



Risk-Based Gas Detection and Respiratory Protection Strategies for High-Hazard Oil and Gas Facilities: Closing the Barrier-Management Lifecycle within the Integrated Risk Management Methodology

Altair Jilgildiyev^{a,*}

^a ex-Executive Sales Director for Caspian Region, The MSA Safety Co, Atyrau, Kazakhstan

ABSTRACT :

Gas detection systems and respiratory protective equipment (RPE) at oil and gas facilities are usually specified as separate technical packages: detection as an instrumentation problem and respiratory protection as a procurement problem. This paper develops an exposure-control framework in which the two are specified together, as the detection layer and the last mitigative layer of a single barrier architecture, driven by one shared risk-tiering step. The gap it addresses is structural rather than technical. Four elements of barrier management — identification and representation of barriers on threat and consequence legs, the independence and sufficiency of credited protection layers, the integrity of the instrumented layer, and indicators of barrier condition and degradation — developed historically in separate documents, are owned by separate teams and are revised on separate cycles. A facility can therefore hold a valid barrier diagram, a completed protection-layer assessment, a verified instrumented function and a maintained indicator set, and still be unable to demonstrate that these describe the same barrier: the lifecycle is open at the joins. Four design blocks are proposed: a hazard and scenario inventory that separates chronic low-level exposure from acute release; a semi-quantitative risk-tiering step using qualitative likelihood and consequence descriptors; a detection layer in which sensing technology, placement, independence and alarm logic are matched to the target gas and to the ambient conditions of the installation; and a respiratory protection layer that specifies device class, fit testing, donning and doffing discipline and escape provision according to the assigned tier. A fifth, connecting element carries barrier-condition indicators from both layers back into periodic re-verification. The framework is a module-level extension of the author's previously published Integrated Risk Management Methodology. Two analytical results follow: risk concentrates not in the substances of highest nominal toxicity but in the combinations of substance, task and location where warning properties are poor and monitoring resolution is coarse; and alarm credibility, not catalogue sensitivity, is the limiting design parameter, since an incredible alarm silently invalidates the protection-layer credit taken for the detection barrier.

Keywords: gas detection; respiratory protection; barrier management; safety barriers; occupational risk assessment; oil and gas facilities; semiconductor gas sensors; safety climate

1. Introduction

1.1 The exposure problem at upstream and midstream installations

Upstream and midstream oil and gas installations concentrate several inhalation hazards within a single work envelope. Hydrogen sulfide from sour crude and sour gas processing, hydrocarbon vapours from tank gauging and sampling, amine and glycol mists from treatment units, ammonia from refrigeration circuits, welding fume, drilling-fluid aerosols and oxygen-deficient atmospheres in tanks, sumps and pits coexist with high-energy fire and pressure risk.

Hydrogen sulfide remains the reference hazard of the group. Its acute toxicity, the olfactory paralysis that removes the warning property at precisely the concentrations at which warning matters most, and the steepness of its dose–response relationship place it in a category of its own among routinely encountered industrial gases (Guidotti, 2010). The health objective, however, cannot be confined to acute poisoning. A joint statement of the American Thoracic Society and the European Respiratory Society concluded that workplace inhalational hazards remain common worldwide and contribute materially to the burden of airway disease, interstitial fibrosis, hypersensitivity pneumonitis, other granulomatous lung disease and selected respiratory infections, while emphasising that these contributions are in principle preventable (Blanc et al., 2019). Irritant gases and hazardous air pollutants matter additionally for workers who already have airway disease, since sensitised individuals may respond to remarkably low concentrations and irritants can lower the bronchoconstrictive threshold (Leikauf, 2002).

Exposure at these installations is rarely steady. It is dominated by short excursions tied to specific tasks — line breaking, filter changes, pigging, confined space entry, turnaround maintenance. Task-resolved measurement in the offshore petroleum sector illustrates the magnitude of the effect: in an assessment of benzene exposure on the Norwegian continental shelf over 2002–2018, exposure levels separated sharply between task groups on the same installations, with tank work and maintenance on process equipment standing well clear of routine operation (Ridderseth et al., 2023). A control

strategy built around long-term average concentrations therefore systematically underestimates the hazard experienced by the individual worker, and a monitoring architecture built only around fixed points cannot recover information that was never collected at the right place and time.

Two implications run through the whole of this paper. Control must address chronic low-level exposure as well as acute release; and the susceptibility profile of the workforce is part of the risk equation rather than an administrative detail.

1.2 Detection technology has outgrown detection specification

Semiconductor devices based on metal oxides, conducting polymers, carbon nanotubes and two-dimensional materials now cover a wide range of target analytes, each combination of material and transducer carrying a characteristic profile of sensitivity, selectivity, stability and operating temperature (Nikolić et al., 2020). Nanostructured titanium dioxide illustrates how morphology can be tuned during fabrication to modify sensing response, with demonstrated application to hydrogen, acetone, ammonia and ethanol (Xu et al., 2021). Ammonia sensing alone has generated an extensive body of work across material classes and operating principles (Kwak et al., 2019).

These developments enlarge the design space available to the facility engineer. What they do not supply is the decision rule: which device belongs at which location, at what alarm threshold, with what demonstrated credibility, and — decisively — what must happen to the people present when it alarms. The last question is not an instrumentation question at all.

1.3 Respiratory protection presents the mirror-image problem

Reviews of filtering devices against particulate and bioaerosol challenges have identified persistent gaps between laboratory filtration performance and protection achieved in use (Rengasamy et al., 2004). A systematic review of personal protective equipment for highly infectious exposures found that equipment type and the method of donning and doffing both influence contamination, yet the certainty of the underlying evidence was limited (Verbeek et al., 2020). Evidence from industrial rather than clinical practice points the same way: simulated workplace protection factor measurements on half-facepiece respiratory protective devices show a wide distribution of achieved protection across wearers and exercises, with a substantial fraction of measurements falling well below the nominal performance of the device class (Duling et al., 2007). Face-seal leakage, wear time and user behaviour, not filter efficiency, dominate the outcome.

The consequence for specification is immediate: a device class assigned without fit testing, seal checking and a written donning and doffing sequence is a nominal control and cannot legitimately be credited as a barrier.

1.4 The organisational condition

Technical layers operate as designed only under organisational conditions that are themselves measurable. Longitudinal benchmarking across North Sea oil and gas installations showed that safety climate dimensions were predictive of self-reported accidents and that substantial, measurable improvement with convergence between installations was attainable within a year (Mearns et al., 2001). Climate is therefore treated here not as context but as a monitored parameter of the control system.

1.5 The structural gap: four elements of barrier management, four documents

The technical shortfalls above are real but they are not the binding constraint. The binding constraint is structural, and it is a consequence of how barrier management developed.

Four elements are required to manage a safety barrier over its life. The barrier must be identified and represented in relation to the threats it interrupts and the consequences it limits; the guidance for this developed around the bow-tie representation and is codified in the process-safety concept literature (Center for Chemical Process Safety [CCPS] & Energy Institute, 2018; de Ruijter & Guldenmund, 2016). The barrier must be shown to be independent and sufficient in combination with the other layers credited for the same scenario; this is the domain of protection-layer analysis (CCPS, 2001). Where the barrier is an instrumented function, its integrity must be specified and verified against the functional-safety standards (International Electrotechnical Commission [IEC], 2010, 2016). And its condition must be monitored, so that degradation is visible before demand rather than after failure; this element entered mainstream practice through the barrier-management regimes developed in the petroleum sector, in which barrier status and performance are followed as standing indicators (Petroleum Safety Authority Norway, 2013).

Each element is sound. The difficulty is that they were written in different documents, at different times, by different constituencies, and in ordinary facility practice they are owned by different functions and revised on different cycles: the barrier diagram by process safety, the protection-layer assessment by the risk group, the integrity verification by instrumentation and controls, and the indicator set by HSE. Reviews of barrier classification, assessment and management in the process industries note precisely this heterogeneity — that barrier concepts and their assessment methods have proliferated in parallel rather than convergently (Yuan et al., 2022) — and the documented limitations of the bow-tie as an analytical rather than communicative device are part of the same picture (de Ruijter & Guldenmund, 2016).

The practical consequence is a set of joins that no single document owns. A facility may hold a current barrier diagram, a completed protection-layer assessment, a verified instrumented function and a maintained indicator set, and still be unable to demonstrate:

- that the detection barrier credited as an independent layer is the same physical and functional barrier whose availability the indicator set tracks;
- that the alarm on which the credit depends is credible in the ambient conditions of the installation, rather than one whose false-alarm behaviour has already eroded the response it assumes;
- that the respiratory protection credited as a mitigative barrier on the consequence leg has a current fit-test record for the personnel actually assigned to the task;

- that a measured degradation in any of the above re-enters the assessment that credited the barrier, rather than terminating in a maintenance record. These are the joins, and they are where the exposure actually occurs. Closing them is the subject of this paper.

1.6 One line of work: relationship to the author's Integrated Risk Management Methodology

This paper is not a free-standing proposal. It continues a single line of work.

The Integrated Risk Management Methodology (IRMM) for hydrogen sulfide at high-risk oil and gas facilities (Jilgildiyev, 2026) was developed to close exactly this loop at facility-architecture level. It treats hazard identification, risk assessment, barrier management, monitoring, emergency preparedness and organisational learning as one closed cycle in which the output of each element is the input to the next; it classifies protective barriers into three functional categories — preventive, detection and control, and mitigative; it requires that the layers credited for a critical scenario be as independent of one another as practicable; and it requires that for every critical barrier a function, a performance criterion, an owner, a verification method and frequency, criteria for an unacceptable condition and the action on degradation be defined. Its extended eight-level formulation and the associated application tools develop the same architecture into a working procedure.

The IRMM therefore already carries all four elements identified in Section 1.5, and carries them as one system rather than four. What it defines at the level of architecture, it does not specify at the level of the individual barrier. The present paper supplies that specification layer for the two barrier classes through which inhalation exposure is controlled — the detection layer and the respiratory protection layer — and extends their application from hydrogen sulfide alone to the mixed inhalation-hazard inventory of a working installation, a direction of development identified as a priority in the parent methodology (Jilgildiyev, 2026).

Consistent with that continuity, this paper introduces no analytical instrument that is not already declared in the parent methodology. Barrier representation, layer independence, functional-safety integrity and barrier-condition indicators are used here in the same sense and in the same combination in which the IRMM defines them; what is added is the specification content for two barrier classes and the rules that keep the four elements coupled at those two barriers. Table 1 states the correspondence element by element.

Table 1 – The four elements of barrier management: source tradition, usual owner, and what this paper adds at the detection and respiratory protection barriers.

Element	Source tradition	Usual owner in facility practice	Position in the IRMM (Jilgildiyev, 2026)	Specification added here
Barrier identification and representation	Bow-tie concept literature (CCPS & Energy Institute, 2018; de Ruijter & Guldenmund, 2016)	Process safety	Level 1, hazard identification; barrier register	Exposure archetypes that fix the required temporal and spatial resolution of monitoring (Block 1, Section 2.3)
Independence and sufficiency of credited layers	Protection-layer analysis (CCPS, 2001)	Risk / process engineering	Principle of layered protection; independence of protection layers	Three-layer monitoring architecture with an explicit non-substitution rule; conditions under which a detection layer may be credited (Block 3, Section 2.5)
Integrity of the instrumented layer	Functional-safety standards (IEC, 2010, 2016)	Instrumentation and control	Detection and control barriers; automatic action	The interface: what the framework delivers to, and requires back from, the instrumented function (Section 2.5.4)
Barrier condition and degradation indicators	Petroleum-sector barrier management (Petroleum Safety Authority Norway, 2013); process safety indicators (Hopkins, 2009)	HSE	Principle of barrier performance monitoring; barrier statuses; leading and lagging indicators	A single indicator set spanning both barrier classes, with defined feedback into re-verification (Sections 2.7 and 3.7)

1.7 Aim, contribution, scope and article type

The aim is to specify a framework in which gas detection and respiratory protection are designed together, driven by a common risk-tiering step, with the four elements of barrier management kept coupled at both barriers, and to state the technical and organisational conditions under which such a framework can be expected to work.

The contribution claimed is methodological and consists of four items:

- Closure of the barrier-management lifecycle at the two exposure-control barriers. Identification, layer independence, integrity and condition monitoring are specified in one document, for the same barrier, with defined handovers between them and a defined return path from indicator to re-verification. The novelty is not in any one element — each is established and none is claimed as new — but in their closure into a single cycle at the barrier level, which is where the joins currently fail.

- A single tiering step that drives two barrier classes. Detection requirements and respiratory protection requirements are derived from the same tier assignment for the same scenario, so that a demanding detection package cannot be coupled to a permissive protection regime, or the reverse.
- Exposure archetypes that determine monitoring resolution. Assigning each release mechanism to an archetype converts the question “how many detectors” into the question “at what temporal and spatial resolution”, which is the question that determines whether the measurement can represent the exposure at all.
- Alarm credibility as a design parameter of equal standing with detection limit — and as a condition of layer credit. Cross-sensitivity, drift, humidity dependence, calibration interval, bump-test frequency and a documented response per alarm class are specified alongside the detection limit, and a detection layer whose alarm credibility is not demonstrated is not credited as an independent layer.

Scope. The framework is formulated for the class of high-hazard oil and gas installations — upstream, midstream and associated processing and storage. It is facility-class and not project-specific: it neither derives from nor describes the circumstances of any individual operator, asset or engagement, and no proprietary data are used. The worked scenarios are generic by construction.

Article type. This is a conceptual and methodological paper. Its object is a specification framework and its output is a set of design rules, not an empirical finding. It is not a literature review: published evidence is the input, and the output is a specification and a verification procedure. It is not an empirical research article either, and does not present itself as one. Section 2.1 states the evidential basis and the reasons why no facility-attributable dataset is reported.

2. Materials and Methods

2.1 Article type, evidential basis and confidentiality

The framework is derived from published evidence, appraised internally, and informed by the author’s professional practice in atmospheric monitoring and respiratory protection at oil and gas installations. No facility data were collected, extracted or analysed, no exposure or incident dataset is reported, and no quantitative performance claim is made anywhere in this paper.

The absence of a facility-attributable dataset is a property of the work, not an omission, and the reasons are stated here because they are otherwise open to misreading.

Genre. The deliverable is a specification method intended to be applied across a class of installations. A dataset drawn from any one installation would bind the framework to that installation’s hazard inventory, ventilation regime and organisational maturity — the failure mode the framework exists to prevent. Framework development, methodological integration and synthesis are established article types in the safety sciences, and are evaluated on the coherence, traceability and applicability of the method rather than on the size of an accompanying sample.

Heterogeneity of the evidence base. The knowledge domains a coupled detection and protection strategy must draw on — clinical epidemiology, sensing materials, protective-equipment performance, barrier theory and organisational research — are heterogeneous by construction. No numerical values were extracted or pooled across them, because a quantitative synthesis over material of this kind would not be defensible: pooling a sensor response curve with a protection-factor distribution and an accident-rate series produces a number with no referent.

Confidentiality and infrastructure sensitivity. Barrier-level work of the kind described here is carried out on operating assets. The identity of operators and installations, the scenario sets developed for them, the assessed condition of their barriers and the performance figures arising from that work are subject to commercial confidentiality, and the last two additionally constitute vulnerability information relating to critical infrastructure. They are therefore not reported, and neither the framework nor its worked scenarios are traceable to any asset. This is the ordinary reporting convention of the field rather than a limitation specific to this paper; its methodological price is that the reader cannot verify the framework against the practice that informed it, and it is for exactly that reason that Section 4 states an application and verification procedure by which any operator can test the framework on its own installation and generate its own evidence.

Where a design rule rests on engineering reasoning rather than on a cited result, this is stated at the point where the rule is introduced, so that the two kinds of warrant are not confused.

2.2 Source selection

Sources were selected to span the domains the framework must draw on: the toxicology of the reference hazard and the health consequences of occupational inhalation exposure; task-resolved exposure measurement; gas sensing materials and devices; respiratory and personal protective equipment performance; safety-barrier theory and the four barrier-management traditions identified in Section 1.5; and organisational safety performance. Each source was characterised by study type, population or material studied, outcome measured, and the design decision it could legitimately inform.

2.3 Block 1 — Hazard and scenario inventory

For each process unit, target substances are listed with their release mechanisms and assigned to one or more exposure archetypes: continuous background emission; task-linked excursion; containment failure; and confined-space accumulation or oxygen displacement.

The archetype determines the temporal resolution that monitoring must achieve. Continuous background emission can be characterised by fixed sampling at representative points. Task-linked excursion cannot: the exposure is defined by the worker’s position relative to a transient source, and the task-resolved measurement literature indicates how large the discrepancy between the two can be (Ridderseth et al., 2023). Containment failure requires

detection sited against dispersion behaviour rather than against occupancy. Confined-space accumulation requires pre-entry and continuous in-space measurement including oxygen.

The inventory is scenario-based rather than substance-based: a substance appears once per release mechanism per location, not once per facility. Where the facility maintains barrier diagrams, each scenario in this inventory corresponds to a top event on one of them, and the detection and respiratory protection barriers specified in Blocks 3 and 4 are the same barriers that appear on its consequence leg (CCPS & Energy Institute, 2018). The inventory is thus not a parallel register but the specification-level view of one already in use — the first of the joins in Section 1.5.

2.4 Block 2 — Risk tiering

Likelihood and consequence are expressed through qualitative descriptors rather than invented numbers. Likelihood descriptors combine release frequency, the number of workers who can plausibly be present, and the duration of the task. Consequence descriptors combine acute toxicity and warning properties of the substance, the possibility of oxygen deficiency, and the health endpoints attributable to occupational inhalation exposure, including chronic airway disease and interstitial processes (Blanc et al., 2019). Substances with poor warning properties or with steep dose–response behaviour at low concentrations are escalated — the hydrogen sulfide case is definitional here (Guidotti, 2010) — as are scenarios involving workers with pre-existing airway disease, consistent with the reactivity of sensitised individuals to very low irritant levels (Leikauf, 2002).

The output is a tier assigned to each scenario, not to each substance in isolation. The tier scale is that of the parent methodology's risk-ranking step (Jilgildiyev, 2026), so that a facility operating the IRMM does not maintain two incompatible schemes.

2.5 Block 3 — Detection layer

2.5.1 Sensing technology selection

Selection follows the comparative characteristics reported in the materials literature rather than catalogue sensitivity. Metal oxides are attractive for ease of fabrication, low cost, high sensitivity and stability, but are limited by selectivity and by high operating temperature; conducting polymers operate at low temperature and respond to many organic vapours but are humidity-affected; carbon nanotube devices are chemically and mechanically stable with sensitivity toward nitric oxide and ammonia but commonly require doping or modification (Nikolić et al., 2020). Where morphology-dependent performance is exploitable, fabrication conditions become part of the procurement specification rather than a supplier detail (Xu et al., 2021). For ammonia service, the breadth of material and transduction options argues for explicit selection criteria against operating temperature, expected concentration range and cross-interferent profile rather than default choice (Kwak et al., 2019).

2.5.2 Architecture and layer independence

Three monitoring layers are specified and are not interchangeable: fixed detection for containment monitoring and automatic action; area monitors for temporary work zones; personal monitors for task-linked exposure. Fixed detection cannot establish personal exposure; personal monitoring cannot protect an unattended plant. Each scenario is assigned to at least one layer by its exposure archetype.

The distinction matters for more than coverage. Where a detection function is credited as an independent protection layer for a scenario, the credit assumes that the layer is independent of the other layers credited for the same scenario and that it is effective for the specific initiating event (CCPS, 2001). Two conditions are therefore attached explicitly in this framework, and both are conditions on the credit rather than on the equipment: the layer must not share a common cause with another credited layer — a shared power supply, a shared logic solver, a shared calibration gas and a shared technician all defeat the assumption — and its alarm must be credible in the ambient conditions of the installation in the sense defined in Section 2.5.3. A detection layer that does not satisfy both is documented in the barrier register as present but not credited. This is the second join in Section 1.5, and stating it as a specification rule is what keeps it from being lost between the risk group and the instrument group.

2.5.3 Alarm logic and credibility

Alarm logic distinguishes advisory, action and evacuation thresholds, each with a documented required response and a defined recipient. Cross-sensitivity, drift and humidity dependence are recorded in the specification as known failure modes, to be managed by calibration schedule, bump-test frequency and, where justified, sensor-array redundancy or voting. Detector availability is specified as a value with a verification method, not assumed. These parameters appear in the same specification section as the detection limit, for the reason set out in Section 3.3.

2.5.4 Interface with the instrumented function

Where a fixed detection function initiates automatic action — executive action, emergency shutdown, deluge — that function forms part of a safety instrumented system, and the integrity with which it must act is determined under the functional-safety standards (IEC, 2010, 2016). This framework does not duplicate that determination. It defines the interface in both directions, which is the third join in Section 1.5 and the one most often left implicit:

- Delivered to the instrumented function: the scenario and its tier; the target substance, its expected concentration range and its interferent profile; the required detection architecture and placement basis; the alarm classes and their intended consequences; and the assumed availability with the verification method that will demonstrate it.
- Required back from it: confirmation that the assumed availability is met in service; notification when measured availability, proof-test results or demand behaviour depart from the design assumption; and an explicit statement of which credited layers the instrumented function is and is not independent of.

Where the second of these is not returned, the availability assumed in Section 2.5.3 and the credit taken in Section 2.5.2 are unverified, and the barrier is recorded accordingly.

2.6 Block 4 — Respiratory protection layer

The same tiers map onto respiratory protection requirements. Device class is specified by tier and by whether the atmosphere can be immediately dangerous to life or health, with escape sets and self-contained breathing apparatus reserved for scenarios in which detection cannot provide reliable advance warning — a criterion that follows directly from the warning-property term in the consequence descriptor.

Filtration performance is a necessary but insufficient criterion, since laboratory efficiency does not translate directly into workplace protection (Rengasamy et al., 2004) and measured protection across wearers and activities is distributed rather than fixed (Duling et al., 2007). Fit testing, seal checking, and written donning and doffing sequences are therefore specified as mandatory at the higher tiers, because equipment type and donning or doffing method both affect outcome (Verbeek et al., 2020). Training and supervision requirements are attached to each tier explicitly.

Respiratory protection is placed last in the control order by construction, and is not permitted to substitute for a detection or engineering control that the tiering step has identified as required. Where it appears as a mitigative barrier on the consequence leg of a scenario, it is credited only for personnel whose fit-test record is current and whose donning and doffing procedure has been observed; absent those records the barrier is documented as present but not credited, on the same principle as Section 2.5.2. A barrier whose performance is not verified does not meet the definition of a barrier (Sklet, 2006).

Table 2 – Structure of the tier-driven specification (schematic).

Tier	Typical scenario character	Detection requirement	Respiratory protection requirement	Verification establishing the credit
Low	Routine operation, well-ventilated, good warning properties, continuous fixed monitoring	Fixed detection at representative points; advisory and action thresholds	Standard workwear; escape provision per facility rule	Scheduled calibration; periodic review
Medium	Recurrent task in an area with fixed detection but variable occupancy	Fixed plus area monitoring for the duration of the work	Specified filtering device class; fit test on record; seal check per use	Calibration plus bump test; fit-test register
High	Task-linked excursion; poor warning properties; transient source near the breathing zone	Personal monitoring mandatory; area monitoring; documented response per alarm class	Device class by atmosphere; mandatory fit testing; written donning/doffing sequence; supervision	Bump test before each task; observed doffing; task-level records
Highest	Confined-space entry, line breaking, potential IDLH or oxygen deficiency; detection cannot give reliable advance warning	Continuous in-space monitoring including oxygen; pre-entry test; attendant	Supplied-air or SCBA; escape set; rescue provision specified before entry	Pre-entry verification; drill-tested rescue capability

The table states the structure of the mapping and the class of requirement at each tier. Numerical thresholds, device designations and calibration intervals are facility-specific and are set during application against the substance inventory, the applicable national regulation and the manufacturer's data. They are deliberately not prescribed here.

2.7 The connecting element: barrier condition indicators and the return path

The three blocks above specify barriers. The fourth element makes the specification a cycle rather than a document.

For each barrier specified in Blocks 3 and 4, the framework requires the five attributes that the parent methodology requires of any critical barrier (Jilgildiyev, 2026): function, performance criterion, owner, verification method and frequency, and the criterion of unacceptable condition with the action that follows from it. Condition is then tracked through the indicator set given in Section 3.7, drawing on the barrier-management practice in which barrier status is followed as a standing indicator rather than reconstructed after an event (Petroleum Safety Authority Norway, 2013).

The return path is what closes the loop, and it is specified as a rule rather than left to good practice: a measured departure from the performance criterion of a credited barrier requires re-entry into the assessment that credited it. Concretely — a sustained rise in false-alarm rate at a location, a measured availability below the value assumed in Section 2.5.3, or a lapse in fit-test currency for personnel assigned to a high-tier task each trigger review of the corresponding tier assignment and of the layer credit taken in Section 2.5.2 or Section 2.6, not merely a maintenance work order. This is the fourth join in Section 1.5, and its absence is the mechanism by which a facility's documented barrier set and its actual barrier set diverge without anyone observing the divergence.

2.8 Internal appraisal

Appraisal was limited to internal checks, and the limits of what such checks establish are stated in Section 5.

Consistency was tested by working the framework through generic scenarios representing sour gas processing, tank farm operation, ammonia refrigeration and confined-space entry, and by verifying that each design rule traced back either to a cited source or to an explicitly stated engineering assumption. Coupling was tested by confirming that no scenario could receive a detection requirement and a protection requirement from different tiers, and that every credited barrier carried the five attributes required in Section 2.7 together with a defined return path. Organisational feasibility was considered through the safety-climate benchmarking evidence, which supports periodic measurement of climate dimensions alongside technical audit (Meams et al., 2001).

No validation against exposure measurements or incident records was performed within this study. That constraint is carried forward into Sections 4 and 5.

3. Results and Discussion

3.1 Where risk concentrates

Applying the tiering step to the generic scenarios produced a consistent pattern: risk concentrated not in the substances with the highest nominal toxicity but in the combinations of substance, task and location where warning properties were poor and monitoring resolution was coarse. Confined-space entry and line breaking repeatedly rose to the highest tier, while routine operation of well-ventilated units with continuous fixed monitoring fell to lower tiers even where the inventory included highly toxic gases.

The practical consequence is a reallocation rule: resources spent adding fixed detection points in already well-characterised areas yield less risk reduction than resources spent on task-linked monitoring and on protection discipline during short, high-exposure operations. This is consistent with the task-resolved exposure evidence, in which between-task variance dominates (Riddersteth et al., 2023), and it is a conclusion a substance-by-substance assessment cannot reach, because the unit of assessment is wrong.

3.2 Sensor selection and the risk of over-specification

Sensor selection is the most technology-dependent part of the framework and the part most vulnerable to over-specification. No single technology satisfies all requirements of a mixed-hazard facility: metal oxide devices offer cost and sensitivity at the price of selectivity and elevated operating temperature, while conducting polymers offer low-temperature operation and organic vapour response but remain humidity-sensitive (Nikolić et al., 2020). For an installation with sustained high humidity — the ordinary condition of a coastal or offshore facility — humidity dependence is not a laboratory caveat but a determinant of alarm credibility.

Morphology-tuned nanostructured titanium dioxide shows that response can be engineered through fabrication rather than only through signal processing, which supports writing material and structural requirements into procurement specifications for demanding services (Xu et al., 2021). Where ammonia is a principal hazard, the diversity of sensing approaches indicates that selection should be justified explicitly against operating temperature, expected concentration range and cross-interferent profile (Kwak et al., 2019).

3.3 Alarm credibility is the limiting parameter — and the point at which layer credit fails silently

Alarm credibility, rather than raw sensitivity, was the recurring limiting factor in every scenario worked, and the mechanism is as much organisational as technical. Cross-sensitivity and drift produce false positives; false positives erode response discipline; eroded discipline converts a functioning detection system into a nominal one. The system continues to satisfy its design documentation while no longer protecting anyone.

The significance of this for barrier management is greater than its significance for instrumentation, and it is the clearest illustration of why the four elements must be coupled. A detection layer credited as an independent protection layer carries an implicit assumption that it will be acted upon (CCPS, 2001). Erosion of response discipline invalidates that assumption without altering a single item of hardware, without generating a maintenance finding, and without appearing anywhere in the integrity verification of the instrumented function — because the integrity verification addresses whether the function acts, not whether anyone acts on it. Under an open lifecycle the credit therefore remains on the books indefinitely. Under the coupling rule of Section 2.7, a sustained rise in false-alarm rate is a condition indicator that re-enters the assessment and puts the credit in question.

Two consequences follow for specification. Calibration interval, bump testing and a documented response per alarm class are treated as design parameters of equal standing with detection limit and specified in the same document section rather than deferred to a maintenance procedure. And detector availability is specified as a value with a verification method and a notification obligation when service behaviour departs from it, as set out in Section 2.5.4.

3.4 Respiratory protection as the last layer

The framework subordinates respiratory protection to engineering and administrative control, then specifies it rigorously where it is still required. Filtration efficiency measured against test aerosols overstates protection achieved in service, because face-seal leakage, wear time and user behaviour intervene (Rengasamy et al., 2004). Measured protection is a distribution rather than a value, so the protection actually delivered to a given worker on a

given task cannot be read off the device class (Duling et al., 2007). And systematic review evidence indicates that equipment type and donning/doffing procedure both influence outcome, while cautioning that certainty in this literature is low (Verbeek et al., 2020).

Low certainty is an argument for local verification, not for indifference: fit testing, observed doffing and incident review generate facility-specific evidence where the published literature cannot. This is why the framework treats the fit-test register and the donning/doffing record as constitutive of the barrier rather than as documentation about it (Sklet, 2006), and why an expired fit test removes the credit rather than generating a training action only.

3.5 Closing the loop between monitoring data and health objectives

If the objective were only to prevent acute poisoning, detection thresholds tied to short-term exposure limits would suffice. The respiratory literature indicates a broader objective, since occupational inhalation exposures contribute to chronic airway disease, interstitial fibrosis, hypersensitivity pneumonitis and other non-malignant conditions (Blanc et al., 2019).

Cumulative low-level exposure therefore requires surveillance in its own right, handled here by trending personal monitoring data across tasks rather than by alarm counting: an alarm count measures excursions above a threshold and is structurally blind to the accumulation of sub-threshold exposure. Susceptibility is treated explicitly, with conservative tier assignment for workers with pre-existing airway disease (Leikauf, 2002).

3.6 Organisational conditions

Longitudinal benchmarking of safety climate across offshore installations demonstrated both that climate elements relate to self-reported accidents and that substantial absolute improvement, with convergence between installations, is attainable over a year (Mearns et al., 2001). This supports embedding climate measurement into the same review cycle as detector calibration and respirator fit testing, so that technical and behavioural indicators are interpreted together rather than by different functions on different schedules. A facility with drifting sensors and a permissive attitude to seal checking is protected only nominally in precisely the scenarios the tiering step identifies as critical, and each deficiency makes the other harder to detect.

3.7 The closed loop in operation: one indicator set across four elements

The parent methodology requires that barrier performance be monitored through leading and lagging indicators (Jilgildiyev, 2026). The framework makes that requirement operational for the two barrier classes it specifies, and the indicator set is deliberately single: detection and protection indicators are reviewed together, because a tier assignment met on one side and not the other is a failure of the specification even when both sides appear compliant in isolation.

Table 3 – Indicator set and its return path.

Barrier-management element	Leading indicator	Lagging indicator	Return path when the criterion is not met
Identification and representation	Proportion of scenarios in the inventory with both barriers specified and owned	Scenarios found during review to be absent from the inventory	Re-open the inventory and the corresponding barrier diagram
Layer independence and sufficiency	Proportion of credited detection layers with documented independence and demonstrated alarm credibility	Demands on which a credited layer did not act as assumed	Re-examine the credit; record the layer as uncredited until restored
Integrity of the instrumented layer	Measured availability against the value assumed at design; proof-test completion	Failures on demand; spurious trips	Notify under Section 2.5.4; revise the assumed availability and the dependent credit
Barrier condition — detection	Calibration currency; bump-test completion before high-tier tasks; false-alarm rate per detector-month by location; alarm-response completion by class	Alarms with no recorded response; exposures above tier assumptions	Review tier assignment and layer credit, not only the maintenance record
Barrier condition — respiratory protection	Fit-test currency for personnel assigned to high-tier tasks; observed seal-check and doffing compliance; proportion of high-tier tasks with personal monitoring deployed	Unplanned escape-set deployments; near misses in high-tier tasks	Remove the credit for affected personnel; review task assignment
Organisational condition	Safety-climate dimension scores measured on the calibration and fit-testing cycle	Incidents	Joint technical and organisational review

The asymmetry between the leading and lagging columns is intentional and follows the established critique of process safety indicators: low-frequency, high-consequence events cannot demonstrate that a system is working, so an absence of incidents is not evidence of control and the weight of evaluation must fall on the leading set (Hopkins, 2009).

4. Application and Independent Verification

The framework is presented for application and independent verification at operating installations. Because Section 2.1 states that the practice informing it is not reportable in facility-attributable form, this section specifies how an operator generates its own evidence, and what result would count against the framework.

Application. The tiering step is applied to the full scenario register; detection and respiratory protection specifications are rewritten from the resulting tiers; the five barrier attributes of Section 2.7 are recorded for every credited barrier; and the indicator set of Table 3 is placed on the existing review cycle. A before-and-after comparison at a single installation is the minimum design; a parallel comparison across two or more installations of the same operator, with existing practice retained at one, supports stronger inference where it is organisationally feasible.

Verification metrics. Monitoring deployment at high-tier tasks and task-resolved exposure coverage should rise. False-alarm rate per detector-month and alarm-to-response time, with its variance, should fall. Measured detector availability should converge on the value assumed at design. Fit-test currency and observed seal-check compliance for high-tier task personnel should rise. The number of scenarios whose detection and protection requirements originate from different tiers, and the number of credited barriers lacking any of the five attributes, should go to zero. The ratio of expenditure on task-linked monitoring and protection discipline to expenditure on additional fixed detection should rise without an increase in uncontrolled scenarios. Lagging measures are interpreted with the caution in Section 3.7.

Testing the mitigative layer. Scenario-based drills are the only practical means of testing mitigative barriers, since the events of interest are too rare to be observed. Target times for detection, alarm, isolation, evacuation, personnel accounting and mobilisation are set before the drill and compared with actual times afterwards; a significant deviation is treated as information about the status of the corresponding barrier and enters the return path of Section 2.7. A drill in which detection and alarm meet target while isolation and evacuation do not has identified a barrier deficiency, not produced a pass. Any such scenario is documented as simulated, and its outputs are not reported as operational performance data.

What would count against the framework. If, after application, high-tier tasks continue to lack task-resolved monitoring; if false-alarm rates do not fall despite specified calibration and bump-test regimes; if fit-test currency does not translate into observed seal-check compliance; if condition indicators do not in practice re-enter the assessments that credited the barriers; or if the reallocation of resources from fixed detection to task-linked monitoring is accompanied by a rise in uncontrolled scenarios — then the central claim, that coupling the four elements at the barrier captures risk that they miss separately, is not supported.

5. Limitations

The framework derives from a heterogeneous source set spanning clinical epidemiology, sensing materials, protective-equipment research, barrier theory and organisational studies. Several design rules rest on engineering reasoning rather than on direct empirical support in oil and gas settings; these are identified as such at the point of use, and the citation of a source adjacent to a rule should not be read as evidence that the rule has been tested.

A significant part of the protective-equipment evidence originates in healthcare and infection-control contexts, where the hazard is contaminated body fluid or bioaerosol rather than toxic gas or oxygen deficiency (Verbeek et al., 2020). The industrial protection-factor evidence cited (Duling et al., 2007) partially offsets this, but transfer from the clinical literature remains qualitative.

The sensing literature reports laboratory performance under controlled conditions (Kwak et al., 2019; Nikolić et al., 2020; Xu et al., 2021). Field behaviour in the presence of hydrocarbon interferents, salt aerosol and temperature cycling is not thereby established, and that gap is precisely what Section 3.3 identifies as decisive.

The safety-climate evidence (Mearns et al., 2001) is drawn from the North Sea offshore sector; its transferability to onshore installations in other regulatory and cultural settings is not established here.

The coupling rules in Sections 2.5.2, 2.5.4 and 2.7 are specification requirements, not empirical findings. The claim that coupling the four elements reduces the divergence between documented and actual barrier sets is argued from the mechanism and from the reported heterogeneity of barrier practice (Yuan et al., 2022), and has not been measured. Section 4 states how it can be.

As set out in Section 2.1, the practice informing this framework is not reportable in facility-attributable form, so the reader cannot verify the framework against it. No quantitative performance claim is made and none should be inferred.

6. Conclusion

Gas detection and respiratory protection at high-hazard oil and gas facilities are best specified as two layers of one barrier architecture, coupled through a common risk-tiering step. The framework presented here links a hazard and scenario inventory to qualitative risk tiers and derives both detection and respiratory protection requirements from the same tier assignment, so that the two cannot drift apart.

Its central contribution is structural. The four elements required to manage a barrier over its life — identification and representation, independence and sufficiency of credited layers, integrity of the instrumented layer, and condition indicators with a return path — are established individually and are claimed as new by no one. They are, however, held in separate documents, by separate functions, on separate cycles, and the exposure occurs at the joins. This framework closes them into a single cycle at the two barriers through which inhalation exposure is controlled, which is the same closure that

the author's Integrated Risk Management Methodology performs at facility-architecture level and which this paper extends to specification level and to the mixed inhalation-hazard inventory.

The analytical results follow from that closure. Risk concentrates in task-linked, short-duration operations with poor warning properties, so monitoring resolution and protection discipline at those moments matter more than the aggregate number of installed detectors. Alarm credibility, not catalogue sensitivity, is the limiting design parameter, and an incredible alarm invalidates the layer credit taken for the detection barrier silently — a failure that only a closed lifecycle detects. Respiratory protection remains the last layer and is credited only where fit testing and procedural discipline are verified. Health objectives encompass chronic as well as acute endpoints, and susceptible workers are assigned conservatively. Technical layers require an organisational one: safety-climate measurement belongs on the same review cycle as calibration and fit testing.

The framework is offered for application and independent verification. Section 4 states how an operator generates its own evidence and what result would count against it.

Acknowledgements

Not applicable.

Conflict of Interest

The author declares no conflict of interest.

Ethical Statements

Not applicable. The study involved no human participants, human data or animals.

Data Availability

No new data were created or analysed in this study. All sources supporting the framework are cited in the reference list. Facility-attributable information is not reported, for the reasons stated in Section 2.1.

REFERENCES

1. Blanc, P. D., Annesi-Maesano, I., Balmes, J. R., Cummings, K. J., Fishwick, D., Miedinger, D., et al. (2019). The occupational burden of nonmalignant respiratory diseases: An official American Thoracic Society and European Respiratory Society statement. *American Journal of Respiratory and Critical Care Medicine*, 199(11), 1312–1334. <https://doi.org/10.1164/rccm.201904-0717ST>
2. Center for Chemical Process Safety. (2001). Layer of protection analysis: Simplified process risk assessment. American Institute of Chemical Engineers.
3. Center for Chemical Process Safety, & Energy Institute. (2018). Bow ties in risk management: A concept book for process safety. John Wiley & Sons and AIChE. <https://doi.org/10.1002/9781119490357>
4. de Ruijter, A., & Guldenmund, F. (2016). The bowtie method: A review. *Safety Science*, 88, 211–218. <https://doi.org/10.1016/j.ssci.2016.03.001>
5. Duling, M. G., Lawrence, R. B., Slaven, J. E., & Coffey, C. C. (2007). Simulated workplace protection factors for half-facepiece respiratory protective devices. *Journal of Occupational and Environmental Hygiene*, 4(6), 420–431. <https://doi.org/10.1080/15459620701346925>
6. Guidotti, T. L. (2010). Hydrogen sulfide: Advances in understanding human toxicity. *International Journal of Toxicology*, 29(6), 569–581. <https://doi.org/10.1177/1091581810384882>
7. Hopkins, A. (2009). Thinking about process safety indicators. *Safety Science*, 47(4), 460–465. <https://doi.org/10.1016/j.ssci.2007.12.006>
8. International Electrotechnical Commission. (2010). IEC 61508:2010. Functional safety of electrical/electronic/programmable electronic safety-related systems. IEC.
9. International Electrotechnical Commission. (2016). IEC 61511-1:2016. Functional safety — Safety instrumented systems for the process industry sector — Part 1: Framework, definitions, system, hardware and application programming requirements. IEC.
10. Jilgildiyev, A. (2026). Integrated risk management methodology for hydrogen sulfide at high-risk oil and gas facilities. *Scientific Culture*, 12(5), 873–892. <https://doi.org/10.5281/zenodo.20678843>
11. Kwak, D., Lei, Y., & Marić, R. (2019). Ammonia gas sensors: A comprehensive review. *Talanta*, 204, 713–730. <https://doi.org/10.1016/j.talanta.2019.06.034>
12. Leikauf, G. D. (2002). Hazardous air pollutants and asthma. *Environmental Health Perspectives*, 110(Suppl. 4), 505–526. <https://doi.org/10.1289/ehp.02110s4505>
13. Mearns, K., Whitaker, S. M., & Flin, R. (2001). Benchmarking safety climate in hazardous environments: A longitudinal, interorganizational approach. *Risk Analysis*, 21(4), 771–786. <https://doi.org/10.1111/0272-4332.214149>
14. Nikolić, M. V., Milovanović, V., Vasiljević, Z. Ž., & Stamenković, Z. (2020). Semiconductor gas sensors: Materials, technology, design, and application. *Sensors*, 20(22), 6694. <https://doi.org/10.3390/s20226694>

15. Petroleum Safety Authority Norway. (2013). Principles for barrier management in the petroleum industry. Petroleum Safety Authority Norway.
16. Rengasamy, A., Zhuang, Z., & BerryAnn, R. (2004). Respiratory protection against bioaerosols: Literature review and research needs. *American Journal of Infection Control*, 32(6), 345–354. <https://doi.org/10.1016/j.ajic.2004.04.199>
17. Riddersest, H., Daltveit, D. S., Hollund, B. E., Kirkeleit, J., Kromhout, H., Krüger, K., et al. (2023). Benzene exposure from selected work tasks on offshore petroleum installations on the Norwegian continental shelf, 2002–2018. *Annals of Work Exposures and Health*, 67(2), 228–240. <https://doi.org/10.1093/annweh/wxac067>
18. Sklet, S. (2006). Safety barriers: Definition, classification, and performance. *Journal of Loss Prevention in the Process Industries*, 19(5), 494–506. <https://doi.org/10.1016/j.jlp.2005.12.004>
19. Verbeek, J. H., Rajamaki, B., Ijaz, S., Sauni, R., Toomey, E., Blackwood, B., et al. (2020). Personal protective equipment for preventing highly infectious diseases due to exposure to contaminated body fluids in healthcare staff. *Cochrane Database of Systematic Reviews*, 2020(4), CD011621. <https://doi.org/10.1002/14651858.CD011621.pub5>
20. Xu, T., Cui, X., Lai, T., Ren, J., Yang, Z., Xiao, M., et al. (2021). Gas sensors based on TiO₂ nanostructured materials for the detection of hazardous gases: A review. *Nano Materials Science*, 3(4), 390–403. <https://doi.org/10.1016/j.nanoms.2021.05.011>
21. Yuan, S., Yang, M., Reniers, G., Chen, C., & Wu, J. (2022). Safety barriers in the chemical process industries: A state-of-the-art review on their classification, assessment, and management. *Safety Science*, 148, 105647. <https://doi.org/10.1016/j.ssci.2021.105647>