International Journal of Scientific Research and Engineering Development-– Volume 9 Issue 6, Nov – Dec 2026
Available at www.ijsred.com

RESEARCH ARTICLE OPEN ACCESS

# Beyond Operational Security: Aggregation Failure, Interface Ownership and the Design of Layered Risk Assurance in Complex Organisations

Prajwal Bhattarai
*Comparative Legal Researcher, Public Policy Consultant and Regulatory Governance Advisor Kathmandu, Bagmati Province, Nepal*
Email: prajwal7bhattarai@gmail.com
Phone No.: +977-9851143549
ORCID: https://orcid.org/0009-0004-3313-7772

**********

## Abstract:

Organisations that run large public operations distribute risk management across contracted security providers, operational managers, digital teams, executives and public authorities. Enterprise risk management, high-reliability organising and resilience engineering offer partially competing perspectives on prioritisation, interpretation and adaptation in these arrangements. This article asks how organisations should structure layered risk assurance across corporate operations, public gatherings, contracted services and digital infrastructure. It conducts a structured documentary analysis of the published findings of an official public inquiry into a mass-casualty attack at a UK arena, read alongside the published scopes of the principal risk and continuity standards and intergovernmental mass-gathering guidance. The inquiry documents failures of risk assessment and coordination; the article interprets them through two proposed design mechanisms. The first is aggregation failure: scoring instruments that average or net across hazards arithmetically suppress low-frequency, high-severity hazards, a mechanism visible in a counter-terrorism assessment rated low while the national threat level was severe. The second is interface non-ownership: responsibilities were allocated to parties while the junctions between them were allocated to none. The article argues these defects are architectural rather than attitudinal, a reading the inquiry's own refusal to find deliberate compromise of safety supports and that enterprise risk management and high-reliability theory give contradictory instructions on aggregation. A Multi-Layer Risk Assurance Model specifies severity-gated registers, named interface owners and verification modes. The analysis rests on one inquiry in one jurisdiction and establishes no counterfactual.

***Keywords*** — enterprise risk management; organisational resilience; high-reliability organising; risk aggregation; mass gatherings; business continuity; interface governance; assurance design

**********

## 1. Introduction

Risk management in large organisations is layered by necessity. A venue operator contracts stewarding to a specialist supplier, retains an in-house operations function, buys information-security services and operates alongside police and transport authorities whose decisions it does not control. Each layer may be competently staffed and separately documented. The question this article pursues is what happens at the junctions and whether the standard apparatus of enterprise risk management is capable of seeing them.

The prevailing diagnosis of failure in such systems is insufficiency: too few controls, too little documentation, inadequate investment. That diagnosis has the advantage of suggesting an obvious remedy. It also has a specific empirical weakness which is that catastrophic failures in heavily regulated, heavily documented environments frequently occur in organisations whose paperwork was formally complete. Where that is so, the insufficiency diagnosis cannot

 ©IJSRED: All Rights are Reserved 

explain the outcome and prescribing more of the same apparatus that failed is unlikely to help.

This article develops an alternative account from primary documentary evidence. Its research question is: how can organisations structure layered risk assurance across corporate operations, public gatherings, contracted services and digital infrastructure? Its thesis is that the inquiry illustrates two potentially generalisable assurance mechanisms, aggregation failure and interface non-ownership, operating alongside documented deficiencies in expertise, training and risk assessment. The first is aggregation failure, in which the arithmetic of a scoring instrument suppresses the hazards the instrument exists to surface. The second is interface non-ownership, in which every substantive responsibility is allocated to some party and the junctions between parties are allocated to none.

These mechanisms can persist despite conscientious compliance with a procedure but the inquiry also identified failures of care, oversight and professional judgment. Design and human conduct are complementary explanatory levels rather than mutually exclusive diagnoses. If failure is attitudinal, the remedy is supervision and sanction. If it is architectural, supervision will produce more conscientious completion of an instrument that cannot register the hazard which is precisely the pattern the documentary record displays.

The evidence base is Volume 1 of the Manchester Arena Inquiry (2021), an official examination of security arrangements drawing on testimony and contemporaneous documents. Its detailed findings allow mechanism tracing but the retrospective inquiry neither provides a representative sample of organisations nor establishes a counterfactual for the proposed redesign.

The article's contribution is threefold. Theoretically, it identifies a direct conflict between enterprise risk management which requires commensuration of risks in order to prioritise them and high-reliability organising which counsels reluctance to simplify interpretations; aggregation failure is what occurs when simplification obscures information relevant to the second. Empirically, it traces both defects to specific, attributed findings in the inquiry record rather than asserting them. Practically, it specifies a Multi-Layer Risk Assurance Model whose distinguishing features are a prohibition on severity-diluting aggregation and the explicit assignment of interface ownership.

## 2. Four traditions and one unresolved conflict

The relevant literature divides into four traditions that are usually presented as complementary. Reading them as complementary obscures the conflict this article relies on.

Enterprise risk management treats risk as a portfolio to be managed in relation to strategy and performance and the influential COSO framework positions risk management as integrated with strategy-setting rather than operating as a separate compliance function (Committee of Sponsoring Organizations of the Treadway Commission, 2017). The framework's operational requirement is comparison: to decide where to allocate finite attention, an organisation may rank exposures using comparable information. Some scoring systems reduce qualitatively different hazards to comparable quantities but other tools preserve categorical constraints. This is not a criticism of enterprise risk management but a statement of what it must do to function.

High-reliability organisation theory describes the practices of organisations that operate hazardous technologies with unexpectedly low failure rates. Weick and Sutcliffe (2007) characterise these as preoccupation with failure, reluctance to simplify interpretations, sensitivity to operations, commitment to resilience and deference to expertise. The second of these is directly relevant. Reluctance to simplify is an injunction against exactly the reduction that commensuration performs: it holds that collapsing a complex hazard into a category or a number destroys the information needed to detect its approach.

The tension is architectural when a single scoring instrument is expected to perform both resource prioritisation and sensitive hazard detection. Portfolio decisions commonly require simplification, whereas high-reliability organising cautions against simplifying away operational signals. These are competing design pressures, not logically incompatible properties of all enterprise risk management systems. An organisation operating both, as large venue operators typically do, holds two prescriptions that may be poorly satisfied by one undifferentiated numerical score.

 ©IJSRED: All Rights are Reserved 

The literature does not resolve this, in part because the two traditions address different audiences, enterprise risk management speaking to boards and high-reliability theory to operators. The resolution proposed in Section 6 is architectural: the functions require separate registers rather than a reconciled one.

Resilience engineering contributes a third position that cuts across both. Hollnagel's (2014) distinction between safety understood as the absence of adverse outcomes and safety understood as the capacity to succeed under varying conditions implies that counting controls or incidents measures the wrong object. On this view assurance should assess adaptive capacity which is not registered by a risk matrix and is not readily audited, because what is being assessed is a disposition to respond to the unanticipated rather than conformity to a specified arrangement.

A rival position must be registered before the frame is adopted, because it questions whether design improvement is available at all. Perrow's (1999) analysis of high-risk technologies holds that where a system combines interactive complexity, meaning that components interact in unplanned and unforeseeable sequences, with tight coupling, meaning that there is little slack between components and processes cannot be halted midway, serious accidents are a normal property of the system rather than a consequence of error. On this account improving instruments and clarifying responsibilities will reduce some failures while generating new interactions, so that the residual accident rate reflects the system's structure and not its management. A large venue hosting a mass gathering adjacent to a transport interchange has both characteristics in some measure. The implication for the present argument is that the two defects identified below, if real, must be shown to be correctable features of specific instruments rather than restatements of irreducible systemic complexity. Section 6 addresses this objection directly and the distinction it turns on is whether a defect is a property of a scoring rule which can be rewritten or of the coupling itself which cannot.

Rasmussen (1997) supplies the fourth and, for present purposes, the most useful frame. He describes risk management as a socio-technical system spanning legislators, managers, work planners and system operators, observes that each level is typically studied separately by a different discipline and argues for treating risk management as a control problem spanning all societal levels for a given hazard category rather than decomposing it structurally. He advocates functional abstraction over structural decomposition and directs attention to the boundaries of acceptable performance within which behaviour is shaped. The prediction that follows for layered organisational arrangements is precise: if each level is governed by its own discipline, its own documents and its own assurance regime, the couplings between levels will be the least examined part of the system, because no discipline owns them.

Event-sector research bears out the organisational pattern without having isolated the mechanism. Nguyen and Kearns (2026) document divergent risk frames and resource negotiations among event-security stakeholders and Hutton, Robertson and Ranse (2025) examine mass-gathering safety through the differing perspectives of event managers, police and medical responders, finding that the same event is understood differently by those responsible for parts of it. Ntovoli et al. (2026) report difficulties among running-event organisers in connecting expressed awareness to integrated strategy and evaluation and Keller, Mátyás and Kökény (2026) examine risk perception and reduction in incentive travel. These studies establish that multi-stakeholder events generate incompatible risk representations. What they do not establish is whether the resulting failures are located in the representations, in the instruments that aggregate them or in the junctions between the parties holding them. That is the gap this article addresses and addressing it requires a documentary source with access to the internal instruments themselves.

### *2.1 Distributed cognition and the missing joint picture*

The literature on distributed cognition supplies a mechanism missing from a purely organizational chart. Hutchins (1995) shows how cognitive work can be accomplished through coordinated people, representations and artifacts rather than by one actor holding a complete mental model. In a layered safety arrangement, the operational picture similarly depends on how observations travel between stewarding staff, the operator's control function, police and transport

 ©IJSRED: All Rights are Reserved 

services. An interface is not simply a line dividing contracts; it is a transformation point at which one group's observations must be translated into a representation another group can act upon. A responsibility matrix can name every participant and still fail to specify who performs this transformation.

The resulting proposition is narrower than a claim that 'no one knew' about the danger. One actor may hold a threat assessment, another an operational incident log and a third a decision right. The system-level picture is incomplete when those artifacts cannot be jointly read in time to affect the decision. This formulation yields a concrete audit question: which role can reconcile the hazard list, control status, external intelligence and live operational observations into a single actionable account and what happens if a contradiction is found? It also yields a potential alternative explanation: an information interface may be adequately designed but fail because of poor execution, turnover or conflicting priorities. The documentary record must discriminate between design and execution wherever it can.

## 3. Two architectural defects and their signatures

The framework formalises the two defects and specifies their observable signatures.

*Aggregation failure* arises when a scoring instrument combines hazards using an operation that is insensitive to the distribution of severity. Three operations have this property. Averaging across hazard categories produces a composite in which a single catastrophic exposure is offset by numerous benign ones. Multiplication of likelihood by consequence produces equal scores for hazards with radically different profiles, so that a frequent minor hazard and a rare catastrophic one become indistinguishable. Netting controls against exposures credits mitigation that has not been verified as effective for the specific hazard. In each case the instrument is correctly completed and its output is wrong in a direction that is systematic rather than random: it understates precisely the hazards whose low frequency and high consequence place them at the tail.

The observable signature of aggregation failure is therefore distinctive and distinguishing it from ordinary error matters for diagnosis. Ordinary error produces scattered inconsistency. Aggregation failure produces a specific pattern: an instrument that is internally consistent, completed by competent staff and yields a low rating for a hazard that external authoritative sources rate as high. The divergence between an internal rating and an external threat assessment is thus not merely evidence that someone was wrong. It is the diagnostic marker of a mis-specified aggregation rule, because an instrument without such a rule would propagate the external assessment rather than dilute it.

*Interface non-ownership* arises when responsibility is allocated by activity rather than by junction. Contractual and organisational allocation naturally proceeds by activity: stewarding to the supplier, fabric to the operator, policing to the force. The junctions between activities, meaning the handover of information, the sequencing of actions, the resolution of conflicting instructions and the definition of mutual expectations, are not activities and therefore attract no allocation. Each party may discharge its own duties fully while the junction remains unexamined and no party's assurance regime covers it because each regime is scoped to its holder's activities.

The two defects interact in a way that is worth stating explicitly, because it explains why they are difficult to detect separately. Aggregation failure removes the hazard from the register; interface non-ownership removes the party who would have noticed its absence. An organisation with only the first defect might have its error caught by a counterparty with a different instrument. An organisation with only the second might still register the hazard internally. Together they produce a system in which the hazard is invisible on every instrument and no one is responsible for comparing them.

Four propositions follow and are testable in principle. First, divergence between internal hazard ratings and authoritative external threat assessments indicates aggregation rules rather than individual judgement and should therefore persist across personnel changes. Second, hazards whose mitigation requires joint action by two parties are systematically under-controlled relative to hazards mitigable by one. Third, assurance regimes scoped to individual duty-holders cannot detect either defect, because both are properties of the relation between holders.

 ©IJSRED: All Rights are Reserved 

Fourth, instruments that prohibit severity-diluting aggregation will surface more tail hazards and will therefore be experienced as generating unmanageable registers which is a cost rather than a refutation.

## 4. Documentary method: the inquiry record as primary source

The study is a structured documentary analysis of a single official inquiry record, supplemented by comparative analysis of the published scope of relevant standards and intergovernmental guidance. It is neither a systematic review nor a comparative case study and it reports no quantitative estimate.

*Primary material.* The principal source is Volume 1 of the Manchester Arena Inquiry report which addresses security for the arena (Manchester Arena Inquiry, 2021). This source was selected for three reasons rather than for convenience. It concerns a system with exactly the layered structure the research question addresses, comprising a venue operator, a contracted stewarding supplier and a police force with concurrent responsibilities. It was produced under statutory powers and therefore records internal instruments, scores and approval chains that are otherwise inaccessible. And its findings are individually attributed and reasoned which permits analysis of instrument design rather than of outcome alone.

*Secondary material.* The published scopes of the principal risk, continuity, crowd-management and information-security standards were examined to establish what each addresses and, critically, whether conformity with it is capable of third-party certification (International Organization for Standardization, 2018, 2019, 2022, 2024, 2026). Intergovernmental guidance on mass-gathering risk assessment and risk communication was examined for the same purpose (World Health Organization, 2023, 2025). The UK regulator's published event-safety guidance was used as a bounded example of how organiser responsibilities are articulated in one jurisdiction (Health and Safety Executive, n.d.).

*Unit of analysis.* The unit is the risk-assessment instrument, defined as a document producing a hazard rating or control determination, together with the interface, defined as a junction at which two parties' responsibilities meet. This differs from the usual unit in inquiry-based research which is the decision or the actor. Taking the instrument as the unit permits the design question to be separated from the question of individual conduct and the separation is the analytical point of the article rather than a methodological convenience.

*Procedure.* Published findings were extracted where they bore on one of four matters: the content and scoring of a risk-assessment instrument; the propagation of a finding from one document or party to another; the allocation of responsibility for a junction; and the inquiry's own reasoning about causes. Extracted findings were then classified as evidence of aggregation failure, interface non-ownership, both or neither. Findings bearing on individual conduct were recorded separately and are reported in Section 5.4, because they constitute the principal internal challenge to the architectural thesis and were not set aside.

The inquiry provides a retrospective account of one operation, not a representative sample of organisational risk systems. It supports analysis of recorded assessment and coordination failures without estimating their prevalence or establishing what would have occurred under alternative controls. The severity-gated register and proposed interface-owner arrangements are design propositions rather than tested interventions. Technical standards are used for their publicly described scope, not for unexamined clause-level requirements. Application in other organisational settings would require comparable evidence or prospective operational assessment.

## 5. Findings from the inquiry record

### *5.1 Aggregation failure in the documentary record*

The inquiry's findings on risk assessment furnish a documented instance of the mechanism specified in Section 3 and the specificity of the findings is what makes the instance analytically useful.

The primary inquiry distinguishes several deficiencies that must not be collapsed into a single aggregation mechanism. Its assessment of SMG discusses inadequate counter-terrorism attention and shortcomings in written risk assessment (Manchester Arena Inquiry, 2021, paras. 6.25–6.42). Its treatment of Showsec separately records missing attack methodologies,

 ©IJSRED: All Rights are Reserved 

weak oversight and reliance on unsupported answers about whether security recommendations had been actioned (paras. 6.55–6.58). An analytically defensible explanation therefore needs both numerical and non-numerical pathways: a scoring rule can suppress an extreme hazard; a missing scenario or incorrect documentary assertion can bypass the scoring process entirely. The proposed severity gate addresses the first, not automatically the second. Independent scenario review and evidence checks are complementary controls, not consequences of the gate.

The Chairman found that the venue operator's general risk assessment failed to adopt a rigorous approach to assessing the risk of a terrorist attack, did not identify steps to reduce that risk and did not form part of planning. Two features of the instrument itself are recorded. Within its explosions scenarios, different risk profiles received the same score. And no control measures were directed at terrorist attacks (Manchester Arena Inquiry, 2021).

The first of these is the signature of a multiplicative scoring rule. When likelihood and consequence are multiplied, scenarios with different profiles converge on identical outputs and the resulting score cannot distinguish a frequent minor event from a rare catastrophic one. The instrument was not completed incorrectly; it was completed correctly and returned outputs that could not discriminate between the cases it was being used to compare. The second feature follows from the first with some force. If the instrument cannot distinguish a catastrophic scenario from an ordinary one, the absence of controls directed at the catastrophic scenario is not an additional oversight but a predictable consequence, because controls are allocated against scores and the score contained no signal to allocate against.

The contracted supplier's counter-terrorism assessment of January 2017 provides the clearer case and the inquiry's findings on it are unusually precise. The Chairman found the assessment inadequate. Its overall terrorism rating was 12, categorised as low, at a time when the national threat level was severe. The Chairman found that averaging across hazards diluted the most likely threats. A person-borne improvised explosive device was omitted as a hazard altogether. And answers confirming that counter-terrorism security adviser recommendations had been actioned were unsupported (Manchester Arena Inquiry, 2021).

This is aggregation failure in the specific sense defined above and three aspects warrant attention. The contrast between an internal rating of low and the severe national threat level is significant context, not a direct comparison of equivalent measures: the national level assesses terrorism threat, whereas the venue assessment used its own scoring criteria. The stronger finding is the Chairman's identification of averaging that diluted serious hazards, independently of the difference between the two ratings. And the arithmetic produces an outcome that no participant intended which is why the finding is one of inadequacy rather than of dishonesty. The omission of the specific device type compounds the effect, since an averaging rule applied to a hazard set from which the relevant hazard has been excluded cannot recover it at any weighting.

The inquiry's own conclusions on method align with this reading. The Chairman concluded that severity should carry the greatest weight and that the risk of attack should not be dismissed because attacks are statistically rare (Manchester Arena Inquiry, 2021). Both propositions are corrections to an aggregation rule rather than to an attitude. Weighting severity above frequency changes the arithmetic; declining to discount by rarity prevents the tail from being netted away. Read as instrument design rather than as exhortation, the Chairman's conclusions specify what an aggregation rule must not do.

A further finding shows that the defect is not confined to scoring. A document titled Manchester Arena Security Risk Analysis identified a potential high risk from attacks near the building but was not used in day-to-day risk assessment (Manchester Arena Inquiry, 2021). Here the organisation generated a correct assessment and the assessment did not reach the instrument that governed operational decisions. The hazard was not undetected; it was detected and then failed to propagate. This is a distinct pathology and it indicates that correcting the aggregation rule alone would be insufficient without attention to whether an instrument's output reaches the decisions it is meant to inform.

 ©IJSRED: All Rights are Reserved 

It is worth being precise about why this matters for the thesis rather than treating it as an additional shortcoming. An aggregation defect and a propagation defect have different remedies and different detection signatures. A mis-specified scoring rule can be identified by inspecting the instrument alone, without reference to any other document. A propagation failure can only be identified by comparing two documents and observing that a finding present in one is absent from the other which requires someone whose remit spans both. The second defect is therefore an instance of the interface problem developed in Section 5.2 occurring inside a single organisation rather than between organisations: the junction between the analytical function that produced the risk analysis and the operational function that conducted daily assessment belonged to neither.

The event-specific layer completes the pattern. The Chairman found that the assessment for the concert in question never identified terrorism as a hazard and had descended into a box-ticking exercise and that no written terrorism risk assessment existed for the event at all (Manchester Arena Inquiry, 2021). The phrase is the Chairman's and it is precise: a box-ticking exercise is an instrument completed in form while discharging none of its function. Its outputs are available for audit and contain no information.

Two further findings concern scope rather than arithmetic and are reported here because they bear on the same instrument family. The supplier's assessment was employee-focused and the absence of any assessment directed at the safety of event-goers was described as a substantial and serious omission. And a counter-terrorism awareness document recognised patrolling as a key mitigation while omitting patrols of the adjacent City Room from its control measures which the Chairman described as a serious failure, with the inclusion of a pre-egress sheet characterised as misleading (Manchester Arena Inquiry, 2021). In both instances the instrument's declared scope excluded the population or the location at which the hazard materialised, so that complete and accurate completion within scope could not surface it.

### *5.2 Interface non-ownership in the documentary record*

The inquiry's findings on coordination establish the second defect directly. The Chairman found that the venue operator and the contracted supplier did not achieve the necessary level of communication, coordination and co-operation and that neither knew the deployment plans of the British Transport Police for each event (Manchester Arena Inquiry, 2021).

The second finding is the more analytically significant because it concerns an absence of knowledge rather than a failure of execution. Three parties held concurrent responsibilities for an interconnected space. Each presumably understood its own deployment. The composite picture, meaning who would be where and therefore which areas were covered by whom at which times, existed in no party's possession. It was not withheld by anyone. It was simply not a document anyone was responsible for producing, because its production was not an activity within any party's allocated scope.

This is interface non-ownership as specified in Section 3 and it supports the third proposition of that section. An assurance regime scoped to the venue operator would examine the operator's arrangements. One scoped to the supplier would examine the supplier's. Neither would examine the junction, because the junction belongs to the relation rather than to either party. Rasmussen's (1997) prediction that couplings between levels are the least examined part of a socio-technical system is realised here in a documented and specific form.

The pattern recurs in the finding on expert advice. The Chairman found that the operator should have sought expert help and that the absence of such help was a key weakness and separately that answers confirming counter-terrorism security adviser recommendations had been actioned were unsupported (Manchester Arena Inquiry, 2021). The interface with the external advisory function was thus incomplete in both directions: advice was not sufficiently sought and the record of whether received advice had been implemented did not correspond to what had occurred. Neither failure is attributable to the advisory function or to the organisation alone; both are properties of the junction between them.

### *5.3 What assurance is available: the certifiability asymmetry*

The standards architecture produces a structural constraint on layered assurance that is independent of any organisation's choices and it is

 ©IJSRED: All Rights are Reserved 

best seen by comparing standards by form rather than by subject.

The principal risk-management standard is published as guidelines (International Organization for Standardization, 2018). Guidance for crowd management likewise takes the form of guidelines, published in a first edition in August 2026, addressing a management framework, event and crowd risk profiling, capacity assessments, the management of arrival, circulation and departure, crowd management plans for planned operations and emergencies and stakeholder engagement, while expressly excluding crowd control and interventions directed at crowds exhibiting undesirable or dangerous behaviours (International Organization for Standardization, 2026). Business continuity management systems and information security management systems, by contrast, are published as requirements specifications (International Organization for Standardization, 2019, 2022), as is event sustainability management (International Organization for Standardization, 2024). The distinction is consequential: requirements specifications are the form against which third-party certification is conducted, whereas guidance documents are not designed to support it.

The asymmetry this produces is not a matter of subject-matter importance. Independent third-party assurance is readily available for information security and business continuity and is not available in the same form for risk management generally or crowd management specifically. An organisation seeking to demonstrate assured competence is therefore steered by the form of the available standards rather than by the severity of its hazards. It can present a certificate covering its information security management system and cannot present an equivalent covering its crowd-safety arrangements, not because the latter matters less but because the relevant document is guidance. The exclusion noted above sharpens the point: the one international standard addressed directly to crowd management is both non-certifiable in form and expressly silent on intervention against dangerous crowd behaviour which is the circumstance in which an operator most needs an authoritative benchmark.

Three consequences follow. First, assurance portfolios will tend to over-represent certifiable domains, so that the visible assurance of an organisation is a poor guide to its coverage of the most severe hazards. Second, the non-certifiable domains depend on regulatory inspection, client assurance or self-declaration which vary in rigour and are not peer-recognised. Third and most relevant to the present argument, cyber-physical dependencies fall between the regimes. A dependency in which a digital system supports a physical safety function, such as access control, communications or surveillance, is partly within the scope of a certifiable information-security standard and partly within non-certifiable crowd-management guidance. The inquiry record is consistent with the general point without demonstrating it in the digital domain: the Chairman found that proper assessment would have exposed weaknesses in closed-circuit television monitoring alongside the security perimeter, patrolling and the pre-egress check (Manchester Arena Inquiry, 2021) and surveillance monitoring is precisely such a dependency. Because the inquiry concerned physical security, this article treats the digital layer as a structural inference from standards architecture and not as a documented finding.

Intergovernmental guidance addresses some of the gap. The all-hazards risk assessment instrument for mass gatherings provides a published methodology directed at the hazard class rather than at the management system (World Health Organization, 2023) and the risk communication and community engagement toolkit addresses the information-flow dimension (World Health Organization, 2025). Neither is a certification scheme, so they improve the available method without altering the assurance asymmetry.

### *5.4 The challenge from the record: why the defects are architectural rather than attitudinal*

The strongest internal challenge to the architectural thesis is that the documented failures are better explained by commercial pressure and insufficient care which would make the remedy supervisory rather than structural. The inquiry record addresses this directly and its findings cut both ways. Reporting them fully is necessary, because the architectural reading would be unsupported if the record had found what the simpler explanation predicts.

 ©IJSRED: All Rights are Reserved 

On commercial pressure, the Chairman observed that it was inevitable that the operator and the supplier would look to save money when they could and that there are very few private and public organisations who do not do this. Staffing reviews took place and management undoubtedly hoped they would produce savings. Critically, after a warning that cutting staff would reduce security, no reductions were made. The Chairman found that the organisations failed to take steps to improve security that they should have taken and then expressly declined to go further: that does not mean that they deliberately risked the safety of event-goers and he did not find that they knowingly risked the safety of event-goers to increase their profits. He concluded that a finding that safety was deliberately compromised to save money is justified was not supported, even on the balance of probabilities. The justified criticism was that they did not take a number of necessary steps, some of which would have involved the spending of additional money (Manchester Arena Inquiry, 2021).

This is a considered rejection of the venality explanation by the body with the fullest access to the evidence and the detail that no staffing reductions were made following a security warning is directly inconsistent with a profit-over-safety mechanism in the one instance where the record permits the comparison. The explanation that survives is that necessary steps were not identified as necessary which returns the analysis to the instruments that were supposed to identify them.

On individual care, the record does contain findings of fault. The Chairman identified a lack of care and attention and a worrying lack of oversight at a senior level in relation to the supplier's assessment and attributed responsibility by name for authorship, approval and failure to notice error. Named individuals at the operator accepted that the assessments were inadequate (Manchester Arena Inquiry, 2021).

These findings must be given their weight and doing so qualifies rather than defeats the architectural thesis. Individual fault and architectural defect are not competing explanations of the same variance; they operate at different levels. The architectural claim is that an averaging rule applied to a hazard set omitting the relevant device type will return a low rating however carefully it is completed which is a property of the instrument. The finding of inadequate oversight explains why the error was not caught which is a property of the review arrangements. The two together constitute a stronger account than either alone and the first proposition in Section 3 is what distinguishes them empirically: if the defect is architectural, the divergence between internal and external ratings should persist across personnel changes, whereas if it is attitudinal it should not. The inquiry record does not test this and the proposition therefore remains open.

*5.5 Two controlled stress tests of severity-gating*

The architecture can be evaluated analytically without pretending that a new empirical evaluation has taken place. The following two scenarios are explicitly constructed stress tests. They are not reconstructions of the Manchester Arena attack and the assigned numbers are illustrative only. A conventional likelihood-by-severity grid can assign identical products to hazards whose policy consequences differ. For example, one threat scored likelihood 2 and severity 5 and another scored likelihood 5 and severity 2, both produce a composite score of 10. That equality may be reasonable for a limited prioritisation purpose but it erases the distinction between a potentially catastrophic consequence and a more frequent limited one. It does not show that every enterprise risk method makes this mistake or that arithmetic scores are intrinsically invalid. It shows what information a particular product operator loses.

Scenario A starts with both hazards correctly entered on the register. The proposed severity gate retains a separate marker for any hazard meeting a predefined consequence threshold, regardless of its composite product; a possible threshold might be multiple fatalities or loss of a critical safety function. The analyst must define the threshold before seeing the scores and must preserve the original likelihood estimate. The gate changes visibility and routing, not the probability that the event will occur. Decision-makers can still compare feasible mitigations under explicit resource constraints. In this bounded example, the gate repairs one form of information loss: equal composite numbers no longer make the two consequence classes indistinguishable.

 ©IJSRED: All Rights are Reserved 

Scenario B changes the input rather than the scoring rule. An event-specific assessment excludes a relevant attack methodology altogether. A severity gate applied to the existing register detects nothing, since the omitted event has no consequence score or entry. The remedy instead requires a scenario-generation step drawing on threat intelligence, plausible attack paths, adjacent-space responsibilities and challenge by a reviewer not responsible for the original list. Separating missing hazards from mis-ranked hazards is essential: a system that conflates them will present one improved scoring formula as a solution to two distinct failure mechanisms. The inquiry material described in Section 5.1 contains examples of both omission and problematic numerical treatment; the stress tests merely specify why different controls are needed.

A third failure mode concerns propagation. Suppose a counter-terrorism adviser identifies a relevant weakness but that finding remains in a separate report rather than being connected to the event-specific action register. Both documents may be accurate and neither a revised score nor a complete responsibility chart will fix the absent transfer. The model therefore requires an interface owner to maintain a linkage from source assessment to action, due date, verification and escalation. The decisive evidence is not the existence of two records but a demonstrable path between their findings and the operational decision. The relevant governance metric is the proportion of critical findings whose action status is reconciled across the affected teams, not the number of hazard registers completed.

## 6. Discussion: separating detection from prioritisation

The analysis resolves the conflict identified in Section 2 by denying its premise, namely that a single register can serve both prioritisation and detection. Commensuration can facilitate cross-risk prioritisation yet obscure important distinctions during hazard detection if the composite score is the only information retained. The architectural implication is separation: an organisation needs a severity-gated register that cannot be netted, maintained alongside the commensurated register that supports allocation decisions. Hazards above a defined consequence threshold enter the first register and remain visible there irrespective of likelihood, control status or composite score. The second register continues to rank and compare and draws on the first without absorbing it.

Table 1 specifies the resulting model. Its distinguishing features are the aggregation rule column which states what operation is prohibited at each layer and the interface owner column which names a party responsible for each junction rather than for an activity.

 ©IJSRED: All Rights are Reserved 

TABLE 1. MULTI-LAYER RISK ASSURANCE MODEL: LAYER, PROHIBITED AGGREGATION, INTERFACE OWNERSHIP AND AVAILABLE ASSURANCE.

| Assurance layer | Primary hazard class | Prohibited aggregation operation | Junction requiring an owner | Interface owner proposed | Available third-party assurance |
|---|---|---|---|---|---|
| Corporate governance | Strategic and reputational exposure | Netting of severity-gated hazards into composite board reporting | Board register to operational register | Board risk committee chair | None specific; internal audit only |
| Operational management | Crowd dynamics, fabric, egress | Averaging across hazard categories; likelihood-consequence multiplication for tail hazards | Operator to contracted supplier | Named operations director, with a jointly signed composite plan | Guidance only; not certifiable in the same form as requirements standards |
| Contracted services | Stewarding, screening, patrolling | Crediting unverified control effectiveness against exposure | Supplier to public authority | Joint deployment-picture holder, designated by contract | Client assurance and inspection; no peer-recognised certification |
| Digital and cyber-physical | Access control, communications, surveillance | Treating a digital control as mitigating a physical hazard without verifying the dependency | Information security to physical safety | Named cyber-physical dependency owner | Certifiable for information security management; not for the physical dependency |
| Continuity and recovery | Disruption, loss of capability | Averaging recovery capability across dissimilar scenarios | Continuity plan to live incident command | Incident commander, pre-identified by role | Certifiable for business continuity management systems |
| Public authority interface | Concurrent statutory responsibilities | Assuming coverage from another party's presence without a shared record | Organisation to police, transport and health authorities | Multi-agency liaison owner on the organisation's side | None; regulatory oversight and published guidance |

Three features of the model respond to the documentary findings rather than to general principle. The joint deployment-picture holder in the third row addresses the finding that no party knew the police deployment plans: the remedy is not better communication which was already an expectation but a named owner of a specific composite document whose absence is detectable. The cyber-physical dependency owner in the fourth row addresses the certifiability asymmetry, since the dependency falls between two regimes and will otherwise be assured by neither. The prohibition on crediting unverified control effectiveness in the third row addresses the finding that confirmations of implemented advice were unsupported.

The handoff between the two registers requires specification if separation is not to become duplication. Three rules suffice. Entry to the severity-gated register is determined by consequence alone, assessed against a defined threshold such as potential for multiple fatalities or loss of a critical safety function and entry is not conditional on likelihood, control status or cost of mitigation. Movement within that register is permitted only on evidence that the hazard class no longer applies to the operation, not on evidence that it has been controlled, since controlled hazards remain hazards and the register's function is visibility rather than accounting. Transfer to the commensurated register is additive rather than substitutive: a hazard appears in both, ranked in the second for allocation purposes while remaining unranked in the first. The arrangement is deliberately redundant and redundancy is the point, since the failure mode it addresses is the disappearance of a hazard from the only place it was recorded.

Five objections deserve response and the first two are the most serious.

*Severity-gating will flood the register and paralyse decision-making.* If every high-consequence hazard remains permanently visible regardless of likelihood, the register becomes long, undifferentiated and practically unusable and attention will be drawn from frequent hazards that cause most actual harm. This objection is well founded and constitutes the principal cost of the proposal. Two things limit it. The severity-gated register is a detection instrument and is not asked to allocate resources; allocation remains with the commensurated register which continues to rank. And the number of hazards above a well-set consequence threshold in a given operation is bounded and small, since catastrophic hazard classes for a venue are few even though their specific manifestations are many. The objection nonetheless identifies a real trade-off and the fourth proposition in Section 3 concedes it: the model will be experienced as generating unmanageable registers and whether that cost is worth bearing is an empirical question this article does not settle.

*More controls do not produce better outcomes.* Additional controls can generate congestion, delay, privacy intrusion and operational complexity and each of these can itself create hazard. The response is that the model adds no controls. It changes an aggregation rule and assigns ownership of existing junctions. Both are specification changes rather than additions to the control inventory and the model is in this respect deliberately parsimonious.

A serious rival explanation is hindsight bias. Inquiry findings are reached knowing the outcome and an instrument that failed to surface the hazard that materialised will always appear defective in retrospect.

 ©IJSRED: All Rights are Reserved 

This objection has force and constrains what is claimed. The claim advanced here is not that the instruments would have prevented the attack which the record cannot establish and which is not asserted anywhere in this article. It is that the instruments contained identifiable design features, namely averaging across hazards, identical scores for different risk profiles, omission of a hazard type and scope confined to employees which are defective independently of the outcome. An averaging rule that dilutes severe hazards is mis-specified whether or not a severe hazard eventuates and this is why the analysis was conducted at the level of instrument design rather than of causation.

The complexity objection is substantial. Perrow (1999) argues that tightly coupled, interactively complex systems can generate serious failures despite competent management; eliminating one register defect does not remove structural uncertainty. The proposed model has the narrower aim of fixing inspectable process defects rather than predicting that all hazards can be controlled. A scoring rule that hides severe consequences, a missing scenario or an undocumented handoff can be identified before an incident and examined independently of its outcome. These are in principle correctable, whereas the entire interaction network is not. The model's value should therefore be judged through prospective testing of detection, decision latency and ownership of interfaces, not by a claim that it would necessarily have prevented this attack.

*A single inquiry cannot support a general account.* It cannot and no general empirical claim is made. The inquiry demonstrates that the mechanisms are real and documents their operation in one system with precision unobtainable by other means. Whether they are common is a separate question requiring a different design, such as coded comparison of risk-assessment instruments across venues which Section 7 specifies.

### *6.1 Costs, prioritisation conflicts and criteria for abandoning the model*

A severity gate creates costs of its own. First, a register of extreme but remote consequences can grow until managers again treat it as background noise. The control therefore needs an explicit review cycle, a defensible definition of critical consequence and a method for closing or consolidating entries without erasing the underlying hazard. Second, a severe-hazard flag can place long-horizon threats ahead of frequent preventable injuries; the gate must not be read as an automatic priority ordering. A board may reasonably allocate resources to both, using separate decisions about salience and cost-effectiveness. Third, separate detection and prioritisation registers can produce conflicting owners or incompatible action dates. An agreed conflict-resolution authority and escalation time are necessary to prevent the assurance architecture from reproducing the interface defect it was designed to cure.

The proposal would be weakened if it produced large volumes of unactionable escalations, delayed safety-critical decisions or reduced attention to frequent hazards without materially improving recognition of catastrophic exposures. A proper evaluation should compare the old and proposed approaches on a defined set of plausible scenarios, record missed hazards and avoidable escalations and test whether staff can identify the responsible interface owner quickly. Such an evaluation is not reported here. The contribution is a control-design hypothesis disciplined by inquiry findings, not evidence that a new register would have prevented the historical attack.

### *6.2 External qualitative evidence and contextual transfer*

Recent research provides an independent point of comparison without converting the Manchester record into a representative sample. Hutton, Robertson and Ranse (2025) examined mass-gathering safety through the perspectives of event managers, police and medical personnel, identifying communication and ownership as practical concerns. Nguyen and Kearns (2026), drawing on observations, interviews and event documents from two large US events, describe how actors operating with different risk frames negotiate resource requests and responses. These studies support treating representational and coordination gaps as analytically important but they do not independently validate the present severity-gate design. The inquiry supplies a documented instance of risk misclassification and incomplete transfer of information; the other research makes the cross-boundary mechanism plausible in additional settings, while leaving causal effectiveness open.

The engineering implication is to separate a non-compensatory *screening rule* from an *allocation decision*. A severity threshold should trigger additional examination and explicit ownership, not automatic investment in every extreme hypothetical. The screening stage preserves information about consequential but infrequent hazards; a subsequent documented decision weighs evidence about exposure, uncertainty, feasible controls and competing claims on

 ©IJSRED: All Rights are Reserved 

finite resources. An independently reviewable exception process is preferable to an absolute veto rule, because otherwise the number of flagged hazards could overwhelm operational attention. A minimum record should explain the hazard scenario, evidence of plausibility, the responsible interface owner, the decision and when the decision must be reconsidered.

## 7. Implications for operators, contracts and standard-setting

For operators, the first implication is to audit the aggregation rule rather than the register. A review that asks whether hazards have been identified and controls assigned will pass an instrument whose arithmetic suppresses the hazards it should surface. The diagnostic question is narrower and answerable: for each hazard class, does the instrument's scoring operation permit a severe consequence to be offset, averaged or netted against anything else and does any internal rating diverge from an authoritative external assessment of the same hazard. Divergence is a defect in the instrument until shown otherwise and the accountable actor is the holder of the risk-management framework rather than the completer of the form.

The second implication is that junction ownership must be written into contracts. Where responsibilities are divided between an operator and suppliers, the contract allocates activities and will not allocate junctions unless it does so expressly. Contracts should therefore name a holder for each composite document the junction requires, specify what the document contains and state the period for which it must be current. Authority rests on the contract itself which is a stronger instrument for this purpose than policy because it binds the counterparty. The safeguard is that the holder is identified by role rather than by individual, so that the obligation survives personnel change. The indicator is whether the composite document exists and is current which is verifiable without assessing anyone's judgement.

The third implication is for regulators and standard-setting bodies and follows from the certifiability asymmetry. Where the most severe hazard classes in a sector are governed by guidance and lesser classes by certifiable requirements organisations will accumulate assurance where it is obtainable rather than where it matters. Addressing this requires either requirements-form standards for the severe classes, with the attendant costs and the risk of converting judgement into conformity or explicit regulatory recognition that certificates in adjacent domains carry no implication for crowd safety. The choice between these is a policy question beyond the scope of this analysis but the asymmetry should not be left implicit, since its present effect is to make the visible assurance of an organisation systematically unrepresentative of its coverage.

A fourth implication concerns how organisations assure contracted providers in domains where certification is unavailable. Where no peer-recognised scheme covers the supplier's core activity, the buyer's assurance options are inspection, client audit and self-declaration and each has a known weakness: inspection samples, client audit is conducted by a party with a commercial interest in continuity and self-declaration is unverified. The documentary record is instructive here, since confirmations that advisory recommendations had been actioned were found to be unsupported. The practical response is to assure the evidence rather than the assertion: for each control the supplier claims, require the record that would exist if the control were operating, specify in the contract what that record is and verify a sample against it. This shifts verification from the supplier's account of its practice to artefacts generated by the practice itself and it is available to buyers who cannot obtain certification because none exists for the activity.

For research, the propositions in Section 3 are testable with a design this article has not conducted. A documentary study could obtain risk-assessment instruments from a defined population of venues, code their aggregation operations against pre-specified definitions and compare internal ratings for a common hazard class against the contemporaneous external threat assessment. Inter-coder reliability would be reported and the analysis would be descriptive rather than causal, since venues with better instruments plausibly differ in other respects. A separate strand could examine whether hazards requiring joint mitigation are under-controlled relative to those mitigable by one party, using control inventories rather than outcomes, since outcome-based comparison across venues is confounded by exposure. Neither design would establish that improved instruments reduce harm which would require variation in instruments across comparable settings and a follow-up period long enough to observe rare events, a requirement that may place the question beyond reach.

## CONCLUSION

The inquiry illustrates two architectural failure mechanisms in layered risk assurance; their prevalence elsewhere remains untested. Aggregation failure occurs when a scoring instrument averages, multiplies

 ©IJSRED: All Rights are Reserved 

or nets across hazards in a manner that suppresses low-frequency, high-severity exposures, so that a correctly completed instrument returns a systematically wrong answer. Interface non-ownership occurs when responsibility is allocated by activity, leaving the junctions between parties unallocated and therefore unexamined by every party's assurance regime. The inquiry records a low-rated terrorism assessment despite a severe national threat level, documented omissions and coordination failures. The model distinguishes scoring error from omitted scenarios and unverified handoffs rather than reducing them to one cause.

That these defects are architectural rather than attitudinal is supported by the inquiry's own refusal to find that safety was deliberately compromised for profit and by its finding that proposed staffing reductions were abandoned once their security consequence was raised. What remained was that necessary steps were not identified as necessary which directs attention back to the instruments that existed to identify them.

The answer to the research question is therefore structural. Organisations should separate the register that detects from the register that prioritises, because commensuration is necessary for the second and destructive of the first; should prohibit severity-diluting aggregation in the detection register; and should allocate ownership of junctions by name and by contract, not merely expect cooperation across them. The principal cost of the first measure is register inflation which is real and unresolved. The principal limitation of the analysis is that one inquiry establishes that these mechanisms operate without establishing how often they do.

## REFERENCES

Committee of Sponsoring Organizations of the Treadway Commission. (2017). *Enterprise risk management: Integrating with strategy and performance*. https://www.coso.org/enterprise-risk-management

Health and Safety Executive. (n.d.). *Managing an event*. https://www.hse.gov.uk/event-safety/managing-an-event.htm

Hollnagel, E. (2014). *Safety-I and safety-II: The past and future of safety management*. Ashgate.

Hutchins, E. (1995). Cognition in the wild. MIT Press.

Hutton, A., Robertson, M., & Ranse, J. (2025). Exploring safety at mass gathering events through the lens of three different stakeholders. *Frontiers in Public Health, 12*, Article 1451891. https://doi.org/10.3389/fpubh.2024.1451891

International Organization for Standardization. (2018). *ISO 31000:2018: Risk management, guidelines*. https://www.iso.org/standard/65694.html

International Organization for Standardization. (2019). *ISO 22301:2019: Security and resilience, business continuity management systems, requirements*. https://www.iso.org/standard/75106.html

International Organization for Standardization. (2022). *ISO/IEC 27001:2022: Information security, cybersecurity and privacy protection, information security management systems, requirements*. https://www.iso.org/standard/27001

International Organization for Standardization. (2024). *ISO 20121:2024: Event sustainability management systems, requirements with guidance for use*. https://www.iso.org/standard/86389.html

International Organization for Standardization. (2026). *ISO 22353:2026: Security and resilience, guidelines for crowd management*. https://www.iso.org/standard/50259.html

Keller, K., Mátyás, S., & Kökény, L. (2026). When rewards turn risky: The hidden dangers and risk reduction strategies in case of incentive travels. *International Journal of Event and Festival Management*. https://doi.org/10.1108/IJEFM-08-2025-0120

Manchester Arena Inquiry. (2021). *Report of the public inquiry into the attack on Manchester Arena on 22nd May 2017: Volume 1, security for the arena*. https://manchesterarenainquiry.org.uk/report-volume-one/

Nguyen, T., & Kearns, E. (2026). Event security: Interdisciplinary risk communication. *International Journal of Event and Festival Management*. https://doi.org/10.1108/IJEFM-08-2025-0124

Ntovoli, A., Karagiorgos, T., Polyakova, O., Anoyrkati, E., Papadopoulou, S., & Alexandris, K. (2026). Exploring sustainable management in running events: Insights from Greek organizers. *International Journal of Event and Festival Management*. https://doi.org/10.1108/IJEFM-09-2025-0164

Perrow, C. (1999). *Normal accidents: Living with high-risk technologies* (updated ed.). Princeton University Press.

Rasmussen, J. (1997). Risk management in a dynamic society: A modelling problem. *Safety Science, 27*(2–3), 183–213. https://doi.org/10.1016/S0925-7535(97)00052-0

 ©IJSRED: All Rights are Reserved 

Weick, K. E., & Sutcliffe, K. M. (2007). *Managing the unexpected: Resilient performance in an age of uncertainty* (2nd ed.). Jossey-Bass.

World Health Organization. (2023). *The generic all-hazards risk assessment tool for mass gathering events*. https://www.who.int/publications/i/item/WHO-2023-Generic-Mass-gatherings-All-Hazards-RAtool-2023-1

World Health Organization. (2025). *Risk communication and community engagement readiness and response toolkit: Mass gatherings*. https://www.who.int/publications/i/item/9789240109148

 ©IJSRED: All Rights are Reserved 