

# Unlocking Intelligence in Industrial Data

How GILT transforms dormant operational data into a strategic asset for mission-critical industries

## The Opportunity

Mission-critical facilities — LNG plants, mines, refineries, renewable energy hubs, power stations and petrochemical complexes — generate vast quantities of data across their lifecycle. Design documentation, process historians, SCADA telemetry, OT network configurations, regulatory submissions and maintenance records accumulate over decades of operation. Yet this data remains overwhelmingly stale, static and siloed. It sits in disconnected systems, decaying in relevance, inaccessible to the engineers and operators who need it most.

The consequences are quantifiable. Engineers spend 30–50% of their time searching for information rather than applying it. Knowledge is permanently lost during personnel changes and operator transitions — often representing millions of dollars of accumulated institutional memory. Compliance obligations are tracked in spreadsheets that become outdated within months, creating regulatory exposure. Production issues are diagnosed reactively, drawing on individual experience rather than systematic data analysis. OT network configurations drift from baseline without detection, creating security vulnerabilities in environments where a breach could have physical safety consequences.

GILT's thesis is simple: this data is not worthless — it is dormant. With the right architecture, it becomes one of the most valuable assets an operator possesses. The technology to unlock this value exists today. The constraint is no longer technology — it is adoption.

## How GILT Works

GILT operates across three distinct stages, each engineered for the unique characteristics of mission-critical industrial data:

- **Ingest** — We connect to your existing data sources — P&IDs, datasheets, process historians, SCADA systems, regulatory documents, CMMS platforms, vendor files and OT network configurations — and parse them into structured representations using domain-specific extraction pipelines. Our ingestion handles multi-format data including PDF, DWG, DXF, scanned TIFFs, Excel, CSV, XML and proprietary historian formats. Fine-tuned vision models recognise P&ID symbols, instrument bubbles, line numbers and annotation blocks. Every entity is tagged with a unique identifier and linked to its document context.
- **Organise** — Data is assembled into a multi-layer knowledge graph that preserves the relationships between entities. A control valve exists simultaneously on a P&ID, in a datasheet, on a line list, in a C&E diagram, in HAZOP documentation, in maintenance records, in spare parts inventory and in live process data. GILT's knowledge architecture preserves and exploits all of these relationships. Domain-adapted embeddings ensure that industry-specific terms carry their correct meaning during retrieval. Hybrid search combines semantic similarity with exact-match keyword retrieval for precision.
- **Deliver** — Natural-language query interfaces let engineers, operators and managers interact with their entire data estate conversationally. Every response is grounded in source documentation with full provenance traceability — document name, page number, revision. The system supports conversational Q&A, automated reporting, interactive dashboards, proactive alerting and integration with existing tools like Excel and PowerBI.

## Key Applications

- **Intelligent Digital Twins** — Queryable facility representations built from existing engineering documentation in weeks, not months. No 3D modelling required. Every tagged item on a P&ID is linked across all document boundaries, creating a navigable, searchable representation of the entire facility — even for brownfield assets where traditional digital twins were previously considered infeasible.
- **Document Intelligence & Knowledge Recovery** — Automated classification, semantic deduplication and natural-language search across thousands of documents. When projects change hands or experienced engineers retire, GILT captures and preserves institutional knowledge so it remains accessible to the incoming team from day one.
- **Compliance Tracking & Scheduling** — Every compliance obligation is extracted from approval documentation and structured into trackable registers with automated scheduling, Gantt charts and calendar reminders. Multi-jurisdictional obligations are managed in a single, unified view.
- **AI-Assisted HAZOP & LOPA** — Historical HAZOP reports are ingested and cross-referenced against the facility's P&IDs and design basis. Common hazard scenarios are identified automatically, LOPA rulesets determine SIL requirements with full audit trail, and workshop time is reduced by 40–60% while improving consistency.
- **Automated Fault Finding & Troubleshooting** — When equipment fails or a plant trips, GILT interrogates trip data, alarm sequences, historian trends and equipment documentation to identify the most probable root cause — reducing diagnosis time from hours to minutes.

## Security & Data Sovereignty

GILT's security posture is informed by IEC 62443 (Industrial Automation and Control Systems Security), ISO 27001 (Information Security Management), NIST SP 800-82 (ICS Security), the Australian Essential Eight and the AESCSF (Australian Energy Sector Cyber Security Framework). Client data is stored within agreed jurisdictions — for Australian clients, all processing occurs on Australian-hosted infrastructure. Data is encrypted in transit (TLS 1.3) and at rest (AES-256) with per-client key management. Processing environments are logically isolated per client. GILT does not retain client data or intellectual property beyond the engagement period — all data, embeddings, vector indices and knowledge graph contents are securely deleted upon project completion with cryptographic verification.

## Why Now

The technology to transform dormant industrial data into live intelligence exists today. Large language models provide the reasoning capability. Retrieval-augmented generation provides the grounding. Knowledge graphs provide the structure. Domain-specific engineering provides the accuracy. GILT brings these components together with the one ingredient that no generic AI platform can replicate: deep understanding of mission-critical operations and the problems that matter most to the people who run them. The operators who move first will compound their advantage as their knowledge bases grow, learn and deliver increasingly powerful insights over time. The future of industrial intelligence is not a roadmap item — it is deployable now.