OCTAVE

WHITE PAPER

Context in, insights out.

Octave InConcert Assistant turns contextualized engineering data into insights — and becomes the AI assistant your project and operations teams can rely on.





The challenge

Engineering project teams and asset owners of industrial facilities are under pressure on three fronts at once. Capital projects overrun their budgets by 30–45% on average, with information gaps and document-handover failures consistently identified as the largest avoidable category.[1] Engineers — in EPCs and inside operating organisations alike — lose 1.8–2.5 hours per person per day searching across five to eight systems for information that already exists in the enterprise.[2] New engineering hires take 12–18 months to reach full productivity, because the asset-specific knowledge they need is held in the heads of individuals rather than in an accessible, chronicled foundation.[3] In every case the pattern is the same: the right data exists somewhere in the enterprise — it just is not accessible quickly enough to act on.

Why generic AI does not solve the industrial data problem

It is reasonable to ask whether general-purpose assistants — ChatGPT, Microsoft Copilot, Gemini — could close the gap. For engineering work today, the answer is no. Generic assistants see a folder of documents; they cannot know that a particular datasheet revision applies to one tag and not the analogous tag on a parallel circuit, or that an inspection finding from the most recent shutdown has or has not been resolved in the latest datasheet revision, or which line-class specification governs a specific tie-in. Without that asset-context layer, general purpose assistants fill gaps with statistically likely values rather than admitting they do not know — acceptable in marketing copy, but clearly unacceptable when an engineer is verifying that the design pressure on the datasheet matches the tag property and the line specification. Bolt-on AI connectors that index documents into a separate vector store create a parallel problem: a duplicate copy of sensitive engineering data outside the data lineage and access controls of the original system, with stale snapshots that lag the source of truth.

---

[1] Ernst & Young / Independent Project Analysis (IPA), oil & gas megaproject overruns of 30–45%.

[2] McKinsey & Company / IDC research on knowledge-worker information retrieval (~1.8–2.5 hrs/day on search).

[3] Octave + customer interviews; engineer onboarding 12–18 months in process and power.

Introducing Octave InConcert Assistant

Octave InConcert Assistant is the AI assistant built into Octave InConcert (formerly HxGN SDx2). AI assistants are also available for other Octave products — including Attune EAM (formerly HxGN EAM) and Sequence Enterprise (formerly EcoSys), with the family growing — each tailored to its product's data foundation. This paper focuses on the InConcert variant. It is purpose-built on an agentic framework — a composition of specialised agents that reason across the asset data foundation rather than a single monolithic chatbot. Engineers ask questions about any asset — pump, vessel, valve, line, instrument — in natural language and receive a contextualized, source-cited answer in seconds. Project engineers verify that design properties are consistent across datasheets, P&IDs, isometrics, and inspection reports. Capital projects teams check approval status and the state of open change requests. Inside owner-operators, reliability engineers investigate recurring trips and look for analogous failures in maintenance history; plant operators check the latest inspection of a valve before isolating it. Asset owners run engineering projects too — InConcert Assistant works the same way for these audiences. Time consuming workflows that previously required navigation across multiple systems reduce to a question.

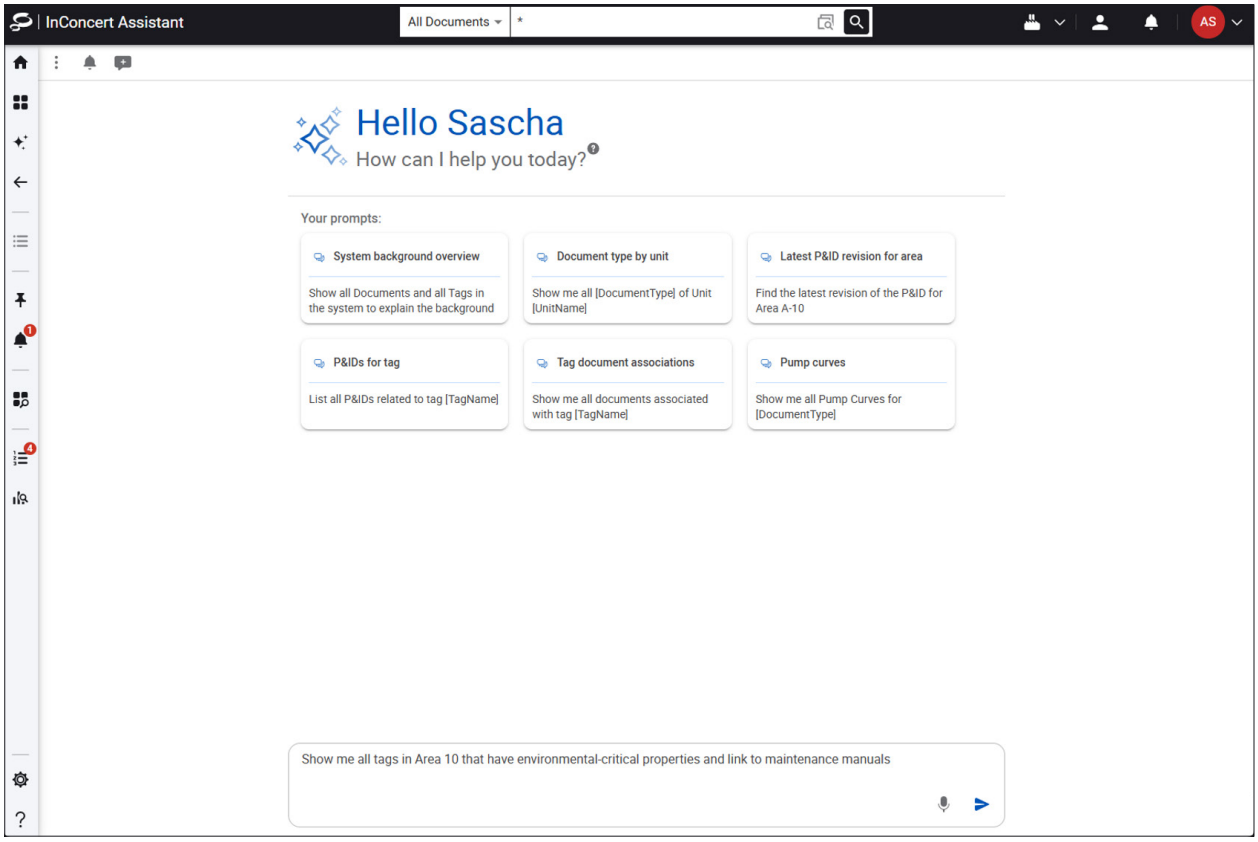


Figure 1. InConcert Assistant in action

Why contextualized data is the prerequisite

It is tempting to describe this argument with the well-worn aphorism: garbage in, garbage out. That formulation is too narrow. The problem with general-purpose AI in engineering is not that the data is garbage. It is that the data is disassociated — a heap of documents that may individually be high-quality, but whose relationship to specific assets, tags, and operating contexts has been lost. Octave InConcert is best understood as an asset-contextualized data model: when a document is ingested, the platform extracts the tags, properties, and references it contains and creates persistent links between the document and every asset it describes (see the InConcert blog, Capturing Information and Making it Intelligent).[4] The customer's investment in the foundation compounds — every datasheet, P&ID, isometric, or inspection report ingested deepens the network of relationships available for any subsequent question.

When an engineer asks InConcert Assistant about pump P-9000, it does not run a keyword search across a flat document store. It resolves the question to the asset, walks the existing links from that asset to its connected documents and properties, and synthesises an answer from only that connected material. Each cited source — datasheet revision, P&ID, isometric, inspection report — appears in the response with its document name and revision identifier, so the engineer can verify any answer against the original record before acting on it.

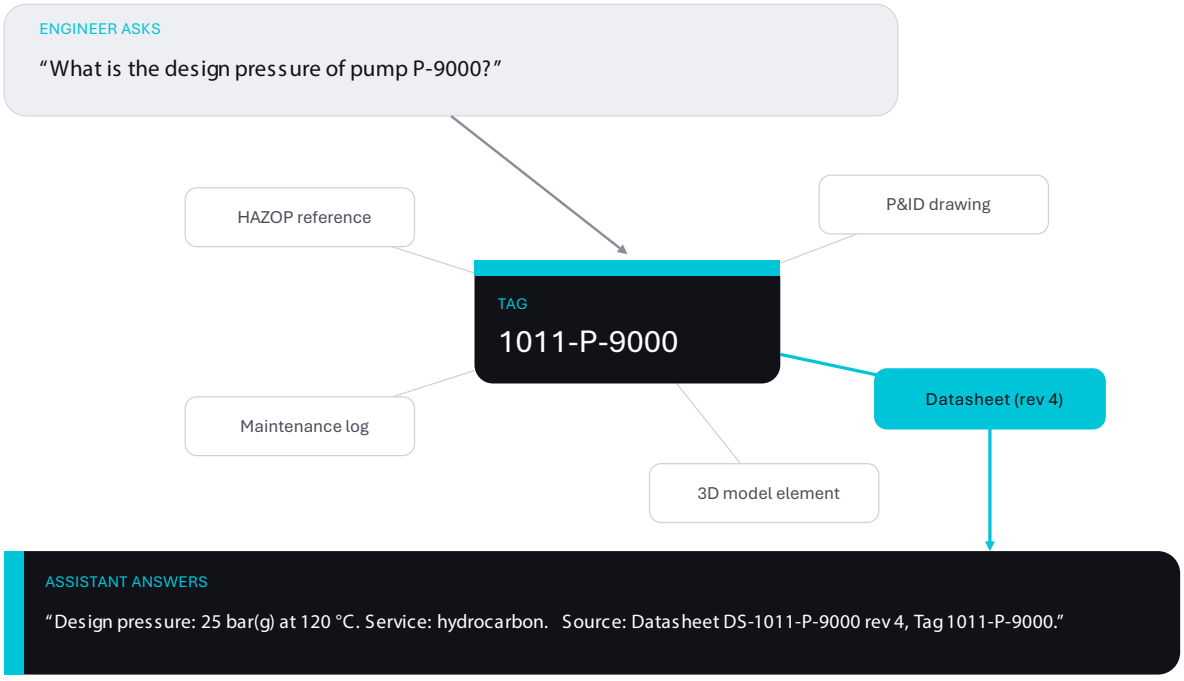


Figure 2. InConcert Assistant walks the connections between a tag and its linked records — datasheet, P&ID, 3D model, isometric, inspection report — answering only from material that is contextually connected to the asset.

[4] Octave Resources, "InConcert: Capturing information and making it intelligent."

Trust by design

Trust in an engineering decision-support tool is crtitical and a property of its design. InConcert Assistant is intended to deliver reliable, accurate answers, not ceding autonomous control. Four properties make its answers safe to use in safety-adjacent work:

- **Source-cited by default** — every response carries the documents, revisions and tags Assistant relied on, so the engineer can verify any answer against the original record before acting on it.
- **Asset-grounded, not generative-grounded** — Assistant answers only from data connected to the asset and says when the data is missing, rather than inventing it.
- **Need-to-know by inheritance** — Assistant honours the access-control policy of InConcert: every user sees only the documents and tag data their existing role permits. The assistant never bypasses, broadens, or duplicates entitlements.
- **Human in the loop** — the decision (to release a work order, approve a change, isolate a valve) remains with the engineer accountable for it. Assistant proposes; the human commits.

Together these properties place the human at the centre of every decision. InConcert Assistant's job is to surface verifiable context faster than any individual could; the engineer's job — and accountability — remains the same. That human–agent collaboration is the point.

Where InConcert Assistant will take the conversation

Today, Assistant retrieves. Future releases are intended to extend it from retrieval into action — turning Assistant into the central conversational entry point of InConcert, through which engineers can not only retrieve information but trigger actions: drafting a work order, opening a change request, generating a handover package. Each action will remain under explicit human authorization. Capabilities will be packaged as composable skills — units that combine specific instructions, a purpose-built UI for that workflow, and (where determinism is required) executable code, often built from multiple agents in an observable, auditable pipeline. Skills are not part of Assistant today; they describe how the platform will grow. The first candidates we are evaluating with design partners are document-quality review and specification compliance checks. Each will ship only once it has been validated against real engineering data. The architectural commitments — native to the foundation, human-in-the-loop, fully traceable — hold for everything we ship.

About Octave

Octave is a leader in enterprise software, turning data into decisive action and intelligence into your edge. Our software solves for and simplifies complexity, from the design and build to operations and protection of people, property and assets – for any scope, at any scale. For decades, we've partnered with customers to sharpen performance, elevate efficiency, and amplify results. From factory floors to entire cities, our solutions are tuned to scale up what's possible from day one onward.

©2026 Octave Intelligence plc and/or its subsidiaries and affiliates.
All Rights Reserved.