INDEPENDENT POLICY WHITE PAPER

AI Assurance for Guyana's Extractive Sector

A risk-based governance blueprint for mining, environmental compliance and worker protection

Core thesis. Guyana should allow AI adoption in mining to advance quickly, while requiring stronger assurance as the consequence of an AI-supported decision rises. The governing principle should be simple: no algorithm should weaken legal accountability, environmental evidence, worker rights, or the duty of a named human decision-maker.

Prepared by Ragunauth Ramsaroop

Georgetown, Guyana | 14 September 2026

Evidence cut-off: 14 September 2026. All time-sensitive claims were checked against sources available by this date.

Executive Summary

Guyana is moving toward a more digital state at the same time that its extractive industries are expanding. Budget 2026 identifies AskGov, the MyGuyana Citizen Portal and the e-ID system as core digital initiatives. In November 2025, the Government and Cerebras Systems announced a memorandum of understanding for an AI data centre of up to 100 MW at Wales. These developments show national intent to build digital capacity, but they do not by themselves create a governance regime for AI used in mines, regulators, workplaces or environmental decision-making.

Mining is a strong candidate for early AI adoption because the sector combines large data flows with costly equipment, complex processing, remote operations, environmental obligations and safety-critical decisions. Global operators already report AI use in plant control, computer vision, predictive maintenance, digital twins, hazard reporting and automation. BHP reports that AI-powered plant control at Escondida has saved 3 billion litres of water and 118 GWh of energy [16] since FY2022. This is a company-reported result and should be treated as attributed evidence rather than an independently verified benchmark.

The policy issue is therefore not whether AI belongs in mining. The issue is how Guyana preserves accountability when an AI system influences a blast plan, a worker-risk alert, a tailings decision, an environmental anomaly, an ore estimate, a procurement screen or a regulatory recommendation. Existing mining and environmental rules continue to assign duties to licence holders, agencies and people. AI should strengthen those duties, not blur them.

This paper proposes a Guyana Extractive AI Assurance Framework built around six controls: consequence classification, named human accountability, data provenance, testing and monitoring, explainability and contestability, and incident reporting. It also proposes a four-tier risk model. Low-consequence analytical tools would face light documentation. Material systems would require formal validation. High-consequence systems affecting physical safety, environmental compliance, worker rights or regulatory decisions would require independent assurance and a human decision gate. Certain uses should face red-line restrictions unless law and institutional capacity support them.

The recommended approach is deliberately standards-aligned rather than technology-specific. ISO/IEC 42001 [11] provides an AI management-system model. The OECD AI Principles [12], updated in 2024, emphasise human oversight, transparency, robustness and accountability. NIST's AI Risk Management Framework [13][14] structures action around Govern, Map, Measure and Manage. Guyana can adapt these models to the realities of a resource economy without importing another jurisdiction's entire legal regime.

Five decisions that matter now

1	**Issue sector guidance before high-consequence deployment** GGMC, EPA and the Ministry of Labour should define minimum AI assurance expectations for mining and mineral processing.
2	**Classify by consequence, not by model type** A simple spreadsheet model that influences a tailings response may deserve more scrutiny than a large language model used for drafting.
3	**Keep a named human accountable** Every material AI-supported decision should retain an identified person with authority to accept, reject or stop the output.
4	**Require evidence trails** Data sources, model versions, key inputs, validation results and material decisions should be logged and retrievable.
5	**Build local capability with industry** A small regulatory AI assurance unit, supported by universities and operators, should develop sector competence before rules become more prescriptive.

1. Why this question belongs on Guyana's agenda now

Guyana's policy setting has shifted quickly. The 2026 Budget projects 17.6 percent growth in mining and quarrying [1], with gold and bauxite projected to expand by 5.4 percent and 19.3 percent respectively. The 2026 Budget Speech estimates that mining and quarrying expanded by 21 percent in 2025 [2] and that gold declarations reached 484,321 ounces [2]. These are government estimates and projections, not final audited national accounts.

The same state is investing in digital public infrastructure. Budget 2026 lists AskGov, the MyGuyana Citizen Portal and e-ID-linked services among the planned digitalisation measures. A September 2026 government release continued to describe digital transformation, connectivity and domestic innovation as national priorities. Separately, the Government and Cerebras Systems announced an MOU in November 2025 for an AI data centre of up to 100 MW in Wales [9][10]. The project should be described as an announced partnership, not as an operating facility.

This combination matters. A growing extractive sector, expanding digital infrastructure and a state that is openly adopting AI create a short policy window. Rules written after a technology becomes embedded are harder to enforce, and governance added after an incident tends to be more costly than governance designed into procurement, operating procedures and permits.

What existing Guyanese rules already imply

Mining regulation already assumes identifiable responsibility. GGMC's published process for a mining licence requires technical and financial capability, a work programme and an EPA environmental permit before a licence is issued. Medium-scale permits also require environmental bonds and relevant environmental permits. GGMC describes the EPA as the agency empowered under the Environmental Protection Act to regulate activities that may affect the environment and to grant permits for developmental projects. None of these duties disappears when software supports a decision.

Parliament enacted the Data Protection Act 18 of 2023 [6][7]. Section 1 states that the Act comes into operation on a day appointed by ministerial order, with different dates allowed for different provisions. The public sources reviewed for this paper did not identify a commencement order as of 14 September 2026. This paper therefore treats the Act as enacted legislation with commencement dependent on ministerial order, rather than stating that every provision is already in force.

Evidence status used in this paper

VERIFIED FACT	Directly supported by legislation, an official government publication, an international standard body, or another primary source.
ATTRIBUTED CLAIM	Reported by a company, minister, institution or project sponsor. The attribution is retained.
PROJECTION	Forward estimate, target, planned project or announced timetable. It is not presented as a completed outcome.
POLICY PROPOSITION	Author recommendation or analytical judgement. It is separated from factual description.

2. Where AI creates value, and where it creates exposure

The strongest case for AI in mining is not novelty. It is the ability to combine sensor data, operating history, images, maintenance records and process variables at a speed that supports earlier action. Global mining examples already cover the full value chain.

Use area	Potential value	Primary governance exposure
Exploration and geology	Target generation, core logging support, geospatial pattern detection	Model bias can steer capital toward false positives or underweight local geological knowledge.
Mine planning and processing	Grade prediction, plant optimisation, recovery forecasting, digital twins	Poor data or model drift can affect production, energy, reagent use and reserve confidence.
Mobile equipment and autonomy	Collision avoidance, autonomous haulage, remote operation	Failures can create direct physical safety consequences.
Maintenance	Failure prediction, condition monitoring, computer vision	False negatives can defer maintenance. False positives can waste scarce parts and labour.
Environment and ESG	Water, emissions, biodiversity and land-disturbance monitoring	Models can miss anomalies or create false confidence if sensors, baselines or thresholds are weak.
Workforce and safety	Hazard reporting, fatigue detection, access control, scheduling	Worker privacy, unfair profiling, surveillance and opaque employment decisions become material risks.
Compliance and permitting	Document screening, anomaly detection, inspection prioritisation	Automated recommendations can bias regulatory attention or conceal the reasoning behind a decision.
Supply chain and procurement	Vendor risk, fraud flags, demand forecasting	Poor data can unfairly exclude local suppliers or create opaque risk scores.

Table 1. Illustrative AI use cases and governance exposure. This is an analytical classification, not a statement that every listed use is currently deployed in Guyana.

What global practice already shows

BHP stated in May 2025 that AI-powered plant control at Escondida had saved 3 billion litres of water and 118 GWh of energy [16] since FY2022. In January 2026, BHP also described computer-vision monitoring of conveyors and rail-loading systems and a voice-to-text hazard reporting application linked to historical incident data. These examples show potential operational value, while also illustrating why data quality, alert thresholds, override logic and accountability need formal controls.

The ILO's 2025 global report [15] on AI and digitalisation at work reaches a balanced conclusion. Automation and intelligent monitoring can reduce hazardous exposures and support prevention, but new risks require proactive, evidence-informed and participatory policy responses. For mining, worker involvement is especially important where AI affects fatigue monitoring, productivity scoring, scheduling, access control or disciplinary decisions.

3. A four-tier risk model for extractive AI

The proposed model classifies an AI use by consequence rather than by technical sophistication. The test is not whether a system uses machine learning, generative AI or computer vision. The test is what happens when the system is wrong, unavailable, manipulated or misunderstood.

Tier	Consequence	Examples	Minimum control
Tier 1	Low consequence	Drafting, translation, routine analytics, non-binding summaries	Owner approval, basic data controls, usage log, confidentiality rules.
Tier 2	Material operational	Predictive maintenance, process optimisation, geospatial targeting, supplier risk support	Validation dataset, performance thresholds, drift monitoring, change control, named owner.
Tier 3	High consequence	Autonomous equipment, safety alerts, environmental compliance signals, worker scoring, permit recommendations	Independent validation, human decision gate, fail-safe mode, incident reporting, auditable logs, periodic assurance.
Tier 4	Restricted / exceptional	Fully automated enforcement, unreviewed dismissal decisions, autonomous override of critical safety controls, irreversible regulatory action without accountable human review	Do not deploy unless legislation, due process, technical assurance and competent authority explicitly support the use.

Table 2. Proposed risk tiers. Policy proposition by the author.

The consequence test

- Physical safety: Could a wrong output expose a person to injury or remove a safety barrier?
- Environmental consequence: Could a wrong output delay detection of pollution, land disturbance, water stress or tailings risk?
- Rights and employment: Could the output affect hiring, discipline, termination, surveillance or access to benefits?
- Regulatory consequence: Could the output influence whether a permit, sanction, inspection or enforcement action occurs?
- Financial materiality: Could the output materially affect reserve reporting, production commitments, procurement or public revenue?
- Reversibility: Is the decision easy to reverse, or does a wrong decision create lasting harm?

4. The Guyana Extractive AI Assurance Framework

The framework below draws on ISO/IEC 42001 [11], the OECD AI Principles [12] and NIST AI RMF, then adapts those concepts to the operating reality of mines and regulators. ISO/IEC 42001 [11] provides a management-system structure for responsible AI. The OECD principles emphasise human oversight, transparency, robustness, safety and accountability. NIST organises practical risk management around Govern, Map, Measure and Manage.

1. Consequence classification Every AI use receives a tier before deployment. Reclassification occurs when the use, data or decision authority changes.	**2. Named human accountability** A responsible person owns the use case and has authority to reject, stop or override AI-supported decisions.
3. Data provenance and integrity Material data sources, ownership, quality checks, transformations and retention rules are documented.	**4. Testing, monitoring and change control** Performance is tested before deployment, monitored after deployment and revalidated after material changes.
5. Explainability and contestability Affected people and regulators receive enough information to understand the role of AI and challenge material outcomes where appropriate.	**6. Incident response and assurance** Material failures, near misses, data breaches and unsafe outputs trigger defined escalation, evidence preservation and review.

Figure 1. Six-pillar assurance model. Policy proposition by the author.

Minimum evidence pack for Tier 3 systems

- System purpose, decision boundary and prohibited uses.
- Named accountable owner and named technical custodian.
- Data inventory, provenance record and access controls.
- Validation plan, test results, known limitations and failure thresholds.
- Human override design and safe fallback procedure.
- Model or system version history and change approvals.
- Operational monitoring, incident log and periodic review record.
- Worker, community or regulator communication where a use materially affects them.

5. How Guyana can govern without slowing responsible innovation

A new standalone AI regulator is not the first requirement for extractives. The first requirement is coordinated competence inside institutions that already hold mining, environmental, labour, data and standards mandates. A sector protocol can begin as guidance, procurement requirements and permit-linked expectations, then mature as evidence accumulates.

Institution / actor	Proposed role
GGMC	Lead sector inventory of AI uses in mining. Embed assurance expectations in technical reviews, inspection protocols and guidance for licence holders.
EPA	Set expectations for AI used in environmental monitoring, EIA support, emissions, water, biodiversity and compliance evidence. Require traceable human sign-off for material conclusions.
Ministry of Labour and Manpower Planning	Address worker monitoring, algorithmic management, occupational safety and consultation. Preserve worker voice where AI affects working conditions.
GNBS	Support standards alignment, including awareness of ISO/IEC 42001 [11] and related AI management and risk standards.
Data protection authority / responsible ministry	Clarify lawful data handling as the Data Protection Act moves through commencement and implementation, especially biometric and worker data.
Operators and contractors	Maintain AI inventories, tier systems by consequence, preserve logs, validate high-consequence uses and report material incidents.
Universities and technical institutes	Build local capability in data science, assurance, cybersecurity, mining engineering and environmental analytics.

6. A 12-month implementation path

Period	Action
0-90 days	Create a joint extractive AI working group. Issue a two-page interim advisory. Ask large operators to submit confidential inventories of Tier 2 and Tier 3 AI uses. Identify worker and environmental data categories needing special treatment.
3-6 months	Publish a draft Guyana Extractive AI Assurance Protocol. Pilot the four-tier classification at two large operations and a selected regulatory workflow. Train inspectors and HSE personnel on AI evidence and failure modes.
6-9 months	Introduce a standard AI assurance annex for relevant environmental and mining submissions. Establish a material AI incident reporting channel and a small regulator testing capability.
9-12 months	Evaluate pilot evidence. Publish final guidance. Define where independent third-party assurance is required. Start an annual anonymised sector report on AI use, benefits, incidents and lessons.

What success looks like

Success is not a count of AI systems deployed. Success is a sector in which useful systems scale, high-consequence systems face stronger evidence requirements, failures are detected early, workers know when AI materially affects them, regulators can reconstruct important decisions, and executives retain clear accountability.

7. Policy positions for debate

No black-box excuse for legal duty. A company or agency should not defend a material failure by saying that the model produced the answer. Legal and managerial accountability stays with the responsible institution and person.

Human review must be meaningful. Human approval is not a safeguard when the reviewer lacks time, information, competence or authority to reject the output.

Local suppliers should not be excluded by opaque scoring. AI-assisted procurement should preserve a review path for vendors, especially where local-content objectives or SME participation are affected.

Worker data deserves stricter treatment. Biometrics, fatigue indicators, location data and behavioural scores should face purpose limitation, access control, retention limits and a clear employment-policy basis.

Environmental AI should produce evidence, not replace evidence. Remote sensing, anomaly detection and predictive systems should strengthen inspections and monitoring. They should not erase the need for calibrated instruments, field verification and accountable professional judgement.

Regulators need their own technical capacity. A regulator that depends entirely on operator-supplied models cannot independently test claims, assumptions or system limitations.

8. Factual scrutiny audit

The following audit records the most material factual checks performed before publication. It is included to separate verified evidence from projections, attributed claims and author proposals.

Claim	Status	Audit result
Mining and quarrying projected growth of 17.6% in 2026	Projection	Verified against Guyana Budget 2026 at a Glance. Presented as a projection, not an outcome.
Mining and quarrying estimated growth of 21% in 2025; gold declarations 484,321 oz	Government estimate	Verified against the 2026 Budget Speech, paragraphs 3.10-3.12.
AskGov, MyGuyana Portal and e-ID are 2026 digital initiatives	Government programme	Verified against Budget 2026 and February 2026 public-service briefing. No claim made that all were fully operational by the evidence cut-off.
Guyana-Cerebras project is up to 100 MW	Announced MOU	Verified against Government and Cerebras releases dated 12 November 2025. Described as an MOU/announced project, not an operating data centre.
Data Protection Act 18 of 2023 [6][7] was enacted	Verified law	Verified against Parliament and Official Gazette. Section 1 makes commencement dependent on ministerial order. No commencement order was identified in the public sources reviewed.
Mining licence process requires EPA environmental permit	Verified administrative requirement	Verified against GGMC published mining licence process.
BHP reports 3bn litres water and 118 GWh energy savings from AI plant control	Attributed company claim	Verified against BHP release dated 27 May 2025. Kept as a BHP-reported outcome.
OECD AI Principles [12] updated in 2024	Verified international policy source	Verified against OECD. Principles include human oversight, transparency,

		robustness and accountability.
ISO/IEC 42001 [11] published December 2023	Verified international standard	Verified against ISO. Described as an AI management-system standard, not as a Guyanese legal requirement.
NIST AI RMF functions are Govern, Map, Measure, Manage	Verified framework	Verified against NIST AI RMF resources. Described as voluntary guidance.
ILO says digitalisation can improve OSH while introducing new risks	Verified international research	Verified against ILO 2025 global report. Used as a balanced policy finding.
Four-tier risk model and six-pillar framework	Author proposition	Original analytical framework. Not represented as existing Guyanese law or policy.

References

1. Government of Guyana, Ministry of Finance. Budget 2026 at a Glance. January 2026. Extractive industries and digitalisation sections. https://finance.gov.gy/

2. Government of Guyana, Ministry of Finance. Budget Speech 2026. 26 January 2026. Paragraphs 3.10-3.12 and 4.260-4.261. https://finance.gov.gy/

3. Guyana Geology and Mines Commission. Application for Mining License. Updated 19 November 2025. https://www.ggmc.gov.gy/ggmc-service/application-for-mining-license/

4. Guyana Geology and Mines Commission. Mining Permit [MP]. Accessed 14 September 2026. https://www.ggmc.gov.gy/mining-permit-mp/

5. Guyana Geology and Mines Commission. Environmental Protection Agency. Accessed 14 September 2026. https://www.ggmc.gov.gy/agencies/environmental-protection-agency/

6. Parliament of Guyana. Data Protection Act 18 of 2023 [6][7]. Published 21 August 2023. https://parliament.gov.gy/publications/acts-of-parliament/dataprotectionact18of2023

7. Official Gazette of Guyana. Act No. 18 of 2023, Data Protection Act 2023. 16 August 2023. https://officialgazette.gov.gy/

8. Government of Guyana, Department of Public Information. Gov’t expands digital access with new citizen apps, e-ID system and online services. 4 February 2026. https://dpi.gov.gy/

9. Government of Guyana, Department of Public Information. Guyana and Cerebras forge historic AI partnership to launch a 100MW Data Center. 12 November 2025. https://dpi.gov.gy/

10. Cerebras Systems. Guyana and Cerebras forge historic AI partnership to launch a 100MW Data Center. 12 November 2025. https://www.cerebras.ai/press-release/guyana

11. ISO. ISO/IEC 42001 [11]:2023, Information technology - Artificial intelligence - Management system. Published December 2023. https://www.iso.org/standard/42001

12. OECD. OECD AI Principles [12]. Adopted 2019, updated 2024. https://www.oecd.org/en/topics/ai-principles.html

13. NIST. Artificial Intelligence Risk Management Framework 1.0. 26 January 2023. https://www.nist.gov/itl/ai-risk-management-framework

14. NIST. AI RMF Playbook. Accessed 14 September 2026. https://www.nist.gov/itl/ai-risk-management-framework/nist-ai-rmf-playbook

15. International Labour Organization. Revolutionizing health and safety: the role of AI and digitalization at work. 2025. DOI: 10.54394/KNZE0733.

16. BHP. BHP launches its first Industry AI Hub in Singapore to accelerate AI adoption in the mining and resources sector. 27 May 2025. https://www.bhp.com/

17. BHP. AI is improving performance across global mining operations. 30 January 2026. https://www.bhp.com/

18. UNESCO. Recommendation on the Ethics of Artificial Intelligence and related governance resources. Adopted 2021. https://www.unesco.org/ethics-ai/

19. Government of Guyana, Department of Public Information. Press Release: PM Phillips reaffirms commitment to building Digital Guyana. 7 September 2026. https://dpi.gov.gy/

20. Guyana Bureau of Statistics. National accounts and production datasets. Accessed 14 September 2026. https://statisticsguyana.gov.gy/

Author note

This paper is an independent policy analysis. References to government programmes, companies and international frameworks do not imply endorsement by those institutions. Policy recommendations are the author's own and are intended to support informed debate on responsible AI adoption in Guyana's extractive sector.