June 2026
ISSUE
02
Volume 20

THE ASSET JOURNAL



ASSET MANAGEMENT COUNCIL

DOES AI IMPROVE MAINTENANCE?

An Empirical Comparison of Traditional Machine Learning and Large Language Models for Incident Classification in Electrical Networks

Risk and Value-Focused Asset Management: Probabilistic Cost-Benefit Analysis for Railway Slopes

AI-Powered Maturity Model for Sustainable Asset Management and Data-Driven Strategy

A Probabilistic Approach to Asset Degradation Modelling Using Expected Value

Includes AMPEAK 2026 Supplement

CONTENTS

Editors Column	3
National Chair's Letter	4
Technical Article 1: An Empirical Comparison of Traditional Machine Learning and Large Language Models for Incident Classification in Electrical Networks, Oscar Gibson & Thomas Gordon, Essential Energy	5
Technical Article 2: Risk and Value-Focused Asset Management: Probabilistic Cost-Benefit Analysis for Railway Slopes, Kevin Hang, Sydney Trains	15
Technical Article 3: AI-Powered Maturity Model for Sustainable Asset Management and Data-Driven Strategy, Ramakrishnan Vijayakumaran, CQU	33
Technical Article 4: A Probabilistic Approach to Asset Degradation Modelling Using Expected Value, Hanbit Cho, AssetFuture & Jeremy Ngo, Woolworths	44
Star Profile: Shishir Tripathi CSAM	50
Star Profile: Rihannah Hesami CSAM	52
Star Profile: Ali Ahmadi CPAM	54
Chapter News	56
Member Application	60
Corporate Members	63
AMPEAK26 Supplement	65

Publisher: The Asset Management Council
Chief Editor: Ernst Krauss
Publication Design: AMCouncil
ISSN: 1834-3864

For all enquires please contact:
publications@amcouncil.com.au
Asset Management Council
PO Box 2004 Oakleigh Vic 3166
Tel 03 9819 2515



ERNST KRAUSS

EDITOR IN CHIEF

AI for Maintenance sounds like the answer many of us have been waiting for. But what should we realistically expect from AI in this context? Too often, we look for a 'universal solution' that tells us everything we think we want to know, rather than what we truly need to know, and we assign that role to AI. Before we apply AI across the business, we need to define clearly what Maintenance needs to know and what the wider business needs to know about Maintenance.

Many practitioners and analysts have found that AI itself is often not the problem. Failures usually stem from limited capability in using it effectively and a poor understanding of where it should be applied. In Maintenance, we need methods and tools that provide meaningful insight. This includes access to key information on asset performance, utilisation, condition, and reliability and more, ideally at the touch of a button. It also includes visibility of workflow issues familiar to planners, supervisors, and superintendents. At the same time, finance teams want a clear view of maintenance expenditure so they can support better investment decisions. The key question is: what exactly should AI be searching for and answering?

In my view, defining the information we need is only part of the challenge; we must also ensure the data is complete and reliable. Many organisations still work with fragmented, inconsistent datasets spread across multiple systems, which can lead AI—whether agentic or anthropic—to produce biased or misleading results unless the boundaries of the data being queried are clearly set. Claims that AI can reduce maintenance effort by 10% to 30% are certainly attractive, but we need to ask what those reductions would actually mean in practice. Could risk increase? Would the savings be sustainable? How do we guarantee performance? These are some of the fundamental questions for anyone considering AI in Maintenance. It should also go without saying that every analytical result, including AI-generated opportunities, still requires human judgement to validate the actions to be taken.

Given the complexity of Maintenance, particularly if it is to become a future profit centre, we first need to understand the foundations required for AI to deliver good results. Asset registers and asset information must follow a consistent structure and syntax across the data repositories and the organisation so that people can access the right information reliably. That will require effort, as will defining clear boundaries for AI and aligning purpose across departments to protect the integrity of the information retrieved. Before expecting AI to deliver its full promise, we may achieve early value by using it to identify inconsistencies in maintenance strategies and tasks for similar assets across the business. It may also help reveal planning, scheduling, and other workflow issues that often remain hidden. Used this way, AI could become a practical training ground for engineers, maintainers, support teams, and managers as they shape the future of Maintenance in the age of AI.

As always, we value your experiences and feedback and trust you will enjoy this edition of "The Asset".



WENDY MCPATE

NATIONAL CHAIR

Artificial Intelligence is rapidly becoming one of the most discussed topics in our profession. This edition of The Asset Journal explores a question many organisations are now asking: Does AI improve maintenance?

Asset management professionals have long used data, analytics and condition monitoring to support maintenance decisions. The emergence of artificial intelligence represents something different. Rather than simply providing more data, AI has the potential to change how we access information, develop insights and make decisions.

During recent discussions with infrastructure and government leaders in Washington DC, one theme emerged consistently. The focus was not on replacing people or automating judgement. Instead, the conversation centred on how AI can help organisations better utilise information, improve productivity and support more informed decision making.

This is particularly relevant for asset management. Maintenance is ultimately about sustaining asset performance, managing risk and delivering value throughout the asset lifecycle. While AI has the potential to provide faster access to information and new insights, responsibility for balancing cost, risk and performance continues to rest with people.

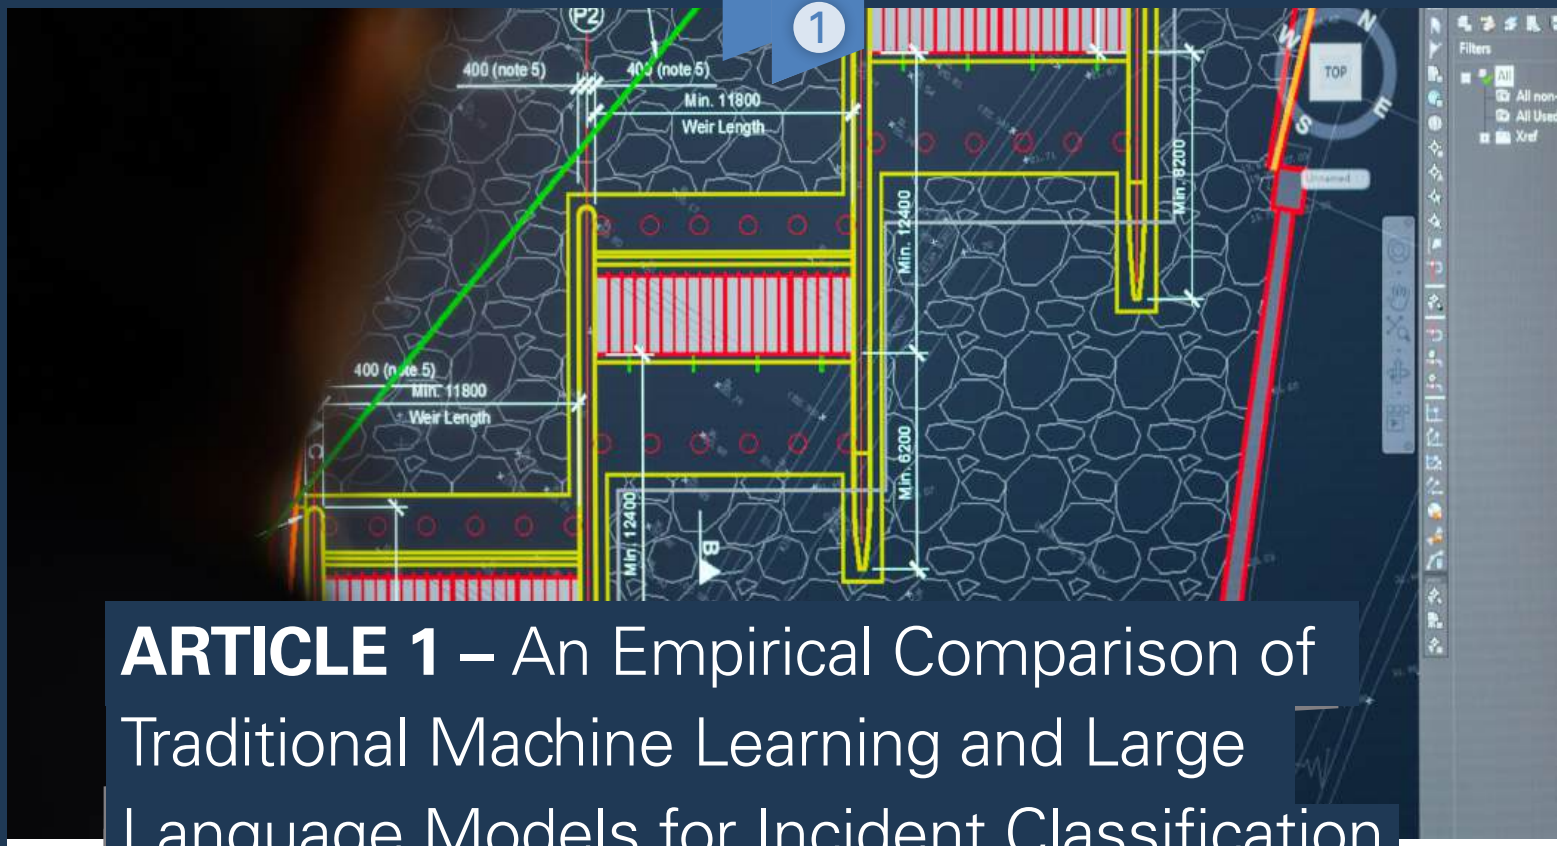
Many organisations are also discovering that the greatest opportunities may lie beyond maintenance itself. AI is enabling faster access to knowledge, reducing time spent searching for information, supporting analysis and allowing people to focus more of their effort on strategy, judgement and decision making. At the same time, it is raising important questions about governance, accountability, trust and workforce capability.

As asset managers, perhaps the question we should be asking is not whether AI improves maintenance, but how we ensure it improves decisions.

Across Australia and internationally, organisations are exploring how AI can enhance productivity, resilience and asset outcomes. As this technology continues to evolve, our challenge will be to embrace its opportunities while maintaining the principles of sound asset management that have always underpinned long term success.

I hope this edition provides practical insights, diverse perspectives and useful examples that help inform your own thinking as we navigate this rapidly evolving landscape together.

TECH 1



ARTICLE 1 – An Empirical Comparison of Traditional Machine Learning and Large Language Models for Incident Classification in Electrical Networks

Oscar Gibson and Thomas Gordon, Essential Energy

SUMMARY

As organisations invest in Artificial Intelligence (AI) for asset management, a critical question emerges: do newer Large Language Model (LLM) approaches justify replacing established machine learning (ML) techniques? This study provides an empirical comparison of BERT-based text embeddings with neural network classifiers versus modern LLM models for classifying electrical switching incidents, highlighting practical trade-offs to inform AI investment decisions.

Essential Energy records network incidents using both structured fields (designed for authoritative reporting and analytics) and unstructured free-text comment fields entered by field operators. Initially, a rules-based automation approach was used to identify switching incidents from unstructured comments through keyword matching and value filtering. While this achieved high recall, it produced many false positives. To address this, a two-stage hybrid system was implemented: rules-based filtering for comprehensive candidate identification, followed by ML-based false-positive verification. This second stage compared a trained neural network classifier against an LLM-based classifier using prompt engineering.

Evaluation was conducted on 100 rules-identified comment records, manually validated by subject-matter experts to establish ground truth. Performance was assessed using precision, recall, and F1-scores, reflecting real- world operational decision thresholds. Both approaches incorporated Named Entity Recognition (NER) to anonymise text and meet data governance and privacy requirements.

Results demonstrate complementary strengths. The neural network, trained on limited data, achieved 75% precision but catastrophically low 10.2% recall, missing 90% of genuine incidents. It had an average inference time of 44 milliseconds and an estimated processing cost of 0.01 cents per incident.

The LLM-based approach achieved 76.8% precision and 89.8% recall, provided interpretable reasoning, handled ambiguous language more flexibly, and required no model training, at the cost of higher inference latency (1.2 seconds) and processing cost (0.1 cents per incident). Compared to a rules-only baseline precision of 59%, the LLM approach significantly improved correct identification of switching incidents, albeit with increased computational expense.

These findings support evidence-based AI investment decisions by illustrating trade-offs between accuracy, latency, interpretability, and total cost of ownership when applying AI to unstructured, contextual operational data, without extending conclusions beyond the intended data scope.

Keywords: Artificial intelligence, large language model, machine learning, neural networks, incidents

1 INTRODUCTION

Asset management organisations investing in AI face a critical decision: do modern Large Language Models (LLMs) justify replacing established machine learning techniques? This paper presents an empirical comparison of three classification approaches for Essential Energy's incident data: a rules-based baseline, a neural network classifier, and an LLM based classifier.

Essential Energy owns and operates the electricity distribution network across approximately 95% of New South Wales and parts of Queensland, delivering electricity to more than 900,000 customers. The network comprises approximately 1.4 million power poles and over 200,000 kilometres of powerlines. To maintain a reliable and safe electricity supply, the network is subject to continuous maintenance and upgrade activities. A critical component of this work involves switching operations, which temporarily de-energise sections of the network. These operations consist of the planned isolation and reconfiguration of electrical assets to enable safe access for construction, maintenance, and inspection activities.

There are strict protocols governing switching operations, designed to mitigate asset operational discrepancies, unintended outages, and potential risks. Despite these controls, the scale and geographic extent of the network means that incidents do occur and must be systematically identified, investigated, and analysed to support continuous network improvement.

Essential Energy records all incidents, both switching and non-switching related, in a central incident management system. Incident records include free-text fields that allow employees to describe events and provide contextual detail. While this unstructured text is rich in qualitative information, it is not designed for direct analytical use, making use of it for automated incident classification challenging and a substantial manual workload. This field can contain high-value contextual information that is not captured in structured variables.

Historically, a rules-based automation system was used to identify and report switching incidents at scale. This approach combined keyword matching within free-text descriptions (e.g., the presence of terms such as "switching") with structured filters, such as incidents recorded by System Control staff, to flag potential switching events. This method achieved high recall, successfully capturing the majority of true switching incidents. However, it also generated a substantial number of false positives. Expert review confirmed that 59% of incidents flagged by the rules-based system were genuine switching incidents, limiting the approach's effectiveness for reliable downstream analysis.

To reduce the number of false positives, two machine learning approaches were implemented and evaluated: a traditional neural network classifier built on BERT text embeddings, and a Large Language Model (LLM)–based approach using prompt engineering. Both approaches demonstrated potential for improving precision. However, they introduced distinct operational trade-offs. The neural network approach delivered low-latency inference at scale but required a substantial volume of labelled examples for training. In contrast, the LLM-based approach enabled zero-shot classification (Brown, et al., 2020) without the need for labelled training data but incurred higher computational and inference costs. When classifying incidents at scale, these trade-offs needed to be considered to ensure operational requirements are met in production. Therefore, these

differences necessitated a rigorous empirical comparison to inform production deployment decisions in a high-volume operational environment.

This study is framed around practical questions faced by utilities investing in AI-enabled operational analytics: what levels of accuracy improvement justify increased computational cost? Under what conditions is the investment in labelled training data warranted? How does model interpretability contribute to operational decision-making and downstream value?

The contributions of this paper are as follows:

- An empirical performance comparison across 100 expert-validated incidents, evaluating precision, recall, and F1-score.
- A computational cost analysis, including per-incident inference cost and implications for total cost of ownership.
- A qualitative assessment of model interpretability and handling of edge cases
- Practical production deployment insights
- A framework for asset management firms to assess the feasibility of using LLM in production.

The remainder of this paper is organised as follows. Section 2 describes switching incidents, data characteristics, and the evaluation methodology. Section 3 presents classification performance results, computational metrics, and qualitative findings. Section 4 discusses observed trade-offs and provides guidance for practitioners. Section 5 concludes with implications for AI investment in asset management.

2 PROBLEM AND METHODOLOGY

2.1 Switching Incidents in Electrical Networks

Switching operations are planned activities undertaken to isolate, reconfigure, or restore sections of the electricity distribution network. These operations support safe maintenance, fault response, and network upgrades.

In the 2024-25 financial year, approximately 17,000 planned outages and 30,000 unplanned outages were managed across Essential Energy's distribution footprint. Each outage required switching operations to be performed in accordance with established operational and safety procedures. Incident reporting processes are in place to identify, record, and learn from any deviations that may occur during these activities.

For each switching operation, field operators and system control staff follow a structured, step-based process that emphasises verification at critical points. This process includes the following key steps:

1. Reading the switching instruction.
2. Taking the instruction to the point of operation.
3. Checking the asset label against the instruction.
4. Preparing to operate.
5. Re-checking the asset label against the instruction.
6. Operating the device.
7. Confirming the device has operated as intended.
8. Marking the step as complete on the instruction.

This process is designed to reduce the likelihood of switching errors through systematic verification and confirmation at each stage of the operation.

A switching incident is defined as any unintended or incorrect action, outcome, or system behaviour during switching operations. These can happen in planning, preparation, or execution phases of the sequence. Examples of switching incidents might include:

- **Asset Operational Discrepancies:** Occurrences involving the unintended operation of switching devices or observed deviations from the prescribed switching sequence.
- **Information Exchange Gaps:** Occurrences involving discrepancies or ambiguities in the communication of operational data and instructions between field operators and system control centres.
- **Procedural Variance:** Observed departures from standard sequential verification steps or deviations in the alignment of physical isolation measures with planned configurations.
- **Technical Performance Anomalies:** Instances of unexpected asset behaviour or hardware response characteristics that deviate from anticipated operational parameters.

The potential consequences of switching incidents include risks to personnel safety, unplanned customer outages, extended restoration times, and non-compliance with regulatory or procedural requirements. Reducing the frequency and severity of such incidents is therefore a critical objective for network operators.

To support this objective, incident frequency and trends must be monitored over time to enable informed oversight of network safety and reliability performance and to support continuous improvement initiatives. However, trend analysis alone is insufficient without an understanding of the underlying factors contributing to individual incidents. Identifying these contributing factors requires consistent and accurate incident classification. While structured incident fields capture key operational attributes, they do not always fully describe the circumstances surrounding an event. In a subset of cases, free-text comments provide additional contextual detail that can clarify intent, sequence of actions, or environmental conditions relevant to the incident.

Accurate classification therefore depends not only on the availability of structured data, but also on a clear understanding of how unstructured data can be selectively leveraged to extract relevant information. Recognising both the value and the limitations of these data sources is essential to ensure that insights derived from incident analysis are reliable, interpretable, and appropriately applied to process improvement, system design, and targeted training interventions.

2.2 Data Characteristics

When an incident is reported within Essential Energy, the reporting staff member may additionally provide a free-text description of the event. Incident reports cover a broad range of occurrences across the network, including safety events, third-party interactions with network assets, and asset-related incidents. Switching incidents represent only a subset of these reports and are not captured through a dedicated structured field, instead being identifiable primarily through the free-text description.

The unstructured nature of this text presents challenges for automated classification. Switching incidents must be distinguished from a large volume of non-switching incidents based solely on narrative descriptions, which may vary in terminology, level of detail, and clarity.

Historically, switching incidents were identified using a rules-based approach. This involved filtering reports by relevant personnel and task types, combined with keyword matching within the description field. While effective as an initial screening method, this approach achieved a precision of approximately 59%, with a proportion of non-switching incidents incorrectly identified. This reflects factors such as the diversity of operational tasks performed by staff, the use of overlapping terminology (for example, "switching" in non-switching contexts), and cases where switching activity is implied rather than explicitly stated in the description.

These characteristics highlight the limitations of purely rules-based methods and motivate the exploration of more advanced ML and LLM approaches for incident classification.

2.3 Approaches Evaluated

Three approaches to automating incident classification were evaluated, reflecting a progression from rules- based heuristics, through traditional supervised machine learning methods, to more recent large language model techniques. Each approach involves different trade-offs in terms of classification performance, computational requirements, data dependencies, and interpretability.

Rules-Based Baseline – The existing rules-based approach applies SQL queries that combine keyword matching within free-text incident descriptions with filters on structured attributes. These filters include relevant personnel roles (for example, incidents recorded by System Control staff) and associated activity types. Keywords such as "switching" are used as indicators to identify incidents that may relate to switching activities. The approach is intentionally configured to prioritise high recall, with the objective of capturing the majority of genuine switching incidents, while accepting a higher rate of false positives at the initial screening stage. As described in Section 3.2, this baseline approach achieved a precision of approximately 59% when evaluated against incidents validated by subject matter experts.

Traditional ML: BERT + Neural Network – This approach applies a traditional supervised machine learning pipeline in which a pre-

trained BERT-base-uncased model (Devlin, et al., 2019) is used to generate text embeddings from incident descriptions. These embeddings are then classified using a feedforward neural network. The model was trained on a dataset of 838 labelled incidents, of which 699 were classified as EXCLUDE, reflecting the natural class imbalance present in incident reporting data. This approach requires a curated labelled dataset for training but supports low-latency inference (approximately 44 milliseconds per incident) and low per- incident computational overhead once deployed.

Large Language Model: Two-Stage Prompt Engineering – This approach uses an LLM to perform zero- shot incident classification through structured prompt design, without requiring task-specific model training. The solution implements a two-stage process intended to first identify potential switching incidents and then extract relevant attributes from those identified as genuine switching incidents. In this study, the LLM component used the databricks-gpt-oss-120b model API.

Stage 1 performs an initial screening and produces three outputs:

1. A classification outcome (INCLUDE or EXCLUDE)
2. A qualitative confidence rating (HIGH, MEDIUM, or LOW)
3. A natural-language explanation describing the basis for the classification.

The prompt incorporates subject matter expert definitions of switching incidents, examples of relevant incident characteristics, and a predefined response structure to promote consistency. Incidents classified as EXCLUDE are removed at this stage, while those classified as INCLUDE proceed to the second stage.

Stage 2 assigns a set of structured attributes to incidents identified as switching-related. These attributes include the point of occurrence (Planning, Preparation, or Execution), adherence to the 8-step switching process (Yes or No), task complexity (Low, Medium, or High), location (Office or Field), and planning status (Planned or Unplanned). The resulting structured data supports downstream activities such as trend analysis, training needs assessment, and compliance-related reporting. This approach does not require labelled training data and provides interpretable, case-specific reasoning alongside each classification. However, it involves higher inference latency (approximately 1.2 seconds per incident) and higher per-incident processing cost when compared with the neural network–based approach.

2.4 Evaluation Methodology

To evaluate the three classification approaches, a reference ("ground truth") dataset was created by randomly sampling 100 incidents from those initially flagged by the rules-based method. Subject matter experts manually reviewed each sampled incident to determine whether it represented a genuine switching incident or a false positive. This review process produced a labelled dataset used for comparative performance evaluation.

Prior to evaluation, all incident descriptions were anonymised to protect employee privacy and to support a fair comparison between models. This preprocessing step is particularly important when working with safety- related incident data, as free-text descriptions may include identifying information about personnel involved. In addition to privacy considerations, anonymisation helps mitigate the risk of unintended bias in model behaviour. Without anonymisation, a model could inadvertently learn spurious associations between specific names and incident classifications (for example, inferring higher task complexity or execution-phase involvement based on name frequency rather than incident characteristics). Removing personal identifiers ensures that classification decisions are driven by incident context rather than individual identifiers.

The anonymisation process used SpaCy Named Entity Recognition (NER) to identify and replace personal names and employee identifiers with a generic "PERSON" label. SpaCy identifies entities based on linguistic structure and contextual cues, including references to people, organisations, and locations. While this approach provides a reasonable level of privacy protection, it has practical limitations. Automated detection may not capture identifiers specific to operational contexts, such as crew designations, work group references, or informal workplace terms. Conversely, some entities may be incorrectly identified, resulting in partial over-anonymisation. The anonymisation approach

was therefore designed to balance preservation of incident context with protection of employee privacy. Importantly, the same anonymisation pipeline was applied consistently across both the neural network and LLM-based approaches prior to evaluation. This ensures that observed performance differences reflect model behaviour rather than variations in data preprocessing.

Classification performance was assessed using multiple metrics to capture both analytical quality and operational relevance. Precision measures the proportion of incidents classified as switching-related that are genuine, providing an indication of false positive reduction. Recall measures the proportion of genuine switching incidents correctly identified, which is particularly important from a safety and learning perspective, as missed incidents may represent lost opportunities for insight. The F1-score provides a balanced measure that combines precision and recall.

3 RESULTS

3.1 Classification Performance

Table 1 summarises the classification performance metrics for the three approaches evaluated using the expert- validated dataset. Within this evaluation, the LLM-based approach recorded the highest precision at 76.8%, representing a 17.8 percentage point improvement over rules-based baseline (59%). However, the neural network exhibited severe performance degradation, achieving only 10.2% recall despite 75% precision. These results are computed on a validation set consisting solely of incidents that were first identified by the rules-based filters, so they reflect performance on rules-flagged candidates rather than all recorded incidents.

The neural network's low recall is attributed to severe training data limitations. Of the 838 training incidents, only 139 were INCLUDE examples (16.6%), creating pronounced class imbalance. Despite class weighting and resampling during training (Chawla, et al., 2002), the model became heavily biased toward predicting EXCLUDE. Of 100 validation incidents, the neural network predicted only 8 as INCLUDE (6 correct, 2 false positives), missing 53 genuine switching incidents. This 89.8% false negative rate renders the approach operationally unviable, as missing genuine switching incidents defeats the safety and compliance objectives of incident tracking.

In contrast, the LLM leveraged pre-trained representations of technical language and operational context, enable effective zero-shot classification without task-specific training examples. The LLM correctly identified 53 of 59 true switching incidents (89.8% recall) while reducing false positives from 41 (rules-based) to 16, achieving the best F1-score (0.828).

The rules-based baseline achieved perfect recall (100%) by design, flagging all 100 sampled incidents for manual review. However, this resulted in 41 false positives, a 41% false positive rate that originally motivated the ML-based approaches.

Table 1: Performance metrics on 100 expert-validated incidents sampled from incidents flagged by rules-based filters

Metric	Rules-Based	BERT + NN	LLM
Precision	59%	75%	**76.8%**
Recall	100%	10.2%	**89.8%**
F1-score	0.742	0.179	**0.828**
Accuracy	59%	45%	**78%**
Specificity	0%	**95.1%**	61%
True Positives	**59**	6	53
False Positives	41	**2**	16
True Negatives	-	**39**	25
False	-	53	**6**

Recall performance demonstrates critical differences between the approaches. The neural network's 10.2% recall indicates a catastrophic failure to identify genuine switching incidents, missing nearly 90% of true positives. The LLM maintained high recall (89.8%), missing only 6 incidents that lacked explicit switching terminology or contained genuinely ambiguous descriptions.

F1-scores demonstrate that the LLM achieved the best balance between both precision and recall (0.828), followed by the rules-based approach (0.742). The neural network's F1-score of 0.179 reflects its inability to balance the two metrics, prioritising false positive avoidance at the expense of missing genuine incidents.

3.2 Computation Performance

Table 2 compares the computational characteristics of the ML and LLM classification approaches. The neural network–based approach recorded the lowest inference latency, at approximately 44 milliseconds per incident, making it well suited to high-throughput or batch processing scenarios.

The LLM-based approach exhibited higher inference latency, at approximately 1.2 seconds per incident. This reflects the use of a two-stage evaluation process, involving sequential calls for initial screening and subsequent attribute extraction for incidents classified as INCLUDE, as well as the computational demands associated with large-scale language models.

Per-incident processing costs reflect the observed differences in computational requirements. The neural network–based approach incurs minimal per-incident cost (approximately $0.0001), making it suitable for high- volume processing scenarios.

The LLM-based approach involves higher per-incident costs (approximately $0.001), representing an order-of-magnitude increase relative to the neural network. While higher on a per-incident basis, these costs remain small in absolute terms within the scope of the evaluation.

Table 2: Computational performance metrics

Metric	BERT + NN	LLM (Two-Stage)	Ratio (LLM/NN)
Mean inference time	44ms	1200ms	27x
Cost per incident	$0.0001	$0.001	10x
Throughput (incidents/hour)	~82,000	~3,000	0.04x
Token usage per incident	N/A	~800 tokens (average)	-

Table 3: Representative LLM classifications with reasoning

Incident	Ground	LLM	Confidence	Reasoning Excerpt	Outcome
1	INCLUDE	INCLUDE	High	"The incident involves a switching instruction that is not executed correctly..."	INCLUDE
2	EXCLUDE	EXCLUDE	High	"The incident involves a vehicle colliding with a fence while parking..."	EXCLUDE
3	INCLUDE	EXCLUDE	Medium	"The incident describes a software / communication issue..."	INCLUDE
4	EXCLUDE	INCLUDE	Medium	"The incident describes a physical failure (bottom bolt falling out) that occurred..."	EXCLUDE

3.3 Qualitative Analysis

A key advantage of the LLM approach is the provision of natural language reasoning for each classification. Table 3 presents representative examples demonstrating how LLM explanations support operational decision- making.

For correctly classified incidents (Examples 1–2), the LLM's reasoning demonstrates an understanding of domain-specific concepts, including the identification of specific protocol violations (e.g., the eight-step rule) and the ability to distinguish electrical switching incidents from other uses of the term "switching."

In incidents that were misclassified (Examples 3–4), the associated confidence scores often appropriately reflect model uncertainty, with both errors occurring at medium or low confidence levels. This behaviour supports operational workflows in which low-confidence predictions are routed for manual review, enabling a hybrid approach that combines the efficiency of LLM-based classification with human expertise in ambiguous cases.

For incidents classified as INCLUDE in Stage 1, the LLM's Stage 2 outputs provide structured attributes supporting root cause analysis. Table 4 shows example Stage 2 classifications.

These attributes enable operational analytics, including identifying whether incidents cluster during the planning, preparation, or execution phases; tracking compliance with the eight-step rule; and correlating incident frequency with task complexity. Such insights can inform targeted training interventions and process improvements.

Analysis of the 49 incidents that were correctly classified by the LLM but missed by the neural network reveals several consistent patterns. The LLM successfully identified incidents involving:

- **Implied switching terminology:** Descriptions referencing actions such as "isolated feeder" or "CB operation" without explicitly using the term "switching."
- **Complex narratives:** Multi-sentence incident descriptions requiring contextual understanding across the full text.
- **Ambiguous terminology:** Cases requiring distinction between electrical switching and other uses of the term (e.g., "switching to a backup generator").

Conversely, the neural network correctly classified 16 incidents that were misclassified by the LLM. These cases were predominantly ones where the neural network classified as EXCLUDE and the LLM flagged them as false positives, demonstrating the bias toward classifying as EXCLUDE by the neural network.

Both approaches failed to correctly classify six incidents characterised by minimal descriptive detail, excessive technical jargon without sufficient context, or genuinely ambiguous descriptions requiring domain expertise beyond what either model could reliably infer.

4 DISCUSSION

4.1 Key Trade-offs

While the LLM demonstrated superior performance compared to the neural network in classifying safety incident data, several important trade-offs should be considered when selecting an appropriate classification approach. LLMs offer significant advantages in terms of interpretability and flexibility. They can articulate their reasoning by referencing specific incident details, which is particularly valuable when users must understand and trust

Table 4: Example two-stage classifications showing incident attributes

Incident ID	Point of Occurence	8-Step Rule	Complexity	Location	Planned Status
1	Planning	No	Medium	Office	Planned
2	Execution	Yes	High	Field	Unplanned

classification outcomes. LLMs are also well suited to complex classification tasks, leveraging contextual understanding to assess nuanced criteria, such as determining compliance with the eight-step rule by identifying subtle procedural deviations within narrative text.

Traditional neural networks present a distinct set of trade-offs that may be better aligned with certain operational requirements. They provide more reproducible and deterministic classifications, which can be important for auditing and legal defensibility. However, they offer limited explanatory capability and often require either multiple models for different classification dimensions or a single, complex model with many output nodes. This structure can result in sparse training data for individual class combinations, particularly in domains where serious incidents are rare and labelled examples are limited. Neural networks also lack the domain knowledge and contextual reasoning capabilities inherent to LLMs, although they avoid the non-deterministic behaviour that can cause LLMs to produce different classifications for identical inputs across runs.

Taken together, these characteristics highlight that no single approach is universally optimal. The relative suitability of each method depends on the intended use case, data availability, performance requirements, and governance considerations.

Overall, LLM-based approaches are generally preferable for incident classification due to their interpretability, adaptability, contextual reasoning capabilities, and efficient use of limited training data. In most safety-critical applications, these benefits outweigh higher computational costs and non-deterministic behaviour. Nevertheless, organisations processing extremely high volumes of incidents or requiring strict classification consistency may find neural network–based approaches more appropriate.

4.2 Guidance for Practitioners

When selecting between traditional ML and LLM approaches for classification, practitioners should consider the following factors:

- **Training data availability:** Traditional ML models typically require substantial numbers of labelled examples per class. In this evaluation, 139 positive examples proved insufficient for effective neural network training. In contrast, LLMs support zero-shot or few-shot learning, making them valuable when labelled data is limited or costly to obtain.
- **Recall versus precision requirements:** For safety-critical applications where missing incidents is unacceptable, high recall is essential. This evaluation demonstrates that LLMs can maintain high recall while simultaneously improving precision. By comparison, insufficient and imbalanced training data caused the neural network model to sacrifice recall in favour of precision.
- **Interpretability requirements:** In contexts where stakeholders require explanations for classification decisions, LLM-generated reasoning can provide significant value. While neural networks may offer feature importance metrics, they often lack case-specific explanations that are accessible to non-technical audiences.
- **Computational constraints:** In this use case, LLMs incur approximately 10× higher per-incident costs and 27× longer inference times. For extremely high-volume applications (e.g., millions of classifications), these differences may become material. However, for moderate volumes, the absolute cost difference may be modest relative to the operational value delivered.
- **Operational adaptability:** LLMs can be updated through prompt refinement without retraining, enabling rapid adaptation to procedural changes or emerging incident types. In contrast, neural networks typically require retraining with newly labelled data, resulting in longer deployment cycles.
- **Total cost of ownership:** Evaluation should extend beyond inference costs to include data labelling effort, model maintenance, and manual review workload. This analysis suggests that improved LLM precision can reduce downstream review burden, potentially offsetting higher computational costs.
- **Prompt governance and maintainability:** Because much of the task logic is encoded in prompts, organisations should manage prompts as governed artefacts (with version control, change approval, and traceability)

and periodically review them as procedures, definitions, or regulatory requirements evolve.

4.3 Limitations

Several limitations should be considered when interpreting these results. The evaluation dataset comprised 100 incidents from a single organisation, validated against incidents flagged by the rules-based baseline. This sample size enabled detailed expert review but limits generalisability to other utilities, incident types not captured by rules-based filtering, or different operational contexts. The neural network's poor performance reflects severe training data constraints (only 139 positive examples among 838 total). Organisations with larger labelled datasets may achieve better traditional ML results. The LLM evaluation used a specific model implementation; performance, cost, and latency characteristics may differ substantially with alternative LLM providers or architectures.

Incident reporting patterns and procedural definitions evolve over time, necessitating periodic model re- evaluation. Finally, the SpaCy-based anonymisation approach may affect both privacy protection and classification accuracy in ways not independently quantified in this study.

Assessing readiness for broader safety-critical deployment would require replication on larger, out-of-sample datasets spanning different utilities, reporting practices, and incident types.

5 CONCLUSION

This study provides empirical evidence to inform AI investment decisions in asset management by comparing rules-based, neural network, and large language model (LLM) approaches for electrical switching incident classification. Results show that, on a 100-incident expert-validated dataset drawn from rules-flagged candidates within a single organisation, LLMs achieved superior balanced performance (76.8% precision, 89.8% recall, F1- score of 0.828) compared with rules-based systems (59% precision) and neural network approaches (75% precision, 10% recall).

The neural network's recall failure, attributable to insufficient labelled training data, highlights a critical consideration for traditional machine learning adoption: adequate labelled data is essential for effective model performance. In contrast, the LLM's zero-shot learning capability enabled effective classification without task- specific model training.

Clear trade-offs emerge. LLMs provide improved interpretability, contextual reasoning, and training data efficiency, but incur approximately 10× higher per-incident costs and 27× longer inference times compared with neural networks. For moderate-volume applications in similar operational contexts, where recall and explainability are prioritised, LLM-based approaches appear to offer compelling value despite these computational premiums; however, broader claims about safety-critical deployment should be validated through cross-context studies and larger evaluation sets.

These findings offer guidance for practitioners facing similar AI investment decisions. Rather than focusing solely on per-incident computational costs, organisations should assess training data availability, recall requirements, interpretability needs, and total cost of ownership. In scenarios where labelled data is limited and missed classifications carry operational or safety consequences, LLM-based approaches warrant serious consideration.

6 REFERENCES

[1] Brown, T. B. et al., 2020. Language Models are Few-Shot Learners. *Advances in neural information processing systems,* Volume 33, pp. 1877-1901.

[2] Chawla, N. V., Bowyer, K. W., Hall, L. O. & Kegelmeyer, W. P., 2002. SMOTE: synthetic minority over- sampling technique. *Journal of artificial intelligence research,* Volume 16, pp. 321-357.

[3] Devlin, J., Chang, M.-W., Lee, K. & Toutanova, K., 2019. BERT: Pre-training of deep bidirectional transformers for language understanding, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies,* 15 12, Volume 1, pp. 4171-4186.

TECH 2



ARTICLE 2 – Risk and Value-Focused Asset Management: Probabilistic Cost-Benefit Analysis for Railway Slopes

Kevin Hang, Sydney Trains

SUMMARY

Wet-weather hazards are widely recognised in slope asset management, yet rarely embedded in lifecycle investment decisions. Available data remains underutilised, and proactive interventions struggle for legitimacy because of uncertainty and the invisibility of avoided failures.

This paper proposes a probabilistic data-to-decision framework for wet-weather railway slopes that elevates risk-based practice to risk + value by linking event-conditioned climate risk to investment choices through interpretable, auditable logic. Machine-learning models trained on historical failures, asset attributes and local rainfall-exceedance signals estimate future slope risk profiles and annual failure probabilities at weather-station granularity. These inform a lifecycle valuation that recognises residual risk and monetises consequences comprehensively, including emergency and permanent repairs, resource opportunity costs and user costs from closures and temporary speed restrictions. Monte Carlo simulation generates distributions of net present value and benefit-cost ratios, together with percentile and tail-risk metrics, while a counterfactual baseline supports ex ante appraisal and ex post evaluation.

A real-world case study shows that rainfall-exceedance exposure is the dominant driver of portfolio value and that ridership patterns shape where proactive intervention delivers the greatest risk reduction per dollar, moving portfolios from reactive fixes toward strategic packages that reduce expected and tail losses.

Keywords: probabilistic cost-benefit analysis; Monte Carlo simulation; machine learning; rainfall-exceedance exposure; lifecycle investment; railway slope asset management

1 INTRODUCTION

Extreme wet-weather hazards pose growing challenges for transport asset management, particularly for slopes where failures lead to major service disruptions and safety risks. In March 2022, multiple landslips between Katoomba and Emu Plains severely damaged the Blue Mountains Line, causing staged closures of up to 18 days for freight and 31 days for full passenger operations. At a system level, Network Rail estimates weather-related delays and cancellations cost about £50-£100 million per year, with heavy rainfall and flooding as major contributors (Network Rail, 2021). These cases exemplify a wider pattern in which localised slope failures propagate into network-scale reliability losses, with significant economic and social consequences, and are expected to intensify as extreme rainfall becomes more frequent in a warming climate (IPCC, 2021).

Despite these impacts, slope management in many rail systems remains organised largely around surveillance, wet-weather patrols, Temporary Speed Restrictions (TSRs) and emergency interventions, with structural risk-reduction measures often deferred under constrained budgets. Disaster-risk and public-finance studies show that expenditure is commonly skewed toward post-event relief and reconstruction, with comparatively limited investment in ex ante mitigation (Productivity Commission, 2014; Vorhies, 2012). This systematic underinvestment, sometimes described as a recovery bias (Meyer & Kunreuther, 2017), creates a legitimacy gap for proactive resilience projects whose benefits are probabilistic, diffuse and realised over long horizons. At its core, this gap reflects both the highly variable operational consequences of slope failures and persistent fragmentation across geotechnical, operational and economic data and models.

At a technical level, hazard-exposure-vulnerability (HEV) framing, as recognised in international disaster risk governance (Cardona et al., 2012; UNISDR, 2017), has typically been operationalised through classical HEV scoring, deterministic stability analyses and GIS-based susceptibility mapping, which rarely capture uncertainty, time-varying processes such as rainfall exceedance, or the propagation of slope failures into network-wide operational and user impacts (Huang, 2025; Qian et al., 2025; Pant et al., 2016; Pregnolato & Dawson, 2018; Zhou et al., 2021). Operational consequences are further complicated by non-linear disruption effects, including timetable knock-on delays, diversion capacity limits and prolonged recovery under constrained access and resources. Although catastrophe modelling and recent machine learning work offer probabilistic tools for landslide risk and fragility (Grossi & Kunreuther, 2005; Liu et al., 2023; Elansari et al., 2024), there is still limited practical guidance on how rail operators can connect these methods to value-based asset management and investment decisions (AECOM, 2025; World Economic Forum, 2024).

Research on wet-weather hazards, network resilience, machine learning and cost-benefit analysis is extensive but often siloed, so agencies still struggle to link climate exposure and asset failure to network disruption and investment value in a governance-ready way. These gaps raise a practical question for operators and funders: how can wet-weather slope failure be quantified, linked to operational and user losses, and assessed within a transparent probabilistic cost-benefit framework? Against this backdrop, this paper makes four contributions:

- *Event-conditioned failure modelling.* Machine learning on slope attributes, risk records and rainfall-exceedance exposure to estimate future risk profiles and the associated annual failure probabilities for each slope.
- *Operational and user impact modelling.* A link from slope failures to block-section passenger and freight operations, quantifying closures, TSRs and associated user costs.
- *Probabilistic cost-benefit framework.* A Monte Carlo-based probabilistic cost-benefit analysis that generates financial and reliability metrics, including tail-risk measures and marginal packaging of interventions.
- *Programme design and learning.* A comparison of reactive, risk-based and value-based intervention programmes against a counterfactual baseline, supporting ex-ante appraisal and ex-post evaluation.

Overall, it operationalises Risk + Value-Focused Asset Management

by linking rainfall-exceedance exposure (i.e., rainfall surpassing trigger thresholds) to slope risk, network disruption and probabilistic cost-benefit value.

2 OVERVIEW OF THE FRAMEWORK

Figure 1 illustrates a geospatial multilayer integration of hazard, vulnerability and exposure in rail slope risk management. Rainfall hazard is represented as a spatial grid of rainfall-exceedance cells, each linked to affected slopes and, in turn, to the block sections and corridors they influence. This connects weather (hazard) to slope instability (vulnerability) and to passenger and freight operations (exposure[1]), allowing closures, speed restrictions and delays to be translated into a coherent cost-value-benefit view within a single holistic framework.

The proposed framework (Figure 2) establishes a comprehensive pathway linking extreme wet weather hazards to slope failures, transport network disruptions and economic performance. It does not predetermine investment outcomes; instead, it clarifies the system structure and the relationships among hazard, vulnerability, exposure and value that shape different decisions. This allows decision makers to apply their own objectives, constraints and risk tolerance within a transparent quantitative model.

At the top of the framework, extreme wet-weather events are represented by location, intensity-frequency-duration (IFD) statistics, historical rainfall exceedance and projected trends. These variables are treated both as a long-term deterioration driver of slopes and as the immediate trigger for individual slope failures.

On the left-hand branch, the interventions module includes slope stabilisation works, drainage improvements and other geotechnical risk-mitigation measures. These interventions, together with the extreme wet-weather deterioration driver, are combined with slope attributes and historical intervention records to form the asset profile, using machine learning models. This profile captures physical and functional characteristics, condition information and likelihood and consequence ratings, and is used to estimate annual failure probabilities for different slope types and risk classes, providing a probabilistic description of vulnerability that is consistent with the HEV concept introduced in the literature review.

On the right-hand branch, extreme wet-weather events lead to asset failures that are specified by location, frequency and direct engineering consequences such as Temporary Speed Restrictions or line closures. These failures propagate into transport operations, which represent exposure. The framework focuses on operational responses that are material for cost and service outcomes, including replacement bus operation, passenger trip loss or mode shift, and freight rerouting, mode shift

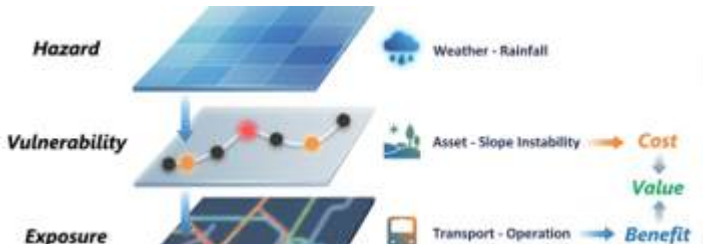


Figure 1: Geospatial multilayer integration of hazard, vulnerability and exposure in slope risk management

1 Here, 'exposure' is used in the transport systems sense, referring to the degree to which network operations and passenger or freight services are affected by slope failures, distinguishing it from climate risk exposure, which refers only to the hazard facing assets.

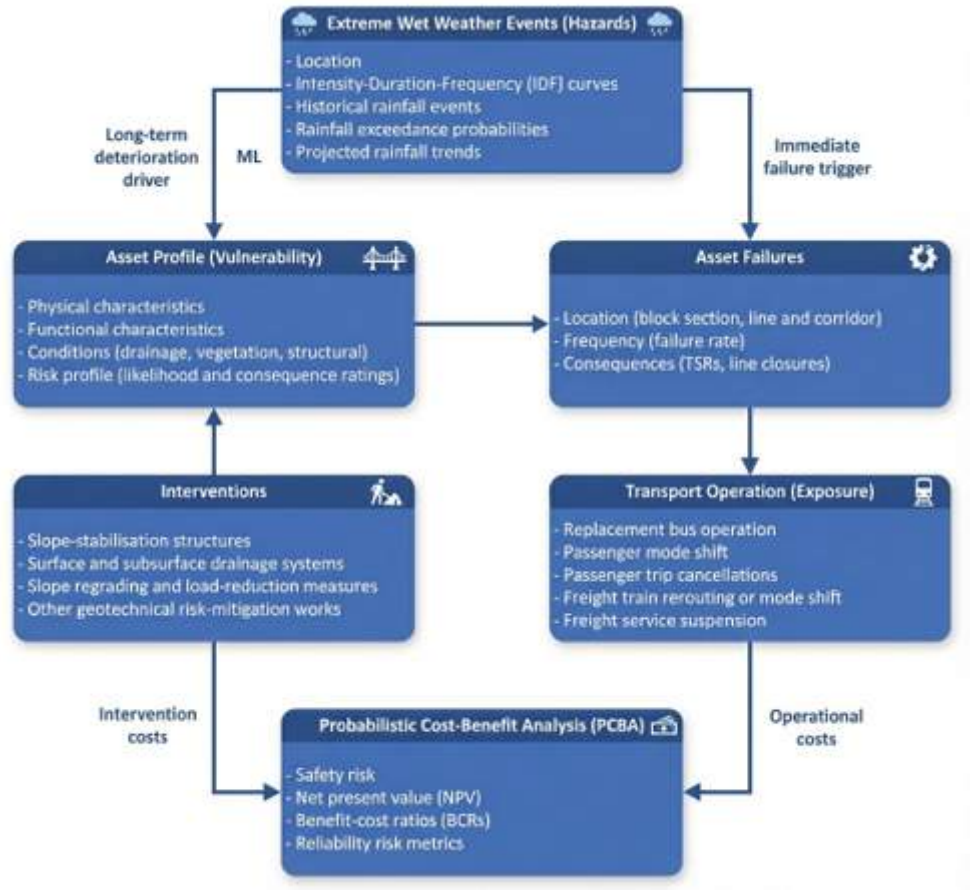


Figure 2: Framework for probabilistic cost-benefit analysis of railway slopes

or suspension. These responses generate operational costs in the form of additional operating expenditure, delay costs and revenue impacts. Expressing exposure at block-section level allows operational impacts to be linked back to specific segments of the network and, where data permit, to specific passenger and freight trips.

The lower part of Figure 2 introduces the probabilistic cost-benefit analysis (CBA) module, which builds on the event-conditioned failure probabilities and operational loss model. In its core formulation, the probabilistic CBA module uses a probabilistic mathematical model to combine annual failure probabilities, per-event loss components and intervention effects, and to derive expected values and, where tractable, distributions of net present value (NPV), benefit-cost ratios (BCRs) and reliability metrics such as closure days. Monte Carlo simulation is then used as an additional tool for sensitivity analysis and for scenarios where interactions, non-linearities or time-varying hazards make closed-form evaluation impractical. Overall, the framework offers a coherent and auditable path from rainfall statistics and slope attributes to operational losses and economic performance measures, and is designed to be transferable to other wet-season-sensitive civil assets.

3 DATA AND METHODS

3.1 Extreme Wet-Weather Events

Wet-weather hazard is represented using Intensity-Frequency-Duration (IFD) design rainfall and long-term observations from Bureau of Meteorology (BoM) stations across the network (BoM, 2016). For each station, IFD curves provide design depths for selected durations and exceedance probabilities

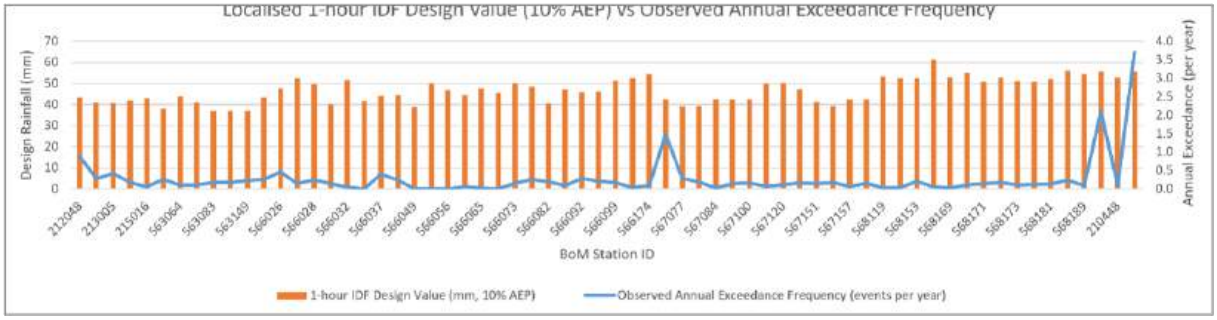


Figure 3: Comparison of design 1-hour 10% AEP rainfall with observed exceedance frequencies

(e.g., 3-day 5% AEP) consistent with Australian Rainfall and Runoff practice (Ball et al., 2019). Historical records are screened for exceedances of these thresholds and averaged to derive local annual exceedance counts that characterise extreme wet-weather exposure (Wasko & Sharma, 2017; Wasko, 2021). These rainfall-exceedance metrics are parsimonious, scalable proxies for triggering conditions rather than full hydrological models. Figure 3 illustrates the 1-hour 10% AEP design rainfall and corresponding average annual exceedance counts.

Each slope is mapped to its nearest BoM station and inherits that station's exceedance metrics, a common approach in infrastructure-scale hazard attribution (Upadhyay, 2024). Material-specific thresholds are used so that the hazard definition reflects dominant failure mechanisms, with shorter, more intense bursts for rock cuttings and longer accumulation for soil cuttings and embankments (Corominas et al., 2014). Future trends are represented by applying scenario factors from regional climate projections to station-level exceedance metrics, modifying the frequency or magnitude of threshold exceedance over time (CSIRO & BoM, 2015; IPCC, 2021). These adjusted exposure pathways are then propagated through the failure and cost-benefit models to explore the implications of more frequent extreme rainfall and non-stationary design conditions.

3.2 Asset Profile and Machine Learning

The asset-profile model provides a consistent, data-driven vulnerability characterisation for each slope, expressed as likelihood and consequence ratings and the derived risk class.

Table 1: TfNSW geotechnical risk assessment matrix

			Likelihood					
	Likelihood of geotechnical event affecting track		Incredible Not expected to occur at any time	Rare, improbable Not expected to occur or only under exceptional conditions	Unlikely May only occur under very adverse weather conditions at some time	Possible It is possible event will occur within 12 months in adverse weather conditions	Likely Will probably occur within 12 months in most circumstances	Almost certain Expected to occur within 12 months in most circumstances
Consequence			I1	I2	I3	I4	I5	I6
> 10 fatalities	Disastrous	C6	B-, P1	B+	A	A	A	A
2 - 10 fatalities	Catastrophic	C5	C+, P1	B-, P1	B+	A	A	A
1 fatality (2 - 10 serious injuries)	Critical	C4	C-	C+, P1	B-, P1	B+	A	A
1 major injury	Major	C3	B	C-	C+, P1	B-, P1	B+	A
1 or more minor injuries	Minor	C2	B	B	C-	C+, P2	B-, P2	B+
Illness, first aid treatment or injury not requiring treatment	Negligible	C1	B	B	B	C-	C+, P2	B-, P2

Targets are taken from historical risk records: for each slope and inspection year, recorded likelihood and consequence ratings, the composite risk class and any observed failures. Separate ordinal models are trained for likelihood and consequence, and predictions are mapped back through the existing risk matrix to remain governance-compatible. Machine learning is used to capture non-linear interactions across correlated predictors and to produce calibrated probabilities, providing a consistent projection of how existing likelihood and consequence bands may evolve under changing wet-weather exposure rather than replacing expert judgement.

Predictors are grouped into four broad categories: geometry and setting, ground and drainage condition, wet-weather exposure, and intervention history. Geometry and setting cover material type (rock cutting, soil cutting, soil embankment), slope height and length, batter angle, curvature and proximity to the track. Ground and drainage condition summarise inspection observations such as seepage, erosion, cracking, slaking, vegetation cover and the performance of surface and sub-surface drains. Wet-weather exposure is represented by the rainfall-exceedance metrics from Section 3.1, linked to each slope via its nearest BoM station and material-specific thresholds. Intervention history records stabilisation and drainage works, their type, extent and age, so the model can learn differences between treated and untreated slopes and infer any reduction in vulnerability after works.

Tree-based ensemble classifiers are fitted by material group, using a stratified train–test split by risk class and cross-validation within the training data. Model performance is summarised in Table 2, which reports accuracy, precision, recall, and macro-F1, with predicted probabilities calibrated to ensure the distribution of modelled ratings aligns with observed ratings and failure frequencies. The results indicate that likelihood models achieve strong discrimination (F1 around 0.76-0.82), whereas consequence models perform more modestly (F1 around 0.57-0.65). Because consequence ratings are inherently harder to predict and the models show only moderate discrimination, these outputs are used to refine and cross-check existing banded consequence ratings rather than replace them, ensuring conservative and governance-aligned use in subsequent economic analysis. The economic analysis remains anchored on the observed bands and their associated cost ranges.

The calibrated models are then used to project likelihood, consequence and composite risk ratings for future analysis years, conditional on assumed rainfall-exceedance exposure and intervention pathways. Shapley Additive Explanations (SHAP)

Table 2: ML model performance by asset type and target

Modelling Category	Accuracy	Precision	Recall	F1-score
Rock Cutting Likelihood	0.83 (248)	0.82	0.83	0.82
Rock Cutting Consequence	0.59 (248)	0.6	0.59	0.59
Soil Cutting Likelihood	0.81 (81)	0.77	0.81	0.79
Soil Cutting Consequence	0.58 (81)	0.61	0.58	0.57
Soil Embankment Likelihood	0.80 (232)	0.75	0.8	0.76
Soil Embankment Consequence	0.65 (232)	0.64	0.65	0.65

are applied across the trained models to generate prioritised factor rankings and diagnostic dependence plots. Figure 4 presents a representative example for rock cuttings, illustrating how attributes such as height, drainage condition, rainfall exceedance and intervention history most strongly drive higher likelihood and consequence ratings and supporting engineering review of the model behaviour. These consistently high-influence attributes are subsequently adopted as key explanatory variables in the asset-failure models developed for the cost-benefit analysis.

3.3 Asset Failure Modelling

Annual failure rates (AFR) are estimated at slope level using material type, likelihood rating, rainfall-exceedance exposure and historical failure data. Each slope is mapped to its nearest Bureau of Meteorology (BoM) weather station and assigned annual exceedance counts for material-specific IFD thresholds (for example, different thresholds for rock and soil cuttings).

For each material-likelihood combination, baseline AFRs are calibrated from observed failures and then scaled by a rainfall-exceedance multiplier so that higher local wet-weather exposure increases the estimated failure rate:

$$\mathrm{AFR}_i = \lambda_{m\ell}\,\kappa_i \qquad \textit{Equation 1}$$

where is the baseline AFR for material *m* and likelihood band ℓ. The rainfall multiplier *Ki* is defined using a small set of exceedance indicators (typically two to three per material type):

$$\kappa_i = \mathrm{clip}\left(\sum_{k=1}^{K_m} w_{m,k}\,\frac{\mathrm{RE}_{i,k}}{\overline{\mathrm{RE}}_{m,k}},\ \kappa_{\min},\kappa_{\max}\right) \qquad \textit{Equation 2}$$

where $\mathrm{RE}_{i,k}$ is the annual exceedence count for indicator *k* at slope i, $\overline{\mathrm{RE}}_{m,k}$ is the mean exceedance for material *m*, and the non-negative weights $w_{m,k}$ sum to 1. The indicators and weights are informed by prior interpretable machine-learning analysis (SHAP), with higher weights assigned to rainfall thresholds that show greater and more stable influence on failure likelihood. Bounding *Ki* ensures numerical stability and physical realism, avoiding zero failure rates where exceedances are rare and limiting unrealistic amplification under extreme local exceedance ratios. This weighted multi-indicator formulation preserves the transparency and governance readiness of a simple multiplicative AFR model aligned with existing risk ratings, while better reflecting the empirical finding that wet-weather slope failures are often driven by more than one rainfall characteristic.

3.4 Interventions

Interventions are permanent geotechnical measures designed to reduce wet-weather slope risk. Typical solutions include retaining and stabilisation structures with associated drainage upgrades, consistent with current practice for

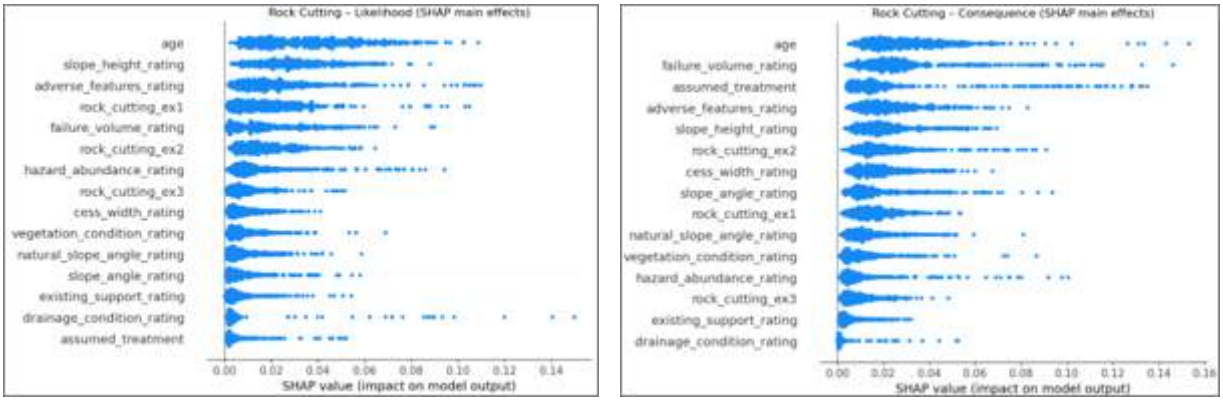


Figure 4: Drivers of rock-cutting risk (SHAP: Likelihood & Consequence)

transport earthworks management (Corominas et al., 2014; Upadhyay, 2024). The option selected for each slope is governed by height and geometry, footprint constraints, ground conditions and the dominant failure mechanism. Intervention extent is specified at slope level (for example, treated length and height), and the capital cost of planned work, denoted I_0, is taken from engineer's estimates prepared for each intervention type and extent. These estimates combine quantities for structural works, earthworks, drainage, access and traffic management with assumed unit rates and allowances for overheads, possessions and contingency.

For each slope, the cost if a failure occurs in year t is

$$C_t^{\text{event}} = ER_t + PR_t + OC_t + OP_t(a) \qquad \textit{Equation 3}$$

Define yearly components as: ER_t emergency repair, PR_t permanent remediation on the reactive path, OC_t opportunity cost from capacity constraints, subcontracting premiums and deferred-work inefficiencies, and $OP_t(a)$ operational cost for event afrom line closure or TSR on the affected block section. In the cost-benefit model, asset-level costs are all realised only when a failure occurs. ER_t covers temporary works to restore service, for example debris removal, temporary support and drainage. PR_t is typically higher than planned capital I_0 for an equivalent solution due to reactive procurement, access constraints and scope escalation. OC_t reflects diversion from planned work, including idled labour and plant and deferred maintenance. Residual risk is represented via the post-intervention AFR, enabling comparison of risk reduction per dollar across options.

3.5 User Costs: Passenger and Freight Impacts

In this paper, a block section b is defined as the operational segment over which a closure or TSR is applied, with boundaries reflecting station turn-around capacity and the ability to stage replacement buses and staff. It is the spatial unit used to model passenger and freight disruption.

User costs are captured through an operational impact term at block-section level, denoted $OP_{t,b}$. The expected operational impact of a slope-failure event a on block section b in year t is

$$OP_{t,b}(a) = q_a\left[(L_{t,b} + F_{t,b})D_a^{close} + T_{t,b}D_a^{TSRclose}\right] + (1 - q_a)\left[T_{t,b}D_a^{TSRopen}\right] \qquad \textit{Equation 4}$$

where q_a is the probability that event a triggers a block-section closure; D_a^{close} is closure duration; $T_{t,b}$ is TSR cost per day; and $D_a^{TSRclose}$ and $D_a^{TSRopen}$ are TSR days with and without closure. $L_{t,b}$ and $F_{t,b}$ use passenger and freight parameter sets from current Transport for NSW values (Transport for NSW, 2025). A block-section inventory links each section to its passenger area and freight corridor and records key attributes (line role, demand, tonnage, etc.). Under residual risk, expected user cost in year t combines $OP_{t,b}(a)$ with the event probability distribution and is carried into the probabilistic CBA. The operational response model is stylised to capture first-order disruption rather than detailed timetable dynamics. Costs are represented as generalised economic costs consistent with standard appraisal practice, intended for decision support rather than a full welfare decomposition.

For each block section b, passenger cost per closure day is

$$L_{t,b} = P_{\text{bus},b}c_{\text{bus}}^{\text{area}(b)} + P_{\text{car},b}c_{\text{car}}^{\text{area}(b)} + P_{\text{cancel},b}c_{\text{cancel}}^{\text{area}(b)} \qquad \textit{Equation 5}$$

Where b denotes the block section and "area" (b) is its associated area; $P_{\text{bus},b}$, $P_{\text{car},b}$ and $P_{\text{cancel},b}$ are the numbers of passengers in block section b who respectively transfer to bus, switch to car, or cancel their trip during a line closure; $c_{\text{bus}}^{\text{area}(b)}$ is the unit bus transport cost in that area, determined by block-section length, the bus-rail detour factor, the mix of city-bus and operator-provided replacement services, bus/train speeds, waiting time, value of travel time and bus operating costs; $c_{\text{car}}^{\text{area}(b)}$ is the unit car cost, determined by fuel and non-fuel operating costs, parking cost, the car-rail detour factor and car occupancy; and $c_{\text{cancel}}^{\text{area}(b)}$ is the economic value of a foregone public transport trip, reflecting lost benefits such as productivity and access to services.

Freight cost per closure day is

$$F_{t,b} = F_{\text{rail},b}c_{\text{rail}}^{\text{corr}(b)} + F_{\text{road},b}c_{\text{road}}^{\text{corr}(b)} + F_{\text{susp},b}c_{\text{susp}}^{\text{corr}(b)} \qquad \textit{Equation 6}$$

Where b denotes the block section, and corr (b) is the freight corridor associated with that block section; $F_{\text{rail},b}$, $F_{\text{road},b}$ and $F_{\text{susp},b}$ are the tonnes of freight in block section that are respectively rerouted by rail, transferred to road, or suspended during a closure; $c_{\text{rail}}^{\text{corr}(b)}$ is the cost per tonne of rerouting freight by rail within that corridor; $c_{\text{road}}^{\text{corr}(b)}$ is the cost per tonne of transferring freight to road; and $c_{\text{susp}}^{\text{corr}(b)}$ is the economic loss per tonne associated with freight suspension in that corridor.

The area-specific unit costs $c_{\text{bus}}^{\text{area}(b)}$, $c_{\text{car}}^{\text{area}(b)}$, $c_{\text{cancel}}^{\text{area}(b)}$

reflect inner-city, suburban and regional conditions (replacement-bus share, speeds, parking costs and the value of foregone trips), while the corridor-specific unit costs $c_{\text{rail}}^{\text{corr}(b)}$, $c_{\text{road}}^{\text{corr}(b)}$, $c_{\text{susp}}^{\text{corr}(b)}$ reflect freight detour options, road freight capacity and port access. Passenger responses are therefore treated as area-based and freight responses as corridor-based, within a generalised cost framework consistent with standard transport appraisal practice and intended for decision support rather than full welfare decomposition (Boardman et al., 2018; Pant et al., 2016; Transport for NSW, 2025). TSRs are handled analogously, with $OP_{t,b}$ driven primarily by delays arising from speed restrictions. Passenger cancellation and freight suspension costs are treated as mutually exclusive outcomes within each disruption state, avoiding double counting with delay or diversion costs.

3.6 Probabilistic Cost-Benefit Analysis

Expected costs over the analysis horizon T are computed separately for the reactive baseline and the proactive (treated) case. Probabilistic CBA is used to reflect uncertainty in failure timing, residual risk after repair, and intervention effectiveness, so investment decisions are informed by a distribution of outcomes rather than a single-point estimate. In the baseline (reactive) case, present value is

$$PV_{\text{baseline}} = \sum_{t=1}^{T} S_{t-1} p_t \left[\frac{C_t^{\text{event}}}{(1+r)^t} + \sum_{u=t+1}^{T} \frac{p_u^R C_u^{\text{event}}}{(1+r)^u}\right] \qquad \textit{Equation 7}$$

where r is the discount rate, p_t, is the probability of first failure in year t, S_{t-1} is the survival probability up to the start of year t, and $p_u^R = \beta p_u$ is the residual post-repair failure probability $(\beta \in [0,1])$ reflecting that repairs reduce but do not eliminate risk.

In the proactive case, a one-off planned intervention cost I_0 is incurred at $t = 0$ and residual failure probabilities are reduced by a factor $\phi \in [0,1]$. The present value of residual costs is

$$PV_{\text{proactive}} = \sum_{t=1}^{T} \frac{\phi p_t C_t^{\text{event,pro}}}{(1+r)^t} \qquad \textit{Equation 8}$$

where $C_t^{\text{event,pro}}$ is the residual event cost under treatment (typically with reduced ER_t, PR_t, OC_t and/or OP_t).

Benefits are defined as

$$PVB = PV_{\text{baseline}} - PV_{\text{proactive}} \qquad \textit{Equation 9}$$

with net present value

$$NPV = PVB - I_0, BCR = \frac{PVB}{I_0} \qquad \textit{Equation 10}$$

This analytical formulation is applied at slope and later block-section level to obtain expected NPV and BCR for each intervention package, which form the basis for deterministic ranking before stochastic analysis. A 7% discount rate and 50-year horizon are used in this study for consistency with current Australian public-sector appraisal practice, noting that alternative assumptions will change absolute values and can affect rankings at the margin, but the framework is designed to make the underlying risk-value trade-offs transparent.

3.7 Monte Carlo Analysis and Tail Metrics

Commuter rail systems are sensitive to tail risks because rare, severe disruption events (e.g., multi-day wet-weather closures) can drive disproportionate operational, economic and reputational impacts. Conventional rail appraisals typically emphasise expected values and can obscure this downside structure.

To characterise uncertainty, the framework is embedded in a Monte Carlo (MC) scheme. The purpose of the Monte Carlo analysis is not to predict realised outcomes, but to expose the sensitivity of investment decisions to uncertainty and downside conditions that are typically hidden in expected-value appraisal. In each run, annual failures are sampled for every slope from the AFR model (Section 3.3), conditioned on rainfall-exceedance exposure and material type, and block-event logic (Section 3.5) ensures operational impact is applied at most once per event per block section. Discounted baseline and proactive cost streams then yield distributions of costs, benefits, NPV and BCR, summarised using percentiles and tail metrics such as Conditional Value-at-Risk (CVaR), the mean of the worst per cent of outcomes (Janeková & Onofrejová, 2015; Baral & Shahandashti, 2022). This exposes downside performance (e.g. P10 and CVaR) and supports marginal packaging of interventions, using diminishing NPV and BCR as a transparent stopping rule. The probabilistic outputs are used as diagnostics of downside exposure and robustness, rather than as optimal decision rules.

Sensitivity tests indicate that while absolute NPV and BCR vary with discount rate and unit-cost assumptions, relative rankings and the identification of high-impact block sections are broadly preserved across plausible ranges, supporting robust prioritisation under uncertainty.

3.8 Counterfactual Baseline and Ex-post Evaluation

The probabilistic CBA is anchored to a do-minimum reactive baseline with no planned works at t=0, where slopes are managed through emergency repair, follow-up remediation and operational responses. Baseline failure probabilities use the untreated AFR model (Section 3.3), with residual post-repair risk scaled by . All other inputs (traffic, rainfall exposure, unit costs and block-section logic) are held constant so differences in present value reflect the intervention package alone (Boardman et al., 2018; NSW Treasury, 2018).

For ex-post evaluation, the same baseline model is re-run using realised data over the evaluation period (for example, actual rainfall exceedance, observed demand and updated economic parameters). The resulting distribution of counterfactual failures, closure days and costs is compared with observed post-intervention outcomes, allowing realised PVB, NPV and reliability improvements to be estimated and key parameters (such as AFRs or repair multipliers) to be updated. This provides a consistent backbone for both forward-looking appraisal and learning about the effectiveness of wet-weather resilience investments.

4 CASE STUDY AND RESULTS

4.1 Overview of Case Study

This framework is demonstrated on the Sydney Trains network using a dataset of 2,000+ cuttings and embankments across ~2,000 km of track.

Slopes are linked to 250+ stations, aggregated into 70 operational block sections, and assigned to 64 BoM weather stations with IFD design rainfall and up to 77 years of minute-level observations. These linkages combine slope attributes, risk ratings and rainfall-exceedance exposure at slope and block-section level. The pipeline feeds observed data into the machine-learning asset profile, then into the AFR model, and finally into probabilistic CBA and Monte Carlo analysis to compare proactive and reactive investment against a no-intervention counterfactual.

The asset-profile model randomly splits the historical dataset into 80% training and 20% test data, stratified by risk class, and generates calibrated predictions for all slopes. The cost-benefit analysis then evaluates two baseline vulnerability profiles over the 50-year horizon (2026-2075): (i) a fixed 2025 ratings baseline that holds recorded likelihood, consequence and risk ratings constant, and (ii) a predicted-ratings baseline that applies the machine-learning projections over time. Both baselines are passed through the AFR and cost-benefit modules to compare reactive and proactive wet-weather slope interventions under alternative representations of the risk profile.

4.2 Key Modelling Outputs

Figure 5 shows, for slopes ordered by 2025 risk rating, how discounted operational cost at 7% splits between passengers and freight. The stacked bars indicate strong variation along the network: freight dominates in some sections, especially freight-intensive corridors like the Central Coast &

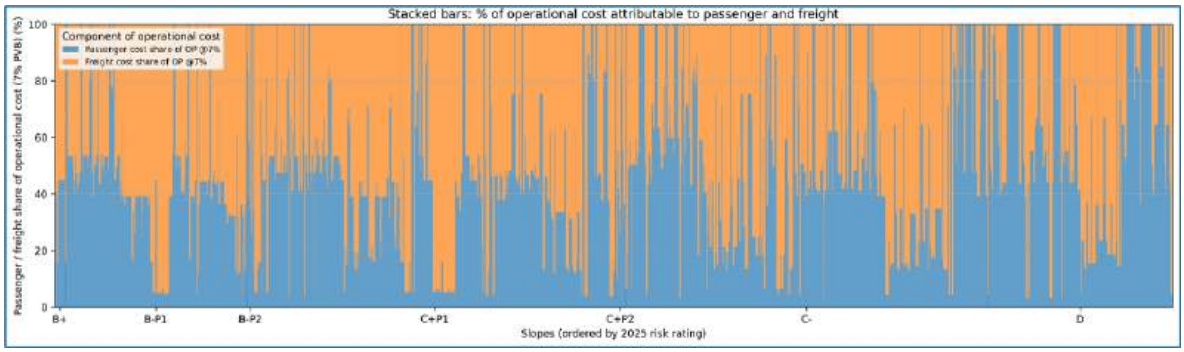


Figure 5: Passenger and freight shares of operational cost, per slope, by 2025 risk rating

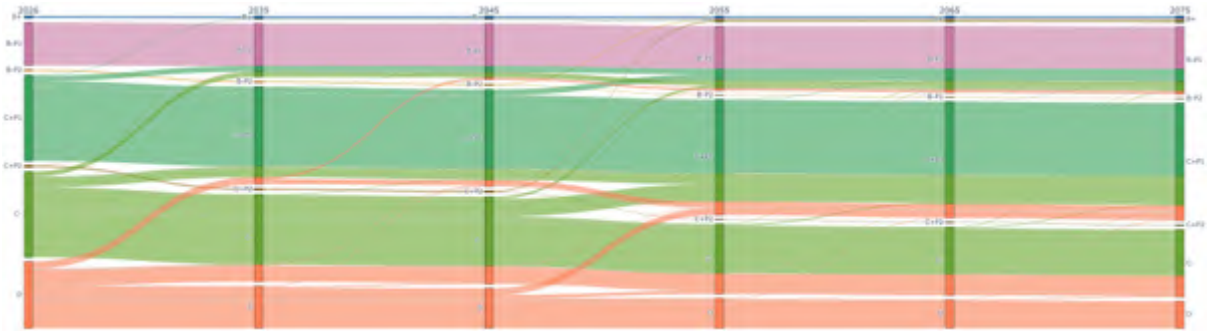


Figure 6: Predicted slope risk-rating transitions, 2026-2075

Newcastle (CCN) Line, while inner-suburban sections are mainly passenger-driven. This matches past CCN wet-weather disruptions, where rerouting constraints, path loss and suspensions hit freight hardest, providing an external plausibility check. It also shows why impacts should be assessed at block-section level, rather than assuming passenger effects always dominate.

Figure 6 shows the predicted evolution of slope risk ratings from 2026 to 2075 using the machine-learning model under a no-intervention counterfactual. Each coloured band represents a risk class (from B+ to D), and the connecting ribbons show how slopes move between these classes over time. Most assets remain in their original class, but a noticeable minority shift into higher-risk bands as wet-weather exposure and condition evolve. For comparison, the alternative '2025 rating baseline' assumes that risk ratings remain fixed at their 2025 values, so class memberships would not change through time and are therefore not shown as a separate figure.

Figure 7 plots per-slope NPV and BCR by 2025 risk rating for discount rates of 3%, 7% and 10%. The portfolio is dominated by medium and low risk slopes: roughly one sixth of assets are in the high-risk bands (B+, B-P1, B-P2), about one third in C+P1/P2, another

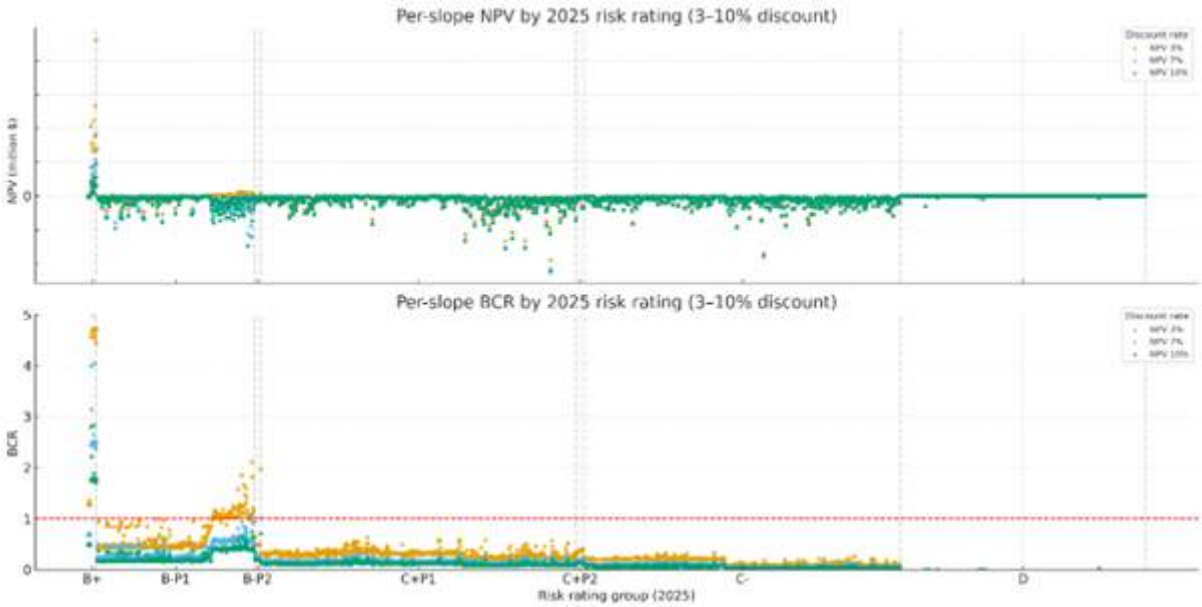


Figure 7: Per-slope NPV and BCR, 3-10% discount rates (2026-2075, 2025 risk ratings)

third in C-, and just under a quarter in D. Consistent with this, both panels show a pronounced long-tail pattern. A small group of slopes, mainly B+ and B-P1, deliver clearly positive NPVs and BCRs above 1, especially at 3%, but the majority cluster around zero NPV with BCRs below 1. In practice this means that most slopes are not economically justified for proactive treatment, and even some B+ slopes fail the economic test.

Two key insights emerge. First, only a very small subset of slopes delivers positive NPV and BCR greater than 1, confirming that proactive programmes must be highly selective rather than broad-brush. Second, several slopes in the highest risk band (B+) exhibit weak or negative economic returns, demonstrating the limits of relying solely on risk ratings and reinforcing the need for an explicit value-based screening lens.

Figure 8 repeats the analysis using predicted future risk profiles from the machine-learning model. The same long-tail structure appears, but some of the strongest NPV and BCR results now arise from slopes currently rated in C bands, reflecting deterioration under more frequent extreme wet-weather exposure. This indicates that a forward-looking, event-conditioned probabilistic CBA can both deselect high-risk slopes with weak economics and surface lower-rated assets that become attractive candidates for proactive intervention. The no-intervention projections imply that ageing assets and wetter conditions will push more slopes into higher risk classes over time, so a portfolio-level view of future funding needs, grounded in both risk and cost-effectiveness, is essential for staging interventions and managing long-term slope risk.

4.3 Sensitivity Analysis

A one-way sensitivity analysis was undertaken using the same 50-year cost-benefit framework and predicted risk ratings that underpin the Monte Carlo simulation. Programme comparisons are intended to illustrate how different prioritisation logics shape portfolio composition under uncertainty, rather than to assert the universal superiority of any single strategy. Each parameter was perturbed independently by ±20% while other variables were held constant, and the resulting change in portfolio present value of benefits at a 7% discount rate was computed. Parameters are: RE (rainfall-exceedance multiplier at year level), PC (line-closure probability), DC (line-closure duration in days), TSR (TSR days), PTV (passenger trip volume) and FRT (freight volume).

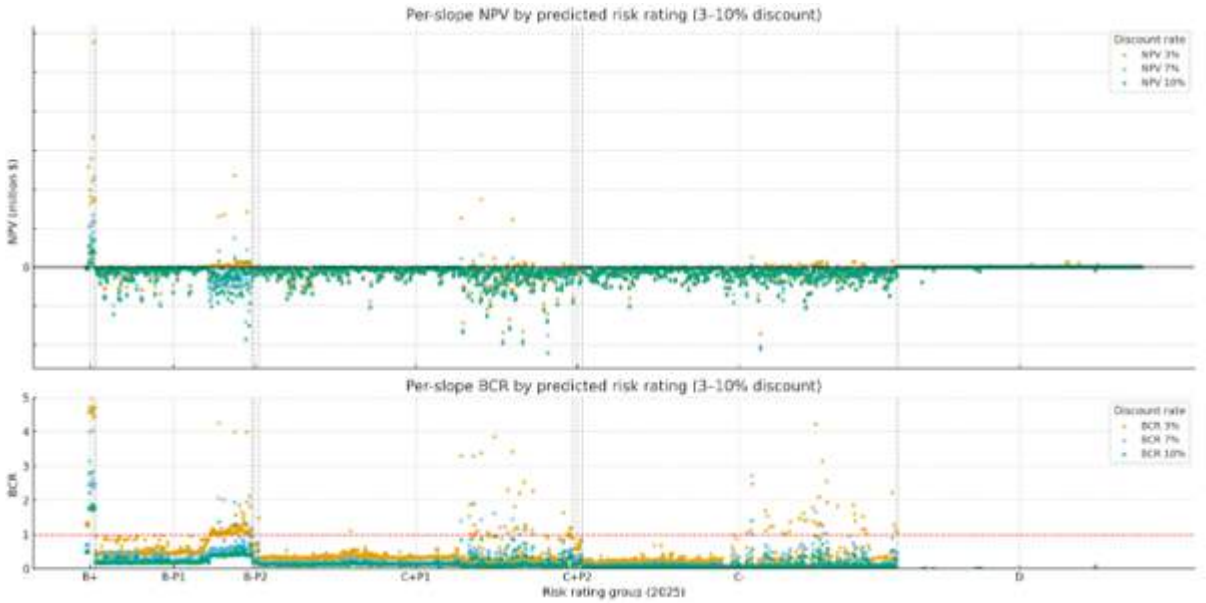


Figure 8: Per-slope NPV and BCR, 3-10% discount rates (2026-2075, predicted risk ratings)

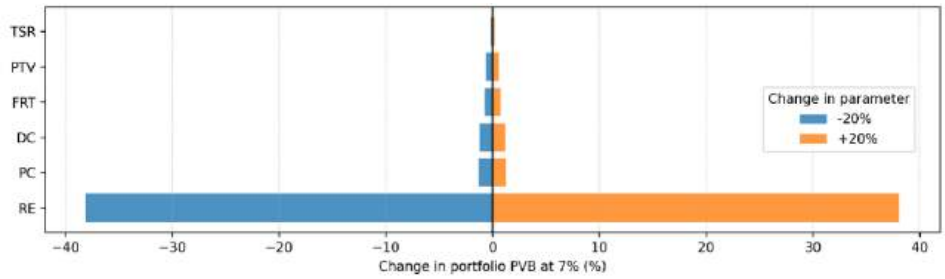


Figure 9: Percentage Change in Portfolio PVB under One-way Parameter Variations

Figure 9 shows that the RE multiplier has by far the greatest influence on portfolio PVB, while the other inputs have comparatively minor effects. This result reflects the fact that RE operates at the hazard-failure interface rather than as a linear cost or demand scaler. Here, RE does not represent total rainfall but the frequency of rainfall patterns that exceed thresholds associated with slope failure, driven by short-duration intensity bursts, storm clustering and multi-day accumulation. A 20% increase in RE therefore increases the frequency with which slopes cross failure thresholds and trigger block-section disruptions, leading to non-linear amplification of portfolio impacts through event occurrence, closure duration and tail-cost dominance. The resulting sensitivity reflects threshold-driven changes in failure behaviour rather than proportional scaling of operational inputs.

Parameter uncertainty is represented using bounded triangular distributions defined by minimum, most-likely and maximum values, consistent with contemporary probabilistic practice in transport infrastructure appraisal where expert judgement is available but historical evidence is incomplete (Transport and Infrastructure Council 2020). A Monte Carlo simulation with 10,000 iterations then propagates these distributions through the 50-year model to generate probability distributions for PVB, NPV and BCR, in line with emerging guidance on probabilistic assessment of asset resilience and disruption risks (Dong et al. 2022). Dependence between uncertain parameters is captured using a Gaussian copula, which permits flexible marginal specifications while reproducing realistic correlation structures, following recent applications in climate-related infrastructure vulnerability studies (Li, Dong & Zhu 2020).

4.4 Monte Carlo Analysis of Portfolio Uncertainty

Building on the sensitivity tests, a comprehensive Monte Carlo simulation is developed to explore joint uncertainty in key drivers of portfolio performance. For rainfall, each BoM station's record is first converted to an annual rainfall-exceedance index using a peaks-over-threshold approach fitted with Generalised Pareto and Poisson models, consistent with contemporary extreme-value practice in Australian flood estimation (Ball et al., 2019; BoM, 2016). Passenger trip volumes are modelled via realistic weekday/weekend patterns rather than arbitrary shocks, with weekend patronage typically 50-60% of weekday demand, while closure and TSR durations, freight volumes and key economic parameters follow bounded distributions aligned with current transport appraisal guidance (Transport and Infrastructure Council, 2020). The Monte Carlo model also applies a 1% per annum real growth in passenger value of time, together with 2% per year passenger-demand growth and 1.8% per year freight-demand growth, consistent with current economic-appraisal practice.

In each iteration, joint random draws are propagated through the AFR and cost-benefit modules over the 50-year horizon to produce distributions of PVB, NPV, BCR and reliability metrics (e.g., closure days), summarised by P10, P50 and P90. P10, P50 and P90 are the 10th, 50th and 90th percentiles of simulated outcomes; for benefit metrics (BCR, NPV, PVB) lower values are more adverse, while for cost metrics (failures, closure days, user costs) lower values are more favourable. This provides central and downside views of portfolio performance that are rarely available in conventional rail appraisals focused on expected values. For slopes with baseline BCR > 1 at 7%, the simulation reports per-slope

and portfolio BCR percentiles, showing which interventions remain defensible under downside conditions that drive disruption and funding scrutiny. Reductions in closure days and other tail outcomes are also expected to lower reputational exposure and improve customer confidence in wet-season reliability, although these effects are not monetised in the base results.

Figure 10 compares each slope's deterministic BCR at a 7% discount rate with its P10 BCR from the Monte Carlo simulation, where P10 is used as a downside measure representing the 10th-percentile BCR under relatively adverse combinations of rainfall, failure rates and demand. The deterministic BCR already approximates the central, P50, case, while optimistic P90 values are less relevant for funding decisions. Many slopes with deterministic BCRs close to or slightly above 1 have materially lower P10 values once joint uncertainty is taken into account, indicating that their expected benefits are not robust under unfavourable conditions, whereas only a small cluster maintains P10 BCRs above 1. This probabilistic view highlights which candidate projects remain justified on a downside basis and which are more sensitive to variability, providing a level of robustness information that is seldom reported in traditional rail cost-benefit analyses and offering a clearer basis for selecting proactive treatments.

5 INTERVENTION PACKAGING AND DECISION RULES

Building on the slope-level BCR distributions and downside metrics reported in Section 4.4, this study compares alternative intervention programmes assembled from the same pool of candidate works.

The current asset strategy of Sydney Trains is to secure funding to future-proof slopes prioritised by their risk rating, while also repairing newly failed slopes to restore train services. In practice, however, funding constraints mean that many B+ slopes have not been treated proactively; instead, their risk is often managed temporarily down to B-P1 through hazard-management measures such as speed restrictions, surveillance and monitoring.

Using the comprehensive Monte Carlo results, the same pool of candidate works was reassembled into three alternative programmes: (i) a reactive-only approach with no proactive treatment; (ii) a risk-rating-based programme that proactively treats all predicted B+ slopes; and (iii) a value-based programme that proactively treats all slopes with a median BCR greater than 1 at a 7% discount rate, regardless of risk band.

Figure 11 summarises Monte Carlo results from 2,000 runs for P10, P50 and P90 outcomes, where the P-values represent increasingly conservative conditions (P10 better-than-expected, P50 central estimate, P90 worse-than-expected). The three strategies are: (a) a reactive baseline with no proactive works, (b) a rating-based proactive programme that treats all B+ slopes, and (c) a value-based programme that treats all slopes with BCR > 1. Compared with taking no proactive action, both proactive strategies reduce risk and total cost. The rating-based programme, while intervening on only a small subset of the network, reduces total failures by about 6%, line-closure days by about 12% and passenger-plus-freight operational cost by roughly 10%, and lowers total capital expenditure (CAPEX) by about 13%. Extending treatment to all slopes with BCR greater than 1 strengthens these

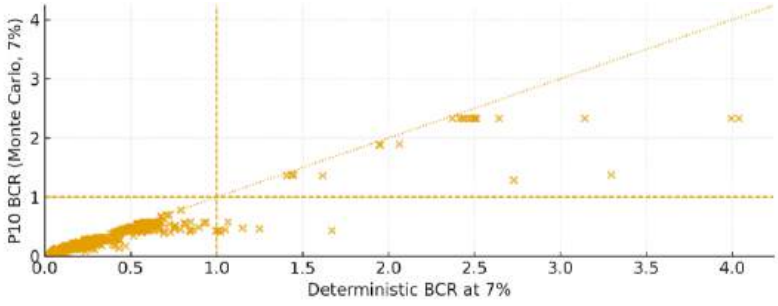


Figure 10: Robustness of slope-level BCR under combined uncertainty

gains, with failures down around 8%, closure days down about 20%, operational cost down about 17% and CAPEX down about 18%, within a similar lifetime envelope. In both proactive strategies, the present value of planned works is only around 2-3% of the total investment envelope I0, indicating that a relatively small, well targeted share of capital can deliver disproportionately large reductions in expected failures, closure days and user costs. Nevertheless, a sizeable portion of cost remains unavoidable because fundamentals such as geology, corridor geometry and the frequency of extreme wet seasons cannot be changed; proactive intervention manages these risks but cannot eliminate them entirely.

When assessed probabilistically, the value-based programme delivers markedly better performance, particularly on downside outcomes. The scope of treatment increases by more than half, capital investment rises only moderately, and the total present value of benefits increases by roughly one third, while P10 results show larger reductions in failures, closure days and operational cost than the rating-based programme. Although the combined BCR of the selected slopes decreases slightly as additional moderate-return projects are included, the value-based strategy concentrates investment on slopes whose benefits remain strong under adverse combinations of rainfall, failure rates and demand, whereas the rating-based programme includes several slopes whose P50 or even P10 BCR falls below 1. In other words, a value-based selection rule not only delivers greater expected customer and network benefit per dollar invested, but also improves resilience to tail-risk conditions that matter most for customers, operators and funders.

6 IMPLEMENTATION AND CONCLUSIONS

The framework is intended for routine use rather than a one-off study. Once the slope inventory, block sections, rainfall-exceedance metrics and cost parameters are established, the analysis can be re-run efficiently when new rainfall seasons, updated IFDs or revised risk ratings become available. This allows slope investment priorities to be refreshed regularly so decisions reflect current wet-weather experience rather than a single baseline year. By linking slopes to specific block sections, passenger areas and freight corridors, the model expresses engineering works directly in terms of reliability and user-cost changes for customers and freight.

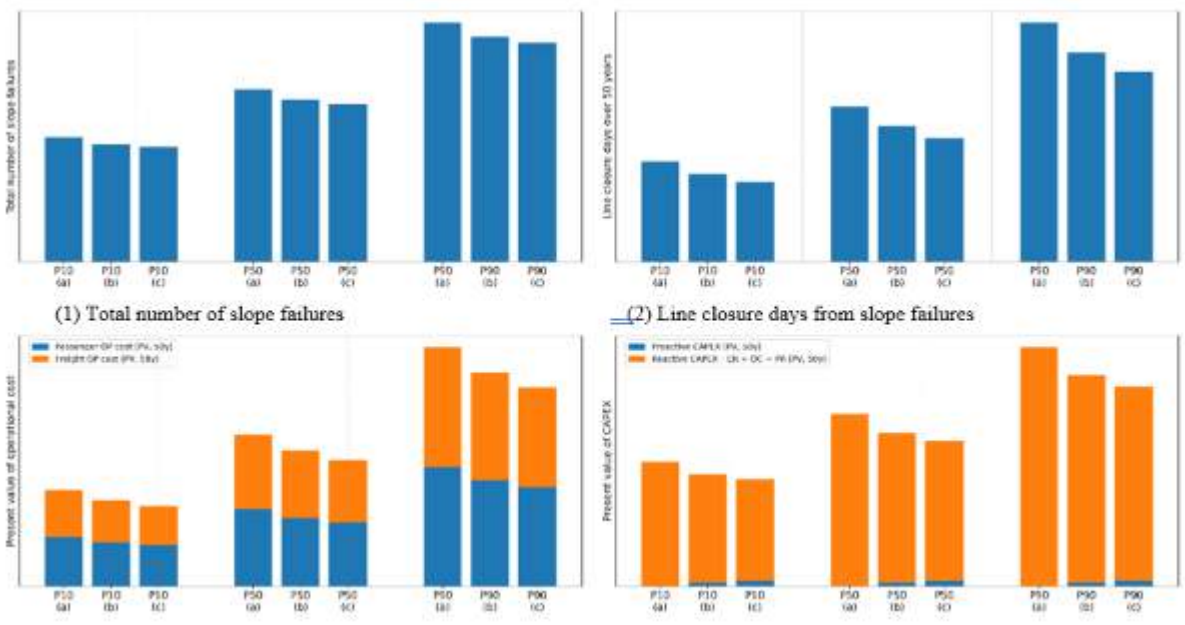


Figure 11: Portfolio outcomes over 50 years under three intervention strategies

A key strength of the framework is that it operates at the level of individual slopes rather than relying on corridor- or section-level total scores. By incorporating local network characteristics, demand patterns and rainfall-exceedance metrics for each slope-block-section pair, it represents operational impacts more precisely than traditional corridor frameworks and helps target scarce funding where interventions deliver the greatest reduction in service disruption. In contrast to prevailing practice, which typically prioritises slopes using qualitative risk scores based on likelihood multiplied by consequence, the framework quantifies the marginal economic value of each intervention and reveals where risk-based priorities diverge from value-based priorities. Some high-rated slopes deliver weak or negative net benefits, while lower-rated assets can generate strong economic returns once operational and user impacts are monetised. By embedding interventions in a probabilistic CBA and framing them against an explicit counterfactual baseline, the approach supports not only better forward-looking selection but also ex-post accountability for how resilience funding changes expected failures, closures and user costs at a network level.

Implementation is supported by an assumption log and parameter workbook aligned with current appraisal guidance, together with basic version control for data and code. The framework has already supported multiple Sydney Trains business cases and rail-resilience funding applications, showing that data-driven, event-conditioned probabilistic CBA can identify high-value interventions, avoid uneconomic treatments, and concentrate scarce funding where it delivers the greatest reduction in expected and tail-risk closures. The current implementation rests on several simplifying assumptions, including time-stationary annual failure-rate distributions, rainfall-exceedance exposure proxied by the nearest BoM station, deterministic block-section closure patterns and a fixed real growth rate in passenger value of time; these are recorded in the assumption log and can be relaxed as richer data and more dynamic operational processes become available.

The framework relies on several simplifying assumptions, including station-based rainfall attribution, stylised operational responses and time-stationary annual failure-rate structures. These choices reflect current data availability and governance practice rather than theoretical optimality, and highlight consequence modelling and recovery dynamics as priority areas for future research. The analysis is descriptive rather than causal, with projections conditioned on observed relationships and not interpreted as structural estimates of treatment effects.

As extreme rainfall intensifies, this integrated framework offers a scalable and transferable approach for embedding climate-resilience economics into asset management across railways and other wet-weather-sensitive civil infrastructure. Although demonstrated for railway slopes, the same hazard-vulnerability-exposure-value logic can be applied to assets such as culverts, surface drainage and retaining structures, and to precinct-level resilience planning and adaptive portfolio management as climate signals, asset condition and operating patterns evolve. Extension to other asset classes requires mode-specific hazard proxies and operational response definitions, but the underlying logic remains unchanged. Its event-conditioned probabilistic structure also aligns with digital engineering and scenario-based planning, enabling stress-testing of future network performance under compound climate and demand pressures. More broadly, the study shows how investment decisions can be framed as an interpretable risk-and-value problem rather than one guided solely by condition scores, risk ratings or engineering judgement. The framework is positioned as a decision-support tool to inform internal prioritisation and business-case development, rather than as a substitute for formal policy appraisal or funding allocation processes. For practitioners, this implies routinely linking asset risk to operational performance and customer impacts, supported by ex-post monitoring of realised benefits relative to reactive management. For policymakers, it reinforces the need to account for whole-of-system passenger and freight disruption, including diversion and trip suppression during prolonged closures, not only internal delay metrics. For researchers, it identifies consequence modelling as a key limitation and priority for future work, while demonstrating that probabilistic, value-focused appraisal can still support credible decisions under deep uncertainty.

7 ACKNOWLEDGEMENTS

The author gratefully acknowledges the support for this work provided by the Asset Management Council through the 2025 Asset Management Research Scholarship. The author would also like to thank the asset management and network maintenance teams at Sydney Trains for their productive discussions and for providing the initial data that inspired this research.

8 REFERENCES

[1] AECOM (2025) Risk, cost, and performance: a balanced approach to infrastructure asset management. AECOM Technical Report.

[2] Ball, J.E., Babister, M., Nathan, R., Weeks, W., Weinmann, E., Retallick, M. and Testoni, I. (eds) (2019) *Australian rainfall and runoff: a guide to flood estimation.* Canberra: Engineers Australia.

[3] Baral, A. and Shahandashti, S.M. (2022) 'Risk-averse rehabilitation decision framework for roadside slopes', *Journal of Infrastructure Preservation and Resilience,* 3(1), 5.

[4] Boardman, A.E., Greenberg, D.H., Vining, A.R. and Weimer, D.L. (2018) *Cost-benefit analysis: concepts and practice.* 5th edn. Cambridge: Cambridge University Press.

[5] BoM (Bureau of Meteorology) (2016) *Intensity-Frequency-Duration (IFD) data system.* Canberra: Bureau of Meteorology.

[6] Cardona, O.D., van Aalst, M.K., Birkmann, J., Fordham, M., McGregor, G. et al. (2012) 'Determinants of risk: exposure and vulnerability', in Field, C.B., Barros, V., Stocker, T.F. and Qin, D. (eds) *Managing the risks of extreme events and disasters to advance climate change adaptation.* Special Report of the Intergovernmental Panel on Climate Change (IPCC). Cambridge: Cambridge University Press.

[7] Corominas, J., van Westen, C., Frattini, P., Cascini, L., Malet, J.-P., Fotopoulou, S. and Catani, F. (2014) 'Recommendations for the quantitative analysis of landslide risk', *Bulletin of Engineering Geology and the Environment,* 73(2), pp. 209-263.

[8] CSIRO and BoM (Bureau of Meteorology) (2015) *Climate change in Australia: information for Australia's natural resource management regions.* Canberra: CSIRO and Bureau of Meteorology.

[9] Dong, B.-X., Shan, M. and Hwang, B.-G. (2022) 'Simulation of transportation infrastructures resilience: a comprehensive review', *Environmental Science and Pollution Research,* 29, pp. 8845-8869.

[10] Elansari, K., Idrissi, A. and Tifernine, H. (2024) 'Machine learning to predict railway infrastructure defects', in *Advances in Intelligent Systems and Computing.* Cham: Springer.

[11] Huang, J. (2025) 'Slope stability analyses: from factor of safety to quantitative risk assessment', *CSEE Proceedings.*

[12] IPCC (Intergovernmental Panel on Climate Change) (2021) *Climate change 2021: the physical science basis.* Contribution of Working Group I to the Sixth Assessment Report of the IPCC. Cambridge: Cambridge University Press.

[13] Janeková, J. and Onofrejová, D. (2015) 'Implementation of Monte Carlo simulation in investment decision making', *Acta Simulatio,* 1(2), pp. 27-32.

[14] Li, Y., Dong, Y. and Zhu, D. (2020) 'Copula-based vulnerability analysis of civil infrastructure subjected to hurricanes', *Frontiers in Built Environment,* 6, 571911.

[15] Liu, Y., Zhang, L., Chen, G., Zhang, S. and Li, T. (2023) 'Machine learning-based fragility curves for landslide risk assessment', *Natural Hazards.*

[16] Meyer, R. and Kunreuther, H. (2017) *The ostrich paradox: why we underprepare for disasters.* Philadelphia: Wharton Digital Press.

[17] Network Rail (2021) *Weather resilience and climate change adaptation plan.* London: Network Rail.

[18] NSW Treasury (2018) *NSW Government guide to cost-benefit analysis.* Sydney: NSW Treasury.

[19] Pant, R., Hall, J.W., Blainey, S.P. and Preston, J. (2016) 'Vulnerability assessment framework for interdependent critical infrastructures under flood hazards', *Risk Analysis,* 36(2), pp. 227-239.

[20] Pregnolato, M. and Dawson, R.J. (2018) 'Benchmarking urban flood resilience in transportation networks', *Transportation Research Part D: Transport and Environment,* 63, pp. 288-302.

[21] Productivity Commission (2014) *Natural disaster funding arrangements.* Inquiry Report no. 74. Canberra: Productivity Commission.

[22] Qian, L., Zhang, J., Wang, X. and Chen, Y. (2025) 'Dynamic risk assessment of collapse geological hazards on highway slopes', *Atmosphere.*

[23] Sydney Trains (2014) *t-hr-ci-12100-st Geotechnical risk assessment and hazard management.* Sydney: Sydney Trains.

[24] Transport and Infrastructure Council (2020) ATAP T7: *Risk and Uncertainty Assessment.* Australian Transport Assessment and Planning Guidelines. Canberra: Transport and Infrastructure Council.

[25] Transport for NSW (2025) *Economic parameter values for cost-benefit analysis.* Sydney: Transport for NSW.

[26] UNISDR (United Nations International Strategy for Disaster Reduction) (2017) *Words into action guidelines: national disaster risk assessment - governance system, methodologies, and use of results.* Geneva: UNISDR.

[27] Upadhyay, V. (2024) *Landslide hazard risk and vulnerability monitoring: GIS based approach.* Cham: Springer.

[28] Wasko, C. (2021) 'A global assessment of trends in annual maximum rainfall', *Water Resources Research,* 57(2).

[29] Wasko, C. and Sharma, A. (2017) 'Global assessment of flood and storm extremes with increased temperatures', *Scientific Reports,* 7, 7945.

[30] World Economic Forum (2024) *Implementing a life-cycle approach to infrastructure: a policy roadmap for cities.* Geneva: World Economic Forum.

[31] Zhou, Y., Wang, J., Yang, X. and Hu, S. (2021) 'Resilience of urban transport networks under flood hazards', *International Journal of Disaster Risk Reduction,* 57, 102192.

TECH 3



ARTICLE 3 – AI-Powered Maturity Model for Sustainable Asset Management and Data-Driven Strategy

Ramakrishnan Vijayakumaran, CQU

SUMMARY

Achieving sustainability and resilience in asset-intensive industries requires integrating operational performance data and sustainability objectives into asset management practices. However, essential artefacts within an asset management system, such as Policies, Strategic Asset Management Plans (SAMPs), and Asset Management Plans (AMPs) are often developed and reviewed as static documents, resulting in a loss of line of sight between declared asset management intent and day-to-day operational reality. This fragmentation leads to limited alignment with evolving performance data and sustainability frameworks. This research explores the potential of Artificial Intelligence (AI) to address this challenge through the development of an AI-powered maturity model for sustainable asset management. The proposed research investigates how Generative AI and Agentic AI can be applied to interpret asset management artefacts and reconcile declared intent with operational and sustainability evidence, enabling more frequent and data-informed evaluation of alignment. The model integrates principles from ISO 55001, the Sustainable Development Goals (SDGs), and circular economy frameworks to assess organizational maturity in aligning strategy, data, and sustainability objectives. By examining case studies within asset-intensive industries, the study aims to identify practical applications of AI that support transparency, energy efficiency, and decarbonization through improved governance and continuous alignment between strategy and execution. The expected outcomes include a validated maturity assessment approach and a supporting framework that enables adaptive, sustainability-ready asset management strategies. This research contributes to advancing sustainable infrastructure planning

by demonstrating how emerging AI technologies can support responsible, evidence-based decision-making that enhances long-term environmental and societal value.

Keywords: Asset management maturity, Loss of line of sight, Sustainable asset management, Data-driven decision-making, Generative Artificial Intelligence (GenAI), Agentic Artificial Intelligence, ISO 55001, Sustainability integration

1 INTRODUCTION

Asset-intensive organizations rely on structured asset management system artefacts—such as asset management policies, Strategic Asset Management Plans (SAMPs), and Asset Management Plans (AMPs)—to translate organizational objectives into operational decision-making. International standards emphasise that these artefacts should maintain alignment between strategy, planning, and execution to enable value realisation from assets (ISO, 2024). In practice, however, such artefacts are frequently developed and reviewed as static documents, weakly connected to evolving operational performance and sustainability outcomes. This fragmentation contributes to a persistent loss of line of sight between declared asset management intent and day-to-day execution (Woodhouse, 2014).

Loss of line of sight is further intensified by differences in decision-making cadence across organizational levels. Strategic intent is typically reviewed annually, tactical plans are updated periodically, and operational and maintenance decisions are executed daily. In dynamic operating environments, this mismatch allows strategic intent to drift from operational reality long before formal governance processes detect it (Hastings, 2010; IAM, 2024). Although asset-intensive industries now generate large volumes of operational, maintenance, and energy-related data, this data is rarely used in a systematic and continuous way to verify whether asset management intent is being realised in practice (ISO, 2024).

The implications of this gap are particularly evident in sustainability and ESG initiatives. While sustainability objectives are increasingly articulated at an organizational level, they are often implemented as parallel reporting or compliance activities rather than being embedded within asset management governance and decision-making. This separation has historically reflected a structural limitation rather than a lack of intent: asset management intent is predominantly expressed through unstructured narrative documents that are difficult to interpret and reconcile with operational evidence at scale.

Recent advances in Artificial Intelligence, particularly Generative AI (Segall, 2025) and agent-based approaches, present an opportunity to revisit these limitations. Generative AI enables interpretation of unstructured textual artefacts, while agent-based approaches support repeatable and adaptive analytical processes (Gartner, 2023). This research investigates how such AI capabilities can be applied to examine alignment between asset management artefacts, operational performance data, and sustainability objectives, and how they may support earlier detection of misalignment as conditions evolve.

Within this context, this paper proposes and explores an AI-powered maturity model for sustainable asset management and data-driven strategy. The model evaluates how alignment between declared intent and execution evidence progresses across different levels of maturity. Rather than positioning AI as a replacement for established asset management practices, the research investigates its role as an analytical support mechanism that enables more frequent, evidence-informed evaluation of asset management artefacts consistent with ISO-aligned governance principles.

The remainder of this paper proceeds by outlining the problem statement and research objectives, reviewing relevant literature, presenting the proposed maturity model and methodology, discussing results and practical implications, and concluding with key contributions, limitations, and directions for future research.

2 PROBLEM STATEMENT

Asset management in asset-intensive industries is formally structured around a hierarchy of strategic, tactical, and operational artefacts intended to maintain a clear line of sight between organizational objectives and day-to-day execution. At the strategic level, asset management policy and objectives are

translated into a Strategic Asset Management Plan (SAMP), which informs tactical artefacts such as Asset Management Plans (AMPs), lifecycle plans, and work management strategies. Collectively, these artefacts represent the organization's declared intent for how assets should deliver value across their lifecycle, consistent with the principles outlined in ISO 55000 and ISO 55001 (ISO, 2024).

In parallel, operational environments generate extensive execution evidence through maintenance activities, asset performance management processes, and day-to-day operations. This evidence includes information related to reliability, availability, cost, safety, risk, energy consumption, and environmental performance. ISO 55001 emphasises that such evidence should inform decision-making and support continual improvement across the asset management system (ISO, 2024). In practice, however, the feedback loop between operational evidence and strategic or tactical intent is often weak or discontinuous.

As illustrated conceptually in Figure 1 , a persistent gap exists between declared asset management intent and operational execution. Strategic and tactical artefacts are typically reviewed on annual or periodic cycles, while operational and maintenance decisions occur daily. This misalignment in review cadence has been widely recognised as a contributor to strategic drift, where asset management intent progressively diverges from operational reality in dynamic environments (Hastings, 2010; IAM, 2024). Once embedded, such drift may remain undetected until formal reviews or audits are conducted.

This gap is further reinforced by the differing nature of asset management artefacts and operational data. Strategic and tactical intent is predominantly captured in unstructured narrative documents, whereas operational evidence is held within structured digital systems such as EAM, CMMS, and performance analytics platforms. Historically, organisations have lacked scalable mechanisms to systematically interpret intent expressed in documents and reconcile it with operational evidence on a continuous basis, relying instead on manual reviews and point-in-time assessments (Woodhouse, 2014).

The implications of this limitation are particularly evident in the treatment of sustainability and ESG objectives. Although sustainability goals are increasingly embedded in corporate policies and strategies, they are often implemented as parallel compliance or reporting initiatives rather than as integral components of asset management

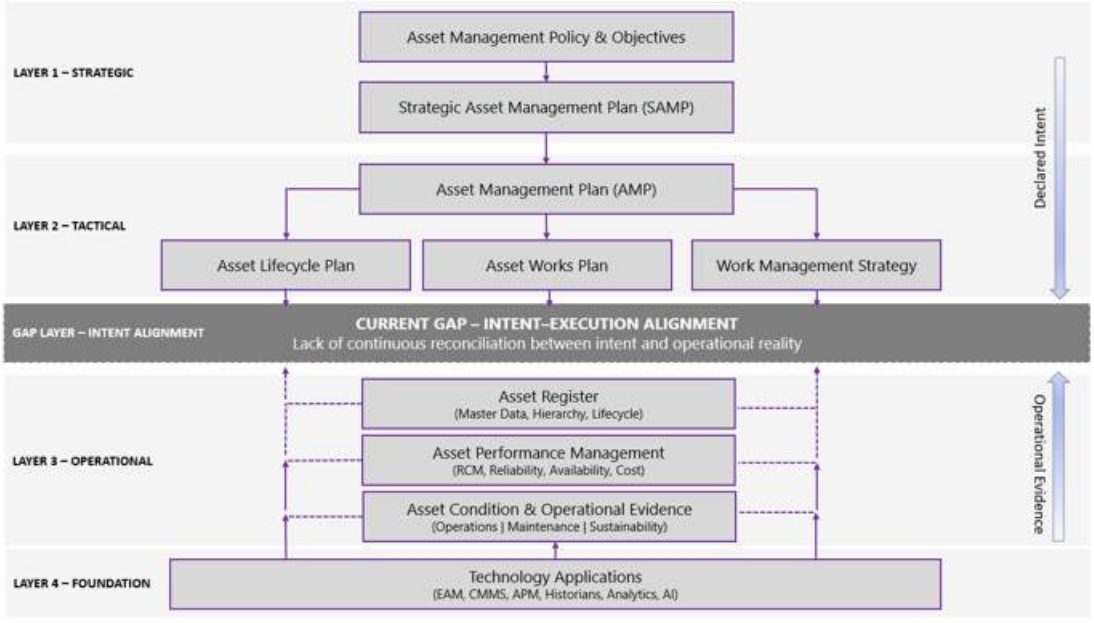


Figure 1: Gap between declared strategic intent and operational evidence

governance. As a result, sustainability objectives are not consistently translated into asset-level decision criteria or performance thresholds, limiting their influence on lifecycle optimisation and baseline performance improvement (ISO, 2024).

In operating contexts characterised by increasing asset complexity, digitalisation, and sustainability expectations, the absence of a continuous mechanism to reconcile asset management intent with execution evidence represents a fundamental limitation of current practice. This research addresses this problem by investigating how emerging Artificial Intelligence capabilities, specifically Generative AI and Agentic AI, can be applied to support continuous reconciliation between asset management artefacts and operational evidence, enabling earlier detection of misalignment and improved integration of sustainability objectives.

3 RESEARCH OBJECTIVES AND QUESTIONS

The objective of this research is to investigate how emerging Artificial Intelligence capabilities can be applied to address the persistent loss of line of sight between asset management intent and operational execution in asset-intensive organizations. Consistent with the principles of alignment, evidence-based decision-making, and continual improvement articulated in ISO 55000 and ISO 55001, the study examines how Generative AI and Agentic AI may support systematic reconciliation between asset management artefacts, operational performance data, and sustainability objectives (ISO, 2024).

While current asset management practice recognises the importance of maintaining alignment across policy, strategy, planning, and operations, maturity assessments and governance reviews are typically conducted periodically and rely heavily on qualitative judgement (IAM, 2024; Woodhouse, 2014). In increasingly data-rich and dynamic operating environments, such approaches offer limited capability to detect emerging misalignment or strategic drift in a timely manner (Hastings, 2010). This research therefore investigates whether AI-enabled interpretation of unstructured artefacts and adaptive analytical processes can enhance the effectiveness and frequency of alignment verification.

A further objective is to examine how sustainability (Dong et al., 2024) and ESG objectives—often implemented as parallel compliance or reporting initiatives—can be more effectively embedded within asset management governance and asset-level decision-making. Prior studies indicate that sustainability objectives struggle to influence operational outcomes when they are not systematically integrated into asset management artefacts and performance evaluation processes.

Guided by these objectives, the research addresses the following questions:

- How are asset management system artefacts currently structured and governed, and what factors contributing to loss of line of sight between declared intent and operational execution (ISO, 2024; IAM, 2024)?
- How can Generative AI be applied to interpret unstructured asset management artefacts and extract intent relevant to performance and sustainability objectives (Woodhouse, 2014)?
- In what ways can Agentic AI support repeatable and adaptive evaluation of alignment between asset management intent, operational evidence, and sustainability outcomes (Hastings, 2010; ISO, 2024)?
- How can an AI-powered maturity model support progressive improvement in alignment and sustainability integration across different levels of asset management maturity (IAM, 2025; ISO, 2024)?

By addressing these questions, the research seeks to advance understanding of how AI-enabled analytical capabilities can enhance asset management maturity, strengthen governance, and support the integration of sustainability objectives into data-driven asset management practice.

4 BACKGROUND AND RELATED WORK

Asset management standards and professional guidance emphasise alignment between organizational objectives, asset management policy, strategy, planning, and operational execution as a prerequisite for value realisation from assets (ISO, 2024). Asset management artefacts such as policies, Strategic Asset Management Plans (SAMPs), and Asset Management Plans (AMPs) serve as the primary mechanisms through which strategic intent is articulated and governed. Professional bodies,

including the Institute of Asset Management, have further operationalised these principles through maturity models that describe progressive capability development across governance, planning, risk, and performance dimensions (Institute of Asset Management, 2015; IAM, 2024).

While maturity models provide valuable structure, their application is typically based on qualitative self-assessment and periodic reviews. Prior research indicates that such approaches offer limited capability to continuously verify alignment or respond dynamically to changing operational conditions, allowing strategic drift to persist between formal assessment cycles (Woodhouse, 2014; Hastings, 2010).

Sustainability and ESG considerations are increasingly embedded in organizational strategies in response to regulatory, investor, and societal expectations (United Nations, 2015). However, existing literature highlights that sustainability objectives are often treated as parallel reporting or compliance activities rather than being systematically integrated into asset management governance. When managed separately from asset management artefacts and lifecycle planning, sustainability objectives have limited influence on asset-level decisions and baseline performance improvement (Schuman and Brent, 2005).

Digitalisation has significantly increased the availability of operational, maintenance, and performance data in asset-intensive industries. Asset performance management systems, condition monitoring technologies, and analytics platforms are widely used to support operational optimisation. However, their application has primarily focused on improving reliability, cost, and risk outcomes, rather than on validating or adapting strategic and tactical asset management intent (Hastings, 2010). As a result, the potential of operational data to inform governance-level alignment and sustainability integration remains underutilised.

Recent advances in Artificial Intelligence, particularly Generative AI and agent-based approaches, present new opportunities to address these limitations. Generative AI enables interpretation of unstructured textual artefacts, while agent-based approaches support repeatable and adaptive analytical processes (Gartner, 2023). Despite this potential, the application of AI to asset management artefacts, maturity evaluation, and alignment verification remains an emerging area with limited guidance in existing literature.

In summary, existing asset management maturity models are predominantly static, sustainability frameworks are weakly integrated into asset management governance, and digital technologies are largely applied at the operational level. The intersection of these domains—continuous, evidence-based alignment between asset management intent, operational execution, and sustainability outcomes—remains underdeveloped. This research addresses this gap by investigating how Generative AI and Agentic AI can support an AI-powered maturity model that enhances alignment and enables adaptive, sustainability-ready asset management.

5 AI-POWERED MATURITY MODEL FRAMEWORK

5.1 Framework Overview

Asset management standards emphasise that value is realised when organizational objectives, asset management policy, strategy, planning, and operational control are coherently aligned and subject to continual improvement (ISO, 2024). In practice, maintaining this alignment is challenging due to fragmented artefacts, periodic review cycles, and the increasing volume and complexity of operational and sustainability data. The proposed maturity model addresses this challenge by introducing an AI-enabled mechanism to interpret declared intent, analyse execution evidence, and evaluate alignment on a more continuous basis rather than at discrete points in time.

As illustrated in Figure 2 , the framework combines interpretation of asset management intent, verification of execution against that intent, and governance-oriented decision support within a structured maturity model. Artificial Intelligence is investigated as an enabling mechanism that supports interpretation, comparison, and adaptive evaluation, rather than as a replacement for established asset management governance processes.

The framework is structured around two core elements:

- Maturity dimensions, which define the capability areas required to maintain alignment between intent and execution.

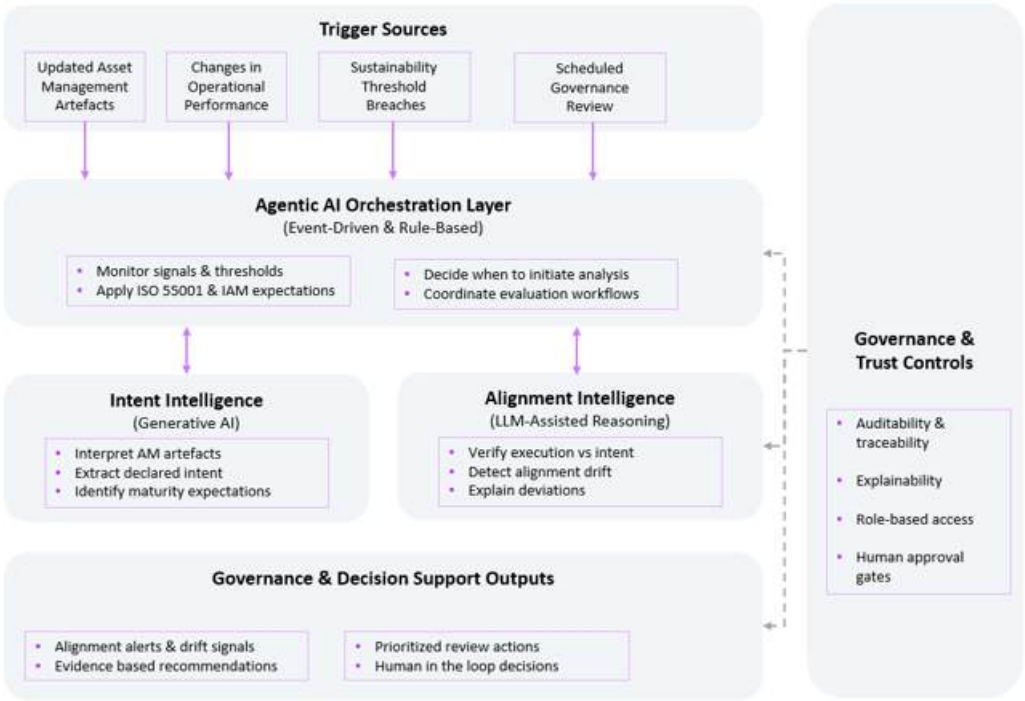


Figure 2: AI-Enabled Alignment Evaluation Framework for Asset Management Maturity

- Maturity levels, which describe progressive stages of capability development across those dimensions.

5.1.1 Maturity Dimensions

The proposed maturity model (Mahmood et al., 2014) evaluates alignment across four interrelated dimensions reflecting key asset management and sustainability principles as illustrated in Table 1.

5.1.2 Maturity Levels

The progression of maturity levels used in this study is summarised in Table 2. The maturity levels defined in this study are intentionally technology-agnostic and aligned with established asset management maturity concepts. Artificial Intelligence is investigated not as a defining characteristic of maturity, but as an enabling capability that enhances how alignment between intent and execution can be evaluated, particularly at higher maturity levels. As maturity increases, AI-enabled alignment evaluation supports greater frequency, consistency, and adaptability in governance processes, without altering the underlying maturity structure.

5.1.3 Maturity Levels

Within the framework, Artificial Intelligence supports maturity evaluation through two complementary analytical functions. Generative AI is investigated as a means of interpreting unstructured asset management artefacts, enabling extraction of declared intent, identification of sustainability-related expectations, and detection of inconsistencies across documents. Agentic AI is examined as an orchestration mechanism that enables repeatable and adaptive evaluation, triggering analysis in response to changes in artefacts, operational performance, or sustainability indicators. Together, these capabilities support continuous verification of alignment between intent and execution and enable earlier detection of maturity drift, while maintaining human oversight and governance controls.

6 METHODOLOGY

This research adopts a qualitative, design-oriented methodology to investigate how Generative AI and Agentic AI (Bandi *et al.,* 2025) may support improved alignment between asset management artefacts,

Table 1: Maturity Dimensions

Maturity Dimension	Description
Strategic Intent and Governance	Assesses the clarity and coherence of asset management intent as expressed in policies, objectives, and the Strategic Asset Management Plan, including the extent to which sustainability and ESG objectives are embedded and supported by governance mechanisms consistent with ISO 55001 (ISO, 2024).
Tactical Planning and Lifecycle Integration	Examines how strategic intent is translated into Asset Management Plans, lifecycle strategies, and work management approaches, and whether performance and sustainability expectations are reflected in asset-level planning and remain aligned with evolving operating conditions (IAM, 2024).
Operational Evidence and Data Integration	Focuses on the availability, quality, and use of operational, maintenance, and sustainability-related data to validate declared intent and support learning and improvement, rather than being confined solely to operational optimisation (Hastings, 2010).
Review, Learning, and Adaptation	Evaluates the mechanisms through which alignment is reviewed and improved over time, including review frequency, use of evidence in governance processes, and the ability to detect emerging misalignment before it results in degraded performance or missed sustainability outcomes (Woodhouse, 2014).

Table 2: Maturity Levels

Maturity Level	Characteristics
Level 1 – Initial	Asset management system artefacts exist in an ad hoc or compliance-driven form. Strategic and tactical intent is weakly defined, largely static, and disconnected from operational and sustainabili Alignment is continuously monitored and adaptively managed. Declared intent is dynamically reconciled with operational and sustainability evidence, enabling early detection of drift and supporting resilient, sustainability-ready asset management (IAM, 2025). ty evidence. Alignment verification is informal and reactive.
Level 2 – Managed	Core asset management system artefacts are established and governed. Alignment between intent and execution is assessed periodically through manual reviews, with limited use of operational or sustainability data to inform decision-making.
Level 3 – Defined	Strategic, tactical, and operational artefacts are consistently structured and linked. Operational performance data is used to inform reviews, but alignment verification remains largely periodic and retrospective.
Level 4 – Integrated	Alignment between asset management intent, operational execution, and sustainability objectives is actively managed. Evidence-informed reviews are conducted regularly, and sustainability considerations are embedded into asset-level planning and decision criteria.
Level 5 – Adaptive	Alignment is continuously monitored and adaptively managed. Declared intent is dynamically reconciled with operational and sustainability evidence, enabling early detection of drift and supporting resilient, sustainability-ready asset management (IAM, 2025).

operational performance data, and sustainability objectives. The approach is consistent with applied asset management research, where frameworks and models are examined to address practical governance challenges rather than to achieve statistical generalisation (Hastings, 2010; Woodhouse, 2014).

6.1 Research Approach

The study follows a design science research logic, in which a structured artefact—an AI-powered maturity model—is proposed and investigated to address the identified loss of line of sight between strategic intent and operational execution. Design science is well suited to asset management contexts as it enables systematic exploration of governance and decision-support mechanisms grounded in real-world practice (Schuman and Brent, 2005). The research focuses on feasibility and practical relevance, examining whether AI-enabled interpretation and evaluation of asset management artefacts can enhance maturity assessment and support more frequent alignment verification.

6.2 Development of the Maturity Model

The maturity model was developed through synthesis of established asset management principles and sustainability considerations. Core expectations were derived from ISO 55000 and ISO 55001, with emphasis on alignment between organizational objectives, policy, strategy, planning, and operational control (ISO, 2024). Guidance from the Institute of Asset Management informed maturity concepts and progression across capability levels (IAM, 2024). Sustainability considerations were incorporated by embedding asset-relevant themes, such as energy efficiency, emissions reduction, lifecycle optimisation, and resilience, into the maturity dimensions (United Nations, 2015).

6.3 AI-Enabled Evaluation Design

The methodology investigates two complementary AI roles. Generative AI is examined as a mechanism for interpreting unstructured asset management artefacts, enabling extraction of declared intent, identification of sustainability-related expectations, and detection of inconsistencies across documents. Agentic AI is explored as an orchestration mechanism that supports repeatable and adaptive evaluation, triggering analysis in response to changes in artefacts, operational performance, or sustainability indicators (Gartner, 2023). AI outputs are treated as analytical inputs, with final interpretation remaining grounded in asset management principles and professional judgement.

6.4 Case-Based Application and Analysis

An illustrative case-based approach is used to investigate application of the maturity model in an asset-intensive context. Data sources include asset management artefacts, selected operational and maintenance performance indicators, and sustainability-related data relevant to asset performance. Analysis is qualitative and interpretive, focusing on identifying patterns of alignment and governance implications rather than numerical scoring.

6.5 Scope and Limitations

The methodology is designed to investigate feasibility and practical value rather than to provide statistically generalisable results. Findings are influenced by the availability and quality of artefacts and data within the case context. These limitations highlight opportunities for future research involving multi-organisation studies and longitudinal maturity assessment.

7 RESULTS AND PRACTICAL IMPLICATIONS

This section summarises key observations from applying the proposed AI-powered maturity model within an illustrative asset-intensive context and outlines implications for asset management practice. Findings are qualitative and exploratory, consistent with the design-oriented methodology adopted in this study.

7.1 Observed Maturity and Alignment Characteristics

Application of the maturity model revealed partial alignment across strategic, tactical, and operational layers. Core asset management artefacts—such as the asset management policy, Strategic Asset Management Plan, and Asset Management Plans—were generally well structured and aligned with formal governance expectations, consistent with ISO 55001-aligned practice reported in prior studies (ISO, 2024; IAM, 2024).

However, limited evidence was found of systematic reconciliation between declared intent and operational performance data. While operational and

maintenance data were actively used for day-to-day decision-making, they were rarely applied to validate or refine strategic and tactical intent. Consequently, emerging performance trends, risks, and sustainability impacts were not consistently reflected in asset management artefacts, reinforcing the loss of line of sight identified in earlier sections (Hastings, 2010; Woodhouse, 2014).

7.2 Sustainability Integration Outcomes

The analysis highlighted a pronounced gap in the integration of sustainability and ESG objectives (Saeed et al., 2025). Sustainability targets were visible at policy or corporate strategy levels but were weakly translated into asset-level decision criteria or performance thresholds. Sustainability-related data, including energy consumption and emissions proxies, was commonly managed through separate reporting mechanisms rather than embedded within asset management planning and review processes. This finding aligns with existing literature indicating that sustainability initiatives often remain compliance-driven when not structurally integrated into asset management governance.

7.3 Contribution of AI-Enabled Analysis

The investigation demonstrated that Generative AI could support more systematic interpretation of unstructured asset management artefacts, enabling identification of inconsistencies where sustainability intent was articulated in policy but not reinforced in tactical or operational plans. Such gaps were not readily apparent through traditional manual review processes (Woodhouse, 2014).

The conceptual application of Agentic AI further illustrated potential to shift maturity evaluation from a periodic activity to a more adaptive capability. Repeatable, event-triggered analysis offers a mechanism to detect emerging misalignment earlier in the governance cycle, addressing limitations of conventional maturity assessments that rely on infrequent, retrospective reviews (IAM, 2024).

7.4 Practical Implications for Asset Management Practice

The findings suggest several implications for asset-intensive organizations:

- Strengthening line of sight: The maturity model supports systematic verification of whether asset management intent is being realised in practice, reinforcing ISO 55001 alignment principles (ISO, 2024).
- Embedding sustainability into governance: Evaluating sustainability as part of maturity encourages integration of environmental objectives into asset-level decision-making rather than parallel ESG reporting.
- Improving governance effectiveness: AI-enabled interpretation of artefacts enables more frequent, evidence-informed reviews, improving transparency and reducing reliance on manual assessments.
- Supporting progressive maturity improvement: The framework enables incremental improvement across maturity levels, rather than binary compliance outcomes.

For asset managers, the framework provides a practical means to link strategic intent with operational evidence. For sustainability leaders, it offers a mechanism to embed environmental objectives into asset management governance. For digital and analytics teams, it highlights an opportunity to extend the value of existing data platforms beyond operational optimisation toward strategic alignment.

7.5 Implications for Future Adoption

The primary value of the proposed approach lies not in assigning a single maturity score, but in enabling continuous learning and adaptation. As asset-intensive organizations face increasing uncertainty related to energy transition, regulation, and asset complexity, AI-enabled maturity evaluation offers a pathway to detect alignment drift early and support adaptive, sustainability-ready asset management consistent with established standards.

8 CONCLUSION AND FUTURE DIRECTIONS

This paper investigated the challenge of maintaining line of sight between asset management intent and operational execution in asset-intensive organizations, with particular emphasis on the integration of sustainability objectives into asset management governance. Despite widespread

adoption of formal asset management standards and increasing availability of operational data, the findings indicate that asset management artefacts often remain fragmented and weakly connected to execution evidence. This disconnect limits the ability of organizations to continuously verify alignment, detect strategic drift early, and translate sustainability intent into asset-level outcomes.

To address this challenge, the research proposed and explored an AI-powered maturity model for sustainable asset management and data-driven strategy. Grounded in ISO 55000 and ISO 55001 principles, the model investigates how Generative AI and Agentic AI can support systematic evaluation of alignment across strategic, tactical, and operational layers (ISO, 2024). The study demonstrates how AI-enabled interpretation of unstructured artefacts and adaptive evaluation processes can enhance the frequency and effectiveness of alignment verification beyond traditional, periodic review mechanisms.

The illustrative application of the maturity model highlights its potential to strengthen governance, improve transparency, and embed sustainability considerations into asset management decision-making. Rather than treating ESG initiatives as parallel compliance activities, the proposed approach enables sustainability objectives to be evaluated alongside performance, risk, and lifecycle outcomes. The primary value of the model lies not in point-in-time scoring, but in supporting continuous learning and adaptive improvement through sustained alignment verification.

This study has several limitations. The research adopts a qualitative, case-based approach intended to examine feasibility and practical relevance rather than to provide statistically generalisable results. Findings are influenced by the availability and quality of artefacts and data within the case context, and further work is required to formalise evaluation logic, governance controls, and integration with enterprise asset management systems.

Future research should examine application of the maturity model across multiple organizations and asset classes, supported by longitudinal studies to observe how alignment and maturity evolve over time. Further investigation is also warranted into the development of quantitative alignment indicators, integration with sustainability reporting and assurance processes, and the role of AI-enabled maturity models in supporting regulatory compliance and accountability.

In an environment characterised by increasing complexity, digitalisation, and sustainability expectations, maintaining continuous line of sight between intent and execution is becoming a critical capability. By investigating how Artificial Intelligence can support this capability within an established asset management framework, this research contributes toward advancing more adaptive, evidence-based, and sustainability-ready asset management practice.

9 REFERENCES

[1] Hastings, N.A.J., 2010. Physical asset management. 2nd ed. London: Springer.

[2] Segall, R.S. (2025) 'Generative AI (Artificial Intelligence): What Is It? & What Are Its Inter- And Transdisciplinary Applications?', Journal of Systemics, Cybernetics and Informatics, 23(7), pp. 116–125. Available at: https://doi.org/10.54808/jsci.23.07.116.

[3] Gartner, 2023. Emerging technologies: Generative AI and autonomous agents. Stamford, CT: Gartner Research.

[4] Dong, Y. et al. (2024) 'A Framework for Developing Information Requirements for Built Asset Management Based on Organisational Sustainability Objectives', in. springer nature switzerland, pp. 310–320. Available at: https://doi.org/10.1007/978-3-031-56121-4_30.

[5] Saeed, L. et al. (2025) 'The Impact of Sustainability Initiatives on Corporate Financial Performance: A Quantitative Analysis', The Critical Review of Social Sciences Studies, 3(2), pp. 1983–1997. Available at: https://doi.org/10.59075/v4wzdp20.

[6] Bandi, A. et al. (2025) 'The Rise of Agentic AI: A Review of Definitions, Frameworks, Architectures, Applications, Evaluation Metrics, and Challenges', Future Internet, 17(9), p. 404. Available at: https://doi.org/10.3390/fi17090404.

[7] Mahmood, M.N. et al. (2014) 'Towards an Integrated Maturity Model of Asset Management Capabilities', in. springer, pp. 431–441. Available at: https://doi.org/10.1007/978-3-319-06966-1_39.

[8] Institute of Asset Management (IAM), 2024. Asset Management – An Anatomy. Version 4. Bristol: Institute of Asset Management.

[9] Institute of Asset Management (IAM), 2025. Self Assessment Methodology 2025 (SAM2025). Bristol: Institute of Asset Management.

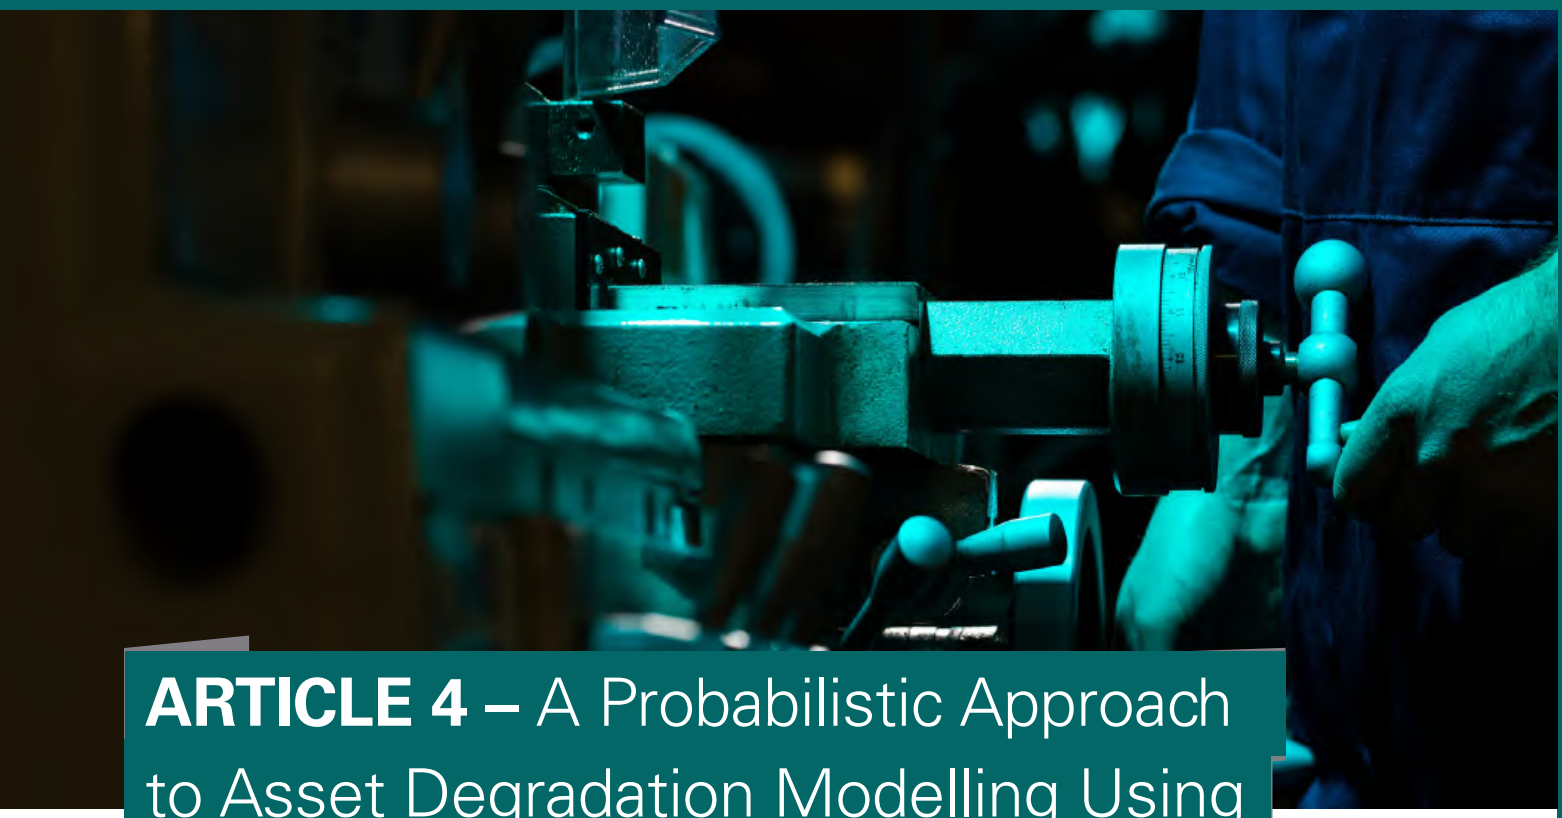
[10] International Organization for Standardization (ISO), 2024. ISO 55000:2024 Asset management — Vocabulary, overview and principles.

[11] International Organization for Standardization (ISO), 2024. ISO 55001:2024 Asset management — Asset management system — Requirements.

[12] Schuman, C.A. and Brent, A.C., 2005. Asset life cycle management: Towards improving physical asset performance in the process industry. International Journal of Operations & Production Management, 25(6), pp.566–579.

[13] United Nations, 2015. Transforming our world: The 2030 Agenda for Sustainable Development. New York: United Nations.

[14] Woodhouse, J., 2014. Asset management decision-making. London: ICE Publishing.



ARTICLE 4 – A Probabilistic Approach to Asset Degradation Modelling Using Expected Value

Hanbit Cho, AssetFuture & Jeremy Ngo, Woolworths

SUMMARY

Predicting the degradation of assets in Asset Management Systems (AMS) has generally followed a deterministic approach, where the prescribed design life is used to determine when assets will reach their End-of-Life (EOL). As a result, the probabilistic nature of assets, which allows for the possibility of failure occurring before or after the end of their design life, is often overlooked. An alternative probabilistic degradation modelling-based approach, using Expected Value (EV) to calculate the costs associated with the probability of asset failure occurring earlier or later than its expected operational life, is proposed. This method was tested using historical data to determine the probability distribution of asset failure via a Cumulative Distribution Function (CDF), based on the condition ratings of assets. The EV for each asset was then calculated to estimate the costs associated with the probability of asset failure, given the condition of the assets and the generated asset failure distribution. Results from cost forecasting show that, when combined with an adjustable risk threshold, asset portfolios can be maintained more economically than under traditional methods that rely on total asset replacement costs. This outcome is most evident when only the costs of assets associated with a high probability of the failure are considered. This approach provides an alternative perspective on asset management by shifting from deterministic to probabilistic degradation modelling. It uses EV, a well-established concept for quantifying the cost of risk in asset failure to forecast asset replacement costs.

Keywords: AAsset management, Probabilistic degradation, Replacement cost, Design life

1 INTRODUCTION

Asset Management Systems (AMS) have generally adopted a deterministic approach to predicting asset degradation by relying on the design life prescribed in the Original Equipment Manufacturer (OEM) guidelines. As a result, these systems will indicate that assets will reach their End-of-Life (EOL) once the prescribed design life has expired. This operational framework, which adheres to OEM recommendations, delivers high levels of asset

reliability, safety and performance (Okirie et al., 2024). However, using a deterministic degradation model as the basis for capital planning is inconsistent with the probabilistic nature of assets, which accepts the possibility that an asset may fail earlier or later than its expected design life.

To overcome the uncertainty in predicting when assets will reach their EOL state, extensive research has been conducted to estimate degradation using probabilistic models. Many of these studies focus on key structural components such as reinforced concrete, where various deterioration factors, including corrosion and fatigue, are modelled to account for uncertainty in determining asset condition state over time (see Val et al., 2025 for a literature review). A variety of statistical methods have also been demonstrated, including Monte Carlo Simulation (MCS) to model deterioration at different points in time (Sudret, 2008), Bayesian inference to update prior assumptions about degradation uncertainties (Ma et al., 2013), and machine-learning techniques such as the Probabilistic Neural Network model to identify the factors that contribute most to degradation (Tran et al., 2006).

Studies using time-dependent reliability analysis provide further insights into how degradation can be probabilistically modelled at a more granular scale. For example, Stewart and Rosowsky (1998) used MCS to calculate the probability of failure in yearly increments over the lifetime of a reinforced concrete bridge. Furthermore, time-dependent reliability analysis can be integrated with Life-Cycle Cost (LCC) estimation, where the LCC component is used to calculate the total expected cost of asset ownership. Combined, these methods can help determine the most cost-effective maintenance strategy over the life of an asset, while managing risk through a threshold based on a reliability index (Sajedi and Huang, 2019).

Collectively, these studies demonstrate that probabilistic models provide a methodologically robust approach for addressing uncertainty in asset degradation. However, they also primarily focus on individual high-value infrastructure assets, such as bridges, rather than the more common and simpler asset typically managed by organisations (e.g., lights, carpet, tables). Moreover, some probabilistic methods are not easily transferable to different asset types. For example, a probabilistic model developed for corrosion in reinforced concrete cannot be applied to assets made from different materials.

As a solution, this paper proposes an approach that incorporates aspects of the time-dependent life-cycle cost analysis but can be applied to a wider range of assets. Specifically, the likelihood of assets being at EOL is calculated using a Cumulative Distribution Function (CDF) based on the condition ratings at which assets were replaced. This likelihood is then used to calculate the Expected Value (EV) of each asset. A risk threshold is also applied to determine which assets will be considered for replacement. Lastly, through slight modifications to the cost calculation method, where the EV is substituted by the full replacement value of the asset, different scenarios are explored to address varying levels of risk, uncertainty, and financial requirements.

2 METHOD

2.1 Data

To represent a hypothetical asset portfolio, 810 observations were randomly sampled (30 per asset across 27 distinct assets) from a historical dataset obtained through a cyclical condition assessment program conducted by an Australian-based asset management solutions company. The dataset includes a variety of assets from buildings located in rural and urban environments within New South Wales. To calculate the cost associated with the likelihood of assets being at EOL, the condition at replacement, and replacement cost, grouped by each individual asset, were extracted. The condition values were recorded on a continuous scale from 1.00 (As New) to 5.00 (EOL). However, there was no upper limit, as the values increased incrementally in real time while stored in the database after initial capture and before replacement occurred. Note that although the sampled data originates from a real dataset, it is effectively similar to simulated data given that it represents less than ten percent of the full dataset.

2.2 Procedure

To derive the likelihood (P_f) of assets being in an EOL state, the Weibull CDF was used for its reliability and versatility in modelling a wide range of assets (Abernethy, 2006). The equation used is shown below:

$$P_f(x) = 1 - e^{-\left(\frac{x}{\lambda}\right)^k} \qquad \textit{Equation 1}$$

where x represents the condition of the asset at replacement, k represents the shape parameter, and λ represents the scale parameter. The distribution parameters were derived from the mean (μ) and

standard deviation (σ) for each distinct asset based on the condition-at-replacement values using the method of moments. Additionally, *k* was approximated using the coefficient of variation method (Frolov et al., 2024), and λ was estimated employing the Stirling's approximation to enable efficient computational processing, as shown below:

$$k = \left(\frac{\sigma}{\mu}\right)^{-1.0}$$ *Equation 2*

$$\lambda = \frac{\mu}{\Gamma(1+\frac{1}{k})}$$ *Equation 3*

Once the likelihood of asset failure was calculated, a hypothetical risk threshold of 50% was applied to determine which assets to replace. This means that only assets which have a likelihood of failure greater 50% are considered for replacement. The total cost of replacement was then derived from the sum of the EV, where the cost of replacing the asset is multiplied by the risk associated with the asset being EOL. This equation is expressed as:

$$C_{Total} = \sum_{i=1}^{n} (P_i \times C_i) \cdot \mathbb{1}_{\{Pi>X\}}$$ *Equation 4*

where C_{Total} represents the total expected replacement cost for entire portfolio, P_i is the probability of failure, C_i is the cost of replacing asset, and X in $\mathbb{1}_{\{Pi \geq X\}}$ is the risk threshold which equals 1 when P_i exceeds X. This method of calculating the total replacement cost of an asset portfolio aims to efficiently account for the cost associated with the risk of asset failure by pooling funds to address the proportion of assets expected to fail.

As shown in Table 1 below, other cost scenarios were also considered with each varying qualitatively in how risk or uncertainty is addressed as a trade-off against cost. *The Expected Replacement Cost without a Risk Threshold* scenario calculates the cost required to address all levels of risk of asset failure (i.e., > 0%) across the portfolio. In contrast, the *Full Replacement Cost with a Risk Threshold* scenario calculates the total cost of replacing all the items that have a probability of failure greater than the applied risk threshold. The latter scenario removes the uncertainty around whether the predicted amount will be sufficient to replace all assets exceeding the specified risk threshold. For instance, if a portfolio contains large-ticket items that are few in number, but all require replacement, the allocated funds set aside to address the risk of failure (i.e., calculated using the EV method) may not be enough to meet the full cost of replacement.

3 RESULTS

As shown in Table 2 below, the *Expected Replacement Cost with a Risk Threshold* scenario was the most economical of all the scenarios. The *Full Replacement Cost with a Risk Threshold* scenario ranked as the second most economical; however, this outcome was largely driven by the applied risk threshold, which excluded approximately 69% of assets from being considered for replacement.

Table 1: Description of the cost scenarios based on different methods of calculating replacement cost with or without a risk threshold

Scenarios	Description
Expected Replacement Cost with a Risk Threshold	The total replacement cost is calculated as the sum of the expected value of all assets that have a greater likelihood of being at EOL than the set threshold.
Expected Replacement Cost without a Risk Threshold	The total replacement cost is calculated as the sum of the expected value of all assets in the portfolio.
Full Replacement Cost with a Risk Threshold	The total replacement cost is calculated as the sum of the full replacement value of all assets that have a greater likelihood of being at EOL than the set threshold.
Full Replacement Cost without a Risk Threshold	The total replacement cost is calculated as the sum of the full replacement value of all assets in the portfolio.

Notably, when the risk threshold is decreased to 25%, the estimated total replacement cost increases by 110% to $786,577.35, making it the second most expensive option. This significant increase is likely due to the shape parameter being high for most assets ($k > 4$), which reflects an accelerated rise in the likelihood of failure as assets approach their "high-risk" phase. This trend is highlighted in Figure 1, where the data for a general purpose outlet shows a non-linear, S-shaped relationship between condition at replacement and the likelihood of an asset being at EOL derived from the Weibull CDF. Additionally, the scale parameter ($\bar{\lambda} =$ =4.97) indicates that 63.2% of assets are expected to have reached their EOL state by the time their condition reaches 4.97.

The acceleration in the risk of failure as assets move beyond the low-risk phase results in the likelihood

Table 2: Total replacement cost and descriptive statistics of the Weibull cumulative distribution function applied across 27 distinct assets under different cost scenarios

Scenarios	*N*	Condition at replacement		Shape parameter		Scale parameter		Likelihood		Total replacement cost
		M	SD	M	Range [Min, Max]	M	Range [Min, Max]	M	SD	
Expected Replacement Cost with a Risk Threshold (50%)	248	5.82	0.96	4.55	[2.79, 8.04]	5.16	[3.52, 6.91]	0.77	0.15	$281,325.94
Expected Replacement Cost without a Risk Threshold	810	4.24	1.33	4.65	[2.79, 8.04]	4.97	[3.52, 6.91]	0.40	0.27	$521,392,01
Full Replacement Cost with a Risk Threshold (50%)	248	5.82	0.96	4.55	[2.79, 8.04]	5.16	[3.52, 6.91]	0.77	0.15	$373,851.39
Full Replacement Cost without a Risk Threshold	810	4.24	1.33	4.65	[2.79, 8.04]	4.97	[3.52, 6.91]	0.40	0.27	$1,347,660.62

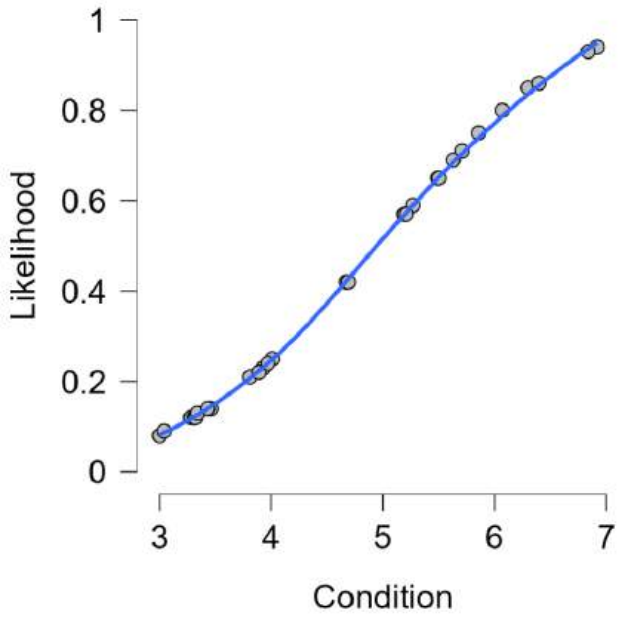


Figure 1: Relationship between asset condition at replacement and the likelihood of an asset being at end-of-life derived from the Weibull cumulative distribution function for a general purpose outlet

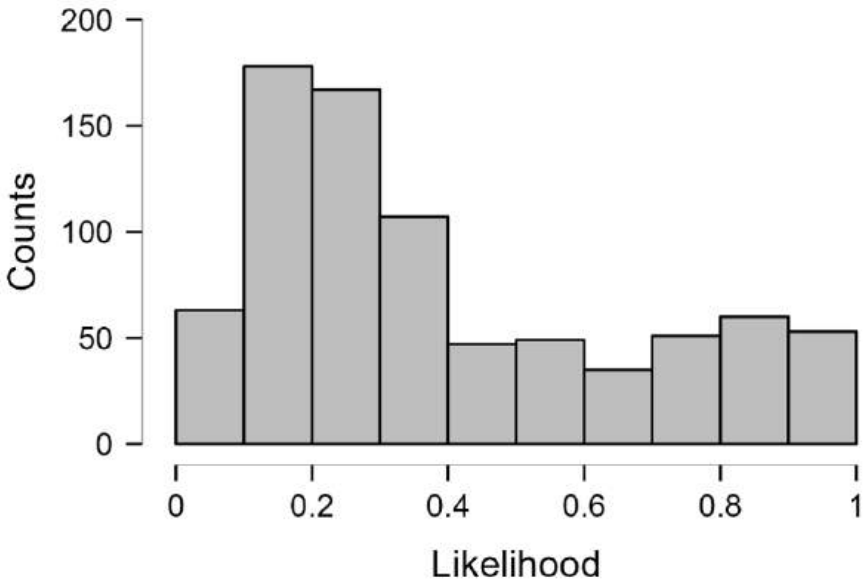


Figure 2: Relationship between asset condition at replacement and the likelihood of an asset being at end-of-life derived from the Weibull cumulative distribution function for a general purpose outlet

distribution being bimodal as shown in Figure 2. The two peaks indicate that the data cluster around two distinct values, which often reflects the presence of two underlying processes. In this case, assets tend to fall into either a low-risk or high-risk states, with relatively fewer in the transitional phase. This pattern is typically observed when the timing of asset replacement is highly predictable, with low variance in the condition values at which assets are replaced.

4 DISCUSSION

Predicting the total replacement cost of a hypothetical asset portfolio based on a probabilistic degradation model was demonstrated by applying the Weibull CDF to estimate the likelihood of assets being in an EOL state and using EV to quantify the cost associated with this risk. The variations in costs simulated across the different scenarios highlights how the model can be tailored to address different levels of risk tolerance or uncertainty in asset replacement. Specifically, calculating total replacement cost using EV provides the most economical outcome, as the cost is directly proportional to the number of assets expected to fail. In contrast, using the full replacement value is more expensive but removes the uncertainty on not knowing which of the assets will fail. The risk threshold offers an additional mechanism to balance cost and risk, as increasing the threshold reduces total replacement cost, while decreasing it raises the cost.

However, it is important to note that the costs generated were based on a random sample used to represent a hypothetical asset portfolio. Because the data come from a subset of items replaced as part of cyclical assessment program, the overall condition profile may be similar to that of a portfolio in which most assets have already deteriorated. Although the differences in costs between scenarios are not expected to change, a more typical asset portfolio, with a less skewed distribution of condition values, would generally exhibit a lower overall likelihood of failure. As a result, the total replacement cost is expected to be attributed to a smaller proportion of assets than reported in this paper.

Another point to consider is that, unlike much of the existing research employing time dependent life cycle cost analysis, this study did not incorporate maintenance or repair activities. This is because many of the assets examined (e.g., lights, taps, fire extinguishers) were relatively low-value, such that the cost of repair would exceed the cost of replacement. Moreover, organisational requirements for maintenance are far more stringent than those for replacement and can incur significant penalties if statutory obligations are not met. Therefore, using AMS to support routine maintenance was not considered as critical an issue as managing asset replacement.

Finally, for probabilistic degradation modelling to be effectively used by organisations for capital planning, the ability to forecast future replacement costs on an annual basis is essential. One way to achieve this is to estimate future condition ratings in yearly increments based on the correlation between condition and time. These projected condition ratings can then be used to calculate the future likelihood of failure using the proposed Weibull CDF. Furthermore, depending on the method used to calculate cost (i.e., EV vs full replacement value), cost forecasts that simulate preventative maintenance (i.e., replacing assets before failure occurs) can also be generated. For scenarios that rely on full replacement cost, the risk threshold can be used to reset the asset's condition to zero to reflect its replacement. For scenarios that use EV, the yearly replacement cost can be used to determine how many assets can be replaced, with each replaced asset having its condition reset to zero and priority given to those with the highest likelihood of failure.

5 CONCLUSION

One of the most common maintenance strategies employed by organisations is to follow OEM guidelines, supported by AMS, to predetermine when assets should be maintained or replaced. However, because asset behaviour is inherently uncertain, a probabilistic approach to modelling degradation using the Weibull CDF was proposed. Specifically, cost scenarios were generated using a risk threshold along with two cost calculation methods: EV derived from the likelihood converted from condition-at-replacement values through the Weibull CDF, or full asset replacement value. Together, these scenarios offered multiple ways to estimate the total replacement cost of a hypothetical asset portfolio under varying levels of risk, uncertainty, and financial outcomes. A method for forecasting future total replacement costs on a yearly basis using the proposed probabilistic degradation model was also outlined, providing a basis for future investigation.

6 REFERENCES

[1] Abernethy, R.B. (2006) *The New Weibull Handbook.* 5th edn. North Palm Beach, FL: Robert B. Abernethy.

[2] Frolov, M., Tanchenko, S. and Ohluzdina, L., 2024. *Parameter estimation of the Weibull distribution in modeling the reliability of technical objects.* Journal of Engineering Sciences, 11(1), pp. A1–A10. Available at: https://doi.org/10.21272/jes.2024.11(1).a1

[3] Ma, Y., Zhang, J., Wang, L. and Liu, Y. (2013) *Probabilistic prediction with Bayesian updating for strength degradation of RC bridge beams.* Structural Safety, 44, pp.102–109. Available at: https://doi.org/10.1016/j.strusafe.2013.07.006

[4] Okirie, A.J., Barnabas, M., Ejomarie, E., Alhassan, J.A. and Yusuf, B.A., 2024. *Understanding deviations from original equipment manufacturers' maintenance recommendations: Reasons, barriers, and benefits.* Journal of Engineering and Applied Science, 71, Article 219. Available at: https://jeas.springeropen.com/articles/10.1186/s44147-024-00563-y

[5] Sajedi, S. and Huang, Q. (2019) *Reliability-based life-cycle-cost comparison of different corrosion management strategies,* Engineering Structures, 186, pp. 52–63. Available at: https://doi.org/10.1016/j.engstruct.2019.02.018

[6] Stewart, M. G. and Rosowsky, D. V. (1998) *Time-dependent reliability of deteriorating reinforced concrete bridge decks,* Structural Safety, 20(1), pp. 91–109. Available at: https://doi.org/10.1016/S0167-4730(97)00021-0

[7] Sudret, B. (2008) *Probabilistic models for the extent of damage in degrading reinforced concrete structures.* Reliability Engineering & System Safety, 93(3), pp. 410–422. Available at: https://doi.org/10.1016/j.ress.2006.12.019

[8] Tran, D.H., Ng, A.W.M., Perera, B.J.C., Burn, S. and Davis, P. (2006) Application of probabilistic neural networks in mo*delling structural deterioration of stormwater pipes.* Urban Water Journal, 3(3), pp.175–184. Available at: https://doi.org/10.1080/15730620600961684

[9] Val, D. V., Andrade, C., Sykora, M., Stewart, M. G., Bastidas-Arteaga, E., Mlcoch, J., Truong, Q. C. and El Soueidy, C. P. (2025) *Probabilistic modelling of deterioration of reinforced concrete structures,* Structural Safety, 113, Article 102454. Available at: https://doi.org/10.1016/j.strusafe.2024.102454

STAR PROFILE – Shishir Tripathi CSAM



Why asset management?

To me, asset management is simply the most logical way of running an organisation. At its core, it is about getting the most value out of assets in the safest and most reliable manner - which is ultimately the reason any asset exists in the first place. What I find compelling is that it's not just a technical discipline. It's a system of governance that brings together stakeholders across all levels of an organisation. In that sense, it's similar to functions like accounting - it applies everywhere, across industries, and influences decisions from the ground level through to executive leadership. Over time, I've come to see asset management less as a function, and more as a way of thinking and operating.

What is your speciality?

My speciality is at the intersection of engineering, rolling stock, and asset management, with a current focus on operator-side asset assurance and engineering governance. My background spans mechanical engineering, reliability engineering, and asset management, primarily within rolling stock environments. That progression has allowed me to approach asset management not just from a framework perspective, but from a practical engineering lens - understanding how assets behave, how decisions are made, and how they play out in operations.

What drew you to explore this particular speciality?

This area sits at the intersection of engineering and decision-making, which is what drew me to it. It requires not only sound technical understanding, but also an appreciation of asset management principles and how they translate into real-world outcomes. I've always been interested in how decisions are actually made - how engineering intent translates into operational reality. This space allows me to contribute not just to technical solutions, but to how those solutions are governed, assured, and sustained over time.

What makes a great asset manager?

A great asset manager understands that asset management is not a one-person/one-team job. While asset managers play a key role in analysing problems and guiding decisions, the real value comes from recognising that asset management is a system - one that spans multiple functions, disciplines, and stakeholders across an organisation. It's not just about producing plans or frameworks, but about embedding a way of thinking that influences how decisions are made at every level.

What is the most exciting trend you've noticed in asset management today?

One of the most encouraging trends I've noticed is the shift in how asset management is being perceived. More and more, it's no longer seen as a compliance exercise or a set of documents that need to be produced. Instead, there is a growing recognition that asset management can actively drive decisions, influence outcomes, and deliver real value. Seeing asset management move from "paper-based obligation" to something that genuinely informs action and strategy is a very positive development.

What is the biggest challenge facing up-and-coming asset managers today?

I would say there are two equally critical challenges. The first is gaining buy-in, particularly from senior leadership. While the value of asset management is increasingly recognised, embedding it into day-to-day decision-making still requires strong alignment and support at the top. The second is data. Asset management relies heavily on data quality and availability, and when data is incomplete or uncertain, it introduces assumptions and risk into decisions. Navigating that balance remains an ongoing challenge for many asset managers.

What advice would you give to an up-and-coming asset manager today?

I would say spend more time with the problem. There's often a tendency to move quickly to solutions without fully understanding what the problem actually is. Taking the time to ask the right questions - whether the problem is clearly defined, whether the data is reliable, what risks are involved, and what happens if it's not addressed - almost always leads to better outcomes. In my experience, the quality of the solution is directly linked to how well the problem is understood.

When you're not busy at work, what do you enjoying doing to unwind/relax/explore?

When I'm not at work, I try to keep things fairly simple. I spend time with my wife and our dog, go for long walks, and make the most of being close to the beach. I also enjoy going for long drives and exploring new places - it's a good way to switch off and clear my head. Alongside that, I enjoy reading and occasionally just slowing things down a bit, which helps create a good balance outside of a fairly structured and demanding work environment.



STAR PROFILE – Rihannah Hesami CSAM



Why asset management?

At its core, asset management is about optimisation — and that is what drew me in. It sits at a rare intersection where engineering rigour meets strategic thinking and long-term decision-making. It is not simply about understanding how assets work; it is about understanding how they create and sustain value over time, balancing cost, risk, performance, and service outcomes across the full lifecycle.

That combination of technical depth and big-picture thinking is what keeps me engaged. Every decision has consequences that play out across decades, and I find that genuinely compelling.

How long have you been working in the asset management sector?

More than 15 years — though for the early part of my career, I did not even know the work I was doing had a name. I was deep in asset analytics, lifecycle planning, and technical decision support before I ever connected it to the formal discipline of asset management. That realisation changed how I saw everything I had already been doing.

What is your speciality?

The intersection of engineering, asset management, and asset intelligence. I am most interested in how technical evidence, data, and increasingly AI can translate into better long-term infrastructure decisions — not just better monitoring, but genuinely better thinking at an organisational level.

What drew you to explore more about this particular speciality?

I am a problem solver by nature, and I wanted to work on bigger problems — not just technical ones. Early in my career, I realised that the most impactful work happened at the boundary between engineering knowledge, data, and organisational decision-making. That is where I could see the gap between what organisations knew and what they actually decided — and I wanted to help close it. That curiosity has never left me.

What's the best career advice you've ever received and who gave it to you?

The advice that has stayed with me most is: do not limit yourself to the boundaries of your current role, be curious, and always ask the right questions. It came through mentors and hard experience rather than one defining moment. What it taught me is that meaningful growth happens when you deliberately build capability beyond what your position formally requires.

For me, that advice meant staying curious beyond my immediate responsibilities — taking on challenges outside my comfort zone, building capability across different disciplines, and not waiting for the perfect conditions to grow. Looking back, that is what allowed me to move across various technical, analytical, and strategic work throughout my career.

What makes a great asset manager?

Breadth of thinking combined with the courage to make decisions under uncertainty. A great asset manager understands risk, cost, lifecycle, performance, and governance — not in isolation, but as a connected system, balancing competing priorities toward the best possible outcome. They translate complexity into clear, practical decisions and bring different parts of an organisation with them.

Technical credibility matters, but so does the ability to see the big picture, ask the right questions, and focus on long-term value rather than short-term fixes.

What is the most exciting trend you've noticed in asset management today?

Without question, it is the integration of AI and advanced analytics into asset management practice. When applied with sound judgement, these tools have the potential to fundamentally shift how organisations understand asset performance, prioritise investment, and manage risk across the lifecycle. What excites me most is not the

technology itself, but what becomes possible when it is combined with strong engineering knowledge and good governance. That said, I am also cautious. Technology vendors often oversell what their tools can do, and organisations frequently buy before they fully understand what they are getting. The opportunity is real — but so is the risk of investing in capability before the culture is ready to use it well.

What is the biggest challenge facing up-and-coming asset managers today?

Developing breadth, not just technical depth. Most emerging professionals enter the field from strong engineering or operational backgrounds — which is valuable — but asset management demands a much wider lens. Lifecycle thinking, risk, cost, governance, and service outcomes all need to work together. On top of that, they are navigating an environment with more data and tools than ever, which creates its own challenge — learning about how to select the right tool for the right task, as well as how to extract genuine insight rather than just producing more output.

What advice would you give to an up-and-coming asset manager today?

Build strong foundations, but do not stay narrow. Seek exposure across engineering, operations, finance, planning, and governance. The professionals who grow fastest are not necessarily the most technically gifted — they are the ones who develop judgement, communicate clearly, and can connect their work to outcomes that matter to the organisation. Stay curious, question assumptions, and learn to communicate. I learned the hard way that the best analysis means nothing if you cannot get it to the right person in the right way. Technical credibility opens the door — but communication is what creates actual change.

What is the biggest challenge facing the asset management sector today?

In my experience across different industries, the common thread is about turning asset management from a recognised framework into a genuinely embedded organisational capability. Many organisations speak the language — lifecycle, risk, value — but there is still a significant gap between policy and practice. In infrastructure, that gap is costly.

The deeper challenge is integrating fragmented data, competing priorities, ageing assets, and rising service expectations into one coherent, evidence-based approach to decision-making. Organisations often have more data than ever but less clarity than they should. Closing that gap — between what is known and what is actually decided — is the work that matters most right now.

What is your proudest career achievement?

One piece of work I return to often is a programme I led connecting field measurement data to asset condition assessment, investment criteria, and long-term planning for a major urban tram network. On the surface it sounds technical, but the real challenge was much broader — translating raw field data into something that operational people understand and trusted, planners could act on, and decision-makers could use with confidence across both annual and long-term horizons.

What made it meaningful was bringing together data acquisition, research, asset management thinking, and — critically — the knowledge of the people actually working in the field every day. No single discipline could have done it alone.

What I am most proud of is not the methodology, but the long-term impact. More broadly, I am proud of building a career that has genuinely bridged engineering depth with strategic asset management and data intelligence. Achieving CSAM is part of that story, not the end of it.

What's next for you?

I want to keep expanding my contribution at the strategic and leadership level — where engineering, asset management governance, and data intelligence converge. I see significant opportunity in helping organisations not just adopt better tools, but develop the capability and culture to use them well — particularly as AI becomes more embedded in infrastructure decision-making.

CSAM is a milestone, but the work I am most interested in is still ahead.

When you're not busy at work, what do you enjoy doing to unwind?

I recharge through time with family and friends, cooking, and interior design. I am also working on a personal project to create digital content across topics I care about, using AI tools — something that started as a creative outlet but has deepened my thinking about how AI can be applied in practical, meaningful ways. I find that stepping into a completely different mode of thinking helps me return to work with a clearer perspective.

STAR PROFILE – Ali Ahmadi CSAM



Why asset management?

I spent the first few years of my professional career in Engineering roles designing concrete and civil structures and carrying out complex investigations. While I enjoyed these roles, I always wanted to understand the bigger picture and how different elements of Asset Management enable organisations achieving their objectives. That is why I picked Asset Management as an area that I wanted to develop in professionally.

How long have you been working in the asset management sector?

8 years

What is your speciality?

My speciality is around designing and implementing context specific Asset Management Systems for organisations as well as infrastructure renewal and growth planning in the water sector.

What's the best career advice you've ever received and who gave it to you?

The best piece of career advice I have received was that "nothing replaces experience." This came from one of the first managers I had, Adam Russo.

What makes a great asset manager?

There are multiple characteristics that make a great Asset Manager. These are the ability to explain complex ideas simply to wider audiences, bringing stakeholders along the journey and the ability to maintain a strategic lens while being able to go into the details where required.

What is the biggest challenge facing up-and-coming asset managers today?

The biggest challenge facing asset managers is that the concept of Asset Management is relatively new and many organisations are early on their maturity journey. As such applying Asset Management principles requires more than just understanding it theoretically but having the ability to tailor it such that it suits the context of the organisation and having the ability to take the organisation on the journey.

What advice would you give to an up-and-coming asset manager today?

I have two pieces of advice. With artificial intelligence increasingly present in our work, critical thinking has become more crucial than ever. My first piece of advice is to protect and sharpen your ability to think critically and do not outsource your judgment to a machine.

My second piece of advice is to invest deeply in your communication skills, encompassing social and emotional awareness, public speaking, and influence. While AI can write for us, it cannot speak for us. The ability to genuinely connect with people and move them to action will become one of the most valuable skills any professional can possess

What is your proudest career achievement?

I have many, however one standout is the work we have been doing at GMB Water. I am responsible for implementing our Asset Management System and progressively increasing our AM maturity, and through strong collaboration across the team, we have made significant strides in both.

When you're not busy at work, what do you enjoying doing to unwind/relax/explore?

Outside of my Asset Management work, I am a professional TEDx and keynote speaker and currently working through my first book. Similar to my Asset Management work, I enjoy the work it takes to communicate ideas in way that is simple and memorable and drives people to action.

The Asset Management

CAREER PLANNING TOOL

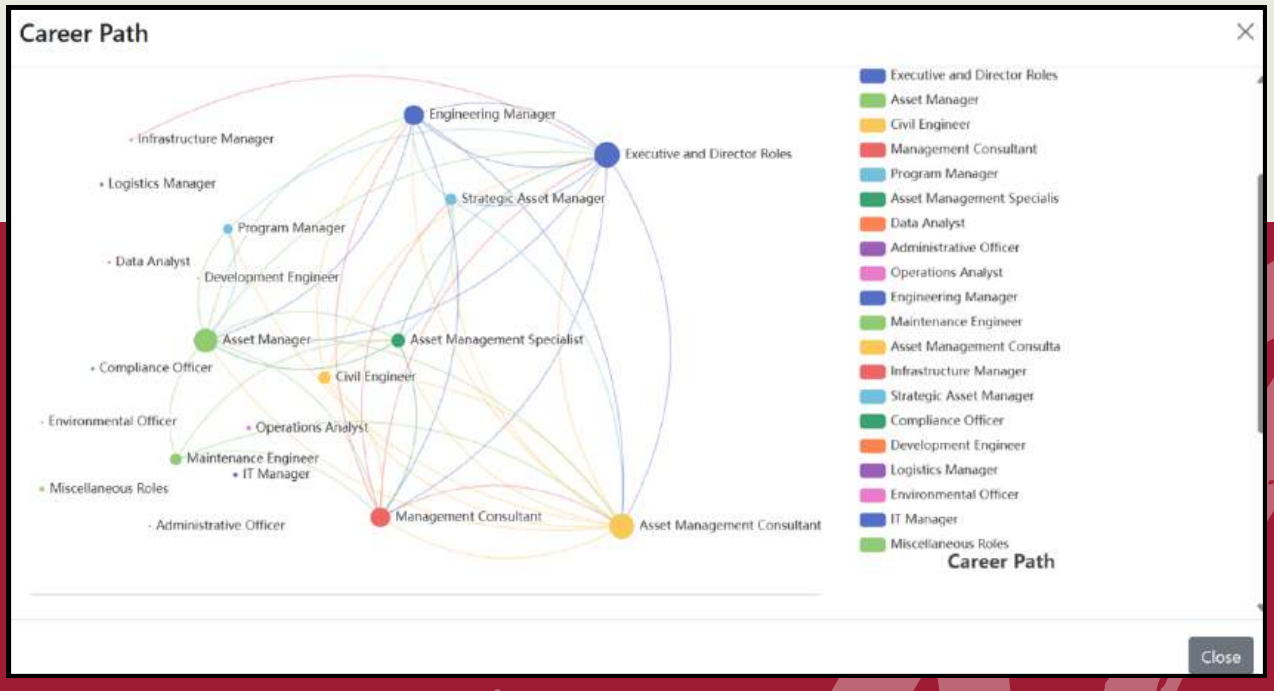
- **Wondering what asset management competencies you still need in your asset management journey?**
- **Looking for a Personalised Asset Management Career Mapping Tool?**
- **Wanting an asset management competency profile?**

Use the online tool to Create your Asset Management Individualised profile. Use your self assessment to map your future career and work on any gaps in asset management competencies.

Available to everyone from the AM Council website. Just login to the AM Council website and navigate to certification pages and AM GFMAM landscape or use this shortcut link to get started: www.amcouncil.com.au/career

There is no particular defined career path in Asset Management, but more a personal journey in collecting Asset Management competencies together with other competencies in your particular chosen field that allows one to get to the position to which you aspire.

Use the career path tool to compare your asset management competencies against those most common for various job roles and plan your future today!



CHAPTER **NEWS**

ADELAIDE

The Adelaide asset management community came together in May for a focused technical session that connected national industry insights with local practice. The event provided attendees with key takeaways from AMPEAK 2026 while also introducing the new Adelaide AMCouncil Chapter Committee.

The session opened with a keynote presentation from Graeme Brown of Infrastructure SA, who shared valuable perspectives on South Australia's State Asset Management Policy and Framework and its role in supporting long-term public value and informed decision-making.

Local practitioners Michael Sinclair and Gerry Mavrogiannis then reflected on their experiences at AMPEAK 2026, highlighting emerging trends, industry challenges, and practical lessons that can be applied within South Australian organisations.

The technical program concluded with two presentations from Adelaide Chapter Chair Phil Kowalczyk. The first explored asset management approaches in complex operating environments, while the second examined lifecycle data management practices aligned with ISO 19650 and ISO 55013, reinforcing the importance of quality information throughout the asset lifecycle.



Thank you to all presenters, attendees, and committee members who contributed to a successful event. We look forward to continuing to strengthen connections across Adelaide's asset management community and sharing knowledge that supports better asset outcomes across the region.

MELBOURNE

The Melbourne Chapter welcomed members and guests to a May lunchtime technical event in the Melbourne CBD, focused on Managing Ageing Infrastructure at Scale: A Local Government Perspective.

The session was presented by Tze-Sian Hor, Coordinator Asset Management & Systems, Environment & Infrastructure at the City of Stonnington, who shared practical insights from local government asset management, including enterprise-wide business process mapping, AMIS requirements development, governance uplift, data standards, asset hierarchy considerations, and the importance of aligning asset processes with finance, works management and broader corporate systems.

Attendees heard valuable lessons on the realities of managing large volumes of ageing physical and digital assets, including system limitations, change management challenges and the importance of cross-department alignment.

Thank you to Tze-Sian for generously sharing his experience and insights, and to Melbourne Chapter Chair Andrew Sarah and the Melbourne Chapter Committee for organising another valuable technical event for the asset management community.





CHAPTER **NEWS**

NEW ZEALAND

June's New Zealand Chapter event in Christchurch brought together asset management professionals for an afternoon of discussion, learning, and firsthand industry experience.

The technical program explored how organisations are building stronger asset management capability in response to changing operational, regulatory, and investment demands. Presenters shared practical insights on capability development, transformation initiatives, rolling stock investment strategies, and the journey towards ISO 55001 implementation.

The conversations highlighted the importance of aligning people, processes, and systems to deliver sustainable asset performance and long-term value.

For attendees joining in person, the event concluded with a rare opportunity to step inside KiwiRail's new rolling stock and maintenance facilities. The guided tour offered a close-up view of the infrastructure and operations supporting the future of rail in New Zealand, providing valuable context to many of the themes discussed throughout the afternoon.

The New Zealand Chapter thanks everyone who attended and contributed to the event, helping to foster knowledge exchange and strengthen connections across the asset management community.

PERTH

The May Women Creating Value event represents a continued step-change in the quality and ambition of the AM Council Perth Chapter's programming. Over recent months, the Chapter has expanded its focus from technical asset management topics to include the strategic dimensions of leadership, governance, and organisational capability — reflecting the evolving needs of asset management professionals across Western Australia's resource and infrastructure sectors.



CHAPTER **NEWS**

The key themes emerging from the presentations reflect closely the status of women in Industry, specifically related to Aset Management:

- AI as a governance and leadership issue, not merely a technology decision — with 95% of generative AI implementations showing no P&L impact due to absent frameworks, designated ownership, and human-centred controls.
- Diverse teams as risk intelligence, not just social good — a direct contribution to asset management governance thinking.
- Credibility and trust as prerequisites for effective change leadership in high-consequence asset environments.
- Alignment as the real sustainability challenge — four teams, one asset, four perspectives requiring deliberate convergence.
- Agentic AI as an immediate frontier requiring organisational protocols, not a future consideration.

The Perth Chapter and organisers of the event received positive feedback about the quality of contributions and the format of the "Value Minute". The quality of speakers, breadth of industry representation, and active audience engagement position the Chapter as a recognised platform for leadership and strategic discourse in Western Australian asset management











ASSET MANAGEMENT COUNCIL

ASSET MANAGEMENT FUNDAMENTALS COURSE

Run online in a live interactive session over two consecutive mornings offers a thorough overview of asset management. The purpose of the Asset Management Fundamentals program is to understand the fundamentals of asset management and how they can provide benefits to the organisation. Please click here to register.
Course learning objectives:

- Define asset management within the context of ISO 5500x suite, the Asset Management Landscape and AMBoK
- Identify the principles that underpin asset management
- Identify available asset management tools and techniques that are applicable within organisational context
- Identify opportunities to apply these learnings to improve individual and organisational performance.

JULY 13 & 14

AUGUST 13 & 14

SEPTEMBER 14 & 15

OCTOBER 12 & 13

NOVEMBER 12 & 13

DECEMBER 7 & 8

Register at www.amcouncil.com.au

How Much KNOWLEDGE & EXPERIENCE Protects Your Assets?



CFAM
Certified Fellow in Asset Management

CSAM
Certified Senior Practitioner in Asset Management

CPAM
Certified Practitioner in Asset Management

CAAM
Certified Associate in Asset Management

ARE YOUR EMPLOYEES ASSET MANAGEMENT COUNCIL CERTIFIED?

Improve Retention ▪ Enhance Recruitment ▪ Increase Portability

Get the facts: www.amcouncil.com.au/certification

Membership Application

ASSET MANAGEMENT COUNCIL

ASSET MANAGEMENT COUNCIL LTD

A Technical Society of Engineers Australia

ABN: 15 141 532 747 www.amcouncil.com.au
Phone: +613 9819 2515 Email: accounts@amcouncil.com.au

Thank you for joining the Asset Management Council. Please complete all sections. Phone or email with any queries.

PERSONAL DETAILS (Please print in BLOCK CAPITALS)

Title (Please circle) Dr Mr Mrs Ms Miss Other (Please specify) Sex (Please circle) F M

Family Name Given Names (in full)

Date of Birth Engineers Australia Membership No

CONTACT DETAILS (Please print in BLOCK CAPITALS)

Preferred Address: ☐ Private Address or ☐ Business Address

Position

Organisation

Postal Address

City State

Country Postcode

Phone Fax

Mobile

E-mail

AREAS OF INTEREST (Please tick)

Technical Topics

- ☐ Reliability
- ☐ Availability
- ☐ Maintainability
- ☐ Performance
- ☐ Spares Planning
- ☐ Maintenance Planning and Scheduling
- ☐ Maintenance Plan development and implementation
- ☐ Maintenance Policy/Strategy development
- ☐ Logistics
- ☐ Shutdown planning and the maintenance interface
- ☐ Asset Management
- ☐ Other:

Issues

- ☐ Skills development
- ☐ Training
- ☐ Other:

Industries

- ☐ Facility Management
- ☐ Consulting
- ☐ Power
- ☐ Transport
- ☐ Defence
- ☐ Oil and Gas
- ☐ Mining and Industry
- ☐ Water and Utilities
- ☐ Infrastructure
- ☐ Other:

Return completed Membership Application with payment to:
Asset Management Council
PO Box 2004, Oakleigh Vic 3166

GROUP AFFILIATION

☐ Young Asset Management Practitioners (18-35 year olds)

CHAPTER AFFILIATION (Please tick one)

☐ Newcastle ☐ Canberra ☐ Sydney ☐ Illawarra ☐ Mackay
☐ Melbourne ☐ Adelaide ☐ Brisbane ☐ Hobart
☐ Darwin ☐ Overseas ☐ Gippsland ☐ Perth

MEMBERSHIP FEES Effective Jan 2015 (Please tick one membership type only)

Individual Annual Fee (including GST)

☐ Member $190.00
☐ Student $35.00

Corporate Band Annual Fee (including GST)

☐ Unlimited $11,500.00 ☐ 11-20 $3,950.00
☐ 6-10 $2,100.00 ☐ 0-5 $1,100.00

GST (10%) does not apply to overseas memberships.

CORPORATE MEMBER NOMINEES

Gold – 10 nominees, Silver – 10 nominees, Bronze – 5 nominees

	Name	Email	Date of Birth (Mandatory)	AM Council Chapter
1				
2				
3				
4				
5				
6				
7				
8				
9				
10				

Contact Asset Management Council to provide more corporate nominee details.

PAYMENT

Method of Payment (please tick one and enclose payment)

☐ Cash

☐ Money Order or Cheque drawn in AUD from an Australian bank) payable to **Asset Management Council Ltd**

☐ International Money Order

☐ Credit Card
(Australian or New Zealand Bankcard only acceptable)

Credit Card Details Please charge my card (tick one card type)

☐ Visa ☐ Bankcard ☐ Mastercard
☐ Diners ☐ American Express

Card no _____ CCV no _____

Expiry _____ Amount $ _____

Name on card _____

Signature _____ Date _____

Corporate Partners and Corporate Members

INDUSTRY PARTNER

ASC
Newmont
Rio Tinto

SUPPORTING PARTNER

Anglo American T&S Services
Ausgrid
AusNet
BAE Systems
BGIS
City of Gold Coast
Department of Transport and Planning (VIC)
Service Stream
Transdev
Transport for NSW
Transpower NZ
Ventia

GOLD

Airservices Australia
Alstom Group
Aurecon
Babcock International Group
Department of Families, Fairness and Housing
Department Of Health
Department of Housing and Works
Energy Queensland
Essential Energy
Evoenergy
Flinders Port Holdings
GHD Advisory
Infrastructure NSW
Jacobs Group Australia
KiwiRail
KPMG
Lucid Consulting Australia
Metro Trains Melbourne
Pacific National
Parks Victoria
Port of Newcastle
QBuild
Spotless Facility Services
Stanwell Corporation
Sydney Water Corporation
TAFE Infrastructure NSW
Territory Generation
Transgrid
Transurban
V/Line Corporation
Western Power
Wood Plc (Australia)
WSP Australia
Xenco

SILVER

ACEREZ M&L
AECOM Australia
AMCL
APA Group
Aurizon Network
BetterAIM
Campbelltown City Council
Comfort Delgro NSW
Court Services Victoria
Defence Equipment Management Organisation (NZDF)
Downer Rail & Transit Systems (RTS)
Evolve FM
Fire and Rescue NSW
GHD
Gladstone Ports Corporation
Greater Western Water
IBM
Infrastructure Advisory Group
ISS Facility Services
Jemena
John Holland Group
KordaMentha
Mackay Regional Council
Napier Port
Norship
NSW Public Works
NSW Telco Authority
Origin Energy
Port Botany Operations
Prometheus Group
Public Transport Authority
SEQWATER
Sodexo Australia
Southern Ports
Stantec Australia
SunWater
Sutherland Shire Council
Sydney Opera House
Tasports
TasWater
Townsville City Council
Unitywater
VicTrack
Water Corporation
Western Australia Police Force

BRONZE

Altrac Light Rail
ANSTO
APP Corporation
Asset Management Associates
Assetivity
Austal Australia
Balance Agri-Nutrients
Beca
Beyond Asset Management
Brisbane Airport Corporation
Bureau Veritas Australia
BWE Waterwell
CBC Facilities Maintenance
CBS Group
Covaris
Cushman and Wakefield
DAS Consulting
Department of Jobs, Skills, Industry and Regions
Department of Planning, Industry & Environment
Dyno Nobel
Egis
Fire Rescue Victoria
Fraser Coast Regional Council
Fremantle Ports
Fresh Hope Communities
Gladstone Area Water Board
Gladstone Regional Council
Goulburn Valley Water
Hardcat
Horizon Power
Hunter Water Corporation
Infrastructure SA
Infrastructure WA
Innovative Thinking IT
Invercargill City Council
IPS-ENERGY AUSTRALIA PACIFIC
Ipswich City Council
Kepa Software
LYB Operations & Maintenance
Melbourne & Olympic Parks
Melbourne Water
Metrix Assets
Metro Trains Sydney
Minset
Murray Irrigation
North East Water
Northern Territory Government Dept. of Infrastructure, Planning & Logistics
Northrop Grumman Integrated Defence Services
NRG Gladstone Operating Services
Ontoit Global
Placemaking NSW
Port of Melbourne Unit Trust
Powerlink Queensland
Quarterbac
Qubist Pty
Queensland Rail
Raytheon Australia
Reeves Group Services
Rockfield Technologies Australia
SA Water Corporation
Saab Australia
Shoalhaven Water
Silver Edge Technologies
SPM Assets
STRUCTURED CHANGE
SUSTECH ENGINEERING
Taronga Conservation Society Australia
TATWEER MIDDLE EAST AND AFRICA L.L.C.
Terotek (NZ)
The Asset Management College
The Exchange Collective
United Energy Services
Verbrec
Verosoft
Wagners Composite Fibre Technologies (CFT)
Wannon Water
WaterNSW
Wesfarmers Chemical Energy and Fertilisers
Worley
Worley Power Service



AMPEAK 2026 Supplement



AMPEAK26 WRAP UP

The AMPEAK conference, held at the Star Perth during April, was a great success! Attendees appreciated the opportunity to come together on site to share knowledge and best practices, create new contacts, and network. The event featured a diverse line up of speakers and numerous award submissions, making it a valuable experience for all involved.

Over three days, delegates got to experience and enjoy:

- **over 100 different sessions**
- **3 varied and interesting key note speakers**
- **interactive workshop sessions**
- **international delegates and speakers from Australia and across the globe**
- **professional development opportunities**
- **poster presentations from universities around Australia**
- **fun and energising social events**
- **many opportunities to network with each other and with sponsors and exhibitors**

Conference Themes

Drawing on the collective wisdom of over 100 speakers, the 2026 conference considered common asset management themes that topics across the theme of Powering Productivity and Prosperity:

- Workforce Transformation and Knowledge Retention
- Data-Driven Decision Making and Digital Innovation
- Strategic Alignment and Value-Focused Asset Management
- Operational Excellence, Reliability and Safety
- Sustainable Asset Management and Societal Value







KEY NOTE SPEAKERS

A line up of impressive keynote speakers kept the audience engaged and provided great insights. Keynote speakers included:



Simon Kuestenmacher

Expert on demographic trends, consumer insights and cultural change in Australia



Mel Thorpe

Strategic Facilitator and Executive Coach: Driving Clarity, Connection & Performance for Leaders and Teams

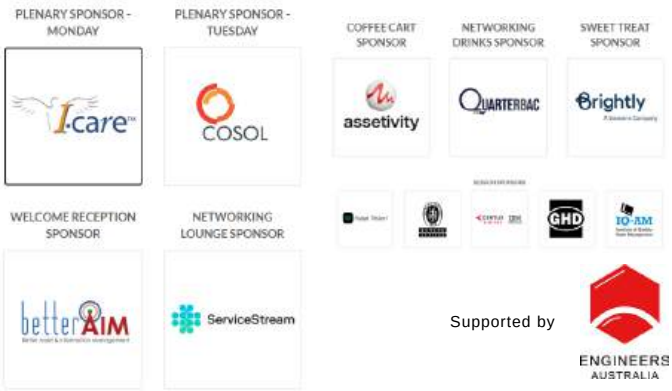


Paul Carter

From Oil Fields to Bestsellers: Paul Carter's Journey Through Chaos, Danger and Laughter

THANKS TO SPONSORS

AMPEAK was also incredibly successful because of the sponsors in particular:







All exhibitors at AMPEAK added an important element to the conference, allowing delegates to network, make new contacts and learn about key companies servicing the asset management field.

Everything about AMPEAK was a reflection of the vision and passion of the AM Council and the commitment to work with and support professionals from across all industries involved in asset management.

The AM Council was delighted with AMPEAK26 and we look forward to seeing you all at next year's conference.

SOCIAL FUNCTIONS

AMPEAK included several social events where delegates got to relax and network. It began with a Welcome Function on the Sunday night in at The Camfield for drinks and canapes to kick of AMPEAK26, where guests were treated to some retro diner entertainment on roller skates.

This was followed by networking drinks on the Monday night and the Gala Dinner Awards Night on the Tuesday evening.

And the winners are...
This year, excellence awards were given in the following categories, recognising excellence in the management of assets and the use of best practice. Congratulations to all our winners.

- **Innovation** - Road Efficiency Group Te Ringa Maimoa - Consistent Condition Data Collection Project
- **Information Management** - Monash University - Asset Data Capture Project
- **Cost/Risk/Performance** - TasWater - Sewer Regional Master Plans
- **Sustainability** - Water Corporation - Embedding Sustainability into Water Corporation Asset Management Practices
- **Resilience** - Essential Energy - Resilience at Scale: Transforming the Network Through Composite Pole Transition
- **Safety** -Transgrid - Gorilla Robot Delivers Smarter and Safer Solutions
- **Rising Star** - April Jewell
- **MESA Medal** - Ernst Krauss
- **Best Paper** - Lalinda Kurunaratne - Data to Decisions: An integrated maintenance investment model for complex infrastructure networks



AM Council presents awards at AMPEAK Gala Dinner



INNOVATION

Road Efficiency Group Te Ringa Maimoa - Consistent Condition Data Collection Project



INFORMATION MANAGEMENT

Monash University - Asset Data Capture Project



COST/RISK/ PERFORMANCE

TasWater - Sewer Regional Master Plans



SUSTAINABILITY

Water Corporation - Embedding Sustainability into Water Corporation Asset Management Practices

AM Council presents awards at AMPEAK Gala Dinner



RESILIENCE

Essential Energy - Resilience at Scale: Transforming the Network Through Composite Pole Transition



SAFETY

Transgrid - Gorilla Robot Delivers Smarter and Safer Solutions



MESA MEDAL

Ernst Krauss



BEST PAPER

Lalinda Kurunaratne - Data to Decisions: An integrated maintenance investment model for complex infrastructure networks



RISING STAR

April Jewell

AM Council presents awards at AMPEAK Gala Dinner



LIFE MEMBER

David Daines



LIFE MEMBER

Mark Mackenzie



AMPEAK

DAILY DOWNLOAD

ASSET MANAGEMENT COUNCIL

DAY ONE | MONDAY 13 APRIL 2026

Last night's opening of AMPEAK26 was a vibrant start to the conference, proudly sponsored by BetterAIM.

Held at The Camfield in Burswood, the 1950s-themed welcome function brought delegates together for a lively evening of connection and fun. Immediate Past Perth Chapter Chair, Adam Rotapel, welcomed everyone to the conference, setting the tone for the days ahead.

Roaming roller skaters dressed as diner girls added a playful touch, while guests enjoyed local beers on tap, an oyster shucking bar, an ice cream mobile, and a delicious spread of food.

Just a comfortable walk from Crown, it was the perfect first event to kick off AMPEAK26.













POSTS OF THE DAY

@AlieenG: Excited to be attending #AMPEAK2026 in Perth and looking forward to a great few days of insights, discussion and connection with asset management peers across Australia on how asset management can power productivity and prosperity, helping organisations focus investment where it delivers the greatest long-term impact.

@SloanT: Heading west to Perth for #AMPEAK26, my fourth consecutive #AMPEAK conference.

@WarrenA: I'm happy to once again be heading to #AMPEAK Asset Management Council in Pert this week.

@ChrisC: Interested to hear other industries perspectives and application. The program looks rich!

AMPEAK

DAILY DOWNLOAD

ASSET MANAGEMENT COUNCIL

DAY TWO | TUESDAY 14 APRIL 2026

CONFERENCE KICKSTART

Day Two of AMPEAK26 opened with a strong focus on productivity, innovation and the future of the profession, with the morning session proudly sponsored by I-care. Dr Yvonne Power began with AI in Asset Management to Increase Productivity, followed by Pieter Van Camp from I-care, who presented Maintenance Autopilot. The morning keynote from Simon Kuestenmacher then explored global demographics and their impact on the asset management sector, offering delegates valuable insight into the trends shaping the future.







NEWS UPDATE

AMCouncil held its AGM on Day One at 1pm. This year, four candidates were up for election for the board, with one board director position available for nomination or renomination. We are delighted to announce that Sarah Hannah has been re-elected for another term as a board director.

POSTS OF THE DAY

@YossiK: Strong start to AMPEAK 26 by Asset Management Council - Day 1 delivered exactly what the industry needs right now.

@KarlH: Great first day at the Asset Management Council AMPEAK 2026 here in Perth. Looking forward to continuing these discussions over the next couple of days and building on the connections made today through Tuesday and Wednesday.

@CassieK: Attending AMPEAK 2026 has been an incredibly rewarding experience.

@JamesG: A gathering of Australia's largest critical infrastructure and asset owners - and without question one of the most engaging industry events I've been part of.

@GopiC:Thanks to AMC for bringing industries, universities and professional bodies together through asset management for making the world a better and safer place to live.

STUDENT POSTERS AT AMPEAK

Be sure to take some time to explore the student poster displays in the main exhibition foyer. Showcasing a range of research projects from universities across the field of maintenance, reliability engineering and asset management, the posters offer a great insight into emerging ideas and future talent in the profession.



Delegates are encouraged to view the displays and cast their vote for Best Student Poster through the Conference app. The winner will be determined by delegate vote and announced at the Excellence Awards Gala Dinner on Tuesday 14 April.

SESSION SNAPS









SOCIAL SPLASH



AMPEAK

DAILY DOWNLOAD

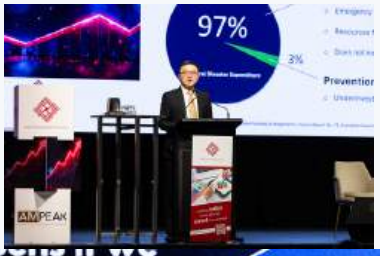
ASSET MANAGEMENT COUNCIL

DAY THREE | WEDNESDAY 15 APRIL 2026

HAPPY WORLD ASSET MANAGEMENT DAY

Day Three began with the plenary session, proudly sponsored by Cosol, and a fitting celebration of World Asset Management Day. Toby Horstead welcomed delegates before Kevin Hang from Sydney Trains shared a research presentation on probabilistic cost–benefit analysis for railway slopes. Following a presentation from Cosol, Mel Thorpe took to the stage with an engaging keynote on what high-performing and productive teams do differently, setting an inspiring tone for the sessions ahead.

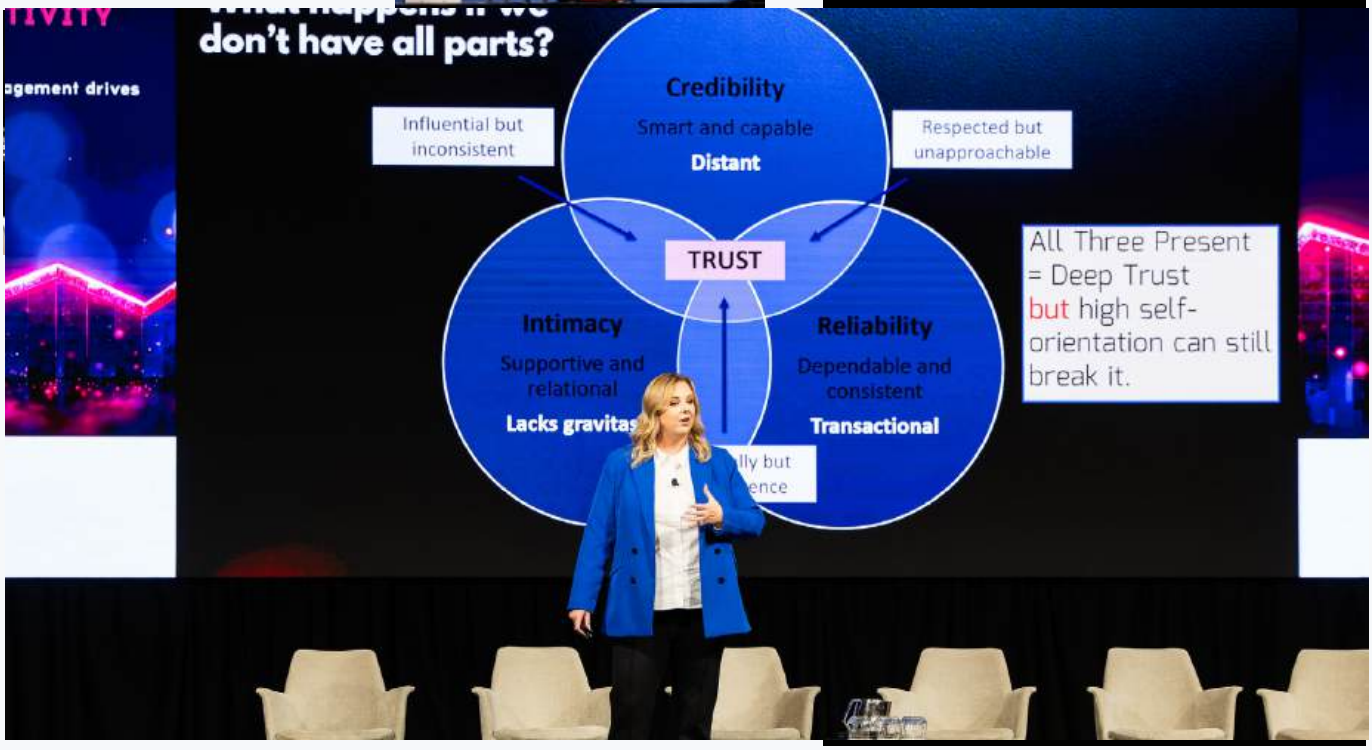




POSTS OF THE DAY

@DaveD: Honoured to join this fine group of AM Professionals by receiving Life Membership of the Asset Management Council. Congratulations to Mark Mackenzie (Life member recipient) and Ernst Krauss CFAM (MESA Medal recipient)

@JiteshP: Great to be back at #AMPEAK 2026 in Perth. It was an amazing first day yesterday, and today is taking over already



2026 EXCELLENCE AWARDS

And the winners are...

Crown Perth set the scene for our much anticipated AMCouncil Excellence Awards last night. Congratulations to winners:

Innovation
Road Efficiency Group Te Ringa Maimoa for the Consistent Condition Data Collection Project

Information Management
Monash University for the Asset Data Capture Project

Cost/Risk/Performance
TasWater for Sewer Regional Master Plans

Sustainability
Water Corporation for Embedding Sustainability into Water Corporation Asset Management Practices

Resilience
Essential Energy for Resilience at Scale: Transforming the Network Through Composite Pole Transition

Safety
Transgrid for Gorilla Robot Delivers Smarter and Safer Solutions

Rising Star
April Jewell

Life Member Awards
Dave Daines
Mark Mackenzie

MESA Medal
Ernst Krauss

Best Paper
Lalinda Karunaratne for Data to Decisions: An integrated maintenance investment model for complex infrastructure networks.



SOCIAL SPLASH

















PICTURES FROM AMPEAK 2026

























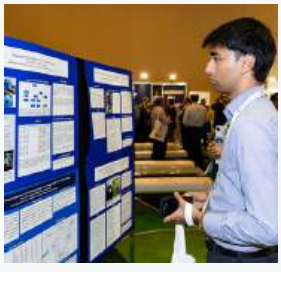












APRIL 2026 PERTH, WA








































ASSET MANAGEMENT COUNCIL

THE ASSET JOURNAL

Asset Management Council
PO Box 2004
Oakleigh Vic 3166
Tel 03 9819 2515
www.amcouncil.com.au