



CUSTOMER CASE STUDY

PI Tag Mapping for Contextualization Agents

Transforming tag-to-equipment mapping into an auditable reasoning workflow

Solution: Contextualization Agents **Industry:** Industrial Operations **Use Case Focus:** PI Tag-to-Equipment Mapping

GOALS

- Establish traceable and auditable tag-to-equipment mapping processes.
- Build a reusable decision-memory layer to scale contextual asset intelligence.

SOLUTIONS

- Deploy Contextualization Agents to expand metadata, compare candidates, and self-reflect.
- Persist full reasoning paths, constraints, and approvals into a structured Context Graph.
- Embed human verification directly into the agent decision loop.

CHALLENGES

- Inconsistent, abbreviated tag metadata lacks necessary process and hierarchical context.
- Traditional manual reviews or static rules discard the engineering reasoning behind decisions.

RESULTS

- Reduced hallucination risks in future mappings by referencing accumulated decision history.
- Achieved end-to-end traceability for simplified compliance and scalable governance.
- Prevented mapping rebuilds when physical plant structures or hierarchies evolve.

EXECUTIVE SUMMARY

PI tag-to-equipment mapping is a complex decision problem hindered by inconsistent metadata and lost engineering context. This article outlines how Contextualization Agents formalize mapping as a reasoning workflow, capturing the full decision path within a structured Context Graph. By preserving alternatives, constraints, and validation steps, this approach replaces static rules with a reusable decision-memory layer, ensuring traceable, scalable, and audit-ready data governance across industrial assets.

The Challenge of Tag Ambiguity

Industrial operators frequently struggle with inconsistent and abbreviated PI tags that lack essential process context. A single tag can often reasonably map to multiple candidate assets. Traditional solutions, such as manual engineering reviews or rigid static rulesets, might yield a temporary mapping but completely discard the underlying logic. Critical decision factors—including why a mapping was triggered, which alternatives were evaluated, and how engineers validated the outcome—are lost. This lack of preserved reasoning makes mappings highly difficult to audit, scale, or maintain as physical plant configurations evolve over time.

Agentic Reasoning Workflows

Contextualization Agents address this problem by formalizing PI tag mapping as an active reasoning workflow rather than a simple database lookup. These intelligent agents expand tag metadata, automatically generate and compare potential equipment candidates, and apply self-reflection before finalized decisions are made. Crucially, human verification is embedded directly within this active decision-making loop. Instead of saving only the final endpoint, the system logs the entire journey—including evaluation contexts, alternative candidates, and constraints—directly into a structured Context Graph database.

Building Reusable Decision Memory

The resulting Context Graph acts as a robust, reusable decision-memory layer for the entire enterprise. Subsequent mapping tasks can reference this accumulated reasoning history instead of restarting the computational inference process from scratch. This drastically minimizes hallucination risks and guarantees high consistency across different plants. Each mapping remains fully traceable from end to end, facilitating seamless data governance, simplifying engineering audits, and building long-term contextual intelligence across all physical assets and software deployments.