Proposal vision for VIR

VIR has built its reputation through specialist expertise, trusted client relationships, and the consistent delivery of high-quality work. This proposal is not about replacing that judgement with automation or tools. It is about strengthening the way VIR captures, structures, reuses, and applies its knowledge across projects.

The opportunity is to create a more reliable operating foundation for knowledge work. Today, valuable project history, assumptions, evidence, methods, and commercial insight sit across folders, documents, and individual experience. If that information is made easier to find, govern, and reuse, VIR can improve delivery quality, reduce duplicated effort, support stronger proposals, and identify new commercial opportunities.

Technology can support this, but only once the underlying work is properly understood. Automations, AI tools, and agents should not be built on messy or unstable information. The first priority is to get the data, workflows, and ownership model into a usable shape, with strong stakeholder input and adoption from the start.

This proposal is built around three practical areas:

1. ***VIR Operating Memory*** - Turning the current SharePoint project archive into structured, searchable project intelligence.
2. **HEOR Commercial and Delivery Workbench** - Supporting proposals, evidence work, assumptions, models, reports, and quality control through better reuse of knowledge and project materials.
3. ***CRO Growth Wedge*** - Exploring targeted support for trial start-up and submission workflows where VIR can create a focused commercial advantage.

The goal is not to outsource thinking.

VIR's expert judgement remains central. The goal is to give that judgement a better operating base: cleaner information, clearer workflows, faster access to prior work, and systems that help people make better decisions without losing the nuance of the work.

SA Office: +27 12 003 6624
Address: Byls Bridge Office Park. C/O Olivenhoutbosch & Jean Ave, Building 14, Block B, Centurion, Gauteng
www.valueinresearch.com | value@valueinresearch.com
VI Research (PTY) LTD | Registration Number : 2006/030051/07

VIR Operating Memory

Purpose

Turn VI's historical SharePoint project archive into usable business intelligence, while creating a future-proof way for every new project to be structured and searchable from day one, with a route to agent-ready workflows once the core actions are proven.

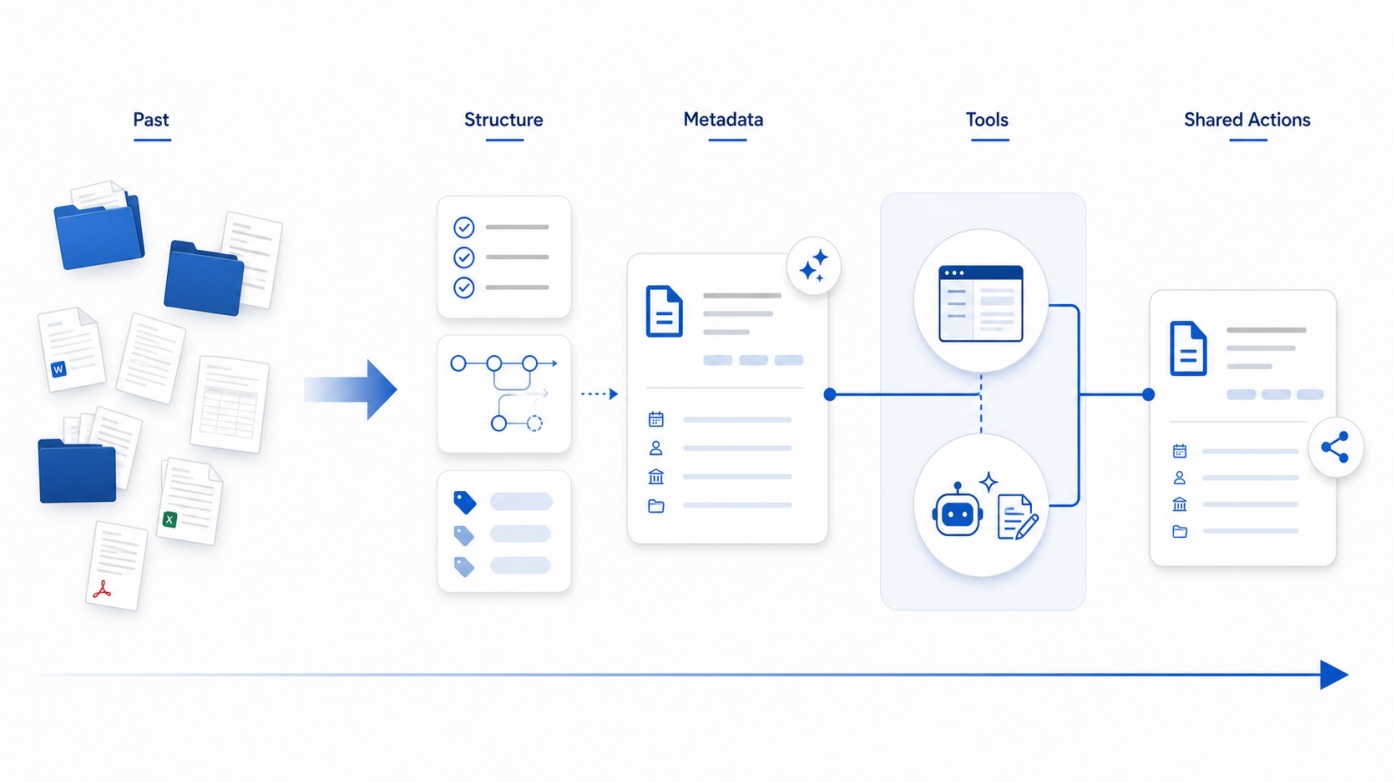

Projects to Build

These projects move from foundation work to later candidate tools. The table below stays summary-level; the full project definitions will sit in the separate VI Project Intelligence Plan.

Project	Thesis	Output	Phase
Project 1: SharePoint Inventory Map	VI needs to see what already exists before building tools on top of the archive.	SharePoint inventory and pattern summary.	**Phase 1**
Project 2: Historical Project Classifier	Old folders contain enough signals to draft useful project records, but uncertainty must stay visible.	Draft historical catalogue with confidence, gaps, and review needs.	**Phase 1**
Project 3: Project Metadata Model	VI needs one shared language for old and new project knowledge.	Metadata model, data dictionary, validation rules, and source-link conventions.	**Phase 1**
Project 4: Historical Project Backfill Pipeline	Historical work must enter the intelligence layer, not remain loose examples.	Backfilled project records linked to original SharePoint artifacts.	**Phase 1**
Project 5: Project Registry for New Work	New work should start as structured project intelligence from day one.	Project registry record linked to SharePoint folders, deliverables, and owners.	**Phase 2**

VI Research
Real-World Evidence, Real Outcomes

Project	Thesis	Output	Phase
Project 6: Project Setup Workflow	Folder setup and starter metadata should be repeatable instead of recreated manually.	Standard project setup workflow with registry record, folder structure, and templates.	**Phase 2**
Project 7: Document Registration and Drift Checker	Project folders need light governance so they do not drift back into chaos.	Project health report and suggested cleanup actions.	**Phase 2**
Project 8: Commercial Automation Layer	Past work and structured metadata can make BD faster and more consistent once the foundation is proven.	Proposal, SOW, quote, timeline, and assumption support with source traceability.	**Later candidate**
Project 9: HEOR Delivery Workbench	Delivery tools should support evidence, assumptions, models, reports, and QC without replacing expert judgement.	Reusable evidence objects linked to projects, documents, models, and citations.	**Later candidate**
Project 10: Agent-Native Project Hub	Agents should operate through stable project actions, not scrape the UI or invent workflows.	Browser UI and MCP/Codex tools over the same project objects.	**Later candidate**

Reference note: The full hypotheses, outputs, phase boundaries, and detailed build sequence should be handled in the separate VI Project Intelligence Plan.

Supporting document: VI Project Intelligence Plan (vi_project_intelligence_plan.docx)

VI Research
Real-World Evidence, Real Outcomes

HEOR Commercial and Delivery Workbench

What would we do differently if every HEOR consultant had an unlimited bench of research assistants, analysts, reviewers, model checkers, writers, and project coordinators?

Purpose

Identify which parts of HEOR work are repeatable, labour-heavy, error-prone, or poorly captured, then build tools that make the expert team faster, cheaper, and more consistent.

Suggested First Experiments

The HEOR opportunity should start with bounded internal experiments that reduce production effort, improve reviewability, and preserve expert sign-off.

Experiment	Thesis	Output	Guardrail / Success Signal
Experiment 1: HEOR Intake + Scope Assistant	Rough client requests should become structured briefs before proposal and delivery work starts.	Intake brief, missing questions, scope risks, draft assumptions/exclusions, work packages, and proposal outline.	Low scientific risk; success is faster proposal drafting and cleaner sales-to-delivery handoff.
Experiment 2: Evidence Table Builder	Evidence extraction is labour-heavy and reusable, but outputs must stay reviewable and source-linked.	Draft evidence table with citations, page references, confidence score, reviewer status, and evidence gaps.	AI drafts extraction only; human approval is required for extracted values.
Experiment 3: Assumption Register	Assumptions are where HEOR expertise lives and should be captured as reusable project knowledge.	Assumption log with rationale, linked evidence, approval status, sensitivity scenarios, and reusable candidates.	Success is fewer undocumented assumptions and easier model/report review.
Experiment 4: Model-Report QC Assistant	Model outputs, assumptions, and report claims should be checked for consistency before final review.	Issue log, severity rating, model/report mismatches, unsupported claims, missing references, and anomalies.	Flags only; the tool must not silently change models, reports, or claims.
Experiment 5: Internal HEOR Memory	Past HEOR project knowledge should become searchable and reusable instead of trapped in files.	Searchable project memory, source library, assumption library, country library, model registry, and similar-project suggestions.	Success is less searching, more reuse, faster onboarding, and reduced duplication.

VI Research
Real-World Evidence, Real Outcomes

CRO Growth Wedge

Purpose

Turn protocol synopses into reviewable startup, submission, site activation, bilingual ICF, TMF, and reusable study-knowledge outputs for Saudi Arabia, Abu Dhabi, and Qatar.

The strongest wedge is not an AI CRO. It is trial startup, submission, site activation, TMF/QC, and bilingual document workflow support: boring, document-heavy, auditable work where small and mid-sized CROs, hospital research offices, and pharma affiliates can review every output before use.

Top 10 Opportunity Areas

The table below keeps the country and regulatory detail at summary level; the full CRO opportunity map should hold the evidence, risk notes, discovery questions, and country-specific constraints.

Opportunity	Buyer	Workflow / Tool Wedge	Prototype / Success Signal
1. Trial startup + submission pack copilot	Local CROs, pharma affiliates, hospital research offices	Protocol-to-checklist extraction, country pack builder, approval tracker, missing-doc alerts, bilingual cover-letter/ICF support.	Saudi + Abu Dhabi + Qatar pack; 30-50% faster startup pack prep and fewer missing-doc queries.
2. Site feasibility + investigator intelligence CRM	CRO BD/feasibility teams, pharma, site networks	Map protocol eligibility to site capability, investigator history, patient-pool questions, pharmacy/lab/EMR readiness, and startup timelines.	25-site Saudi/UAE/Qatar feasibility database; feasibility response in under 72 hours.
3. Bilingual ICF and patient-material QA	CRO regulatory teams, ethics offices, hospitals	Arabic/English concordance checker, plain-language review, risk/benefit consistency, glossary, change log, and reviewer workflow.	ICF QA workflow with comments; fewer IRB language queries and shorter resubmission cycles.
4. TMF/eTMF QC and audit-readiness assistant	CRO QA, sponsor oversight, hospital research offices	Map TMF folders to expected index, detect missing/expired/inconsistent documents, and produce audit-readiness reports.	TMF gap report for live or closed studies; higher critical-doc completeness and fewer inspection-prep hours.
5. Safety/deviation narrative + deadline tracker	CRO safety/PV, sponsor safety teams, PI offices	Structured SAE/deviation intake, narrative draft, source completeness check, SUSAR/DSUR clock, and CIOMS/XML checklist.	Narrative and timeline dashboard; on-time safety submissions and fewer regulator/sponsor follow-ups.
6. Recruitment prescreening workflow	Hospitals, site networks, CRO recruitment teams	Convert eligibility into human-review prescreen criteria, consent/referral tracking, bilingual outreach scripts, and site-governed worklists.	Synthetic-data prescreen dashboard; improved screened-to-enrolled ratio and faster first patient in.
7. RWE / observational study protocol + data dictionary builder	Pharma medical affairs, HEOR teams, CROs, DoH-facing research teams	Study-design assistant, variable dictionary, RWE checklist, data-completeness scoring, and CSR/manuscript shell.	Therapy-area RWE pack with protocol, SAP shell, CRF/data dictionary; faster SOW/protocol turnaround.
8. Monitoring visit report / CAPA copilot	CRO clinical ops, CRAs, study managers	Turn monitor notes into visit report drafts, action trackers, deviation checklists, CAPA status, and evidence links.	Visit report generator from CRA note template; shorter report cycle time and fewer overdue actions.
9. Proposal/SOW estimator and assumptions library	Small/mid CROs, regional CRO BD teams	Protocol synopsis to task list, country assumptions, pass-through fees, timelines, site count, risk flags, and change-order assumptions.	SOW generator for three study archetypes; faster proposal turnaround and fewer change orders.

Opportunity	Buyer	Workflow / Tool Wedge	Prototype / Success Signal
10. Reusable study knowledge base	CRO owners, QA leads, PMs, site networks	Closeout interview, lessons learned, country/site timelines, query patterns, template library, and source-date confidence notes.	Country playbook and study lessons database; higher reuse rate and fewer repeated startup queries.

Timeline and cost

Indicative timelines and costs for the three proposed workstreams. Final scope should be confirmed before each workstream starts.

Workstream	Timeline	Indicative cost	Scope note
VIR Operating Memory	4 weeks	R70k	SharePoint archive mapping, structured project intelligence, and a reusable foundation for new project records.
HEOR Commercial and Delivery Workbench	4-8 weeks	R70k-R150k	Depends on selected scope: proposal support, evidence reuse, assumption registers, QC support, or internal HEOR memory.
CRO Growth Wedge	4 weeks	R70k	Focused first wedge around trial startup, submission, site activation, TMF/QC, or bilingual document workflow support.