 A3.5)	Reference-implementation values from Green Guide density baselines; live deployment pending vSC monitoring upgrade
Frequency-calibrated zone assignment	0.0–1.0 P/S scale (administrative)	Three-zone Safe / Grey / Danger scale with physical thresholds per cell (Annex 3 A3.4)
Individual-level cumulative exposure	Not available	Reference-implementation per-cell reading; per-person integration pending live monitoring
Per-cell DIM ICE coloring with physical justification	Not applicable (framework absent)	Annex 4: nine-cell Event DIM ICE matrix for Coastport United fixture, each cell cross-referenced to Venue P-Report cell
EP curve for insurance actuarial use	Not available	Framework-ready output; not populated in reference implementation because live simulation data are not in scope for the fictional venue

Table 2: Quantitative outputs available in R02 that are absent in R01. The absences in R01 reflect the traditional format's design (administrative-compliance aggregate), not deficient drafting.

Decision	**R01 evidence**	**R02 additional evidence**
P/S factor rating (both at 1.0)	Consensus of consulted bodies; no objections from Police, Ambulance, SGSA, Building Control, EH	P-Report shows six of nine cells Safe, three cells in managed Grey; no cell in Danger except transiently mitigated (Annex 3 Table 2)
Capacity restriction for high-away-attendance matches (Coastport allocation reduced from 4,800 to 4,410)	Safety Officer risk assessment based on prior experience and Police intelligence	Modelled peak P_{max} at full allocation $= 0.81 \times 10^{-3}$, reduced to observed 0.73×10^{-3} under restricted allocation; documents narrow safety margin and effectiveness of mitigation (§1(d) Riverside Stand paragraph)
Riverside Stand concourse capacity assessment	Qualitative: "limited space," "pressure on concourses," "break-out area in use"	Nine-cell disclosure showing vHC and vHE Grey Zone with $P_{\text{max}} = 0.73 \times 10^{-3}$ observed and full Software row (vSI, vSC, vSE) in Grey pending monitoring upgrade (Annex 3)
Temporary Traffic Regulation Order deployment for Coastport fixture	Narrative account: "worked well," "easier egress," "used again where needed"	Quantitative: without TTRO, effective footway width 1.0 m — *Danger*; with TTRO active, effective width 4.8 m — *Safe*. Documents the TTRO as a mandatory mitigation of a Danger Zone condition, not a discretionary operational improvement (§1(d) River Road paragraph; Annex 4 A4.3)

Table 3: Key decisions in R01 and the additional evidence available under R02. The R02 evidence supports the same decisions; it changes the evidentiary basis on which those decisions are defended.

The pattern illustrated by the TTRO row is central to the framework's value proposition. Under R01, the TTRO was documented as a successful operational improvement that "worked well." Under R02, the TTRO is documented as a mandatory mitigation whose absence would leave vHE in Danger Zone. The evidentiary difference matters if the fixture had produced an incident: under R01, the venue would need to argue retrospectively that the TTRO was necessary; under R02, the argument is prospective and quantitative, with the threshold source (Green Guide 5th ed., Section 9, 1.2 m minimum footway width) recorded before the fixture. This is the concrete meaning of Structured Ambiguity (Section 2.5) applied to a certification decision: the boundary is anchored in advance, and the fixture-day management is documented as operating relative to that anchor.

4.5 What the traditional version does well

Before addressing what R02 adds, it is essential to acknowledge what R01 already does well. The Reference A tradition — reproduced faithfully in R01 — has several strengths:

1. **Multi-agency review.** The Safety Advisory Group structure integrates Police, Fire, Ambulance, SGSA, Building Control, and Environmental Health perspectives. No single-source assessment matches this breadth.
2. **Continuity over seasons.** The annual review cycle documents evolution: what changed since last year, what emerged during the season, what remains outstanding. This longitudinal continuity is valuable and is not automatically produced by any quantitative metric.
3. **Legal and procedural clarity.** The Certificate structure ties directly to statutory authority (Safety of Sports Grounds Act 1975) and to enforcement mechanisms (criminal offense for contravention). This legal integration is a mature institutional asset.

R02 does not replace any of these strengths. The two-layer stack preserves the SAG structure, the annual cycle, and the statutory framework. What it adds is a quantitative substrate beneath the qualitative outputs those structures produce.

4.6 What the DIM ICE + P-index version adds

Reading Tables 1–3 together, five specific additions R02 makes beneath the preserved R01 structure emerge:

1. **Physical justification for the P/S factor rating.** The 1.0 rating for both factors is retained in R02, but is now accompanied by nine-cell physical readings (Annex 3) that document *why* the rating is defensible: six cells consistently Safe, three cells in managed Grey, no cell in Danger except under a documented and mitigated interface condition. Under external review, the rating is no longer only a consensus statement of consulted bodies; it is a consensus statement accompanied by physical evidence.
2. **Individual-level exposure documentation in high-density stands.** The Riverside Stand concourse reading of $P_{\max} = 0.73 \times 10^{-3}$ against a Grey/Danger boundary at 10^{-3}

documents a narrow safety margin in physical terms. R01 acknowledges the concourse as a management issue; R02 quantifies the margin against a threshold anchored in the biomechanical literature (Section 2.5).

3. **P_Ω-based prioritisation of the retrofit budget.** Annex 3 recommendation (1) — upgrade of vSC from static heatmap to live $P_\Omega(x, y, t)$ observation on the Riverside Stand concourse — is derived from the nine-cell reading pattern, not from generic operational preference. Three of the four highest-impact recommendations arise from Grey Zone cells that would move toward Safe under specified interventions. The framework thus provides a structured basis for prioritising the Club's limited safety-investment budget.
4. **Insurance renewal negotiation support.** Although the fictional R02 does not populate an EP curve (live simulation data being out of scope), the framework produces the class of output that catastrophe-modelling pipelines and Solvency II Internal Model submissions can ingest: frequency-calibrated ordinal grades per cell, with threshold sources documented. Section 5.3 discusses the reading protocol for insurance underwriters approaching the framework's output.
5. **DIM ICE-structured attention allocation for the SAG's next-season agenda.** Annex 4's per-cell Event DIM ICE matrix for the Coastport United fixture demonstrates how the SAG can, in future reporting cycles, direct pre-fixture attention to cells that are Amber or Red in the fixture-specific plan and cross-reference each to the standing venue P-Report cell. This produces a structured pre-fixture agenda rather than a general narrative discussion.

None of these five additions requires abandonment of any element of the traditional format. Each is available as a supplementary evidentiary layer beneath the existing certification decision. The reference implementation in Appendices A (R01) and B (R02) demonstrates the two formats operating in parallel; the choice to adopt the bipartite framework is available to Certifying Authorities without disruption to the underlying statutory process.

5. Reading the Two Reports Side by Side

5.1 Guidance for the reviewer

Readers approaching this document may occupy any of three roles: (a) a Certifying Authority officer or Safety Advisory Group member evaluating whether the two-layer stack would augment current practice; (b) an insurance underwriter or reinsurer evaluating whether the P-index outputs integrate with catastrophe-modeling pipelines; (c) a crowd safety consultant evaluating whether the framework is compatible with existing DIM ICE-based advisory work.

Each role reads the R01/R02 comparison differently. This section provides guidance for each.

5.2 For the Certifying Authority officer

The Certifying Authority officer approaches this comparison to answer one operational question: does the two-layer stack augment or interfere with the existing General Safety Certificate process? The recommended reading protocol is three-step.

Step one: read R01 for baseline familiarity. R01 (Appendix A) is a faithful reproduction of the current-practice Committee Report format under the Safety of Sports Grounds Act 1975. If the officer's own reporting practice is well-established, R01 will read as recognisable. This is by design: the R02 comparison rests on the fact that the traditional format is not diminished, only supplemented.

Step two: examine R02 §1(d) and §1(h)(i) for the physical grounding of the P and S factors. The traditional P/S factor determination in R02 remains at 1.0/1.0, but is now accompanied by per-topic Venue P-Report readings embedded in §1(d) and cross-referenced in §1(h)(i) under the heading "Relationship to the Venue P-Report." The officer's decision on the rating is unchanged; what changes is the evidentiary layer beneath the decision. The Riverside Stand paragraph in R02 §1(d) is the clearest single illustration: the same operational conclusion (restricted away allocation, break-out area, ongoing monitoring) is now supported by an explicit modelled $P_{\max}$ reading against a documented threshold.

Step three: examine R02 Annex 3 (nine-cell reading table) and R02 Annex 4 (Coastport United worked example). Annex 3 provides the full physical evidentiary layer supporting the P/S determination; Annex 4 demonstrates how the Certifying Authority's fixture-specific consideration (in this case, the TTRO decision) can be documented against a physically anchored threshold. The officer is not required to adopt the framework to read the annexes; the annexes are self-contained and can be consulted only when a specific rating decision comes under external scrutiny.

The officer's institutional question is whether the framework's supplementary evidentiary layer is worth the marginal effort of parallel drafting. That question is answered by considering the venue-specific exposure to external review: certificates supporting venues with active insurance-renewal cycles, prior incident history, or media exposure benefit from the supplementary layer; certificates supporting long-standing well-managed venues in low-scrutiny contexts may find the traditional format sufficient. The framework does not compel adoption; it makes the option available.

5.3 For the insurance underwriter

The insurance underwriter approaches this comparison to answer a different question: does the framework's output integrate with catastrophe-modelling and Solvency II Internal Model pipelines? The reading protocol for this role differs from the officer's protocol above.

Skip R01 for actuarial purposes. R01's qualitative content (P/S factor at 1.0, consensus of consulted bodies, no objections) does not translate cleanly into an actuarial input. The traditional format was not designed for such translation, and forcing the translation risks precisely the point-estimate laundering warned against in Section 2.2 (Consequence 3).

Read R02 Annex 3 directly as the primary actuarial artefact. The nine-cell reading table (R02 Annex 3, Table 2) presents each venue cell as a frequency-calibrated ordinal grade with documented threshold sources (Annex 3 A3.4). The zone-frequency mapping stated in Section 6.5 of this white paper — Safe $\lesssim 2\times10^{-4}$ per fixture-year, Grey between 2×10^{-4} and 5×10^{-3}, Danger

$\gtrsim 5 \times 10^{-3}$ — provides the calibration input for premium banding under the Standard Formula and for undertaking-specific model calibration under the Internal Model pathway. Individual cell readings can be aggregated into a venue-level distribution suitable for exceedance-probability curve derivation, though the reference implementation in R02 does not populate the EP curve itself (live simulation data being out of scope for the fictional venue).

Compare against existing catastrophe-model output formats. Underwriters familiar with RMS, Verisk/AIR, and CoreLogic stadium-event output formats will recognise the ordinal-grade-with-frequency-band structure. The P-Report differs from these in the source of its calibration (physical thresholds versus historical loss experience) but agrees with them in the shape of its output. This shape-compatibility is the specific feature that permits ingestion into existing pipelines without bespoke translation. Section 6.5 documents the four standards frameworks (Solvency II, NFPA 101, ISO 16730, Eurocode 1-7) against which the shape has been designed.

Note the interior-versus-boundary distinction. The underwriter should not attempt to treat a Grey cell as a specific probability of incident. The framework's methodological commitment (Section 2.5) is that the ordinal position is meaningful and the zone-frequency band is meaningful, but the interior point-estimate is not claimed. Underwriting decisions accordingly proceed on zone-band evidence and on the pattern of Grey and Danger cells within a venue, not on synthesised point probabilities derived from ordinal positions. This is the same discipline that credit-portfolio managers apply to rating-band inputs; it is not a limitation of the framework but a specification of what its outputs mean.

5.4 For the crowd safety consultant

The crowd safety consultant approaches this comparison from a third position: is the framework compatible with existing DIM ICE-based advisory work, or does it compete with it? The reading protocol for this role begins from the familiar upper layer and works downward.

Examine R02 Annex 4 first (Event DIM ICE worked example for the Coastport United fixture). This annex uses the standard DIM ICE 3×3 matrix in the form documented by Still (2020) and taught through GCMA and ESI curricula. The consultant should confirm that the matrix structure, the Red/Amber/Green coloring convention, and the phase-influence decomposition are the standard ones. They are; no proprietary modification of DIM ICE is required by the framework.

Then examine the P-index outputs as evidentiary layer beneath each Red or Amber cell. R02 Annex 4 section A4.3 walks through each cell of the Coastport fixture matrix and identifies, for each Amber or Red cell, the corresponding Venue P-Report reading that anchors the coloring. The Design–Egress cell (Red without TTRO, Green with TTRO) is the clearest illustration: the consultant's independent judgement that egress conditions are compromised without the TTRO is supported by a physical measurement (effective footway width 1.0 m falling below the Green Guide 1.2 m minimum), not by narrative assertion. The consultant's judgement remains authoritative; the P-Report provides quantitative documentation for that judgement.

Consider the reverse configuration. If a P-Report reading disagrees with a consultant's cell coloring — for example, a P-Report Grey reading beneath a consultant's DIM ICE Green cell

— the framework treats this as a methodological prompt for review, not as a directive overriding the consultant's judgement. The two outputs come from different measurement regimes, and disagreement is diagnostically valuable: it identifies a cell where either the physical model or the expert assessment merits re-examination. Section 3.4 discusses this compatibility explicitly: the P-index does not require abandonment of DIM ICE, and DIM ICE does not require adoption of the P-index; the two compose but neither compels the other.

Consider the client-facing implication. A consultant offering DIM ICE-based advisory work to a venue that will subsequently seek insurance renewal or face regulatory scrutiny benefits from being able to reference the P-Report layer beneath the DIM ICE matrix. This is not a change to the consultancy deliverable; it is a change to what the deliverable can be defended against under external review. The AI vendor licensing pathway described in Section 6.3 is the specific mechanism by which independent consultants can attach P-Report evidentiary support to their existing DIM ICE outputs without building the physical model from scratch.

5.5 What the comparison does not attempt

This document does not attempt to show that R02 produces different decisions than R01. In the case documented (Riverside FC 2024/25, a season without incidents at a well-managed venue), R01 and R02 arrive at broadly similar recommendations: the Certificate should be renewed, capacity should remain at the current level, specific ongoing concerns should continue to be monitored.

What the comparison does show is that R02 produces different evidence for those decisions. The evidentiary layer available under R02 is not accessible under R01, and the difference in evidentiary layer is the concrete meaning of the "quantification gap" this document addresses.

Readers are invited to evaluate the significance of that difference for their own use case. The evaluation is not made for them by this document.

6. From Reference Implementation to Operational Adoption

6.1 The reference implementation status

The R01/R02 comparison presented in Section 4 is a *reference implementation*, not an operational deployment. It demonstrates the format, structure, and evidentiary content that the two-layer stack produces. It does not, by itself, constitute a certified or validated safety report for the fictional venue described. Operational adoption of the framework requires site-specific model configuration, parameter calibration to venue-specific geometry and attendance patterns, and integration with the existing SAG or equivalent multi-agency review structure.

6.2 Adoption pathways

Three distinct adoption pathways are available, distinguished by which counterparty leads the introduction of the framework into existing workflow. The pathways are not mutually exclusive; a venue may find itself on more than one simultaneously.

(a) Venue-led parallel pilot. A single venue produces both R01 (traditional format) and R02 (bipartite framework) outputs for one full season, running the two documents alongside each other for Safety Advisory Group review. This pathway carries the lowest institutional friction: no regulator or insurer commitment is required, and the SAG retains full discretion over which document is treated as authoritative. It is well suited to venues holding a General Safety Certificate at capacity between 10,000 and 25,000 with active SAG engagement, and to venues whose Certificate Holder anticipates future insurance-renewal or litigation exposure. Cost is bounded (a single season of parallel drafting) and the SAG output at season end is the direct comparison document for further adoption decisions.

(b) Insurer-led adoption. A property or liability insurer underwriting stadium operations conditions premium terms on receipt of R02-format outputs, either as a discount for compliance or as a required disclosure at renewal. This pathway shifts the adoption cost to the insurance-negotiation cycle rather than the certification cycle. It is well suited to insurers whose current stadium book relies on qualitative safety-plan documentation that they cannot cleanly integrate into Solvency II Internal Model submissions or comparable actuarial pipelines. The frequency-calibrated ordinal grade structure of the P-Report (Section 2.5) is the specific feature that motivates insurer-led adoption: it produces the class of output that catastrophe-modelling and capital-calculation frameworks can ingest without ad-hoc translation.

(c) Regulator-led adoption. A Certifying Authority accepts R02-format outputs as supplementary evidence for the P and S factor determinations required under the Green Guide, without altering the statutory basis of the Certificate. This is the pathway demonstrated in the R02 reference implementation itself (see R02 §1(h)(vi)). It is well suited to Certifying Authorities with active SAG-level engagement whose officers currently rely primarily on consensus-of-consulted-bodies documentation for P/S ratings and who anticipate future demand for physical evidentiary support of those ratings under external review. The regulator retains full authority over the rating; the framework provides structured evidence supporting the decision.

Details of each pathway — entry criteria, expected timeline, deliverable schedule, and pricing structure where applicable — are developed in consultation with early counterparties and are not fixed in this document. Readers whose organisation may occupy one of the three positions above are directed to the contact channel at the close of Section 7.

6.3 Engagement structures

The engagement structures under which the framework may be acquired, licensed, or co-developed are described in the companion commercial white paper W01, Section 6. The methodology white paper does not restate those structures; readers considering engagement are directed to W01 and to direct discussion via the contact channel below.

AI vendor licensing pathway. A specific engagement structure worth noting here, because it bears directly on the AI-augmentation transition discussed in Section 2, is the availability of the P-index methodology as a licensable substrate for third-party AI vendors. A vendor of computer-vision density estimation, machine-learning incident prediction, or generative simulation tooling can license the P-index framework as the methodological basis against which its own outputs are expressed. In this arrangement the vendor's continuous outputs (per-frame density, per-scenario incident probability) are translated into the P-Report's frequency-calibrated ordinal grades, with the boundary anchors (Section 2.5) providing the calibration reference. PneumaTheorem, Inc. provides the certified reference implementation of the translation; the vendor retains its own detection, forecasting, or simulation stack. This arrangement is the specific mechanism by which the AI-attaches-below-the-substrate configuration described in Section 2.3 is realised in practice, and it is available under the licensing terms in W01 Section 6.

6.4 Positioning within the AI-augmentation transition

Adoption of the two-layer stack in the near term positions adopting counterparties for the AI-augmentation transition described in Section 2. Once the physical substrate is in place, AI-driven data augmentation attaches to the substrate rather than replacing the logical framework. The Logic-First position, adopted now, becomes the platform on which responsible AI adoption occurs later.

Delaying adoption until after the AI-augmentation wave has established weak-logic quantification as market norm produces the opposite outcome: adopters must displace an entrenched but unvalidated quantitative baseline, at higher cost and against greater institutional inertia. The near-term adoption window is therefore the low-friction window.

6.5 Actuarial and standards mapping

The frequency-calibrated ordinal grade structure introduced in Section 2.5 is not a bespoke construction; it is the standard shape that established actuarial and safety-engineering standards already accept. This subsection maps the two-layer stack onto four such standards, identifying which stack element corresponds to which requirement in each.

Solvency II (EIOPA, European Union). The Solvency II framework governs insurance-undertaking capital requirements across the EU and applies to reinsurers writing UK stadium risk. It admits two calculation pathways for non-life underwriting risk: the Standard Formula (prescribed factors applied to gross premium and reserves) and the Internal Model (undertaking-specific model, subject to supervisory approval under Articles 112–127). Frequency-calibrated ordinal grades attach naturally to both pathways. Under the Standard Formula, the P-Report zone assignment (Safe / Grey / Danger) maps to premium banding within the fire and other damage to property or general liability lines. Under the Internal Model, the frequency band attached to each zone (Section 2.5) is the calibration input: Safe corresponds to $\lambda \lesssim 2 \times 10^{-4}$ per fixture-year, Grey to $2 \times 10^{-4} \lesssim \lambda \lesssim 5 \times 10^{-3}$, Danger to $\lambda \gtrsim 5 \times 10^{-3}$, drawn from crowd-crush incident base-rate analysis in W01 Section 5. These bands are ordinal boundaries, not point estimates; individual venues within a band are ranked but not assigned specific probabilities.

NFPA 101 Life Safety Code (National Fire Protection Association, US). NFPA 101 governs occupancy and egress requirements for assembly occupancies (Chapter 12–13) in the United States. Chapter 5 admits a Performance-Based Design pathway under which a design meets the Code's objectives through documented analysis rather than prescriptive specification. The P-Report cell structure maps directly onto this pathway: each Hardware-row cell (vHI, vHC, vHE) corresponds to a physical-condition performance criterion; each Management-row cell (vMI, vMC, vME) corresponds to a life-safety-management performance criterion. Zone assignment against threshold sources documented in R02 Annex 3 A3.4 provides the analytical basis for the Performance-Based Design submission.

ISO 16730 (Fire Safety Engineering). ISO 16730 specifies procedures for assessment, verification, and validation of calculation methods used in fire safety engineering. Its structure — explicit statement of the physical basis, documented input data ranges, characterised uncertainty — matches the specification-based reporting format required in Section 2.3. The P-index framework's derivation from Hertzian contact, Coulomb friction, and biomechanical lethality (W01 Section 2), together with the anchor documentation in Section 2.5, provides the physical-basis statement ISO 16730 requires. The ordinal zone structure provides the uncertainty characterisation: interior positions within a zone are not offered as point estimates, and the zone boundaries carry the evidentiary weight.

Eurocode 1, Part 1–7 (European Committee for Standardization). EN 1991-1-7 addresses accidental actions on structures, including the assessment of exposure and consequence classes. The P-Report's Ingress / Circulation / Egress phase axis maps to the exposure classification used in the consequence-class assignment; the three-zone reading provides the qualitative consequence descriptor (negligible / significant / severe) that EN 1991-1-7 Annex A requires for consequence-class determination. Adoption of the P-Report as supplementary evidence for consequence-class assignment does not alter the Eurocode framework; it provides quantitative anchoring for a determination that would otherwise rest on structural-engineer judgement.

Across all four standards, the same pattern holds: frequency-calibrated ordinal grades are the accepted output form, boundary anchoring is the accepted physical-basis form, and specification-based (rather than implementation-tied) documentation is the accepted evidentiary form. The two-layer stack was designed against these constraints from the start. The consequence for adopting counterparties is that the stack's outputs enter existing actuarial and standards pipelines without requiring bespoke translation.

7. Positioning with Respect to Existing Traditions

7.1 Relationship to Still and DIM ICE

The P-index framework does not replace, subsume, or diminish DIM ICE. Section 3 of this document has established the two-layer stack in which DIM ICE occupies the upper layer of structured judgment and the P-index occupies the lower layer of physical grounding. Still's original methodological position — that DIM ICE deliberately does not incorporate computer simulation — is preserved: DIM ICE continues to operate as a non-computational framework at its own layer. The P-index provides the computational layer beneath it, meeting the rigor standard that motivated Still's original demarcation.

Any presentation of the framework that positions it as a competitor to DIM ICE misrepresents both frameworks. The two-layer stack is offered in the tradition of methodological composition, not of methodological replacement.

7.2 Relationship to Fruin and density-based tradition

Fruin's 1971 work established density-based service-level assessment as the foundation of crowd management planning. Density remains a necessary — though, as established in W01 Section 1.4, not sufficient — indicator of crowd state. The P-index does not deprecate density measurement; it supplements density with individual-level cumulative-exposure measurement.

In practical terms, a venue's density map and its P_Ω heatmap are complementary displays. Density identifies where the crowd is; P_Ω identifies where within the crowd, cumulative pressure exposure crosses threshold. Neither replaces the other. The Fruin tradition of Level of Service assessment continues to serve its function above the P-index layer, as an operational descriptor of crowd state.

7.3 Relationship to Green Guide and SGSA tradition

The Green Guide P/S factor system, developed and maintained by SGSA, is the mature administrative certification instrument of UK sports ground safety. The P-index framework does not displace this instrument. It provides physical evidentiary support for the P and S factor ratings that SGSA-supervised Certifying Authorities issue. In the two-layer stack, the P/S factor remains the administrative output; the P-index provides the physical measurement that supports the rating decision.

For jurisdictions outside the UK, the framework's specification-based design (compatible with NFPA 101, ISO 16730, Eurocode, and Solvency II performance-based standards, as documented in W01 Section 5) permits similar composition with the local administrative certification instrument. The two-layer stack is not UK-specific.

7.4 Relationship to the Kroll biomechanical tradition

The biomechanical lethality envelope employed by the P-index derives from Kroll et al. (2017), coauthored by Still. This work provides the acute-force threshold ($F_{\text{acute}} \approx 6000$ N at $\tau < 15$ s, corresponding to flail-chest structural failure) that anchors the upper end of the P-index's contact-force scale. The framework does not extend Kroll's biomechanical claims beyond the domain in which Kroll et al. validated them: as documented in W01 Section 4.6.1, Kroll et al. explicitly delimit their model to short-duration loading and note that the model does not address long-duration crowd-crush hypoxia. The P-index's low-force, long-duration threshold (approximately 1100 N at $\tau > 60$ s) is drawn from the compression asphyxia literature independently of the acute-force anchor.

7.5 Independence of development

The framework was developed independently of the proprietary crowd simulation tools currently in commercial use, including Pathfinder (Thunderhead Engineering), MassMotion (Oasys), FDS+Evac (NIST), and Viswalk (PTV Group). No source code, dataset, or proprietary methodology from any of these tools was referenced, incorporated, or relied upon during the framework's development. The independence claim is verifiable through standard due-diligence examination of development records, as described in W01 Section 7.

7.6 Closing statement

The quantification gap in current crowd safety reporting is a structural feature of the existing methodology stack, not a failing of any existing framework. Closing the gap requires a physically grounded, individual-level, cumulative-exposure metric that does not currently exist in operational deployment. The P-index framework provides one such metric. Composed with DIM ICE, it forms a two-layer stack that preserves the existing methodological tradition while extending its evidentiary reach.

The Logic-First position — that physical grounding must precede AI-driven data augmentation — is not a resistance to AI. It is a specification of the order in which methodological components must be introduced if the resulting workflow is to remain trustworthy under adversarial or high-stakes conditions. AI attaches to the substrate; the substrate cannot be replaced by AI.

This white paper is offered as the methodological rationale under which the framework, described in commercial terms in W01 and in litigation-posture terms in E01, is made available for institutional engagement. Subsequent technical or commercial inquiries may be addressed to:

Kevin Tao, Founder & CEO
PneumaTheorem, Inc.
Delaware C-Corporation, incorporated 2026

Email: `kevin@pneumatheorem.org`
Web: `pneumatheorem.org`
ORCID: 0009-0005-8136-3886

References

Roger W. Byard, Regula Wick, Ellie Simpson, and John D. Gilbert. The pathological features and circumstances of death of lethal crush/traumatic asphyxia in adults—a 25-year study. *Forensic Science International*, 159(2–3):200–205, 2006. doi: 10.1016/j.forsciint.2005.08.003.

Theodore C. Chan, Gary M. Vilke, Tom Neuman, and Jack L. Clausen. Restraint position and positional asphyxia. *Annals of Emergency Medicine*, 30(5):578–586, 1997. doi: 10.1016/S0196-0644(97)70072-6. Hogtie/restraint-position study; restrictive pulmonary pattern but no clinically relevant hypoxia in healthy subjects. Not a weight-loading study.

James R. Gill and K. Landi. Traumatic asphyxial deaths due to an uncontrolled crowd. *American Journal of Forensic Medicine and Pathology*, 25(4):358–361, 2004. doi: 10.1097/01.paf.0000147316.62883.8b.

C. A. Hall, A. M. D. McHale, A. S. Kader, L. C. Stewart, C. S. MacCarthy, and G. H. Fick. Incidence and outcome of prone positioning following police use of force in a prospective, consecutive cohort of subjects. *Journal of Forensic and Legal Medicine*, 19(2):83–89, 2012. doi: 10.1016/j.jflm.2011.12.008. Prospective cohort (5000 use-of-force events). Contested: later reviews argue the series is underpowered to detect prone-restraint deaths.

Mark W. Kroll, G. Keith Still, Tom S. Neuman, Michael A. Graham, and Larry V. Griffin. Acute forces required for fatal compression asphyxia: A biomechanical model and historical comparisons. *Medicine, Science and the Law*, 57(2):61–70, 2017.

Bradley A. Maron, Tammy S. Haas, and Barry J. Maron. Sudden death from collapsing sand holes. *New England Journal of Medicine*, 356(25):2655–2656, 2007. doi: 10.1056/NEJMc070913. Correspondence (research letter). 52 cases 1997–2007 (US, UK, Australia, New Zealand); 31 deaths.

Bartholomew A. Michalewicz, Theodore C. Chan, Gary M. Vilke, Stephen S. Levy, Tom S. Neuman, and Fred W. Kolkhorst. Ventilatory and metabolic demands during aggressive physical restraint in healthy adults. *Journal of Forensic Sciences*, 52(1):171–175, 2007. doi: 10.1111/j.1556-4029.2006.00296.x.