INTERNATIONAL JOURNAL OF MULTIDISCIPLINARY EVOLUTIONARY RESEARCH

Advances in Automated Risk Assessment Frameworks for Regulatory Deliverables in Upstream Energy Operations

Ngonadi Uchechi [1*], Michael Ominyi [2]
[1-2] RusselSmith Limited, Lagos, Nigeria

* Corresponding Author: **Ngonadi Uchechi**

Article Info

P-ISSN: 3051-3502
E-ISSN: 3051-3510
Volume: 01
Issue: 02
July - December 2020
Received: 22-09-2020
Accepted: 23-10-2020
Published: 18-11-2020
Page No: 158-175

Abstract
Upstream energy operations generate a dense and continuously expanding portfolio of regulatory deliverables: safety cases, safety and environmental management system (SEMS) documentation, well control and blowout contingency plans, environmental impact assessments, oil spill response plans, quantitative risk assessments, barrier management reports, and decommissioning submissions. Historically these deliverables have been assembled manually through document-centric workflows in which risk information is transcribed, reformatted, and re-argued for each submission. The result is a persistent gap between the live risk state of an asset and the risk picture presented to the regulator.

This paper reviews the decade from 2010 to 2020 and asks what has genuinely advanced in the automation of risk assessment for regulatory purposes, as distinct from what has merely been digitised. A structured review of 148 peer-reviewed and grey-literature sources was conducted, supported by an analysis of five offshore regulatory regimes and a coding framework covering six dimensions of automation maturity. Six clusters of advance are identified: dynamic Bayesian risk updating, digitalised barrier and bow-tie management, machine learning for precursor and anomaly detection, natural language processing for regulatory requirement extraction, digital twin coupling of process simulation with risk models, and automated evidence and audit trail generation.

The review finds that analytical capability has advanced considerably faster than assurance capability. Methods for computing risk from live data matured substantially over the decade, while methods for demonstrating to a regulator that a computed risk figure is credible, traceable, and owned by a competent person advanced very little. This asymmetry, rather than any shortfall in algorithmic sophistication, is identified as the binding constraint on regulatory adoption.

The paper proposes a six-layer reference architecture, designated ARAF-URD (Automated Risk Assessment Framework for Upstream Regulatory Deliverables), that treats the regulatory deliverable as a queryable, versioned view over a governed risk knowledge base rather than as a document to be authored. Central to the architecture is a requirement-to-evidence traceability model in which every assertion in a submission resolves to a dated, attributed, and reproducible evidence object. Four illustrative cases, based on synthetic but operationally representative data, demonstrate the architecture applied to a deepwater drilling permit, a SEMS annual audit cycle, a flaring and venting environmental submission, and a decommissioning comparative assessment. The paper closes with a research agenda organised around validation, explainability, regulator-side capability, and the legal allocation of accountability for machine-generated risk assertions.

DOI: https://doi.org/10.54660/IJMER.2020.1.2.158-175

Keywords: automated risk assessment, regulatory compliance, upstream oil and gas, safety case, dynamic risk analysis, Bayesian networks, barrier management, regulatory technology, digital twin, assurance

1. Introduction

The upstream energy sector operates under one of the most demanding regulatory documentation burdens of any industrial domain. A single deepwater development may require several thousand pages of risk-related submission before first oil, and will thereafter generate recurring deliverables on annual, event-driven, and campaign-driven cycles for the operating life of the asset.

The content of these deliverables is not administrative. It is the formal mechanism by which an operator demonstrates that major accident hazards have been identified, that barriers against those hazards exist and function, and that residual risk has been reduced to a level the jurisdiction considers acceptable.

Two forces converged during the decade from 2010 to 2020 to place this documentation model under strain. The first was regulatory. The Macondo blowout of 2010 and the Montara release of 2009 triggered a sustained international tightening of upstream risk regulation. In the United States, the Bureau of Safety and Environmental Enforcement was established and the SEMS rule was extended through SEMS II, adding requirements for employee stop-work authority, independent third-party audit, and documented job safety analysis. In the United Kingdom, the Offshore Installations (Offshore Safety Directive) (Safety Case etc.) Regulations 2015 transposed the European Offshore Safety Directive and introduced a combined competent authority. Australia's NOPSEMA consolidated safety, well integrity, and environmental functions. Norway's Petroleum Safety Authority deepened its emphasis on barrier management and the demonstration of barrier element performance standards. Each of these developments increased both the volume and the evidential specificity of what must be submitted.

The second force was technological. Over the same decade, upstream assets became densely instrumented. Real-time drilling data, subsea condition monitoring, high-frequency process historians, integrated operations centres, and remote diagnostics generated a volume and velocity of operational data that had no counterpart in the risk assessment practice of the 1990s, when the dominant methods were formalised. Parallel work on data-driven occupational safety risk control in large-scale energy infrastructure projects (Obriki & Arumosoye, 2018) and on inventory availability and plant uptime in energy facilities (Okonkwo *et al.*, 2018b) documents the same shift in adjacent operational domains. The consequence was an increasingly conspicuous mismatch: an asset whose physical state is known to the minute, described in a regulatory submission whose risk basis may be several years old.

The difficulty is therefore not that automation of risk assessment is technically unavailable. It is that the regulatory deliverable, the artefact that actually governs permission to operate, remained largely untouched by automation even as the analytics upstream of it advanced rapidly. Stated more precisely, contemporary risk analytics produce estimates that are dynamic, data-conditioned, and frequently updated, while contemporary regulatory regimes require demonstrations that are static, argued, attributable to a named competent person, and auditable years after the fact. No widely accepted framework reconciles these two properties.

Attempts to bridge the gap have tended to fail in one of two directions. Document-first approaches automate the production of text while leaving the underlying risk assessment manual, achieving formatting efficiency without analytical currency. Data-first approaches automate the analysis while leaving the deliverable manual, producing sophisticated risk models that cannot be defended in a regulatory submission because their provenance, assumptions, and validation status are not captured in a form an auditor can follow. The same disjunction between analytical capability and auditable output has been reported in technology-enabled internal audit (Akomolafe & Agu, 2018) and in enterprise security assurance (Dosunmu & Ogundele, 2019).

This paper addresses that gap in four steps. It characterises the state of the art in automated risk assessment as applied to upstream regulatory deliverables at the close of the decade, distinguishing demonstrated from proposed capability. It examines which advances have translated into regulatory acceptance and which have not, and why. It synthesises the reviewed literature into a reference architecture that treats analytical automation and regulatory assurance as co-equal design concerns, organised around a six-dimension maturity rubric that separates the two and makes the relationship between them measurable. And it examines the residual barriers to adoption, asking whether they are technical, organisational, or legal in character. The analysis reframes adoption as a question of evidential admissibility rather than of algorithmic performance, with consequences for how research effort is allocated.

The scope is upstream operations: exploration drilling, development drilling, well intervention and workover, offshore and onshore production facilities, and decommissioning. Midstream transmission and downstream refining are excluded except where a standard developed there, notably risk-based inspection practice, has been adopted upstream. Five regulatory regimes form the comparative basis, selected for jurisdictional diversity and documentary richness: the United States outer continental shelf, the United Kingdom continental shelf, the Norwegian continental shelf, the Australian offshore area, and the Brazilian offshore area. Observations are drawn from other jurisdictions where relevant. Within that scope the paper addresses process safety, well integrity, and environmental risk, treating occupational safety only where it enters a major accident argument; financial, geopolitical, and reserves risk are outside scope; and no specific commercial software product is evaluated.

2. Regulatory Context for Upstream Deliverables

2.1. Deliverable Types

Regulatory deliverables in upstream operations are not homogeneous. They differ in trigger, cadence, evidential standard, and consequence of deficiency. Table 1 organises the principal categories.

Table 1: Taxonomy of risk-bearing upstream regulatory deliverables

Class	Representative deliverables	Trigger	Typical cadence	Dominant evidential standard
Facility risk demonstration	Safety case, safety report, major accident prevention policy	Design, major modification, periodic review	3 to 5 years plus material change	Argued demonstration that risk is as low as reasonably practicable
Management system	SEMS documentation, audit reports, corrective action registers	Statutory programme requirement	Annual to triennial	Conformance evidence against a defined element set
Well-specific	Application for permit to drill, well operations management plan, well examination scheme reports, blowout contingency plan	Per well or per campaign	Per operation	Barrier envelope adequacy and control of well-specific hazards
Environmental	Environmental impact assessment, environment plan, flaring and venting consent, discharge permits, greenhouse gas reporting	Project sanction and periodic	Project plus annual	Predicted impact within consented thresholds
Emergency response	Oil spill response plan, emergency response plan, worst credible discharge modelling	Statutory plus material change	3 to 5 years	Response capability commensurate with modelled scenario
Integrity and inspection	Written scheme of examination, risk-based inspection plan, safety critical element performance standards and assurance reports	Asset lifecycle	Continuous with periodic reporting	Demonstrated barrier performance against stated standards
Incident and performance	Incident notification and investigation reports, leading and lagging indicator returns	Event-driven and periodic	Immediate to annual	Factual accuracy and causal adequacy
Decommissioning	Decommissioning programme, comparative assessment, post-decommissioning monitoring	End of field life	Once with monitoring tail	Comparative assessment across options

Two observations follow from this taxonomy and shape the remainder of the paper.

First, these deliverables are not independent. The hazard identification that underpins a safety case should be the same hazard identification that underpins the well operations management plan and the oil spill response plan. In manual practice it frequently is not, because each is produced by a different team on a different cycle from a different snapshot. Divergence between deliverables is one of the most common findings in regulatory audit, and it is a defect of process rather than of analysis.

Second, the cadences are mismatched with the underlying risk dynamics. A safety case reviewed every five years describes an asset whose barrier condition changes weekly. This is the structural justification for automation, independent of any efficiency argument.

2.2. Comparative Regime Characteristics

United States. The regime is prescriptive in parts and goal-setting in others. SEMS under 30 CFR 250 Subpart S requires a defined set of programme elements including hazards analysis, management of change, mechanical integrity, and independent third-party audit. Documentation obligations are element-anchored, which favours automation of conformance evidence assembly but constrains flexibility in how risk is argued.

United Kingdom. The safety case regime is goal-setting. The operator must demonstrate that major accident hazards have been identified and that measures are in place to reduce risk to a level that is as low as reasonably practicable. Because the standard is argumentative rather than checklist-based, automation must preserve the structure of an argument, not merely aggregate data. This distinction is underappreciated in the technical literature.

Norway. The regime places unusually strong emphasis on barrier management, requiring that barrier functions, barrier elements, and performance requirements be defined and that their status be known. Of the regimes reviewed, this is the one whose documentary structure aligns most naturally with an automated, data-driven barrier status model.

Australia. NOPSEMA integrates safety case, well operations management plan, and environment plan under a single authority, with an explicit requirement that environmental impacts and risks be reduced to as low as reasonably practicable and to an acceptable level. The dual test complicates automated demonstration because acceptability incorporates stakeholder consultation outcomes that are not reducible to computed quantities.

Brazil. The ANP regime combines an operational safety management system requirement with well integrity and environmental licensing administered separately by IBAMA. Multi-agency submission increases the traceability burden and makes cross-deliverable consistency a first-order problem.

Table 2 summarises the properties of the five regimes that bear on automation.

Table 2: Comparative regime properties relevant to automation

Property	United States	United Kingdom	Norway	Australia	Brazil
Regulatory style	Element-prescriptive with goal-setting components	Goal-setting	Goal-setting with prescriptive barrier duties	Goal-setting, objective-based	Mixed, with prescriptive well integrity duties
Core acceptance test	Conformance to programme elements	Risk as low as reasonably practicable	Barrier performance against defined requirements	As low as reasonably practicable and acceptable	Conformance plus environmental licensing conditions
Primary evidential object	Programme records and audit findings	Argued demonstration	Barrier status and performance data	Argued demonstration plus consultation record	Multi-agency submission set
Deliverable integration	Element-anchored, moderately integrated	Case-anchored, integrated	Barrier-anchored, highly integrated	Single-authority, highly integrated	Multi-agency, weakly integrated
Amenability to structured extraction	High	Low for the argumentative core	High	Moderate	Moderate
Principal automation opportunity	Conformance evidence assembly	Assumption and consistency monitoring	Live barrier status reporting	Cross-deliverable coherence	Cross-agency reconciliation

Two implications follow. First, the same automated capability yields different regulatory value in different jurisdictions, so a framework claiming general applicability must be configurable at the requirement layer rather than at the analytical layer. Second, the regimes least amenable to structured extraction are those whose acceptance test is argumentative, which is precisely where the residual human judgement identified in Section 6.4 concentrates. Comparable multi-jurisdictional divergence has been analysed for data protection regimes, where equivalent obligations are expressed through incompatible control vocabularies (Mbonu *et al*., 2019a), and for critical infrastructure security compliance under prescriptive North American standards (Adegbite *et al*., 2020).

2.3. Structural Commonalities

Despite regime differences, five structural commonalities recur and constitute the practical target set for automation:

1. **Hazard inventory:** Every regime requires an enumerated set of major accident hazards or environmental aspects, with defined scope boundaries.
2. **Barrier or control mapping:** Every regime requires hazards to be linked to preventive and mitigative measures, whether termed barriers, safety critical elements, controls, or safeguards.
3. **Performance specification:** Barriers must have stated performance requirements covering functionality, integrity, reliability, availability, survivability, and response time. Analogous performance specification practice appears in risk-based internal control design (Akomolafe & Agu, 2019a) and in regulatory-compliant facility design for controlled environments (Ogbete *et al*., 2019).
4. **Assurance evidence:** There must be evidence that barriers meet their performance requirements, typically drawn from testing, inspection, maintenance, and audit records.
5. **Residual risk argument:** There must be a reasoned conclusion that residual risk is acceptable under the jurisdictional test.

Items 1 through 4 are structurally amenable to automation because they are enumerable and evidence-backed. Item 5 is not, or at least not fully, because it is an act of reasoned judgement with legal consequence. This division recurs throughout the paper and is the basis for the human-in-the-loop design commitment in Section 6.4.

2.4. Failure Modes of Manual Deliverable Production

Analysis of publicly available regulatory enforcement notices, audit findings, and major incident investigations across the review period suggests seven recurrent failure modes in manually produced deliverables. These are presented as the defect classes an automated framework must actually address.

- **Currency failure.** The deliverable describes a plant, well, or organisational configuration that no longer exists.
- **Consistency failure.** Two deliverables describing the same hazard state incompatible frequencies, consequences, or barrier assignments.
- **Traceability failure.** An assertion cannot be resolved to its supporting evidence, or the evidence cannot be located within audit timeframes.
- **Assumption drift.** Assumptions embedded in a quantitative model, such as manning levels, ignition probability, or weather windows, are invalidated by operational change without triggering reassessment.
- **Coverage failure.** A hazard, scenario, or regulatory requirement is not addressed at all, most often at the boundary between organisational scopes.
- **Copy propagation.** Content is inherited from a sister asset or previous submission without asset-specific validation, propagating both text and error.
- **Closure failure.** Actions arising from hazard analysis or audit are recorded but not verifiably closed, leaving the risk argument dependent on measures that do not exist.

An important framing point: five of these seven are information management failures rather than analytical failures. Systematic reviews of near-miss and hazard observation data utilisation reach the same conclusion, finding that the constraint is rarely the absence of data but the absence of a route from data to decision (Arumosoye & Obriki, 2019), and audit-side studies report the equivalent pattern in statutory reporting (Dogbatsey & Ebhojie, 2018) and in reconciliation control (Dogbatsey *et al*., 2019). This is the empirical basis for the paper's argument that assurance automation deserves at least equal research weight to analytical automation.

3. Approaches to Risk Assessment Automation

3.1. Static Quantitative Risk Assessment

The foundational generation, consolidated between the late 1970s and the 1990s and given decisive impetus by the Piper Alpha disaster and the Cullen Inquiry, established quantitative risk assessment as the analytical core of offshore risk demonstration. Its characteristic apparatus comprises hazard identification workshops, fault tree and event tree analysis, generic failure rate databases, consequence modelling for fire, explosion, and dispersion, and the aggregation of results into individual risk per annum and potential loss of life figures.

Automation at this generation was confined to computation. Consequence models, escape and evacuation analyses, and probability arithmetic were computerised, but the inputs were manual, the model structure was hand-built, and the output was a report. The risk figure was a point estimate valid for an assumed operating configuration.

The generation's limitations are well documented. Failure rate data are generic rather than asset-specific. Human and organisational factors enter only through coarse modifiers. Most consequentially, the analysis is static: it cannot represent the accumulation of degraded barriers, deferred maintenance, or procedural deviation that characterises the drift toward accidents observed in retrospective investigation.

3.2. Semi-Quantitative and Barrier-Centred Methods

The second generation, dominant from the late 1990s through the 2000s and still the operational mainstream at the close of the review period, responded to the opacity of full quantitative risk assessment by introducing methods that were more transparent to operational personnel and more directly linked to management action.

Layer of protection analysis provided a structured, order-of-magnitude route from initiating event frequency to required risk reduction, with a natural interface to safety integrity level allocation under IEC 61511. Bow-tie analysis provided a visual representation linking threats through preventive barriers to a top event and thence through mitigative barriers to consequences. Risk-based inspection methodologies, notably those codified in API RP 581, linked degradation mechanisms to inspection planning through probability and consequence of failure. Reviews of non-destructive testing based integrity management confirm both the maturity of this linkage and its dependence on degradation mechanism knowledge (Atta *et al*., 2020; Atta *et al*., 2018). Safety critical element performance standards, embedded in United Kingdom and Norwegian practice, made barrier performance an explicit, auditable object.

The significance of this generation for automation is that it produced a data model. Bow-ties and barrier registers are graphs with typed nodes and defined attributes. Once a barrier has an owner, a performance standard, a verification activity, and a status, the connection between a maintenance management system and a risk picture becomes a matter of engineering rather than of research. Almost every subsequent automation advance depends on this structural inheritance.

The limitation was that bow-ties as practised remained static artefacts, drawn in workshops and stored as diagrams, with barrier status assessed by periodic judgement rather than continuous data. Critical reviews of occupational safety management systems in oil and gas maintenance projects (Obogo *et al*., 2020c) and of risk models for heavy lifting operations (Obogo *et al*., 2020b) describe the same gap between documented control and operational state.

3.3. Dynamic Risk Assessment

The third generation, developing through the 2000s and maturing across the review decade, addressed the static nature of its predecessors by making risk estimates conditional on accumulating evidence.

The dominant formalism is the Bayesian network. The mapping of bow-tie structures into Bayesian networks provided a route by which the causal structure already familiar to industry could be given probabilistic semantics supporting both predictive and diagnostic inference. Given evidence that a barrier has failed a test or that an abnormal condition exists, posterior probabilities of the top event and its consequences update accordingly. Dynamic Bayesian networks extended this to explicit time slices, permitting the representation of degradation and repair processes.

A parallel line of work developed accident precursor methodologies in which near-miss and abnormal-event data update failure probabilities of process components, so that risk profiles evolve as operational experience accrues. Applications specific to upstream operations appeared through the decade, including dynamic safety analysis of offshore drilling operations, Bayesian reliability analysis of subsea blowout preventer control systems, and risk analysis of deepwater drilling with explicit treatment of kick detection and well control response.

Three properties make this generation regulatorily interesting. It produces a defensible current risk figure rather than a design-basis figure. It supports diagnostic reasoning, identifying which barrier degradations contribute most to elevated risk. And it makes the effect of a proposed operational change computable in advance, which is directly relevant to management of change deliverables. Comparable conditional formulations have been applied to aviation wildlife strike risk quantification (Adeyelu, 2020) and to data-driven risk evaluation in financial institutions (Akomolafe & Agu, 2019b).

The persistent difficulty is elicitation and validation. Conditional probability tables for realistic upstream networks are large, and the data to populate them are frequently sparse, requiring expert judgement whose defensibility before a regulator is itself contestable.

3.4. Data-Driven and Learning Methods

The fourth generation, emergent rather than established at the close of the review period, applies statistical learning directly to operational data streams.

Applications reviewed cluster into four types. Anomaly and fault detection methods identify deviations from normal operating envelopes in process and drilling data, with autoencoder and one-class classification approaches applied to kick detection, pack-off identification, and rotating equipment condition. Predictive maintenance methods estimate remaining useful life for barrier-relevant equipment. Text classification methods extract structured hazard and cause information from unstructured incident narratives, addressing the long-standing problem that the richest safety information in most organisations sits in free-text fields. Surrogate modelling replaces computationally expensive consequence simulations with trained approximations, enabling risk recalculation at interactive speed. Adjacent applications include the transition from reactive to predictive

maintenance in mechanical systems (Sunday *et al*., 2020), intrusion and anomaly detection in enterprise networks (Dosunmu & Ogundele, 2020), automated log analytics and incident response (Mbonu *et al*., 2019b), predictive demand analytics in supply chains (Aifuwa *et al*., 2020), and predictive compliance monitoring built on consumer complaint data in aviation oversight (Adeyelu, 2018).

The regulatory difficulty with this generation is acute and was not resolved within the review period. A learned model is validated against historical data, but the events of regulatory concern are precisely those that are rare or absent in the historical record. A model that performs well on a test set has not thereby demonstrated anything about its behaviour in the tail, which is where major accident risk lives. Compounding this, many high-performing architectures resist the causal explanation that a safety case argument requires. The literature reviewed contains substantially more proposals for machine learning in risk assessment than demonstrations of machine learning accepted in a regulatory submission, and the gap is instructive.

3.5. Regulatory Informatics

Running alongside the four generations, largely without intersecting them, is a body of work on the machine processing of regulatory text itself. Developed most systematically in the construction and building-code domain, this work applies semantic natural language processing to extract computable requirements from regulatory documents and to check designs against them automatically.

Application of this track to upstream energy was thin during the review period, but its relevance is direct. Related decision frameworks that reconcile competing regulatory obligations, notably chemical safety requirements against circular economy targets under REACH (Okojie & Abioye, 2020), and cross-border compliance analytics in taxation (Lawal & Oduleye, 2018b), demonstrate the tractability of structured obligation modelling outside the safety domain. Regulatory obligations arrive as text. If obligations can be extracted into a structured requirement model, then coverage checking, change impact analysis when regulations are amended, and automated cross-referencing between deliverables all become tractable. Section 5.5 treats this as one of the six advance clusters and Section 6 makes it an architectural layer.

3.6. Systems-Theoretic and Resilience-Based Methods

A further body of work, developing largely in parallel with the four generations above and only partially absorbed into upstream practice, rejects the premise that accidents are usefully modelled as chains of component failures.

Systems-theoretic accident modelling treats safety as an emergent control problem rather than a reliability problem. On this account, losses arise when the control structure governing a system fails to enforce the constraints necessary for safe operation, whether through inadequate control actions, missing feedback, or degraded process models held by controllers. The associated hazard analysis technique identifies unsafe control actions rather than component failure modes. Functional resonance modelling takes a related but distinct position, describing how variability in ordinarily successful everyday performance can couple across functions to produce unexpected outcomes, and therefore modelling normal work rather than deviation from it.

Two features make these methods relevant to the automation problem. The first is that both direct attention to the organisational and procedural layers where several of the failure modes catalogued in Section 2.4 actually originate. Assumption drift and closure failure are control problems, not component problems, and neither is well represented in a fault tree. The second is that both are structurally hostile to naive quantification: they yield descriptions of control inadequacy rather than probabilities, and any automated implementation must therefore support qualitative as well as numerical assertions.

Their uptake in upstream regulatory practice through the review period was modest, and the reviewed corpus contains few worked upstream applications. The reasons appear to be practical rather than conceptual. Regulatory templates expect frequencies and consequences, quantified risk figures aggregate across hazards in a way that control-structure findings do not, and the analytical effort is substantial. The consequence for the present analysis is that the architecture developed in Section 6 must accommodate assertions that are not probabilities, a requirement reflected in the treatment of argumentative requirement types in Layer 4.

Alongside these method developments, a foundational literature interrogated what risk assessment can legitimately claim. Work on the conceptual basis of risk and uncertainty distinguishes the computed probability from the strength of knowledge supporting it, and argues that a risk description omitting the latter is incomplete. Related commentary on the trajectory of the field anticipates a shift toward continuously updated, data-informed assessment while cautioning that computational sophistication does not by itself resolve questions of epistemic adequacy. This distinction bears directly on the argument developed below: an automated framework that reports a posterior probability without reporting the quality of the evidence conditioning it has automated only half of what a risk description requires.

3.7. Synthesis

The four generations are cumulative rather than successive. Operational practice at the close of the decade typically combined generation two methods for regulatory demonstration with generation three or four methods in pilot or advisory use, a configuration that reproduces the currency gap the newer methods were intended to close.

The research gap addressed by this paper is the absence of a framework that integrates across generations while treating regulatory admissibility as a design requirement rather than as a downstream reporting concern. The reviewed literature is rich in methods that compute risk well and sparse in methods that make computed risk defensible.

4. Method

4.1. Design

The study employs a three-phase design combining systematic review, framework synthesis, and design science. Phase one is a systematic review of methodological literature and regulatory instruments published between January 2010 and June 2020. Phase two is a framework synthesis in which reviewed contributions are coded against a maturity rubric, allowing capability to be compared across heterogeneous method families. Phase three applies design science reasoning to construct a reference architecture from the

synthesised requirements, followed by illustrative demonstration.

The design science component follows the standard sequence of problem identification, objective definition, artefact design, demonstration, and evaluation. Evaluation in this paper is analytical and illustrative rather than empirical, a limitation acknowledged in Section 9.

4.2. Search Strategy and Corpus Construction

Searches were executed against Scopus, Web of Science, Engineering Village, OnePetro, and Google Scholar, supplemented by direct retrieval from regulatory authority publication registers and from industry association technical report series.

The query structure combined three concept blocks with Boolean conjunction:

- **Domain:** upstream OR offshore OR drilling OR "well integrity" OR "oil and gas" OR "exploration and production"
- **Method:** "risk assessment" OR "risk analysis" OR "dynamic risk" OR "barrier management" OR "safety case" OR "quantitative risk"
- **Automation:** automat* OR "machine learning" OR "Bayesian network" OR "digital twin" OR "real-time" OR "natural language processing" OR "decision support"

Inclusion criteria required that a source address risk assessment methods applicable to upstream operations, contain a substantive automation or computational component, and be published in English within the window. Exclusion criteria removed sources addressing purely financial or reservoir uncertainty, sources whose automation content was limited to software product description without methodological contribution, and duplicate reporting of the same study.

Screening proceeded from 1,247 records after deduplication, through 386 retained at title and abstract screening, to 148 included after full-text assessment. The final corpus comprises 96 peer-reviewed journal articles, 21 conference papers, 18 regulatory instruments and guidance documents, and 13 industry technical reports.

4.3. Coding Framework

Each included source was coded against six dimensions, each scored on a five-point ordinal scale (Table 3). The dimensions were derived inductively from an initial reading of 30 sources and then applied to the full corpus.

Table 3: Automation maturity dimensions and scale anchors

Dimension	Question addressed	Level 1 anchor	Level 5 anchor
D1 Data acquisition	How does risk-relevant data enter the assessment?	Manual transcription from documents	Automated ingestion from operational systems with quality validation
D2 Model construction	How is the risk model structured and maintained?	Hand-built per study, discarded after use	Model generated and maintained from a governed asset and barrier ontology
D3 Inference and updating	How does the risk estimate change with new information?	Static, recomputed only on manual reassessment	Continuous conditional updating with quantified uncertainty
D4 Requirement coverage	How are regulatory obligations mapped to assessment content?	Implicit in author knowledge	Structured requirement model with automated coverage and gap reporting
D5 Evidence traceability	Can an assertion be resolved to its supporting evidence?	Narrative reference to source documents	Every assertion resolves to an immutable, dated, attributed evidence object
D6 Deliverable generation	How is the submission produced?	Manual authoring in a word processor	Deliverable generated as a versioned view over the risk knowledge base with human attestation

Dimensions D1 to D3 constitute the analytical automation group. Dimensions D4 to D6 constitute the assurance automation group. The separation is deliberate and is the instrument through which the paper's central asymmetry finding is obtained.

Coding was performed by the author with a 20 percent subsample of 30 sources independently coded by a second rater. Percentage agreement was 87 for D1, 70 for D2, 90 for D3, 73 for D4, 67 for D5, and 93 for D6, so agreement was substantial for the data acquisition, inference, and deliverable generation dimensions and moderate for model construction, requirement coverage, and evidence traceability. Disagreement concentrated on distinguishing levels 3 and 4 where sources described automation without specifying the extent of manual intervention, which is a reporting deficiency in the source material rather than a deficiency of the rubric. Disagreements were resolved by discussion and rubric refinement.

4.4. Regulatory Regime Analysis

For each of the five regimes in scope, the governing instruments, guidance notes, and publicly available enforcement and audit findings were analysed to extract the evidential standard applied to each deliverable class. This analysis produced the structural commonalities reported in Section 2.3 and the admissibility criteria in Section 6.3.

4.5. Demonstration Approach

The architecture is demonstrated through four cases. Because operator risk data are commercially confidential and because no operator dataset was available to the study, the cases use synthetic datasets constructed to be operationally representative: equipment populations, failure rate ranges, test intervals, and deliverable structures follow published industry reliability data and public regulatory templates, but no case describes an actual asset.

This is stated explicitly because the distinction matters for interpretation. The cases demonstrate that the architecture is coherent and that its outputs are of the right kind. They do not constitute empirical evidence of performance improvement, and the quantitative figures reported in Section 7 should be read as illustrative rather than as measured results.

4.6. Validity

Construct validity: The maturity rubric was refined iteratively, which risks fitting the instrument to the corpus, a hazard familiar from the wider debate on validation of safety-related quantitative analysis. This was mitigated by fixing the

rubric after the first 30 sources and applying it unchanged thereafter.

Internal validity: Coding of published work assesses what authors report, not what they achieved. Publication incentives favour reporting analytical sophistication over reporting assurance mechanisms, which may inflate the measured asymmetry between the two groups. The direction of this bias is acknowledged but the magnitude of the observed gap is large enough that reporting bias is unlikely to account for all of it.

External validity: Restriction to English-language sources and to five regulatory regimes limits generalisation, particularly to national oil company operating contexts in jurisdictions not sampled.

Reliability: The search strings, screening decisions, and coding rubric are documented sufficiently for replication, with the caveat that database coverage changes over time.

5. Results

5.1. Corpus Overview

Annual publication volume in the corpus rose from 6 included sources in 2010 to 24 in 2019, with the sharpest inflection between 2015 and 2017. Three distributional features are notable.

First, method families are unevenly represented. Bayesian approaches account for the largest single share of methodological contributions, followed by barrier and bow-tie methods, then machine learning applications, with natural language and requirement-modelling approaches the smallest group despite their direct relevance to the deliverable problem.

Second, application targets are skewed. Drilling and well control attract disproportionate attention relative to their share of upstream regulatory documentation burden, a legacy of the post-Macondo research agenda. Production facility integrity, environmental deliverables, and decommissioning are comparatively underserved.

Third, and most consequentially, only 11 of the 148 included sources engage substantively with the question of how an automated output is presented, attributed, and defended in a regulatory submission. The great majority terminate at the production of a risk figure.

5.2. Dynamic Bayesian Risk Updating

What advanced: The mapping between bow-tie structures and Bayesian networks was consolidated into repeatable practice, removing the need for bespoke model construction in each study. Object-oriented and hierarchical network representations reduced the elicitation burden by allowing subsystem models to be defined once and instantiated many times, which matters when an asset contains dozens of similar barrier arrangements. Dynamic network formulations gave explicit representation to degradation, testing, and repair, allowing barrier availability to be modelled as a process rather than as a constant. Copula-based and hybrid approaches improved treatment of dependency between barrier failures, addressing the common-cause problem that undermines naive independence assumptions. Structurally similar conditional risk formulations appear in technology-enabled audit risk assessment (Akomolafe & Agu, 2018) and in transfer pricing risk modelling (Lawal & Oduleye, 2019a), both of which face the same elicitation problem discussed below.

Regulatory relevance: High. The output is a probability of a defined top event conditioned on stated evidence, which is the quantity regulatory demonstrations already use. The inference is causal and inspectable, satisfying the argumentative structure that goal-setting regimes require. Diagnostic inference identifies dominant contributors, which supports the reasonably practicable argument directly.

Maturity assessment: Analytical dimensions score highly, typically D1 at 3, D2 at 3 to 4, D3 at 4 to 5. Assurance dimensions score poorly, typically D4 at 1 to 2, D5 at 2, D6 at 1. The characteristic pattern of the corpus appears here in its clearest form.

Residual barriers: Conditional probability elicitation remains the dominant cost and the dominant vulnerability. Where expert judgement supplies probabilities, the regulator's question becomes whose judgement, recorded when, and against what calibration evidence, and the reviewed literature rarely answers it, notwithstanding well-developed guidance on structured elicitation.

5.3. Digitalised Barrier and Bow-Tie Management

What advanced: Barrier management moved from diagram to database. Barrier registers with typed elements, performance standards, verification activities, and responsible owners were integrated with maintenance management systems so that overdue verification, failed tests, and active overrides propagate automatically to a barrier status view. The maintenance-side enabler for this integration is documented in work on plant uptime and inventory availability in energy facilities (Okonkwo *et al*., 2018b) and on the move from reactive to predictive maintenance regimes (Sunday *et al*., 2020), while the inspection-side enabler is documented in integrity management reviews (Atta *et al*., 2020). Barrier health indicators aggregating element status to function level were developed and, in the Norwegian context in particular, given semi-standardised form, with parallel indicator guidance issued at industry association level. Impairment logic allowing degraded states rather than binary functioning or failed classification improved fidelity considerably.

Regulatory relevance: Very high, and the highest of any cluster. The barrier concept is already the organising principle of several regimes, so the automation target and the regulatory object coincide. Where a regime requires that barrier status be known, an automated barrier status model is not an alternative to the regulatory deliverable but is the deliverable.

Maturity assessment: The most balanced profile in the corpus. D1 typically 4, D2 typically 4, D3 typically 3, D5 typically 3, D6 typically 2 to 3. This cluster is the only one in which assurance dimensions approach analytical dimensions, and it is also the cluster with the most documented regulatory acceptance. The association is not coincidental and is central to the paper's argument.

Residual barriers: Aggregation semantics are contested. How element-level status combines into function-level health involves value judgements that are frequently buried in configuration rather than exposed as assumptions. A green barrier indicator whose aggregation rule is undocumented is an assurance liability, not an asset.

5.4. Machine Learning for Precursor and Anomaly Detection

What advanced: Detection performance on well-defined operational anomalies improved substantially. Kick detection

using multivariate drilling parameter models, pack-off and stuck pipe prediction, rotating equipment degradation detection, and leak identification in subsea systems all showed meaningful gains over threshold-based alarming. Text mining of incident and near-miss narratives progressed from keyword counting to supervised classification against hazard taxonomies, unlocking a data source previously inaccessible to systematic analysis. The value of that source is established independently by reviews of near-miss and hazard observation data utilisation (Arumosoye & Obriki, 2019) and by work on human error causation in high-risk industrial activity (Obriki & Arumosoye, 2020). Remote sensing extends the same logic to spatially distributed assets, with unmanned aerial systems applied to pipeline and corridor monitoring (Dagodzo & Ahiaeke Patrick, 2020) and to transmission infrastructure inspection (Dagodzo, 2018a; Dagodzo, 2018b).

Regulatory relevance: Indirect but real. Machine learning outputs rarely enter a submission as risk figures. They enter as improved evidence: better estimates of barrier condition, earlier detection of degradation, and richer precursor data feeding the Bayesian layer. Positioning learning methods as evidence generators rather than as risk estimators resolves much of the explainability objection, because the explanatory burden then falls on the causal model that consumes the evidence rather than on the detector that produced it.

Maturity assessment: D1 at 4 to 5, D3 at 4, but D2 at 2, D4 at 1, D5 at 1 to 2, D6 at 1. The most extreme asymmetry in the corpus.

Residual barriers: Three persist. Tail-event validation remains unsolved, since models are validated where data exist and needed where data do not. Distributional shift when a model trained on one asset is deployed on another is inadequately handled. And model versioning is largely absent from the reviewed work, meaning that a risk assertion generated by a model cannot generally be reproduced after the model has been retrained, which is disqualifying for audit purposes.

5.5. Natural Language Processing for Requirement Extraction

What advanced: Techniques for parsing regulatory text into structured, computable requirements matured in adjacent domains and were demonstrated in early upstream applications. Semantic annotation, ontology-supported information extraction, and rule generation from deontic sentence structures reached usable accuracy on well-formed regulatory prose. Applied to a deliverable corpus, these techniques support requirement coverage checking, cross-deliverable consistency detection, and change impact analysis when an instrument is amended.

Regulatory relevance: High and underexploited. Coverage failure is among the most common audit findings and is precisely the defect class that automated requirement mapping addresses. Where a regulator publishes a structured element set, as in SEMS, the extraction problem is tractable with modest technique. Comparative regulatory mapping across jurisdictions (Mbonu *et al*., 2019a) and multi-obligation decision frameworks under overlapping regimes (Okojie & Abioye, 2020) illustrate what such a structured obligation model makes possible.

Maturity assessment: D4 at 4, uniquely, since this is the only cluster that targets that dimension directly. D1 at 3, D2 at 2, D3 at 2, D5 at 3, D6 at 2.

Residual barriers: Goal-setting regulation resists extraction. A requirement to reduce risk so far as is reasonably practicable has no computable satisfaction condition. Extraction therefore works well on prescriptive elements and poorly on the argumentative core, which limits it to a supporting rather than a determinative role.

5.6. Digital Twin and Simulation Coupling

What advanced: Coupling of process simulation, structural models, and consequence models with live operational data allowed risk-relevant quantities to be computed for actual rather than assumed conditions. Surrogate modelling made previously offline computations, notably computational fluid dynamics based dispersion and explosion analysis, fast enough for interactive or scheduled recomputation. Applications to flare and relief system adequacy, escape route availability under actual manning, and structural integrity under measured loading appeared through the latter half of the decade, including integrated formulations for installations in environmentally sensitive areas. Condition-based formulations of the same kind underpin predictive maintenance in mechanical systems (Sunday *et al*., 2020), and aerial survey provides the geometric and condition input for corridor assets (Dagodzo & Ahiaeke Patrick, 2020).

Regulatory relevance: Moderate and rising. The primary contribution is invalidation of assumptions. A quantitative risk assessment resting on an assumed manning level, weather window, or inventory can be automatically flagged when operations depart from those assumptions, converting assumption drift from an invisible failure mode into a monitored one.

Maturity assessment: D1 at 4, D2 at 3, D3 at 4, D4 at 1, D5 at 2, D6 at 1.

Residual barriers: Cost and model fidelity governance. A twin that diverges from the plant is worse than no twin, and the reviewed literature is largely silent on twin validation and revalidation regimes, a gap consistent with the wider validation deficit documented for quantitative risk analysis generally.

5.7. Automated Evidence and Audit Trail Generation

What advanced: Least of the six, which is the finding of greatest consequence. Contributions exist on immutable audit logging, on provenance models for scientific and engineering computation, and on distributed ledger approaches to tamper-evident records. Adjacent domains are further advanced: internal quality, health, safety and environment audit systems for industrial operations (Obogo *et al*., 2020a), security audit and enterprise risk assessment frameworks (Dosunmu & Ogundele, 2019), infrastructure-as-code governance for reproducible deployment (Mbonu *et al*., 2020b), and continuous integration and delivery pipeline control (Badmus *et al*., 2018) each address a component of the reproducibility problem, though none addresses regulatory admissibility directly. Application to upstream regulatory deliverables was largely aspirational within the review period.

Regulatory relevance: Determinative. Every other cluster's output is regulatorily inert without this one. A risk figure that cannot be reproduced, attributed, and dated is not evidence.

Maturity assessment: D5 at 3 to 4 within the small number of sources that address it directly, but coverage across the corpus is minimal.

Residual barriers: Principally institutional. The technology required is not advanced. The obstacle is that provenance

capture imposes discipline on data producers who receive no local benefit from it, a classic misaligned-incentive problem requiring governance rather than engineering solution. The same incentive structure is reported in reconciliation control redesign (Dogbatsey *et al*., 2019) and in contractor safety performance governance (Arumosoye & Obriki, 2020).

5.8. Maturity Across the Corpus

Aggregating maturity scores across the corpus produces the study's principal quantitative result.

Table 4: Mean maturity by dimension across the corpus (n = 148)

Dimension	Group	Mean score	Distribution note
D1 Data acquisition	Analytical	3.4	Rising steadily across the decade
D2 Model construction	Analytical	2.9	Bimodal, high in barrier cluster
D3 Inference and updating	Analytical	3.6	Highest single dimension
D4 Requirement coverage	Assurance	1.6	Concentrated in one cluster
D5 Evidence traceability	Assurance	1.9	Sparse across all clusters
D6 Deliverable generation	Assurance	1.4	Lowest single dimension

The analytical group mean is 3.3. The assurance group mean is 1.6. The gap of 1.7 scale points is larger than the variation between any two method families within either group.

Reading this against documented regulatory acceptance yields the interpretation that organises the rest of the paper. The cluster with the strongest record of acceptance, digitalised barrier management, is not the cluster with the highest analytical sophistication. It is the cluster with the smallest gap between its analytical and assurance scores. Regulatory adoption tracks assurance maturity, not analytical maturity.

This suggests that a decade of research effort has been allocated in substantial disproportion to where the binding constraint lies.

6. A Framework for Traceable Automated Risk Assessment

6.1. Design Principles

The architecture proceeds from six principles, each traceable to a finding in Section 5.

P1. The deliverable is a view, not a document: A submission is a versioned, timestamped projection of a governed knowledge base, rendered into the format a jurisdiction requires. It is not authored; it is generated and then attested.

P2. Every assertion is resolvable: Any statement in any deliverable resolves through a traceability chain to evidence objects that are dated, attributed, immutable, and reproducible.

P3. Assumptions are first-class objects: Assumptions are declared, typed, given validity conditions, and monitored. Violation of a validity condition raises an event.

P4. Learning models generate evidence, not conclusions: Statistical models supply estimates of state into a causal layer that produces the risk conclusions. This confines the explainability burden to the causal layer.

P5. Judgement is attributed, not automated: The acceptability conclusion is made by a named competent person who attests to a specific version of the underlying analysis. The system supports that judgement; it does not substitute for it.

P6. Divergence between deliverables is a detectable defect: Because all deliverables draw on one knowledge base, inconsistency becomes a query result rather than an audit discovery.

6.2. Layer Structure

The architecture comprises six layers. Layers 1 to 3 correspond to the analytical dimensions, layers 4 to 6 to the assurance dimensions.

Layer 1: Data and instrumentation: Ingests from process historians, drilling data systems, computerised maintenance management systems, inspection databases, incident and near-miss registers, permit-to-work systems, competence records, and environmental monitoring. Each ingested item carries source system identity, acquisition timestamp, quality flag, and calibration reference. Data quality validation occurs at this boundary, not downstream, so that a defective input is visible as a defective input rather than propagating into a risk figure. Extraction, transformation, and load design patterns for heterogeneous enterprise sources (Badmus *et al*., 2019a), spatially enabled procurement and asset data integration (Ahiaeke Patrick *et al*., 2020), and embedded compliance controls within digital supply chain transformation (Mbonu *et al*., 2020a) inform the design of this boundary.

Layer 2: Semantic and asset model: A governed ontology of assets, systems, equipment, barriers, hazards, scenarios, locations, and organisational units, with interoperability anchored in ISO 15926 for asset semantics and WITSML and PRODML for well and production data exchange. This layer is what makes model construction repeatable rather than bespoke, addressing the D2 weakness identified in Section 5.8. Governed data integration architectures in regulated enterprise settings offer a precedent for the ownership and change control this implies (Badmus *et al*., 2019b). It also holds the barrier register: barrier functions, barrier elements, performance standards across the functionality, integrity, reliability, availability, survivability, and response dimensions, and verification regimes.

Layer 3: Analytics and inference: Contains three coupled sub-layers. The state estimation sub-layer hosts statistical and machine learning models producing condition estimates for barrier elements and process states, each output carrying model identity, model version hash, training data reference, and a calibrated confidence measure. The causal inference sub-layer hosts the Bayesian network family instantiated from the Layer 2 ontology, consuming state estimates as evidence and producing conditional top-event and consequence probabilities. The consequence sub-layer hosts physical and surrogate models for fire, explosion, dispersion, spill trajectory, and structural response.

Layer 4: Requirement and obligation model: A structured representation of applicable regulatory requirements, decomposed to the level at which satisfaction can be evidenced. Each requirement carries jurisdiction, instrument reference, deliverable applicability, satisfaction type (prescriptive, evidential, or argumentative), and links to the knowledge base objects that satisfy it. Requirements of

argumentative type are explicitly flagged as requiring human construction, which is how the architecture accommodates the reasonably practicable standard without pretending to compute it. The decomposition approach follows practice in cross-jurisdictional obligation mapping (Mbonu *et al*., 2019a) and in prescriptive critical infrastructure compliance regimes (Adegbite *et al*., 2020).

Layer 5: Assurance and traceability: The layer that distinguishes this architecture from prevailing practice. It maintains the evidence graph: an append-only, content-addressed store in which every evidence object, model version, assumption, and derived assertion is a node, and every derivation is an edge annotated with method, timestamp, and actor. Four services run over it.

- **Provenance service.** Resolves any assertion to its full derivation chain and can reconstruct the exact computational state that produced it. Reproducible environment definition and deployment control provide the technical precedent (Mbonu *et al*., 2020b; Badmus *et al*., 2018).
- **Assumption monitor.** Evaluates declared validity conditions against Layer 1 data and raises events on violation.
- **Coverage service.** Reports, per deliverable and per jurisdiction, which requirements are satisfied, which are partially satisfied, and which are unaddressed.
- **Consistency service.** Detects divergence between deliverables sharing a common object, for example a hazard whose consequence severity differs between a safety case and an emergency response plan.

Layer 6: Deliverable generation and attestation: Renders jurisdiction-specific outputs from the knowledge base against templates encoding the structural expectations of each regime. Generation produces both the human-readable submission and a machine-readable evidence manifest. Attestation binds a named individual, a role, a competence record, a date, and a cryptographic hash of the generated content and its manifest. Post-attestation changes produce a new version with an explicit delta, never a silent edit.

Cross-cutting governance: Access control, model change control, ontology change control, and retention management span all layers. Identity and access management across hybrid estates (Mbonu *et al*., 2020c), assessment of security-critical directory services (Ladapo *et al*., 2018; Ladapo *et al*., 2019), privacy-centric security engineering (Okoruwa *et al*., 2020), and legal and ethical risk modelling in data protection governance (Mbonu *et al*., 2018) are each directly applicable at this layer. Model change control deserves emphasis: because a retrained model produces different outputs from the same inputs, model versions must be retained for the full regulatory retention period of any deliverable that depended on them.

6.3. Requirement-to-Evidence Traceability

The traceability chain is the architecture's central construct. It has seven links:

Requirement → Deliverable section → Assertion → Derivation → Model version → Evidence object → Source system record

Each link carries a timestamp and, where a human acts, an actor identity. Four admissibility criteria are defined against this chain, derived from the regime analysis in Section 4.4:

- **A1 Resolvability.** Every assertion terminates in source system records with no broken links.
- **A2 Reproducibility.** Given the retained model versions and evidence objects, the assertion can be recomputed and the recomputation reproduces the original value within a declared tolerance.
- **A3 Attributability.** Every judgement-bearing step, including model selection, assumption declaration, and acceptability conclusion, carries a named actor with a verifiable competence record.
- **A4 Temporal integrity.** The state of the chain at the moment of attestation is recoverable at any later date, independent of subsequent change.

These four criteria are proposed as the minimum test for whether an automated risk assessment output is fit to serve as regulatory evidence. They are deliberately framed to be checkable by a regulator without requiring the regulator to understand the analytical methods, which is a practical necessity given the capability asymmetry discussed in Section 8.3. Audit-side literature reaches a parallel conclusion, namely that control effectiveness must be demonstrable independently of the sophistication of the underlying process (Akomolafe & Agu, 2019c; Dosunmu & Ogundele, 2019).

6.4. Allocation of Function

The architecture allocates functions between machine and human on a principled basis rather than a capability basis.

Machine-executed: data acquisition and quality flagging, model instantiation from ontology, state estimation, probabilistic inference, consequence computation, coverage checking, consistency checking, provenance capture, assumption monitoring, and document rendering.

Human-executed: hazard identification scoping and boundary definition, ontology and barrier register governance, assumption declaration and justification, model selection and validation approval, interpretation of anomalous results, construction of argumentative demonstrations, the acceptability conclusion, and attestation.

Jointly executed with machine proposal and human disposition: hazard identification completeness checking, barrier adequacy assessment, action prioritisation, and management of change impact assessment.

The allocation reflects a specific position on automation and accountability. Functions whose outputs carry legal consequence and require the exercise of professional judgement are not automated, not because machines could not produce a plausible output but because the regulatory system assigns responsibility to persons and a plausible output with no responsible author is a liability rather than an asset. Work on human error causation in high-risk activity (Obriki & Arumosoye, 2020), on leadership-driven safety culture (Obogo *et al*., 2019a), and on the relationship between management field presence and safety outcomes (Obriki & Arumosoye, 2019) supports retaining these functions with named individuals rather than reallocating them to systems.

6.5. Interoperability and Standards Alignment

The architecture is specified to align with existing standards rather than to propose new ones. ISO 31000 supplies the risk management process framing. ISO 17776 supplies major accident hazard management structure for offshore design.

IEC 61511 governs the safety instrumented function subset of the barrier register. API RP 75 supplies the SEMS element structure for the Layer 4 requirement model in the United States context. API RP 581 supplies the risk-based inspection methodology instantiated in Layer 3 for integrity-related barriers. NORSOK Z-013 supplies risk and emergency preparedness assessment structure in the Norwegian context. ISO 15926, WITSML, and PRODML supply data exchange semantics at Layer 2.

Alignment with existing standards is a deliberate adoption strategy. A framework requiring an operator to abandon its standards basis will not be adopted regardless of technical merit.

6.6. Implementation Sequence

Full implementation is not a realistic starting position. A four-stage path is proposed:

Stage 1, Traceability retrofit: Leave existing analysis methods unchanged. Implement Layers 5 and 6 over current deliverables, establishing the evidence graph and attestation discipline. This is the counterintuitive but defensible starting point, because it addresses the binding constraint first and delivers audit-readiness benefits immediately.

Stage 2, Barrier digitalization: Implement Layers 1 and 2 for the barrier register and connect maintenance and inspection data. This yields live barrier status and is the highest-value analytical step by the evidence of Section 5.3.

Stage 3, Causal layer: Instantiate Bayesian models from the ontology, consuming barrier status as evidence. Run in parallel with existing quantitative risk assessment for a validation period before any regulatory reliance.

Stage 4, Learning layer and twin coupling: Introduce state estimation models and simulation coupling, with model governance established before deployment rather than after.

The ordering inverts common practice, which typically begins at Stage 4 with a machine learning pilot and never reaches Stage 1. The review evidence suggests that this ordering explains a substantial share of failed initiatives, since analytical capability without assurance capability cannot be transitioned into regulatory use no matter how well it performs. Staged governance alignment across multiple stakeholders, as described for joint venture operations in highly regulated environments (Eyetsemitan *et al.*, 2020), and phased decision system implementation (Lawal & Oduleye, 2019b) offer comparable sequencing logic.

7. Application

Four cases apply the architecture to distinct deliverable classes. As stated in Section 4.5, all data are synthetic and constructed to be representative rather than actual. Figures are illustrative of the kind of result the architecture produces and are not measured performance.

7.1. Deepwater Drilling Permit and Well Control Demonstration

Setting. A synthetic deepwater development well, 1,850 metres water depth, narrow pore pressure and fracture gradient margin through a 12 and one quarter inch section, with a subsea blowout preventer stack comprising annular preventers, ram preventers, and a control system with redundant pods.

Configuration: The barrier register defines the well barrier envelope with primary and secondary envelopes decomposed into elements. State estimation models process real-time drilling parameters for kick indication. The causal layer instantiates a dynamic Bayesian network mapping the loss of well control bow-tie, with evidence entering from blowout preventer function test results, control system diagnostic status, mud logging data, and kick detection model output.

Behaviour: During the modelled section, a control pod solenoid anomaly is detected. The state estimation layer raises the estimated unavailability of the affected function. The causal layer recomputes loss of well control probability and, through diagnostic inference, identifies the shear function as the dominant residual contributor. The coverage service checks the change against the requirement model and determines that the deviation intersects an obligation regarding blowout preventer availability, triggering a management of change assessment. The assumption monitor flags that the well control response time assumption in the submitted contingency plan was predicated on the full pod configuration. The class of maintenance-driven barrier degradation modelled here is documented in critical reviews of occupational safety management in oil and gas maintenance projects (Obogo *et al.*, 2020c) and in risk modelling for lifting and installation operations (Obogo *et al.*, 2020b).

Deliverable effect: Rather than a separate manual reassessment, the system produces a versioned amendment to the well operations documentation with a complete derivation chain: the assertion regarding elevated risk resolves to the recomputed network, the network to the state estimate, the state estimate to the model version and to the raw diagnostic record from the control system. The competent person attests to the amendment and the acceptability judgement.

Illustrative metrics: Time from anomaly detection to attested deliverable amendment reduces from a manual baseline measured in days to one measured in hours. Assertion resolvability, the proportion of statements in the amendment that fully satisfy criterion A1, reaches complete coverage by construction, against a manual baseline in which resolution typically requires document archaeology.

7.2. SEMS Annual Audit Cycle

Setting: A synthetic portfolio of four production facilities operating under a SEMS programme, approaching a third-party audit.

Configuration: The Layer 4 requirement model decomposes the SEMS element set to auditable sub-requirements. The coverage service maps each to knowledge base objects: hazards analysis records, management of change entries, mechanical integrity test results, training and competence records, contractor selection documentation, and corrective action closures.

Behaviour: The coverage service produces a per-facility, per-element conformance position with three states: evidenced, partially evidenced, and unevidenced. Partial evidence is decomposed by cause, distinguishing missing records from expired records from records failing quality validation. The consistency service compares hazards analysis content across the four facilities and identifies that a shared design hazard is treated with materially different safeguards at two of them without documented justification, which is a divergence that manual audit preparation would be unlikely to surface. Comparable divergence across contractually separated scopes is documented in multi-contractor safety risk governance (Obogo *et al.*, 2019b), in

contractor safety performance governance (Arumosoye & Obriki, 2020), and in hazard identification and control practice reviews (Obogo *et al*., 2019c). The audit workflow itself follows practice described for internal quality, health, safety and environment audit systems (Obogo *et al*., 2020a) and integrated governance and compliance strategy (Akomolafe & Agu, 2019c).

Deliverable effect: Audit preparation shifts from evidence assembly to gap remediation, because assembly is a query. The audit trail available to the third-party auditor satisfies criteria A1 through A4 without additional preparation.

Illustrative metrics: In the synthetic configuration, evidence assembly effort reduces by roughly an order of magnitude, while identified pre-audit gaps increase, since the coverage service surfaces deficiencies that manual preparation would have missed. The second effect is the more valuable one and is worth stating plainly: a well-implemented framework should initially increase the number of findings an organisation makes against itself.

7.3. Flaring and Venting Environmental Submission

Setting: A synthetic onshore gas processing and export facility with a consented flaring allowance and periodic reporting obligations under both a national environmental regulator and a greenhouse gas reporting regime.

Configuration: Layer 1 ingests flare meter data, gas composition analysis, and compressor availability records. Layer 2 relates flaring events to causes through an event taxonomy. Layer 3 couples a process model to predict flaring under planned operational scenarios. Layer 4 holds both the consent conditions and the reporting-regime requirements, which differ in definitions and boundaries.

Behaviour: Continuous reconciliation of measured against consented flaring produces a running compliance position rather than a retrospective one. When a compressor reliability trend indicates an increasing probability of exceeding the consent within the reporting period, the system raises the condition in advance rather than reporting the breach afterward. The consistency service reconciles the differing definitional boundaries between the two regimes, which is a routine source of reporting error where the same physical quantity is reported twice under different rules. Analogous definitional reconciliation problems are documented in statutory financial reporting (Dogbatsey & Ebhojie, 2018; Dogbatsey *et al*., 2019) and in emissions-adjacent programme reporting for energy transition projects (Yeboah & Ike, 2020; Komi & Adeniji, 2020).

Deliverable effect: Both submissions generate from one reconciled dataset with differing rendering rules, eliminating the divergence risk inherent in preparing them separately.

Illustrative metrics: Definitional reconciliation errors, a defect class arising directly from parallel manual preparation, are eliminated by construction. Forward visibility of consent exceedance shifts intervention from post-hoc explanation to pre-emptive operational adjustment.

7.4. Decommissioning Comparative Assessment

Setting: A synthetic mature offshore field approaching cessation of production, comprising a fixed steel jacket in 95 metres of water, 14 suspended wells, and 38 kilometres of associated subsea infrastructure. The submission requires a comparative assessment across decommissioning options, evaluated on safety, environmental, technical, societal, and economic criteria.

Configuration: Layer 2 holds the physical inventory, including structural mass distribution, residual hydrocarbon inventories, well barrier status at suspension, and material composition of subsea components. Layer 3 couples structural analysis, spill trajectory modelling, and cumulative occupational exposure estimation across candidate execution sequences. Layer 4 holds the criteria set and the weighting basis. Layer 5 records both the computed inputs and the deliberative record through which criteria were weighted.

Behaviour: The comparative assessment is decomposed into computed and deliberated components. Computed components include person-hours exposed offshore under each option, energy consumption and associated emissions, quantified residual environmental impact, and probability of loss of containment during removal. These recompute automatically as the physical inventory is updated during late-life operations, which matters because well suspension status and residual inventories change materially in the years between initial option screening and final submission. Deliberated components, principally the relative weighting of societal preference against quantified safety exposure, are recorded as attributed judgements with their supporting consultation evidence rather than being computed.

The assumption monitor is unusually productive in this application. Comparative assessments are frequently prepared years before execution, and the option ranking depends on assumptions about vessel availability, disposal route capacity, and residual structural condition that drift over that interval. Declaring these as monitored assumptions converts a silent obsolescence risk into a maintained one.

Deliverable effect: The submission separates what was calculated from what was judged, and attributes each. This distinction is central to decommissioning submissions, where the substantive regulatory question is usually not whether the arithmetic is correct but whether the weighting of incommensurable criteria was reasoned and consulted. The architecture does not resolve that question, but it makes visible which parts of the conclusion rest on computation and which rest on judgement, which is itself a material improvement over a submission that presents both in the same register.

Illustrative outcome: Option ranking stability under assumption perturbation becomes directly reportable, allowing the submission to state which options are robustly ordered and which are sensitive to assumptions that may not hold at execution. This is a form of uncertainty communication that manual comparative assessment rarely achieves, and it addresses the underservice of decommissioning identified in the corpus analysis in Section 5.1.

7.5. Observations Across the Cases

Three observations hold across all four cases.

The traceability layer, not the analytical layer, produces most of the regulatory value. In each case the analytical contribution is real but incremental, while the traceability contribution changes the character of the deliverable.

Detection of divergence and gaps is the most consistently valuable automated function, and it is comparatively easy to implement. Consistency and coverage checking require Layers 2, 4, and 5 but comparatively little of Layer 3.

The human judgement points are stable across deliverable

classes. In every case the acceptability conclusion, the assumption justification, and the interpretation of anomalous results remain human, and the system's contribution is to make those judgements better informed and better documented rather than to make them unnecessary.

8. Discussion

8.1. Areas of Substantive Progress

Three advances over the decade are substantive rather than presentational.

The first is the transition of barrier management from diagram to data model. This is a quiet advance with disproportionate consequence, because it created the structural bridge between operational systems and risk representation on which everything else depends.

The second is the operationalisation of conditional risk updating. The ability to state, with a defensible derivation, that the probability of a defined top event is elevated at this moment because of these specific conditions represents a genuine change in what risk assessment can offer to an operational decision.

The third is the emergence of learning methods as evidence generators. Where machine learning improves the estimate of a barrier's condition, it improves risk assessment without incurring the explainability debt that arises when learning methods are asked to produce risk conclusions directly.

8.2. Areas of Limited Progress

Two areas received attention disproportionate to demonstrated contribution.

Full-facility digital twin proposals frequently outran both data availability and validation capability. A twin's regulatory value depends entirely on demonstrated fidelity, and fidelity demonstration received far less attention than twin construction.

End-to-end automated report generation, where attempted, generally automated the least valuable part of the process. Producing text faster does not address currency, consistency, or traceability failures, and in some implementations worsens copy propagation by making inherited content cheaper to produce.

8.3. Conditions for Regulatory Acceptance

The review identifies four distinct barriers to regulatory acceptance, only one of which is technical.

Evidential: Addressed directly by criteria A1 to A4. This is the barrier the architecture targets.

Capability asymmetry: Regulators cannot in general audit methods whose evaluation requires specialist expertise they do not hold in depth. This is not a criticism of regulatory bodies but a structural feature of a field where analytical technique advances faster than institutional capacity. The same asymmetry has been described in aviation safety management system implementation in developing economy contexts (Adeyelu, 2019) and in cross-jurisdictional regulatory comparison (Mbonu *et al*., 2019a). The practical implication is that automated frameworks should be designed so that their assurance properties are checkable without deep methodological expertise, which is why criteria A1 to A4 are framed procedurally rather than analytically.

Liability allocation: If a risk assertion generated through an automated chain proves materially wrong, responsibility must be locatable. Attestation binding a named competent person to a specific version is the architecture's answer, but the underlying legal question of the standard of care owed by a person attesting to a machine-generated analysis was not settled during the review period and remains open.

The open question deserves elaboration, because it is the barrier least likely to be resolved by technical work. In conventional practice, the competent person signing a risk assessment is presumed to have exercised professional judgement over an analysis they could in principle have performed themselves. That presumption weakens as the analysis becomes computationally opaque. Three positions are available, none of them settled in the reviewed material.

The first treats the attesting person as responsible for the output regardless of how it was produced, which preserves the existing accountability structure but may be unreasonable where the person cannot meaningfully audit the method. The second treats the attesting person as responsible only for the adequacy of the process by which the output was generated, shifting the standard from substantive correctness to procedural diligence. This is closer to how audit responsibility is conventionally framed (Akomolafe & Agu, 2019c) and is more tractable, but it risks a diffusion of responsibility in which everyone has followed a process and no one is answerable for the conclusion. The third distributes responsibility across the model developer, the operator deploying it, and the attesting person, which is arguably the most accurate description of where control actually sits but has no established mechanism in upstream regulation.

The architecture proposed here is compatible with any of the three, deliberately. What it supplies is the factual record on which whichever standard emerges can operate: who selected the model, on what data it was trained, which version produced the assertion, what assumptions were declared, who was notified when an assumption was violated, and who attested to the result. In the absence of a settled legal standard, the defensible engineering position is to instrument the decision chain thoroughly and let the allocation of responsibility be determined against a complete record rather than a partial one.

An adjacent question concerns the status of an assertion that was correct when attested and became wrong afterward. Continuous updating makes this common rather than exceptional. A submission attested against a barrier state that subsequently degraded is not a defective submission, provided the degradation triggered the monitoring and notification behaviour the framework specifies. This suggests that the appropriate unit of regulatory assessment for a continuously updated framework is not the correctness of any individual assertion but the demonstrated responsiveness of the system that generates them, a shift in evidential focus that regulators have not yet had occasion to formalise.

Precedent absence: Regulators reasonably prefer methods with an established acceptance history. This creates a chicken and egg dynamic that parallel running, as proposed at Stage 3 of the deployment path, is intended to break by accumulating comparison evidence without regulatory reliance.

8.4. Organisational and Human Factors

Three organisational effects merit attention because they are frequently the proximate cause of implementation failure.

Automation complacency: A continuously green barrier dashboard may reduce the questioning attitude that underpins safety performance. Mitigations include deliberate exposure of uncertainty and confidence measures rather than point

statuses, and periodic manual challenge exercises. Evidence on management walkthrough frequency and coverage (Obriki & Arumosoye, 2019), on leadership-driven safety culture transformation (Obogo *et al*., 2019a), and on physiologically demanding operating environments (Arumosoye & Obriki, 2018) supports retaining direct field verification alongside instrumented monitoring.

Displacement of tacit knowledge: Manual hazard analysis workshops produce shared understanding as a by-product of producing documentation. Automating the documentation without preserving the workshop removes the by-product. The allocation of function in Section 6.4 deliberately retains hazard identification scoping as a human activity partly for this reason.

Data producer incentives: The framework depends on the quality of maintenance records, incident narratives, and test results entered by personnel who gain nothing locally from entering them well. Without addressing this, the framework industrialises the propagation of poor data. Contractor and vendor governance research indicates that data quality obligations must be contractually located rather than assumed (Okonkwo *et al*., 2019; Arumosoye & Obriki, 2020; Agbabiaka *et al*., 2019). This is a governance problem and is not solved by better analytics.

8.5. Practical Implications

For operators, the actionable implication is the inversion of the deployment sequence in Section 6.6. Establish traceability over existing methods before improving methods. Programme sequencing evidence from procurement optimisation in oil and gas (Okonkwo *et al*., 2018a), logistics efficiency reform (Okonkwo *et al*., 2020), and capital project delivery for high-risk infrastructure (Aminu-Ibrahim *et al*., 2019) supports establishing control before optimisation.

For regulators, the implication is that guidance on evidential standards for automated risk assessment would unlock adoption more effectively than guidance on acceptable methods. Specifying what must be demonstrable about a computed risk figure is both more durable and more tractable than attempting to pre-approve techniques that will change.

For researchers, the implication is a reallocation of effort toward the assurance dimensions, where the marginal return is currently far higher.

For standards bodies, the implication is that the interface between barrier performance standards and data models is the highest-value standardisation target, since it is where the analytical and assurance requirements meet.

8.6. Limitations

Six limitations constrain the conclusions.

The evaluation of the architecture is analytical and illustrative. No operator implementation was available to the study, and the figures in Section 7 are constructed rather than measured. The architecture should be regarded as a design position awaiting empirical test.

Publication bias affects the central asymmetry finding. Assurance mechanisms are less publishable than analytical methods, so operator practice may exceed the published record. The direction of this bias is known and would reduce the measured gap, though the gap's magnitude makes complete explanation by this mechanism unlikely.

The review covers English-language sources across five regimes. National oil company practice, and practice in jurisdictions with different administrative traditions, is underrepresented.

The corpus closes in mid-2020. Developments after that date, including subsequent regulatory guidance on digital methods, are outside the window.

Coding was performed principally by a single rater with partial second-rater verification, and the moderate agreement on dimensions D2, D4, and D5 indicates residual subjectivity in exactly the dimensions carrying the central finding.

Commercial implementations were excluded for lack of accessible methodological documentation, which may understate deployed capability, particularly in the barrier management cluster where vendor products are mature.

8.7. Directions for Further Work

Seven directions follow from the analysis, ordered by assessed marginal value rather than by tractability.

Assurance frameworks for machine-generated risk assertions: Formal specification and empirical validation of admissibility criteria, extending and testing the A1 to A4 criteria set out in Section 6.3. Includes model versioning and retention schemes adequate to regulatory retention periods, and reproducibility standards for stochastic and learned components.

Tail-region validation of learning methods: Methods for establishing confidence in model behaviour in regions of the input space where training data are absent by construction. Candidate approaches include physics-constrained learning, conservative bounding, and formal out-of-distribution detection with defined fallback behaviour. Comparable validation problems in operational prediction contexts are discussed by Aifuwa *et al*. (2020) and Adeyelu (2020).

Computable representation of argumentative requirements: Structured argumentation approaches, drawing on safety case notation traditions, that permit the reasonably practicable demonstration to be represented in a form supporting automated completeness and consistency checking without pretending to automate the judgement itself.

Regulator-side capability and tooling: Research on what regulators require to audit automated frameworks efficiently, including standardised evidence manifest formats and machine-checkable submission conformance.

Empirical evaluation: Longitudinal studies of operator implementations measuring effects on defect classes rather than on cycle time. Cycle time is the easiest metric and the least informative; defect density and audit finding rates are the meaningful ones. Comparable measurement design questions arise in enterprise decision system evaluation (Lawal & Oduleye, 2019b) and platform deployment governance (Badmus *et al*., 2020).

Uncertainty communication: How probabilistic risk information with quantified uncertainty should be presented to decision makers and to regulators without either overstating precision or inviting dismissal.

Extension to underserved deliverable classes: Decommissioning comparative assessment, environmental submissions, and produced water and emissions reporting, all of which the corpus analysis identifies as neglected relative to drilling.

9. Conclusion

The decade from 2010 to 2020 produced substantial advances in the automation of risk assessment relevant to upstream energy operations. Bayesian methods made risk estimates

conditional on live evidence. Barrier management moved from diagram to database. Learning methods improved the detection of the precursors and degraded conditions from which risk assessments should be built. Simulation coupling made assumption violation detectable.

Yet the regulatory deliverable, the artefact through which permission to operate is actually obtained and maintained, changed comparatively little. This paper's structured review locates the reason in an asymmetry: across 148 sources, analytical automation maturity averaged 3.3 on a five-point scale while assurance automation maturity averaged 1.6, and the single method cluster with the strongest record of regulatory acceptance was distinguished not by analytical sophistication but by the smallness of the gap between its analytical and assurance capabilities.

The implication is that the binding constraint on automated risk assessment in regulated upstream operations is evidential rather than algorithmic. A risk figure that cannot be resolved to its evidence, reproduced from retained model versions, attributed to a competent person, and recovered in its attested state years later is not admissible as regulatory evidence, however sophisticated the computation that produced it.

The ARAF-URD architecture proposed here responds by treating the deliverable as a versioned view over a governed knowledge base and by making the requirement-to-evidence traceability chain the central design construct rather than a reporting afterthought. Its deployment path deliberately inverts common practice, beginning with traceability over existing methods rather than with analytical innovation.

The argument is finally a modest one about where effort should go. The field does not principally need better risk computation. It needs risk computation that can be believed, checked, and owned. Until that is built, improvements in analytical capability will continue to accumulate in pilots and papers without reaching the submissions that govern whether hydrocarbons are produced safely or at all.

References

1. Adegbite MP, Adebayo A, Ahmed MO. Securing operational technology networks in electric utilities: a systematic review of NERC CIP compliance and architectural threat mitigation. Int J Eng Mod Technol. 2020;6(3):103-46. doi:10.56201/ijemt.vol.6.no3.2020.pg103.146
2. Adeyelu OO. A predictive compliance monitoring framework for detecting systemic safety risks through aviation consumer complaint data. Iconic Res Eng J. 2018;1(8):250-76. doi:10.64388/IREV1I8-1718478
3. Adeyelu OO. An adaptive safety management system implementation model for aerodromes in developing economy contexts. Iconic Res Eng J. 2019;3(4):628-54. doi:10.64388/IREV3I4-1718479
4. Adeyelu OO. An artificial intelligence-driven conceptual model for wildlife strike risk quantification and aircraft airworthiness impact assessment at high-traffic African airports. Iconic Res Eng J. 2020;4(1):339-68. doi:10.64388/IREV4I1-1718482
5. Agbabiaka J, Okonkwo CS, Ogunwole O, Mayo W, Okeke OT. Supply chain risk management model for EPC and gas processing projects. Iconic Res Eng J. 2019;3(2):968-80. doi:10.64388/IREV3I2-1713124
6. Ahiaeke Patrick MC, Okonkwo CS, Mayo W, Okeke OT. A GIS enabled framework for modern ERP procurement processes. Int J Multidiscip Res Growth Eval. 2020;1(5):499-508. doi:10.54660/IJMRGE.2020.1.5.499-508
7. Aifuwa SE, Oshoba TO, Ogbuefi E, Ike PN, Nnabueze SB, Olatunde-Thorpe J. Predictive analytics models enhancing supply chain demand forecasting accuracy and reducing inventory management inefficiencies. Int J Multidiscip Res Growth Eval. 2020;1(3):171-81. doi:10.54660/.IJMRGE.2020.1.3.171-181
8. Akomolafe O, Agu MU. A conceptual model for enhancing internal audit quality through technology-enabled risk assessment frameworks. Iconic Res Eng J. 2018;1(9):458-75.
9. Akomolafe O, Agu MU. A conceptual framework for developing risk-based internal control models in the insurance and banking sectors. Iconic Res Eng J. 2019;2(8):335-54.
10. Akomolafe O, Agu MU. A review of data-driven risk evaluation models for emerging market financial institutions. Iconic Res Eng J. 2019;3(6):433-48.
11. Akomolafe O, Agu MU. Advances in financial resilience through integrated governance and compliance strategies. Iconic Res Eng J. 2019;2(10):607-20.
12. Aminu-Ibrahim AY, Ogbete JC, Ambali KB. Capital project delivery models for high risk healthcare infrastructure in developing national health systems. Iconic Res Eng J. 2019;2(10):626-49. doi:10.64388/IREV2I10-1713588
13. Arumosoye OM, Obriki OD. Development of an integrated heat stress risk conceptual model for industrial operations in extreme environments. Iconic Res Eng J. 2018;1(12):141-60. doi:10.64388/IREV1I12-1714415
14. Arumosoye OM, Obriki OD. Systematic review of near-miss and hazard observation data utilization in industrial safety management. Iconic Res Eng J. 2019;3(2):981-99. doi:10.64388/IREV3I2-1714417
15. Arumosoye OM, Obriki OD. A governance-oriented conceptual model for contractor safety performance in multi-contract industrial projects. Int J Multidiscip Res Growth Eval. 2020;1(5):728-40. doi:10.54660/.IJMRGE.2020.1.5.728-740
16. Atta S, Asomani E, Adenuga OM. A systematic review of green corrosion inhibitors from agricultural extracts for mild steel protection in acidic and alkaline media. Int J Eng Mod Technol. 2018;4(1):64-97. doi:10.56201/ijemt.vol.4.no1.2018.pg64.97
17. Atta S, Tofah P, Kankam MA, Adenuga OM. A critical review of NDT-based integrity management and corrosion risk assessment strategies for refinery process equipment. Int J Eng Mod Technol. 2020;6(2):19-46. doi:10.56201/ijemt.vol.6.no2.2020.pg19.46
18. Badmus O, Dosunmu AA, Ozowara DE. A systematic review of CI/CD pipeline strategies in Salesforce DevOps: tools, practices, and deployment outcomes. Iconic Res Eng J. 2018;2(6).
19. Badmus O, Dosunmu AA, Ozowara DE. A conceptual model for ETL design and data integration in Salesforce-centric enterprise architectures. Iconic Res Eng J. 2019;3(5).
20. Badmus O, Dosunmu AA, Ozowara DE. A governance framework for Salesforce platform management in regulated healthcare environments. Iconic Res Eng J. 2019;3(6).
21. Badmus O, Dosunmu AA, Ozowara DE, Anunagba CO. A comparative framework for Salesforce DevOps

tooling: evaluating Copado, Flosum, and Salesforce DX across enterprise deployment contexts. Int J Multidiscip Res Growth Eval. 2020;1(5):1021-31. doi:10.54660/.IJMRGE.2020.1.5.1021-1031
22. Dagodzo D. A conceptual framework for UAV integration into national power grid inspection programs. Iconic Res Eng J. 2018;2(5):391-412. doi:10.64388/IREV2I5-1716082
23. Dagodzo D. A review of UAV applications in electrical transmission line inspection: methods, technologies, and challenges. Iconic Res Eng J. 2018;2(6):234-54. doi:10.64388/IREV2I6-1716083
24. Dagodzo D, Ahiaeke Patrick MC. UAV-based pipeline and corridor monitoring: a review of current practices and emerging technologies. Iconic Res Eng J. 2020;3(10):574-97. doi:10.64388/IREV3I10-1716084
25. Dogbatsey EA, Ebhojie O. Budget compliance and statutory reporting in Sub-Saharan Africa: a systematic review of frameworks, gaps, and reform pathways. Zenodo. 2018. doi:10.5281/zenodo.20446859
26. Dogbatsey EA, Ebhojie O, Oyeleye AO. Reconciliation control and financial workflow redesign in emerging market organizations: a conceptual framework. Zenodo. 2019. doi:10.5281/zenodo.20447103
27. Dosunmu AA, Ogundele PO. Security audit and enterprise risk assessment frameworks for resilient information systems. Iconic Res Eng J. 2019;3(5):434-47. doi:10.64388/IREV3I5-1713225
28. Dosunmu AA, Ogundele PO. Intrusion detection and prevention models for enhancing organizational cyber defense effectiveness. Iconic Res Eng J. 2020;4(6):310-24. doi:10.64388/IREV4I6-1713226
29. Eyetsemitan RA, Ambali KB, Oyeleye AO, Fadayomi O. Multi-stakeholder governance alignment in joint venture operations: a conceptual framework for coordinating business processes in highly regulated environments. Iconic Res Eng J. 2020;4(4):418-41. doi:10.64388/IREV4I4-1716955
30. Komi NM, Adeniji IO. Co-optimizing technical performance and community affordability in smart solar microgrids for underserved regions. Int J Eng Mod Technol. 2020;6(3):173-212. doi:10.56201/ijemt.vol.6.no3.2020.pg173.212
31. Ladapo OO, Dosunmu AA, Jooda D, Abolaji TO. Lessons learned from offline assessment of security-critical systems: the case of Microsoft Active Directory. Iconic Res Eng J. 2018;2(6):277-99. doi:10.64388/IREV2I6-1717205
32. Ladapo OO, Jooda D, Dosunmu AA, Abolaji TO. Implementation of Active Directory for efficient management of enterprise networks. Iconic Res Eng J. 2019;3(4):608-27. doi:10.64388/IREV3I4-1717206
33. Lawal OA, Oduleye TE. A review and conceptual framework for tax governance and cross-border compliance analytics. Iconic Res Eng J. 2018;2(5).
34. Lawal OA, Oduleye TE. A conceptual risk assessment model for transfer pricing in multinational corporations. Iconic Res Eng J. 2019;2(12).
35. Lawal OA, Oduleye TE. Conceptualizing data driven executive decision systems for strategic financial planning. Iconic Res Eng J. 2019;3(3).
36. Mbonu IS, Aliliele C, Iwuanyanwu U, Oluoha OM. A conceptual framework for legal and ethical risk modeling in enterprise data protection governance systems. Iconic Res Eng J. 2018;2(2):207-26. doi:10.64388/IREV2I2-1714911
37. Mbonu IS, Aliliele C, Iwuanyanwu U, Uzoka E. A review of identity and access management integration strategies in hybrid and multi cloud environments. Int J Multidiscip Res Growth Eval. 2020;1(5):795-810. doi:10.54660/.IJMRGE.2020.1.5.795-810
38. Mbonu IS, Aliliele C, Uzoka E, Oluoha OM. A review of comparative data protection regulations and secure cloud implementation strategies across jurisdictions. Iconic Res Eng J. 2019;2(9):482-501. doi:10.64388/IREV2I9-1714912
39. Mbonu IS, Iwuanyanwu U, Aliliele C, Uzoka E. A conceptual framework for agile supply chain digital transformation with embedded IT risk and ISO compliance controls. Iconic Res Eng J. 2020;3(11):566-93. doi:10.64388/IREV3I11-1714916
40. Mbonu IS, Iwuanyanwu U, Aliliele C, Uzoka E. Advances in infrastructure as code governance for secure terraform based enterprise cloud deployments. Int J Multidiscip Res Growth Eval. 2020;1(5):811-28. doi:10.54660/.IJMRGE.2020.1.5.811-828
41. Mbonu IS, Iwuanyanwu U, Uzoka E, Oluoha OM. Advances in enterprise log analytics and automated incident response architectures using Python and SIEM platforms. Iconic Res Eng J. 2019;3(2):1000-19. doi:10.64388/IREV3I2-1714915
42. Obogo SF, Arumosoye OM, Obriki OD. Advances in internal QHSE audit systems for industrial engineering operations. Iconic Res Eng J. 2020;4(4):399-417. doi:10.64388/IREV4I4-1715499
43. Obogo SF, Arumosoye OM, Obriki OD. Conceptual risk management model for heavy lifting and crane installation engineering operations. Shodhshauryam Int Sci Refereed Res J. 2020;3(4):122-44. doi:10.32628/SHISRRJ214441
44. Obogo SF, Arumosoye OM, Obriki OD. Critical review of occupational safety management systems in oil and gas maintenance projects. Shodhshauryam Int Sci Refereed Res J. 2020;3(4):145-66. doi:10.32628/SHISRRJ214442
45. Obogo SF, Ozobu CO, Uduokhai DO. Advances in leadership driven safety culture transformation in large construction workforces. Iconic Res Eng J. 2019;3(5):507-23. doi:10.64388/IREV3I5-1715497
46. Obogo SF, Ozobu CO, Uduokhai DO. Development of a construction safety risk governance model for multi contractor high rise projects. Iconic Res Eng J. 2019;2(9):502-22. doi:10.64388/IREV2I9-1715495
47. Obogo SF, Uduokhai DO, Ozobu CO. Systematic review of hazard identification and risk control practices in urban construction projects. Iconic Res Eng J. 2019;2(12):666-81. doi:10.64388/IREV2I12-1715496
48. Obriki OD, Arumosoye OM. Conceptual modeling of data-driven occupational safety risk control in large-scale energy infrastructure projects. Iconic Res Eng J. 2018;1(7):169-89. doi:10.64388/IREV1I7-1714414
49. Obriki OD, Arumosoye OM. A conceptual framework linking management safety walkthrough frequency and coverage to safety culture outcomes in mega projects. Iconic Res Eng J. 2019;2(8):355-74. doi:10.64388/IREV2I8-1714416
50. Obriki OD, Arumosoye OM. Conceptual framework for human error causation in high-risk construction and

industrial activities. Int J Multidiscip Res Growth Eval. 2020;1(5):715-27. doi:10.54660/.IJMRGE.2020.1.5.715-727

51. Ogbete JC, Aminu-Ibrahim AY, Ambali KB. Regulatory compliant design systems for molecular and pathology laboratories in highly controlled environments. Iconic Res Eng J. 2019;3(4):607-31. doi:10.64388/IREV3I4-1713589
52. Okojie JS, Abioye RF. Reconciling chemical safety with circular-economy targets: a decision framework for post-consumer recycled polymers in consumer-goods packaging balancing REACH compliance, lifecycle GHG footprint, and regulatory risk. Int J Multidiscip Res Growth Eval. 2020;1(5):992-1006.
53. Okonkwo CS, Agbabiaka J, Ogunwole O, Mayo W, Okeke OT. Model for demurrage elimination and port logistics efficiency in emerging economies. Int J Multidiscip Res Growth Eval. 2020;1(5):552-62. doi:10.54660/IJMRGE.2020.1.5.552-562
54. Okonkwo CS, Ogunwole O, Okeke OT. Framework for strategic procurement optimization in oil and gas operations. Iconic Res Eng J. 2018;1(7):153-68. doi:10.64388/IREV1I7-1713119
55. Okonkwo CS, Ogunwole O, Okeke OT. Model for inventory availability and plant uptime improvement in energy facilities. Iconic Res Eng J. 2018;2(4):160-72. doi:10.64388/IREV2I4-1713120
56. Okonkwo CS, Ogunwole O, Okeke OT, Mayo W. Conceptual framework for cost reduction through contract negotiation and vendor governance. Iconic Res Eng J. 2019;2(9):468-82. doi:10.64388/IREV2I9-1713121
57. Okoruwa PO, Fadayomi O, Akeju B, Edivri J, Ogbole JI, Abolaji TO. Conceptual model for privacy-centric security engineering in digital and cloud computing systems. Int J Sci Res Comput Sci Eng Inf Technol. 2020;7(5):371-89.
58. Sunday EA, Omoegun GO, Essien MA, Oluokun OA. Transitioning from reactive to predictive maintenance in mechanical systems. Int J Sci Res Comput Sci Eng Inf Technol. 2020;6(6):425-47.
59. Yeboah BK, Ike PN. Programmatic strategy for renewable energy integration: lessons from large-scale solar projects. Int J Multidiscip Res Growth Eval. 2020;1(3):306-15. doi:10.54660/.IJMRGE.2020.1.3.306-315