

CUSTOMER STORY

Covering Every Process Domain, Without Double-Checking the AI

How Oran Zucker, head of a chemical and process engineering division, moved from folder archives and supervised chatbot answers to organized, reliable technical research with REACTOR.

Background



Oran Zucker heads the chemical and process engineering division at a small industrial engineering firm that connects manufacturers with specialized suppliers of process equipment and production lines. His remit covers everything process- and chemistry-related that crosses the firm's desk: from bench-scale columns, mixers, and reactors to the specialized production processes of the manufacturers the firm represents.

In a firm that size, technical breadth concentrates in one person. Each new engagement can mean a different chemistry, a different unit operation, a different set of standards — and the division head is the technical filter for all of it. Oran started using REACTOR in 2026 and now works with it weekly.

The Challenge

The job runs on technical knowledge that keeps changing shape. One week the question is a mixing process; the next, a separation column; the next, a specialized process a represented manufacturer needs evaluated. Staying current meant reading heavily — papers, technical news, supplier documentation — and Oran's system for it was manual: read everything himself, file it in folders, and rely on memory to find it again. The knowledge was organized around folders, not around questions.

To speed things up he added a general-purpose AI assistant, ChatGPT, to the workflow. It helped — but at a cost he describes plainly as supervision work. Outputs had to be checked, corrected, and re-verified before they could be used. For casual questions that is tolerable. For process engineering — where an answer feeds real decisions about real equipment and chemistry — an assistant that needs auditing on every technical answer caps how much time it actually returns.

What he needed was not another chat window. It was a knowledge system with a scientific base: something that could gather, order, and output information he could rely on, and support the data analysis and document generation his division produces.

WHAT WAS NEEDED

A knowledge system with a scientific base — one that could gather, order, and output technical information, and support data analysis and document generation, without a supervision layer on every answer.

“REACTOR makes it easier to organize, access, and build research data to perform good and efficient research.”

Oran Zucker, Head of Chemical & Process Engineering Division

How REACTOR Fit In

REACTOR became Oran’s weekly entry point for technical topics: learning a new domain, and gathering and understanding the data around a process or product family. The difference he reports against his previous workflow is not a dramatic before-and-after number — it is the reliability of the technical layer. In his words: *“it is more reliable with technical knowledge.”*

That reliability is a design property, not an accident. REACTOR’s Chemical Intelligence Layer grounds answers in chemistry rather than general language patterns, and every output is referenced — summaries trace back to their sources, and uncertainty is flagged rather than smoothed over. For work like Oran’s, where the relevant information is as likely to sit in a supplier’s technical data sheet as in a journal, REACTOR’s scientific web search extends the same loop beyond the literature — each source saved into the Research Space as a numbered, re-queryable reference instead of a link that disappears.

The folder archive has a working replacement, too: Research Spaces hold project context, sources, and notes together, so research data is organized, accessible, and reusable — built up over time instead of re-assembled from scratch for each question.

THE SUPERVISION TAX

The hidden cost of general-purpose AI in technical work: every answer has to be checked against sources before it can be used, so the expert becomes an editor. The tax scales with how technical the question is — highest exactly where help is needed most. REACTOR is designed to remove that layer: chemistry-native reasoning, answers referenced to their sources, and uncertainty marked rather than papered over.

Unlike generic AI, REACTOR was designed with a chemical intelligence layer, built to reason with the nuances of R&D in mind.

WHAT CHANGED IN PRACTICE

Technical learning that used to be scattered across folders and double-checked chats now runs through one reliable system.

WEEKLY

RESEARCH CADENCE IN
REACTOR

**BENCH →
PLANT**

FROM LAB COLUMNS AND
MIXERS TO PRODUCTION
PROCESSES

ONE

ORGANIZED KNOWLEDGE BASE,
IN PLACE OF FOLDER ARCHIVES

The Outcome

Oran is direct about what changed — and what didn't. He doesn't claim a headline metric. What changed is the texture of the work: topics are easier to learn, data is easier to gather and understand, and research material lives in an organized, accessible base instead of a folder tree and a browser history.

The supervision layer his old workflow required has no equivalent in the new one. The reliability gap he flagged in general-purpose tools is the gap REACTOR was built to close: answers grounded in chemistry, traceable to source, usable in engineering work without a second shift of checking.

Asked whether he would recommend REACTOR, his answer is scoped the way an engineer's answer should be: it is most suitable where research and development is the actual work. For his division — evaluating processes, developing recommendations, ramping up on new technical domains — that is precisely the job.

“With general AI tools I had to supervise every answer before I could use it. With REACTOR, the technical knowledge is reliable — for engineering work, that is the difference.”

Oran Zucker, Head of Chemical & Process Engineering Division

[draft quote — pending Oran's approval]

What This Means for Small Industrial Teams

Small industrial firms run wide technical scope through very few people. The binding constraint is rarely expertise — it is the hours the expert has, and how many of them get consumed becoming current in the next domain. General-purpose AI looks like relief, and in Oran's experience it helps — until the questions turn technical, and supervision costs quietly hand back the time the tool saved.

A chemistry-native system changes what one technical lead can cover. Learning a new domain becomes structured research instead of folder archaeology. Answers arrive with references instead of

homework. And the knowledge assembled for one project stays organized and reusable for the next — giving smaller teams research capabilities previously available only to large organizations.

RELIABLE TECHNICAL GROUND

Chemistry-native answers, referenced to source.

The supervision layer general AI requires for technical work is exactly what REACTOR was designed to remove.

FASTER INTO NEW DOMAINS

Learning a topic becomes structured research.

For a division head whose next project can be a new chemistry, ramp-up is the recurring cost REACTOR reduces.

BUILT FOR R&D WORK

“More suitable for research and development.”

Oran’s recommendation is scoped like an engineer’s: REACTOR fits where research and development is the actual work.

AT A GLANCE

R&D Leader: Oran Zucker, Head of Chemical & Process Engineering Division (industrial engineering firm, name withheld)

Field: Chemical & Process Engineering

Scope: Bench-scale columns, mixers, and reactors; specialized production processes for represented manufacturers

Team: Small firm; technical breadth concentrated in one division

Prior workflow: Manual paper and news review filed in folders; general-purpose AI with heavy supervision

Active in REACTOR: 2026, weekly use

REACTOR CAPABILITIES ENGAGED

- Chemistry-native research and technical Q&A
- Research Spaces for organizing project knowledge
- Literature discovery and analysis for new-domain ramp-up
- Referenced answers — summaries traceable to source
- Research data gathering, analysis, and document generation

REACTOR by Elementics.AI

The R&D workbench for chemical and materials science. Built by scientists, designed for real R&D workflows. reactor.elementics.ai

PRE-PUBLICATION NOTES — INTERNAL — REMOVE BEFORE PUBLISHING

1. **Quote status.** Pull quote (“REACTOR makes it easier...”) is Oran’s own questionnaire answer (Q14), lightly edited for grammar (original: “Reactor make it easier to organize, access and build research data to perform good and efficient research”) — confirm exact wording with him. “It is more reliable with technical knowledge” (body text) is verbatim from Q10. The closing quote is a DRAFT generated from his Q6/Q10 answers and requires his explicit approval before publication.
2. **Anonymity.** Per Q15: name approved, company NOT approved. Company described generically; his more identifying self-description and any geographic identifiers were deliberately avoided. Verify with Oran that the descriptor does not identify the firm. No company logo (Q16); he can supply a headshot — layout slot next to the pull quote is the suggested placement.
3. **Excluded content.** The confidential item Oran listed under Q17 (do-not-mention) appears nowhere in this file — including the aside in his Q13 answer. It is deliberately not named here either.
4. **Title discrepancy.** Q1 says “Head of chemical and process engineering division”; his At-a-Glance confirmation says “Head of chemical and engineering division.” This draft uses the fuller Q1 version — confirm which is correct.
5. **No invented metrics.** Oran answered “not much actually” on measurable change (Q11). The metrics band therefore uses factual, qualitative counters only (weekly cadence, scope, one knowledge base). Do not add time-saved or percentage figures unless he supplies and confirms them.
6. **Open gaps.** Q8 (which parts of REACTOR he uses) and Q9 (one concrete example) were left blank. A 10-minute follow-up would enable a “WHAT REACTOR RETURNED” featured sidebar like the Avrutsky study’s strongest element. Suggested questions are in the companion document: Case-Study-Oran-Zucker-Secondary-Angles.