




EOR Studio

One platform. Every recovery process.

How a reservoir simulator for enhanced oil recovery was built, from A to Z

Architecture, physics, numerics, software and proof — told without equations

Hamza Eldaw

Author and developer · with Claude AI assistance (Anthropic)

October 2026 · companion document to the EOR Studio product-tour video

01 — AT A GLANCE

A complete EOR simulator, built and proven in the open

EOR Studio is a reservoir-simulation platform for enhanced oil recovery. It brings together, in one fully implicit simulator and one desktop application, the processes that are usually spread over several commercial tools: waterflooding, black oil, compositional gas and CO₂ injection, CO₂ storage, thermal steam, chemical flooding (polymer, surfactant, alkali, ASP), geochemistry and microbial EOR. This document tells how it was built — the decisions, the architecture, the method and the proof — in plain language, with no equations.

64,000+ lines of C++ in the simulator, tools and desktop application (216 source files)	8 physics families in one fully implicit framework	11 ready-made recovery-process templates in the Model Designer	SPE 1–11 comparative solution projects rebuilt as studies and compared
138 design decisions logged with evidence (D-register)	200+ modelling assumptions recorded and referenced (A-register)	59 regression runs, re-checked before every installed build	851 pages of manual in four volumes, written alongside the code

Development began on **3 October 2026**. The core engine, all eight physics families, the desktop application, the history-matching engine, the import engines, the study language and the SPE verification suite described here were built between 3 and 10 October 2026, on top of the author's earlier, verified research codes, and with Claude (Anthropic's AI model) as a programming and research assistant.

“Every claim needs evidence.” The project uses three words carefully: **verified** means the same input deck gives the same answer as ECLIPSE or OPM Flow; **benchmarked** means the result falls among the published participants of an SPE comparison; **validated** is reserved for field or laboratory data.

Enhanced oil recovery projects rarely involve only one mechanism. A chemical flood changes viscosity, interfacial tension and rock wettability at the same time; a CO₂ project is both a miscible displacement and a storage problem; a steam project couples heat, phase change and heavy-oil viscosity. In practice, engineers move between several simulators and file formats to model them.

Three requirements shaped every later decision:

- The screenshot displays the EOR Studio - Project Hub interface. On the left, the 'Projects' pane shows a tree view under 'EOR Demo - Showcase' containing a 'Simulator project'. The main area features a file explorer titled 'EOR Demo - Showcase E\AEOR Demo' with columns for Name, Size, Type, and Date Modified. It lists various files and folders, including '01 Chemical EOR - ASP', '02 Thermal - Cyclic Steam', '03 Norne Field - ECLIPSE import', and '04 SPE Benchmarks'. A tooltip over the 'Norme.sres.h5' file reads: 'Run a deck with eorsim. Drop a deck here to submit it.' Below the file explorer is a 'Simulation jobs' table with columns: ID, Job, Type, Status, Progress, Sim. time [y], Elapsed, and Output folder. At the bottom are buttons for Start, Stop, Open results, Report, Open folder, Log file, Remove, and Clear finished. On the right, an 'Applications' section contains icons and labels for Model Designer, Deck Editor, 3D Viewer, Simulator, Report Generator, Post-processing Job, and History Matching.

Hamza Eldaw · with Claude AI assistance

03 — FOUNDATIONS

Starting from verified research code

EOR Studio did not start from a blank page. The author had three research codes from earlier work, each strong in one area:

Earlier code	What it contributed	How it was used
lspf3d	Polymer flooding (viscosity, adsorption, inaccessible pore volume, residual resistance, shear), ion exchange and geochemistry validated against published experiments, wettability alteration	Laws re-checked against OPM and MRST, then ported into the new general core
CCS_SPE11	CO ₂ -storage physics (accurate equations of state for water and CO ₂ , CO ₂ -brine solubility), linear solvers (CPR, algebraic multigrid, FGMRES), HDF5 output, a Qt/VTK results viewer	Solver and viewer adopted as the base of the numerical core and the 3-D Viewer
meor_sim	Microbial EOR kinetics (growth, metabolite production, interfacial-tension reduction, biopugging) with 57 unit checks	Ported as the MEOR module and verified against its source model and an independent ODE solution

The rule was strict: before any piece of physics was written, these codes were searched first; anything found was cross-checked numerically against OPM, MRST or CoolProp, and only then ported — or rejected, with the reason recorded. What the old codes lacked (wells with a bottom-hole-pressure unknown, general multi-component switching, a three-phase framework, multi-threading) became the specification of the new core.

The knowledge base

On the first day a working library of references was assembled next to the code: the source of OPM Flow and open-DARTS (read-only), MRST, CoolProp, the opm-tests reference decks with their ECLIPSE results, the original SPE comparative-solution papers, and the keyword manuals of the commercial thermal and compositional simulators so that the input dialect would feel familiar to their users. Ten living documents were started at the same time and kept up to date ever since: progress log, design decisions, assumptions, references, environment, reuse inventory, provenance, cross-checks, validation and source layout.

04 — THE METHOD

How every feature was built

Each capability followed the same loop. It is slow on paper and fast in practice, because problems are found while they are small:

Step	What happens
1	Read the references first
2	Design and record
3	Re-implement
4	Verify in 1-D, 2-D and 3-D
5	Regress
6	Document

Fixes follow the same rule: they must be general and physically justified, never tuned to make one case pass. A good example came from the SPE 11B CO₂-storage benchmark: runs stalled at gas-water contacts. Reading how OPM evaluates fluid density between two cells led to a general change (a continuous two-cell average, decision D-147). It removed the root cause of the stalls and, being general, applies to every CO₂-brine study, not only SPE 11B.




Figure 2. The Model Designer after a full check by the simulator's own deck reader (bottom panel: "the deck is valid"). The same reader is used by the simulator, the Designer and every importer, so a model that passes here will run.

05 — ARCHITECTURE

How the pieces fit together

The platform is organised in layers. Inputs of any origin are turned into one internal description of the study; that description is solved by one of eight physics families sharing a single numerical core; results are written once, in a standard HDF5 file, and read by every downstream tool.




Figure 3. Layers of EOR Studio. The EOR Advisor (top right) is the next module in development and is not part of the current release.

Design choices that matter

- **One framework, many families.** Each physics family supplies its fluid description and its extra conservation laws; accumulation, flux, wells, Jacobian assembly and the Newton loop are shared. A new process inherits the solver, the wells and the output for free.
- **Compile-time sizing.** The number of equations per cell is fixed when the code is compiled (about 30 variants: black oil with 2 to 14 equations, compositional with 2 to 10, storage, steam, K-value thermal). This gives dense, cache-friendly blocks and exact derivatives.
- **A source split for speed of development.** On 7 October the original 2,760-line model class and its single 15-minute compilation unit were split, by scripts that moved function bodies byte for byte, into topic files (state, flux, thermal, chemical, geochemistry, MEOR, wells, assembly) and parallel runner libraries. The result was checked bit-identical against the pre-split binary.
- **One reader for everyone.** The deck reader, the unit converter and the checks are shared by the simulator, the Model Designer, the Deck Editor and the importers.

06 — FROM A TO Z

The build sequence

The order of work was set by the author: waterflood first, then chemical EOR, CO₂, thermal, geochemistry and MEOR, then the application, then import, history matching and the proof campaign.

Date (2026)	Milestone
3 Oct	Knowledge base and reference tools. Core engine : grids (Cartesian, variable, corner point), rock, automatic differentiation, Newton, time stepping, wells with a pressure unknown, black-oil PVT, three-phase relative permeability, keyword reader, HDF5/CSV output. First verification: Buckley–Leverett, SPE 1, SPE 9.
3 Oct	Chemical EOR : polymer and surfactant verified in 1-D, 2-D and 3-D against OPM Flow and MRST; alkali and ASP verified through exact equivalence with MRST and conservation checks. Compositional : Peng–Robinson equation of state, flash and stability.
3–4 Oct	Thermal : hot water and thermal black oil; steam injection, cyclic steam stimulation and heat loss; radial grids checked against exact solutions; K-value thermal model verified against DuMux.
4 Oct	Geochemistry : local equilibrium identical to PHREEQC; reactive transport, kinetics and low-salinity polymer in carbonates. EOR Studio : Model Designer for every process, Project Hub and jobs, Deck Editor, 3-D Viewer and post-processor. MEOR verified.
6 Oct	EOR Match history-matching and uncertainty engine; installer; geology import of grids, ECLIPSE well schedules, trajectories and RESQML; field-scale low-salinity demonstration with 7 wells.
7 Oct	Group control, VFP/THP tables, multi-segment wells; EORL study language and its C++ port; field case set-up; the source split for fast parallel builds.
8 Oct	Deck import engine for ECLIPSE and CMG-format decks; native compositional decks; UNISIM-IV field benchmark; linear-solver speed-up and Solver Guide; professional 3-D rendering.
9 Oct	SPE Validation Suite : every SPE comparative solution project rebuilt as an EORL study; general fixes D-141 to D-148 found by the suite; adaptive CPR weights (SPE 10 run time 86.7 s → 4.5 s); Studio hardening.
10 Oct	Run Health : step log, health report and settings so that users solve convergence problems without code; project handover; this document and the product-tour video.

07 — PHYSICS

What the simulator models

Each family is a complete, fully implicit model. The right-hand column names the independent reference it was checked against — the evidence is in the verification volume of the manual.

Family / feature	What it represents	Checked against
Black oil	Water, oil, gas with dissolved gas and vaporised oil; Stone and Baker three-phase relative permeability	SPE 1, SPE 9 (ECLIPSE results in opm-tests), Buckley–Leverett
Waterflood	Gravity, capillary pressure, equilibration, rate and pressure controlled wells	OPM Flow cross-sections
Polymer	Viscosity mixing, adsorption, inaccessible pore volume, permeability reduction, shear thinning	OPM Flow 1-D/2-D/3-D decks
Surfactant, alkali, ASP	Interfacial tension, capillary-number desaturation, adsorption, in-situ soap, salinity	MRST 1-D, 2-D, 3-D; alkali by exact equivalence with MRST
EOS compositional	Peng–Robinson / SRK, robust flash near the critical point, miscibility effects	MRST (with ECLIPSE 300 and AD-GPRS), SPE 5
CO ₂ storage	CO ₂ -brine solubility and densities, diffusion, thermal effects	OPM Flow CO2STORE, exact diffusion solutions, SPE 11
Thermal steam	Water/steam properties to international standard, heat loss to over- and underburden, radial grids	Exact conduction and radial solutions, SPE 4
K-value thermal	Multi-component thermal model in the style of commercial thermal simulators	DuMux
Geochemistry	Aqueous speciation, ion exchange, minerals, kinetics, wettability and rock alteration	PHREEQC (batch and transport)
MEOR	Microbial growth, metabolites, interfacial-tension reduction, bioplugging	Source model and an independent ODE solution
Dual porosity / permeability	Matrix–fracture transfer with gravity drainage	SPE 6 (ECLIPSE)
Wells	Multi-segment wells, VFP tables and THP control, group control, schedules and trajectories	OPM Flow and ECLIPSE reference decks
Geomechanics	Pore-volume coupling to stress changes	Exact solutions; Geertsma subsidence (3.3 %)

08 — NUMERICS

Solving it fast and robustly

All families share one solution strategy. Every equation in every cell is solved simultaneously at each time step (fully implicit), which keeps large time steps stable even with strong coupling such as steam condensation or polymer viscosity.

- **Exact derivatives.** Automatic differentiation computes the Jacobian of every property and equation exactly; there is no hand-coded derivative to get wrong. A finite-difference check is built in for testing.
- **Safeguarded Newton.** Variable switching for appearing and disappearing phases, update limits, a minimum number of Newton updates where the physics needs it, and recovery from failed steps by cutting the step instead of stopping the run.
- **Constrained-pressure-residual (CPR) preconditioning.** The pressure part of the system is solved with algebraic multigrid and the whole system with FGMRES, multi-threaded. In October 2026 an **adaptive weight switch** (D-139c) was added: if the default weights fail, the solver switches once to true-IMPES weights. On SPE 10 Model 1 (200 days) this cut the run from 86.7 s to 4.5 s and the linear iterations from 51,600 to 3,342; regression stayed within 2 % on all 215 metrics.
- **Adaptive time stepping** driven by Newton effort and solution change, with OPM-style controls available as settings.
- **Run Health.** Every time step can be logged with its Newton and linear effort; a health report points to the cause of difficult steps and to the setting that fixes it.

The screenshot displays the EOR Studio Project Hub interface. The main window shows a file explorer for the 'EOR Demo - Showcase' project, listing various folders and files. The 'Simulation jobs' table at the bottom shows a job named 'SPE09' in a 'Running' state with 34% progress. The 'Job log' window on the right displays the solver output, including time steps, Newton iterations, and a message indicating a step failure due to a small time step cut.

ID	Job	Type	Status	Progress	Sim. time [y]	Elapsed	Output folder
1	SPE09	Simulation	Running 34 %	34 %	0.8	00:00:13	E:\EOR Suite\output\SPE09

Job log - 1 (SPE09): Running 34 %

```

t = 310.00000 d: step failed, cutting dt to 1.347 d
t = 311.34661 d: step failed, cutting dt to 0.6194 d
t = 311.34661 d: step failed, cutting dt to 0.1546 d
  
```

Figure 4. A real run of the SPE 9 benchmark (26 wells, 8 threads) launched from the Project Hub: the job table shows progress and simulated time, the job log streams the solver output live.

09 — THE APPLICATION

EOR Studio, from model to result

The desktop application is written in Qt with VTK for 3-D graphics. It is built around the engineer's workflow rather than around the simulator's files.

- **Project Hub** — projects, files and a job scheduler; seven application tiles.
- **Model Designer** — form-based model building, section by section, with a live 3-D preview and a full check by the simulator's own reader. Eleven templates: waterflood, black oil, polymer, ASP, hot water, steam, CO₂ storage, compositional CO₂ EOR, dual porosity, low-salinity polymer in a carbonate with geochemistry, and steam with nitrogen in heavy and light oil.
- **Deck Editor** — keyword and EORL editing with outline, keyword help and live checks.
- **3-D Viewer** — properties through time, filters, slices, iso-surfaces, streamlines, wells, summary plots, well logs and histograms.
- **Report Generator and post-processing jobs** — HTML reports and scripted comparisons.
- **History Matching** — the EOR Match engine with study set-up, parameters, observed data, workflow and results pages.




Figure 5. New model: eleven ready-made recovery processes.




Figure 6. Chemical EOR page of the ASP template: polymer and surfactant properties.




Figure 7. The 3-D Viewer with the Norne field grid (46 × 112 × 22) imported from its ECLIPSE files: oil saturation at a report step, with the field oil-rate plot below.

Cyclic steam stimulation alternates injection of steam, a soak period and production in the same well. The example below is a heavy-oil model with SPE 4-type data ($15 \times 15 \times 4$ cells, three one-year cycles). The model is cut through the wellbore so the steam chamber is visible on the section.




Figure 9. Later in the second cycle: the heated zone has spread and cooled; oil-rate peaks below mark each cycle.

- Water and steam properties to the international industrial standard, with phase change in every cell.
- Heavy-oil viscosity as a function of temperature; heat conduction in rock and fluids.
- Heat loss to the over- and underburden.
- Cycle scheduling of injection, soak and production on the same well.

11 — LANGUAGE AND IMPORT

EORL, and reading the decks engineers already have

EORL is EOR Studio's own study language: a readable description of a study — geology, fluids, rock-fluid data, initial state, wells and schedule — that reads like a report. Any native deck can be written as EORL and translated back, keyword for keyword and number for number; this exact round trip is checked automatically, and the C++ translator matches the reference implementation on all 360 parity tests.




Figure 10. An ASP pilot study in EORL in the Deck Editor, with the outline on the left, keyword reference on the right and the live translation check below.

Import engines

- **ECLIPSE / OPM decks** — grids, properties, equilibration, schedules, VFP tables, multi-segment wells and group controls are converted into EOR Studio studies.
- **CMG-format decks** (thermal, compositional and black-oil dialects) — read natively; in the GUI check, 53 of 59 example decks round-trip through the Model Designer, the remaining 6 give a plain message.
- **Geology** — corner-point grids, faults and transmissibility multipliers, well trajectories and RESQML 2.0.1 models, verified against ECLIPSE and OPM.

12 — HISTORY MATCHING

EOR Match: a real run, start to finish

EOR Match runs the simulator as a black box on a template deck with placeholders, and combines established methods into one workflow: one-at-a-time and Morris screening, Latin-hypercube and Sobol sensitivity, a Gaussian-process proxy with expected-improvement search, Levenberg–Marquardt refinement with posterior covariance, global optimisers (CMA-ES, differential evolution, particle swarm) and the ensemble smoother with multiple data assimilation. Every run is cached, so a study can be resumed or extended.

For the video, a complete study was run live from EOR Studio: a 3-D waterflood “twin” experiment ($15 \times 15 \times 3$ cells, 5 wells, 7 uncertain parameters including a choice between two relative-permeability sets, 9 observed series).

404

simulations, 0 failed (406 evaluations)

0.99

best objective, from 198.6 for the base case

99.5 %

misfit reduction

9.5 min

wall time, 8 runs in parallel




Figure 11. History-match plots from the EOR Match report: observed data (dots), base case (dashed), best match (blue) and the posterior P10–P90 band, with the forecast beyond the end of history.

In the verification of the same twin experiment, every true parameter value fell inside the posterior P10–P90 range, and the correct relative-permeability set was recovered by every ensemble member.

13 — PROOF

Verification and benchmarking

Every SPE comparative solution project (SPE 1 to SPE 11) is rebuilt as its own EORL study in the Validation Suite and compared with ECLIPSE, OPM Flow or the published participants. The table shows the status honestly, including what is still open.

Case	Topic	Result
SPE 1	Gas injection, black oil	Verified: within 0.02–0.2 % of ECLIPSE and OPM
SPE 2	Three-phase coning	Benchmarked: 6 of 7 quantities inside the 11 participants
SPE 3	Gas cycling	Black-oil version verified (< 0.1 % vs ECLIPSE); 9-component EOS version open
SPE 4	Steam injection	Case A complete (6 of 6); case B open
SPE 5	Miscible WAG	Benchmarked: scenarios 1 and 2 inside the participants; scenario 3 within 0.5–2.6 %
SPE 6	Dual porosity	Verified: 0.4–1.9 % vs ECLIPSE; gas-injection variant open
SPE 7	Horizontal wells	Benchmarked: 14 of 14 inside the 14 participants
SPE 8	Gridding	Verified: 0.2 % vs OPM
SPE 9	Heterogeneous, 26 wells	Verified: 0.05–0.5 % vs ECLIPSE
SPE 10	Upscaling, large models	Partial: Model 1 short runs; full run and Model 2 open
SPE 11B	CO ₂ storage, 1,000 years	Benchmarked: 0 failed steps; in the 34-group band for most quantities
SPE 11A / C	CO ₂ storage, lab / 3-D	11A partial (finer grid needed); 11C open




Figure 12. SPE 1: field oil rate, EOR Studio (red) on top of ECLIPSE and OPM Flow.




Figure 13. SPE 9: field oil rate, EOR Studio (red) on top of ECLIPSE (black).

The suite is also a bug detector: running it found eight general improvements (D-141 to D-148) — among them the starting variable of water-only cells, dead oil with free gas, robust flash near the critical point, continuous face density for CO₂-brine and interfacial-tension relative permeability — each fixed for every case, regression-checked, and logged with its evidence.

14 — PRODUCT ENGINEERING

From code to something people can use

- **Documentation** — an 851-page manual in four volumes: Technical Description (214 pages), User Guide and Tutorials (299), Keyword Reference (255) and Verification and Validation (93).
- **Installer** — a Windows installer with examples copied to the user's documents on first start.
- **Regression discipline** — a 59-run targeted regression with 215 metrics guards every installed build; earlier binaries are kept so any change can be rolled back.
- **Gentle, reproducible builds** — Ninja builds of the split source tree; the shared source must always compile, because several work streams build from it.
- **Naming and licence hygiene** — the interface uses EOR Studio's own names (Project Hub, Model Designer, EOR Match); commercial product names appear only as file formats or comparison references; every borrowed component is listed in a provenance record with its licence.




Figure 14. EOR Match results page inside EOR Studio after the live run: best value, Levenberg–Marquardt estimate and ES-MDA posterior percentiles for every parameter.

15 — WORKING WITH AI

How Claude assisted, and what stayed human

EOR Studio was developed by Hamza Eldaw with Claude, Anthropic's AI model, working as a pair programmer and research assistant over many working sessions. The division of work was deliberate:

The author decided	Claude assisted with
The goal, scope and order of work; which processes matter and why	Reading reference codes (OPM, open-DARTS, MRST), papers and keyword manuals, and summarising how they solve a problem
The rules of evidence, wording and licence	Writing and refactoring C++, Python tools and the Qt application, and keeping the shared source compiling
Which fixes were acceptable and when a build could be installed	Running verification and regression, comparing results, and reporting differences with numbers
The field data, the domain judgement and the final sign-off	Writing the design and assumption registers, the manual chapters, this document and the product video

The working agreement was simple: *the AI proposes, the simulator proves, and the engineer decides*. Every number in this document comes from a run that was executed and logged, not from a model's memory.

The same principle shapes the next module, the **EOR Advisor**: a deterministic physics rule engine produces every finding about a deck or a run; AI only selects the checks, explains the findings in plain language and proposes changes, which the engineer approves before anything is re-run.

16 — WHAT COMES NEXT

Open items, stated plainly

- Complete the SPE campaign: SPE 10 Model 1 full run and Model 2 (1.1 million cells), SPE 11A on a finer grid, SPE 11C in 3-D, SPE 4B steamflood, the 9-component SPE 3 and the SPE 6 gas-injection variant.
- Field validation — the real proof of maturity — on heavy-oil steam-soak, waterflood and CO₂-WAG field cases with production history.
- Mechanical dispersion for CO₂-storage studies.
- Install the Run Health and adaptive-solver updates after their regression, and release the EOR Advisor.
- Close the remaining import gaps (ECLIPSE radial grids, some thermal end-point options).

About the video

The accompanying product-tour video was recorded from the real application on the author's workstation: every run shown — the SPE 9 simulation and the 404-run history match — was executed live during recording, and the narration is placed on the moment each event happens on screen.

Hamza Eldaw — author and developer of EOR Studio.

Built with Claude AI assistance (Anthropic). Feedback and collaboration enquiries are welcome through LinkedIn.



EOR Studio

Design, simulate, and prove your recovery strategy.

Hamza Eldaw · author and developer
Built with Claude AI assistance (Anthropic)