

CASE STUDY

Oil and Gas Leader Whitecap Unifies Upstream Data with Shakudo

<1 HR

for business users to build custom analytics, replacing multi-week dev cycles

1 TB/MO

of relational data, governed on a foundation sized to double or triple as sources land

375K

barrels of oil equivalent a day, supported by a lean analytics team

COMPANY

Whitecap Resources

INDUSTRY

Climate Energy

PROFILE

Whitecap Resources is the seventh-largest oil and gas producer in Canada and the largest landholder in the Alberta Montney and Duvernay plays. The Calgary-based company produces around 375,000 boe/d across a roughly 60/40 liquids-to-gas mix, operating in four conventional regions — Alberta, West Saskatchewan, East Saskatchewan, and Weyburn — alongside a deep unconventional program in the Montney and Duvernay. Its current footprint is the result of a \$15-billion combination with Veren that closed in May 2025, bringing two of Western Canada's largest light oil and condensate businesses under one operator.

WATCH THE VIDEO

Sixteen years ago, [Whitecap Resources](#) was producing about 1,400 barrels of oil equivalent per day. Today the company runs at roughly 375,000 boe/d, and its [combination with Veren](#) roughly doubled the company in a single transaction. The data underneath that growth has scaled with it.

Each month Whitecap moves about a terabyte of data into its relational systems. That figure does not include the microsecond-resolution drilling, frac, and subsurface streams the company also captures, which can two or three times the monthly footprint on their own. James Wakelin, who started Whitecap's analytics team twelve years ago and now leads a 30-person business intelligence and data group, has spent most of those years building infrastructure to keep up.

For James, the question is not whether AI changes analytics. It is whether the data foundation underneath it is in shape for what the next two years are going to demand.

BUILDING ANALYTICS THAT PREDICT, NOT JUST REPORT

For most of his 28 years in analytics, James Wakelin has watched the industry treat reporting as a backward-looking discipline. Something happened. Visualize it. Analyze it. Hope to do better next quarter. By the time the dashboard answered the question, the operation had already moved on. That is the legacy posture of business intelligence across most of the energy sector, and Whitecap's analytics group is built to leave it behind.

"It started out that we had a ton of data but no information. Everyone saw data, but they couldn't do anything with it. So analytics initially starts turning data into information to give something useful to someone who's going to make a decision. As we get into these larger data sets and new tools, we are hoping to implement predictive analytics where we're predicting what the next thing looks like."

James Wakelin

Director of Business Intelligence @ Whitecap

Whitecap's data volumes are not unique to Whitecap. Wells, pipes, fluids, daily production, frac and drilling streams, vendor feeds, subsurface telemetry. Every operator at this scale collects all of it. The difference is what gets done with it. Most companies report on what the data has already shown. James wants the platform to predict what the data is about to show, and to do it fast enough that operators have time to react.

SCALING WITHOUT THE SUPERMAJOR HEADCOUNT

Whitecap operates with a different model than the supermajors. The largest producers employ specialists in every discipline a data platform might need. Whitecap, even at 375,000 boe/d, runs lean by design, with a focused team paired with strong tools. That choice shapes what good infrastructure has to do. The platform has to let domain experts work faster without turning them into platform engineers, and it has to be the leverage point that lets a smaller team operate at supermajor data volumes.

The vendor landscape does not make that easier. The same physical measurement arrives under different names from different providers. Reconciling those feeds manually is one of the problems that scales linearly with each new well, each new vendor contract, each new acquisition. The Veren combination doubled the surface area of all of it overnight.

Manual reconciliation, ticket-based analytics requests, and long requirements cycles are the patterns that quietly erode analytics teams across the industry. They are also the patterns Whitecap moved early to get ahead of. Stack the operating realities together and the picture looks like this:

- A monthly terabyte of relational data, plus microsecond-resolution time series from drilling, frac, and subsurface operations
- A near-doubling of the business through the Veren combination, with employees, systems, and data needing to integrate quickly
- Vendor feeds where the same physical measurement appears under different names across providers

- Long requirements cycles between business users and IT, vulnerable to scope creep and miscommunication
- The risk of shadow IT as business users build their own prototypes faster than IT can productionize them
- A founding mandate to deliver reliable returns to shareholders, which makes IT cost discipline non-negotiable

FROM REQUIREMENTS DOCUMENTS TO WORKING PROTOTYPES

Long before AI was the headline, Whitecap had been laying the groundwork. About a year and a half ago, the company began working with [Shakudo](#) to build a foundational data layer for its analytics group. AI was not the priority. The priority was a single, well-understood place from which the business could pull clean data fast.

That sequencing turned out to matter. When the Veren combination closed, the data team was not starting from scratch. They had a base that could absorb the new operations, the new vendors, and the new headcount. The analytics group did not double overnight, but the surface area of the data did, and the platform was already designed to scale into it.

"We started out with Shakudo about a year and a half ago as a way to build a foundational data layer for our analytics. We weren't totally contemplating AI at that point, but we knew we needed a solid foundation. Then we went through our business combination and doubled in size, and that foundation needed to double too. What started out as the foundational layer, which we needed, will turn into really an advanced AI tool for our business."

James Wakelin

Director of Business Intelligence @ Whitecap

Built on that foundation, a different operating model became possible. The traditional sequence is the part James thinks AI tools collapse most aggressively. A business user describes a need, a project manager translates it, a developer interprets that translation, and scope creeps from there.

"A POC driven by the business users instead of a requirements document is going to be what changes this space. Where they can citizen code, or vibe code, a small program that gets them an objective and they can turn around and say, this is what I need on a corporate scale. I don't think anyone's ever been really good at defining requirements. When someone can spend a day developing with AI coding tools to get them what they need to see, that becomes a living, breathing requirements document."

James Wakelin

Director of Business Intelligence @ Whitecap

James calls this AI entrepreneurialism. The business owns the idea, proves it at small scale, and keeps ownership of the outcome. IT and analytics turn the prototype into something secure, scalable, and maintainable. Accountability does not move from business users to IT once the work goes to production. It stays with the people who originated the need.

The trade-off is that this only works if IT can move quickly. If the platform team cannot productionize prototypes faster than business users can build them, the prototypes become permanent. Three different groups end up running near-duplicate workflows, performance degrades, and the platform spends its time chasing fragmentation. Whitecap's answer is a tooling stack that keeps the gap between prototype and production short enough that experimentation feeds the platform instead of competing with it. This is the same logic behind a disciplined build-versus-buy approach: build what differentiates Whitecap, and buy the layers that do not.

It is also where tools like Kaji for agentic workflows and natural-language-to-SQL queries earn their place. Shakudo's unified integrations make it possible to add these tools without standing up new infrastructure for each one.

SOVEREIGN BY DESIGN

For Whitecap, the architecture choice came before the AI choice. The company is on-prem by deliberate decision, and that decision predates the current generation of frontier models. It is a posture about jurisdiction, not a posture about technology.

"We're an on-prem company because we didn't want our data exposed to laws that we had no control over. Put it in the cloud and it may be subject to a different country's legal framework and exposure to being opened at any time. The company made a decision to keep as much on-prem as possible, because then we had control over where that information ended up."

James Wakelin

Director of Business Intelligence @ Whitecap

That does not mean rejecting frontier models. It means using them deliberately. Whitecap negotiates strong contractual safeguards with its enterprise model providers, including no-training and no-retention terms, but it also runs open and on-prem models for the work where contract language alone is not the right control. The result is a hybrid posture: a frontier model when the capability gap justifies the exposure, an on-prem model that may be a version or two behind for the majority of work where the capability gap does not.

The same posture extends to how James thinks about output trust. AI is treated as a colleague whose work still gets reviewed. Repeatable tasks earn trust through volume. After 20 or 200 consistent answers, a workflow can move from supervised to lightly supervised. New questions, exceptions, and high-stakes outputs stay in front of a human. Validation is a discipline, not a stage.

For a deeper read on the architecture and contractual choices behind that posture, see Shakudo's perspective on [AI agents in regulated industries](#) and the [on-prem LLM and RAG stacks](#) that support them.

INVESTING IN CUSTOM ANALYTICS

As AI tooling matures, Whitecap is doubling down on a posture that started with a data foundation and has grown into a working platform. Eighteen months in, the substrate has done its job. The relational layer, the time-series layer, and the emerging AI layer all run on the same governed environment. The analytics team has absorbed a doubled business on the same governed foundation. Where the platform stands today:

- One foundational data platform powering relational, time-series, and AI workloads across the combined Whitecap business
- Analytics group of about 30 people supporting the post-Veren operational footprint
- Terabyte-scale monthly relational ingestion alongside microsecond-level operational telemetry on the same platform

- On-prem LLM availability for the majority of internal workloads, with frontier models reserved for capability-gap use cases
- Business-led prototyping pattern in place, with IT productionizing rather than originating
- Foundational layer extensible into agentic and predictive analytics work without a platform rebuild

The endpoint James is building toward is not a larger BI catalog. It is an environment where a business user can ask a new question and get a trustworthy answer in an hour or less, without waiting for a development cycle to deliver it.

"I truly believe custom analytics will be driven by AI. We're not reliant on functionality that exists in current analytic programs. If we come up with an idea for a new analytic, it can be created on the fly and added to the arsenal of what people are analyzing. Independent customization by the individual will really take off, so we need to make sure our data is clean and in the right spot, with the right metadata, that the AI agents can use to give the right answer."

James Wakelin

Director of Business Intelligence @ Whitecap

That future puts more weight, not less, on data hygiene. Metadata, lineage, and structure all matter more when an AI agent is the one assembling an answer. The work the team is doing now is preparing the substrate that determines whether those agents become a force multiplier or a generator of plausible-looking nonsense.

Agentic tooling is the layer that follows. Systems like [Kaji](#) point toward governed action, helping teams move from analytics and recommendations into outcomes a business can act on, while people continue to handle exceptions, oversight, and strategy. For a broader read on how this plays out in energy operations, see Shakudo's [practical guide to AI in oil and gas](#) and the [climate and energy industry overview](#).

Whitecap's approach is useful precisely because it does not romanticize AI. The company is building from operational constraints outward: real data volumes, real jurisdictional concerns, real cost discipline, and a real preference for keeping ownership of outcomes inside the business. If you are leading data and AI initiatives in energy, industrials, or any environment where the data foundation has to come before the AI story, [get a demo of Shakudo and Kaji today](#).

GET IN TOUCH WITH US

Ethan Brooke

Senior Account Executive

ethan@shakudo.io

(437) 344-4349