



Responsible Artificial Intelligence

for Major Projects


WHITE PAPER



BY:

ICI INNOVATIONS INC.
18 ARGYLE STREET,
ST. JOHN'S NL CANADA

ST. JOHN'S – HALIFAX



INTRODUCTION

ICI Innovations developed RIVAS to address the information management challenges inherent to major projects. Multidisciplinary documentation and datasets, stakeholder engagement records, spatial data, engineering drawings, and permitting documents accumulate rapidly, particularly during early project phases when teams are small and lack the resources required to manage this volume of information within an increasingly complex digital environment. RIVAS provides a secure environment for managing this information and delivering it in appropriate context. However, as information volume increases, the effectiveness of RIVAS' constituent modules is similarly tested: ReportViz produces fully digital reports linked to underlying data, but navigating a 1,000page multidisciplinary report remains inherently difficult regardless of presentation. Engagement Manager comprehensively manages stakeholder engagements and permitting, but accurately synthesizing the results of 20 meetings with a single stakeholder conducted over several years still requires substantial manual effort. Document Manager enables efficient metadata based search and filtering, but locating answers distributed across multiple documents still requires manual review of each source, and search performance remains contingent on the accuracy and completeness of metadata entered at upload. As project scale and complexity increase, so too does the effort required to ingest and retrieve information, which is why ICI integrated AI as an optional module that when activated offers powerful efficiencies within other RIVAS modules.

LIMITATIONS OF AI APPLIED TO DECENTRALIZED INFORMATION SYSTEMS

Conventional project information management follows a decentralized model, with content distributed across multiple platforms such as SharePoint Online or comparable cloud content management systems usually operating as online file repositories, structured within folder hierarchies that frequently become disorganized and slow retrieval. Stakeholder engagement data is often managed within a separate platform or an expanding collection of Excel files. Geographic information is another substantial data category, typically maintained in yet another platform for analysis, with outputs delivered as static images embedded in reports and disconnected from the source data. The recent proliferation of artificial intelligence (AI) has been positioned as a solution to these structural inefficiencies. In practice, AI is frequently implemented as an additional layer atop this decentralized architecture, rather than as a resolution to its underlying problems. Representative examples include the use of Microsoft CoPilot to process poorly structured document libraries, the upload of documents to commercial AI platforms such as ChatGPT or Claude for summarization, or reliance on the native AI functionality of PDF readers to interrogate multiple documents submitted as part of project development and approvals (e.g., an environmental assessment process).

While such applications of AI may provide localized utility, their net effect on project outcomes can be materially detrimental as explained in these examples:

- CoPilot was employed to summarize stakeholder engagement with a critical group for submission to a regulatory agency. The source documents included one that had not yet undergone review and approval by the project's engagement lead which contained incomplete information. This document had been placed in a shared folder outside of established document control procedures by a team member, making it accessible to CoPilot for processing. The resulting error caused confusion among stakeholders and delayed regulatory engagement pending correction.

- A project team member uploaded draft versions of an environmental assessment to Claude.ai to evaluate whether the document adequately addressed concerns raised by a key stakeholder. The system indicated that it did; however, the team was unaware that the response incorporated information drawn from external online sources referencing the project, a limitation the team did not identify due to insufficient understanding of the system's operating constraints. The response additionally contained a hallucinated assumption, which further reinforced the team's erroneous conclusion. Following release of the assessment, the stakeholder (and regulator) submitted additional information requests, resulting in project delays.

THE RIVAS AI MODULE ARCHITECTURE

RIVAS mitigates these risks through a centralized information repository, ensuring that AI generated outputs are based on validated, authorized content, restricted to users with appropriate access permissions. ICI's design principle for the AI Module is strict: the system may draw exclusively on information stored within RIVAS, with no capacity to incorporate external sources (i.e., general websites from the internet), and the operational scope is defined explicitly by the user. Hallucinated or unsupported outputs are considered unacceptable; consequently, if the available information is insufficient to answer a query, the system returns no answer rather than generating an unsupported response. Where an answer is provided, the underlying content is fully traceable to its source within RIVAS at the sentence level, ensuring verifiable provenance.

Using RIVAS avoids maintaining a document records management system in addition to duplicating content in a separate AI platform. With Document Manager, AI is built-in and available for analyzing any document. Content ingestion requires uploading documents along with descriptive metadata, a process that demands care and can be resource intensive. To minimize this effort, AI provides an alternative to manual entry, as users may scan documents to enable automated metadata extraction, improving both processing efficiency and search reliability. This process will still require review of scanned results to ensure accuracy.

RIVAS operates as an AI powered, enterprise scale search engine built on a proprietary architecture developed and hosted in Canada. It is not a general purpose AI tool, but a purpose built system engineered for secure integration with RIVAS managed information. Its principal technical differentiator is **Intuitive Explainability**, a proprietary referencing framework that generates citations linked directly to source material for every output. Rather than returning an answer alone, RIVAS provides verifiable evidentiary support for that answer. This architecture enables users to directly verify the accuracy and contextual validity of any output, eliminates the "black box" characteristic common to generative AI systems, and structurally precludes hallucination at the point of citation, as references to source material are created non-generatively and cannot be fabricated.

Due to the computational efficiency of the RIVAS AI architecture, associated costs are limited to implementation and annual subscription fees, with no requirement to manage token consumption or absorb unpredictable usage costs. The majority of computational demand occurs at the point of content ingestion, when documents uploaded via Document Manager or Submissions Manager are analyzed by the AI system. Subsequent query processing requires minimal computational overhead, as the system operates on previously analyzed content rather than provisioning new computational resources for each query.

Engagement Manager maintains continuous synchronization between system data and AI, ensuring that all responses reflect the most current information available. All content processed by the AI Module resides within a dedicated repository that is not used for model training or any purpose beyond its intended function; upon deletion of a document or engagement record from RIVAS, the corresponding content is removed from the AI repository as well. This architecture provides RIVAS users with a high degree of confidence that AI generated outputs are accurate, resource efficient, and consistent with responsible information governance standards.

About ICI Innovations

ICI is a 100% Canadian owned software solutions company based in Newfoundland & Labrador with deployments across Canada and the USA in the following markets:

- Renewable Energy
- Oil and Gas
- Mining
- Transportation and Logistics
- Utilities and Infrastructure
- Interactive Data Visualization and Reporting

Since 2014, ICI has implemented custom Interactive Data Visualization and Reporting solutions for government (e.g. Impact Assessment Agency of Canada - Regional Assessment), NGOs (e.g. Maritime Exchange for the Delaware River and Bay) and industry alike (e.g. Oilco - **Oil and Gas Corporation of Newfoundland and Labrador, Atlas Salt Inc, TriplePoint Resources**). ICI has developed tools, procedures and custom user interfaces leveraging existing technology and platforms (e.g., Microsoft 365) to serve our customers better. RIVAS has emerged as ICI's evolved solution for major project information management, providing software, implementation, and professional support services as a complete solution. RIVAS is purpose built to be implemented quickly, flexible, and configurable, with the ability to evolve and scale to the needs of the project stage whether at planning, engagement, environmental assessment, permitting or through to construction. Backed by ICI's multidisciplinary team with collective experience on hundreds of projects, ICI has the professional support services expertise to scale with project requirements.

RIVAS increases the flow of content within an organization, while meeting information management objectives that are often not easy to achieve. Its easy-to-use tools for data visualization and information management facilitate the exploration and understanding of vast amounts of disparate data from many sources using a single system. ICI's goal is to make information flow better, so its customers achieve their goals faster.