Performance and automation opportunity

GO3D Cloud + Dragonfire | Local AI with Qwen | Portfolio-wide planning report

Less continuous desk work. More repeatable work completed.

The opportunity is to turn repeatable processing into validated jobs: organize data once, run suitable operations locally, reuse unchanged results, and bring a person back for exceptions and acceptance. This report describes that target operating model across Eddie Oliveras's industries and services. It does not certify that every integration is already built.

8x
Modeled speed vs manual

4x
Modeled speed vs AI-assisted

$62.50
Modeled cost per package

ILLUSTRATIVE PLANNING CASE - NOT A MEASURED BENCHMARK.
The example assumes 8 hours manual, 4 hours AI-assisted, and a 1-hour GO3D target for the same accepted processing package. Modeled internal costs are $620, $255 and $62.50: approximately 90% and 75% lower in the GO3D case. All inputs and sensitivities are on pages 7-8.

What is established, and what remains to prove

Recorded setup: Aider connected to Ollama with Qwen2.5-Coder 7B on the MSI. **Recorded workload:** project validation; the inspected R63 script checks files and source hashes. **User-reported timing:** 5-6 minutes reduced to 27 seconds, which would be 11-13x on equivalent work. The complete conversion, output quality and cost comparison have not been independently verified. [1, 6, P1]

Commercial proposition: reduce repetitive handling, keep project context together, run approved batches between review points, and measure cost per accepted deliverable. Future optimization opportunities are research hypotheses until tested. No universal 100-1,000x advantage or technician-free operation is asserted.

System identity and evidence

A clear boundary between the installed local assistant and the proposed processing platform.

Item	Evidence status	Meaning
Local AI stack	Recorded setup	Aider selects **ollama_chat/qwen2.5-coder:7b** at the local Ollama endpoint. GO3D Cloud is the system brand; Qwen is the selected model. [1, P1]
Custom model training	Not established	No training dataset, fine-tuning run, model-weight revision or evaluation was supplied. Configuration and project instructions are not evidence of a newly trained GO3D model.
Parallel execution	Limited evidence	Earlier records describe concurrent validation jobs. The directly inspected R63 script tests viewer files, GLB headers and immutable source hashes. It contains no scan reconstruction or Parasolid export. [6, P1]
27-second conversion	User-reported	The arithmetic is 300/27 to 360/27 = 11.11-13.33x. Equivalent inputs, timing boundaries, geometry quality and output format remain to be checked.
Local data and cost	Depends on setup	Ollama supports local execution. Any enabled cloud model, connector, upload or hosted tool changes the data and billing path. Local files alone do not make cloud AI free. [2, 3]
24/7 / all providers	Design objective	A durable scheduler, recovery, adapters, licenses and resource limits require implementation and acceptance tests. No active unattended production queue is established by this report.

Three different layers: Qwen can help interpret a request or write code. Geometry and application tools perform the numerical processing. GO3D orchestration coordinates inputs, dependencies, logs and review. An AI response saying a file is valid is not a substitute for independent geometry checks.

A faceted solid, a smooth reconstructed BREP, and an editable feature-based CAD model are different deliverables. Parasolid supports facet and classic BREP representations. Import diagnostics may repair faces and gaps, but source fidelity and intended-use acceptance need separate tests. [5, 13]

From attended work to reviewed batches

Proposed operating model. Individual stages can be implemented and tested independently.

01 PRESERVE + INDEX

Original scans, images, drawings, documents and metadata

02 PLAN + SCHEDULE

Qwen selects approved recipes; dependencies and resource limits

03 EXECUTE

Local geometry, OCR, feature extraction and application adapters

04 CHECK

Geometry, units, source deviations, completeness and output reopen

05 REVIEW

Exception queue; qualified acceptance and release decisions

06 DELIVER

CAD / BIM / reports / media / records with provenance

The scheduler runs independent files, regions or checks together only after their inputs exist. Dependent operations wait. Job count is bounded by memory, storage, GPU availability and licensed application capacity. A dependency graph is a standard way to express this relationship; it does not imply unlimited recursive parallelism. [4, 7]

What allows work to continue after the technician leaves

A defined job contract. Each job names its source revision, units, output format, tolerances, required checks and allowed transformations. A vague request cannot establish engineering acceptance criteria.

Checkpointed execution. Write outputs to a new revision, keep a stage log, and resume completed stages after interruption. Retry limits prevent a bad file from consuming the entire queue.

An exception path. Missing inputs, failed fits, altered dimensions and low-confidence extraction become review items. Successful automation reduces continuous supervision; it does not remove responsibility for acceptance.

A controlled overnight queue. Only approved, bounded jobs should proceed unattended. Power, cooling, disk space, licenses, wake/sleep behavior and recovery must be demonstrated on the MSI before describing the service as 24/7.

Capacity is limited by the slowest available stage. A simplified upper bound is the minimum of machine jobs/day, reviewer jobs/day and available input jobs/day. Do not multiply an 8x scenario by a separate 3x "24/7 advantage" without measuring those bottlenecks.

Engineering and the built environment

Portfolio inventory from supplied business history. Every automation listed is a candidate, not a deployment claim. [P1]

Reverse engineering / Scan-to-CAD / PLM

DOCUMENTED SERVICE FAMILY

Inputs: Scans, meshes, drawings, photos and existing CAD.
Automate: Batch segmentation, repair candidates, feature proposals, conversions, revisions and export checks.
Value: Less repeated preparation; consistent CAD handoffs.
Review: Units, source deviation, topology, surface quality and required editability.

Metrology / aerospace / manufacturing

DOCUMENTED SERVICE FAMILY

Inputs: Tracker/probe data, point clouds, nominal CAD and tolerances.
Automate: Repeat datum alignment, feature fitting, deviation analysis and inspection reports.
Value: Less repetitive measurement analysis across industrial and defense supply-chain work.
Review: Calibration, datums, uncertainty, tolerances and acceptance decisions.

Construction / civil / survey / BIM

DOCUMENTED SERVICE FAMILY

Inputs: TLS/mobile LiDAR, survey control, photos, drawings and schedules.
Automate: Registration checks, object proposals, as-built preparation, changes and progress comparisons.
Value: Fewer transfers between capture, design and project records.
Review: Control, coordinates, dimensions, completeness and qualified review.

Facilities / asset digital twins

SERVICE EXPERIENCE + PLATFORM GOAL

Inputs: Facility scans, equipment tags, manuals, drawings and maintenance records.
Automate: Index assets, link evidence, propose components and compare repeat captures.
Value: Searchable, traceable operational context beyond a model viewer.
Review: Asset identity, geometry, missing information and source attribution.

Theme parks / rides / attractions

DOCUMENTED CAREER AND SERVICES

Inputs: Scans, tracker data, fabrication drawings and design revisions.
Automate: Fit-up and clearance comparisons, revision checks and installation records.
Value: Earlier coordination findings with less repeated desk analysis.
Review: Engineers approve interfaces, tolerances, loads and safety-critical conclusions.

Robotics / sensors / connected devices

PROPOSED CROSS-CUTTING INTEGRATION

Inputs: Device APIs, calibrated measurements, machine states and approved procedures.
Automate: Queue supported acquisition or inspection tasks and reconcile their results.
Value: Fewer disconnected device-to-software steps.
Review: Actual adapters, interlocks, operator control and physical acceptance tests.

Reality capture, experience and research

Portfolio inventory from supplied business history. Every automation listed is a candidate, not a deployment claim. [P1]

Heritage / sculpture / bodies / avatars

DOCUMENTED SERVICE FAMILY

Inputs: Original scans, reference photos, textures and historical source material.
Automate: Cleanup candidates, watertightness checks, textures, scale variants and print preparation.
Value: Faster figurine, likeness and heritage deliverable preparation.
Review: Preserve likeness, original geometry, scale, artistic intent and source revisions.

Forensics / CSI / medical visualization

DOCUMENTED WORK + DEVELOPMENT

Inputs: Scene scans, photos, measurements, case records and authorized medical material.
Automate: Evidence indexing, alignment, measured features and annotated presentations.
Value: Less repeated case preparation with clearer traceability.
Review: Qualified experts own interpretation, uncertainty and reconstruction assumptions.

Property / hospitality / virtual tours

DOCUMENTED SERVICE FAMILY

Inputs: Matterport captures, panoramas, photos, floorplans and property records.
Automate: Room and asset labels, tour packaging, floorplan workflows and media exports.
Value: Faster property documentation and listing preparation.
Review: Dimensions, completeness, authentic imagery and publication approval.

Marine / yachts / boat design

DOCUMENTED SERVICES + DESIGN CONCEPTS

Inputs: Vessel scans, interior imagery, layouts, CAD and specification records.
Automate: Tour packages, component indexing, fit studies and design revisions.
Value: Connected visual and dimensional vessel information.
Review: Source fidelity and appropriate design review; no implied naval certification.

Agriculture / orchards / crop intelligence

PROPOSED / RESEARCH VERTICAL

Inputs: RGB/multispectral imagery, LiDAR, repeat captures and field observations.
Automate: Tree and row segmentation, canopy change, fruit-count candidates and yield experiments.
Value: Repeatable monitoring for fruit, nut, vegetable and vineyard applications.
Review: Crop-specific ground truth, calibration, sampling limits and uncertainty.

Product design / printing / fabrication

SERVICES + EXPANSION CONCEPTS

Inputs: CAD, meshes, materials, machine constraints and product requirements.
Automate: Variants, print preparation, export checks, render queues and production documents.
Value: Faster approved iterations; possible laser-engraving and tooling workflows.
Review: Fit, material/process constraints, machine setup and first-article acceptance.

Operations, rentals and digital services

Portfolio inventory from supplied business history. Every automation listed is a candidate, not a deployment claim. [P1]

Scanner rentals / targets / tooling

SERVICES + RMS DEVELOPMENT

Inputs: Kit lists, serials, QR/RFID records, condition images and bookings.
Automate: Kit reconciliation, labels, checklists, condition comparisons and stock groupings.
Value: Faster kit turnaround and fewer administrative omissions.
Review: Physical condition, completeness, calibration and actual availability.

Consulting / training / technical support

DOCUMENTED SERVICE FAMILY

Inputs: Manufacturer manuals, approved workflows, examples and support records.
Automate: Knowledge retrieval, troubleshooting candidates, training versions and support routing.
Value: Less repeated explanation and more consistent onboarding.
Review: A specialist verifies instructions against real equipment and software.

GOEV / EV + adventure rentals

BUSINESS ACTIVITY + EXPANSION PLANS

Inputs: Vehicle and gear records, bookings, inspections and maintenance history.
Automate: Packages, matching, inspection documents and maintenance flags for bikes, paddling and camping.
Value: Less office work; better prepared rental bundles and operator workflows.
Review: Physical inspection, suitability, release and actual dealer integrations.

New Smyrna / charters / tourism

OUTFITTERS HISTORY + NEWER CONCEPTS

Inputs: Trip bookings, guest preferences, gear inventories and original media.
Automate: Booking administration, trip documents, checklists and photo/video packages.
Value: More time with guests and less repeated administration.
Review: Captains and guides own conditions, operational choices and customer suitability.

GO3D business OS / operations

PLATFORM DEVELOPMENT; INTEGRATION VARIES

Inputs: Authorized CRM, sales, project, finance, inventory, support and booking records.
Automate: Extract fields, match records, prepare tasks, documents, reconciliation candidates and routing.
Value: Fewer duplicate entries across marketing, forms, commerce, procurement and assets.
Review: Owner approval of financial actions, commitments, access and external messages.

Photography / visualization / media

DOCUMENTED WORK + PLATFORM APPLICATIONS

Inputs: Original photographs, approved logos, models, footage and verified project facts.
Automate: Asset indexing, render batches, revision galleries, media variants and draft copy.
Value: Faster production, virtual experiences and consistent delivery.
Review: Preserve authentic imagery/geometry and approved branding; verify claims.

Adjacent business history: Barber and other personal-service businesses can reuse scheduling, inventory and customer-record automation. This is an administrative application, not an active GO3D 3D vertical. Historical experience does not imply current partnerships.

A transparent speed and cost model

Illustrative inputs chosen for planning. No competitor, customer or GO3D production benchmark is represented.

Comparison unit: one repeatable processing package, starting with the same available source data and ending with equivalent accepted output. Acquisition, travel and client sales activity are excluded. Hands-on preparation, supervision, expert review and ordinary acceptance work must fit within the assumed operator time. If they do not, the cost must be increased.

Assumed input	Manual baseline	AI-assisted baseline	GO3D target
Elapsed processing time	8 hours	4 hours	1 hour
Active operator time	8 hours	3 hours	0.5 hour
Loaded operator cost	$75 / hour	$75 / hour	$75 / hour
Common allocation / package	$20	$20	$20
Incremental cloud AI / package	$0	$10	$0 for local run
Automation setup amortization	$0	$0	$5
Modeled internal cost	**$620.00**	**$255.00**	**$62.50**

Reproducible calculation

Cost = operator hours x $75 + $20 common allocation + cloud usage + automation allocation.
Manual: 8 x $75 + $20 = $620.
AI-assisted: 3 x $75 + $20 + $10 = $255.
GO3D: 0.5 x $75 + $20 + $5 = $62.50.

The $20 allocation represents software, workstation, energy and storage per accepted package. The $5 allocation assumes $5,000 of setup investment spread across 1,000 accepted packages. These are placeholders to replace with actual license, hardware, development and maintenance costs. The $10 AI allowance is also an assumption, not a quoted vendor rate.

Modeled result: 8x elapsed-time speed versus manual, 4x versus AI-assisted; 89.92% and 75.49% lower internal processing cost. Active operator time falls by 93.75% and 83.33%. These numbers apply only if the stated one-hour / half-hour target and equivalent output quality are achieved.

This is a conventional **manual-heavy** comparator and one hypothetical AI-assisted comparator. Some existing tools already automate extensively. A fair competitive claim requires a matched test of the actual alternative. Internal costs are not customer prices or profit margins.

Rework, volume and break-even

Operator rework, utilization, licensing and output quality are part of the economics.

GO3D case	Elapsed / operator	Internal cost	Speed vs manual / AI	Cost saved vs manual / AI
Conservative	2.5 h / 1.5 h	$137.50	3.2x / 1.6x	77.8% / 46.1%
Planning	1 h / 0.5 h	$62.50	8x / 4x	89.9% / 75.5%
Stretch	0.5 h / 0.25 h	$43.75	16x / 8x	92.9% / 82.8%

All three rows are scenarios. They are not measured ranges, confidence intervals, or predictions of this MSI. The stretch case needs stronger automation and less human intervention than the planning case.

Rework can consume the savings

Additional operator rework	GO3D total cost	Savings vs AI-assisted
30 minutes	$100.00	60.8%
1 hour	$137.50	46.1%
2 hours	$212.50	16.7%

At fixed non-labor allocations, GO3D reaches cost parity with the assumed AI-assisted case at approximately **3.07 total operator hours** per package, and with manual at **7.93 hours**. These are cost thresholds; they do not establish elapsed-time speed. Count failed attempts and their review time in cost per accepted output.

Volume changes setup cost per package

Accepted packages sharing $5,000 setup	Setup allocation	GO3D planning cost
1,000	$5.00	$62.50
100	$50.00	$107.50
10	$500.00	$557.50

Investment test: using $57.50 recurring GO3D cost before setup allocation, the illustrative $5,000 investment recovers after 9 packages versus manual or 26 versus AI-assisted. Do not charge both the upfront investment and its amortization in the same payback calculation. A low-volume custom job can be more expensive than the AI-assisted baseline.

Optimization experiments 1-5

Performance opportunities to investigate. These are not deployed capabilities or claimed billing exploits.

01 Profile the entire accepted workflow

Opportunity: Find whether time is spent in loading, geometry, model calls, serialization, human review or output repair. Optimize the stage that controls completion, not the most visually busy CPU meter.

Measure / accept: Log per-stage wall time, operator minutes, failures and accepted outputs. Compare the full job before and after; retain the same quality criteria. [4]

02 Reuse unchanged work

Opportunity: Cache expensive stages using source hashes, parameters and software versions. An unchanged registration, extraction or preview may not need to run again. DVC documents reuse based on recorded stage-run signatures.

Measure / accept: Benchmark cold reconstruction separately from warm reuse. Invalidate after any relevant source or parameter change; verify outputs and lineage. A cache is not a backup. [9]

03 Schedule dependencies and useful batches

Opportunity: Run independent files, regions or inspections together. Avoid huge graphs of tiny operations and repeated waits that serialize otherwise independent work. Dask documents both task graphs and batching considerations.

Measure / accept: Test one worker and bounded larger pools with sensible chunk sizes. Record scheduling overhead, memory, completion time and accepted quality. Do not count task count as a speed result. [7, 8]

04 Move numerical loops into geometry tools

Opportunity: Let Qwen help select or maintain a recipe, while compiled array, tensor and geometry operations perform numerical work. Keep large geometry out of repeated conversational context when the tool can process the original file.

Measure / accept: Compare equivalent numerical operations, minimize copies and record tolerances. Validate the result independently; a language-model statement is not a geometry certificate. [8, 10]

05 Test GPU work where the build supports it

Opportunity: Open3D documents CUDA tensor registration and multiscale ICP. This is a candidate stage optimization, not proof that an entire scan-to-CAD process is accelerated.

Measure / accept: Confirm the actual device, supported package and GPU execution on the MSI. Compare fitness, error, memory and full elapsed time. Windows CUDA/XPU packages are experimental in the cited release. [10, 11]

Optimization experiments 6-10

Performance opportunities to investigate. These are not deployed capabilities or claimed billing exploits.

06 Share one CPU / GPU resource budget

Opportunity: Outer workers and numerical libraries may each create threads. Uncontrolled nesting can oversubscribe the machine. Qwen requests and scan arrays also compete for memory.

Measure / accept: Coordinate worker count, native thread limits and RAM/VRAM headroom. Open3D 0.20 uses TBB thread controls; old OpenMP-only settings are insufficient. [2, 11, 12]

07 Use coarse-to-fine work without losing the source

Opportunity: A lighter representation can help search, preview or initial alignment; a full-resolution acceptance stage must still use the required evidence and tolerances. Preserve the original scan and clearly identify derived representations.

Measure / accept: Measure the final deviation and missed-feature rate, not only preview speed. Establish a tolerance budget before accepting decimation or approximate intermediate calculations. [10; proposed workflow]

08 Route knowledge and models by task

Opportunity: Use structured inputs, approved recipes and a local project index to avoid repeatedly rebuilding context. Test local Qwen for supported coding/planning tasks; use dedicated image, OCR or geometry tools for their respective data.

Measure / accept: Track model/tool calls, errors and operator correction time. Changing providers requires adapter and acceptance tests; do not claim all products are interchangeable. [1, 2; proposed workflow]

09 Automate repeatable outputs and recovery

Opportunity: Generate reports, exports, previews and record updates from a validated result. Use versioned outputs, completion manifests, bounded retries and resumable stages to reduce repeated handling after interruption.

Measure / accept: Interrupt a test job and verify recovery without duplicated work or overwritten originals. Reopen final files in the intended application and retain the evidence for each release. [Proposed workflow]

10 Investigate the next bottleneck

Opportunity: Profile real jobs for reusable features, selective recomputation, better data layouts, specialized kernels or application APIs. Unknown optimizations have no defensible multiplier.

Measure / accept: Change one variable, compare accepted outputs and retain the experiment. Do not multiply stage improvements into a whole-workflow promise. Prioritize the largest measured bottleneck. [4; proposed experiments]

The proof plan

Turn an attractive planning case into a defensible claim for a defined workload.

First priority: recover the source, commands, timing log and output for the reported 27-second operation. Determine whether the timer covered reconstruction, conversion, export, diagnostics or a validation rerun. Preserve both the original and the output before tuning anything.

Representative workload	Acceptance evidence
Human scan to solid	Source hash, units, exact representation, body count, reopen in target CAD, source-to-output deviation and repaired-region review.
Mechanical feature / inspection	Known datums and tolerances, independent feature/measurement comparison and documented uncertainty assumptions.
Facility / BIM package	Control and coordinates, completeness, dimensions, component identities and links back to capture evidence.
Document / asset extraction	Ground-truth field set, missing/incorrect field counts, cited source locations and operator correction time.
Rental / operational report	Known inventory and conditions, correct matching, exception handling and authorized actions only.
Agriculture / forensic research	Separate domain-specific validation with field or case evidence. Do not reuse a generic CAD result as proof of accuracy.

Run a fair comparison

Freeze a representative input and an agreed output contract. Compare the current manual process, the actual AI-assisted alternative, and GO3D on equivalent requirements. Run a small repeat set, for example five accepted runs per configuration, and report the median and range rather than only the fastest attempt. Separate cold and warm-cache trials.

Measure elapsed time, active operator minutes, model/provider usage, licenses and machine allocation, peak memory, failures, rework and geometric or extraction quality. Record revisions and parameters. If a faster result fails the quality contract, it is not an accepted speed improvement.

Release wording after validation: "On [defined task and dataset], GO3D completed [same accepted output] in [median time], versus [baseline], using [operator minutes] and [measured processing cost]." State exclusions and sample size. Until that evidence exists, retain "illustrative scenario" and "user-reported timing" labels.

Implementation sequence: establish evidence and source protection; automate one repeatable job; add bounded parallel execution; prove interruption recovery and review; then add industry adapters. A benchmark in one vertical does not validate every other service.

Sources, assumptions and scope

Primary technical sources checked 24 October 2026. Numbered references identify support, not product endorsement.

[1] Aider - Ollama connection. Local model endpoint and ollama_chat configuration.

[2] Ollama - FAQ. Local/cloud behavior, memory-dependent concurrency and server controls.

[3] OpenAI - ChatGPT Work / Codex pricing. Local sessions can still consume hosted-model usage. No rate is used as the scenario's assumed $10 allowance.

[4] Microsoft - Pitfalls in data and task parallelism. Parallel overhead, over-parallelization and measured results.

[5] Siemens - Parasolid with Convergent Modeling. Distinction between facet BREP and classic BREP.

[6] GO3D source - R63 validation script. Read-only source inspection; file and hash checks, not CAD conversion. Repository access may be required.

[7] Dask - Task Graphs. Dependency-based task scheduling.

[8] Dask - Best Practices. Task granularity, shared intermediates and computation structure.

[9] DVC - Run Cache. Reuse of stage results under matching run conditions.

[10] Open3D - Tensor ICP registration. Multiscale registration, CUDA example and registration quality metrics.

[11] Open3D - v0.20 release, 24 September 2026. Windows accelerator builds are experimental; CPU threading uses oneTBB.

[12] Scikit-learn - Parallelism and resource management. Nested native-thread oversubscription and resource controls.

[13] SOLIDWORKS - Import Diagnostics Overview. Repair, knitting and solid creation; separate acceptance criteria remain necessary.

[P1] User-provided evidence and portfolio. Eddie Oliveras's statements, supplied setup history and service descriptions in this conversation and retained project context. The 5-6 minute / 27-second comparison is user-reported. Portfolio coverage reflects documented experience and plans, not an audit of current commercial delivery or partner relationships.

Method: all scenario arithmetic was calculated from the explicit page-7 assumptions and independently cross-checked. The comparison image uses those same values. No timing was measured on the MSI for this report, no customer price or actual bill was obtained, and no future optimization was assigned an evidence-based speed multiplier.