PILLAR 03 · THE ORGANIZATION

Tacit Knowledge as Critical Infrastructure

Methodology for capturing expert operational knowledge before it leaves with the retiring engineer — and the architecture for the system to inherit it

01

01 · Executive Summary

01 · Executive Summary

85% of the most experienced ranks in the mining industry will retire in the next eight years. That figure — shared directly by the HR director of a major Chilean mining company — is not an abstract demographic projection. It is a countdown to the destruction of operational capacity that no asset management system has in its inventory.

An operator with 15 years at the same mine knew that crusher 2 behaved differently from crusher 3 when ore came wet from the northern sector after three days of rain. He knew that the lining 'vibrates differently' when it has a few hours left. That knowledge was not in any system. It did not appear in SAP. It was not in the manual. When he left, that mental model left with him — and three months later crusher 2 had an unplanned shutdown that technical analysis could not have predicted.

This whitepaper makes an uncomfortable argument: **the tacit knowledge of an expert operator is critical infrastructure.** It has the same status as a physical asset, deserves the same management rigor, and its loss has a calculable cost in tonnes, hours, and probability of unplanned failure.

And it proposes a concrete solution: a three-stage capture methodology — elicitation, structuring, and transfer — backed by a technical architecture that converts that knowledge into a resource the operational intelligence system can actively use, not just archive.

WHO THIS DOCUMENT IS FOR

- **VP of operations:** who has in their organization between 3 and 12 people whose retirement will represent a loss of capacity that no replacement can absorb in 12 months of onboarding.
- **HR and talent management managers:** who need a concrete framework to convert 'knowledge transfer program' from a measurable management asset into something other than a storytelling workshop.
- **Innovation and digital transformation managers:** who want to implement AI in operations and understand that without captured tacit knowledge, the system will only learn from the data incentives allowed to be recorded — not from what actually happens in the operation.

02

02 · The Problem: What the Industry Does Not Have in Its Asset Inventory

02 · The Problem: What the Industry Does Not Have in Its Asset Inventory

A modern mine has a precise inventory of its physical assets. It knows the book value of each truck, the remaining useful life of each crusher, the replacement cost of each kilometer of conveyor belt. It has preventive maintenance programs, vibration analysis, degradation models.

It has nothing equivalent for its operational knowledge.

It does not know how much of the correct diagnosis of a failure in crusher 2 depends on a specific operator being on shift. It does not know what percentage of blending decisions in the concentrator plant depend on the judgment of someone trained over 18 years of observing how that specific flotation responds with that specific ore from that specific sector. It does not have an economic valuation of the risk that each of their retirements represents.

The study *Mining Autonomy: Beyond Technology*, conducted by NTT DATA and MIT Technology Review in Spanish (2025), identified that lack of specialized talent is the second most cited barrier for advancing in autonomy, mentioned by 27.54% of surveyed executives. But that category groups two distinct problems that deserve different treatment:

PROBLEM	DESCRIPTION	SOLUTION
A: LACK OF NEW TALENT	Mining has difficulty attracting young profiles with digital capabilities. It is the visible problem — the one that appears in surveys and HR strategies.	Training programs, Digital Academies, upskilling. Horizon: 2–4 years.
B: LOSS OF EXISTING TALENT	The industry is actively losing, shift by shift, the knowledge that took decades to accumulate. It is the invisible problem — the one that does not appear in surveys because no one has inventoried it.	Capture methodology and transfer architecture. This document.

The urgency of Problem B is greater than that of Problem A for a simple reason: **Problem A can start being solved tomorrow. Problem B has a time window that closes with every retirement.**

2.1 · THE COST NO ONE IS MEASURING

When a veteran operator's successor cannot replicate the early diagnosis that operator made before the alarm threshold, the operational cost is direct and calculable:

CONSEQUENCE OF THE LOSS	OPERATIONAL EFFECT	TYPICAL COST
Late diagnosis of equipment failures	Unplanned shutdowns the predictive model did not anticipate because it does not incorporate expert judgment	USD 500K–2M per shutdown depending on equipment type and duration
Loss of ore blending capacity	Suboptimal blend decisions the system cannot calibrate without the expert's judgment history	1–3% plant recovery per suboptimal shift
Greater operational variability	The replacement operates with greater uncertainty, makes more conservative decisions, follows fewer early signals	10–25% increase in diagnostic time per shift
Degradation of the AI learning loop	The AI system learns from the replacement's decisions, which are systematically different from the expert's — without capturing the superior knowledge	The system converges toward operational mediocrity in 12–18 months

REFERENCE CASE — MAJOR COPPER CONCENTRATOR · NORTHERN CHILE
Challenge: Maximize concentrator plant production by incorporating expert operational judgment into AI systems — rather than relying solely on historical process data. **NTT DATA Solution:** AI recommender implementation with GenAI agents trained with the plant experts' operational judgment, capturing the tacit knowledge that was not visible in any sensor or historian. **Result: Over USD 10M in annual additional fine copper production, with 50% system adoption by operators within the first year.**

03

03 · What Operational Tacit Knowledge Is — and What It Is Not

03 · What Operational Tacit Knowledge Is — and What It Is Not

The term 'tacit knowledge' has a bad reputation in the mining industry because it has been associated with anecdotes, organizational storytelling, 'lessons learned' workshops that no one consulted again. That association is the main obstacle to treating it with the rigor it deserves.

The operational tacit knowledge we want to capture is not the expert's general wisdom — it is something much more specific and much more valuable:

> *"A veteran operator is not labor with seniority. They are a living predictive system, trained with thousands of hours of specific operation in that mine, that machine, those conditions."* — NTT DATA · Mining & Resources, 2025

3.1 · THE FOUR TYPES OF TACIT KNOWLEDGE WITH OPERATIONAL VALUE

TYPE	DESCRIPTION	EXAMPLE	WHY IT IS CAPTURABLE
Local correlations	Relationships between variables that apply only to this machine or specific condition, not to the general category	Crusher 2 loads when ore comes wet from the north — crusher 3 does not. Not in the manual because the manufacturer does not know the north sector.	They are causal-structured patterns. Can be expressed as conditional rules: IF condition A AND condition B THEN behavior C.
Pre-threshold signals	Early indicators the expert detects before sensors exceed the configured alarm threshold	The lining 'vibrates differently' — finer, more anxious — 4–6 hours before the vibration sensor triggers the alarm.	They have temporal structure. Can be expressed as patterns: 'when X is in normal range but Y changes this way, expect Z in T minutes'.
Exception criteria	Conditions in which the standard procedure does not apply and requires adjustment based on context	With rain in the north sector for more than two days straight, crusher 2 processing time increases even if granulometry is normal.	They are explictable contextual rules: IF operational context C THEN modify procedure P in this way.
Interdependency models	How variables from different systems that operationally appear independent interact	Fragmentation at bench 4 affects hydraulic consumption of the shovel with a 15-minute lag — which affects availability with a 2-hour lag.	They are causal relationships with temporal lag. They are exactly the type of relationship that feeds the dynamic causal graph of the intelligence system (see Pillar 02 · The Intelligence).

WHAT IS NOT OPERATIONAL TACIT KNOWLEDGE (AND WE DO NOT CAPTURE)

- Opinions about organizational management or leadership decisions
- Personal preferences about tools or systems
- Stories about the operation's past without causal structure applicable to the present
- Generic industry knowledge documented in manuals or training courses

04

04 · Three-Stage Capture Methodology

04 · Three-Stage Capture Methodology

The NTT DATA Knowledge Condenser methodology structures the capture of operational tacit knowledge in three sequential stages. Each stage has a distinct purpose, a concrete deliverable, and a verifiable quality criterion.

The program name — KODO, from Japanese 鼓動, heartbeat — reflects the design premise: expert knowledge is what gives the operation its pulse. The objective is for that pulse to keep beating when the expert is no longer there.

STAGE 01 · ELICITATION — MAKING VISIBLE WHAT THE EXPERT DOES NOT KNOW THEY KNOW

The main obstacle in capturing tacit knowledge is that the expert does not have conscious access to much of what they know. An operator with 15 years on the same equipment cannot describe the vibration in technical terms — he describes it in bodily terms. A veteran concentrator operator cannot list the variables he observes to diagnose flotation state — he only knows that 'the froth is talking to him'.

Elicitation is not an interview. It is a structured provocation process — making the expert make decisions in real time while the system records what they were observing, what variables they were considering, and what their criterion was.

Elicitation techniques by knowledge type:

- **Think-aloud in real operation:** the expert verbalizes their reasoning while operating. An analyst records without intervening. Captures local correlations and pre-threshold signals.
- **Deviation case analysis:** past situations where they deviated from the procedure are presented to the expert. 'What did you see that the system was not seeing?' Captures exception criteria.
- **Critical scenario simulation:** synthetic scenarios with real historical data are presented. 'How would the system behave under these conditions?' Captures interdependency models.
- **Contrast dialogue:** cases where the system and the expert would have differed are shown. 'Why would you disagree with this recommendation?' Captures the current model's boundary.

Capture tool: PWA mobile + STT

Capture is performed with a PWA mobile application with speech recognition (Whisper STT). The expert speaks — the system transcribes, classifies by knowledge type, and associates with the active operational context at that moment (shift, face, condition).

Cognitive load for the expert is minimal: no writing, no forms. They talk while working.

Stage 01 Deliverable — TACIT KNOWLEDGE MAP:

- **Local correlations inventory:** list of conditional relationships specific to this operation, with the conditions that activate them
- **Pre-threshold signal catalog:** each signal with its description, typical lag until the formal alarm, and the action the expert takes when they detect it
- **Exception criteria tree:** conditions under which the expert modifies the procedure and how they modify it
- **Interdependency graph:** the cross-system causal relationships the expert has learned to observe, with their estimated temporal lags

Quality criterion: A Tacit Knowledge Map is useful when a second expert from the same operation, upon reading it, says: 'yes, that is correct — I would also do that under those conditions'. If the second expert says 'I'm not sure that is generalizable', the elicitation was incomplete.

STAGE 02 · STRUCTURING — CONVERTING KNOWLEDGE INTO RULES THE SYSTEM CAN USE

The Tacit Knowledge Map produced in Stage 01 is in natural language. It cannot be directly consumed by the operational intelligence system. Stage 02 converts that map into formal structures the system can interpret, query, and apply.

From natural language to formal structure:

Raw knowledge (Stage 01):

> 'Crusher 2 behaves differently when ore comes from the north and starts to need more attention to the lubricating oil temperature — which can rise faster than normal even though the load seems normal.'

Formal structure (Stage 02):

- **Type:** local correlation + pre-threshold signal
- **Activation condition:** source_sector = North AND precipitation_72h > 20mm
- **Alert variable:** lubricating_oil_temperature_crusher_2

- **Anomalous behavior:** rate of increase > 1.5°C/hour with load < 85%
- **Action threshold:** activate cooling inspection at T+30min
- **Lag to formal alarm:** 2–4 hours from condition onset
- **Source:** Veteran operator, 15 years in operation. Validated by [expert 2]. 2026.

Technical destinations of the formal structure:

- **Diagnostic flowchart (MOV L3):** high-confidence rules incorporated into the system's structured decision tree. When activation conditions are met, the flowchart guides the operator.
- **Dynamic causal graph (MOV L4):** causal relationships with temporal lag that feed the anticipation model. Allows the system to detect the temperature signal 2–4 hours before the formal alarm.
- **Institutional memory RAG (Qdrant + LlamaIndex):** the complete structured knowledge in a semantic retrieval database. When the system faces a situation not anticipated in the flowchart, it queries this memory to find the most similar case.

> **PRACTICAL RESULT**
>
> When the veteran operator is no longer there, the system can tell the successor: 'Attention: ore is coming from the north sector in wet conditions. Under these conditions, crusher 2 lubricating oil temperature tends to rise faster than normal even though the load seems normal. Inspect cooling now, before the temperature sensor triggers the alarm.'
>
> The system does not replace the expert. It gives the successor access to that expert's criterion at the exact moment they need it.

STAGE 03 · TRANSFER — THE OPERATOR SEES THEMSELVES REFLECTED, NOT REPLACED

Stage 03 has a dual objective: verify that the captured knowledge is correct and complete, and make the expert themselves part of the process of teaching their successor. This second objective is not optional — it is the condition that determines whether the expert will actively cooperate in Stage 01.

A consistent pattern observed across multiple operations: veteran operators with the most accumulated knowledge — the most experienced person in each operation — are paradoxically the ones who initially show the most apprehension about AI systems, because their value in the organization is built on exactly what the system is trying to capture. If the capture process is perceived as a threat to their relevance, the expert will cooperate with minimal depth.

Stage 03's design inverts that logic: **the expert does not teach the system to be replaced — the expert validates that their criterion remains available when they are gone — and so that future operators make better decisions thanks to what they know.**

STEP	ACTIVITY	PARTICIPANTS	RESULT
01	The expert reviews the structured Tacit Knowledge Map and points out errors or ambiguities	Expert + knowledge analyst	Map validated by the expert themselves — first quality guarantee
02	The expert works alongside the successor using the system as a guide. When the system recommends something incorrect, the expert explains why	Expert + successor + system active in real operation	Model gap identification — enrichment opportunity
03	The expert's corrections are incorporated into the model with a traceable identifier (role, tenure, and validation year) and date as source	Expert + data architect	Complete traceability: each rule has a responsible human author, anonymized in external representations
04	The successor operates with the system for 30 days. Their deviations from the model generate automatic questions to the expert before retirement	Successor + expert as remote consultant	'Co-authorship' period that transfers context elicitation did not capture
05	Final evaluation: does the successor make decisions in critical conditions comparable to the expert's, with the system as support?	VP operations + shift supervisor	Program closure criterion — not retirement date

05

05 · Technical Architecture: How the System Inherits the Knowledge

05 · Technical Architecture: How the System Inherits the Knowledge

Knowledge captured and structured in the previous stages needs a technical architecture to be operationally useful. It is not enough to store it — the operational intelligence system must be able to query it, apply it, and update it in real time during the operation.

MOV's L3 architecture (Knowledge Layer) implements three complementary mechanisms:

5.1 · DIAGNOSTIC FLOW GRAPH (STRUCTURED FLOWCHARTS)

High-confidence rules from the Tacit Knowledge Map — local correlations with a clear threshold, pre-threshold signals with a known lag — are incorporated as additional nodes in the structured diagnostic decision tree. Each node has:

- **Activation condition:** the variables and their ranges that activate the rule
- **Recommended action:** what the operator should do when it activates
- **Source and date:** an identifier for the expert who contributed the knowledge (role, tenure, and validation year) — stored with full traceability in the system, anonymized in external representations
- **Confidence index:** what percentage of historical cases in which it was applied resulted in the expected outcome — updated automatically

This mechanism specifically resolves the problem of Step 1 in the causal chain described in the Pillar 02 whitepaper: the earliest signals — the ones a veteran operator detects before any threshold triggers — are now signals the system actively seeks.

5.2 · INSTITUTIONAL MEMORY RAG (QDRANT + LLAMAINDEX)

Exception criteria and interdependency models — the most contextual and difficult-to-formalize knowledge — are stored in a semantic retrieval vector database. When the operator or system faces a situation not found in the flowchart, institutional memory searches for the most similar case in the history of expert criteria.

Retrieval is not by keyword — it is by semantic similarity of the operational context. 'What was an expert's criterion in a situation similar to this one?' is the question the system answers automatically before escalating to the operator.

> **EXAMPLE IN OPERATION**
>
> Scenario: the night shift flotation operator detects unusual behavior in cell 3. Not in the flowchart. They write to the copilot: 'the froth in cell 3 is more viscous than normal even though the pH is in range'.
>
> Institutional memory retrieves: 'Under similar conditions (normal pH, viscous froth in cell 3), a veteran flotation operator adjusted the frothing agent dosage by -8% and maintained aeration. Result: recovery stabilized in 20 minutes. Ore temperature that night: 18°C — similar to the current one.'
>
> The operator has that criterion available at 3 AM in 2028, three years after that expert retired.

5.3 · BIDIRECTIONAL UPDATE LOOP

Captured knowledge is not static. The operation evolves — faces change, ore composition changes, equipment is replaced. The system needs mechanisms to detect when captured knowledge starts to be less predictive and to update it.

Knowledge degradation signals:

- A flowchart rule that had 85% accuracy drops to 60% in the last 30 shifts
- The successor consistently deviates from a specific rule and their results are good — signal that the captured knowledge was incomplete
- Activation conditions of a local correlation no longer occur with the same frequency — the operational context changed

Knowledge update sources:

- **Expert deviations of the successor:** when the successor makes decisions better than the model, their criterion is incorporated — continuing the capture process beyond the formal period
- **Operational condition changes:** when equipment is changed, a new face is opened, or mineralogical composition changes, affected rules are flagged for review
- **Periodic technical team validation:** monthly review of confidence indexes and update of rules that have degraded

> *Knowledge is not an archive — it is a living asset that requires active management to maintain its predictive value.*

06

06 · When to Start and How to Prioritize

06 · When to Start and How to Prioritize

The simple answer is: before the expert announces their retirement. The practical answer requires a prioritization criterion — not all people in an operation have the same level of undocumented tacit knowledge.

6.1 · CRITERIA FOR IDENTIFYING WHO TO CAPTURE FIRST

CRITERION	OBSERVABLE INDICATOR	RELATIVE WEIGHT
Tenure in the same equipment or process	More than 10 years operating the same specific system	High — the local specificity of knowledge increases with time in the same context
Differential performance under atypical conditions	The expert has a better correct diagnosis rate under abnormal conditions vs. their team	Very high — the knowledge that matters is the one that applies when the procedure is insufficient
Frequency of procedure deviations with good results	Their records show deviations the system evaluated negatively but that proved correct	High — each successful deviation is knowledge the model does not have
Proximity to retirement	Horizon of 2 years or less	Urgency — the knowledge has an expiration date, not a relative weight
Knowledge uniqueness	There is no other expert with the same level of specialization in that equipment or process	Very high — if three experts share the same knowledge, losing one is less critical

6.2 · MINIMUM CAPTURE TIME

Elicitation of a high-complexity expert requires between 8 and 16 weeks of active work — not exclusive dedication, but structured sessions integrated into the operation. This means the program must start at least 6 months before the retirement date to include the transfer period.

THE MOST FREQUENT CONDITION THAT MAKES THE PROGRAM FAIL

The organization starts the capture program 4 weeks before retirement. The expert can complete Stage 01 partially. Stage 02 (structuring) does not have enough time. Stage 03 (transfer with co-operation) does not exist.

The result is exactly the 'lessons learned' workshop that no one will consult again — the same format that gave the concept of tacit knowledge capture its bad reputation in the industry.

The difference is not the methodology. It is the time.

DIAGNOSTIC QUESTION *Is there someone in your operation whose retirement in the next 24 months would represent a loss of diagnostic capacity that no manual or replacement can absorb in 12 months?*

07

07 · What This Methodology Does Not Resolve

07 · What This Methodology Does Not Resolve

LIMIT	WHY AND WHAT IT REQUIRES
Does not capture what the expert does not know they know and cannot verbalize under any technique	There is a fraction of tacit knowledge that is genuinely pre-verbal or inaccessible to elicitation — it is bodily, motor knowledge, that Polanyi described as 'we can know more than we can tell'. The methodology minimizes this fraction but does not eliminate it. For that residual knowledge, the only solution is extended co-operation between expert and successor.
Does not work if the expert does not cooperate actively	If the expert perceives the capture process as a threat to their relevance — as preparation for their replacement — Stage 01 quality will severely degrade. Active cooperation requires the organization to have resolved incentive design to make capture beneficial for the expert, not just the organization. This is a precondition the organization must address before the methodology begins — no methodology survives a hostile incentive structure.
Does not replace successor training in explicit knowledge	The methodology captures what manuals do not have. It assumes the successor already masters what manuals do have. If the successor arrives without the basic technical training, the system can give them access to the expert's criterion but the successor will not have the framework to correctly interpret it.
Quality of captured knowledge depends on quality of historical operational data	Confidence indexes for each rule are calculated against the case history. If historical data is contaminated by incentive biases — operators documenting what was acceptable rather than what actually occurred — rules may have apparently high confidence indexes that do not reflect operational reality. Data integrity (Pillar 01 · The Floor, Pillar 02 · The Intelligence) is a non-negotiable prerequisite.

08

08 · Related Reading and Next Steps

08 · Related Reading and Next Steps

RELATED CONTENT IN THIS SERIES

- **Pillar 01 · The Floor:** PISO Framework — the OT infrastructure that enables field knowledge capture (mobile PWA + STT) to function without depending on stable connectivity; and the data integrity baseline without which captured rules produce false confidence indexes
- **Pillar 02 · The Intelligence:** Causal reasoning and hybrid intelligence — how captured tacit knowledge integrates into the dynamic causal graph (interdependency models) and the institutional RAG memory (exception criteria)
- **Pillar 04 · IT/OT Convergence:** The connected worker and the mindset shift — the organizational adoption conditions that determine whether the capture program generates cooperation or resistance

All content available at biztalksnttdata.com

NEXT STEPS

1. **Knowledge loss risk diagnostic:** identification of the 3–5 experts with the greatest undocumented tacit knowledge and their retirement horizon
2. **Capture pilot:** 12-week program with a selected expert, including the three methodology stages
3. **MOV integration:** L3 Layer configuration to incorporate captured knowledge into the active operational intelligence system

---

> *The most valuable knowledge in any operation is the kind that never made it into the manual. With the correct methodology and architecture, it no longer has to leave when the expert does.*

NTT DATA · Mining & Resources · 2026 · biztalksnttdata.com · Pillar 03 · The Organization

For information about implementing the Knowledge Condenser methodology at your operation, contact: Jaime Rebolledo, Director - Mining · jrebollc@nttdata.com