EON AI Ventures

The Mining Situation Room

From incident evidence to role-relevant operational learning

Table of Contents

Executive summary

The Mining Situation Room is a proposed digital front door for learning about high-consequence mining events and dangerous operating moments. Its purpose is not to turn tragedy into spectacle, and it is not a maintenance course. It is an operation-focused experience spanning surface pits, underground mines, and processing and tailings facilities.

The concept begins with a simple problem: important safety knowledge is often distributed across investigation reports, regulator alerts, standards, procedures, equipment documentation, and the experience of people who have worked through abnormal conditions. Each source can be valuable, but the learner still has to connect an incident to a place, a machine, a role, a decision, and an action. The Situation Room is intended to make those connections visible without overstating what the evidence proves.

Phase 1 is a bounded 30-card pilot connected to an existing Mining Living Library of 50 equipment classes and 350 mapped subcomponents. Those figures describe the planned product scope and the supplied BuildPack; they are not safety-performance statistics. Each publishable card must pass two separate gates: an adversarial evidence review and a safety, dignity, and media review.

Seven learning elements are represented on each card: EON Life, Ask Brainy, Course Material, the EON XR Simulator, multiple-choice questions, an oral exam, and certification. In Phase 1 these are clearly marked as coming online unless the responsible system and owner have actually delivered them. The cards create a coherent route into those experiences; they do not pretend that unfinished simulations, assessments, or credentials already exist.

1. Why a Mining Situation Room

Mining safety learning has to bridge two kinds of knowledge. Explicit knowledge is written down: the investigation finding, the standard, the procedure, the equipment specification, or the regulator's warning. Tacit knowledge is more situational: what an experienced operator notices first, which weak signal changes the decision, when normal work begins to feel wrong, and what a supervisor asks before allowing work to continue.

The Mining Situation Room is designed around that bridge. A learner first encounters a specific situation in a credible mining setting. The card then routes the learner to the evidence, the relevant equipment or site system, role-specific implications, and future practice or assessment experiences. The intended learning question is not merely "What happened?" It is "What would I need to notice, verify, communicate, or stop before this situation develops further?"

The product frame is operation plus catastrophes. Landmark events provide an evidence-rich view of system failure and consequence. Recurring operating patterns provide the everyday moments in which recognition and intervention matter. Maintenance and stored energy appear where relevant, but they remain one cross-cutting shelf rather than defining the whole product.

2. Three mining worlds, two routing dimensions

The catalog is organized first by world and then by shelf.

World answers: Where am I?

- Surface pit
- Underground
- Processing & Tailings

Shelf answers: What kind of danger is this?

Surface shelves include haul and mobile equipment, drilling and blasting, pit-wall and ground conditions, and dust, visibility, and roads. Underground shelves include ground control, gas, fire and explosion, ventilation and inrush, and underground mobile equipment. Processing and tailings shelves include crushing and milling, conveyors and material handling, slurry and pumping, and tailings facilities. Cross-cutting shelves cover maintenance and stored energy, electrical isolation, shift handover, and emergency response.

A third layer of controlled human-factor tags answers: Why might people miss the danger? Candidate tags include fatigue, complacency, line of fire, blind spot, isolation not verified, normalized deviance, poor visibility, communication gap, production pressure, and handover gap. These tags support navigation and reflection; they are not diagnoses of individuals or unsupported conclusions about a named event.

This architecture prevents category errors. A tailings facility is treated as a site system, not misrepresented as a machine. A haul-truck event can link directly to a truck class. A highwall event can be routed through a geotechnical site hazard while also linking to relevant supporting equipment. Cards may cross worlds where the evidence requires it.

3. The 30-card pilot

The pilot consists of candidate anchors for cited corpus development. They are not yet a claim that these are the 30 most frequent or most severe mining events. Final titles, causal wording, rankings, standards, and named organizations must be frozen only after primary-source research and adversarial review.

Landmark catastrophes and flagship cases

1. Brumadinho B1 Tailings-Dam Failure
2. Soma Underground Mine Fire
3. Upper Big Branch Methane and Coal-Dust Explosion
4. Pike River Methane Explosion
5. Fundão and Mariana Tailings Failure
6. Buffalo Creek Impoundment Failure
7. Crandall Canyon Ground Collapse

8. Sunshine Mine Fire
9. Underground Inrush and Inundation — final anchor case to be selected from an official report
10. Open-Pit Highwall Collapse — final anchor case to be selected from an official investigation

Recurring operating-pattern candidates

1. Haul Truck and Light Vehicle Blind-Spot Collision
2. Haul Truck Runaway on a Downgrade
3. Truck Reverses Over a Dump Edge
4. Conveyor Entanglement During Operation or Cleaning
5. Caught in Moving Crusher or Feeder Machinery
6. Fall of Roof or Back During Active Work
7. Fall of Face, Rib, Side or Highwall
8. LHD or Underground Truck Strikes a Pedestrian
9. Unexpected Machinery Movement During Work
10. Electrical Contact or Arc at Plant Equipment

World-balancing operational candidates

1. Blast Misfire or Flyrock
2. Dozer Too Close to a Highwall or Dump Edge
3. Water Cart Loses Control on a Haul Road
4. Loss of Primary Ventilation
5. Roof Support Installation Under Changing Ground
6. Underground Vehicle Fire and Evacuation
7. Crusher Blockage Cleared While Feed Can Restart
8. SAG Mill Reline Stored-Energy Event
9. Slurry Line or Pump Release Under Pressure
10. Tailings Water-Balance Warning Ignored

The first group brings official investigations and inquiries into a dignified learning context. The second group is intended to become evidence-based pattern content, but no “most frequent” rank may be published until a reproducible aggregate query stores its numerator, denominator, time period, mine population, and accident classification. The third group ensures balanced operational coverage across the three worlds. Pattern cards must be labeled as patterns; placeholder flagship cases must not ship.

4. Personas and content routing

Six personas provide a transparent v1 routing layer:

- Haul-truck operator — runaway, collision, blind spots, dump edges, roads, and visibility
- Underground miner — roof and rib, ventilation, gas, inrush, mobile interaction, and escape
- Plant operator — crushers, conveyors, mills, slurry, electrical upsets, and process indications
- Shift supervisor — handover, emergency decisions, production pressure, permits, and exclusion-zone control
- Maintenance / LOTO — electrical, hydraulic, gravity, mechanical, and pressure energy
- Geotechnical / TSF engineer — highwall movement, monitoring, water balance, tailings governance, and emergency response

Persona routing changes sequence and emphasis, not the underlying truth. The same card can surface different questions for different roles. A haul-truck runaway card might first ask an operator about grade, speed, braking response, and an escape ramp. For a supervisor, it might foreground traffic rules, road condition, dispatch communication, and emergency control. For maintenance / LOTO, it might route to inspection and stored-energy material only when the evidence makes that relevant.

The initial demonstration may default to the haul-truck operator persona. That default is a presentation choice, not an assertion that one role matters more than another. Future personalization should remain explainable: users should be able to see why a card was recommended.

5. Ask Brainy as a tacit-knowledge layer

Ask Brainy should not be described only as a test of remembered facts. Its more valuable role is to help articulate the practical knowledge surrounding an incident or operating pattern.

For each card, Ask Brainy can be briefed on the cited case, the linked equipment class, the applicable standard or guidance, and clearly attributed practitioner prompts. It can help a learner explore questions such as:

- What weak signals might appear before the obvious alarm?
- What would an experienced operator verify before continuing?
- Which observation would change this from “monitor” to “stop work”?
- What information must be handed over to the next shift?
- Where can normal production pressure distort judgment?
- Which answer comes from the official record, and which is a site-dependent practice that must be verified locally?

This is tacit knowledge made discussable, not tacit knowledge made infallible. Ask Brainy must distinguish sourced facts, standard-based controls, reasoned operational prompts, and local procedures. It must not invent eyewitness experience, speak as though it personally worked at the incident, replace competent supervision, or give site-specific instructions without the

governing local procedure. Every important factual answer should route back to a source or explicitly state its limitation.

6. The seven learning elements on a card

Each Situation Room card is a hub with seven visible destinations:

1. **EON Life** — a future playable mission for practicing the dangerous moment. The Coordinator owns playable missions; Phase 1 may carry a future link and promise copy only.
2. **Ask Brainy** — a cited conversational layer for evidence, equipment context, and tacit operational questions.
3. **Course Material** — a standards-mapped deep dive containing cited reading, diagrams, procedures, or role-specific explanation.
4. **EON XR Simulator** — a represented path to a future equipment or scenario simulation. It must not be described as operational until the responsible product has delivered and verified it.
5. **MCQ** — a future multiple-choice knowledge check. No scores, pass thresholds, or validated assessment claims appear until an assessment contract exists.
6. **Oral Exam** — a future spoken assessment experience. Its rubric, recording rules, privacy handling, and human-review policy must be defined before use.
7. **Certification** — a future credential destination. The card cannot claim certification, accreditation, or training compliance until an authorized credential owner has approved the program and its assessment basis.

During Phase 1, unfinished elements are visibly labeled “Coming online.” This honest representation lets the card show the intended learning journey without fabricating functionality.

7. Evidence, honesty, and adversarial review

Every publishable card begins with an appropriate primary source: an official investigation, regulator alert or dataset, standard, royal commission, court or inquiry finding, or equivalent authority. Secondary sources may help locate primary evidence but cannot anchor a publishable fact.

The evidence model separates incident confidence from frequency confidence. An investigation can verify what happened in one event; it does not prove how often the pattern occurs. A standard can support an expected control; it does not establish that a particular omission caused an event unless the source says so. When no public quantification exists, the product says “not quantified” rather than estimating a rate.

Each important number must carry a URL and a locator such as a page, table, paragraph, dataset field, or reproducible query. A separate adversarial reviewer attempts to disprove the draft: wrong event, wrong toll, overstated causal language, incorrect standard, unsupported frequency, excessive attribution to an operator, or internal inconsistency. The card is revised until it passes or is removed.

A second gate covers safety, dignity, legal risk, and generated media. It rejects victims, bodies, gore, identifiable affected people, distress spectacle, operator logos, unsupported blame, misleading simulation promises, and mechanically implausible imagery. Real catastrophes are shown through terrain, infrastructure, machinery, monitoring conditions, empty space, or restrained aftermath. Technical accuracy does not excuse exploitative treatment; dignity is tested separately.

Every card should expose its confidence and provide a path to the official source. Until validation is complete, generated representations carry a “plausible — not certified” status.

8. The Mining Living Library as the spine

The Situation Room is not a disconnected gallery of incidents. Its equipment spine is the supplied Mining Living Library BuildPack: 50 authoritative equipment classes, organized across ten equipment families, with seven mapped Level-2 subcomponents per class for a total of 350 mapped subcomponents.

That spine gives equipment cards a stable identity. A haul-truck event can link to the rigid-frame haul-truck class. A conveyor event can link to the belt conveyor. A ventilation-loss card can link to the primary ventilation fan. A slurry-release card can link to a centrifugal slurry pump. Site hazards such as tailings governance or highwall movement remain site hazards, with supporting equipment linked only where factually appropriate.

The value of this connection is consistency. The same equipment vocabulary can support the card catalog, explanatory material, future XR objects, and future EON Life missions. It also creates a publish blocker: an equipment class not present in the authoritative BuildPack cannot quietly enter the product under an improvised name.

9. Roadmap and gates

The proposed roadmap is deliberately staged:

- Freeze the Mining schema, controlled vocabulary, and compiler rules.
- Build three end-to-end cited sample cards: a tailings catastrophe, a haul-truck operating moment, and a conveyor/isolation case.
- Run adversarial evidence and dignity reviews before art direction is scaled.
- Test a small media batch, then generate approved posters, locked-camera clips, and narration within the authorized budget ceiling.
- Build and dual-gate the full 30-card corpus.

- Assemble the local storefront with filtering, shelves, deep links, Briefing content, and one-at-a-time media playback.
- Review locally before any separate deployment decision.
- Hand any playable EON Life mission specification to the Coordinator; do not implement it inside the Situation Room lane.

Media follows a restrained doctrine: human-free cinematic mine environments, mechanically plausible equipment, no branding, and no sensational disaster treatment. Motion uses locked cameras and narrowly permitted environmental or machine-native changes so that generated equipment does not melt, rotate, or change geometry.

10. Limitations and non-claims

The Mining Situation Room is a proposed learning and communication product. It is not a substitute for site-specific risk assessment, competent supervision, emergency plans, statutory training, manufacturer instructions, or applicable law.

The 30-card list is a pilot candidate set, not a universal ranking of mining risk. Two flagship slots still require selection of a specific officially investigated case. “Top 10 — Most Frequent” must remain unpublished until its aggregate basis is reproducible. Incident facts and frequency claims require different evidence.

Generated images and clips are illustrative representations. They are subject to mechanical-fidelity, safety, dignity, and branding review. They do not reconstruct an incident with forensic certainty.

Ask Brainy can organize cited information and prompt operational reflection, but it cannot possess personal experience, determine site conditions, authorize work, or replace local procedures. Course material, XR simulation, MCQ, oral examination, EON Life missions, and certification remain future or external capabilities until their owners deliver and validate them.

No statement in this draft claims a reduction in incident rates, a financial return, regulatory acceptance, accreditation, certification validity, or market leadership. Those claims would require separate evidence and authorization.

Source basis for this draft

This draft is based on the approved Mining Situation Room specification, the supplied Mining Living Library BuildPack, the Mining section of the EON product-suite material, and the specification’s source register. That register includes MSHA accident classifications, fatality reports and investigation materials; the Global Industry Standard on Tailings Management; Brazilian ANA material; the Upper Big Branch investigation record; and the Pike River Royal Commission. Each finished card requires its own source-level verification and locator before publication.