Valva

A Federated Gateway and Sovereign Agentic Layer for AI in Oilfield Operations

Prepared by Daniel Hill | Novus Nexum Labs
May 21, 2026

A Federated Gateway and Sovereign Agentic Layer for AI in Oilfield Operations

White Paper from Novus Nexum Labs

Version 1.0, May 2026

Executive Summary

The oilfield is one of the last industries where the constraint on AI deployment is not the model. It is the infrastructure that sits between the model and the systems the model needs to act on. A modern operation, whether on a drilling rig, a completion spread, a production field, a midstream pipeline, or a refinery control room, runs on signal from dozens of vendor systems. Surface and downhole sensing, electronic recorders, historians, distributed control systems, SCADA, alarm management, dispatch, and safety platforms. Each system is a separate vendor integration, a separate compliance perimeter, and a separate audit boundary. An AI agent that needs to reason across them at the speed an operation actually moves runs into three structural problems that current AI platforms do not solve.

The tool catalog for that many vendor systems exhausts the model's context window before the agent begins reasoning. The regulatory and data-sovereignty requirements that operators carry across jurisdictions, from the Norwegian Continental Shelf under GDPR through US BSEE and PHMSA, Canadian CER and AER, UK HSE and NSTA, and the safety regimes of the Middle East and Asia-Pacific, are not enforceable in any consistent way across vendors. And the institutional knowledge that actually drives operational decisions, the patterns held in the heads of drilling supervisors, completion engineers, production operators, pipeline controllers, and refinery panel operators, lives in human memory rather than in any system the agent can call.

Valva is the infrastructure layer built for this problem. It is a federated gateway that lets a single AI session reach an arbitrary number of vendor tools across cloud and edge boundaries through one semantic-indexed surface, with regulatory policy enforced on every call and a defensible audit trail. Above that gateway sits a sovereign agentic runtime designed to run inside the operator's perimeter, on the rig, in the field, or at the plant. The combination opens a category of agentic systems for oilfield operations that does not exist today as a productized offering anywhere in the industry. Agents that encode the institutional knowledge of operators and field crews as enforceable policy, run locally without cloud round-trip, and remain auditable and portable across assets under operator sovereignty rules. Valva is provider-agnostic. The gateway, the policy engine, and the audit trail do not depend on any single model vendor's safety surface to function. The operator chooses the model, frontier or open-weights, and the operator owns the policy and the evidence base.

1. Introduction: The State of AI in Oilfield Operations

Across the oilfield, AI is already in production. Surface vision systems at the shaker. Real-time wellbore-stability inference. Drilling automation that closes the loop on weight on bit and rotary speed. Stimulation models that predict frac geometry from treating pressure. Production optimizers that adjust artificial lift in response to inflow trends. Pipeline integrity systems that infer leak signatures from supervisory data. Refinery soft sensors that fill in for missing analyzer streams. Each of these systems

was built to do one thing well, and most of them do. They were not built to interoperate with each other through a single reasoning surface, and they were not built to participate in an agent loop where one model decides, in real time, that a particular combination of telemetry, alarm state, and document context warrants a specific operational action that the safety system has to sign off on before the field acts.

The reason such an agent loop is hard today has very little to do with the model. Frontier models can reason about these scenarios when they have the right tools in front of them. The reason it is hard is that the tools are scattered across vendor systems with no shared call surface, no shared compliance gate, and no shared audit posture. Every operator has built or bought integration glue to make slices of this work, but the glue is fragile, per-asset, and almost never independently auditable. AI deployments that promise to act on this fabric usually do so by shipping all of the data to a vendor cloud, training a model, and shipping back inference. That model is inadequate for jurisdictions and operators that require data sovereignty, for incident timeframes that do not tolerate cloud round-trip, and for the level of regulatory evidence that European operators in particular increasingly demand. It is also a posture that large service companies and OEMs cannot offer their operator customers without re-architecting their own products around someone else's cloud, which is rarely a commercial fit.

2. Why Current AI Architectures Fail in Heavy Industry

2.1 The context-window ceiling

Frontier LLMs are deployed against an external tool set through protocols such as the Model Context Protocol (MCP), which exposes each tool's name, description, and JSON schema as a context-window resident artifact. The cost is small per tool, a few hundred to a few thousand tokens depending on schema complexity. The cost is also linear in the number of tools, and in operations the practical ceiling is reached well before a hundred tools are attached to a session. Past that ceiling the agent's reasoning quality degrades, tool selection becomes unreliable, and the model frequently calls the wrong tool or hallucinates parameters because the relevant tool descriptions are buried in noise.

A drilling rig with surface, downhole, mud-logging, EDR, dispatch, and safety systems exceeds that ceiling. A production field with SCADA, lift control, separator and treater telemetry, flow metering, asset performance management, and document systems exceeds that ceiling. A midstream control center with line monitoring, compressor station telemetry, terminal automation, leak detection, and alarm management exceeds that ceiling. A refinery DCS with thousands of tags, soft sensors, lab data, and procedure libraries exceeds it many times over. Vendors hide the problem by shipping a vertical slice of pre-curated tools, which works only inside that slice. The moment an operator wants the agent to reason across slices, the architecture breaks down.

2.2 The compliance gap

Tool calls in stock deployments are unauditable in any defensible sense. There is no consistent way to assert that a tool call is permitted under HIPAA, HITRUST, ISO 20022, GDPR, the EU AI Act, DORA, or a sector-specific framework, and no consistent way to produce evidence after the fact that a given call

was made in compliance with a given profile. In oilfield, where API specifications, ISO 29001 quality management for the petroleum supply chain, IEC 62443 industrial control system cybersecurity, and ESG reporting obligations all touch the same infrastructure, the absence of a consistent enforcement and evidence layer is a blocker, not a nice-to-have. The same is true for the US side, where BSEE for offshore operations and PHMSA for pipeline and storage operations expect a defensible record of how automated and assisted decisions were made.

2.3 The sovereignty gap

European operators operate under GDPR Article 28 processor obligations, and operators on the Norwegian Continental Shelf carry an even stricter posture on data residency and processor disclosure. An operator on the NCS has to be able to demonstrate that every party touching well or production data, including AI vendors, complies with processor responsibilities, audit rights, sub-processor disclosure, and data-residency rules. National oil companies in the Middle East and parts of Asia-Pacific carry comparable expectations on national sovereignty over operational data. The cloud-only AI deployment pattern, where field data is shipped to a vendor cloud for inference, breaks this without significant additional contractual and architectural work. Even when it can be made compliant, every additional vendor in the stack is another processor relationship the operator has to defend to its regulator.

2.4 The institutional-knowledge gap

The most valuable assets in an oilfield operation are not the racks, the cameras, the analyzers, or the contracts. They are the diagnostic patterns held by the people who run the asset. The drilling supervisor who recognizes the signature of an incipient pack-off in a particular combination of surface and downhole signal before the standpipe pressure confirms it. The completion engineer who reads a stimulation pressure response and knows the geometry is going planar instead of complex twenty minutes before the rate decline shows up. The veteran production operator who reads a separator pressure pattern and knows which well is loading up. The pipeline controller who has learned that a specific upstream pressure trend, in combination with a particular ambient temperature regime, warrants a deliberate block-valve sequence rather than a wait-and-see. The refinery panel operator who recognizes a precursor pattern in furnace draft and pass temperatures that the alarm system does not flag for another ten minutes. These patterns are real, they are repeatable, and they are largely undocumented in any form a model can call. They live in memory and in informal notes. They do not scale across assets, they do not survive crew rotations and retirements, and an acquirer cannot put them on a balance sheet. No current AI architecture in oilfield provides a structured way to capture this knowledge as enforceable agentic policy.

3. The Valva Architecture

Valva is built as five cooperating subsystems behind a single call surface. The architecture is model-vendor neutral by design. Operators choose the model, frontier or open-weights, and Valva is the consistent policy, federation, and evidence layer beneath whatever model is in use.

3.1 The federated gateway

At the front of Valva is a gateway through which every AI session reaches every tool. The agent never connects directly to a vendor system. It addresses the gateway through a single MCP-compatible endpoint, and the gateway federates downstream to an arbitrary number of vendor MCP servers, REST APIs, CLI binaries, and proprietary integrations. Federation in this sense is not a marketing word. It is the literal architecture. Valva carries a directory of every downstream server and tool, routes calls to the correct downstream, enforces policy on the way in and the way out, and returns results to the agent as if they came from a single unified provider. A session that touches twelve vendor systems is, from the model's point of view, a session that talked to one provider.

3.2 The semantic index

The piece that decouples catalog size from context-window pressure is the semantic index. When the agent issues a request, the gateway runs the request against an index of every available tool, ranks the tools by relevance, and returns only the top-N most relevant tool descriptors into the agent's context window. The model never sees the full catalog. It sees an accurate representation of every tool relevant to the task at hand, drawn from the full catalog rather than a truncated slice. The ranking is hybrid, combining lexical, embedding-based, and outcome-aware signals, and the index is updated continuously as new servers join or change.

The result is a hard architectural removal of the context-window ceiling. A Valva deployment that exposes thousands of tools across a federation imposes the same context cost on the agent as a Valva deployment that exposes thirty. Our own computation platform, Functo, demonstrates this in practice with 717 tools and 732 validated scientific models across 29 namespaces in a single MCP server, addressable from a single AI session through Valva without overrunning the context window.

3.3 The compliance gate

Every call routed through Valva passes through a compliance gate before it leaves the gateway. The gate is configurable per deployment and per session, and it evaluates each call against a regulatory profile that names which frameworks apply, which data classes are touchable, which destinations are reachable, which transformations are required (for example, secret redaction or PII removal), and which evidence must be captured. Profiles in production today include HIPAA, HITRUST, ISO 20022, GDPR, the EU AI Act, and DORA. The same profile model is extensible to oilfield-relevant frameworks including API specifications, ISO 29001, IEC 62443, and ESG reporting standards. When a call is denied by the gate, the gate returns a structured denial that the agent can reason about and the operator can review. When a call is permitted, the gate stamps it with the policy decision and forwards it. In either case the decision and its inputs are recorded. Policy enforcement is the gateway's responsibility, not the model vendor's, which is why the same operator deployment can swap between models without re-negotiating its regulatory posture.

3.4 The audit trail

Every call routed through Valva, permitted or denied, lands in an append-only audit trail with the operator identity, the regulatory profile in force, the inputs, the result, the policy decision, and a defensible record an auditor can reconstruct from. The audit format is structured, signed end-to-end, and exportable into the operator's evidence systems on a schedule the operator controls. For oilfield deployments the audit is the same artifact a regulator, an insurer, or an incident investigator would reach for after a well-control event, a process-safety event, or a pipeline release. A complete, ordered, defensible record of what the AI did and why.

3.5 The on-rig and in-perimeter agentic runtime

Above the gateway sits a sovereign agentic runtime designed to run inside the operator's perimeter rather than in a vendor cloud. On a rig it runs on rig hardware. In a production field it runs at the field office or in a regional operations center. On a pipeline it runs in the control center or in a hardened edge node along the line. In a refinery it runs in a control-room-adjacent compute environment behind the plant firewall. The runtime hosts the model (frontier or open-weights, operator choice), the agent loop, and the policy execution engine that turns codified field-knowledge patterns into enforceable agentic behavior. The runtime calls Valva-gateway tools through the federation, makes decisions within the operator-written policy, and produces auditable actions. No part of the agent loop is required to leave the operator's perimeter. Data exits the perimeter only when the operator explicitly permits it.

4. The Oilfield Reference Deployment

The Valva deployment shape is common across oilfield phases, but the systems it federates and the workflows it carries differ by phase. The four subsections below describe the reference Valva deployment for the four families of operations our partners and prospective customers most often ask about.

4.1 Drilling operations

In a drilling deployment Valva federates the systems that already live on or around the rig. The federation typically includes:

- Surface monitoring streams from cameras, mud-logging instruments, and rig-floor sensors
- Downhole telemetry through MWD and LWD
- Electronic drilling recorder feeds, including hookload, RPM, weight on bit, standpipe pressure, ECD, TVD, and inclination
- WITSML streams from the operator and contractor
- Mud-system telemetry and laboratory mud-property updates
- Predicted rock-strength and formation models
- Well-program documents, bit programs, casing programs, and pore-pressure plans
- Dispatch and safety platforms with second-level latency expectations

Valva sits in the operator's chosen perimeter, on the rig for offshore deployments where Starlink Maritime is the wide-area transport, or in the operator's regional cloud for onshore deployments. The agentic runtime above the gateway hosts on-rig workflows the operator owns, ranging from pack-off and hole-cleaning protection to bit-wear inference and real-time wellbore stability adjustment. The transport is designed for offshore link economics. Ranked tool descriptors and result handles move over the link, not raw payloads, which means the federation remains usable on a Starlink Maritime connection without paying twice for the same bits through cloud egress.

4.2 Well completion and intervention

In a completion or intervention deployment Valva federates the systems that the completion operator, the stimulation provider, and the wireline or coiled-tubing crews already run. The federation typically includes:

- Frac fleet telemetry, including pump rates, treating pressure, slurry concentration, and chemical rate
- Perforating gun-string telemetry and depth correlation
- Coiled tubing and wireline surface readouts
- Distributed acoustic and distributed temperature sensing where deployed
- Completions data systems and stage plans
- Cement bond and pressure-test records
- Stimulation design models and predicted geometry
- Well integrity and barrier-management documents

The agent loop above the gateway carries workflows that depend on cross-system reasoning. Recognizing the signature of near-wellbore tortuosity in treating pressure and slurry profile, catching screen-out precursors earlier than the chart will, comparing actual to designed geometry stage by stage, and coordinating between fleet, wireline, and operator engineers without having to wait for human translation between systems. The same compliance gate and audit trail that cover the drilling deployment cover this one without re-architecture.

4.3 Production operations

In a production deployment Valva federates the systems that run a producing asset, onshore or offshore. The federation typically includes:

- SCADA tag streams and historians
- ESP, rod-pump, and gas-lift telemetry
- Separator, treater, and free-water knockout level and pressure data
- Multiphase and single-phase flow metering
- Asset performance management and reliability platforms
- Chemical-injection and flowline-integrity data
- Document systems carrying operating envelopes, alarm rationalization, and procedure libraries

The agentic runtime carries workflows that are notoriously hard to scale across operators because they depend on senior-operator judgement. Identifying which wells are loading up before the production report shows it. Recognizing when a separator's level trend warrants a setpoint change rather than a continued manual override. Catching ESP or rod-pump precursors of an unscheduled trip far enough in advance to schedule a controlled intervention. Coordinating lift, choke, and chemical changes inside an envelope the operator already wrote. The audit trail produced by these workflows is the same artifact production engineering, reliability, and HSE would reach for in a routine review or in an investigation after a release or a shutdown.

4.4 Midstream and refining

In a midstream or refining deployment Valva federates the systems that run a pipeline, a compressor station, a terminal, or a refinery process unit. The federation typically includes:

- Pipeline supervisory data, line pressure and temperature, and meter-station telemetry
- Compressor station vibration, performance, and emissions monitoring
- Terminal automation, tank gauging, and movement scheduling
- Refinery DCS tag servers and historian access for process units
- Alarm management and rationalization records
- Process safety management documentation, including operating procedures, MOC records, and incident histories
- Leak detection, pig tracking, and right-of-way monitoring

The agentic runtime carries workflows that turn pattern recognition into auditable action. Catching a pipeline pressure-and-temperature trend that warrants a block-valve sequence and surfacing the controller's intended sequence with the supporting evidence inside the controller's existing console. Reading a compressor performance shift against ambient conditions and flagging the maintenance window before mechanical signature confirms it. Reasoning across refinery unit telemetry, alarm history, and the operating procedure on file to surface the precursor pattern a panel operator has learned to watch for. The compliance gate is configured for the applicable framework set, which in practice means PHMSA and API for US pipelines, IEC 62443 for the control system perimeter, and the operator's process safety management profile for the plant. The audit trail is built for the same investigators who would walk in after a release or a process-safety event.

5. Codifying Institutional Knowledge as Sovereign Agentic Workflows

The most important capability Valva adds to an oilfield deployment is not the gateway and not the audit trail. Those are prerequisites. The most important capability is what becomes possible above them, which is the structured capture of operator institutional knowledge as enforceable, auditable, in-perimeter agentic workflows.

5.1 The asset hiding in operator memory

An operator's edge over its peers is rarely the hardware. Every operator on the same shelf, in the same basin, on the same pipeline corridor, or running the same refinery configuration has access to the same vendors, the same instrumentation, and the same procedures. What differentiates one operator from another is the accumulated diagnostic capacity of its people. The patterns the drilling supervisor has learned to read, the completion engineer's instinct for when geometry is going wrong, the production operator's feel for which wells will fail next, the pipeline controller's calm in the minutes before a confirmed event, the refinery panel operator's eye for furnace and pass-temperature precursors. This asset is the operator's most valuable training signal for any AI deployment that would actually move the operation, and today it is the asset most invisible to any AI architecture in the field. It lives in heads and notebooks, transfers slowly through apprenticeship, and walks out the door when a crew rotates or a senior operator retires.

5.2 From observation to enforceable agentic policy

Valva's in-perimeter agentic runtime makes this asset codifiable, and the development pattern is concrete. Our engineers travel to the asset and sit with the people doing the work. We observe the diagnostic patterns in practice, codify each pattern as an explicit agentic policy, run the policy as a Valva-mediated agentic workflow on the operator's hardware, and let the engineer or operator verify the agent's reasoning matches the human reasoning before the policy is signed off as production. The result is not an agent that approximates the field expert. It is an agent that encodes the field expert's exact decision boundaries, expressed as policy a human at the asset wrote and signed off on, executed inside a runtime the operator owns, with every action captured in an audit trail the operator can defend.

5.3 Sovereignty in the strict sense

The agent is sovereign in the precise technical sense that matters for regulators and operators. The model runs inside the operator's perimeter. The policy is the operator's. The data feeding the agent never leaves the perimeter unless explicitly released. The audit trail is the operator's evidence base. No part of the loop requires a cloud round-trip. When the operator chooses to share aggregated learning across assets in their portfolio, the sharing happens through Valva federation under explicit rules the operator wrote, not by exfiltrating raw data to a vendor.

5.4 IP boundaries that hold across operator, service company, and platform

The development model creates clean IP boundaries that every party can defend. Domain perception and analytics IP stays with the vendor that produced it. Field knowledge stays with the operator, expressed as the operator's policy artifacts and never as data shipped to anyone else. Orchestration, sovereignty, and policy-engine IP stays with Valva. The productized result of a co-developed deployment is jointly ownable on terms agreed in writing before the development begins. The development cadence is intense by design. In-field engineering with our team, the relevant service or analytics vendor, and the operator at the asset when that is where the knowledge needs to be captured, followed by a sustained scalable release cycle once the patterns are codified and ready to fan out across the federation.

6. Compliance and Audit for Regulated Operations Worldwide

The most stringent compliance target in oilfield AI today is the European operator on the Norwegian Continental Shelf operating under GDPR. Designing for that target produces an architecture that other jurisdictions can adopt as a relaxation rather than as a redesign.

6.1 The NCS and GDPR baseline

When an operator deploys Valva on the NCS, Valva operates as a processor under GDPR Article 28 with respect to any personal data that flows through the gateway. Valva's compliance profile for GDPR encodes the processor obligations directly. Documented sub-processor disclosure, audit rights granted to the operator, configurable data-residency constraints that the compliance gate enforces on every call, and explicit retention boundaries on the audit trail itself. Norwegian operators increasingly require that operational data remain inside the operator's sovereignty perimeter for the entire processing lifecycle. Valva supports this directly. The gateway, the policy engine, and the audit trail can all run inside the operator's perimeter, on the rig or platform, in an in-country regional cloud, or in an operator-owned data center, with no external dependency that would require data to cross a national or contractual boundary.

6.2 European Union frameworks

For European operators outside the NCS, GDPR and the EU AI Act apply in their general form. DORA applies to operational-resilience-relevant integrations such as financial reporting and trading-adjacent systems that many integrated operators run alongside their physical operations. Valva carries the GDPR, EU AI Act, and DORA profiles in production today, which lets an EU operator stand up a Valva deployment with the European regulatory posture configured rather than constructed.

6.3 United States frameworks

For US operators the relevant frameworks include BSEE for offshore operations on the Outer Continental Shelf, PHMSA for pipelines and storage, API recommended practices and specifications across upstream and midstream, and ISO 29001 for the petroleum quality management system. IEC 62443 applies to the industrial control system perimeter regardless of jurisdiction. ESG reporting obligations and the disclosure regimes that depend on them apply across the portfolio. The Valva compliance profile model is policy-driven, which means adding a US framework is a configuration exercise, not a re-architecture. The audit trail is the same artifact a BSEE or PHMSA investigator, an insurer, or an internal HSE team would reach for.

6.4 Canada, the United Kingdom, the Middle East, and Asia-Pacific

Canadian operators carry CER and AER expectations on pipeline and upstream safety and on environmental performance. UK operators carry HSE and NSTA expectations, plus the UK GDPR and the post-Brexit regulatory inheritance that aligns closely with the EU baseline. National oil companies in the Middle East and parts of Asia-Pacific carry national sovereignty over operational data and asset-specific safety regimes that resemble the NCS posture in spirit even where the specifics differ. Across all of these, the Valva architecture asks the same questions in the same way. Which data classes are

touchable, which destinations are reachable, which transformations are required, and which evidence must be captured. Adding a new jurisdiction profile is a configuration step.

6.5 The audit trail as evidence base

The audit trail is not a log. It is a structured, signed, append-only evidence base built so the same artifact can be produced for a regulator, an insurer, an internal incident investigator, or a downstream acquirer. Each entry carries the operator identity, the active regulatory profile, the inputs and outputs of the call, the policy decision, and a cryptographic chain back to a known root. In an incident reconstruction the operator can show, second by second, what tools the agent called, why each call was permitted, what the agent did with the result, and which human signed off on the policy that allowed it.

7. Deployment Patterns

Valva supports four deployment patterns. Each is appropriate for a different commercial relationship.

7.1 Operator-direct

The operator deploys Valva inside its own perimeter, federates its existing vendor and internal systems through it, applies the appropriate regulatory profile, and runs its own in-perimeter agentic workflows. This is the canonical deployment and the one that delivers the full sovereignty story without depending on any vendor relationship. It is the right shape for operators whose digital strategy treats the AI orchestration layer as core infrastructure rather than as a vendor product.

7.2 Embed

A large service company, OEM, or analytics vendor embeds Valva as the orchestration, compliance, and sovereignty layer inside its own product. The vendor's customers, which are typically operators, get Valva-mediated AI capabilities without having to integrate Valva separately. The vendor extends the AI surface of its product without taking on the full burden of building a federated gateway, a policy engine, and an audit trail from scratch, and the vendor's product becomes deployable into the strictest regulatory perimeters its customers operate in. This pattern is the right shape for vendors with their own brand to defend and their own product roadmap to protect.

7.3 Integrate

A vendor's product is exposed as a federated tool through Valva, so any Valva-equipped operator can call that product's pipeline alongside other operator systems with no per-vendor integration work. This is the right pattern for vendors who want to be reachable inside an operator's AI fabric without giving up the boundaries of their own product. It is also the right pattern for mid-tier vendors who want their product reachable from the AI orchestration layer their customers are already standing up.

7.4 Build on

Operator, vendor, and Novus Nexum Labs co-develop a deployment profile specific to a domain (drilling, completion, production, midstream, refining) with a defined IP split agreed in writing. The deployment

profile is then jointly marketable across operators who want a turn-key compliant AI stack for that domain. This pattern is the right shape for mid-tier oilfield digital and edge-AI vendors who want federation infrastructure underneath their own product, and for service companies whose strategy is to ship a domain-specific compliant AI offering that they can sell into operators without re-building the platform piece for each customer.

8. Capabilities Today and Roadmap

8.1 Today

Federated MCP gateway with semantic-index tool selection. Compliance gate with profiles in production for HIPAA, HITRUST, ISO 20022, GDPR, the EU AI Act, and DORA. Append-only signed audit trail. Provider-agnostic agent surface compatible with frontier and open-weights models. A reference computation platform, Functo, exposing 717 tools and 732 validated scientific models across 29 namespaces in a single MCP server, demonstrating the semantic index at scale. An AI development environment in active development, shaped like a regulated-industry IDE, that enforces regulatory profiles at the editor surface so regulated engineering teams can build without standing up their own compliant environment first.

8.2 Oilfield roadmap

ATEX and IECEx zone awareness for on-rig and in-plant packaging and runtime behavior. Starlink Maritime wire-protocol tuning to minimize offshore-link cost. On-rig and edge-node runtime hardening for unattended operation in unmanaged environments. Additional regulatory profiles for API specifications, ISO 29001, IEC 62443, and ESG reporting. Reference integrations for the common surface and downhole vendors operators already run, and for the common production, midstream, and refining systems those operators depend on. A library of operator-codified agentic patterns for the highest-impact workflows across drilling, completion, production, and midstream and refining operations.

8.3 Cross-industry continuity

The same gateway, policy engine, and audit trail that serve oilfield serve healthcare, financial services, and any other regulated industry where the same architectural problems repeat. Operators benefit from the engineering investment of every other industry on Valva, and the patterns Valva proves in one industry harden the platform for the next.

9. About Novus Nexum Labs

Novus Nexum Labs builds infrastructure for AI in regulated industries. Valva is the company's primary platform. The company is pre-revenue and patent pending on the federated-gateway, semantic-index, compliance-gate, and in-perimeter agentic runtime architecture described in this paper.

Daniel Hill is cofounder and inventor of Valva. Jan B. Tawakol, MD is cofounder and CEO of Plessen Healthcare, the practicing medical group whose operations informed the regulated-industry-first design philosophy that shapes the platform. The combination, an inventor focused on the orchestration and

sovereignty architecture together with an operating physician executive focused on regulated-industry operations and dispatch, is reflected in how the platform handles both the technical and the operational sides of a regulated AI deployment.

The company is headquartered in the United States Virgin Islands and operates with US, EU, and offshore-deployment customers in mind.

10. Engagement

For operators, service companies, OEMs, and analytics vendors considering Valva, the engagement model is staged so that the first step carries no obligation.

10.1 Architecture review

One hour, mutual NDA, our team and yours. We walk through your current AI architecture, the systems that would federate through Valva, the regulatory profile that applies, and the agentic workflows that matter most. Output is a shared one-page sketch of a Valva-mediated deployment. If either side decides there is nothing here, both sides walk away with the sketch and no further commitment.

10.2 Pilot

A scoped deployment on a defined perimeter (one rig, one production unit, one pipeline segment, one process unit) with one or two in-perimeter agentic workflows running through the federation. The pilot proves the gateway, the policy enforcement, the audit posture, and the sovereignty model in the partner's environment. Cost-sharing and IP boundaries are agreed in writing before the pilot begins.

10.3 Productized partnership

With pilot results in hand, the partnership takes one of the four deployment patterns in section 7. The choice is informed by the partner's commercial model and by which IP and operational boundaries the deployment most naturally respects.

10.4 Co-marketed deployment

For partnerships that produce a jointly marketable deployment profile, Novus Nexum Labs and the partner take the profile to the partner's target customer set with a unified message, a documented compliance posture, and a defensible reference architecture.

Contact

I am the right person on our side for Phase 1.

Daniel Hill Cofounder and inventor, Valva Novus Nexum Labs Daniel@NovusNexumLabs.com