RaNDTek LLC Agent OS

Pitch Deck

Everything. Anything. Together.

A desktop operating layer for intelligent software interlinking: all software, all available coding agents, local files, literature, databases, and human review working as one system.

Start Here: Live Links

Agent OS Site
https://agentos.randtek.org/

Investor FAQ
https://agentos.randtek.org/investor-faq

Technical FAQ
https://agentos.randtek.org/technical-faq

Case Card Movies
https://agentos.randtek.org/#movies

Two-second takeaway

Agent OS is a desktop operating layer for large numbers of interoperating agents: one Agent OS, thousands of time-proven software packages, one dedicated agent per package, and workflows that run them together under human review.

MULTI-SOFTWARE WORK ACCELERATION. LESS TIME AND MONEY.

The Team

Rand Dannenberg, Ph.D.

Founder, RaNDTek LLC. 30 years in optics, photonics, materials science, engineering, chemistry, and machine learning. Agent OS grew from real consulting work.

Seeking investors, advisors, board members, executives, software engineers, and technical staff.

agentos.randtek.org | rand@randtek.org | 727-480-9186

THE SCALE PROBLEM

100 million to 1 billion pre-existing software packages already run the world.

AI is not going to rewrite them.

They do everything from forensic accounting, engineering, and drug discovery to quantum gravity. The Age of Agents needs an operating layer that can coordinate existing software, not pretend the installed software universe disappears.

Forensic accounting

Quantitative workflows, files, evidence, reports, and review gates.

Engineering

CAD, FEA, CFD, optics, analysis, validation, and deliverables.

Drug discovery

Models, chemistry codes, databases, literature, and experiments.

Quantum gravity

Pure science workflows, literature, symbolic work, and simulation.

What will run them all in the Age of Agents? RaNDTek Agent OS: thousands of software packages, one desktop operating layer, one dedicated agent per package, coordinated under human review. Near-term focus: quantitative disciplines and media/marketing.

FLAGSHIP MOVIES

Two movies explain the product in seconds.

The first shows the operating thesis: run all software and coding agents on Earth. The second shows modular disciplines that can be sold, installed, populated, expanded, and reused.

Movie link: https://agentos.randtek.org/media/better-movies/flagship-product-tour.mp4

Flagship product tour

Run all software and coding agents on Earth: one desktop layer, one dedicated agent per package, human setup, automated execution, and human review.

Movie link: https://agentos.randtek.org/media/better-movies/part-11-materials-analysis.mp4

Modular discipline suites

Pre-populated discipline suites assemble software packages, files, standards, literature, prompts, workflows, databases, and 1:1 software-agent pairs.

Click either movie URL or frame to watch.

agentos.randtek.org | rand@randtek.org | 727-480-9186

WHAT NO CURRENT OFFERING DOES ALL AT ONCE

No other offering does all of these simultaneously under one roof.

Agent OS is designed and evolved for technical work: local desktop operation with your files, real software, 1:1 software-agent pairs, teams, graphical workflows, security, literature, databases, lower-cost account paths, and any agent from flagship to local.

Local desktop operation + your files Files, folders, drives, and private data stay central.	**Time-proven software** Runs proven local or SaaS tools instead of replacing them.	**1:1 software-agent pair** One dedicated agent can be bound to each software package.	**Any agent, flagship to local** Flagship coding, enterprise, local, or future agents.
Built for technical work Designed and evolved for quantitative, rigorous disciplines.	**Graphical workflows** Easy to build, inspect, revise, monitor, and reuse.	**Looping workflows** Repeated passes toward convergence, specs, and review.	**Decision nodes** Explicit tests, gates, branches, reroutes, and approvals.
Creative nodes Exploration, synthesis, and invention steps can be formalized.	**Agent Teams** Multiple agents coordinate work, status, and deliverables.	**Workflow subagents** Subagents run serial or parallel workflow steps.	**Human review gates** Approve, pause, redirect, rerun, reject, or release.
Multi-level security Permissions, sandboxes, logs, restrictions, OS containment.	**Supported install/setup** Code/channel install, binding, setup, and onboarding.	**Help docs + starter prompts** Docs extraction and prompts for each software path.	**Auto-growing skills** Best practices, prompts, examples, and workflow memory improve.
Literature RAG Papers, books, links, manuals, standards, and notes.	**Custom databases** Private records and databases become workflow context.	**Lower-cost account CLI** OAuth/account access where available; not API-only tokens.	**Integration + discipline packs** Threads reduce MCP burden; packs can be pre-populated.

Detailed competitive landscape: Agent OS does all of these together.

Investor FAQ: https://agentos.randtek.org/investor-faq

agentos.randtek.org | rand@randtek.org | 727-480-9186

WORKFLOW PROOF

The answer is not another prompt. It is an operating workflow.

A real Agent OS workflow preserves the mission, software channels, execution logic, evidence, and review path.

What the movie shows

A structured workflow in Agent OS: the human defines the technical mission, the workflow stores the logic, and software channels execute repeatable steps.

Why it matters

The work no longer disappears into a chat transcript. It becomes a reusable operating object that can be rerun, inspected, improved, and packaged.

This is the bridge from a brilliant temporary agent to a durable desktop operating layer.

Movie link: https://agentos.randtek.org/media/better-movies/what-are-workflows-for-v1.mp4

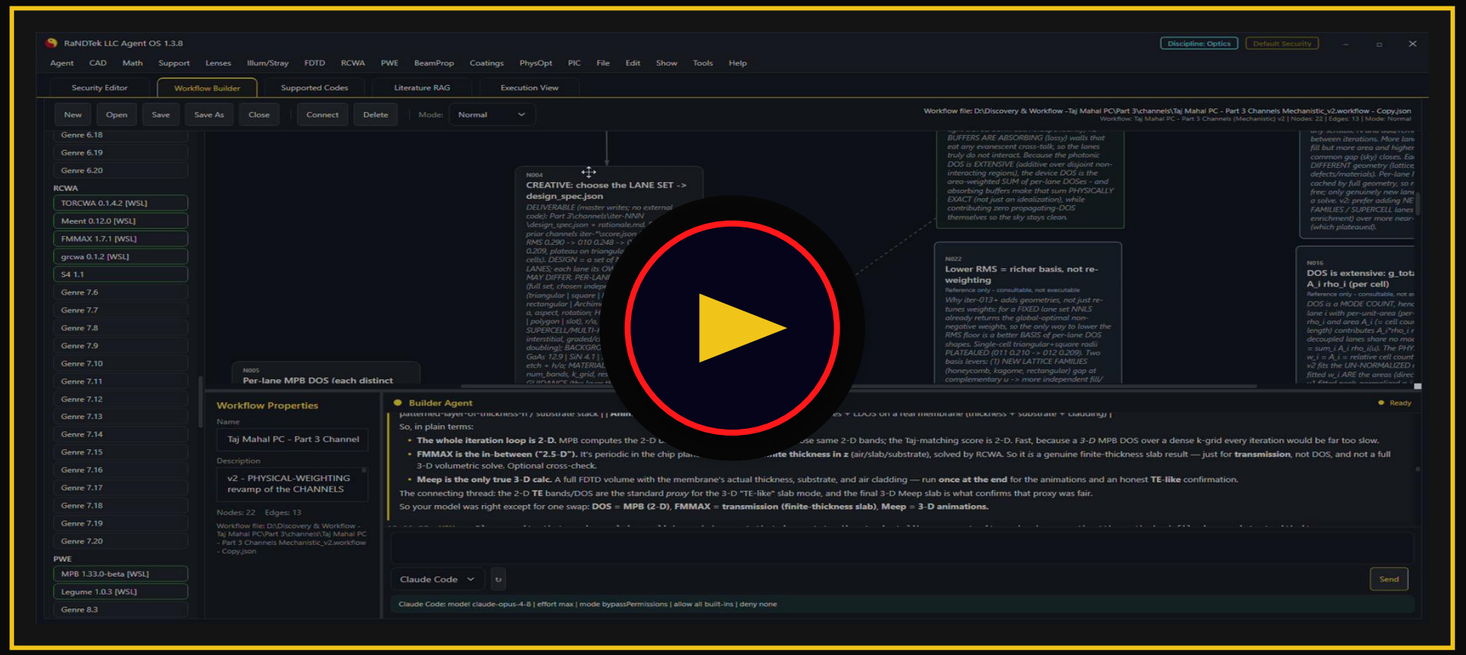

Movie evidence: structured workflows for technical work

Click to watch movie: https://agentos.randtek.org/media/better-movies/what-are-workflows-for-v1.mp4

agentos.randtek.org | rand@randtek.org | 727-480-9186

MARKET FOCUS

Quantitative work first.

Bottom-up, not top-down: ~63M professionals worldwide must produce numerically correct, traceable results across many tools. The right-side model prices that narrow ICP as professional software.

Adoption filter: heavy agentic workflow users

Function	%
Software / IT / data / technical engineering	3.4
Management / ops	1.8
Finance / pro services	1.3
Research / academia	1.0
Marketing / creative	1.0
Admin / clerical	0.9
Sales / customer	0.8
Trades / owner-ops	0.7
Health / legal	0.6
Education	0.5
Non-users	88.0
Total	**100**

Why this ICP pays

Correctness-sensitive, tool-heavy, iterative work.

These users do not buy Agent OS because it is another chatbot. They buy it because real work already lives across multiple technical tools, files, databases, papers, checks, and deliverables. The first answer is rarely final; the value is the loop: set up, run, compare to specifications, revise, document, and repeat.

Launch wedge: optics, photonics, materials, chemistry, mechanical/aerospace engineering, finance/trading, analytics, drug discovery, quantum computing, and other hard quantitative disciplines using real design/simulation software.

You don't want to get on a roller-coaster designed by a "physics-aware surrogate" AI that attempts to replace hardcore, rigorous design/simulation with neural networks that don't.

The ICP is the rigorous worker and organization.

Bottom-up market model

TAM

$113B market

~63M quantitative professionals x $1,800/year; conservative floor $72B.

SAM

$30B serviceable wedge

Windows-first, English-first, RaNDTek launch disciplines: ~15-22M seats.

SOM

$36M ARR

~20,000 paid seats by Year 4. Consulting-led pilots become seats.

Long-term expansion path

500K seats = ~$400-$500M ARR 1M seats = ~$0.7-$1.0B ARR 2M seats = ~$1.4-$2.0B ARR

Blended ACV declines as the broader iterative-work market mixes in.

Key takeaway: Agent OS reaches $36M ARR on ~20,000 seats, not 100,000.

GO TO MARKET

One year to alpha/beta market entry.

Consulting is the springboard: real customer problems become paid pilots, discipline packs, workflow evidence, and the first customer-ready Agent OS deployments.

0-3 months

Harden installer, workflow persistence, logging, permission profiles, and alpha UI.

3-6 months

Run founder-led consulting pilots; convert the strongest use cases into packaged workflows.

6-9 months

Alpha customers in quantitative disciplines and media/marketing; collect usability and reliability data.

9-12 months

Beta-ready product, paid pilots, support process, customer success motion, and next financing evidence.

Year 1 planning projection

500 paid seats; $180K product revenue; $150K services/support; $330K total revenue. Current burn is about $5K/month; post-financing plan is about $120K/month.

Business model

Desktop licensing, professional/enterprise editions, discipline packs, paid alpha/beta access, support contracts, custom discipline setup, security editions, onboarding, and consulting.

The first year is about turning founder-led consulting proof into repeatable product, customers, revenue, and evidence.

agentos.randtek.org | rand@randtek.org | 727-480-9186

THE ASK

The ask: $2,000,000 Pre-Seed / Seed.

$2,000,000 on a $15,000,000 post-money SAFE to fund product hardening, alpha/beta launch, security, installer/UI, and team formation.

$2M
raise amount

$15M
post-money valuation cap

13.33%
investor ownership after round

30.00%
people reserve for team formation

Post-round ownership

56.67% Rand / founder. 13.33% investors. 15.00% founding-team reserve. 15.00% employee/advisor option pool.

What the money buys

Team formation, software engineering, product hardening, installer/UI, security, reliability, documentation, customer deployment, and alpha/beta testing. Estimated post-financing burn: about $120,000/month.

The 30% people reserve

15% founding-team formation plus 15% employees/advisors. It recruits the people needed to build the company.

Alpha and beta path

Real customer deployment, feedback, discipline packs, Year 1 revenue target of $330K, and evidence for the next financing round.

agentos.randtek.org | rand@randtek.org | 727-480-9186

INVESTOR MOAT

The engine-makers can't build the vehicle.

Every frontier lab is shipping powerful agent primitives — coding agents, skills, memory, sandboxes. None can build the layer that runs all of them, because their business models forbid it. That layer is the company.

1 — Structural neutrality. Incumbents are barred, not behind.

A surface that treats every agent as an interchangeable brain can only come from a party that owns no model. Anthropic will never make Codex first-class; OpenAI will never make Claude. Each vendor consolidates inward; the world's work spreads outward across a hundred million installed software packages no lab will ever bind. Every shift in frontier leadership makes the neutral layer more valuable.

2 — Domain curation. Labor no lab will staff.

Binding real professional software takes expert judgment per discipline: which solvers to trust, which outputs to validate, what a correct workflow looks like. That curation is slow, knowledge-dense work that doesn't scale by capital alone — accumulated into discipline packs a generic competitor can't fake and a frontier lab won't bother to build.

3 — Customer-owned compounding. Switching cost that lives on their machine.

With use, each deployment accumulates local operating memory — workflows, control files, best practices, databases tuned to that customer's stack. It grows more valuable every run, and it cannot be absorbed by any model, because it isn't in any model.

The barrier to entry is the maintenance burden these three locks imply — and that burden is exactly what this raise funds.