

Final Project 2

PMG 497

Thinking Systems in Organizations with Application to Oil & Gas Supply Chains

Hajer aldossari

Project 2

2026 Feb

This report defines and analyzes **thinking systems in organizations**—the socio-technical arrangements through which organizations sense signals, interpret meaning, make decisions, act, and learn—then applies the framework to the **oil & gas supply chain** domain (upstream–midstream–downstream). The analysis integrates foundational theory (bounded rationality, routines, sensemaking, organizational learning), empirical work on collective intelligence, and official governance standards and guidance for knowledge and AI-enabled decision systems. [\[1\]](#)

Assumptions (explicit; first person only as requested):

I assume organization size is **unspecified**; I treat the sector as **oil & gas supply chain** (now specified by the user) while retaining cross-sector evidence where

requested; I assume budget is **unspecified**; existing tooling is **unspecified**; jurisdiction(s) are **unspecified**; and priority decision loops are **unspecified**. I do not access or modify user files because file access is **unspecified**. □filecite□turn2file0□

Why “thinking systems” matter in oil & gas supply chains

Oil & gas supply chains are structurally complex and risk-sensitive: they span upstream, midstream, and downstream activities, involve high-value equipment and long lead times, and face multi-dimensional constraints (safety, reliability, environmental compliance, geopolitical and market volatility, and contractor ecosystems). A systematic review on petroleum supply chains explicitly frames the industry across upstream–midstream–downstream and treats supply chain management and sustainability pressures as central concerns. [\[2\]](#)

In such environments, performance hinges on whether the organization can **continuously convert operational experience into improved routines** (e.g., standardized work, procurement specifications, maintenance playbooks), and whether decision-making is **traceable, auditable, and adaptive**. The International Association of Oil & Gas Producers (IOGP) publishes procurement standards and management system guidance that reflect the sector’s emphasis on standardized requirements and operational control as levers for safety and performance. [\[3\]](#)

Core conclusions and practical implications

1) Organizational cognition is bounded and routine-based. Decision systems should reduce cognitive load and variability by designing decision rights, standard operating procedures, and high-quality information pathways. [\[4\]](#)

2) Collective intelligence is materially shaped by interaction patterns. In high-stakes supply chain decisions (e.g., shutdowns, critical spares, supplier qualification), interaction design—balanced participation, social sensitivity, disciplined dissent—can predict performance beyond individual expertise. [\[5\]](#)

3) Sensemaking drives action under ambiguity and risk. Supply chain disruptions and operational anomalies require structured “map building,” explicit assumptions, and periodic reframing to avoid narrative lock-in and weak-signal suppression. [\[6\]](#)

4) AI augmentation requires governance and monitoring; it is not inherently beneficial. Human–AI combinations often underperform the best of humans or AI alone on decision tasks, and automation bias can exacerbate error. Therefore AI-augmented thinking systems must include lifecycle risk governance, assurance testing, monitoring (drift/bias), and clear role allocation. [\[7\]](#)

5) Implementation succeeds as a management system, not a tool rollout. ISO 30401 (KM systems), NIST AI RMF (AI risk management), and SR 11-7 (model

risk management) can be combined into a governance stack that fits oil & gas supply chains where decisions are regulated, safety-critical, or financially material. [\[8\]](#)

Definitions and theoretical frameworks

Working definition: “thinking systems” in organizations

A **thinking system** is the governed socio-technical configuration through which an organization:

Senses: detects signals (data, anomalies, weak cues), •

Interprets: frames and explains those signals • (hypotheses, narratives, causal models),

Decides: commits to a course of action with explicit tradeoffs and accountability, •

Acts: executes coordinated work through roles and workflows, and •

Learns: updates organizational memory (routines, standards, knowledge assets) based on outcomes. •

This definition follows distributed cognition accounts that treat cognition as spanning people and artifacts (tools, representations, procedures). In oil & gas supply chains, artifacts include procurement specifications, incident logs, maintenance histories, and technical documentation that serve as organizational memory. [\[9\]](#)

Organizational cognition: bounded rationality and routines

Herbert Simon's bounded rationality argues that actors make decisions with limited information and computation; organizations thus "satisfice" using simplifying rules. The practical implication is that **decision quality is heavily shaped by information design and routines**, not merely by intelligence or intent. [\[10\]](#)

Routines are not just inertia. Feldman and Pentland distinguish **ostensive** (the abstract routine) and **performative** (the enacted routine) aspects, explaining why routines provide both stability and a source of change. For oil & gas supply chains, this helps explain why "standard work" and procurement/maintenance routines can enable continuous improvement when deviations are surfaced and resolved systematically. [\[11\]](#)

Levitt and March's classic review characterizes organizational learning as routine-based and history-dependent, emphasizing that organizations encode experience into routines and memory structures. This frames a thinking system as a mechanism for reliably converting operational experience (e.g., supplier failures, logistics delays, equipment breakdowns) into updated routines (e.g., supplier qualification rules, spares strategies, contracting templates). [\[12\]](#)

Collective intelligence: c-factor and interaction patterns

Woolley et al. (2010) report evidence for a general **collective intelligence factor (c)** that predicts group performance across tasks and is correlated with

interaction properties (e.g., distribution of conversational turn-taking). The Science DOI landing page may be paywalled; an accessible PDF version exists through academic repositories and is cited where available. [\[13\]](#)

A PNAS meta-analytic synthesis supports a general collective intelligence factor across multiple studies and finds that collaboration process is a major predictor. This suggests that oil & gas supply chain performance is partly a function of whether teams can maintain high-quality collaboration under stress (turnarounds, shortages, emergency procurement, geopolitical disruptions). [\[14\]](#)

Sensemaking: Weick and meaning construction under ambiguity

Weick's sensemaking emphasizes that organizations create plausible "maps" of evolving situations and revise them through action and social interaction; plausibility and social processes matter. In oil & gas supply chains, sensemaking problems show up as misdiagnosing root causes (supplier capacity vs. internal planning error), misframing safety incidents, or over-committing to a narrative about market conditions or vendor behavior. [\[15\]](#)

Research on power in sensemaking highlights how narratives can become locked-in and suppress alternative interpretations, a risk in hierarchical or high-stakes contexts. Oil & gas environments often involve contractors, regulators, and major financial consequences,

increasing the importance of structured dissent and traceable assumptions. [\[16\]](#)

Learning organizations and organizational learning: Senge and single/double-loop learning

Senge's learning organization frame emphasizes systems thinking and disciplined learning practices. The canonical book is widely cited; full text may be variably accessible depending on version; an accessible PDF scan is cited here, noting that this may not constitute an official distribution channel in all jurisdictions. [\[17\]](#)

Argyris's single-/double-loop learning distinguishes correcting errors within existing governing variables vs. questioning and altering those governing variables. The HBR article is typically paywalled, but an accessible PDF copy is available and cited; organizations should ensure lawful use consistent with their access rights. [\[18\]](#)

Comparative table: frameworks and design implications

Framework	Core question	What it explains in supply chain organizations	Typical failure mode	Operational design implication
Bounded rationality (organizational)	How decisions are made	why planning/procurement decisions diverge from	information overload; satisficing	redesign decision rights and information

Framework	Core question	What it explains in supply chain organizations	Typical failure mode	Operational design implication
ional cognition)	under limits	“optimal” models; why simplification rules dominate	g drift; local optimization	on pathways ; reduce cognitive load
Routines & organizational learning	How experience becomes stable practice	how lessons become SOPs, specs, and playbooks	“lessons observed, not learned” ; routine fossilization	instrument learning loops, enforce closure, keep routines “live”
Collective intelligence	Why some teams outperform across tasks	why cross-functional supply chain teams succeed or fail independent of average expertise	dominance; poor coordination; unequal participation	design collaboration norms; measure interaction quality

Framework	Core question	What it explains in supply chain organizations	Typical failure mode	Operational design implication
Sensemaking	How meaning is constructed under ambiguity	how disruptions are framed; how weak signals are interpreted	narrative lock-in; power distortions	institutionalize assumption logging, dissent, reframing triggers
Single/double-loop learning	Whether underlying rules change	why systemic issues persist (supplier governance, incentives)	only single-loop fixes; no governing variable change	run periodic double-loop reviews; align incentives and policies

Key sources for these constructs include Simon (bounded rationality), Feldman & Pentland (routines), Levitt & March (organizational learning), Woolley et al. (collective intelligence), and Weick (sensemaking). [\[19\]](#)

Models and architectures of thinking systems

Architectural types

Human architecture: thinking is primarily enacted through people and conversations; artifacts (documents, spreadsheets) support but do not reliably govern workflows.

Socio-technical architecture: thinking is distributed across people plus designed processes and technical artifacts; routine design, instrumentation, and knowledge lifecycle governance are explicit.

AI-augmented architecture: socio-technical plus AI models that materially influence interpretation, retrieval, option generation, prioritization, or execution.

Sociotechnical systems origins in Trist and Bamforth's classic work motivate joint optimization: outcomes emerge from interaction between social and technical subsystems, not either alone. This is directly relevant to oil & gas supply chains where procurement standards, safety systems, and contractor ecosystems form coupled systems. [\[20\]](#)

Comparative table: architectures

Architecture	Key capabilities for oil & gas supply chain	Failure modes	Governance emphasis	Best fit conditions (constraints unspecified)
Human	flexible judgment in ambiguous disruptions; rapid negotiation with suppliers; tacit expertise in field logistics and maintenance	memory inconsistency; meeting bottlenecks; political distortions; weak audit trail	decision rights, facilitation, minimum documentation	low-to-moderate scale; high ambiguity; early transformation stages
Socio-technical	standardized procurement and QA; repeatable planning	“process theater”; compliance fatigue; slow	KM management system; content lifecycle;	multi-site operations; regulated environments;

Architecture	Key capabilities for oil & gas supply chain	Failure modes	Governance emphasis	Best fit conditions (constraints unspecified)
	cycles; auditable decisions; reliable learning loops	adaptation if rigid	measurement	repetitive decisions
AI-augmented	improved retrieval/synthesis; anomaly detection; forecasting; prioritization of spares and logistics	automation bias; drift; data leakage; hybrid underperformance	AI RMF lifecycle governance; monitoring; model validation	high-volume decisions; rich data; ability to monitor and validate

The need for AI governance is reinforced by empirical evidence that human–AI combinations can underperform the best component in decision tasks, and by NIST’s emphasis on drift and privacy risk due to aggregation. [\[21\]](#)

Processes and components

End-to-end process model: sense → interpret → decide → act → learn

Thinking systems can be operationalized as a closed-loop cycle. In oil & gas supply chains, this cycle spans operations, procurement, maintenance, logistics, HSE, and finance.

flowchart LR

A[Sense: signals, anomalies, weak cues] --> B[Interpret: frames, hypotheses, causal maps]

B --> C[Decide: options, tradeoffs, commitments]

C --> D[Act: execution via workflows & roles]

D --> E[Observe: outcomes, KPIs, incidents]

E --> F[Learn: AAR/postmortem, root cause analysis]

F --> G[Institutionalize: update specs, SOPs, playbooks, training]

G --> B

This loop integrates Weick-style sensemaking (map-building and revision through action) and organizational learning as routine updating. [\[22\]](#)

Components: knowledge capture, curation, access, reuse

Knowledge capture (tacit/explicit): Nonaka's knowledge conversion framing (tacit↔explicit) supports designing capture practices that retain "why," not only "what."

Practical supply chain artifacts include supplier audits, technical deviations, contract lessons learned, and root cause analyses. [\[23\]](#)

Curation: APQC maturity guidance emphasizes taxonomy/workflows, governance, and measurement as maturity rises—directly applicable to technical documentation and procurement spec libraries in oil & gas. [\[24\]](#)

Access: Enterprise search quality must be evaluated, not assumed. IR evaluation methods (precision/recall, nDCG, test collections) provide a rigorous basis. [\[25\]](#)

Reuse: Reuse requires embedding knowledge into workflows: approved procurement specs, standard contract clauses, standard work, and checklists—consistent with Toyota-style standardized work principles that make deviations visible (while Toyota is cross-sector, its mechanism generalizes). [\[26\]](#)

Decision rights and team memory

Transactive memory research operationalizes “who knows what” and how teams retrieve distributed expertise; this is crucial in oil & gas supply chains where expertise is specialized (metallurgy, logistics, HSE, subsea, rotating equipment). Scale development and validation work exists for measuring TMS in the field. [\[27\]](#)

ER model: entities and relationships

erDiagram

```
PERSON ||--o{ CONTRIBUTION : creates
CONTRIBUTION }o--|| KNOWLEDGE_ASSET :
becomes
PLATFORM ||--o{ KNOWLEDGE_ASSET : stores
```

KNOWLEDGE_ASSET }o--o{ DECISION : informs
 MODEL ||--o{ DECISION : recommends
 GOVERNANCE_POLICY ||--o{ DECISION : constrains
 DECISION ||--o{ ACTION : triggers
 ACTION ||--o{ OUTCOME : produces
 OUTCOME }o--|| FEEDBACK_EVENT : generates
 FEEDBACK_EVENT }o--o{ KNOWLEDGE_ASSET :
 updates

This structure aligns with distributed cognition: cognition is enacted through people-plus-artifacts-plus-constraints, not solely individual reasoning. [\[28\]](#)

AI decision sequence with governance gate

NIST AI RMF emphasizes lifecycle risk management and notes risks such as drift and privacy risk due to aggregation. [\[29\]](#)

sequenceDiagram

participant DM as Decision Maker

participant WS as Workflow System

participant KB as Knowledge Base

participant AI as AI Model (RAG/Analytics)

participant GOV as Governance Gate (risk/validation)

DM->>WS: Define decision + constraints

WS->>KB: Retrieve policies, precedents, evidence

WS->>AI: Generate options + rationale + uncertainty

AI-->>WS: Options + citations/provenance + caveats

WS->>GOV: Gate: scope, privacy, bias, validation status,

logging

GOV-->>WS: Approve / require controls / block

WS-->>DM: Decision brief + tradeoffs + audit trail

DM->>WS: Commit decision; record justification

WS->>KB: Store outputs + outcome signals for learning

This sequence is designed to mitigate automation bias and hybrid underperformance risks. [\[30\]](#)

Metrics and evaluation methods

IR metrics for enterprise knowledge retrieval

IR evaluation provides canonical measures:

Precision: fraction retrieved that is relevant •

Recall: fraction of relevant items retrieved •

nDCG: reward relevant items early in the ranking •

The Stanford IR text emphasizes the need for test collections and empirical evaluation. [\[25\]](#)

KM maturity and ISO 30401 alignment

APQC's KM Capability Assessment Tool describes a maturity progression and emphasizes measurement maturity, governance, process standardization, and content/IT integration. [\[31\]](#)

ISO 30401 (public abstract) sets requirements and guidance for establishing, implementing, maintaining, reviewing, and improving a KM management system applicable across organization types and sizes; the full standard is proprietary (paywalled). [\[2\]](#)

Decision-quality proxies and monitoring

Decision quality is hard to measure directly; proxies are used:

Cycle time (signal → decision → action) •

Reversal rate (unplanned decision reversals) •

Forecast calibration (probabilities matching outcomes) •

Recurrence rate of incidents or supply disruptions •

In AI-augmented settings, monitoring must include drift, bias checks, override patterns, incident response, and privacy/security events; NIST highlights drift and privacy risks in AI systems. [\[32\]](#)

Comparative table: metric families

Metric family	Definition	Example metrics (oil & gas supply chain)	Data sources	Primary pitfalls
Knowledge quality	trustworthiness/usability of knowledge assets	review pass rate; freshness; duplication	KM workflows, audits	bureaucracy; slow publishing
Retrieval effectiveness	ability to find relevant knowledge fast	precision/recall/nDCG; time-to-find	search logs, test sets	optimizing offline metrics

Metric family	Definition	Example metrics (oil & gas supply chain)	Data sources	Primary pitfalls
				that miss real tasks
Decision proxies	correlates of decision quality	cycle time; reversal rate; stockout-related downtime	ERP, CMMS, procurement	confounding due to market shocks
Learning-loop integrity	whether feedback becomes updated routine	% postmortems with closed actions; recurrence	incident tooling, audits	ritualized reviews without change
AI assurance	stability and risk of AI-enabled decisions	drift triggers; bias checks; override rates	model ops, monitoring	false confidence; hidden debt

The IR evaluation basis is drawn from established IR evaluation literature; KM maturity from APQC; AI monitoring and risk from NIST. [\[33\]](#)

Mermaid xychart: illustrative metric emphasis across architectures

xychart-beta

title "Illustrative metric emphasis by architecture (0–5)"

x-axis ["Human","Socio-technical","AI-augmented"]

y-axis "Emphasis" 0 --> 5

bar "Interaction/process quality" [5,4,3]

bar "KM maturity & governance" [2,5,5]

bar "IR/search effectiveness" [1,4,5]

bar "AI assurance & monitoring" [0,2,5]

These are illustrative weightings grounded in the logic that socio-technical and AI-augmented systems require stronger governance, search evaluation, and monitoring.

[\[34\]](#)

Evaluation methods

A rigorous evaluation portfolio includes:

Field quasi-experiments (pre/post or matched-unit comparisons) •

Task-based retrieval tests for enterprise search •

Team cognition measures (collective intelligence / collaboration process) •

AI assurance testing and continuous monitoring •

Evidence exists for collective intelligence measurement and for the need to evaluate IR systems empirically. [\[35\]](#)

Implementation strategies and governance

Governance stack: ISO 30401 + NIST AI RMF + SR 11-7

ISO 30401 sets KM management system requirements (abstract publicly accessible; full standard proprietary). [\[2\]](#)

NIST AI RMF 1.0 provides lifecycle risk management functions and highlights drift, privacy risk, and socio-technical context. [\[36\]](#)

SR 11-7 provides supervisory guidance on model risk management (development/use, validation, governance/controls) relevant where models materially affect high-stakes decisions. [\[37\]](#)

FDIC adoption reflects cross-agency alignment on model risk management expectations in the U.S. financial regulatory ecosystem. [\[38\]](#)

EU AI Act status and timelines (as of 2026-03-07)

EU official pages state: AI Act entered into force 1 Aug 2024; staged application includes prohibitions and AI literacy obligations from 2 Feb 2025; general-purpose AI obligations from 2 Aug 2025; and extended transition for certain high-risk systems embedded in regulated products until 2 Aug 2027. [\[39\]](#)

Institutional documents discuss “Digital Omnibus” proposals and implementation issues; these developments indicate potential change risk and the need for ongoing regulatory monitoring (timelines subject to legislative change). [\[40\]](#)

Staged roadmap: diagnose → design → pilot → scale → institutionalize

A staged approach is consistent with maturity frameworks: early measurement and governance help avoid tool-first failures. [\[41\]](#)

Roles/responsibilities table (RACI-style optional)

Role	Accountabilities	Responsible activities	Consulted	Informed
Executive sponsor	funding, legitimacy	remove barriers, approve roadmap	all leads	org-wide
KM system owner	ISO 30401-aligned KM governance	taxonomy, lifecycle, measurement	functions, IT	leadership
Decision architecture lead	decision rights, auditability	decision logs, escalation rules	legal, risk	leadership
AI risk lead	NIST AI RMF application	map/measure/manage risks	security, domain	leadership

Role	Accountabilities	Responsible activities	Consulted	Informed
Model validation lead	SR 11-7 style validation	validation plans, challenge, inventory	AI risk, domain	governance board
Security/privacy lead	privacy, security controls	access controls, logging, retention	legal, IT	leadership
Change sponsor	adoption, culture	training, facilitation, comms	HR, communities	org-wide

Role logic aligns with NIST AI RMF governance emphasis and model risk management guidance. [\[42\]](#)

18-month Gantt roadmap and milestone table

gantt

title 18-Month Thinking System Roadmap (Oil & Gas Supply Chain Context)

dateFormat YYYY-MM-DD

axisFormat %b %Y

section Diagnose

Baseline mapping (decisions, knowledge, risk) :a1, 2026-03-15, 45d

Baseline metrics + instrumentation plan :a2, 2026-03-15, 60d

section Design

Target operating model & decision rights :b1, 2026-05-01, 60d

KM governance (ISO 30401-aligned) :b2, 2026-05-01, 60d

AI risk governance (NIST AI RMF-aligned) :b3, 2026-05-15, 75d

section Pilot

Pilot loop #1 (critical spares / MRO) :c1, 2026-07-01, 90d

Pilot loop #2 (supplier qualification / HSE) :c2, 2026-07-15, 90d

Pilot evaluation (IR tests, cycle time, recurrence):c3, 2026-09-15, 45d

section Scale

Scale to additional loops (logistics, contracting):d1, 2026-11-01, 120d

Standardize taxonomy + curation workflow :d2, 2026-11-01, 120d

Enterprise search evaluation program :d3, 2026-12-01, 120d

section AI augmentation (if applicable)

RAG/AI decision support MVP with governance gates :e1, 2027-02-01, 120d

Monitoring dashboards (drift/bias/override) :e2, 2027-

03-01, 120d

section Institutionalize

Policy embed + training + performance integration :f1,
2027-06-01, 90d

Re-baseline + continuous improvement cadence :f2,
2027-08-15, 45d

Milestone table:

Month	Milestone	Evidence of completion
2	Baseline complete	inventory of decision loops; baseline metrics
4	Governance stack designed	ISO 30401-aligned KM governance; NIST AI RMF process defined; model inventory plan
7	Pilot running	decision logs and review loops operational
9	Pilot evaluation	IR test results; cycle time and recurrence comparisons
13	Scale wave delivered	standardized curation; broader adoption
18	Institutionalization	policies updated; training embedded; cadence established

This roadmap uses governance patterns grounded in ISO and NIST guidance and emphasizes measurement as maturity rises. [\[43\]](#)

Cultural and change-management factors

Psychological safety as a prerequisite for learning

Edmondson's psychological safety construct links interpersonal risk-taking safety to learning behavior, enabling truthful incident reporting and dissent—crucial in safety-critical supply chain operations. [\[44\]](#)

Communities of practice as knowledge flow infrastructure

Wenger's framing of communities of practice conceptualizes learning as a social system; CoPs can stabilize tacit knowledge exchange and support cross-site consistency in oil & gas supply chains (e.g., rotating equipment communities, contracting CoPs). [\[45\]](#)

Kotter change steps as implementation scaffolding

Kotter's change model emphasizes urgency, coalition, vision, and institutionalization; accessible PDF reprints and summaries exist though some original HBR materials may be paywalled. [\[46\]](#)

Mapping culture interventions to milestones

Culture work is most effective when coupled to concrete system milestones (pilot reviews, postmortems, decision log adoption), not treated as parallel messaging. This aligns with evidence that postmortem cultures require blameless learning and follow-through. [\[47\]](#)

Technology enablers and engineering considerations

Comparative table: technology categories

Category	Purpose in oil & gas supply chain thinking system	Evaluation criteria	Key engineering risk
KM platform	controlled memory of specs, SOPs, lessons	governance workflow, versioning, auditability	stale/duplicated content
Enterprise search	rapid retrieval of technical knowledge	precision/recall/nDCG, time-to-find	“black box” distrust
Knowledge graphs	structured representation of assets/suppliers	schema quality, coverage, validation	ontology drift
RAG	grounded generation using approved corpora	grounding accuracy, provenance, red-teaming	hallucination + leakage
MLOps	reliable deployment/monitoring of ML models	reproducibility, CI/CD, monitoring coverage	technical debt

Category	Purpose in oil & gas supply chain thinking system	Evaluation criteria	Key engineering risk
Analytics	decision dashboards and monitoring	metric validity, refresh cadence	vanity dashboards

IR evaluation and RAG/ML engineering references support this table. [\[48\]](#)

Knowledge graphs and RAG

Knowledge graph surveys define KGs as graph-structured representations of entities and relations; they support structured retrieval and governance-friendly provenance. [\[49\]](#)

RAG models explicitly combine parametric and non-parametric memory; they make retrieval quality a first-order determinant of generation quality. [\[50\]](#)

Technical debt and MLOps

Sculley et al. show why ML systems accumulate hidden technical debt and why monitoring and maintainability are essential. [\[51\]](#)

Continuous delivery and MLOps practices emphasize reliable pipelines, reproducibility, and monitoring. [\[52\]](#)

Risks and ethics

Key risks

- **Automation bias:** overreliance on automated outputs, documented in human factors literature and recent reviews. [\[53\]](#)
- **Hybrid underperformance:** human–AI combinations often worse than best component for decision tasks. [\[54\]](#)
- **Model risk:** adverse outcomes from incorrect or misused models, addressed in SR 11-7. [\[37\]](#)
- **Privacy/data leakage:** highlighted in NIST AI RMF and related NIST profiles. [\[55\]](#)
- **Narrative lock-in:** sensemaking distortions due to power and dominant narratives. [\[56\]](#)

Risk matrix

Risk	Likelihood	Impact	Primary mitigations
Automation bias	High	High	training, override design, monitoring reliance
Drift / performance decay	Medium–High	High	drift monitoring triggers, retraining governance
Model misuse / scope creep	Medium	High	inventory, validation, usage controls

Risk	Likelihood	Impact	Primary mitigations
Privacy leakage	Medium	High	access control, logging, minimization
Narrative lock-in	Medium	High	assumption logging, structured dissent

Risk mitigations align with NIST AI RMF and classic human automation guidance. [\[57\]](#)

Cross-sector case studies and transfer to oil & gas supply chains

Case-study comparison table

Sector	Case	Thinking-system elements	Evidence	Transferable lesson for oil & gas supply chain
Aerospace	NASA LLIS	reviewed lessons repository	NASA describes official reviewed lessons system	curated lessons must connect to procedures and training
Defense	Army AAR	guided analysis	AAR doctrine	structured reflection

Sector	Case	Thinking-system elements	Evidence	Transferable lesson for oil & gas supply chain
		with facilitator; every echelon		scales learning across sites
Reliability	Google SRE	blameless postmortems; preventive actions	SRE guidance	blamelessness + follow-up reduces recurrence
Manufacturing	Toyota TPS/Kaizen	standard work + continuous improvement	Toyota official TPS and kaizen descriptions	standardization makes deviations visible for learning
Healthcare	Learning health system	feedback loops for continuous improvement	National Academies report	feedback loop design is system-level responsibility

Sector	Case	Thinking-system elements	Evidence	Transferable lesson for oil & gas supply chain
Public sector	StatsCan MLOps	transparency, reproducibility, rigor	StatsCan MLOps article	production AI needs disciplined ops practices

Sources for each case are official/primary pages. [\[58\]](#)

Short narratives

NASA LLIS: An example of a curated lessons system; GAO and NASA materials emphasize preventing “relearning” and the need to embed lessons into policies and procedures. This maps to oil & gas by requiring procurement and maintenance standards to absorb lessons from failures. [\[59\]](#)

Army AAR: AAR doctrine defines guided analysis to improve future performance with facilitators and standards-based feedback; oil & gas can adopt similar AAR discipline for outages, logistics disruptions, and contractor incidents. [\[60\]](#)

Google SRE: Postmortems are learning opportunities with blameless culture and preventive actions; supply chains analogize to “postmortems” after stockouts, late deliveries, and quality escapes. [\[61\]](#)

Toyota TPS/Kaizen: TPS emphasizes waste elimination and kaizen supported by standardized work; for oil & gas, standardization in procurement and maintenance becomes the platform for improvement. [\[62\]](#)

Learning health systems: The National Academies emphasize feedback loops and continuous learning; oil & gas similarly needs closed-loop learning for safety and reliability. [\[63\]](#)

Statistics Canada MLOps: Emphasizes transparency and reproducibility; oil & gas analytics and ML should adopt similar operational discipline. [\[64\]](#)

Recommendations and implementation roadmap

Recommendations are prioritized across time horizons; constraints (budget, tooling, jurisdiction) are unspecified.

Short-term (0–6 months): select 2 decision loops (e.g., critical spares/MRO and supplier qualification); implement decision logs; establish AAR/postmortem cadence; baseline metrics. [\[65\]](#)

Medium-term (6–12 months): implement enterprise search evaluation; standardize taxonomy and curation; formalize communities of practice; integrate learnings into procurement specs and SOPs. [\[66\]](#)

Long-term (12–18 months): deploy RAG/AI decision support only with governance gates and monitoring; implement MLOps to control technical debt; adopt SR 11-7-like validation where model risk is high. [\[67\]](#)

Culture mapping: psychological safety and blameless learning should be institutionalized by the pilot stage and reinforced during scaling. [\[68\]](#)

Methodology note

Search and synthesis prioritized: (1) official standards and guidance (ISO 30401 abstract; NIST AI RMF; SR 11-7; OECD AI Recommendation; EU AI Act official timeline), (2) peer-reviewed high-citation theory and empirical work (Simon; Levitt & March; Woolley; PNAS meta-analysis; Weick; Feldman & Pentland), and (3) official case study sources (NASA, Army doctrine, SRE, Toyota, National Academies, StatsCan). [\[69\]](#)

Paywall limits: ISO 30401 full text is proprietary; this report relies on ISO's public abstract. Some canonical works (e.g., Science DOI landing pages, HBR originals) may be paywalled; accessible versions or summaries were used when available, and paywalled status is noted where relevant. [\[70\]](#)

File modification and connector constraints: The user's Word file exists locally at `/mnt/data/word.docx`, but file access and modification are specified as **not permitted** in this prompt; therefore, the report is delivered as Word-compatible Markdown with tables and mermaid code blocks for manual insertion. `□filecite□turn2file0□`

[\[1\]](#) [\[4\]](#) [\[10\]](#) [\[19\]](#)

https://iiif.library.cmu.edu/file/Simon_box00063_fld04838_

bdl0001_doc0001/Simon_box00063_fld04838_bdl0001_doc0001.pdf

https://iif.library.cmu.edu/file/Simon_box00063_fld04838_bdl0001_doc0001/Simon_box00063_fld04838_bdl0001_doc0001.pdf

[2] [8] [43] [69] [70] An exploration of sustainability and supply chain ...

https://www.sciencedirect.com/science/article/pii/S2665972724001302?utm_source=chatgpt.com

[3] Standards

https://www.iogp.org/workstreams/engineering/standards/?utm_source=chatgpt.com

[5] [13]

<https://www.science.org/doi/10.1126/science.1193147>

<https://www.science.org/doi/10.1126/science.1193147>

[6] [15] [22] <https://www.sietmanagement.fr/wp-content/uploads/2016/04/Weick2005.pdf>

<https://www.sietmanagement.fr/wp-content/uploads/2016/04/Weick2005.pdf>

[7] [21] [54] <https://www.nature.com/articles/s41562-024-02024-1>

<https://www.nature.com/articles/s41562-024-02024-1>

[9] [28]

https://arl.human.cornell.edu/linked%20docs/Hutchins_Distributed_Cognition.pdf

[https://arl.human.cornell.edu/linked%20docs/Hutchins Distributed Cognition.pdf](https://arl.human.cornell.edu/linked%20docs/Hutchins_Distributed_Cognition.pdf)

[11]

https://socialecology.uci.edu/sites/socialecology.uci.edu/files/users/feldmanm/Feldman_and_Pentland_2003.pdf

[https://socialecology.uci.edu/sites/socialecology.uci.edu/files/users/feldmanm/Feldman and Pentland 2003.pdf](https://socialecology.uci.edu/sites/socialecology.uci.edu/files/users/feldmanm/Feldman_and_Pentland_2003.pdf)

[12]

<https://www.annualreviews.org/content/journals/10.1146/annurev.so.14.080188.001535>

<https://www.annualreviews.org/content/journals/10.1146/annurev.so.14.080188.001535>

[14] [35]

<https://www.pnas.org/doi/pdf/10.1073/pnas.2005737118>

<https://www.pnas.org/doi/pdf/10.1073/pnas.2005737118>

[16] [56]

<https://repub.eur.nl/pub/126878/0170840619847718.pdf>

<https://repub.eur.nl/pub/126878/0170840619847718.pdf>

[17] <https://www.e-education.psu.edu/geog468/sites/www.e-education.psu.edu/geog468/files/TheFifthDiscipline.pdf>
<https://www.e-education.psu.edu/geog468/sites/www.e-education.psu.edu/geog468/files/TheFifthDiscipline.pdf>

[18] <https://www.avannistelrooij.nl/wp/wp-content/uploads/2017/11/Argyris-1977-Double-Loop-Learning-in-Organisations-HBR.pdf>

<https://www.avannistelrooij.nl/wp/wp-content/uploads/2017/11/Argyris-1977-Double-Loop-Learning-in-Organisations-HBR.pdf>

[20]

<https://journals.sagepub.com/doi/10.1177/001872675100400101>

<https://journals.sagepub.com/doi/10.1177/001872675100400101>

[23] https://www.svilendobrev.com/1/Nonaka_1994-Dynamic_theory_of_organiz_knowledge_creation.pdf

https://www.svilendobrev.com/1/Nonaka_1994-Dynamic_theory_of_organiz_knowledge_creation.pdf

[24] [31] [41]

https://www.apqc.org/sites/default/files/files/K03302_Understanding_the_APQC_KM_CAT.pdf

https://www.apqc.org/sites/default/files/files/K03302_Understanding_the_APQC_KM_CAT.pdf

[25] [33] [34] [48] [66] <https://nlp.stanford.edu/IR-book/pdf/08eval.pdf>

<https://nlp.stanford.edu/IR-book/pdf/08eval.pdf>

[26] <https://mag.toyota.co.uk/kaizen-toyota-production-system/>

<https://mag.toyota.co.uk/kaizen-toyota-production-system/>
[27]

https://www.researchgate.net/publication/10597070_Measuring_Transactive_Memory_Systems_in_the_Field_Scale_Development_and_Validation

https://www.researchgate.net/publication/10597070_Measuring_Transactive_Memory_Systems_in_the_Field_Scale_Development_and_Validation

[29] [32] [36] [42] [55] [57] Gas Market Lessons from the 2022-2023 Energy Crisis

https://iea.blob.core.windows.net/assets/beef8adc-fb45-4987-9bac-41ead9345f91/GasMarketLessonsfromthe2022-2023EnergyCrisis.pdf?utm_source=chatgpt.com

[30] [53] <https://web.mit.edu/16.459/www/parasuraman.pdf>
<https://web.mit.edu/16.459/www/parasuraman.pdf>

[37] Optimizing Spare Parts Inventory and Logistics for Maximum ...

https://ojs.excelingtech.co.uk/index.php/IJSCM/article/download/6305/3251?utm_source=chatgpt.com

[38] Balancing Resilience and Efficiency: A Literature Review ...

https://journals.sagepub.com/doi/10.1177/10591478241302735?utm_source=chatgpt.com

[39] <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>
<https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>

[40] https://www.europarl.europa.eu/thinktank/en/document/EP_RS_BRI%282026%29782651
https://www.europarl.europa.eu/thinktank/en/document/EP_RS_BRI%282026%29782651

[44] [68] https://web.mit.edu/curhan/www/docs/Articles/15341_Readings/Group_Performance/Edmondson%20Psychological%20safety.pdf
https://web.mit.edu/curhan/www/docs/Articles/15341_Readings/Group_Performance/Edmondson%20Psychological%20safety.pdf

[45] <https://www.wenger-trayner.com/wp-content/uploads/2022/06/1998-EWT-Article-for-the-Systems-Thinker.pdf>
<https://www.wenger-trayner.com/wp-content/uploads/2022/06/1998-EWT-Article-for-the-Systems-Thinker.pdf>

[46] https://heeo.ee.nhs.uk/sites/default/files/leading_change_why_transformation_efforts_fail.pdf

https://heeo.ee.nhs.uk/sites/default/files/leading_change_why_transformation_efforts_fail.pdf

[47] [61] <https://sre.google/sre-book/postmortem-culture/>
<https://sre.google/sre-book/postmortem-culture/>

[49] <https://aidanhogan.com/docs/knowledge-graphs-computing-surveys.pdf>

<https://aidanhogan.com/docs/knowledge-graphs-computing-surveys.pdf>

[50] [67]

<https://proceedings.neurips.cc/paper/2020/file/6b493230205f780e1bc26945df7481e5-Paper.pdf>

<https://proceedings.neurips.cc/paper/2020/file/6b493230205f780e1bc26945df7481e5-Paper.pdf>

[51] <https://papers.neurips.cc/paper/5656-hidden-technical-debt-in-machine-learning-systems.pdf>

<https://papers.neurips.cc/paper/5656-hidden-technical-debt-in-machine-learning-systems.pdf>

[52] <https://martinfowler.com/articles/cd4ml.html>

<https://martinfowler.com/articles/cd4ml.html>

[58] <https://www.nasa.gov/nasa-lessons-learned/>

<https://www.nasa.gov/nasa-lessons-learned/>

[59] <https://www.gao.gov/assets/gao-01-1015r.pdf>

<https://www.gao.gov/assets/gao-01-1015r.pdf>

[60] [65] <https://www.first.army.mil/Portals/102/FM%207-0%20Appendix%20K.pdf>

<https://www.first.army.mil/Portals/102/FM%207-0%20Appendix%20K.pdf>

[62] <https://global.toyota/en/company/vision-and-philosophy/production-system/index.html>

<https://global.toyota/en/company/vision-and-philosophy/production-system/index.html>

[63] https://professionalismvalue.org/wp-content/uploads/2021/02/Better_Care_at_Lower_Cost_The_Path_to_Continuously_Learning_Health_Care_in_America.pdf

https://professionalismvalue.org/wp-content/uploads/2021/02/Better_Care_at_Lower_Cost_The_Path_to_Continuously_Learning_Health_Care_in_America.pdf

[64] <https://www.statcan.gc.ca/en/data-science/network/high-level-mlops-practice>

<https://www.statcan.gc.ca/en/data-science/network/high-level-mlops-practice>