OPERATIONAL TECHNOLOGY · MINING · AI READINESS

The Floor That Now Has to Think

Why Operational Technology Must Arrive First at the AI Conversation in Mining

A framework for the OT architects, OT managers, CTOs, and VPs of operations who will live with the consequences of the next AI proposal in LATAM mining

AUDIENCE
OT leaders, architects, CTOs, VPs Ops

FOCUS
OT-led conditions for mining AI

REGION
LATAM Mining

NTT DATA · Natural Resources Practice · 2026 · biztalksnttdata.com

Contents

01

THIRTY YEARS SILENT, NOW ASKED TO THINK

Seven of ten mining executives in Latin America, asked about the autonomy and AI projects their operations have already run, acknowledge the same thing: the projects only partially achieved the value objectives they were approved against. The cause is almost never the model. It is the floor the model was asked to run on. That is the claim this paper defends, and it is the claim most AI proposals arriving at a LATAM mine in 2026 are structured to avoid.

For thirty years, operational technology (OT)[1] has held the mine together. Distributed control systems[2] running the concentrator. Programmable logic controllers[3] running the crushing line. A process historian logging every minute of every shift into an archive that the corporate side has not opened in fifteen years. The floor sustained the mine. It was reliable, undocumented in business terms, and largely invisible to the CIO and the CFO.

Today the CFO is asking when the floor becomes intelligent.

The question is not academic, and it does not arrive empty-handed. It arrives as a deck full of vendor logos and a project timeline that assumes the operation already has the substrate to support autonomous decisioning. The integrator promises sub-second optimization on flotation recovery. The corporate CIO endorses the proposal because the board asked for an AI answer this quarter. The OT manager is invited to the third meeting, after the architecture has been drawn, to confirm that the historian “exposes an API.”

This is the contradiction the paper opens with. The floor that has sustained the mine for thirty years is now expected to host the most consequential technology decision of the next decade, and the people who built and maintain that floor are arriving to the conversation second. The cost of that sequencing is not abstract. NTT DATA’s research with MIT Technology Review on autonomy in the mining business — a 2025 survey of mining executives across Chile, Brazil, Peru, and Mexico — found that 72% acknowledge their projects only partially achieve the defined value objectives. The figure is directional, not surgical, and §04 treats its limits honestly, including the fact that it is our own research. The direction is not in doubt: most projects deliver less than they promised, and when the projects are opened up, the shortfall clusters below the model, in the floor. Call that population the seven-in-ten: the projects that delivered only part of what they promised.

Two adjacent arguments will try to absorb this argument, and both should be refused at the door.

---

[1] OT — Operational Technology. The hardware, software, and protocols that control physical industrial processes, including PLCs, DCSs, SCADA systems, and the industrial network connecting them.

[2] DCS — Distributed Control System. The principal control architecture for continuous process operations such as concentrators and refineries.

[3] PLC — Programmable Logic Controller. The principal control device for discrete and hybrid industrial automation tasks, including crushing, conveying, and materials handling.

The first is the question of whether AI belongs in mining. It does, and it already arrived; autonomous haulage settled that debate. The live question is not whether but *who defines the conditions* under which AI captures value, and on whose timeline.

The second is the IT/OT convergence narrative. Convergence describes a destination where both estates meet in the middle. This paper describes who arrives at that destination holding the agenda. The distinction is not semantic. Convergence-led means the party that brought the framework sets the terms, and in the operations the practice has reviewed, the party with the framework was IT. OT-led means OT defines the preconditions the AI must satisfy, in writing, before the architecture is drawn. A mine can be fully converged and still join the seven-in-ten, because convergence is an org-chart event and the failure is an architecture event. None of this is an argument against convergence. It is the argument about what OT must carry into convergence, so that the meeting in the middle happens between peers rather than by absorption. A companion whitepaper in this series makes the case for the convergence destination; this paper is about what OT must bring on the way there.

The argument advanced here is structural. OT cannot wait for an AI proposal to arrive and then review it. It must arrive at the conversation first, carrying the framework that defines what "ready" means. Absent that framework, the conversation is authored by parties who do not own the failure surface, and the operation joins the seven-in-ten on schedule.

What follows is that framework, the three tensions it exists to resolve, an honest accounting of what the intelligent floor actually costs to stand up, and a test the OT manager can run before signing anything.

02

THE MORNING THREE VENDORS ARGUED AND NOTHING WAS FIXED UNTIL LUNCH

Make it real before making it general. A copper concentrator on the LATAM Pacific belt, day shift, a Wednesday in late October. At 10:42 the dispatch system stops receiving truck-position updates. The plant does not stop. It continues to run on the last known state of the mine, which is to say, on a picture of reality that is now minutes old and getting older. Setpoints that assume a feed blend from a fleet that has since moved. The concentrator keeps turning ore that the dispatch layer can no longer see.

The plant manager calls the OT control room. The OT lead checks the industrial network first, because that is where the OT lead always checks first. The fiber backbone is up. The microwave link to the remote dispatch node shows green. The radios are alive. The data is not arriving.

The cause is one layer up and invisible from the network view. The gateway that translates haul-truck telemetry from the field protocol into the dispatch system's message format has stopped acknowledging packets. It is not down. It is accepting connections and silently dropping payloads, which is the worst failure mode an OT architect knows, because every monitoring light stays green while the data stops.

The OT lead opens a ticket with the SCADA[4] vendor, because the gateway sits inside the SCADA contract's scope on the architecture diagram. The SCADA vendor inspects the link, finds it healthy, and routes the ticket to the network integrator, because "the payload loss is a transport problem." The network integrator measures the transport, finds zero packet loss on the link itself, and routes the ticket to the dispatch vendor, because "the data is leaving our boundary intact." The dispatch vendor confirms it is receiving nothing and routes the ticket back toward the SCADA gateway. Ninety minutes elapse. Three vendors, each one correct inside its own boundary, each one's correctness ending precisely at the contract line where the next vendor's begins.

The gateway is restarted at 12:15 by the OT lead, who finally power-cycles it on a hunch rather than on a diagnosis. By then the plant has run the better part of two hours blending against a stale mine model. The loss does not show up as a dramatic line on a dashboard. It shows up as a slightly worse recovery number that quarter, distributed invisibly across hundreds of truck cycles that went to the wrong stockpile.

No vendor was at fault. Each performed exactly to scope. The fault lived in the intersection of three contracts, and the intersection had no owner. The OT lead had, in principle, the mandate to act there. In practice the OT lead had a desk phone, three vendor support lines, a stopwatch, and no contractual authority to compel any of the three to treat a problem none of them owned as urgent.

Hold this scene. When the AI recovery-optimization proposal arrives the following month, promising sub-second setpoint adjustment on this exact circuit, the question that matters is not whether the model is good. The question is who resolves the next silent gateway failure inside the sub-second window the AI architecture assumes exists. The proposal will not answer that question, because no one who has ever placed that 10:42 phone call was in the room when the proposal was written.

This is the substrate the AI conversation actually lands on. Until the substrate is named with this much precision, the conversation is fiction with a budget.

03

WHAT A PLANT VETERAN WROTE ON THE BACK OF A NOTEBOOK

The case anchors here, with a person rather than a system. The VP of Operations of a major LATAM copper producer, thirty years in the wet circuit, came up through process control before "process control" reported to anyone with "digital" in their title. He was presented with an AI-driven recovery optimization proposal. The project was sponsored by the corporate CIO and authored by an integrator partnered with two AI platform vendors. The promised lift was expressed in points of recovery, as a recurring annual benefit in the millions of US dollars.

[4] SCADA — Supervisory Control and Data Acquisition. The supervisory layer that aggregates data from PLCs and DCSs and presents it to operators.

The VP asked three questions in the kickoff meeting, and the project lost its momentum in the time it took to ask them. The questions are reproduced here because they are the closest thing to a portable diagnostic this paper can hand a reader, and because every one of them is an architecture question wearing the clothes of a business question.

> "First. Who owns the data quality of the pH probes the agent is reading? Not the historian that stores the number. The physical calibration program in the field. Name the person and tell me when each probe was last calibrated.""Second. When the agent recommends a lime setpoint at three in the morning, and the operator follows it, and the bank goes off-spec, who sits on the safety review the following week? The operator who pushed the button? The integrator? The model? Show me the name on the line that says 'authorized.'""Third. What happens to this architecture when your contract ends in eighteen months? Does the value transfer to my team, or does it walk out the door with your engineers and their laptops?"
>
> — VP of Operations, copper producer, LATAM

None of the three could be answered in the room. The integrator promised to come back with answers, which is what integrators promise when a question has exposed something the proposal did not contain. The OT manager, brought in to validate the technical architecture, recognized all three questions instantly. They were the questions he had been unable to articulate in his own meetings, the unease he had carried into every AI conversation for two years and never managed to compress into language a CIO would take seriously.

He asked the VP for a working session. They spent an afternoon at a whiteboard, and the OT manager left with a one-page sketch on the back of a notebook: four conditions, listed top to bottom, with arrows showing what each one demanded from the one below it. Not four technologies. Four conditions that any technology would have to satisfy.

That sketch became the working draft of what NTT DATA's mining practice now calls PISO[5], the Platform for Intelligence on Operational Systems. The acronym is a deliberate pun. *Piso* is the Spanish word for floor, the same floor the executive summary of every PISO whitepaper opens with: *the floor has sustained the mine for thirty years; it is time for it to also think.* The sketch names what it needs to think with.

The VP did not authorize the original AI proposal. He authorized a foundation program against the four conditions, six months, with the AI project deliberately resequenced into the second year behind it. The decision cost the operation a quarter of "AI progress" it could report to the board, and the CIO was not pleased. Eighteen months later the operation deployed the recovery optimization at scale, on a substrate that could defend it. The promised lift materialized and held. The architecture survived the integrator's contract ending, because the foundation program had been written so that it would.

This is the recognition moment, and its lesson is not that frameworks are useful. It is that this framework was not invented in a research lab and handed down to operators. It was reverse-engineered from the questions a thirty-year veteran asks when the proposal lands

[5] PISO — Platform for Intelligence on Operational Systems. The NTT DATA Mining & Resources framework articulating the four codependent layers required for sustainable operational intelligence in a mining operation. The acronym is a deliberate pun on the Spanish word for floor.

on his desk and he is the one who will sit on the safety review when it goes wrong. The framework is OT's own knowledge, finally written in a form that survives the meeting.

04

WHY THE FAILURE LIVES BELOW THE MODEL

Step back from the case to the pattern, because one VP's good instincts are an anecdote and the claim of this paper is structural. The pattern recurs across the operations the practice has worked with over the past three years. The proposal arrives written in the language of the model: training-data volumes, inference latency, recovery uplift, optimization horizon. The operation accepts the proposal in the same language, because the model's language is the one the board memo was written in. Eighteen months later the project is deactivated or quietly degraded, and the post-mortem returns the same uncomfortable finding: the model performed exactly as specified on the inputs it was given. The inputs were the problem.

The failure was below the model. It was in the substrate.

A word about the seven-in-ten figure, because a careful OT reader is right to interrogate it before building anything on top of it. It comes from *Autonomy in the Mining Business: Beyond Technology, the Challenge for Talent and Culture*, research NTT DATA conducted with MIT Technology Review across mining operations in Chile, Brazil, Peru, and Mexico, published in late 2025. The finding, precisely stated: 72% of the executives surveyed acknowledge that their autonomy and AI projects only partially achieve the defined value objectives. Three honest limits attach to it. It is survey data — executives reporting on their own portfolios, not an audited measurement of project outcomes. "Partially achieve" is a threshold, not a physical constant; it counts the project that delivered half its business case alongside the one that delivered nearly all of it. And it is our own research: NTT DATA sells remedies of the kind this paper describes, and a reader should weight the figure the way §09 recommends weighting any proposing party's claim. None of this rescues the optimist. Independent cross-industry measurements land in the same territory: a 2025 study of more than 1,250 senior executives finds that around 60% of companies generate no material value from AI despite continued investment, and that only about 5% create substantial value at scale. Whether the true rate in mining is sixty percent or seventy-five, it is high.

The more durable finding in the study is not the percentage at all. It is what the same executives named when asked what blocks implementation. Resistance to change led the list, named by 27.5% of respondents. Lack of specialized talent followed at 14.5%. Technological reliability — the thing the proposals are written about — trailed at 11.6%. The executives running these projects are reporting, in their own words, that the dominant barriers are organizational, not technical. The discriminating variable is not in the model and not in the vendor. It is in the state of a small number of conditions in the operation *receiving* the project. The pattern the practice has seen in its own engagements is consistent with what the executives report: where those conditions were coherently active, the project landed; where one or more was absent, the project failed in ways that read as model failures on the dashboard but were not.

The phrase that carries the weight is **coherently active.** Not "present." Not "on the roadmap." Not "documented in the master systems list." Coherently active means three things at once: the condition is operational in the day-to-day, it has a single named owner with the authority to keep it that way, and it is observable, meaning someone can show you its state today without scheduling a workshop. A data-integrity program that exists as a slide is not coherently active. A governance committee that convenes quarterly is not coherently active. A calibration regime that was current in 2022 is not coherently active in 2026. The condition has to live in the operation, not in the deck.

This is the architectural claim PISO is built on, and it is the claim that decides whether the next project joins the seven-in-ten. The conditions that determine whether an AI project in a mine captures value are not in the model and not in the vendor. They are in the floor. Until the floor is described with the same precision the model is described with, the conversation cannot be productive, because the two sides are not discussing the same system.

Figure 1 The four PISO layers, stacked from the operational substrate up to the learning loop. Each layer depends on the one beneath it. Absence in any single layer propagates upward and disables the value of every layer above it, which is why a strong model on a weak substrate fails as reliably as a weak model does.

The next section names the conditions, and gives each one a hook the OT architect can carry into a vendor evaluation on Monday.

05

THE FLOOR AS A PLATFORM: THE FOUR LAYERS

PISO names four layers, ordered from the most foundational to the most agent-adjacent. Each is defined by what it requires to be coherently active, not by the product that implements it, because the product varies by operation and the condition does not. Each carries a vendor-evaluation hook: the specific thing to demand of any proposal that claims to address the layer, phrased so that a vague answer is visibly vague.

The four-layer count is inherited from the source framework — the PISO Pillar 01 whitepaper, NTT DATA Mining & Resources — rather than invented here; it is the canonical structure, not an authorial taxonomy.

Layer 1 — Data integrity of the operational substrate. Every reading the agent consumes must trace to a physical measurement with a current calibration record, a documented uncertainty band, and a fault-detection regime that distinguishes a real zero from a dead sensor reporting zero. Master data (equipment, material, location, process step) must be canonically identified across every system the agent reads, so that "Mill 3" in the historian, "ML-003" in the maintenance system, and "Molino 3" in the dispatch layer are provably the same machine. The substrate must not lie to the model, and when it has no answer it must say "no data," not "0." The scale of what operations do not know about their own substrate is routinely underestimated: when one industrial facility mapped its operational data flows in preparation for AI, it found 139 OT data connections, many of them previously unknown to the organization. *Vendor-evaluation hook: ask the proposing party to show the calibration coverage and the staleness distribution of the exact tag list the agent will read. If they have never seen your tag list, they are proposing against an imagined substrate.*

Layer 2 — Governance and mandate. Every AI-driven action must trace to a named decision-maker with written authority to authorize it at the relevant level of autonomy. Operator, supervisor, plant manager: the mandate is written before the project starts, not improvised when the bank goes off-spec at three in the morning. Safety, environmental, and ESG[6] accountabilities map to roles, not to systems, because a system cannot sit on a safety review. An AI agent without a governance mandate is a recommendation engine impersonating an action engine, and the impersonation is discovered at the worst possible moment. *Vendor-evaluation hook: ask who is named on the line that authorizes an autonomous setpoint change, and what happens to that authorization when the recommendation is wrong. If the answer is a system and not a person, Layer 2 is empty.*

Layer 3 — Sovereign connectivity and decision architecture. This is the layer the OT architect knows in the body, and the one AI proposals understand least. The operation must be able to act on its own data, at its own latency, without depending on a vendor cloud whose service-level agreement was written for a SaaS dashboard, not a control loop. "Fiber to the pit" is not a slogan; it is a bandwidth budget, a redundancy topology, and a latency figure under fault. The reference point is real and regional: BHP's Copper Advanced

[6] ESG — Environmental, Social, and Governance. The reporting and accountability regime increasingly requiring auditable lineage from physical measurement to regulatory submission.

Services centre in Santiago — opened in October 2024 at an investment of US$48.3 million — coordinates Spence and Escondida from more than a thousand kilometres away, ingesting some 5.4 terabytes of operational data a day over dedicated fibre between Antofagasta and Santiago. An architecture like that is not a connectivity purchase. It is a decision architecture: engineered fallback paths, and a latency budget per decision loop. The architecture question is not "is there connectivity." It is "what is the decision-loop latency from sensor to setpoint when the primary fibre is cut and traffic fails over to the secondary path, and again when the last-mile link drops at 4 AM in a storm." A vendor cloud that round-trips in ninety seconds cannot be the substrate for a control action that must close in one. Decision rights have to be architected per loop: the control loop acts locally and autonomously inside a bounded envelope, the supervisory loop reaches further, the planning loop can tolerate the cloud's latency because it does not run in seconds. *Vendor-evaluation hook: ask the proposing party to draw the data path for one setpoint decision and annotate it with latency under primary-link failure. If the drawing assumes the cloud is always reachable at low latency, the architecture has never met your pit.*

Layer 4 — Learning loop and continuous calibration. The agent must observe the corrections human operators apply to its recommendations and fold them back into its calibration. Without this loop the agent never learns what it is getting wrong; it accumulates a private model that drifts from the operation's actual practice, and the drift stays invisible until a cumulative metric forces an investigation that arrives a quarter too late. The learning loop is the mechanism that keeps the agent's confidence anchored to physical reality rather than to its own history. *Vendor-evaluation hook: ask how an operator override on Tuesday changes the recommendation on Wednesday. If overrides are logged but never close the loop, the agent is not learning; it is repeating.*

The codependency is the part that the executive summary always drops, and it is the part that matters most. Layer 4 without Layer 1 learns lies, confidently and at scale. Layer 3 without Layer 2 acts at speed without the authority to be acting at all. Layer 2 without Layer 1 governs noise, signing off on actions taken against measurements no one can vouch for. The layers are not a menu of pillars to select from according to budget. They are conditions that must be coherently active together, or the whole system collapses to the failure mode of its weakest layer.

This is the central architectural claim of the paper, and it is the line worth quoting in the next leadership meeting. **No layer delivers value in isolation.** An OT technology arsenal that does not sit on all four at once is not an intelligent floor. It is infrastructure waiting to be misused at the speed of a control loop.

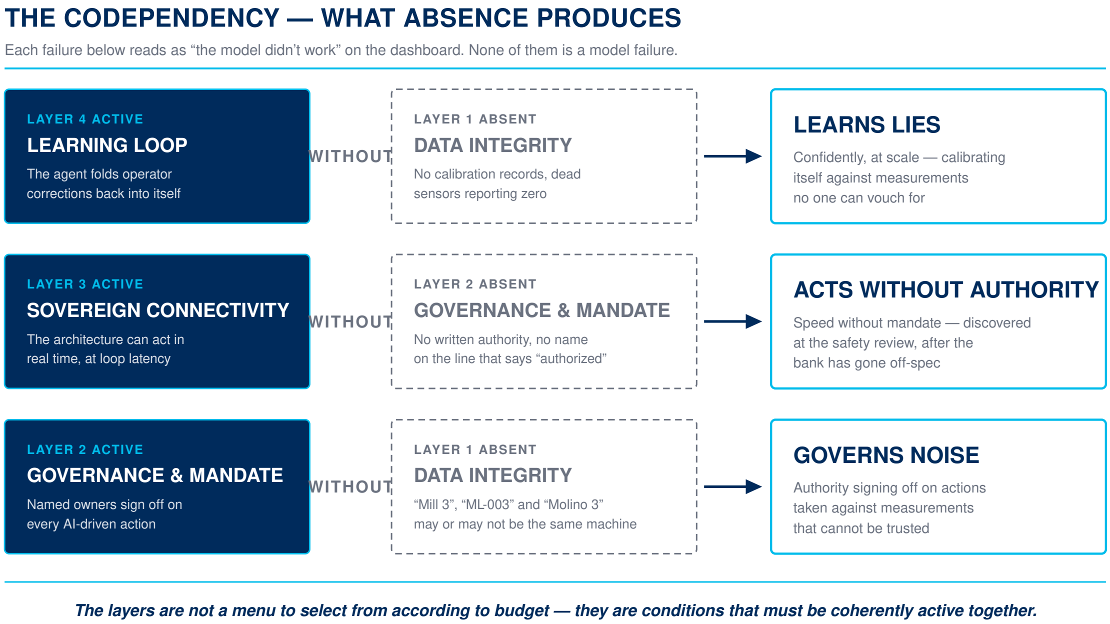

Figure 2 The codependency matrix. Each layer is annotated with the specific failure mode that emerges when the layer below it is absent: learns lies, acts without authority, governs noise. The off-diagonal cells are the failures most commonly mistaken for “the model didn’t work.”

06

WHERE THE MANDATE BREAKS DOWN

PISO names the conditions that must be active. This section names what specifically tends not to be, in LATAM mines, and why no AI vendor will fix it on the operation’s behalf. There are three of them. None is a technology problem. All three are problems of organizational architecture and governance, and that distinction is not a rhetorical flourish; it determines the remedy. A technology problem dissolves when you install the right technology. These do not. They are structural, rooted in how contracts and authority are arranged, and they persist through any number of platform upgrades until someone changes the arrangement. That is the difference between a problem you can buy your way out of and one you have to architect your way out of.

Fragmentation. The industrial network is one contract. The SCADA is a second. The dispatch system is a third. The historian is sometimes inside the SCADA, sometimes its own contract, sometimes a tenant in the corporate IT estate. When something fails at the intersection of two or more, as the gateway did in §02, the problem circulates between vendors, each performing inside its scope, the intersection owned by no one. The remedy is not a “Director of IT/OT Convergence” added to the org chart; a new box owning nothing changes nothing. The remedy is a contract clause. The next time any of those three contracts is renewed, it carries an intersection-accountability term: a named integration owner, a maximum time-to-act on cross-boundary faults, and a clause that makes "it's the other vendor's boundary" a breach rather than a defense. Fragmentation is solved at the procurement table, not the conference table.

Velocity. The business now expects response times the operation's escalation structure cannot produce. A fault that touches three vendors and a corporate IT helpdesk resolves in hours. The control loop the AI is meant to optimize closes in seconds. The gap is not technological; it is a decision-architecture gap, and it is the same gap Layer 3 exists to close: who can act, on which system, at what level of autonomy, without routing through a chain whose latency invalidates the operation's relationship to real time. An AI architecture deployed over an unaddressed velocity gap simply inherits the operation's escalation latency, which means it does not work, regardless of how fast the model infers. The remedy is an explicit autonomy table, written and signed: for each decision class, the actor, the bounded envelope inside which they may act without escalation, and the trip-wire that forces escalation. The table is OT's instrument, and writing it is OT's work.

Identity. This is the tension OT most needs to hear about itself, and it is the one that frames the others as either solvable or terminal. IT arrives with an AI proposal because IT has been building proposals for two years; proposing is its native motion. OT arrives with process knowledge, operational instinct, and thirty years of evidence about what holds on the floor and what fails at 4 AM. What OT does not arrive with is a proposal. When the meeting happens, the side holding the proposal sets the agenda, and OT is cast as the technical reviewer of a project it did not design. The operation then accepts a sequencing that hands the AI vendor the platform before the floor has been described, and the seven-in-ten begins compiling the next case study. Recall what the executives in the regional study named as the leading barrier: resistance to change, ahead of talent, far ahead of technology. The identity tension is that barrier seen from inside — an estate that knows the floor better than anyone, structurally positioned to react rather than to author.

The remedy for the identity tension is emphatically not for OT to learn to out-talk IT in AI vocabulary. OT is not a subset of IT with extra latency constraints; it is a distinct engineering discipline with its own culture, its own failure modes, and its own thirty-year memory, and the moment it tries to win on IT's terms it loses the judgment that made it valuable. Nor is the remedy to resist the convergence the rest of the business is rightly asking for: a convergence in which one estate authors and the other reviews is not convergence — it is absorption with better branding. The framework is what turns the same meeting into a negotiation between equals. The remedy is for OT to arrive at the meeting holding the framework that names the conditions the operation requires the AI to satisfy, so that the proposal is evaluated against the floor's reality rather than the floor being evaluated against the proposal's assumptions. PISO is one such framework. It is not the only one that could work. What is non-negotiable is that OT arrives with *a* framework, that it is written down before the meeting, and that the direction of evaluation runs from the floor to the proposal, not the other way around.

07

THE WINDOW THAT CLOSES WHEN AUTONOMY LANDS

The case for OT to arrive with a framework has been true for a decade. What changed in 2026 is the cost of arriving without one. Three forces converge over the next two to three years, and each one hardens the asymmetry against operations that wait.

The first is autonomy crossing from haulage into the rest of the operation. Autonomous haulage was the proof of concept the industry has now absorbed; the industry's own playbook for it is public, in the Global Mining Guidelines Group's implementation guideline for autonomous systems, now in its second version. The frontier has moved to integrated, AI-assisted operation of the whole value chain, and the regional evidence is no longer hypothetical: BHP's Chilean site-tour materials (November 2024) set out a pathway to over 500 kilotons of incremental copper at Escondida through FY31 via mine-plan optimization and concentrator initiatives; in parallel, the BHP Operating System — the operational-excellence framework that AI-assisted operation plugs into — has been credited with roughly a billion US dollars in cumulative cost savings at Escondida since its deployment. The architectural fact that matters: once an autonomous control loop is in production, retrofitting governance to it costs several times what installing governance ahead of it would have. The Layer 2 mandate has to be written before the Layer 3 architecture is built and energized, or the operation inherits a running autonomous system it has no documented authority to override. PISO treats this as the canonical sequencing case, because it is the one that becomes unrecoverable fastest.

The second is the reporting regime around ESG telemetry. The direction of travel across Chile, Peru, and Brazil is toward auditable lineage: the expectation that a figure submitted to a regulator or an investor can be traced back through the systems that computed it to the physical measurement that originated it. An AI agent that produces recommendations against unaudited substrate also produces, downstream, disclosures against unaudited substrate. The cost of that exposure stops being hypothetical the day the first regulatory or investor challenge grounded in AI-driven decision lineage sets the precedent, and the rest of the industry spends the following two years retrofitting toward it. Layer 1 is not only a precondition for the model; it is becoming a precondition for the license to operate.

The third is the corporate AI mandate itself. The boards of LATAM mining operations receive a version of the same memo every quarter: deploy AI, capture the productivity, or explain to the investors why the peer operation did and you did not. The pressure is no longer abstract, and it is being signed at the very top of the industry: in November 2025, Codelco — the leading state copper producer — signed a strategic "Future of Mining" alliance with NTT DATA in Tokyo, spanning advanced connectivity, generative AI, robotics, autonomous operations, and digital twins. (Disclosure applies in both directions: that alliance is ours, and it is cited here as evidence of where the agenda is being set, not as independent validation.) When agreements of that scope are being signed at board level, the operations below them inherit timelines. The window in which OT can set the agenda is the narrow one between the board memo and the first project signature. Once the signature is on the page, the agenda is set in someone else's framework, and the floor is in permanent catch-up.

The convergence of these three is why this paper is written in 2026 and not 2024. In 2022, the cost of arriving without a framework was a delayed pilot. In 2027, the cost is an operation that cannot defend its own autonomous decisions to a regulator, to a safety board, or to the integrator's successor. The price of waiting is no longer measured in lost upside. It is measured in non-defensibility.

08

WHAT THE INTELLIGENT FLOOR ACTUALLY COSTS TO STAND UP

The proposals are precise about the benefit and vague about the bill. They quote the recovery uplift to a decimal and leave the cost of the substrate that makes the uplift possible as an implementation detail. An OT architect who has run capital programs knows that the implementation detail is where the project lives or dies, so this section does the accounting the proposal skips, with on-premise inference as the worked example, because on-prem is what Layer 3 sovereignty usually requires and it is the line item most often understated.

There are three cost components, and they correspond to a single capital outlay against a recurring operating burden that the capital outlay quietly creates.

The first is the silicon, and it is the only one the proposal usually names. The inference appliance at the edge, the GPU or accelerator, the ruggedized enclosure rated for the dust and temperature swing of a plant room rather than a data center. It is a capital number, it is finite, and it is the smallest of the three over a five-year horizon. An operation that budgets only this item has budgeted the tip.

The second is sustaining engineering, and it is the one that turns a successful pilot into a stranded asset. An on-premise inference stack is not a microwave oven. It needs patching on a schedule that does not break the control environment, model recalibration as the ore body and the circuit drift away from the conditions the model was trained on, and a monitoring regime that detects silent degradation before a cumulative metric does. The plant superintendent did not sign up to run an MLOps practice, and if no one is funded to run it, the model decays in place while every status light stays green, which is precisely the §02 failure mode relocated from the network to the model. This is recurring OPEX, it scales with the number of deployed models, and it is the component the proposal almost never quantifies.

The third is the second source, and it is the one OT instinctively understands and the proposal almost never includes. A single edge appliance with a single vendor's stack and a single specialist who knows how to recover it is a single point of failure dressed as innovation. Sovereignty in the Layer 3 sense means the operation can keep deciding when the primary inference path is down: a failover appliance, a documented degraded-mode that hands control back to the deterministic control logic, and a second engineer who can stand the system back up at 4 AM without a support ticket to another continent. This is part capital, part the harder cost of building internal capability that does not walk out with the integrator's contract.

The point is not that on-premise inference is too expensive. It is frequently the right architecture, and the latency arithmetic of Layer 3 often forces it. The point is that the honest five-year cost is the silicon plus the sustaining engineering plus the second source, and a proposal that quotes only the silicon has not costed the intelligent floor. It has costed a demonstration. The difference between the two is the difference between a project that survives the integrator's contract ending and one that becomes the expensive object in the plant room that everyone has quietly stopped trusting.

DIAGNOSTIC QUESTION

Over a five-year horizon, what is the fully loaded cost of keeping this model coherently active — the silicon, the sustaining engineering, and the second source — and which line of our budget carries it after the integrator leaves?

09

THE QUESTION YOUR INTEGRATOR SHOULD BE ABLE TO ANSWER

The single most useful instrument an OT manager can carry into the next AI proposal meeting is one question, specific enough that a vague answer is audibly vague and structured enough that the shape of the answer reveals whether the proposing party has thought about the floor at all.

> *"What conditions must our operation satisfy, in measurable form, before this project captures the value you are projecting? Name the conditions. Name the measurement for each. Name the owner of each."*

The answer sorts the room. If it is "good data and stakeholder alignment," the proposing party has not thought about the floor, and the operation is on the conveyor toward the seven-in-ten. If it is a list of conditions, each mapped to a measurement, each mapped to a named owner, the proposing party has done the work, and the conversation can productively continue toward sequencing.

The OT manager's only job in that moment is to keep asking until the answer is specific, and to treat each retreat into generality as data. The question is not adversarial; it is the most respectful question available, because it treats the proposal as a serious multi-year commercial commitment rather than a slide rehearsal. Integrators who do good work welcome it, because it lets them show the work. Integrators whose model depends on signing first and clarifying later resent it. Both reactions are exactly the information the OT manager came to gather.

PISO supplies the conditions in portable form. An operation that has internalized the four layers walks into the meeting with the checklist already in hand and does not have to invent the question under pressure. The checklist does not replace the integrator's expertise; it anchors the conversation to the operation's reality instead of the vendor's deck, and it shifts the direction of evaluation so that the proposal answers to the floor rather than the floor auditioning for the proposal.

10

REBRANDING, THE ARSENAL, AND THE PILOT THAT LIVES FOREVER

Some operations believe they already have OT-led AI and do not. Three patterns account for most of the false positives. Each one looks like the real thing from across a boardroom. Each one fails the same examination, and each one has a one-question detection test the OT manager can apply without a workshop.

The rebranding. A corporate IT team renames itself "industrial digital" or "operational data" and assumes it now holds the OT mandate. The team is staffed with capable engineers and does useful work, but the cultural center of gravity never leaves the IT estate. The decisions that matter are still taken in the IT review board, with OT consulted rather than authoring. *Detection: ask who owns the runbook the night the SCADA goes down at three in the morning. If the answer is anyone other than the OT lead with the matching authority and budget, the rebranding is paint.*

The arsenal. An operation buys edge inference appliances, hires data scientists, and signs licenses with three AI platform vendors. The technology stack photographs beautifully. There is no governance layer, no written mandate, and no documented data-integrity program for the arsenal to operate against. The arsenal performs the work the vendors specified and does not deliver the value the operation projected, because the conditions for value were never described before the hardware arrived. This is §05's central claim made flesh: a strong arsenal on an absent substrate is theatre with a maintenance contract. *Detection: ask who has the standing mandate to override an autonomous action the moment it is wrong. If the answer is unclear or is a committee, Layer 2 is empty and the arsenal is staged.*

The pilot that lives forever. A single isolated use case (predictive maintenance on one haul-truck fleet, vibration analysis on one mill) becomes the operation's permanent answer to "what are you doing in AI." The use case may genuinely work. It does not generalize, because the conditions that carried it were specific to its narrow scope, and nobody built the substrate that would let a second use case stand on the first. Two years on, the pilot is still the pilot, the second use case never landed at scale, and the organization's mental model of "AI in mining" has hardened around a demo. *Detection: count the cross-system data flows the AI architecture supports today, then count them as of eighteen months ago. If the number has not moved, the pilot is permanent and the framework was never built.*

None of the three is recoverable by adding technology, which is the instinct each one will reach for. Each is recoverable only by returning to the four conditions and naming, out loud, which one is missing. The work is subtraction and accountability, not procurement.

11

WHAT TO VERIFY BEFORE YOU SIGN THE NEXT PROPOSAL

Here is the test you can run on your own operation, this week, with the colleagues you already have. Five verifications, each returning a binary answer, each designed to be completable inside a working week by an OT manager with access to the right people. This is the falsification test for everything argued above: if your operation passes, the thesis of this paper does not apply to you, and you should sign with confidence. The pattern across the five tells you whether you are inside the seven-in-ten or outside it.

Verification 1 · Substrate ownership. Name the single person responsible for the data integrity of the substrate the proposal depends on: physical-sensor calibration, master-data reconciliation across systems, gap-flagging in the time series. If the responsibility is spread across three roles with no single owner who can show you the calibration coverage today, Layer 1 is not coherently active and the project will inherit the failure mode of its weakest sub-layer. *Pass / fail.*

Verification 2 · Response time. On the most recent operational anomaly that crossed three vendor contracts, measure the elapsed time from the moment the anomaly registered to the moment a corrective action was actually taken. If that time is consistently north of thirty minutes, the operation has a velocity gap no AI architecture can close on its own, and the project's latency assumptions are fiction until the autonomy table is written. *Pass / fail.*

Verification 3 · Proposal authorship. Pull the last three AI proposals presented to the operation and name the authoring party for each. If all three were authored by integrators or by corporate IT, with no OT-authored counter-framework anywhere in the file, the operation has an identity gap and the next signature will land on someone else's framework. *Pass / fail.*

Verification 4 · Precondition specificity. Ask the sponsor of the active or proposed project to name the conditions the operation must satisfy, in measurable form, for the value projection to hold. Generic answer (good data, stakeholder alignment): the sponsor has not done the work. Specific answer (calibration coverage above 95% on the wet-circuit pH probes, master data reconciled across MES[7] and CMMS[8] for the affected equipment classes, sensor-to-setpoint latency bounded and measured under primary-link failure): the sponsor is engaged. *Pass / fail.*

Verification 5 · Mandate map. Draw the four PISO layers and assign a single named owner to each one in the operation as it exists today. If the OT manager and the plant manager,

[7] MES — Manufacturing Execution System. The system that manages production execution between the corporate ERP and the operational floor.

[8] CMMS — Computerized Maintenance Management System. The system that manages maintenance planning, work orders, and asset history.

working together, cannot fill all four within an hour, the substrate is not ready, and the next AI project belongs behind a foundation program rather than alongside one. *Pass / fail.*

These five are not a maturity model, and they do not average into a score. They are pass or fail. An operation that fails three or more should treat the next proposal accordingly, which means resequencing it behind the foundation rather than launching both in parallel and hoping the substrate catches up. It will not catch up on its own. Nothing on the floor ever has.

12

THE COST OF SHOWING UP WITHOUT A FRAMEWORK

The closing argument is an asymmetry, and the asymmetry is the whole case. The cost of arriving at the AI conversation with a framework is real and bounded: a foundation program of six to twelve months, an OT manager who spends time authoring rather than reviewing, a leadership conversation that resequences the next pilot behind the foundation and absorbs one quarter of slower "AI progress." The investment is measurable, the timeline is finite, and the work is reusable across every AI project the operation will face for the next decade. You build the framework once and evaluate every proposal against it forever.

The cost of arriving without a framework is unbounded, and it compounds. Each project signed ahead of the framework becomes a sunk cost the framework would have prevented. Each post-mortem retells the same lesson and changes nothing, because the lesson lives in a slide and not in a written condition with an owner. Each succession of integrator contracts leaves the operation marginally more dependent on parties who do not own the failure surface and will not sit on the safety review. The seven-in-ten is not a property of the technology. It is the cost line of operations that never authored their own framework, accumulating one case study at a time.

A LATAM mining operation in 2026 has two roles available to it. It can be the operation an AI vendor describes in next year's case study as the cautionary lesson, or the operation that authored the framework against which next year's proposals are judged. The first role is comfortable and short. The second is harder and pays compound returns, and it is the only one of the two that is still available to the floor rather than to the parties circling it.

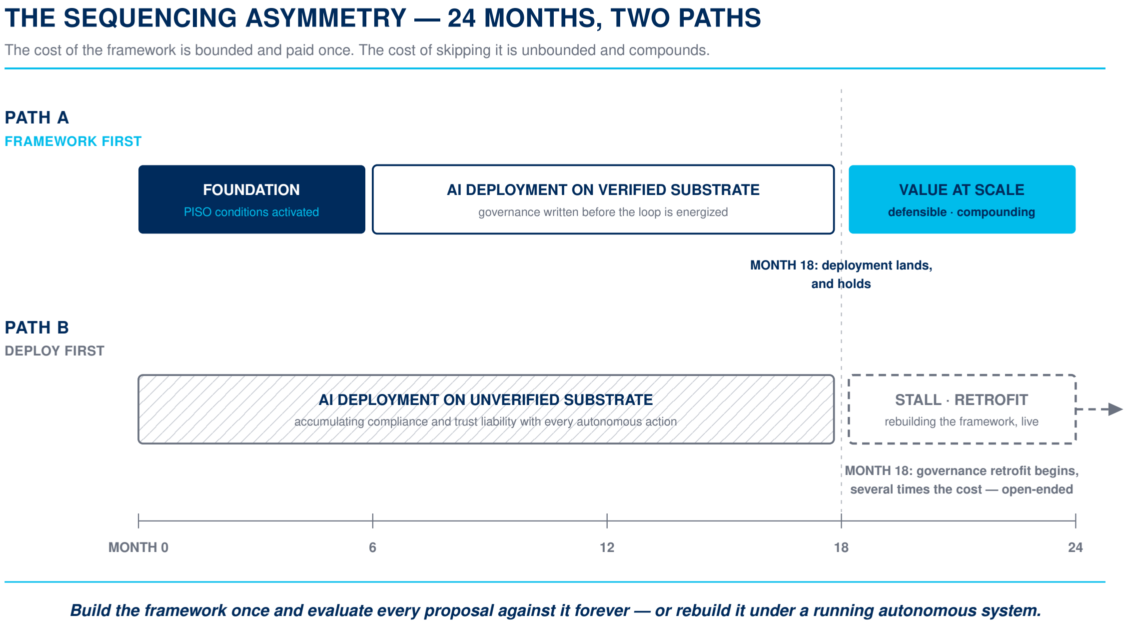

Figure 3 Two sequencing strategies across twenty-four months. The framework-first path lands a defensible AI deployment by month eighteen and compounds from there. The deploy-first path reaches month eighteen carrying an accumulating compliance and trust liability, then spends the following year rebuilding the framework that should have been written first.

The floor has sustained the mine for thirty years. It is the asset the entire AI conversation depends on, and the OT team is the only party in the operation holding thirty years of evidence about what holds on that asset and what fails at 4 AM. The conversation is finally asking what the floor thinks. Silence is an answer, and it is the one that leaves the operation in permanent catch-up. The other answer is a framework.

> *The floor sustained the mine for thirty years. The conversation that asks it to think is the one OT was built to lead. Arrive with the framework, or the framework arrives without you.*

13

KEY INSIGHTS

1 Seven of ten mining executives in LATAM acknowledge their AI and autonomy projects only partially achieved the value objectives they were approved against — and the shortfall clusters below the model, in the floor. That is the whole asymmetry: arriving with a framework costs six to twelve months of foundation; arriving without one is unbounded and compounds, one signed-too-early project at a time.

2 The barriers the executives themselves name are organizational, not technical: resistance to change and talent land far ahead of technological reliability. The proposals are written about the technical barrier. The projects die on the organizational ones.

3 Convergence is an org-chart event; failure is an architecture event. A mine can be fully converged and still join the seven-in-ten. The question is not whether IT and OT meet in the middle — it is who arrives at that meeting holding the framework.

4 PISO names four codependent conditions — data integrity, governance and mandate, sovereign connectivity, learning loop. No layer delivers value in isolation: Layer 4 without Layer 1 learns lies; Layer 3 without Layer 2 acts without authority; Layer 2 without Layer 1 governs noise.

5 "Coherently active" is the test, and it is three things at once: operational in the day-to-day, owned by a single named person with authority, and observable today without scheduling a workshop. A program that exists as a slide fails all three.

6 The honest five-year cost of the intelligent floor is the silicon plus the sustaining engineering plus the second source. A proposal that quotes only the silicon has not costed an intelligent floor. It has costed a demonstration.

7 One question sorts every proposal meeting: "What conditions must our operation satisfy, in measurable form, before this project captures the value you are projecting — name each condition, its measurement, and its owner." A generic answer is the data.

14

SOURCES

The public claims in this paper, with their sources. The §02 gateway incident and the §03 executive story are anonymized composites drawn from LATAM natural-resources engagements; they illustrate the diagnostic, not a named client. The "seven-in-ten" figure originates in NTT DATA's own research and is flagged as such where used.

- NTT DATA × MIT Technology Review, *Autonomy in the Mining Business: Beyond Technology, the Challenge for Talent and Culture*, 2025 — survey of mining executives in Chile, Brazil, Peru, and Mexico; the 72% partial-achievement finding and the barrier rankings (resistance to change 27.5%, specialized talent 14.5%, technological reliability 11.6%). Coverage: Mexico Business News (Dec 2025) · TI Inside (Dec 2025).
- Independent 2025 cross-industry executive study (>1,250 firms) — corroborating figures on AI value capture (the ~60% / ~5% pattern cited in §04). Specific reference available on request.
- International Mining, *BHP Copper Advanced Services opens new modernised ops centre in Santiago*, October 15, 2024 — US$48.3 million investment, 5.4 terabytes of operational data daily, dedicated Antofagasta–Santiago link, serving Spence and Escondida — International Mining.
- BHP, 2024 Chilean copper site tour, 18 November 2024 — Escondida: pathway to >500 kt incremental copper through FY31 via mine-plan optimization and concentrator initiatives, and the BHP Operating System credited with ~US$1 billion in cumulative cost savings at Escondida since deployment — BHP (PDF). FY2025 production (Escondida 1,305 kt, +16%, record concentrator throughput) — SEC EDGAR.
- NTT DATA, *Codelco and NTT DATA sign "Future of Mining" strategic alliance*, November 20, 2025 — first-party announcement; scope spans advanced connectivity, generative AI, robotics, autonomous operations, and digital twins — NTT DATA Newsroom.
- Global Mining Guidelines Group, *Guideline for Implementation of Autonomous Systems in Mining*, Version 2, August 2024 — the industry's reference playbook for planning, testing, and deploying autonomous systems — GMG (PDF).
- IIoT World, *Secure OT Data Flows Before Scaling AI*, May 28, 2026 — the 139 previously unknown OT data connections found in a pre-AI data-flow mapping; historian data as critical infrastructure — IIoT World.
- NTT DATA BizTalks, *IT/OT Convergence in Mining Operations*, April 2026 — companion whitepaper in this series; the convergence question this paper deliberately reframes — BizTalks.

NTT DATA · Natural Resources Practice · 2026 · biztalksnttdata.com · Pillar 01 · The Floor

For more on this topic, contact: ***Jaime Rebolledo****, Director Value Proposition Mining, NTT DATA* · jaime.rebolledocampos@nttdata.com.