A MANIFESTO

The Last App

Why the Application Era is Ending and What Comes Next

The tools will change. Your thinking is the constant.

Ron Lior
Fusália — AI Consulting & Education
Marvão, Portugal · 2026

Introduction: A Jules Verne Moment

Anything one man can imagine, other men can make real.

JULES VERNE

We are living in a Jules Verne moment. Not because we can see the future with perfect clarity, but because the trajectory is unmistakable. Just as Verne described submarines, space travel, and video conferencing a century before they existed, we can now see the outlines of a computing paradigm that will make today's software landscape look as quaint as a telegraph office.

This manifesto is not a technical whitepaper. It is not a product roadmap. It is a declaration of direction — an argument that the era of applications as we know them is ending, and that something fundamentally different is taking its place.

The destination: an AI-native operating system where you never open an app, never learn a new interface, never adapt your workflow to someone else's design decisions. You simply state what you want, and it happens.

This is not science fiction. The building blocks already exist. What's missing is the integration — and the collective imagination to see where these pieces are heading.

01

The App Was Always a Compromise

Every application you have ever used is, at its core, someone else's opinion about how you should work. Photoshop is Adobe's opinion about image editing. Excel is Microsoft's opinion about data manipulation. Salesforce is someone's opinion about how you should manage customer relationships.

We accepted these opinions because we had no alternative. The application was the only bridge between human intent and machine capability. You couldn't tell a computer what you wanted in plain language, so developers built pre-structured pathways — menus, buttons, forms, workflows — that translated your actions into machine instructions.

But think about what this actually means. Every time you open an application, you are surrendering your workflow to someone else's design. You learn their interface, adopt their mental model, and constrain your thinking to their feature set. You don't edit images the way your brain wants to — you edit them the way Adobe decided you should.

THE HIDDEN COST OF APPLICATIONS

Learning curves. Every new app requires hours or weeks of adaptation. The average knowledge worker uses 9–12 applications daily, each with its own logic, shortcuts, and limitations.

Context switching. Moving between apps means constantly translating your intent from one interface language to another. Copy from here, paste there, export in this format, import in that one.

Feature bloat. Applications grow endlessly because they must serve millions of different users with one fixed interface. Most people use less than 20% of any app's features, yet they bear the complexity of the other 80%.

Vendor lock-in. Your data, your workflows, your institutional knowledge — all trapped inside someone else's product decisions. Switching costs keep you captive even when better options exist.

The application model was the best we could do with the technology we had. It is not the best we can do now.

02

The Trajectory is Clear

The history of human-computer interaction is a story of progressive friction removal. Each paradigm shift brought us closer to natural communication with machines:

Command Line (1960s–1970s). You had to speak the machine's language. Every instruction had to be typed in precise syntax. One typo meant failure. Only specialists could operate computers.

Graphical User Interface (1980s–1990s). The mouse and desktop metaphor made computers visual and spatial. You could point at things instead of typing commands. Computing became accessible to millions.

Touch Interface (2000s–2010s). Smartphones removed the intermediary device. You touched what you wanted directly. A child could operate a computer. Computing became universal.

Voice & Conversational AI (2010s–2020s). Siri, Alexa, and early chatbots let you speak naturally. Limited, often frustrating, but a clear signal of where things were heading.

Intent-Based Computing (2020s–). Large Language Models can now understand context, nuance, and complex multi-step requests. The gap between human intent and machine action has narrowed to nearly nothing.

Each transition followed the same pattern: remove a layer of abstraction between what the human wants and what the computer does. The OSAI is simply the next logical step in this progression — and arguably the final one, because there is no layer left to remove after natural language intent.

The best interface is no interface.

GOLDEN KRISHNA

03

Meet OSAI — What the Future Actually Looks Like

It needs a name. You can't start a revolution by saying "the AI-native operating system paradigm." People will fall asleep before you finish the sentence. So let's call it what it is: OSAI. The OS we all know, fused with AI at its core. Four letters, whole paradigm. Easy to say, impossible to ignore. Think of it as your operating system finally graduating from being a glorified file cabinet to something that actually understands what you're trying to do.

This is not abstract futurism. Let me paint a concrete picture of what daily computing looks like when the application layer dissolves and OSAI takes over.

MORNING

You sit down at your desk. There are no app icons, no dock, no taskbar. No cheerful little squares begging for your attention like needy puppies. There is a clean surface — perhaps a subtle prompt, perhaps just your voice. You say: "What's my day look like?"

The OSAI doesn't open a calendar app. It synthesizes your schedule, pending emails, project deadlines, weather (you have an outdoor meeting), and traffic conditions into a single, coherent briefing. It knows you prefer visual timelines in the morning. It knows your 2pm call was rescheduled because it processed that email at 11pm last night.

WORKING

You say: "Prepare the quarterly report for the board. Use last quarter's format but update all numbers, add the new product line performance, and flag anything that's more than 10% off forecast."

There is no "open Excel, pull data, switch to PowerPoint, format slides, curse at PowerPoint, reformat slides, wonder why the chart broke, fix the chart, realize it's the wrong data, start over." The OSAI understands the task holistically. It pulls data from your company systems, applies your established formatting preferences, generates the analysis, and presents it for your review. If it needs a chart, it creates a chart. If it needs a table, it creates a table. It doesn't think in terms of applications. It thinks in terms of outcomes.

CREATING

You say: “I want to edit the product photos from yesterday’s shoot. Make the backgrounds consistent, adjust the color temperature to match our brand guidelines, and prepare web and print versions.”

No Photoshop. No Lightroom. No batch processing scripts. No spending forty-five minutes watching YouTube tutorials on how to remove a background. The OSAI has image manipulation capabilities built into its core. It knows your brand guidelines because they’re part of your persistent context. It processes all images simultaneously and shows you the results for approval. You just saved three hours and a minor existential crisis.

THE KEY PRINCIPLES

No app boundaries. Capabilities flow together seamlessly. A task that today requires five different applications happens in one continuous interaction.

Persistent context. The OSAI remembers your preferences, your projects, your history. You never configure, set up, or onboard. It already knows.

Adaptive presentation. Information appears in whatever form serves you best — visual, textual, tabular, spatial. The format is generated to match the moment, not fixed by an app’s design.

Proactive intelligence. The OSAI doesn’t just respond — it anticipates. It notices patterns, flags anomalies, suggests actions before you think to ask.

YES, OSAI KNOWS YOU. THAT’S THE POINT.

Now, before the privacy alarms start ringing so hard your neighbors call the police — let’s be honest about what’s already happening. Your computer today stores your files, your browser history, your passwords, your photos, your half-finished documents you swore you’d get back to. It knows plenty about you. It just does absolutely nothing useful with that knowledge. Your data sits in dead folders like books nobody reads.

OSAI changes one thing: it makes that data alive. A local OSAI — running on your machine, with your data, building your profile — doesn’t send your life to some distant server farm. It simply becomes aware of what it already has. It understands your preferences, your patterns, your work context. Not because it’s spying on you, but because that awareness is what makes it useful. A personal assistant who doesn’t know your name, your schedule, or your preferences isn’t an assistant — it’s a stranger with a clipboard.

Yes, this is power. And like any power, it deserves open eyes. But the answer isn’t to fear OSAI knowing you — it’s to ensure you own what it knows. Your cognitive profile, your workflow DNA, your professional context — that’s yours. The moment any company tries to hold it host-

age, you should be able to walk away and take it with you. Data portability isn't a nice-to-have in the OSAI era. It's a human right.

04

The Resistance Will Be Economic, Not Technical

If the trajectory is so clear, why isn't it happening faster? The answer has nothing to do with technology and everything to do with money.

The global software industry generates over $600 billion annually. Adobe, Microsoft, Salesforce, Google, SAP — these companies exist because you need their applications. Their entire business model depends on you opening their product, day after day, subscription after subscription.

An OSAI that absorbs application capabilities is not a new feature for these companies. It is an existential threat. It's the equivalent of Netflix threatening Blockbuster — not a better version of the same thing, but a fundamentally different model that makes the old one irrelevant.

HOW THE RESISTANCE WILL MANIFEST

Walled gardens. Companies will try to keep their AI capabilities locked inside their own ecosystems. Adobe's AI only works on Adobe files. Microsoft's Copilot only works with Microsoft products. The strategy is to make their AI a feature of the app, not a replacement for it.

Data moats. Your documents, your projects, your history — all stored in proprietary formats inside proprietary clouds. Switching becomes even harder as AI learns from your usage patterns.

Regulatory capture. Expect arguments about safety, reliability, and accountability. "You can't trust a general OSAI with critical tasks — you need our specialized, certified application." Some of these concerns will be legitimate. Many will be protectionist.

Gradual absorption. Rather than building a true OSAI, incumbents will add AI features to existing apps. This preserves their business model while appearing innovative. It's the equivalent of Blockbuster offering mail-order DVDs. "Look, we innovated!" No, you put a stamp on a dying format.

But economic resistance has never stopped a genuine paradigm shift. It only slows it down. The question is not whether the application era ends, but how messy the transition will be.

05

What Dies, What Survives, What's Born

WHAT DIES

The general-purpose productivity application. The expense tracker, the project management tool, the simple CRM, the form builder, the basic website creator. All those apps cluttering your dock like forgotten houseplants — still there, still taking up space, and you haven't opened half of them since 2023. Why subscribe to a project management tool when the OSAI can generate one that's perfectly tailored to your specific project, then dissolve it when you're done?

Software as a packaged, sold product — for many categories — stops making sense. The SaaS model, which depends on you needing someone else's pre-built solution, loses its foundation. Software becomes ephemeral: generated on demand, shaped to the moment, discarded after use.

WHAT SURVIVES

Deeply specialized, performance-critical, and safety-critical systems. You don't want an AI casually generating aircraft control software or nuclear plant monitoring systems. Game engines, medical imaging, financial trading platforms, embedded systems, industrial control — these require deterministic precision, rigorous certification, and extreme reliability. They will continue to exist as purpose-built systems, though their interfaces may be increasingly mediated by AI.

The infrastructure layer also survives and grows. Someone has to build the capabilities that the OSAI orchestrates. Cloud computing, data systems, security frameworks, networking — the plumbing becomes more important, not less.

WHAT'S BORN

Capability architects. A new discipline of designing modular, machine-consumable capabilities that AI systems can discover and compose. Less about building user interfaces, more about building reliable primitives.

Context engineers. Specialists who design and maintain the persistent knowledge layers that make OSAI systems personal and effective. Your brand guidelines, your workflow preferences, your institutional knowledge — all structured for AI consumption.

Intent translators. People who help organizations bridge the gap between how humans think about their work and how OSAI systems need to understand it. Part consultant, part trainer, part systems thinker.

Trust architects. As OSAI systems handle increasingly critical tasks, someone needs to design the verification, auditing, and accountability frameworks that make this trustworthy.

06

The Human Side

This is where most future-of-technology discussions fail. They describe the technology and forget the humans.

Having spent years in enterprise technology — at Meta, PayPal, Cognizant — and now working directly with professionals navigating AI, I can tell you: the biggest barrier to this future is not technical. It is cognitive. It is emotional. It is deeply human.

THE COMPETENCE CRISIS

Most professionals have built their careers on mastering specific tools. The Excel wizard who can write a VLOOKUP in their sleep. The Photoshop expert who knows seventeen ways to remove red-eye. The Salesforce administrator whose LinkedIn profile is 40% certifications. Their professional identity is intertwined with their tool proficiency. When you tell them the tools are going away, you're not just describing a technology shift — you're threatening their sense of professional worth.

I see this in every workshop I run. The fear isn't about AI. It's about relevance. "If anyone can do what I do just by asking the AI, what am I for?"

THE ANSWER: THINKING BECOMES THE SKILL

In a world where execution is handled by the OSAI, the value shifts decisively to thinking. Knowing what to ask for. Understanding what good looks like. Having the judgment to evaluate outputs. Seeing connections that the AI doesn't. Asking the second question, the harder question, the question nobody else thought to ask.

This is profoundly good news for humans — but only if they prepare for it. The transition from "I can operate this tool" to "I can think clearly about what needs to be done" is not automatic. It requires new mental models, new habits, and frankly, a kind of professional courage that most educational systems don't cultivate.

BUT HOW DO YOU ACTUALLY TEACH THINKING?

Here's the uncomfortable truth: saying "thinking becomes the skill" is easy. Actually teaching people how to think in outcomes is hard. You can't just hand someone an OSAI and say "go think." That's like handing someone a piano and saying "go compose." The instrument isn't the skill. The skill is knowing what music you want to make.

What does thinking in an OSAI world actually require? Three things. First, systems-level reasoning — the ability to see how pieces connect, to understand that asking for a "quarterly report" involves data, narrative, audience, format, and timing all at once. Second, precise problem decomposition — knowing how to break a vague goal into clear, actionable intent. "Make my business better" is useless. "Identify which customer segment has the highest churn rate and show me what they have in common" is powerful. Third, quality evaluation — recognizing when an output is genuinely good versus when OSAI has produced something that looks impressive but is actually nonsense in a nice font.

None of these are natural human defaults. Most education systems don't teach them well even with existing tools. Adding OSAI doesn't automatically solve this — it makes the gap more visible and the consequences of not closing it more severe. The people who master these thinking skills will thrive. The rest will have a very expensive machine producing very confident garbage.

THE COMFORT TRAP

And here's the paradox nobody wants to talk about. The better OSAI gets, the more tempting it becomes to stop thinking altogether. When your OSAI manages your schedule, drafts your emails, prepares your reports, and anticipates your needs — why bother doing any of it yourself? It's like having a personal chef: wonderful, until you realize you've forgotten how to boil an egg.

Skill atrophy is real. Dependency is real. And the accountability question — when OSAI makes a bad decision that costs your company money, who signs the check? — is not trivial. This isn't a reason to reject OSAI. It's a reason to engage with it deliberately, eyes wide open, brain fully switched on. The goal isn't to let OSAI think for you. It's to think with it. The humans who understand this difference will be the ones who flourish.

THE PREPARATION GAP

Right now, there is an enormous gap between what AI can do and what most people understand it can do. This gap is not closing — it is widening. Every month, AI capabilities advance faster than human comprehension catches up. The people who understand this transition early — who start thinking in terms of intent and outcomes rather than tools and interfaces — will have an extraordinary advantage.

Closing this gap is not a technology problem. It is an education problem. A human problem.

07

A Call to Action

This manifesto is not a prediction of what might happen. It is a description of what is already happening. The question is not whether the application era ends. It is whether you will be ready when it does.

FOR PROFESSIONALS

Stop investing your learning time in mastering the next app. Start investing in understanding what AI can do and how to direct it effectively. Your career will not be defined by which buttons you can press, but by the quality of your thinking and the clarity of your intent.

FOR BUSINESS LEADERS

Start questioning every software purchase. Ask: "Could an AI handle this workflow without a dedicated application?" The answer will increasingly be yes. Begin building organizational knowledge and processes that are AI-ready — structured, well-documented, and tool-agnostic.

FOR EDUCATORS

We are still teaching people how to use software. We should be teaching people how to think about outcomes, evaluate AI outputs, and direct intelligent systems. The curriculum needs to shift from tool proficiency to thinking proficiency.

FOR DEVELOPERS

The future of your craft is not building the next app. It is building the capabilities that OSAI systems will orchestrate. Think modular, think machine-consumable, think about reliability and composability. The interface layer is dissolving — what remains is the engine underneath.

We are in the early days. The submarine is still on the drawing board. But the ocean is real, and the direction is clear.

The last app has not been built yet. But it will be the last.

Power to OSAI.

Keep Asking and Rethinking.

RON LIOR · FUSÁLIA — AI CONSULTING & EDUCATION
MARVÃO, PORTUGAL · 2026