AURUS INSTITUTE FOR RESOURCE DEVELOPMENT

18

WHITE PAPER · NUMBER EIGHTEEN

Zero Is an Engineering Problem

Critical Control Management and the reversal of mining's fatality curve

TECHNICAL PAPER · SHARE

July 2026 · ~55-min read

Aurus Technical Committee

AURUS INSTITUTE FOR RESOURCE DEVELOPMENT

Evidence before assertion.

MINING · INFRASTRUCTURE · ENGINEERING · ENVIRONMENT

Contents

The curve turned the wrong way, and it did so for reasons we understand

In 2024 the mining industry's own benchmark recorded 42 fatalities across the 24 member companies of the International Council on Mining and Metals, up from 36 in 2023 and 33 in 2022. That is not a statistical wobble. It is the second consecutive annual increase in fatalities off the 2022 low, a reversal of a decade-long decline, and it happened while recordable injuries kept falling. The two curves separated, and the separation is the story: the systems that reduce sprains and cuts are not the systems that stop people being killed.

This paper reads the public safety record of the world's largest mining companies as an engineering problem rather than a moral exhortation. It is written for the boards, operators and lenders who are asked to accept "zero harm" as an aspiration and who deserve to be shown, instead, the causal architecture that makes zero a design target: the small number of major-hazard events that kill, the critical controls that stop them, and the failure mode, execution in the operating environment, that accounts for most of the deaths.

The reversal is real and measured. Fatal incidents rose 23.3 per cent in 2024, to 37 from 30 the year before, with a fatal-frequency rate of 0.015 per million hours worked against 0.013 in 2023. Only 9 of 24 member companies recorded zero fatalities in 2024, down from 12 of 25 the year before. Over the long series, total recordable injury frequency fell from 5.07 in 2012 to 2.29 in 2024, genuine, sustained progress on injuries, even as the fatality line refused to follow it down.

The geography is uncomfortable and must be named. Operations in Africa accounted for 50 per cent of 2024 fatalities, 21 of 42, on 24 per cent of hours worked, a fatal-frequency rate of 0.031, the highest of any continent. South Africa alone recorded 15 fatalities, 14 of them underground; all 18 underground fatalities in 2024 occurred on just two continents, with 15 in Africa. This is not a reason to disengage from African mining. It is the reason to engineer for it, at the standard the hazard demands.

The causes converge on execution. In 2024, failure to implement effective critical controls accounted for 83 per cent of fatalities; across 2021–2024, 67 per cent of control failures were failures of execution in the operating environment, not of design and not of the control being absent. The controls, in most cases, existed. They were understood, resourced and documented, and on the day they were not effective where the work was done. That single finding reorganises the whole safety problem: from writing better controls to verifying that the ones we have are working, every shift, at the face.

The killers have names, and most can be engineered out. Mobile equipment was the leading hazard in 2024, at 21 per cent of fatalities, seven of nine at operations in Africa; fall of ground was second, at five fatalities, four in South African operations. Three of the five fall-of-ground incidents involved work before full mesh-and-bolt support was installed or repaired. Mobile-equipment interactions and fall of ground together account for roughly two-thirds of underground fatalities. These are precisely the exposures that engineering, collision-avoidance systems, exclusion zones, support-before-entry discipline, is built to remove.

The framework caught up in 2026. The ICMM's Critical Control Management Good Practice Guide was comprehensively updated in 2026, its first revision in a decade, merging the 2015 good-practice and implementation guides into a single nine-step process and, for the first time, positioning critical-control management inside enterprise risk management. This paper reads the 2024 data and the 2026 framework together, because the framework is the answer to the data.

> **Critical control** · a control that is crucial to preventing a fatal event or mitigating its consequences, and whose absence or failure would significantly increase the risk despite the existence of other controls. The word "critical" is doing analytical work: it names the few controls that carry the fatal risk, so that verification effort concentrates where a failure kills.

What the record supports: five findings

- **Fatalities and injuries have decoupled.** Injury frequency more than halved over 2012–2024 while the fatality line reversed upward from 2022. A safety programme optimised for injury rates can be improving and failing at the same time.
- **Execution, not design, is the dominant failure mode.** 67 per cent of control failures over four years were failures of execution in the operating environment, and 83 per cent of 2024 fatalities involved a failure to implement effective critical controls.
- **The fatal risk is concentrated.** A small set of hazards, mobile equipment, fall of ground, falling objects, and a small set of settings, underground, Africa, carry most of the deaths. Concentration is opportunity: it tells engineering where to aim.
- **Organisational factors are under-recorded.** Failure to follow rules or procedures was the top organisational factor in 2024, yet only one fatality was formally classified under "safety culture", a sign that cultural drivers are being missed in investigation, not that they are absent.
- **The 2026 framework is fit for the diagnosis.** The updated CCM guide's nine steps, maturity assessments and readiness checks are built precisely to attack execution and verification, the failure mode the data identifies.

Who should read this paper

This report addresses two audiences. **Operators and their boards** will find in Chapters 1 through 6 a causal reading of the fatal record and, in Chapter 7, a critical-control operating system to hold a site against. **Lenders, insurers and regulators**, the parties who price and permit mining risk, will find a defensible, source-anchored account of where the fatal exposure actually sits, and a vocabulary for asking whether an operator's controls are verified rather than merely written.

The method is deliberately conservative. Every statistic is transcribed from a cited public source, principally the ICMM's own benchmarking report and good-practice guides, and archived in the paper's evidence dossier; ranges are preserved as published; and where a figure is read from a chart rather than a data table, it is flagged as such and never used as a false precision. The back-matter page Method and evidence records the discipline in full, including the paper's treatment of chart-read values and its register of open parameters.

How to use this paper

Three reading paths serve three purposes. Read **front to back** for the argument: from the reversed curve to the framework that answers it. Read **Chapter 7 first** for the instrument: the CCM-based operating system and its leading-indicator dashboard, then trace each element back to the evidence that earns it. Or read **by exhibit**: the seventeen exhibits and six stat tiles, each with its source line burned in, carry the paper's entire quantitative content, and the exhibit index in the back matter maps every one to its references. The framework centrepiece, the critical-control bowtie of Chapter 3, is labelled as a capability instrument and carries no project data.

Six figures that frame the fatality reversal

42

FATALITIES ACROSS ICMM'S 24 MEMBER COMPANIES, 2024 (36 IN 2023, 33 IN 2022)

ICMM, Safety Performance Report, 2025

83%

OF 2024 FATALITIES INVOLVED A FAILURE TO IMPLEMENT EFFECTIVE CRITICAL CONTROLS

ICMM, Safety Performance Report, 2025

50%

OF 2024 FATALITIES AT OPERATIONS IN AFRICA, ON 24% OF HOURS WORKED

ICMM, 2025 (Africa fatal-frequency rate 0.031, highest of any continent)

67%

OF CONTROL FAILURES, 2021–2024, WERE FAILURES OF EXECUTION IN THE OPERATING ENVIRONMENT

ICMM, 2025 (Graph 10)

21%

OF 2024 FATALITIES FROM MOBILE EQUIPMENT: THE LEADING SINGLE HAZARD

ICMM, 2025 (seven of nine at operations in Africa)

45%

OF 2024 FATALITIES WERE CONTRACTORS, WORKING 58% OF HOURS (FATAL RATE 0.012, BELOW EMPLOYEES' 0.020)

ICMM, 2025 (contractor share fell from 69% in 2023)

How to read this page. Each figure is transcribed from the ICMM's own benchmarking publication and re-appears, in context, in the chapters that follow. Ranges and qualifiers are preserved as published; values read from a chart rather than a data table are flagged where they are used.

How this paper is organised

CHAPTER	THE QUESTION IT ANSWERS	ANCHOR EVIDENCE
1 · The uncomfortable data	What did the curve actually do, and where are the deaths concentrated?	ICMM, 2025
2 · Pyramids to precursors	Why did injuries fall while fatalities rose? The SIF decoupling	ICMM, 2025; SIF concept, gap-flagged
3 · Critical Control Management	What has to be true for a control to count as critical, and to work?	ICMM CCM Guide, 2026
4 · The big killers	Which hazards kill, and how does engineering remove the exposure?	ICMM, 2025; EMESRT
5 · Field execution	Why do good controls fail at the face, and how are contractors aligned?	ICMM, 2025 (Graphs 8, 10)
6 · Health beyond safety	What kills slowly, and does not show up in a fatality count?	ICMM, 2025 (DPM commitment)
7 · The operating system	How does a site hold itself to all of the above? A CCM-based instrument	Chapters 1–6, cited record

What this paper is, and is not

It is a causal reading of the mining industry's public fatal-safety record, built from the ICMM's own benchmarking data and good-practice guides, and a capability framework argued from that record. It is not an audit of any named company, a ranking, or a claim of delivered corporate-HSE mandates by the author. Every figure carries its institutional source; chart-read values are flagged; and the closing operating system is a capability instrument, framed from the data discipline of a delivered field campaign and from published good practice, never from safety-performance data the author does not hold.

1

SHARE · THE UNCOMFORTABLE DATA

The uncomfortable data

Injuries kept falling. Fatalities have turned back up for a second straight year off their 2022 low. The two curves have separated, and the gap is where this paper lives.

42	+23.3%	50%
FATALITIES, ICMM 24 MEMBER COMPANIES, 2024 (ICMM, 2025)	RISE IN FATAL INCIDENTS, 2024 VS 2023 (ICMM, 2025)	OF 2024 FATALITIES AT OPERATIONS IN AFRICA, ON 24% OF HOURS (ICMM, 2025)

For most of the last decade the mining industry could tell a clean story about safety: the numbers were getting better. Then the fatality line stopped cooperating. In 2024 the International Council on Mining and Metals recorded 42 fatalities across its 24 member companies, up from 36 in 2023 and 33 in 2022, a second consecutive annual increase off the 2022 low and a reversal of the long-term decline. Two straight years of rising deaths is a trend, not an accident, and the industry's own benchmark now carries it in black and white.

The detail is worse than the headline. Fatal incidents rose 23.3 per cent, to 37 in 2024 from 30 in 2023; three of those incidents each caused more than one fatality; and the fatal-frequency rate rose to 0.015 per million hours worked from 0.013. The share of member companies achieving zero fatalities fell from 12 of 25 in 2023 to 9 of 24 in 2024. Whatever the industry was doing to hold the line, in 2024 it did not hold.

A decade of progress, and a line that will not follow it down

None of this erases genuine progress on injuries. The ICMM's long series runs from 2012, when it standardised disclosure across a common indicator set, and over that period total recordable injury frequency fell from 5.07 in 2012 to 2.29 in 2024, against 2.81 billion hours worked in the latest year. Recordable injuries themselves fell to 6,446 in 2024 from 7,278 in 2023, a total recordable injury frequency of 2.29 against 2.59. Injuries are the one curve that behaved. Exhibit 1 places the two lines on one axis, and the divergence is the argument of this paper made visible.

EXHIBIT 1

Injuries halved; the fatal line reversed. The two curves have separated

ICMM member companies, indexed indicators, 2012–2024. Left axis: total recordable injury frequency (per million hours). Fatalities plotted as annual counts

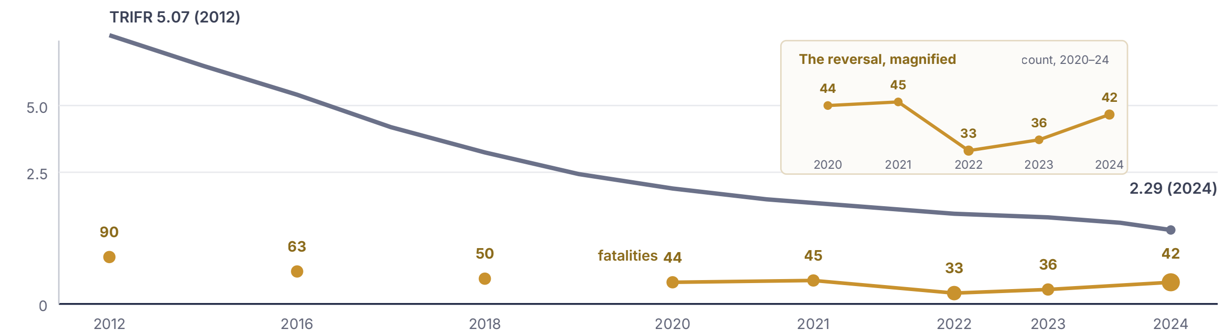

Source: ICMM, Safety Performance Report: Benchmarking Progress of ICMM Company Members in 2024, July 2025 (Table 1 long series; injury and fatality counts). The 2019 spike (287 fatalities, including the 250 Brumadinho tailings-collapse deaths) is excluded from this view to keep the trend legible and is discussed in the text. TRIFR line is schematic between the two published anchor values (5.07 in 2012; 2.29 in 2024); fatality markers are plotted at published annual counts on a separate implied scale, direction only. The inset magnifies the 2020–24 counts on a suppressed-zero scale to make the trough-and-rise reversal legible; every value is labelled.

The 2019 outlier belongs in the record but not on that line: 287 fatalities that year included the 250 deaths at Brumadinho, a tailings-dam failure whose lesson this series takes up in its tailings-governance work rather than here. Strip that catastrophe out and the underlying pattern is the one Exhibit 1 shows: injuries falling steadily, fatalities bottoming around 2020–2022 and turning back up. A programme measured on injury frequency alone would have reported success straight through the reversal.

The geography nobody wants to name

The fatal risk is not spread evenly across the industry's footprint, and honesty requires saying where it sits. In 2024, operations in Africa accounted for 50 per cent of fatalities, 21 of 42, on 24 per cent of hours worked, a fatal-frequency rate of 0.031, the highest of any continent; South America ran 0.009 on 41 per cent of hours, North America 0.029, and the global rate was 0.015, while Europe recorded no fatalities for the fourth year running. Exhibit 2 sets the continents side by side on the only comparison that controls for exposure: fatalities per hour worked.

EXHIBIT 2

Africa carried half the deaths on a quarter of the hours

Fatal-frequency rate by continent, 2024 (fatalities per million hours worked); share of total fatalities and of hours worked annotated

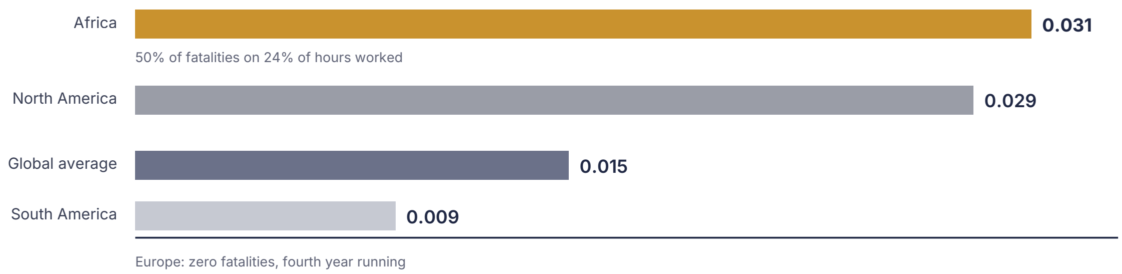

Source: ICMM, Safety Performance Report, 2025. Bars drawn to a common rate scale; the Africa bar is annotated with its share of fatalities and of hours worked, the two figures that make the rate meaningful. Lower-rate continents (Asia, Australia and Oceania) are omitted for legibility.

> FATAL-FREQUENCY RATE (FFR) · fatalities per million hours worked. Unlike a raw count, it controls for exposure, so a region with fewer hours and more deaths, as Africa was in 2024, shows the elevated risk a count alone would hide.

The concentration sharpens further at country level. In 2024 South Africa recorded 15 fatalities, a fatal-frequency rate of 0.038, with 14 of the 15 underground; Brazil and Canada recorded 5 each, Canada at 0.042; Australia 3; and among smaller producers Argentina's two fatalities gave it the highest country rate at 0.075, with Tanzania at 0.054 and the DRC at 0.028. All 18 of the year's underground fatalities occurred on just two continents: 15 in Africa (14 in South Africa, one in the DRC) and three in South America. Exhibit 3 tabulates the country picture as published.

EXHIBIT 3

The country table names where the fatal exposure sits

Reported fatalities and fatal-frequency rate by country, ICMM members, 2024 (all fourteen countries with a fatality; FFR per million hours worked)

COUNTRY	FATALITIES	FFR	NOTE
South Africa	15	0.038	14 of 15 underground
Brazil	5	0.009	
Canada	5	0.042	includes one aviation incident
Australia	3	0.010	
Argentina	2	0.075	highest country FFR
Tanzania	2	0.054	
DR Congo	2	0.028	one underground
United States	2	0.018	
Guinea	1	0.026	
Colombia	1	0.024	
Kazakhstan	1	0.017	

COUNTRY	FATALITIES	FFR	NOTE
Ghana	1	0.016	
Peru	1	0.007	underground
Chile	1	0.003	lowest published FFR

Source: ICMM, Safety Performance Report, 2025 (country table). Fatal-frequency rate is per million hours worked, published for every country with a 2024 fatality; the table reproduces all fourteen. A high count with a low rate reflects large hours worked, and the reverse a small workforce, which is why the rate, not the count, ranks the exposure. Counts are as reported by member companies. The Canada note cross-references the report's aviation category (four aviation fatalities in a single 2024 incident, in Canada), which, set against Canada's five-fatality country total, accounts for four of the five (share computed from the two published totals).

The five-year picture, not the one-year shock

A single year can mislead in either direction, so the honest frame is the multi-year one, and it does not soften the conclusion. Across 2020–2024, underground work averaged 48 per cent of fatalities, and underground plus other process areas ran near 90 per cent of the total; the fatal exposure has sat in the same settings, on the same hazards, for five years. What changed in 2022–2024 was not the location of the risk but the industry's grip on it: the share of members achieving zero fatalities slipped, the fatal-incident count rose, and three incidents each killed more than one person. Read across five years, 2024 is not an anomaly to be explained away; it is the visible surfacing of a control problem that the injury statistics had been masking.

The long decline that preceded the reversal is worth understanding, because it explains why the reversal is so disorienting. From 2012 the industry professionalised its safety management: standardised reporting, fatal-risk protocols, and heavy investment in the visible controls, training, procedures, personal protective equipment, that sit low in the hierarchy of control (Chapter 4). Those investments genuinely reduced injuries. But the low-hierarchy controls they favoured are precisely the ones most exposed to execution failure, and a decade of injury-rate success built a confidence the fatal data did not earn. The reversal is what it looks like when the easy gains are exhausted and the hard, high-hierarchy engineering, the controls that remove exposure rather than manage behaviour, has not yet been done at the scale the hazard demands.

Why this is an engineering problem, not a morality play

It would be easy to read these numbers as a call for more resolve, more posters, more "commitment to zero." That reading is the reason the curve reversed. The data does not describe a shortage of intent; it describes a shortage of effective controls at the point of work. The rest of this paper follows the causal chain the ICMM's own analysis draws: from the hazards that kill (Chapter 4), through the failure mode that lets them (Chapter 5: execution), to the management system built to close the gap (Chapters 3 and 7). The African concentration is not an argument for leaving the continent's mining to others; it is the clearest possible instruction on where the engineering standard has to be highest.

> **A safety programme optimised for injury rates can be improving and killing more people at the same time. That is not a paradox. It is a measurement error.**
>
> THE ARGUMENT OF CHAPTER 2

2

SHARE · FROM PYRAMIDS TO PRECURSORS

From pyramids to precursors

Heinrich's triangle promised that fewer cuts meant fewer deaths. The 2024 data is the counter-example, and the science of serious-injury precursors explains why.

5.07 → 2.29
TRIFR, 2012 TO 2024, INJURIES MORE THAN HALVED (ICMM, 2025)

6,446
RECORDABLE INJURIES, 2024, DOWN FROM 7,278 (ICMM, 2025)

42
FATALITIES THE SAME YEAR, THE CURVES MOVED APART (ICMM, 2025)

The intellectual foundation of a century of safety practice is a triangle. Reduce the base of minor incidents, the theory runs, and the fatal apex shrinks in proportion. It is intuitive, it is teachable, and on the evidence of 2024 it is, for fatal risk, wrong. The industry drove its injury base down for a decade and the fatal apex moved the other way. The reconciliation is not that safety science failed; it is that fatal events have different causes from the injuries beneath them.

The decoupling, stated plainly

Set the two published series against each other and the point is unarguable. Total recordable injury frequency fell from 5.07 to 2.29 across 2012–2024, and recordable injuries fell year-on-year into 2024; over the same window the fatality count bottomed and reversed, ending 2024 at 42. If minor injuries reliably predicted fatalities, these lines would be parallel. They are not. Exhibit 4 states the decoupling as an index, both series set to 100 at the start of the window so the divergence is read directly.

EXHIBIT 4

The safety triangle broke: injuries down, fatal risk not following

Indexed trajectories, ICMM members. Injury-frequency index vs fatality-count index, schematic between published anchors (start of series = 100)

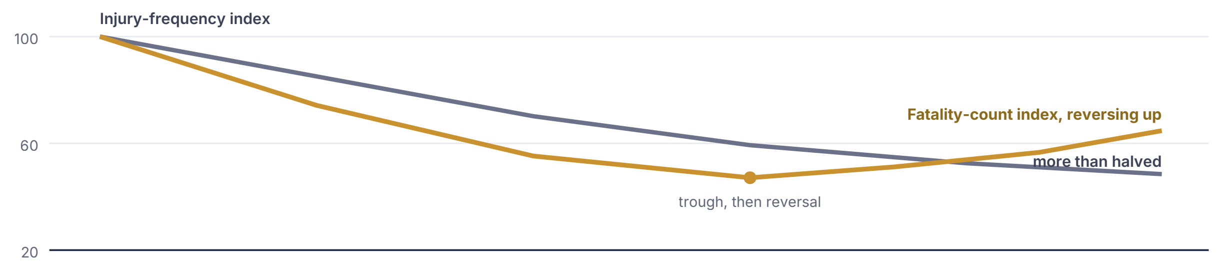

Source: ICMM, Safety Performance Report, 2025 (Table 1). Both series indexed to the start of the published window and drawn schematically between the report's anchor values; the exhibit shows direction and divergence, not point estimates. Injury index derived from the published TRIFR series (5.07→2.29); fatality index from the published annual counts.

> SIF (SERIOUS INJURY AND FATALITY) · the class of events that kill or permanently disable. SIF science holds that these events have their own precursors, distinct from the minor-injury base, so preventing them requires watching for the near-misses that *could* have been fatal, not just counting the injuries that were minor.

What the precursor science actually says

The modern reading, developed in the serious-injury-and-fatality literature, is that a minority of incidents carry fatal potential and that these have identifiable precursors: high-energy exposures, absent or failed critical controls, and specific operational settings. A slip on a walkway and a person on foot in the path of a haul truck are both "incidents," but only one has fatal energy behind it. Counting them together, the logic of the triangle, buries the signal that matters. The practical consequence is a shift in what a safety system watches: from the frequency of all injuries to the presence and effectiveness of the controls that stand between a worker and a high-energy event.

This is not an academic distinction. The ICMM's own 2024 analysis lands in exactly the same place from the data side: 83 per cent of fatalities involved a failure to implement effective critical controls, and over 2021–2024 two-thirds of control failures were failures of execution in the operating environment. Precursor science says fatal events turn on a small set of controls; the fatal record says those controls were present but not effective. The two describe the same problem from opposite ends, and they meet at the management system Chapter 3 sets out.

What a near-miss is worth

The practical yield of the precursor view is a change in what an organisation treats as information. Under the triangle, a near-miss is a minor event to be logged and closed. Under the precursor view, a near-miss with fatal potential, a haul truck that stopped two metres short, a rock that fell into an area a crew had just left, an isolation that was found incomplete before energisation, is the most valuable safety signal an operation generates, because it is a fatality that did not happen for a reason the organisation can still learn. The 2024 data makes the point by omission: with only one fatality formally classified under "safety culture" despite failure-to-follow-procedures being the leading organisational factor, the industry is plainly not yet reading its own precursors well. A system that harvested high-potential near-misses as diligently as it counts injuries would see the fatal risk moving while there was still time to act on it, which is exactly what the leading-indicator dashboard of Chapter 7 is built to do.

> READING THE SCIENCE
>
> **Why Heinrich's ratio does not hold for fatal risk**
>
> The 1930s triangle attributed a fixed ratio between minor incidents, serious injuries and fatalities, implying that suppressing the base suppresses the apex. Two things break the ratio for modern mining. First, **energy**: the events that kill involve high-energy exposures, heavy machinery, unsupported ground, stored energy, whose potential is categorically different from the low-energy events that dominate the injury base, so reducing one does little to the other. Second, **controls**: minor injuries are mostly governed by low-hierarchy controls (attention, footwear, housekeeping), while fatal events turn on high-hierarchy critical controls (elimination, engineering, isolation), and an organisation can excel at the first set while the second quietly degrades. The 2024 decoupling is the ratio breaking in public. The response is not to abandon injury reduction, which is worthwhile in itself, but to stop treating it as a proxy for fatal-risk control, and to measure the fatal risk directly.

EXHIBIT 5 · FRAMEWORK

From counting incidents to controlling energy

The conceptual shift the data forces: a capability instrument, no member data

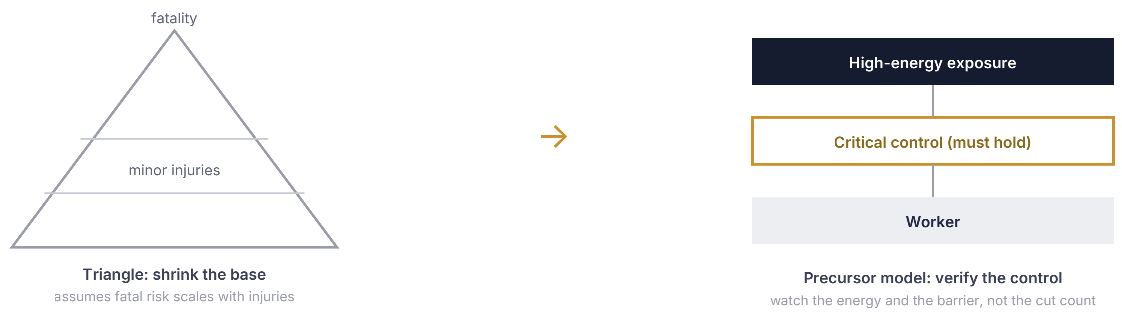

A capability instrument, not an exhibit of member data. It restates the serious-injury-and-fatality reframing that the ICMM's 2024 control-failure findings (Exhibits 1, 4 and Chapter 5) support empirically.

What to measure instead

If injury frequency is the wrong steering signal for fatal risk, the obvious question is what the right one is, and the precursor view answers it directly: measure the state of the critical controls, not the tally of outcomes. Four measures do most of the work. **Coverage**, what fraction of the critical controls were verified this period, tells you whether the system is actually looking. **Effectiveness**, what fraction of those verifications the control passed, tells you whether what it found was sound. **Restoration**, how long a control found degraded stayed degraded, tells you whether the organisation responds or merely records. And **high-potential near-misses**, events with fatal potential that did not cause harm, tell you where the next fatality is being rehearsed. None of these is a body count; all of them move before anyone is hurt; and together they are the leading-indicator set that Chapter 7 builds into a dashboard. The shift they represent is from asking “how many people were harmed?” to asking “how many of the barriers that stand between our people and death are we certain are working?”, a question the 2024 data suggests very few operations could currently answer.

The chapter’s conclusion is uncomfortable for any organisation that reports safety on a single rate: the number that has been improving is not the number that has been killing people. A board that reads only the injury frequency is being told the truth about the wrong risk. Chapter 3 turns to the framework that measures the right one.

> Fatal events are not big injuries. They are a different kind of event, with a different set of causes, and they demand a different kind of control.

3

SHARE · CRITICAL CONTROL MANAGEMENT

Critical Control Management

In 2026 the industry's good-practice guide was rewritten for the first time in a decade. It is the answer to the 2024 data, and its logic is a bowtie.

2026
FIRST COMPREHENSIVE CCM GUIDE UPDATE IN A DECADE (ICMM, 2026)

9 steps
ACROSS THREE PHASES: PLANNING, DEVELOPMENT, IMPLEMENTATION (ICMM, 2026)

83%
OF 2024 FATALITIES: A CRITICAL CONTROL WAS NOT EFFECTIVELY IMPLEMENTED (ICMM, 2025)

Critical Control Management is the discipline of identifying the few controls that stand between a workforce and a fatal event, defining what each must do, and verifying, continuously, at the point of work, that it is doing it. In 2026 the ICMM comprehensively updated its CCM Good Practice Guide, its first revision in a decade, merging the 2015 good-practice and implementation guides into a single document. The timing is not incidental: the framework was rewritten to attack the failure mode the 2024 data had just exposed.

The bowtie: the shape of a controlled hazard

The organising image of critical-control practice is the bowtie. At the centre sits the unwanted event, the release of energy that harms. To the left are the threats that could cause it and the preventive controls that stop each threat from reaching the event; to the right are the consequences and the mitigating controls that limit the harm if the event occurs. Critical controls are the subset, on either side, whose failure would materially raise the risk of a fatality regardless of what else is in place. Exhibit 6, the signature exhibit of this paper, draws the bowtie for a representative major-hazard event and marks which barriers are the critical ones.

EXHIBIT 6 · SIGNATURE EXHIBIT · FRAMEWORK

The critical-control bowtie: preventive barriers left, mitigating barriers right, the fatal event at the knot

A representative major-hazard event drawn to the CCM anatomy. A capability instrument, no member data, no project data

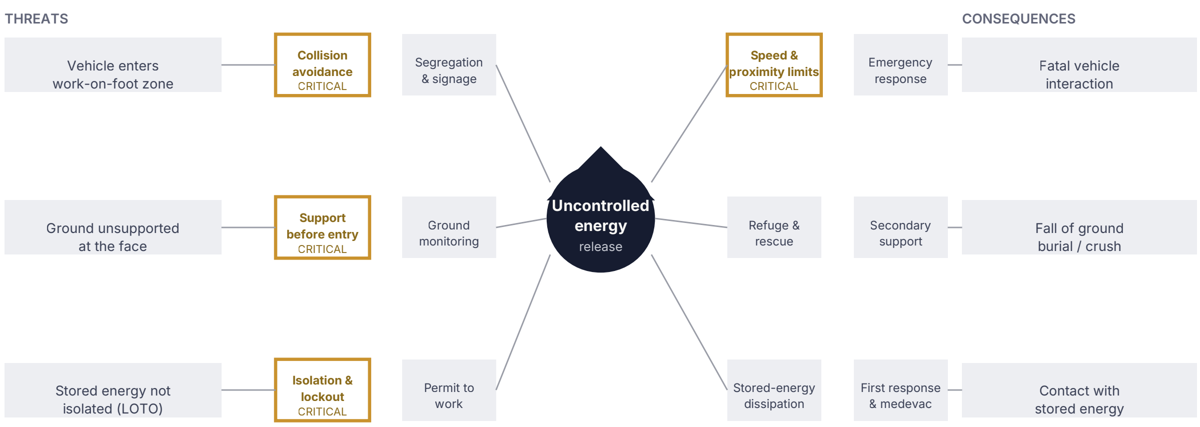

Gold-outlined barriers are the CRITICAL controls: their failure raises fatal risk regardless of the other barriers present.

A capability instrument, drawn to the CCM anatomy in ICMM, Critical Control Management: Good Practice Guide, 2026. It carries no member or project data; the threats, consequences and controls shown are representative of the mining major-hazard set (mobile-equipment interaction, fall of ground, uncontrolled stored energy) that Chapter 4 shows the 2024 fatalities concentrate in.

> BOWTIE · a risk model in which threats and preventive controls sit to the left of an unwanted event and consequences and mitigating controls to the right. It makes explicit which barriers are load-bearing, the critical controls, and where a single failure defeats the whole diagram.

Nine steps, three phases

The 2026 guide sets out a nine-step process split between three phases, with a feedback and improvement loop running across the steps rather than closing them off. **Planning** is a single step: the organisation plans the process and its scope (step 1). **Development** is the analytical core (steps 2–6): it identifies the fatal hazards and unwanted events to be managed, identifies the controls for each, selects which are critical, defines each critical control's performance and how it is verified, and assigns accountability by name. **Implementation** then carries the controls into the operation (steps 7–9): it implements them to local conditions, verifies them in the operating environment, and evaluates and improves them, reopening the earlier steps as the evidence requires. The 2026 update adds maturity assessments and readiness checks, and, the structurally important change, positions critical-control management inside enterprise risk management rather than beside it. Exhibit 7 sets the nine steps out as the loop they form.

EXHIBIT 7

The 2026 CCM process: nine steps that close into a feedback and improvement loop

ICMM Critical Control Management Good Practice Guide, 2026, nine-step process across three phases: Planning, Development, Implementation

Source: ICMM, Critical Control Management: Good Practice Guide, 2026, Table 2 and Figure 2 (nine-step process; three phases, Planning, Development and Implementation; feedback and improvement loop across the steps); full guide archived to this paper's evidence dossier, and cross-confirmed by trade summaries, Mining Weekly, 29 April 2026, and im-mining, 10 July 2025. Step 4, the selection of the critical controls, is highlighted as the hinge of the process; steps 8 and 9, verification and improvement, carry the loop the fatal record demands.

The guide's deep move is step 8. A control that is written, resourced and trained but not verified in the operating environment is, on the 2024 evidence, exactly the kind of control that fails: recall that 67 per cent of control failures over four years were failures of execution, and 83 per cent of fatalities involved a critical control not effectively implemented. Verification is the step that converts a control from a document into a barrier. Chapter 5 is about why that step is the hardest one, and Chapter 7 builds it into an operating system.

One event through the nine steps

The process is easiest to grasp applied to a single material unwanted event. Take the fatal vehicle interaction of Exhibit 6, the leading 2024 hazard. In development, the site names the event (step 2), enumerates its controls, segregation, collision-avoidance, speed limits, traffic-management plans, high-visibility clothing (step 3), and selects as critical the two that carry the fatal risk independently of the others: positive segregation and a collision-avoidance system with automatic intervention (step 4); it then specifies what "effective" means for each, segregation verified before shift, collision-avoidance tested and its intervention function confirmed (step 5), and assigns an accountable owner for each control, by name, not by committee (step 6). In implementation, it implements to the site's actual haul-road geometry rather than a generic template (step 7), verifies both controls in the field on a defined cadence and reports the results as leading indicators (step 8), and when a verification fails, a segregation barrier down, an intervention function disabled, evaluates the failure and improves the control with a defined action and a restore clock, not a note in a log (step 9). Exhibit 6's bowtie is the picture; the nine steps are the verb.

THE 2026 ADDITIONS

Maturity assessment: knowing whether the system is real

The 2026 guide adds two instruments the 2015 versions lacked: maturity assessments and readiness checks. Both exist to answer a question the 2024 data makes unavoidable, is the critical-control system actually operating, or does it exist only on paper? A maturity assessment scores each element of the nine-step process against defined levels, from "documented but not verified" through to "verified, reported and driving response." A readiness check tests whether a specific critical control is fit for the conditions it will meet before work begins, rather than after an incident. Together they turn CCM from a static framework into a self-auditing one, and they are the reason the 2026 update matters more than a routine refresh: they build the verification discipline the fatal record demands into the framework's own machinery. Positioning CCM inside enterprise risk management completes the move, a control found immature or unready is now visible to the board as a live business risk, not filed under health and safety.

READING THE FRAMEWORK

Four tests for whether a control is really "critical"

Not every important control is a critical control, and treating all of them as critical dilutes the verification effort that the label is meant to concentrate. The 2026 guide's selection step (step 4) turns on four questions a site can ask of any candidate. **Does its failure credibly lead to a fatality?** If the worst realistic outcome is a lost-time injury, it is important but not critical. **Is it the barrier that carries the risk?** A control that only matters when three others have already failed is defence-in-depth, not a critical control. **Can its performance be specified and verified?** A control whose effectiveness cannot be checked cannot be managed as critical. **Is it independent of the threats it guards against?** A control defeated by the same event it is meant to stop is not a barrier. The discipline is subtractive: name the few that pass all four, and verify those relentlessly.

The framework, in short, is not new doctrine; critical-control thinking has been in the industry's vocabulary for a decade. What the 2026 update does is sharpen the doctrine at exactly the joint where the 2024 data says it broke, verification and response, and lift it out of the safety silo into enterprise risk, where a failing control can no longer be quietly absorbed. The chapters that follow test the framework against the specific hazards it must handle (Chapter 4) and the specific reason it fails (Chapter 5), before Chapter 7 rebuilds it as an operating system a site can actually run.

> **The 2026 guide did not invent critical-control thinking. It moved the discipline to the one joint where the 2024 data says it broke: proving, at the face, that the control still works.**
>
> WHAT THE UPDATE ACTUALLY CHANGED

A critical control is not the one that matters most on paper. It is the one whose failure, on its own, kills, and there are fewer of those than any risk register admits.

THE SELECTION DISCIPLINE OF CHAPTER 3

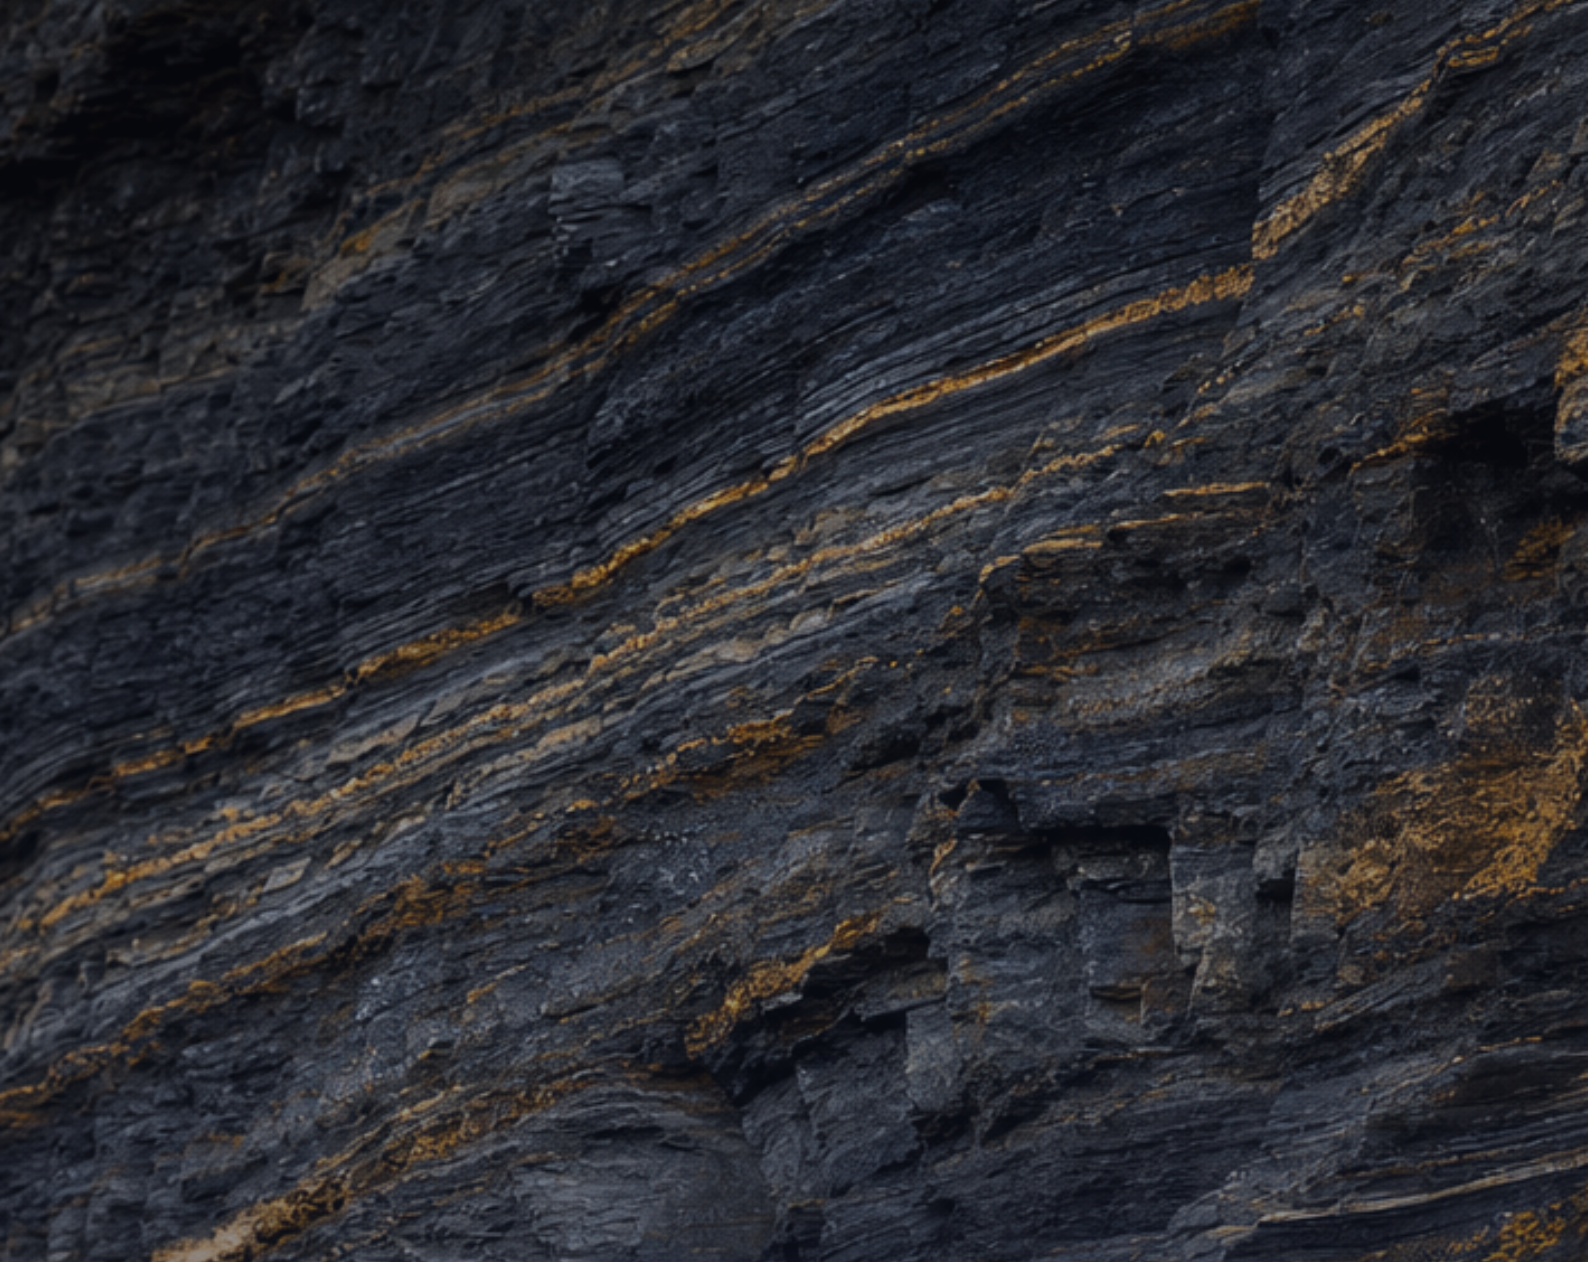

4

SHARE · THE BIG KILLERS

The big killers, engineered out

A short list of hazards carries most of the deaths. Each is a design problem before it is a behaviour problem, and the hierarchy of control says so.

21%
OF 2024 FATALITIES FROM MOBILE EQUIPMENT: LEADING HAZARD (ICMM, 2025)

5
FALL-OF-GROUND FATALITIES, FOUR IN SOUTH AFRICAN OPERATIONS (ICMM, 2025)

≈2/3
OF UNDERGROUND FATALITIES: MOBILE EQUIPMENT + FALL OF GROUND (ICMM, 2025)

The fatal record has a short, stubborn list of causes, and the shortness is the good news: a hazard that recurs is a hazard that can be designed against. In 2024, mobile equipment was the leading cause of fatalities, at 21 per cent, nine deaths, seven of them at operations in Africa and five underground. Fall of ground was second, at five fatalities, four in South African operations. Falling objects caused four, up from zero the year before, each a lapse in secondary retention or no-go-zone discipline; and a single aviation incident in Canada caused four more. Exhibit 8 sets the 2024 hazard profile out.

EXHIBIT 8

A handful of hazards carry the fatal risk

Leading causes of fatality, ICMM members, 2024, by count (and share of the 42 total)

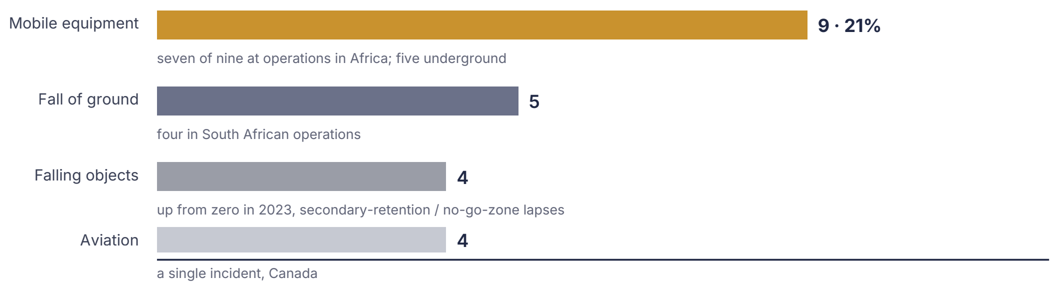

Source: ICMM, Safety Performance Report, 2025. Counts as reported; shares computed on the 42-fatality total and labelled as such. The four hazard classes shown accounted for the majority of 2024 fatalities; the remainder are distributed across a longer tail of causes.

> MOBILE EQUIPMENT INTERACTION · a collision or crush event between a person or light vehicle and heavy mobile machinery; the leading fatal hazard of 2024, and an interaction to be engineered out rather than supervised away. (Ch. 4)

Mobile equipment: the interaction is the hazard

Mobile-equipment fatalities are, overwhelmingly, interaction events: a person or a light vehicle in the path of a heavy machine whose operator cannot see them, or a machine on ground that fails beneath it. The recurrence is why the industry's vehicle-interaction reference framework, the Earth Moving Equipment Safety Round Table's control set, is built around a hierarchy that puts eliminating the interaction above controlling the operator's behaviour. That ordering is the whole argument of this chapter. The most reliable control is the one that removes the exposure, separating people from machines by design, and machines from each other by proximity-detection and intervention systems that act without waiting for a human decision. The least reliable is the one that asks a tired operator, at the end of a long shift, to see what the cab's blind spots hide.

EXHIBIT 9 · FRAMEWORK

The hierarchy of control ranks reliability, and mining's spend often inverts it

Hierarchy of control applied to the mobile-equipment interaction. A capability instrument; fatal-cause context is Exhibit 8

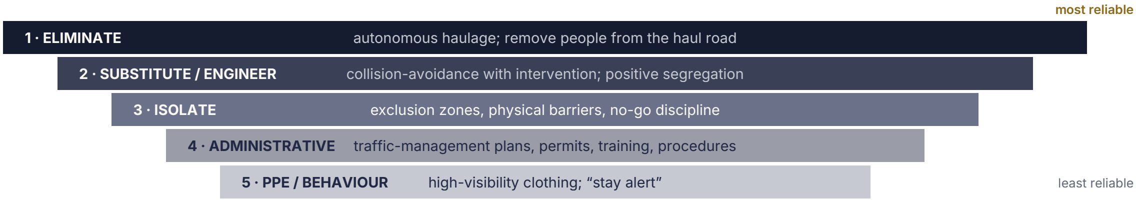

Reliability falls from top to bottom; controls near the base depend on human performance, which fails under fatigue, pressure and darkness.

A capability instrument. The hierarchy of control is standard occupational-safety doctrine; the mobile-equipment application follows the EMESRT vehicle-interaction approach. No member or project data; the fatal-cause context is Exhibit 8 (ICMM, 2025). The EMESRT source pull is pre-registered in this paper's evidence dossier.

> HIERARCHY OF CONTROL · the ranking of control types by reliability, from elimination (most reliable, removes the hazard) down through engineering and isolation to administrative controls and personal protective equipment (least reliable, depends on human performance). Critical controls are chosen high in the hierarchy wherever the fatal energy allows.

Fall of ground: support before entry, or not at all

Fall-of-ground fatalities have an even sharper signature. Of the five in 2024, three involved work carried out before full mesh-and-bolt support had been installed or repaired, and two were seismic bursts in deep, high-stress stopes soon after blasting, a pattern that echoes the Minerals Council South Africa's own fall-of-ground root-cause findings. The first three are a sequencing failure: entry permitted into ground the support system had not yet caught up with. The second two are a design-and-scheduling problem: re-entry timing after blasting in seismically active rock. Neither is primarily a matter of worker carelessness; both are matters of what the system permits. Exhibit 10 sets the triggers against their controls.

EXHIBIT 10

Fall-of-ground deaths cluster on two failures the system controls

Triggers of the five 2024 fall-of-ground fatalities and the critical control each defeats

CASES	TRIGGER, AS REPORTED	CRITICAL CONTROL DEFEATED
3	Work before full mesh-and-bolt support installed or repaired	Support-before-entry: no access until ground is supported to standard
2	Seismic bursts in deep, high-stress stopes soon after blasting	Re-entry timing and seismic exclusion after blasting

Source: ICMM, Safety Performance Report, 2025, citing Minerals Council South Africa Fall of Ground Action Plan, 2021 (root-cause echo). Cited via the ICMM report; the Minerals Council finding is not independently re-attributed here. Both triggers defeat a support/sequencing control rather than a personal one.

Falling objects and aviation: the tail that still kills

Two further 2024 categories deserve naming because each carries a specific lesson. Falling objects caused four fatalities, up from zero in 2023, and the report attributes each to a lapse in secondary retention or no-go-zone discipline: an object not tethered against a second line of defence, or a person permitted beneath a suspended load. Both are critical-control failures of the most preventable kind, the control (secondary retention; exclusion below suspended loads) is unambiguous, cheap and well understood, and its failure is pure execution. Aviation caused four fatalities in a single incident in Canada, a reminder that a mine's fatal exposure extends to the fixed-wing and rotary transport that serves remote operations, and that the same critical-control logic (is the operator's safety-management system verified, not merely contracted?) applies to aviation providers exactly as it applies to haul-truck fleets.

Engineering the exposure out

The unifying move across every hazard on the list is to climb the hierarchy of control: to prefer the control that removes the exposure over the one that manages the person inside it. For mobile equipment that means, at the top of the hierarchy, autonomous or remotely operated haulage that takes people out of the cab and off the haul road, and, one step down, collision-avoidance systems that intervene automatically rather than alerting an operator who may already be committed. For fall of ground it means support-before-entry designed into the mining sequence so that access is physically gated on the ground being supported, and seismic-exclusion timing built into the blast-and-re-entry cycle rather than left to judgement. For falling objects it means secondary retention specified as a design requirement and exclusion zones enforced by physical barrier rather than signage. None of this is speculative technology; all of it is deployed somewhere in the industry today. The gap the 2024 data measures is not one of invention but of adoption, and, as Chapter 5 shows, of verifying that what has been adopted is actually working where the work is done.

The African concentration, read as an engineering brief

Chapter 1 named the geography; Chapter 4 is where it becomes an instruction. Seven of the nine mobile-equipment fatalities were at operations in Africa, and 14 of the 15 South African fatalities were underground, with all but three of the year's underground deaths on the continent. The pattern is not evidence that African mining is inherently more dangerous; it is evidence that the continent's operations concentrate the two hazards, deep, seismically active underground mining and dense mobile-equipment interaction, that carry the highest fatal energy. Read as an engineering brief, that concentration is precise and actionable: it says that the highest-return safety engineering in the industry today is the deployment of support-before-entry discipline, seismic-exclusion timing and collision-avoidance-with-intervention at exactly these operations. The moral of the geography is not to withdraw from African mining, which would export the risk to operators less equipped to manage it, but to hold the engineering standard highest where the hazard is highest.

The underground picture ties the two killers together. Some 43 per cent of 2024 incidents were underground, and across 2020–2024 underground work averaged 48 per cent of fatalities; within the underground total, mobile-equipment interactions (around a third) and fall of ground (around a third) together make up roughly two-thirds, and underground plus other process areas accounted for 88 per cent of 2024 fatalities, near 90 per cent across the five-year window. The message for a developer choosing between an open-pit and an underground method, or for a lender pricing either, is that the underground fatal exposure is concentrated in two hazards that mature engineering controls already address, when those controls are verified at the face. Which is the subject of Chapter 5.

> **Every hazard on the fatal list is a design problem before it is a behaviour problem. The question is never only "was the worker careful?" but "did the system let the worker be exposed at all?"**
>
> THE ENGINEERING PREMISE

The control existed. It was written, resourced and trained. On the day, at the face, it was not effective. That sentence describes most of the deaths.

THE FINDING THAT REORGANISES THE SAFETY PROBLEM

5

SHARE · FIELD EXECUTION

Field execution

Most controls that failed in 2024 were not missing or badly designed. They failed in the operating environment. Closing that gap is a field problem, and much of it is a contractor problem.

67%	45%	>25%
OF CONTROL FAILURES, 2021–2024: INEFFECTIVE EXECUTION (ICMM, 2025)	OF 2024 FATALITIES WERE CONTRACTORS, ON 58% OF HOURS (ICMM, 2025)	OF THE 37 FATAL INCIDENTS: FAILURE TO FOLLOW RULES/ PROCEDURES, THE TOP FACTOR (ICMM, 2025)

Here is the finding that should reorganise every mining safety budget. Across 2021–2024, 67 per cent of control failures were failures of execution, the control was ineffective in the operating environment, rather than failures of design or the control being absent altogether. In 2024 specifically, 83 per cent of fatalities involved a failure to implement effective critical controls. The controls were largely there. They did not work where the work was done. Exhibit 11 sets out the control-failure typology.

EXHIBIT 11

The failure is execution, not design or absence

Control-failure typology, 2021–2024 (the report's Graph 10 is a four-year stacked series, not a single year). The 67% execution share is text-verbatim for 2021–2024; the per-category counts are read from that graph and flagged as such

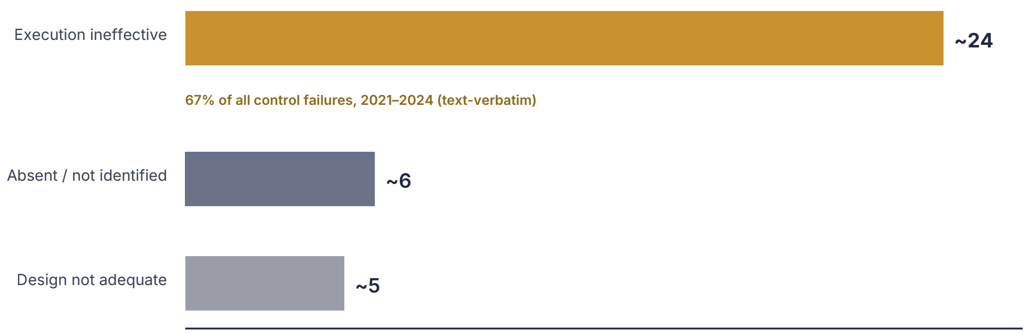

Source: ICMM, Safety Performance Report, 2025 (Graph 10). The headline 67% execution-ineffective share (2021–2024) and the 83% of 2024 fatalities involving a failure to implement effective critical controls are text-verbatim from the report. The per-category counts (~24 / ~6 / ~5) are read from the graph's bars and are approximate; this paper flags them as chart-read and does not treat them as exact.

> EXECUTION FAILURE · a control that exists on paper and in training but is not effective at the point and moment of work, not applied, not applied correctly, degraded, bypassed under production pressure, or defeated by conditions. On the 2024 evidence this is the dominant way critical controls fail.

Why good controls fail at the face

Execution failure has organisational roots, and the ICMM's 2024 analysis names them. Using an incident-cause-analysis method, the leading organisational factor in 2024 was failure to follow rules or procedures, implicated in roughly ten fatal incidents, more than a quarter of the 37, followed by risk management (around six) and management of change (around four); tellingly, only one fatality was formally classified under "safety culture," which the report reads as cultural drivers being under-recognised in investigations rather than absent. A related and specific warning: energy-isolation (lockout/tagout) breaches rose in 2024 and were flagged as a reinforcement priority. Exhibit 12 sets the organisational factors out.

> LOCKOUT / TAGOUT (LOTO) · the energy-isolation discipline that renders stored energy, electrical, hydraulic, mechanical, gravitational, safe before work begins and proves the isolation before anyone re-energises; breaches rose in 2024 and were flagged as a reinforcement priority.

EXHIBIT 12

The organisational factors point at procedures, risk work and change

Organisational factors in 2024 fatal incidents (Graph 8). Counts approximate, read from the report

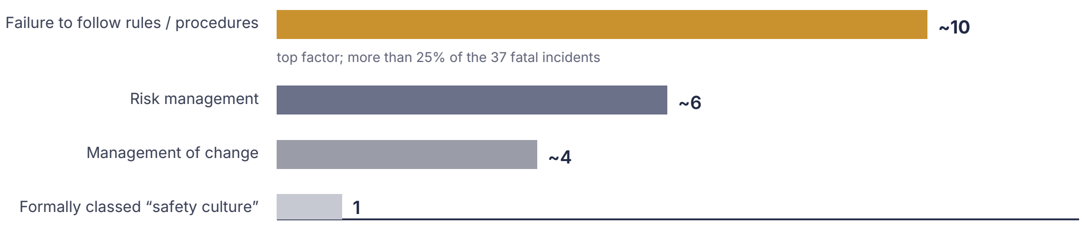

Source: ICMM, Safety Performance Report, 2025 (Graph 8, incident-cause-analysis-based). Counts are read from the graph and are approximate. The report itself reads the single "safety culture" classification as under-recognition of cultural drivers in investigation. Energy-isolation (LOTO) breaches rose in 2024 and were separately flagged as a reinforcement priority.

The under-recorded factor: culture read as compliance

One number in the organisational data deserves to be dwelt on, because it is a measurement artefact masquerading as a finding. In 2024, failure to follow rules or procedures was implicated in roughly ten fatal incidents, the top factor, while only one fatality was formally classified under "safety culture". Read naively, that says the industry has a compliance problem, not a culture problem. Read carefully, it says the opposite. When investigators default to "the worker did not follow the procedure," they record the last link in the chain and stop; the questions that would surface culture, why was the procedure not followed, was it followable under the production pressure of that shift, had it been quietly bypassed for months with tacit approval, did the reward system value output over the pause, go unasked. The single "safety culture" classification is not evidence that culture rarely kills; it is evidence that investigations rarely look for it. The practical consequence for an operating system is that verification must capture both whether a control was in place and whether the conditions of work allowed it to be used as designed, a distinction Chapter 7's dashboard is built to hold.

The same reading applies to the year's specific warning on energy isolation: lockout/tagout breaches rose in 2024 and were flagged as a reinforcement priority. An isolation breach is almost never ignorance of the rule; it is a rule defeated by time pressure, by an isolation that was awkward to apply, or by a normalised shortcut that had not yet been caught. Treating it as a training gap invites more training; treating it as an execution-and-culture signal invites the field verification and the honest question about working conditions that actually close it.

The contractor question

No account of field execution is complete without the contractor picture, because in modern mining the contractor is much of the field. In 2024, contractors accounted for 45 per cent of fatalities, 19 deaths, down sharply from 69 per cent, or 25 deaths, in 2023, while working 58 per cent of hours; and the contractor fatal-frequency rate, at 0.012, was below the employee rate of 0.020. That last figure deserves emphasis because it contradicts a common assumption: on the 2024 exposure-adjusted numbers, contractors were not more dangerous per hour than employees. Exhibit 13 sets the comparison out.

EXHIBIT 13

Contractors worked most of the hours, at a lower fatal rate than employees

Contractor vs employee, ICMM members, 2024: share of fatalities, share of hours worked, and fatal-frequency rate

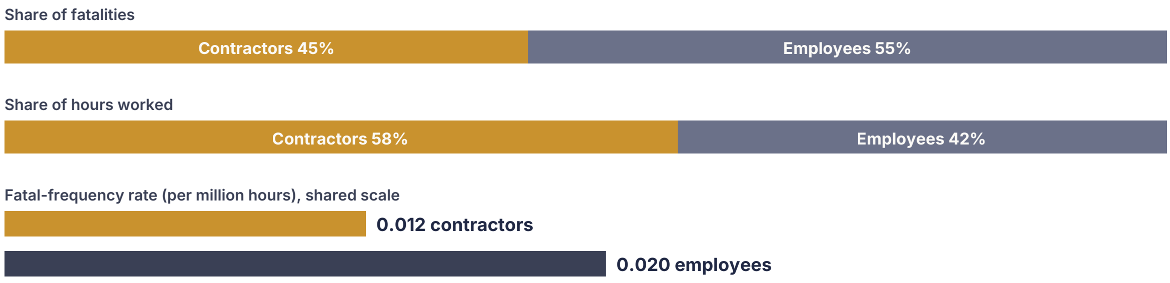

Contractors carried more hours at a lower fatal rate; the 2023→2024 fall in the contractor share (69% to 45%) is the year's largest single shift.

Source: ICMM, Safety Performance Report, 2025. The contractor rate (0.012) below the employee rate (0.020) is as published; it does not license complacency, since contractors still carried the majority of hours and nearly half the deaths. Rates and shares are exposure-adjusted where the report provides hours.

What verification actually looks like at the face

If execution is the failure mode, verification is the fix, and it has to be concrete enough to schedule. Verifying a critical control in the operating environment is not an audit of paperwork; it is a physical check, at the point of work, that the control is present and doing its job: the segregation barrier is up and unbreached, the collision-avoidance system's intervention function is live and was not disabled to keep production moving, the ground is supported to standard before anyone is under it, the isolation is complete and proven before energisation. Each check has an owner, a cadence and a pass/fail standard, and each failure has a defined response and a clock. This is deliberately unglamorous work, and it is the work the 2024 data says was missing: the controls existed; the verification that they were effective, that day, did not. A useful discipline is to separate what is verified from how confident the verification is, a self-check by the crew doing the work is weaker evidence than an independent field verification by someone accountable for the control, and a critical control should not rest on the weakest form.

The year-on-year swing in the contractor share, from 69 per cent of fatalities in 2023 to 45 per cent in 2024, is the single largest movement in the 2024 dataset, and it should be read with care rather than celebration. A fall of that size in one year is unlikely to be a step-change in contractor safety; it is more likely the noise of a small denominator, the shifting mix of who did the most hazardous work in a given year, and the year's particular incidents. The durable finding is not the swing but the structure beneath it: contractors carried 58 per cent of hours, so the majority of the industry's fatal exposure sits with people who are not the operator's direct employees, managed under a contract that may or may not carry the operator's critical-control regime across the boundary. That is the structural risk, and it does not move with the annual share.

The contractor finding sets the design brief for any HSE operating system: it must extend the same critical-control regime, unchanged, across the organisational boundary between employer and contractor, because the hours, and therefore the exposure, sit mostly on the contractor side. That means shared control definitions, shared verification, shared reporting and shared accountability, written into the contract rather than bolted on after mobilisation. The discipline that makes this possible is not exhortation; it is the same field-protocol rigour that a well-run technical mandate already imposes on data. Aurus frames the field protocols structured into our mandates on exactly that principle: a control, like a datum, counts only when its integrity can be verified at the point it is captured.

AURUS PRACTICE NOTE

Data discipline is control discipline

The habit that makes a critical-control system work, specify the standard, verify at the point of work, reject anything that fails the check, is the same habit that makes a field data campaign trustworthy. On its delivered exploration-and-geophysics campaign, the firm's consultants worked to CRIRSCO-family Competent-Person supervision with 100 per cent data-quality acceptance and a verification discipline over every station reading and record. That is a capability adjacency, not a safety-performance claim: this paper reports no delivered corporate-HSE mandate, and the critical-control operating system of Chapter 7 is argued from the ICMM's published good practice and from that data-verification discipline, never from safety data the firm does not hold.

The execution gap, stated as a design requirement

Read together, the 2021–2024 control-failure typology and the 2024 contractor picture set a single design requirement rather than a list of interventions. The controls that fail are, in the main, controls that exist: written, resourced, trained, and then not effective at the point and moment of work. The people most exposed to that failure are as likely to be contractors as direct employees, because contractors carry the majority of the hours, 58 per cent of them in 2024. An operating system that answers this evidence therefore has to do two things at once, and do them in the field rather than on paper: verify that each critical control is effective where the work actually happens, and extend that same verification, unchanged, across the contract boundary. Everything Chapter 7 proposes is an attempt to make those two obligations schedulable, ownable and visible to a board, so that the execution gap the 2024 data exposes is closed by routine rather than rediscovered by the next investigation.

> The safest control is the one you can prove is working right now, at the face, on the contractor's shift as much as your own.

6

SHARE · HEALTH BEYOND SAFETY

Health beyond safety

A fatality count measures the deaths that happen on a shift. It does not measure the exposures that kill years later. The industry has begun to say so.

Dec 2024
ICMM MEMBER COMMITMENT ON UNDERGROUND DIESEL PARTICULATE MATTER (ICMM, 2025)

2012
COMMON SAFETY-INDICATOR SET IN PLACE SINCE; HEALTH METRICS STILL DEVELOPING (ICMM, 2025)

88%
OF 2024 FATALITIES UNDERGROUND + PROCESS AREAS, THE SAME SETTINGS CARRY THE HEALTH LOAD (ICMM, 2025)

Everything in the preceding chapters concerns acute safety: the events that harm in an instant. But the same underground and process settings that carry the fatal risk also carry a slower burden, respirable dust, diesel exhaust, noise, heat, chemical exposure, whose consequences are measured in careers, not shifts, and which a fatality count cannot see. This chapter is deliberately shorter and more cautious than the others, because the public data is thinner: the ICMM's common indicator set has run since 2012 for safety, but industry-wide occupational-health metrics are still under development, and no member-wide health dataset yet exists on the safety benchmark's footing. Where the numbers are not there, this paper says so rather than inventing them.

Diesel particulate matter: the one hard commitment

The clearest recent signal is on diesel particulate matter. In December 2024, ICMM members committed to further measures on underground diesel-particulate-matter exposure, accompanied by a dedicated DPM good-practice guide. DPM is a classified carcinogen: the International Agency for Research on Cancer classifies diesel-engine exhaust as a Group 1 human carcinogen (IARC Monographs, Volume 105, 2012). Underground, where ventilation is finite and diesel plant is dense, it is one of the defining chronic exposures of the industry. The commitment matters both for its content and for its form: it is the health equivalent of a critical control, a named exposure, a defined standard, and a mechanism to verify performance against it. The same logic the safety chapters advance applies here, one exposure at a time.

EXHIBIT 14 · FRAMEWORK

The exposures a fatality count cannot see

Chronic occupational-health exposure classes in the settings that carry the acute fatal risk. Qualitative; the one dated public anchor is the December 2024 DPM commitment

Below the waterline: chronic exposures whose harm is measured in careers, not shifts, and for which industry-wide metrics are still developing.

Gold outline marks the one exposure carrying a dated public ICMM commitment; the others are named qualitatively, not quantified here.

Source: ICMM, Safety Performance Report, 2025 (DPM commitment, December 2024; occupational-health metrics under development). A qualitative capability instrument: only the DPM item carries a dated public anchor; the other exposure classes are named from standard occupational-hygiene practice and are not quantified in this paper, which holds no member-wide health dataset.

> DIESEL PARTICULATE MATTER (DPM) · the fine soot and associated compounds in diesel-engine exhaust, a classified carcinogen and a defining chronic exposure of underground mining, where ventilation is finite and diesel plant dense. The subject of a December 2024 ICMM member commitment and a dedicated good-practice guide.

Why the health data lags, and why that is not an excuse

There are good reasons the occupational-health picture is thinner than the safety one, and none of them is a reason to wait. Chronic exposures manifest over years or decades, often after a worker has left the industry, so the outcome data is inherently lagged and hard to attribute to a single employer or exposure. Measurement is technically harder than counting incidents: it requires exposure monitoring, hygiene sampling and longitudinal health surveillance rather than an incident report. And the common indicator set that standardised safety reporting from 2012 has no fully developed health equivalent yet. All of that is real. But the critical-control logic does not depend on outcome data; it depends on control data, which is available now. An operation does not need a member-wide silicosis dataset to know whether its dust-suppression control is specified, owned and verified this month. Waiting for the epidemiology is waiting for the wrong evidence; the actionable evidence (is the control working?) is the same evidence the safety chapters demand, and it is collectable today.

Why health belongs in the same operating system

The temptation is to treat occupational health as a separate discipline with its own department, its own reporting line and its own annual report. The critical-control logic argues against that separation. A chronic exposure is a hazard with a critical control, ventilation design and monitoring for DPM, dust suppression and respiratory protection for silica, engineering and enclosure for noise, and it fails in the same way acute controls fail: in execution, in the operating environment, on the shift. Folding health exposures into the same critical-control register, with the same verification and the same dashboard, is the structurally honest move, and it is the one Chapter 7's operating system is built to make. Until the industry's health metrics mature to the safety benchmark's standard, the discipline has to be carried by method rather than by comparison.

The measurement gap is itself a governance risk

There is a governance point hiding in the thinness of the health data. A board that receives a mature safety dashboard and a sparse health picture may reasonably conclude that health is under control, when in fact it is merely unmeasured. The absence of a member-wide occupational-health dataset is not evidence of low exposure; it is evidence of low visibility, and the two are easy to confuse at board level precisely because the safety numbers are so much better developed. The disciplined response is to treat the measurement gap as a named risk in its own right, to ask, for each material chronic exposure, whether the operation can state the control, the standard and the verification, and to record "not yet measured" as a finding rather than a silence. The December 2024 diesel-particulate-matter commitment is valuable partly because it converts one exposure from unmeasured to managed; the task is to extend that conversion, one exposure at a time, without waiting for an industry benchmark that does not yet exist.

That is the through-line from Chapter 1 to here. Whether the hazard is acute or chronic, the failure mode is the same, a control that is not verified as effective where the work is done, and so the answer is the same: name the control, specify the standard, verify at the point of exposure, and report what you find before someone is harmed. Chapter 7 assembles that answer into a single operating system.

> **A chronic exposure is just a critical-control failure in slow motion. The control is ventilation, or suppression, or enclosure, and it fails the same way the acute ones do: in execution.**
>
> FOLDING HEALTH INTO THE SYSTEM

Measure the barrier, not the body count. The number that matters moves before anyone is hurt.

THE INSTRUMENT THIS PAPER CLOSES ON

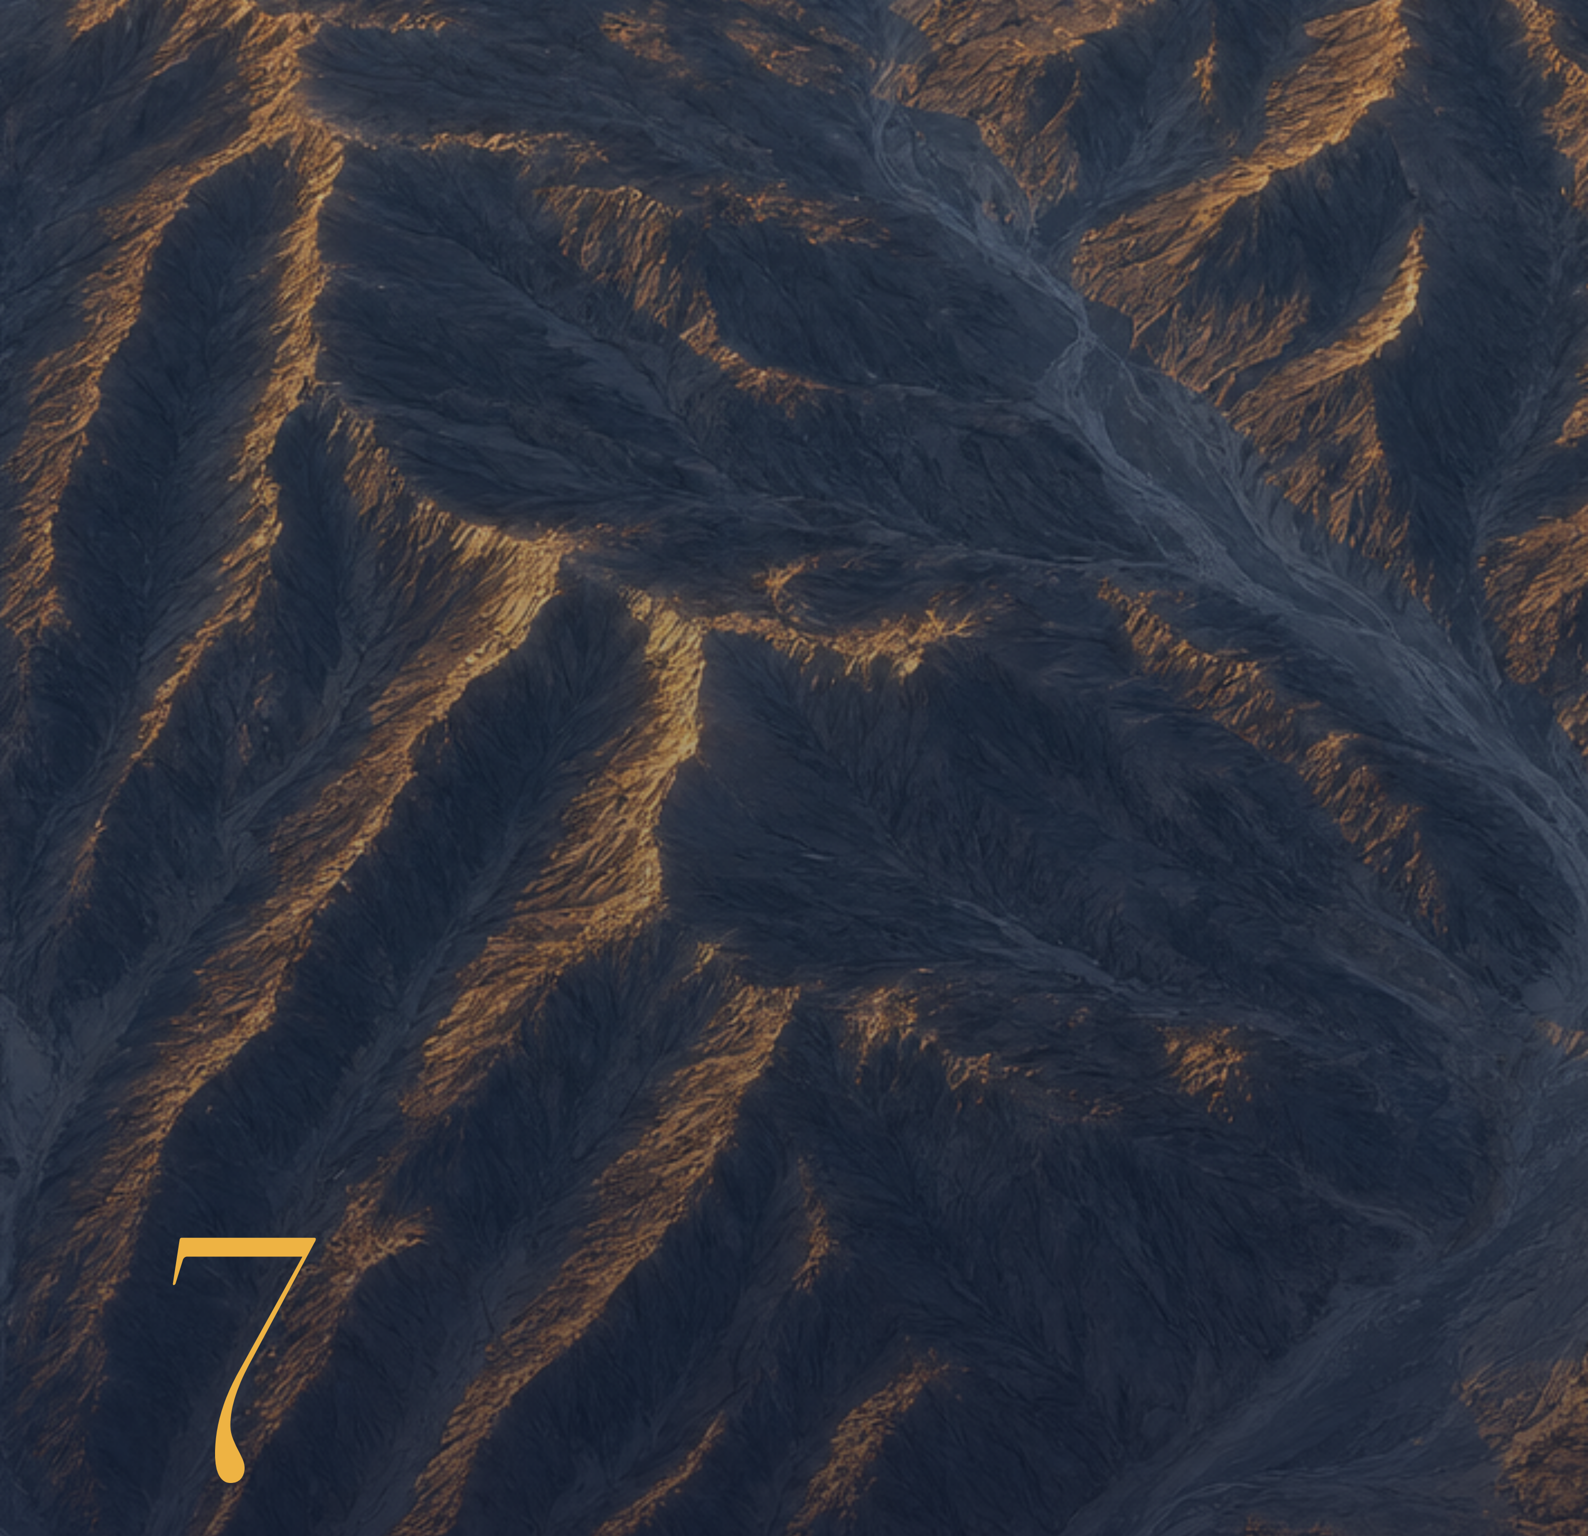

7

SHARE · THE AURUS HSE OPERATING SYSTEM

The Aurus HSE operating system

A capability instrument built from the record: critical controls, verified at the face, dashboarded on leading indicators, nested in enterprise risk. No delivered corporate-HSE mandate is claimed.

9
CCM STEPS THE SYSTEM OPERATIONALISES (ICMM, 2026)

Step 8
VERIFY & REPORT: THE HINGE THE 2024 DATA IDENTIFIES (ICMM, 2026)

Leading
INDICATORS, NOT LAGGING COUNTS: MEASURE THE BARRIER, NOT THE BODY COUNT

This closing chapter turns the preceding six into an instrument. It is a capability framework, argued entirely from the cited public record and from the data-verification discipline of delivered field work; it is explicitly not a claim of delivered corporate-HSE mandates, which the author does not hold. Its purpose is to give a site, a board or a lender a structured way to ask the one question the 2024 data makes decisive: are our critical controls verified as effective, at the point of work, every shift?

The system, in four layers

The operating system has four layers, and each answers a failure mode the earlier chapters named. A **critical-control register** (answering Chapters 3–4: name the few controls that carry the fatal energy, chosen high in the hierarchy of control). A **field-verification routine** (answering Chapter 5: verify each critical control at the point of work, on the contractor's shift as much as the employer's, because that is where 67 per cent of failures live). A **leading-indicator dashboard** (answering Chapter 2: measure control effectiveness, not injury frequency, so the system watches the risk that kills). And **enterprise-risk integration** (answering the 2026 guide's structural change: CCM nested inside enterprise risk management, so a degraded control reaches the board as a business risk, not a safety footnote). Exhibit 15 sets the layers out.

EXHIBIT 15 · FRAMEWORK

The operating system: four layers, each closing a failure mode the data named

A capability instrument. Each layer maps to a chapter finding; no member or project data

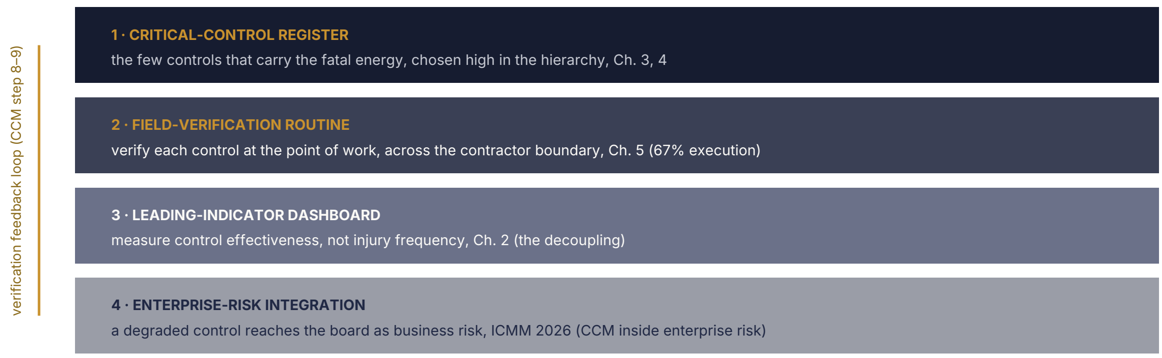

The layers are a loop, not a stack: layer 3 feeds layer 4, which re-tasks layer 1. The 2024 data lives in the join between layers 1 and 2.

A capability instrument built on ICMM, Critical Control Management: Good Practice Guide, 2026, and the fatal-cause findings of Chapters 1–5 (ICMM, 2025). No member or project data; the framework is argued from the cited public record, never from safety-performance data the author holds.

The dashboard measures barriers, not bodies

The single most important design choice in the system is what the dashboard counts. A lagging dashboard counts injuries and fatalities, events that have already happened, and Chapter 2 showed why that is a dangerous instrument on its own: the injury line fell for a decade while the fatal line reversed. A leading dashboard counts the state of the critical controls: verification coverage, verification pass rate, the count of controls found degraded and the time to restore them, and near-miss reports with fatal potential. It answers "how many of our critical controls did we verify as effective this week, and how many did we find failed?", a question whose answer moves before anyone is hurt. Exhibit 16 mocks the leading-indicator view.

EXHIBIT 16 · FRAMEWORK

A leading-indicator dashboard: what moves before anyone is hurt

Illustrative leading-indicator set for a critical-control operating system. A capability instrument; values are placeholders, not data

Verification coverage	Verification pass rate	Controls degraded	High-potential near-misses
%	%	n	n
your metric here	your metric here	your metric here	your metric here
critical controls checked this period	of checks the control passed	open, with restore time	events with fatal potential

Lagging counts (fatalities, TRIFR) remain, but as outcomes to explain, never as the primary steering signal.

The dashboard's test: does the number move before someone is hurt, or only after?

Placeholder tiles: this is a capability instrument, not a report of any site's performance.

A capability instrument. The indicator set operationalises steps 8 and 9 of the 2026 CCM guide (verify the critical controls; evaluate and improve them). Values are placeholders; no site data is shown or implied.

What the system asks of an operator, a board and a lender

For an **operator**, the system is a build order: name your material unwanted events, select the critical controls that carry each, specify how each is verified, and stand up the dashboard before the next audit, not after the next fatality. For a **board**, it is a governance test: can you see, this quarter, how many critical controls were verified and how many failed, and does a degraded control reach you as a business risk? For a **lender or insurer**, it is a diligence lens: an operator who can show verified critical controls and a leading-indicator record is demonstrably managing the fatal risk the 2024 data exposed; one who can show only an injury rate is reporting the wrong number. The consultants who advise on this work bring combined decades of experience in geology, geophysics, geochemistry, drilling, ESIA, scoping studies, PFS, DFS, NI 43-101, JORC and AACE-class estimating; the HSE operating system extends that same field-verification discipline, framed as capability, to the critical-control problem.

Standing it up: a hundred days, not a five-year programme

The objection to any framework is that it takes years to implement and so never starts. The critical-control operating system resists that objection because its logic is subtractive: it does not ask a site to control everything better, only to identify the few controls that carry the fatal risk and verify those relentlessly. That can begin inside a quarter. A defensible hundred-day sequence runs: in the first month, name the material unwanted events and select the critical controls for the top three, almost certainly including the mobile-equipment interaction and fall of ground that Chapter 4 identifies; in the second, specify what "effective" means for each, assign a named owner, and design the field-verification check; in the third, run the verifications, stand up the leading-indicator dashboard, and report the first cycle to the board as enterprise risk. What that hundred days produces is not a finished system but a working one, a small number of critical controls being verified and reported, which is already more than the 2024 fatal record suggests most operations could evidence. Maturity comes later; verification comes first, because verification is what was missing.

READING THE ROADMAP

Three failure modes of a CCM rollout, and how to avoid them

Critical-control programmes fail in predictable ways, and each maps to a chapter of this paper. **Control inflation:** naming forty "critical" controls so that none is truly verified, defeated by the four-test selection discipline of Chapter 3, which is subtractive by design. **Paper verification:** a verification routine that checks documents rather than the control at the face, defeated by Chapter 5's insistence that verification is a physical check with an accountable owner and a pass/fail standard. **The contractor gap:** a regime that stops at the employer's boundary while 58 per cent of hours sit on the other side of it, defeated by writing the same control definitions, verification and reporting into the contract before mobilisation. A rollout that avoids these three is already ahead of the record.

AURUS PRACTICE NOTE

The firm reports no delivered corporate-HSE mandate, and this paper claims none. Its relevant capability is the field-protocol discipline it applies to technical work, specify a standard, verify at the point of capture, reject what fails, with field protocols structured into our mandates. That discipline, evidenced by 100 per cent data-quality acceptance on a delivered field campaign under Competent-Person supervision, is the adjacency from which the critical-control operating system is offered as a capability instrument, not a track record.

EXHIBIT 17 · FRAMEWORK

Ten questions that separate a verified system from a documented one

A diligence checklist for operators, boards and lenders. A capability instrument; each question maps to a finding of this paper

#	THE QUESTION TO ASK	EARNED BY
1	Are your material unwanted events named, and are the critical controls for each selected by a defined test?	Ch. 3
2	Is the number of critical controls small enough that each is actually verified?	Ch. 3
3	For each critical control, is there a written standard of what "effective" means?	Ch. 3
4	Does each critical control have a single accountable owner, by name?	Ch. 3
5	Is verification a physical check at the point of work, not a document review?	Ch. 5
6	Does the same regime apply, unchanged, to contractors carrying most of the hours?	Ch. 5
7	Does your dashboard show control state (leading), not just injury counts (lagging)?	Ch. 2, 7
8	When a control is found degraded, is there a defined response and a restore clock?	Ch. 3, 7
9	Do high-potential near-misses reach the same scrutiny as recordable injuries?	Ch. 2
10	Does a degraded critical control reach the board as an enterprise risk?	Ch. 7

A capability instrument. Each question operationalises a finding of this paper against the 2026 CCM guide (ICMM, 2026); an operation that can answer all ten is managing the fatal risk the 2024 data exposed. No member or project data.

An operator, a board or a lender who works down that list learns quickly where a safety programme sits on the spectrum from documented to verified. The uncomfortable likelihood, given a fatal record in which 83 per cent of deaths involved a critical control not effectively implemented, is that many programmes today would answer "yes" to the first four questions, the ones about documents, and "not really" to the last six, the ones about verification and response. That gap is the whole distance between the industry's intent and its outcomes, and closing it is the work this paper argues is both necessary and, on the evidence of the 2026 framework, entirely within reach.

What "zero" actually asks for

The argument of this paper reduces to a single reversal of emphasis. The industry's fatal record did not stall because operators stopped caring, stopped training or stopped writing controls; on the 2024 evidence, the controls were largely present, and 83 per cent of deaths involved one of them not being effectively implemented rather than being absent. It stalled where those controls were not verified as effective at the face, on the shift, on the contractor's hours as much as the operator's own. A critical-control operating system is not a new safety programme layered over the existing one; it is the discipline of proving, this quarter, that the few controls carrying the fatal risk are working, and of putting that proof in front of the board as an enterprise risk rather than an injury statistic. That is the whole distance between a target and a slogan, and on the evidence of the 2026 framework it is buildable now, from published good practice, rather than pending on any future dataset.

The test this sets is not whether an operation cares about safety. Every operation says it does, and the 2024 record is not a record of indifference; it is a record of controls that existed and were not proven effective where the work was done. The test is whether an operation can show, this quarter and to a board or a lender, that the handful of controls carrying its fatal risk were verified in the field, and what it did when one was found wanting. On the evidence assembled here, that is the single discriminator between an operation that is managing the fatal risk and one that is only documenting it, and it is the discriminator this paper has tried to make legible, buildable and, above all, checkable. The industry already knows how; the 2026 framework is proof of the how. What the 2024 record asks for is not more knowledge but the discipline to verify what is already known, relentlessly, at the face.

> Zero is not a slogan. It is what you get when every critical control is verified as effective, at the face, on every shift, and when you measure that, not the funerals.

References

Citation policy of this series: every statistic in this paper is transcribed from a cited public source and attributed by institution, publication and year, in the sentence, the exhibit source line or this list, and traces to an archived source in the paper's evidence dossier. Values read from a published chart rather than a data table are flagged as chart-read where they are used and are never treated as exact. Statistics cited by the ICMM from third parties are reported as "ICMM (2025), citing [source, year]" and are not independently re-attributed. Ranges and qualifiers are preserved as published; no point figure is derived from a published range, and the few shares computed from published totals are labelled as computed where they appear.

1. International Council on Mining and Metals, Safety Performance Report: Benchmarking Progress of ICMM Company Members in 2024, July 2025 (primary source; pp.1–14 transcribed to the evidence dossier).
2. International Council on Mining and Metals, Critical Control Management: Good Practice Guide, 2026 (comprehensive update; nine-step process).
3. Minerals Council South Africa, Fall of Ground Action Plan, 2021, cited within reference 1 as the root-cause echo for the 2024 fall-of-ground triggers; not independently re-attributed.
4. Earth Moving Equipment Safety Round Table (EMESRT), Vehicle Interaction Control Effectiveness (the Nine Layer Model of Control Effectiveness), the recognised reference for the mobile-equipment control hierarchy of Exhibit 9; named as a framework, with no statistic attributed to it. Archived to this paper's evidence dossier.
5. Mining Weekly, note on the 2026 ICMM Critical Control Management guide update, 29 April 2026 (trade summary; cross-confirmation of reference 2's nine-step structure).
6. im-mining, note on the ICMM CCM guide update following the fatality rise, 10 July 2025 (trade summary; cross-confirmation).
7. International Council on Mining and Metals, Health and Safety Performance Indicators guidance, 2021, the common indicator set underlying the benchmarking series since 2012.
8. International Agency for Research on Cancer (World Health Organization), Diesel and Gasoline Engine Exhausts and Some Nitroarenes, IARC Monographs on the Evaluation of Carcinogenic Risks to Humans, Volume 105, 2012, the established scientific classification of diesel-engine exhaust as a Group 1 human carcinogen; the source for the diesel-particulate carcinogen statement in Chapter 6 and the glossary. Cited as an attributed established classification, not a benchmarking statistic; no headline figure depends on it.

Whitelist note. References 5 and 6 are trade-press items and sit outside this series' institutional source whitelist (ICMM, EMESRT and the Minerals Council South Africa). They are admitted here only as non-whitelisted cross-checks: neither carries an attributed statistic, and both serve solely to corroborate the nine-step structure of the 2026 CCM guide (reference 2), which remains the primary source for that material. Reference 8 (the IARC Monographs classification of diesel-engine exhaust) is an institutionally published scientific reference and sits inside the whitelist; it is carried as an established classification, not a benchmarking statistic, and no headline figure depends on it.

Exhibit source index

Each exhibit's primary sources, by reference number above. Framework exhibits (5, 6, 9, 14, 15, 16, 17) are capability instruments and carry no member or project data by design.

EXHIBIT	SUBJECT	REFERENCES
1	Injury vs fatality trend, 2012–2024	1
2	Fatal-frequency rate by continent, 2024	1
3	Fatalities and FFR by country, 2024	1
4	The injury/fatality decoupling (indexed)	1
5	From counting incidents to controlling energy	(framework)
6	Critical-control bowtie (signature)	2 (framework)
7	The 2026 CCM nine-step process	2, 5, 6

EXHIBIT	SUBJECT	REFERENCES
8	Leading causes of fatality, 2024	1
9	Hierarchy of control, mobile equipment	4 (framework)
10	Fall-of-ground triggers and controls	1, 3
11	Control-failure typology, 2021–2024	1
12	Organisational factors, 2024	1
13	Contractor vs employee, 2024	1
14	Chronic occupational-health exposures	1 (framework)
15	The Aurus HSE operating system	2 (framework)
16	Leading-indicator dashboard	2 (framework)
17	Ten-question diligence checklist	2 (framework)

Citing this paper

Cite as: Aurus Technical Committee, Zero Is an Engineering Problem: Critical Control Management and the Reversal of Mining's Fatality Curve, Aurus Institute for Resource Development, Technical Paper WP-18, July 2026. Statistics quoted from this paper should carry their original institutional source as listed above, principally the ICMM, rather than this paper, as the primary attribution; this paper transcribes, it does not originate.

Exhibit conventions. Across all exhibits: gold marks the load-bearing element, the critical control, the leading hazard, the execution-failure share; grey marks context and comparison; outlined shapes mark critical controls in the bowtie and the framework diagrams; and any value read from a published chart rather than a data table is flagged as chart-read at the point of use. Framework exhibits carry no data by design and say so in their source lines.

Freshness. The next ICMM benchmarking report (2025 data) is expected around July 2026; readers should treat the 2024 figures here as a snapshot at press time. Where this paper reports a dated commitment, the December 2024 diesel-particulate-matter measures, the 2026 CCM guide update, the date is carried as published.

Source archive and reproducibility. The headline evidence of this paper is read from a single primary artifact, the ICMM 2024 benchmarking report, archived to the paper's evidence dossier with the transcribed pages recorded. The three framework sources are archived in full alongside it: the ICMM 2026 Critical Control Management guide, whose Table 2 and Figure 2 are the source for the nine-step process of Exhibit 7 and are cross-confirmed by the two trade summaries at references 5 and 6; the EMESRT vehicle-interaction control-effectiveness guide behind Exhibit 9; and the IARC Volume 105 monograph behind the Chapter 6 diesel-exhaust classification. No cited source of this paper is unarchived. A reader who wants to test a figure can move from the sentence to its source line, from the source line to the dossier row that transcribes the value in its published form, and from that row to the archived document, and should find the figure unchanged at each step. Every value read from a published chart rather than a data table is marked chart-read where it is used and again in the Method note that follows; where the public record is silent, the silence is named in the gap register rather than bridged by estimation. This is the same test the paper asks operators to apply to a control: not whether it is written down, but whether it holds when checked.

The benchmarked population. The figures in this paper aggregate the ICMM's member companies, 24 in 2024 and 25 in 2023 before the Newmont and Newcrest consolidation, reporting on a common health-and-safety indicator set in place since 2012. That population is large-company mining, and it is the strength and the limit of the benchmark at once: it is consistent, exposure-adjusted and comparable year on year, and it does not capture the artisanal and non-member fatality tail, which this paper names in the gap register rather than estimating. Where a figure is described as an industry finding, it should be read as a finding for the benchmarked members, which is the population the source measures.

A note on small numbers. Because the benchmarked population records only a few dozen fatalities in a year, a single-year movement in any one sub-category can be large without being meaningful, and the report itself sets its annual figures against multi-year context for that reason. This paper follows the same discipline: it reads a year-on-year swing, such as the fall in the contractor share of fatalities from 2023 to 2024, as a signal to be interpreted against the structure beneath it rather than as a trend in its own right, and it says where a figure is a single-year reading against a small base. The durable findings of this paper, the decoupling of injury from fatality, the dominance of execution failure, and the concentration of fatal risk in a few hazards and settings, are the ones that hold across the multi-year window rather than in any single year.

Method and evidence

Every figure in this paper was transcribed, archived and audited before it was typeset. This page records how, so that a reader, or a lender's technical adviser, can test the paper the way the paper asks operators to test their controls.

The evidence chain

Each statistic traces through three layers. First, a **source line** in the sentence or exhibit where it appears, naming institution, publication and year. Second, an entry in the paper's **evidence dossier**, which transcribes the figure in its exact published form and records every correction made during review. Third, an **archived artifact**: for this paper, principally the ICMM benchmarking PDF, downloaded from the report's own content-delivery network and read directly, pages 1–14, with the full 2026 CCM Good Practice Guide, the EMESRT vehicle-interaction guide and the IARC Volume 105 monograph archived alongside it and read to the page. The ICMM is admitted to this series' source set as an institutionally published, attributed benchmarking body under the recorded whitelist-extension protocol.

Chart-read values, flagged

The report's two most-quoted structural findings, that 67 per cent of control failures over 2021–2024 were failures of execution, and that 83 per cent of 2024 fatalities involved a failure to implement effective critical controls, are text-verbatim in the primary source and are used as such. Several per-category counts, however, are read from the report's graphs rather than a published data table: the control-failure typology counts (execution ~24, absent ~6, design ~5) and the organisational-factor counts (rules/procedures ~10, risk management ~6, management of change ~4). These are flagged as chart-read wherever they appear (Exhibits 11 and 12) and are treated as approximate; the text-verbatim percentages carry the argument, and the chart-read counts illustrate its shape. One period point is worth stating plainly, because it is easy to get wrong: the typology graph (Graph 10) is a stacked four-year series covering 2021–2024, not a single 2024 year, so the ~24/~6/~5 heights are cumulative over that window and Exhibit 11 is labelled 2021–2024 accordingly; only the 83 per cent figure and the organisational-factor graph (Graph 8) are single-year 2024 readings.

VALUE CLASS	TREATMENT	WHERE
Text-verbatim percentages (67%, 83%, 50%, 21%, 45%)	Used as published, exact	Throughout; At a glance
Graph-read category counts (~24/~6/~5; ~10/~6/~4)	Flagged chart-read; approximate	Exhibits 11, 12
Shares computed on the 42-fatality total	Labelled as computed	Exhibit 8
Indexed trend series	Schematic between published anchors	Exhibits 1, 4

Open parameters and gaps, named

Where the archived public record is silent, this paper states the silence rather than filling it. Every source cited in this paper is archived on disk, including the three framework references, the 2026 CCM guide, the EMESRT vehicle-interaction guide and the IARC Volume 105 monograph, added in full for this edition; no cited source is unarchived. What remains open is not a missing document but a class of data the public record does not yet publish. As of this edition that open set, carried in full in the evidence dossier's gap register, is three items, none of them bridged by estimation:

- **Serious-injury-and-fatality precursor literature.** Chapter 2 argues the reframing from acute counts to precursor control qualitatively, pending an archived peer-reviewed primary pull; no consultancy secondary source is admitted in its place.

- **The artisanal and non-member fatality tail.** Named as sitting outside the ICMM member benchmark and deliberately not quantified here, because no whitelisted institutional source for it is yet archived; the claim is narrowed to the benchmarked population rather than estimated across the whole industry.
- **Industry-wide occupational-health metrics.** Chapter 6 is deliberately qualitative because no member-wide health dataset yet exists on the safety benchmark's footing; the one dated public anchor, the December 2024 diesel-particulate-matter commitment, is reported as such and the other exposure classes are named from standard occupational-hygiene practice, not quantified.

What this paper never does

No figure is remembered, averaged from undisclosed sources, or inferred. Ranges and qualifiers are preserved exactly as published. The closing operating system (Chapter 7) is a capability instrument argued from published good practice and from the firm's field-data discipline; it makes no claim of a delivered corporate-HSE mandate, and no safety-performance data of the firm's own is reported, because none is held.

How to challenge a figure in this paper

Because the argument turns on a small number of load-bearing statistics, this paper is written to be checked rather than trusted, and a reader who wants to contest a number should be able to do so in four moves. First, find the figure's source line, in the sentence, the exhibit or the reference list, and confirm the institution, publication and year. Second, decide whether the figure is used as text-verbatim or as chart-read: the headline shares, the 42 fatalities, the 21 per cent mobile-equipment share, the 67 and 83 per cent control findings, the 45 per cent contractor share, are transcribed word for word from the primary report and should match it exactly, while the per-category counts of Exhibits 11 and 12 are read from graphs and are offered only as approximate. Third, check the period: several readings span 2021–2024 rather than the single 2024 year, and the paper labels each accordingly. Fourth, where a value is absent, read the gap register above rather than inferring one, because a silence here is deliberate. A figure that survives those four checks is doing the work the paper claims for it; one that does not should be treated as the paper's error, not the source's, and corrected against the source.

Every correction made to a figure during this paper's review is recorded in the evidence dossier against its row, so the published values are the reconciled ones rather than the first drafts. Two are worth naming here because they are the kind of error the method exists to catch: the 2019 fatality spike was reconciled to its published components, of which the great majority sits in a single catastrophic tailings event, so it is not read as a trend; and the control-failure typology of Exhibit 11 was re-labelled to its true 2021–2024 period after an early draft mis-stated it as a single 2024 year. Neither correction changed a headline conclusion, which is the point: the discipline is meant to be invisible in the result and auditable in the record.

A figure written to be checked is a figure worth trusting. Every number in this paper is built to survive the four moves above.

Glossary of critical-control terms

The vocabulary used in this paper, stated precisely. Terms marked with a chapter number also appear as margin definitions at first use.

Bowtie · a risk model with threats and preventive controls to the left of an unwanted event and consequences and mitigating controls to the right; makes explicit which barriers are the critical ones. (Ch. 3)

Critical control · a control crucial to preventing a fatal event or mitigating its consequences, whose failure would significantly raise the risk despite other controls. (Exec. summary; Ch. 3)

Critical Control Management (CCM) · the discipline of identifying, specifying, implementing and verifying critical controls; the subject of the ICMM's 2026 Good Practice Guide. (Ch. 3)

Diesel particulate matter (DPM) · the classified-carcinogen soot of diesel exhaust; a defining underground chronic exposure and the subject of a December 2024 ICMM commitment. (Ch. 6)

Execution failure · a control that exists on paper and in training but is not effective at the point and moment of work; the dominant 2024 failure mode. (Ch. 5)

Fall of ground · the collapse or ejection of rock from a face, back or wall; the second-leading fatal hazard in 2024, concentrated in South African underground operations. (Ch. 4)

Fatal-frequency rate (FFR) · fatalities per million hours worked; the exposure-adjusted measure that reveals regional and contractor risk a raw count hides. (Ch. 1)

Hierarchy of control · the ranking of control types by reliability, from elimination down to personal protective equipment; critical controls are chosen high in it wherever the energy allows. (Ch. 4)

High-potential near-miss · an event that did not cause harm but could credibly have been fatal; a leading indicator of critical-control state. (Ch. 7)

Lagging indicator · a measure of harm that has already occurred, such as a fatality count or injury frequency; necessary but insufficient as a steering signal. (Ch. 2, 7)

Leading indicator · a measure of control state that moves before harm occurs, such as verification coverage and pass rate. (Ch. 7)

Lockout / tagout (LOTO) · the energy-isolation discipline that renders stored energy safe before work; breaches rose in 2024 and were flagged as a reinforcement priority. (Ch. 5)

Material unwanted event (MUE) · an event with the potential for a fatality or catastrophic loss; the anchor of the CCM process's planning phase. (Ch. 3)

Mobile equipment interaction · a collision or crush event involving heavy mobile machinery; the leading fatal hazard in 2024. (Ch. 4)

Mitigating control · a control that limits the harm of an event after it has begun; the right-hand side of the bowtie. (Ch. 3)

Preventive control · a control that stops a threat reaching the unwanted event; the left-hand side of the bowtie. (Ch. 3)

Serious injury and fatality (SIF) · the class of events that kill or permanently disable, with precursors distinct from the minor-injury base. (Ch. 2)

Support before entry · the underground discipline of installing ground support to standard before access is permitted; the control defeated in three of five 2024 fall-of-ground deaths. (Ch. 4)

Total recordable injury frequency (TRIFR) · recordable injuries per million hours worked; the industry's standard injury measure, which fell from 5.07 (2012) to 2.29 (2024). (Ch. 1)

Verification · the act of checking, in the operating environment, that a critical control is present and effective; step 8 of the 2026 CCM process and the hinge of the whole system. (Ch. 3, 7)

AURUS INSTITUTE FOR RESOURCE DEVELOPMENT

Zero Is an Engineering Problem: Critical Control Management and the Reversal of Mining's Fatality Curve
A Technical Paper in the Aurus white paper library, on the Share pillar. Prepared by the Aurus Technical Committee, July 2026.

Every statistic in this paper traces to the cited public sources listed in the References, principally the International Council on Mining and Metals; chart-read values are flagged as such; ranges are preserved as published. Nothing in this paper draws on client-derived data, none of the cited figures is Aurus analysis, and no delivered corporate-HSE mandate is claimed.

Set in Fraunces (display), Source Serif 4 (text) and Inter (captions and data). Typographic system and exhibits art-directed in-house; this edition carries no photographic or generated imagery by design.

Working draft 1: wave-1 Technical Paper; adversarial certification pending.

MINING · INFRASTRUCTURE · ENGINEERING · ENVIRONMENT