


Beyond Green Ticks: A Human-Centred Method For Evaluating Critical Control Effectiveness In Mining

H Lourens¹, K Genn²

1. **CEO, Stratflow Australia / Senior Consultant, Art of Work**, Brisbane QLD 4000, Australia. Email: hendrik@stratflow.com.au
2. **Founder, Art of Work Solutions**, Melbourne VIC 3000, Australia. Email: kelvin@artofwork.solutions

Keywords: critical controls, control effectiveness, human-centred design, Safety-II, hierarchy of controls, energy damage model, work-as-done, human factors, critical control management, mining safety



Copyright 2026

Abstract

Australian mining safety regulation has shifted decisively from compliance auditing towards risk-based approaches grounded in critical control management (CCM) and high reliability organisation (HRO) thinking. In Queensland, the Resources Safety and Health Legislation Amendment Act 2024 makes CCM a prescribed requirement, with the post-grace-period compliance test commencing on 1 June 2026. Regulators and boards are no longer satisfied with evidence that a procedure exists; they are asking a harder question: how reliable are critical controls under real operational conditions? Traditional reviews remain largely binary: control present or absent, inspection done or not, which creates a false sense of assurance that fatally weak controls appear fully compliant.

This paper describes Control Effectiveness Assessment (CEA), a structured, field-tested methodology that directly addresses that question. CEA combines the Energy Damage Model, the hierarchy of controls, and human-centred design principles to evaluate not only the presence of critical controls but also their practical reliability. Multi-disciplinary workshops with subject matter experts (SMEs) and frontline personnel score each control on four interdependent dimensions: Robustness, Human Dependency, Alignment and Adaptability. The output is a quantitative reliability profile that exposes where controls are fragile, overly human-dependent or not aligned with how work is actually done.

Beyond a single assessment, the workshop architecture turns the workforce into a distributed sensor network for control health, it surfaces drift in control dependability over time as a leading indicator of failure, and provides defensible evidence of 'reasonably practicable' decision-making. Using an anonymised critical-risk case study in mobile equipment and driving safety from a large Australian asset-intensive operator, the paper demonstrates how CEA exposes the gap between work-as-imagined and work-as-done, differentiates true hard controls from supporting administrative measures, integrates digital indicators into a meaningful picture of control health, and offers an effective framework that mining operators can apply to focus assurance efforts on the controls that actually matter.

Introduction

For more than three decades, mining safety assurance has rested on the assumption that a documented control, when applied consistently, is effective. Bowties, performance standards, audit checklists and verification standards have been built around that assumption. The architecture works well at one task: confirming that a control exists. It works much less well at another: confirming that the control reduces hazardous energy under the conditions in which the work is actually done.

That distinction is no longer academic. Australian mining safety regulators have shifted decisively from compliance auditing towards risk-based approaches grounded in critical control management (CCM) and high reliability organisation (HRO) thinking (ICMM, 2015; Hopkins, 2019; Weick and Sutcliffe, 2007). The shift is most visible in Queensland, where the Resources Safety and Health Legislation Amendment Act 2024 (Qld) (State of Queensland, 2024) introduced an explicit statutory framework for CCM. The provisions commenced on 1 June 2025 with a 12-month grace period, ending on 1 June 2026. From that date, Resources Safety and Health Queensland (RSHQ) tests CCM not by the presence of a documented system, but by whether each control actually holds up under production conditions. For coal operations, integration into Principal Hazard Management Plans (PHMPs) is supported by Queensland Guidance Note QGN35 (RSHQ, 2024); for minerals and quarrying operations, no equivalent guidance exists, and the adoption into the safety and health

management system has to be argued through procedures, standards, training, verification and governance.

The post-1-June compliance posture is consequential. RSHQ has signalled that its enforcement approach will combine documentation review with unannounced inspections and spot-checks of operating practices, including in response to worker complaints. Enforceable undertakings are now available as an alternative to prosecution. The question that an inspector will ask of a critical control is not whether it appears in the bowtie. It is, of each control: under what conditions does this control fail, what evidence is there that it has not failed, and what has been done to strengthen it where it is fragile? Most operators arriving at 1 June 2026 will have bowties, performance standards and verification protocols in place. Far fewer will be able to demonstrate effectiveness in those terms.

This paper sets out a method intended to close that gap. Control Effectiveness Assessment (CEA) is a structured, field-tested methodology that produces a quantitative reliability profile for each critical control, derived from multidisciplinary workshops with SMEs and frontline employees, and scored across four interdependent dimensions: Robustness, Human Dependency, Alignment, and Adaptability. The method draws explicitly on the Energy Damage Model (EDM) (Haddon, 1973; Viner, 1991), the hierarchy of controls (NIOSH, 2015; Manuele, 2005), and human-centred design (HCD) principles applied to safety (Norman, 2013; Dekker, 2014). It is consistent with Safety-II thinking (Hollnagel, 2014) and with the wider human and organisational performance (HOP) movement (Conklin, 2019). Applied repeatedly over time, the method also turns the workforce into a dispersed sensor network for control health and surfaces drift in control reliability as a leading indicator of failure.

The case material in this paper is drawn from a critical-risk control effectiveness assessment for the driving-safety portfolio of a large Australian asset-intensive operator with a fleet exposed to public-road driving and field-services work; the case is anonymised here, but the structural findings translate directly to mining critical risks. The remainder of the paper is organised as follows. The next section sets out why traditional binary critical control reviews fail under the new regulatory test. The following sections describe the conceptual foundations of CEA, the four-dimensional reliability frame, the workshop method, the human sensor network framing and the anonymised case study. The paper then translates CEA findings into typical mining critical-risk scenarios, discusses how the method generates evidence of 'reasonably practicable' decision-making, and concludes.

Why Binary Critical Control Reviews No Longer Suffice

Traditional critical control assurance produces binary verdicts. A control is either present or absent; an inspection is either done or not done; a procedure either exists or does not exist. The resulting picture is administratively tidy and audit-friendly. It is also frequently misleading. A control can score 100 per cent compliant on the verification protocol while being routinely worked around at the workplace, fundamentally not aligned with working conditions, or so cognitively demanding that its actual reliability is low. The binary frame cannot see those failure modes.

Three framework deficiencies follow. First, work-as-imagined diverges from work-as-done. The procedures described in the bowtie, the standard operating procedure, the training package, and the audit checklist outline how work is supposed to be performed (Hollnagel, 2014). The work that actually happens, under production pressure, with shift handovers, contractor variability, unplanned weather, equipment availability constraints and competing priorities, is systematically different. Verification activity that focuses only on the

imagined work cannot see the gap, and it is where fatalities and high-severity incidents tend to occur (Dekker, 2014).

Second, hard controls and supporting administrative measures get confused. The hierarchy of controls (NIOSH, 2015) places elimination, substitution and engineering controls above administrative controls and personal protective equipment (PPE). On a typical critical-risk bowtie, however, procedures, training packages, permits, fatigue policies and supervisory routines are routinely labelled 'critical controls' alongside engineering controls. The labelling is not a stylistic preference. It shapes investment decisions: when administrative measures are treated as critical controls, the operator is, in effect, deferring engineering investment while still satisfying the assurance audit. The cost is not abstract. The Brady review of fatal accidents in Queensland mining (Brady, 2020) found that the corrective actions put in place after Serious Accidents were, in 62 per cent of cases, administrative controls only (the weakest form in the hierarchy). Part of the explanation, developed further below, is that operational systems often lack the capacity to implement physical or engineering controls that require sustained planning, so the operator defaults to the lowest tier of the hierarchy. That default is, in effect, what binary assurance produces when the engineering investment case has not been made.

Third, the digital indicators that ought to inform a control health picture are largely treated as compliance artefacts rather than as predictive evidence. Telematics overspeed counts, proximity-detection alerts, gas-monitoring trends, isolation-permit cycle times, and other data streams are reported because the management system requires them, not because they have been integrated into a structured assessment of which controls are degrading and why. A critical control health picture that does not leverage the available digital information leaves significant assurance value on the table.


A method that can address these three weaknesses needs to evaluate the reliability of a control rather than its presence; it needs to draw on how work is actually done, not just how it is imagined; and it needs to integrate digital indicators inside the structured framework. CEA is built for that purpose.

Theoretical Bases

CEA rests on four well-established bodies of work that, taken together, give the method its analytical reach.


The Energy Damage Model

Haddon (1973) framed harm as the unwanted transfer of energy from a source to a recipient that lacks the capacity to absorb it. A control, in EDM terms, must do one of three things: prevent the energy from being released, redirect or attenuate the energy after release, or protect the recipient from the energy that does arrive. The model gives a precise diagnostic question for any candidate critical control: which energy pathway does this control interrupt, and how? Controls that cannot answer that question with a specific mechanism are typically administrative supports rather than hard controls. The EDM lens, applied consistently across a critical-risk portfolio, sharpens the conversation about what is really doing the safety work.




The hierarchy of controls

The hierarchy of controls (NIOSH, 2015; Manuele, 2005) provides a structural ranking of control types by inherent reliability. Elimination, substitution and engineering controls do their work without depending on a person to act in a particular way at a particular moment. Administrative controls and PPE depend primarily on human action and are therefore inherently less reliable. The hierarchy does not say that administrative controls are unimportant; it says that they should not be relied on as the primary control where a higher-order alternative is reasonably practicable. CEA treats the hierarchy as a diagnostic frame: where the primary control for a critical risk sits low in the hierarchy, the reliability profile must be unusually strong elsewhere, or the operator must provide a defensible argument that higher-level options are not reasonably practicable.



Human-centred design and Safety-II

Human-centred design (Norman, 2013) starts from the person doing the work. It asks what the work feels like under normal and abnormal conditions, what cognitive load the control imposes, and whether the control is experienced as a support or a punishment. Applied to safety, HCD shifts assurance from doing safety to people towards doing safety with people (Dekker, 2014). The complementary lens, Safety-II (Hollnagel, 2014), asks why work usually goes right rather than only why it occasionally goes wrong. Together, HCD and Safety-II ground the assessment in the reality of the workforce, surface workarounds and adaptations as evidence rather than transgressions, and provide structured access to the work-as-done picture. Recent extensions in human and organisational performance (Conklin, 2019) and learning teams provide the facilitation grammar for assembling that picture without producing blame.



How the foundations combine

EDM tells the assessment what a control must do. The hierarchy tells it which control types are inherently more reliable. HCD and Safety-II tell it how to see what the control actually does in the field. The combination enables CEA to move beyond presence/absence to a structured, defensible assessment of reliability.

The Four-Dimensional Reliability Frame

Each critical control is scored on four interdependent dimensions. The dimensions are designed so that scoring requires triangulation across evidence sources: design documentation, maintenance and inspection records, structured frontline assessment, near-miss patterns, abnormal-condition walkthroughs and digital indicators. The four dimensions are summarised in Table 1.






Table 1

The four CEA reliability dimensions, primary evidence sources and typical failure modes.

Dimension	What it measures	Primary evidence	Typical failure mode
	Engineering strength of the control, its position in the hierarchy of controls and the organisational support required for it to function as designed.	SME assessment, design documentation, maintenance and inspection records, digital indicators of control health.	Strong on paper but unsupported by maintenance, training access, staffing or spare-parts availability; the control degrades silently between audits.
	The focus, memory and skill required of a person for the control to function under normal and abnormal operational conditions.	Frontline workshop scoring, observation, structured task analysis, fatigue and workload data.	High cognitive load under production pressure; works in steady state but is fragile under variability, congestion or shift change.
	Whether the control helps or hinders the worker doing the work; whether it is experienced as a support or as a punishment.	Frontline workshop scoring, near-miss data, workaround patterns, telematics gaming patterns.	Technologically robust but worked around or gamed because of how it is used or perceived (eg punitive monitoring, intrusive journey-management apps).
	How the control behaves outside the design envelope: contractor variability, night shift, congestion, abnormal weather, operational disruption.	SME and frontline scenario walkthroughs, abnormal-condition reviews, near-miss data, post-incident learning.	Designed for the steady state; breaks down precisely in the conditions where the energy event is most likely.

Each dimension is scored 1 to 5, with 5 representing high reliability on that dimension. The composite reliability score is presented with the four sub-scores visible rather than collapsed into a single number, so that the assessment supports a remediation discussion rather than a single rating. Scoring is anchored by behavioural descriptors at each level, ensuring reliability across assessors and over time.



The four dimensions are not independent. A control that is technically robust but poorly aligned will degrade through gaming and workaround behaviour even when the engineering remains intact. A control with high human dependency will be fragile under adaptability stress, because the same human performance variability that makes it work in a steady state is what fails under abnormal conditions. The dimensions, therefore, have to be scored in the same workshop frame, by the people best placed to see each one, rather than aggregated from separate audits.

The Workshop Method

CEA is delivered through a structured workshop sequence per critical-risk material unwanted event (MUE) in scope. The sequence has four phases.

The mobilisation phase reviews the existing bowtie, performance standards, and management-system integration documentation for the MUE, then pinpoints critical controls in scope by applying the EDM and the hierarchy of controls to separate hard controls from supporting administrative measures. The output is a control list, a participant plan that includes night-shift and contractor representation, and a digital-indicator inventory.

The SME workshop phase assigns a Robustness score to each critical control. SMEs draw on design documentation, maintenance and inspection records, engineering judgement and digital indicators of control health. The output is a Robustness sub-score per control with a written narrative and an evidence trail.

The frontline workshop phase scores Human Dependency, Alignment and Adaptability. Frontline operators, including night-shift workers and contractor representatives where relevant, are facilitated using structured ethnographic prompts that surface workarounds, abnormal-condition behaviour and near-miss patterns. Psychological safety is built in through learning-team principles (Conklin, 2019): non-punitive discovery, worker voice, and explicit acknowledgement that the people who do the work are the people best placed to see how it actually works. The output includes three sub-scores per control, a written narrative, and an evidence trail.

The reconciliation phase brings the SME and frontline assessments together. Divergences between the two views are deliberately interrogated rather than averaged out, because they are usually where the most actionable findings sit. A control that SMEs score 4 on Robustness, yet frontline staff score 2 on Alignment, is signalling a distinct aspect of how the control is being used, and the reliability profile must reflect that. The output is a final reliability profile per control, a remediation pathway for each fragility finding (design intervention, organisational support, training, supervision, or operational system change), and a regulator-ready report linking each finding back to its underlying evidence.

The method is deliberately structured to produce a traceable evidence trail, so that the resulting reliability profile can be defended to a regulator, a board safety committee or a coronial process.

CEA As A Human Sensor Network For Control Health

CEA's value goes beyond the per-engagement reliability profile. The workshop architecture is also a structured way of using the workforce as a distributed, time-stamped sensor network for control health.

Frontline staff as a dispersed sensor network

Conventional regulatory structures rely on two sources of information about control performance: documentation (do the controls exist on paper?) and incident investigation (what failed after something went wrong?). CEA introduces a third source: the structured, periodic assessment of control effectiveness by the people who interact with those controls every shift. Three properties of that third source matter.

First, it surfaces information that is otherwise only available after incidents. After a serious event, investigation teams routinely discover that controls had degraded, been worked around, or had become misaligned with operational conditions well before the event. Frontline staff knew this. It was part of their daily experience. No structured process existed to capture it. CEA creates that capture mechanism proactively, in advance of the incident that would otherwise trigger discovery of the gap.

Second, distributed assessment surfaces site-by-site and time-based variation. A control that scores reliably at one location may score fragile at another, owing to local conditions, equipment age, crew experience or maintenance history. That geographic and temporal variation is invisible to a compliance-based framework that asks only whether the control exists in the PHMP. For an inspectorate, a board safety committee or a multi-site head office, the distributed picture provides a far richer basis for oversight than aggregate compliance statistics.

Third, the workshop conversations themselves change the quality of organisational knowledge. SMEs who design and specify controls gain direct exposure to the real-world realities of how those controls interact with actual work. The conversations have surfaced findings previously unknown to safety leadership, not because the information did not exist, but because no structured process existed to surface it. Audit and inspection processes, however thorough, assess conformance to a standard; they do not seek to understand the worker's experience of working with the control.

Detecting drift as a leading indicator

Repeating the CEA frontline assessment at regular intervals (typically annually) creates a time series for each control's four sub-scores. That time series makes drift visible.

A control that scores 3.8 on Alignment in year one, 3.4 in year two and 2.9 in year three is signalling something specific: the gap between what the control requires and what the work demands is widening, and workers will be increasingly tempted to work around it. The drift signal is a leading indicator. It predicts future control failures before those failures result in incidents.

The concept of drift into failure (Dekker, 2014) describes how organisations move incrementally toward ruinous outcomes through a series of small, individually rational changes that are invisible in the moment. Each adjustment seems reasonable in context; no single step crosses a clear boundary; the accumulation erodes safety margins until the system can no longer absorb a disturbance. By definition, retrospective compliance checks cannot see drift, because each individual state appeared acceptable when assessed in isolation. A downward trend in a frontline CEA score is exactly the kind of weak signal that high reliability organisation thinking (Weick and Sutcliffe, 2007) identifies as the prerequisite for the first HRO principle, preoccupation with failure. The other four HRO principles, namely reluctance to simplify, sensitivity to operations, commitment regarding resilience and deference to expertise, are similarly served by a method that brings frontline judgement into the assurance picture as a primary data source rather than as a post-incident artefact.

Case Study: Mobile Equipment And Driving Safety

The CEA methodology was developed and field-tested through a critical-risk control effectiveness assessment of a driving-safety portfolio at a large Australian asset-intensive operator with a fleet of light vehicles and field-services equipment exposed to public-road driving, regional travel and active-site environments. The reference case is non-mining, but the structural findings are directly translatable to the controls that dominate Queensland mining MUEs in mobile equipment and traffic management. The case is reported here at a level of generality consistent with confidentiality undertakings; specific scores, control descriptors and frontline findings are drawn from the operator's assessment report.

Material unwanted event and controls in scope

The MUE in scope was defined as the activity of driving a company vehicle beyond 400 km or for more than 4 hours, with collision or rollover as the unwanted event and severe to extreme consequences, including multiple fatalities or serious injuries. The hazard category was kinetic energy from a moving motor vehicle. The untreated risk was rated as a severe consequence, almost certain likelihood and 'critical' on the operator's risk matrix.

Six controls were assessed: ANCAP 5-star vehicle selection; fit-for-purpose vehicle specification; GPS driver-monitoring system; journey management system; driver competency system; and fit-for-work and fatigue management. Each control was first classified using a flowchart that distinguishes a hard control (an object or system that interrupts the energy pathway) from a control support (a human act, procedure, training package, or policy that helps a hard control function). Hard controls that were essential to preventing an MUE or mitigating its consequences were classified as critical controls; the remainder were treated as standard controls or as control support.

Workshop sequence and scoring scale

Three workshops were run: a critical risk and critical control current-state evaluation with risk owners and SMEs (workshop 1); an SME control effectiveness evaluation focused on Robustness (workshop 2); and a control-user workshop with frontline drivers and field operators focused on Human Dependency, Alignment and Adaptability (workshop 3). Sessions were two hours each and were delivered live, supported by structured pre-reads. Worker voice was protected through learning-team principles: non-punitive framing, anonymous score capture and explicit invitations to surface workarounds.

Each control was scored 1-5 on each of the four dimensions and aggregated into a composite reliability score. The composite was banded against five reliability descriptors, summarised in Table 2.

TABLE 2

CEA composite score bands and reliability descriptors used in the assessment.

Composite score (range)	Reliability descriptor
Below 1.5	No control
1.5 to 2.0	Unreliable control
2.0 to 3.0	Somewhat reliable control
3.0 to 4.0	Reliable control
4.0 to 5.0	High reliability control

Per-Control Reliability Profile

The four-dimensional sub-scores and composites for the six controls are shown in Table 3. The sub-scores are deliberately kept visible so that the assessment supports a remediation discussion rather than a single rating.

TABLE 3

Reliability profile for the six driving-safety controls assessed in the case study. Each sub-score is on a 1-5 scale; the composite is banded against the descriptors in Table 2. R = Robustness, HD = Human Dependency, A = Alignment, Ad = Adaptability.

Control	Type	R	HD	A	Ad	Composite	Descriptor
ANCAP 5-star vehicle	Critical (separation)	4.1	3.8	3.6	3.5	3.7	Reliable
Fit-for-purpose vehicle	Critical (separation)	3.8	3.7	3.7	3.2	3.6	Reliable
GPS driver-monitoring system	Critical (engineering)	3.5	3.3	2.7	2.6	3.0	Somewhat reliable
Journey management system	Standard (engineering)	2.8	2.9	3.0	2.8	2.9	Somewhat reliable
Driver competency system	Standard (administrative)	2.3	3.5	3.5	2.8	3.0	Somewhat reliable
Fit-for-work and fatigue management	Standard (administrative)	2.2	2.6	3.5	2.9	2.8	Somewhat reliable

Findings: where the binary frame would have missed the failure modes

The two highest-scoring controls were the engineering controls at the top of the hierarchy: vehicle separation using ANCAP 5-star selection (3.7) and fit-for-purpose vehicle specification (3.6). Both were rated robust, with low human reliance and good alignment. Both, nonetheless, revealed a structural coverage gap: neither control applied to the

operator's grey fleet (personal vehicles used for work purposes) or to delivery-partner vehicles, which, between them, accounted for a significant share of the actual driving exposure. A binary review would have ticked these controls as 'in place'; the four-dimensional review made the coverage gap visible as a specific, costable remediation pathway.

The fit-for-purpose vehicle control illustrated the value of separating the SME view from the frontline view. SMEs scored Robustness at 3.8 against a documented vehicle-specification standard. Frontline staff scored Adaptability at 3.2, the lowest of the four dimensions, for a different reason. Vehicles were perceived to be 'rarely fitted out for the work required', modifications were difficult to obtain, a 'one size fits all' approach was reported, and, in some cases, navigation features had been disabled because subscriptions had lapsed. The control was robust in design but degraded in adaptation to actual task demands.

The GPS driver-monitoring system was the clearest illustration of alignment failure as a discrete reliability mode. Robustness scored 3.5: the technology functioned, the data were captured, and the analytics worked. Alignment scored only 2.7 ('poorly aligned'). Frontline staff reported that the data were perceived to be used primarily for performance management rather than for safety, that drivers were 'focusing on the speedometer' rather than on the road, that technical issues and inconsistent application undermined trust, and that the system did not extend to grey-fleet or delivery-partner vehicles where much of the actual exposure sat. The composite settled at 3.0 ('somewhat reliable'). A high-Robustness, high-technology control was eroded by alignment failure that no compliance audit on the technology itself would surface.

The journey management system, which required journey management plans for any trip exceeding 400 km or 4 hours, was scored 2.9 ('somewhat reliable'). The intent was sound, the policy was documented, and SMEs scored the system credibly. Field staff, however, described the supporting application as 'clunky' and difficult to use, notably in remote locations or where diversions were required. The plan-completion count was the available verification metric; it indicated how many plans had been submitted, not whether they had reduced fatigue exposure or been adapted when the journey changed. The control was viewed as a burden rather than as a support, and the verification data did not yet integrate with telematics in a way that could test whether plans matched journeys actually undertaken.

The fit-for-work and fatigue management control received the lowest composite score, 2.8 ('somewhat reliable'), with a Robustness sub-score of just 2.2 ('fragile'). The intent was sound, and the policy framework existed, but the control depended heavily on individual self-assessment and self-management. There was no anticipatory monitoring or technology-supported fatigue indicator; external demands (management expectations, deadlines) and the realities of arduous work and living conditions in remote areas were reported to undermine policy compliance. Drug and alcohol testing supported a related domain but did not address fatigue. The translation to mining is direct: rosters, shift-end fatigue management, and self-assessment-based fitness-for-work tools share the same fragility pattern, and they tend to be the controls that lean hardest on production pressure on the wrong side of the variability curve.

The driver competency system was scored 3.0 ('somewhat reliable'), but with the lowest Robustness sub-score after fatigue management (2.3, 'fragile'). A licence check confirmed compliance with a system state (the driver held a valid licence), but did not test whether the driver was competent to perform the actual driving task, particularly off-bitumen or on long regional drives. Training profiles were managed manually by leaders, with no automated trigger when a vehicle was allocated, and access to practical training such as four-wheel-drive courses, was reported as constrained by cost. Accountability was perceived as reactive: it was activated after incidents rather than before them. The structural lesson is generalisable to any control where a state check is mistaken for a capability check.

Two cross-cutting observations followed from the case. First, the assessment identified three coverage gaps (grey fleet, delivery partners, and manual training profile management) that did not appear in the bowtie at all. The binary frame had not been wrong about the controls in scope; it had been silent about the population not in scope. Second, the available verification metrics (plan counts, licence validity reports, drug and alcohol test rates, telemetry uptime) were strong on system state and weak on safety outcome. Calibrating those metrics against exposure (kilometres driven, journey-plan adherence to the actual route, fatigue-relevant variables such as time of day and trip duration) was identified as a remediation pathway in its own right.

Hard Controls And Control Support: A Structural Separation

The case study also identified a structural finding that directly translates into mining critical-risk portfolios: hard controls and control supports must be kept separate in the assessment, as confusing them leads to investment misallocation. In the driving-safety case, the engineering controls at the top of the hierarchy (vehicle separation via ANCAP 5-star selection, fit-for-purpose vehicle specification and the technical layer of the GPS system) were the hard controls. Training, supervision, journey-management policy, fitness-for-work assessment and competency administration were treated as control support rather than as critical controls in their own right. That separation, anchored on the EDM lens, is structural in CEA. It is not always present in conventional bowties, in which training packages and supervisory routines are commonly labelled as critical controls.

The same architecture applies in mining to fail-safe braking, interlocks, geotechnical support, ventilation, gas-monitoring engineering and collision-avoidance technology: these are the hard controls; the procedures, training packages, supervision routines and permits around them are control supports. A defensible reliability profile for a mining MUE has to make that distinction explicit and direct investment accordingly.

Digital Indicators As Predictive Evidence

The case demonstrated the value of integrating digital indicators into the reliability picture as predictive evidence rather than as compliance reporting. Overspeed counts, harsh-braking and harsh-cornering events, telemetry coverage gaps, journey-pattern data, journey-plan-completion counts and licence-validity reports were available in the operator's safety dashboard. Read alongside the four-dimensional scoring, those data became actionable. Falling alignment scores on the GPS system, for example, were reflected in known device-disabling patterns in coverage-gap regions. Plan-completion counts that did not match the telematics-derived journey patterns suggested that journeys were being re-planned in the field without the policy framework being followed. The dashboard moved from a reporting artefact to a control-health instrument.

Translating CEA to Mining Critical Risks

Five mining critical-risk areas show structural similarity to the case-study findings and benefit visibly from CEA assessment.

Mobile plant and traffic management controls (procedures, spotters, segregation, exclusion zones) depend heavily on how people manage congestion, night shifts and contractor variability. The fatigue-management and alignment failure modes from the case study are directly relevant. CEA scoring usually exposes that the engineering layer (separation, fail-safe braking, collision avoidance) is the hard control and the procedural layer is control support, rather than the reverse.

Geotechnical and strata controls in underground operations and open-cut highwalls often look strong on the PHMP but are routinely worked around to meet production targets when ground conditions vary. Adaptability scoring, anchored on abnormal-condition

walkthroughs, is where these workarounds become visible. The structural lesson from the case study is the same: a control that is adaptable for the design envelope is not necessarily adaptable when the envelope is breached.

Isolation and lockout/tagout controls during shutdowns are typically robust on paper but degrade under time pressure, congestion, and contractor variability. Human Dependency and Adaptability are the dimensions on which they tend to fail. The case-study fatigue findings translate directly to shutdown crew fatigue and to handover dynamics during round-the-clock work, and the case-study competency findings translate directly to the difference between a permit being signed and the work being competently performed.

Gas monitoring and ventilation controls are typically robust at the sensing layer but fragile at the alarm-response, escalation and shift-handover layers. The reliability profile usually exposes the gap between sensor performance (Robustness) and human response (Human Dependency, Adaptability), in a pattern similar to the case-study GPS system's gap between technological capability and frontline alignment.

Collision avoidance and proximity detection systems, increasingly mature in surface and underground mining, are exposed to the same alignment failure mode as the case-study GPS system. A high Robustness score combined with a low Alignment score is a strong indicator that the technology is reliable but its safety value is being eroded through gaming or workaround behaviour.

In each of these areas, the value of CEA is not that it identifies risks the operator does not know about. It is that it produces a structured reliability profile for each critical control, anchored on evidence and traceable to the workforce that does the work, in a form that supports remediation prioritisation and that defends a 'reasonably practicable' decision.

Reasonably Practicable Evidence

The post-1-June compliance regime asks operators to demonstrate not just that critical controls exist, but that they are reliable and that the operator has acted on fragility when it has been found. The phrase 'reasonably practicable' carries the legal weight of that demonstration in Australian work health and safety law.

A CEA reliability profile produces direct evidence for that demonstration. The evidence trail reads: critical controls were identified using EDM and the hierarchy of controls; reliability was scored across four dimensions by SMEs and frontline staff using a structured method; divergences between the two views were interrogated and resolved; fragility findings were linked to specific remediation pathways; and remediation decisions were made on the basis of the reliability profile. That story, supported by workshop minutes, scoring sheets, narrative summaries and a traceability matrix, is the form of evidence the post-1-June framework is built to receive. It also satisfies the parallel governance need: a board safety committee can read the reliability profile, see where the fragilities are, see how they have been addressed and form a defensible view of residual risk.

Australian mining is not the first sector to confront the relationship between operational variability, cognitive load and safety outcomes. Aviation, nuclear power, healthcare, oil and gas and petrochemical industries have each developed regulatory systems that explicitly account for protective capacity and operator cognitive load: mandatory staffing ratios in air traffic control, technical specifications for redundancy in nuclear power, and alarm-management standards (eg ISA-18.2) in petrochemical operations. The shift to risk-based CCM in Australian mining belongs in that tradition. The CEA frame is consistent with the underlying logic of those regulatory architectures: that the reliability of a control under everyday operating conditions, not its presence on paper, protects against the unwanted event.

CEA is deliberately not a one-off audit instrument. The reliability profile is meant to be re-scored at intervals, with the digital indicators tracked between assessments, so that the operator can see drift in control health rather than discovering it at the next incident.

Conclusions

The shift from compliance auditing to risk-based critical control management asks operators a harder question than 'is the control present?'. It asks whether each control is reliable under the conditions in which it has to function. Binary reviews cannot answer that question and, when relied on as the primary assurance method, can mask the failure modes that matter most.

Control Effectiveness Assessment combines the Energy Damage Model, the hierarchy of controls and human-centred design within a four-dimensional scoring frame (Robustness, Human Dependency, Alignment and Adaptability) that triangulates SME and frontline evidence into a structured reliability profile. Repeated over time, the same architecture turns the workforce into a distributed sensor network for control health and exposes drift as a leading indicator of control failure. The case study shows that controls that look robust on paper can be fragile in the field due to human dependency, alignment failure, or adaptability gaps, and that high-technology controls in particular can lose their safety value through alignment failure even when the engineering is sound. Coverage gaps that the binary frame does not surface, namely populations of work that fall outside the controls in scope, are routinely exposed by the four-dimensional review. The translation of mining critical risks into mobile plant, geotechnical control, isolation, gas monitoring, and collision avoidance is direct.

A note on the wider operating context is in order. Where mining operations run with chronic firefighting and supervisory cognitive overload, the strongest reliability profile will not survive long contact with production. The Brady (2020) finding that, after Serious Accidents, 62 per cent of corrective actions in Queensland were administrative controls only is, in part, a finding about that wider platform: when the operation lacks the bandwidth to plan and execute engineering interventions, the path of least resistance leads back down the hierarchy. CEA helps an operator see where its controls are fragile; the complementary work of stabilising the operating platform sits alongside the assurance frame, not inside it.

For Queensland operators preparing for the post-1-June 2026 compliance test, and for Australian mining operators more generally adopting risk-based CCM, the practical implication is the same: assurance effort needs to focus on reliability rather than presence, on work-as-done rather than work-as-imagined, and on a structured frame that produces defensible 'reasonably practicable' evidence. CEA offers one such frame.

Acknowledgements

The Control Effectiveness Assessment methodology described in this paper is the intellectual property of The Art of Work and was developed and field-tested through a critical-risk driving-safety engagement with a large Australian asset-intensive operator. The author thanks the SMEs and front-line workers whose participation made the case study possible, and the Art of Work principals whose facilitation produced the four-dimensional reliability frame. This paper was drafted using a large language model (Anthropic Claude) for outlining and copy-editing assistance, in accordance with the AusIMM Guide to Authors disclosure expectations; all analytical content, case material, structural arguments, and conclusions are the author's original work.

References

- Brady, S, 2020. Review of all fatal accidents in Queensland mines and quarries from 2000 to 2019 (Resources Safety and Health Queensland: Brisbane).
- Conklin, T, 2019. The 5 Principles of Human Performance: A Contemporary Update of the Building Blocks of Human Performance for the New View of Safety (Pre-Accident Media: Santa Fe).
- Dekker, S, 2014. Safety Differently: Human Factors for a New Era, second edition (CRC Press: Boca Raton).
- Haddon, W, 1973. Energy damage and the ten countermeasure strategies, Journal of Trauma, 13(4):321-331.
- Hollnagel, E, 2014. Safety-I and Safety-II: The Past and Future of Safety Management (Ashgate: Farnham).
- Hopkins, A, 2019. Organising for Safety: How Structure Creates Culture (CCH Australia: Sydney).
- International Council on Mining and Metals (ICMM), 2015. Health and Safety Critical Control Management - Good Practice Guide [online]. Available from: <<https://www.icmm.com>> [Accessed: 15 April 2026].
- Manuele, F A, 2005. Risk assessment and hierarchies of control, Professional Safety, 50(5):33-39.
- National Institute for Occupational Safety and Health (NIOSH), 2015. Hierarchy of Controls [online]. Available from: <<https://www.cdc.gov/niosh/topics/hierarchy/>> [Accessed: 15 April 2026].
- Norman, D A, 2013. The Design of Everyday Things, revised and expanded edition (Basic Books: New York).
- Resources Safety and Health Queensland (RSHQ), 2024. Queensland Guidance Note QGN35: Principal Hazard Management Plans and Critical Control Management for Coal Mines [online]. Available from: <<https://www.rshq.qld.gov.au>> [Accessed: 15 April 2026].
- State of Queensland, 2024. Resources Safety and Health Legislation Amendment Act 2024 (Qld) [online]. Available from: <<https://www.legislation.qld.gov.au>> [Accessed: 15 April 2026].
- Viner, D, 1991. Accident Analysis and Risk Control (Derek Viner Pty Ltd: Melbourne).
- Weick, K E and Sutcliffe, K M, 2007. Managing the Unexpected: Resilient Performance in an Age of Uncertainty, second edition (Jossey-Bass: San Francisco).