


OSaaS: The End State of Enterprise Technology

by Mark McNasby & Michał Ogłodek · Co-Founders · June 2026

Why the future of enterprise software is not a portfolio of applications — whether fifty or several hundred — but one sovereign operating system that grows them on demand: **Operating System as a Service**.



The enterprise is drowning in its own good decisions

No CIO ever set out to run fifty applications, much less hundreds. Each one arrived as a rational, defensible, locally correct decision. A tool to handle expense reports. A platform for onboarding. A system for vendor risk. A point solution for recertification, another for credentialing, another for the thing the last two didn't quite cover. Every purchase solved a real problem. And every purchase created two new ones: an integration to maintain, and a seam where the truth about the organization could now hide.

This is not a story about bad procurement. It is a story about entropy. Sprawl is not what happens when an enterprise makes mistakes; it is what happens when an enterprise makes a thousand reasonable choices over a decade, each optimizing locally while the system as a whole decays. The portfolio is the sum of a hundred sensible meetings and the residue of a dozen reorganizations, and no one is to blame, which is precisely why no one can fix it.

And the number is a lie anyway. An enterprise with "hundreds of applications" undercounts the disease by an order of magnitude, because each of those applications is rarely one thing. It is a web app and an iOS app and an Android app and sometimes a desktop client, each a separate codebase, each with its own release train, its own security review, its own signing identity, its own relationship with an app store. Behind every one sits a hidden production stack — the QA tooling, the penetration-testing engagement, the compliance platform, the DevOps pipeline, the team whose entire job is to keep the same intent in sync across four runtimes. The real surface area is not the count of applications. It is the count of applications multiplied by the platforms each must be rebuilt for, multiplied again by the machinery required to keep each one alive and trustworthy.

The enterprise is not running hundreds of things. It is running hundreds of things several times over, and paying a standing army to keep them dusted.

Every cure for sprawl was another symptom of it

The tech industry has been trying to solve sprawl for thirty years, and every solution failed on its own terms.

The mega-suite promised consolidation and delivered captivity: replace your portfolio with ours, and inherit our business logic, our data model, our pace of change, our notion of how your organization ought to work. The integration layer promised to make the seams disappear and instead made them official, a permanent budget line for moving data between stores that should never have been separate. The "single pane of glass" promised one view and delivered one more window in front of the same forty-five systems.

The reason all of them failed is the same, and it is worth saying plainly: **the unit of consolidation was always the application, and applications are rigid by nature.** You cannot consolidate rigid things. You can only dissolve the boundaries between them. Every attempt to herd the applications into a smaller pen left the applications intact, and so left the seams intact, and so left the disease intact.

You can see the failure most clearly in the protocols the industry invented to cope. SAML 2.0 was not a feature anyone ever wanted. Nobody woke up wishing for assertion-passing between identity providers and service providers. It exists for exactly one reason: you have N separate applications, each with its own idea of who a user is, and something has to make them *pretend* to agree. SAML is a treaty between sovereign states. Once you see it that way, you start counting the other treaties. SCIM exists to synchronize identities that should never have been separate. OAuth scopes exist to broker permissions across boundaries that should not exist. The iPaaS market exists to ferry data between stores that should have been one store. And single sign-on — the most damning of all — is the enterprise paying real money to *simulate* the unity it surrendered the day it bought its second application.

Every glue protocol is a receipt for fragmentation the enterprise already paid for. The more federation infrastructure an organization runs, the more separate things it is quietly admitting it has. You can tally the symptoms, and the tally is the diagnosis.

What changed

The reason this could not be fixed before is that the alternative was unaffordable. Building a bespoke, compliant, tested, secured, cross-platform application for a single purpose cost more than buying someone else's near-miss. So you bought the near-miss, hundreds of times.

That economic fact has now inverted. Generating an application — one that arrives tested and scanned, carrying its own compliance evidence, signed, and projected onto every platform at once — has become nearly free at the margin. And when the marginal cost of a purpose-built application approaches zero, the entire logic of the point-solution market collapses. You no longer buy a rigid near-miss and bend your organization around it. You produce the exact thing you need, and you produce it inside your own walls.

This is the deeper version of the thesis that business logic is migrating to a tier above the database. The CRUD layer — store a row, read a row — was never the value. The value was always the logic on top of it, and that logic is now something an operating system can generate, govern, and revise on demand. The application stops being a thing you acquire. It becomes a thing your OS *does*.

OSaaS, defined

OS as a Service is the recognition that the right unit of consolidation was never the application. It was the operating system underneath all of them.

OSaaS arrives as a desktop application — it can also run in the browser — but what you install is a window onto a living substrate, not a frozen product. Behind that window, applications are scaffolded on demand, governed from birth, and dissolved when their purpose ends. An app produced this way is not integrated into the organization's identity, data governance, audit, and policy regime — it is *born inside* them. A point solution is compliant *despite* its origins, through months of review and remediation. An OSaaS application is compliant *by construction*, because the policy boundary is the soil it grows in. There is nothing to bolt on, because there was never a moment when the app existed outside the boundary.

The principle underneath this deserves a name, because it is the law the entire architecture obeys. Call it **sovereignty**. The organization is a sovereign state. Inside its borders there is one identity, one data model, one policy, one law. The treaty languages — SAML, the app stores, external connectors — are spoken only at the border, with an outside world the organization does not control. CISOs already understand and prize *data* sovereignty. OSaaS extends it to the entire stack: **software sovereignty**.

The same principle plays out in three places, and in each one a category of work the old world treated as permanent simply disappears:

Identity stops being a federation problem. Inside the border there is no single sign-on to wire up, because there is no second system to sign on to; SAML survives only at the edge, where the

organization meets a bank, a state portal, or the identity provider of a company it is acquiring mid-merger.

The same app stops being rebuilt for every device. Intent is expressed once, and every surface is a *projection* of it — render targets, not separate builds. The OS emits a web app and a signed desktop binary, and when an application is published it automatically generates native Swift and Kotlin versions. For internal applications — the overwhelming majority — these reach staff through enterprise distribution: managed (MDM) delivery, or surfaced inside the organization's own UpliftOS container app on the devices it administers, with no per-app store submission and no separate mobile team. When an app is genuinely public, it ships through the stores like any other app — but through the organization's single developer relationship rather than hundreds of them. The stores still set the rules; the enterprise just meets them once, and only for the apps that actually face the public.

Integration stops being a job. Applications no longer ferry data between private stores, because there are no private stores to bridge; connectors to genuinely external systems survive only at the port of entry.

The pattern repeats on purpose. Everything that can live inside one trust boundary lives there; everything genuinely external is admitted through a single governed port of entry, on the organization's terms. That is not a feature list. It is one rule, applied three times.


One operating system, many populations

Sovereignty says there is one identity, one data model, one law inside the border. But a border does not imply a monoculture. Look at any real institution and you find not one kind of user but several, each standing in a completely different relationship to it — and the pattern holds with equal force across every institution this paper is written for. A hospital holds patients, clinicians, billing staff, and administrators, who share a building and almost never a screen. A government agency holds citizens, caseworkers, auditors, and field officers, each entitled to a different slice of the same record. A university holds students, faculty, registrars, and development officers — one needs financial aid and a degree audit, another lives entirely in gifts and pledges and donor histories. A large corporation holds employees, customers, finance, and the field, each in a separate software universe. Four kinds of institution, one structure: a single sovereign organization, and a half-dozen populations that have almost nothing in common except the border they share.

The point-solution era answered this with a different vendor for each population, or with one monolith that pretended a clinician and a patient, a caseworker and a citizen, wanted the same product with different checkboxes ticked. OSaaS answers it the way it answers everything else: with projection. Just as one intent projects onto web and desktop and mobile as render targets rather than separate builds, one operating system projects onto each population as a distinct surface drawn from the same substrate. An application is not merely built; it is *provisioned* — scaffolded once, then made visible, on the organization's terms, to exactly the population it serves and no one else. The patient sees the patient's OS; the auditor sees the auditor's. Neither is a different system; both are the same sovereign substrate, wearing the face the population in front of it is entitled to see.

This is where provisioning becomes the operational expression of sovereignty. Identity, inside the border, is not merely a login — it is a statement of which population you belong to, and therefore which

applications, which data, and (as we will see) which agents are yours to use. Add a population and you do not add a system; you add a projection. The OS that serves ten thousand members of the public and four hundred staff is not ten thousand four hundred accounts bolted onto one app. It is one operating system showing each person the institution as it pertains to them, and nothing of the institution as it does not.



Your website is just another app

Follow provisioning to its edge and something quietly radical falls out. If an application is provisioned to a population, then the most external population — the public, the anonymous, everyone who is not yet anyone to the organization — is simply a population like the others. And the application you provision to that population is the thing the enterprise currently buys from an entirely separate industry and calls its *website*.


The corporate website is not a different category of artifact. It is an app on the OS, scaffolded by the same pipeline, born inside the same governance, sharing the same substrate — provisioned to the population called *the public*. Every app on the system can project a web surface if it wants one; the public website is simply the app whose population happens to be everyone. The marketing site, the customer portal, the public service the institution offers the world — these stop being a separate CMS, a separate hosting relationship, a separate vendor with its own seam, and become what they always should have been: applications, owned and alive, on the operating system that runs everything else.

The line between "our internal software" and "our public web presence" was one more receipt for fragmentation. OSaaS dissolves it. There is no inside-software and outside-website. There is one OS, and one of its populations is the world.

The pipeline is the product

Here is the move that separates OSaaS from every low-code, no-code, and AI-app tool in the market: they consolidate the *application* layer. OSaaS consolidates the application layer **and the production layer**, in a single act.

The production layer is the sprawl no one counts, because it hides inside engineering. The penetration-testing firm engaged twice a year. The governance and compliance platform. The QA team and its end-to-end testing tooling. The DevOps platform team. The code-signing infrastructure. Each of these is a separate vendor, a separate cost center, a separate seam — not in *running* software, but in the act of *making* it. OSaaS folds all of it into the OS as the non-negotiable price of admission for every application it emits. Security operations, vulnerability scanning, compliance evidence, automated end-to-end testing, build and release — these are not services the enterprise assembles around its apps. They are stages every app passes through on the way out of the pipeline, every time.



The build-and-release pipeline. Every app passes through Plan, Build, Try, Secure, Comply, and Deploy. Deploy forks to a web app, a signed desktop binary, and automatic native mobile apps, and end-user feedback from every surface flows back into Plan.

It is worth being precise about what those stages do, because the reader this paper is written for will notice if we are not. The per-release pipeline runs static and dynamic analysis, dependency and secret scanning, and full end-to-end tests — automatically, on every build. **Penetration testing is a different discipline, and we treat it as one** rather than rebadging a scanner as a red team: SecOps, the OS's continuous adversarial-testing service, probes deployed applications on an ongoing basis — not once a year, and not as a checkbox at build time. Compliance is handled with the same honesty. SOC 2 and HIPAA are mostly organizational controls, not artifacts a single application can emit; what every release produces is *control mapping and evidence* — the dated, signed audit trail that turns the annual attestation from a fire drill into a formality. The pipeline does not make a company compliant by fiat. It makes every application carry its own evidence, which is the part software has never done before.


This is what kills the reasonable fear that generated software is unreliable software. An OSaaS application is not improvised and shipped. It is generated and then *interrogated* — tested, scanned, evidenced, and signed — before it is ever permitted to touch real data, and adversarially probed by SecOps for as long as it lives after. Generation is cheap. **Verification is the product.** The app is production-legitimate from its first minute of life, not because someone trusted the generator, but because the generator's output had to survive the pipeline.

And once you understand the pipeline as the product, you understand the deepest distinction in this entire paper: the difference between **dead software and living software.**

A point solution is dead the day you buy it. It looked alive in the demo and now it is frozen in the posture it held at signing. Every improvement is a petition to a vendor with four thousand other customers and a roadmap you do not control: six months if you are fortunate, "it's on our radar" if you are not. The enterprise owns a shelf of beautifully preserved specimens and pays a staff to keep them dusted.

An OSaaS application lives. Friction observed on Monday becomes a build request the same day, a fast revision through the same generation-and-verification pipeline that bore it, and a re-shipped, re-tested, re-signed application before the week is out. The enhancement loop closes in hours, not in vendor quarters. *Do not wait six months for a vendor to fix it. Get it tomorrow.* That sentence is the whole emotional argument of OSaaS, and it is only deliverable because the pipeline that makes apps is the same pipeline that keeps them alive.

Living software has a full lifecycle, and the third act is the one no incumbent can offer. An OSaaS app is free to be born — scaffolded, tested, secured, signed, cross-platform, trustworthy from minute one. It is alive while useful — metabolizing through the pipeline, improving on the organization's clock rather than a vendor's. And it is **free to die**. When you no longer need it, you delete it. There is no contract tail, no decommissioning project, no renegotiation, no vendor permission required. And deletion is not amnesia: the application goes, but its records remain — retained, governed, and queryable — for exactly as long as the law or the organization requires — the death of the app without the death of the data, which is the precise opposite of the migration projects that lose both.



The full lifecycle: born verified and signed, kept alive by re-shipping through the same pipeline in days, and free to die with no contract tail. Birth and death are both nearly free.

That third act matters more than it sounds. The dirty secret of every large enterprise is the shelfware it cannot turn off — tools no one opens, kept alive only because the contract outlives the usefulness and killing them means a migration project. The enterprise has lost the ability to perform a clean death. A healthy living system sheds what it no longer needs as a sign of health; the disease state is the thing that *refuses to die* and metastasizes across the organization, consuming resources, unkillable. Sprawl, seen clearly, is an organization that has lost the ability to let software go.

When birth and death are both nearly free, software stops being a *capital asset* — depreciated over years, committed to, defended because it was expensive — and becomes a *consumable*. Spin an app up for a single quarter's compliance push; delete it when the push ends. Software becomes cattle, not pets. That is a change not in tooling but in the entire relationship between an organization and its own software, and it is available only when the death is as costless as the birth — which is exactly the half no contract-based vendor can ever sell you.

Governed, transparent, and human in the loop

The word "generated" frightens enterprises for a reason: it conjures a black box that writes code no one reads and ships it where no one looks. OSaaS is the opposite of a black box, by construction. Generation is the starting point, not the authority. There is as much human in the loop as the organization wants — from fully autonomous builds for low-stakes internal tools to mandatory human authorship and review for anything that touches sensitive data. Everything the system produces is transparent and inspectable, and the generated code is not a sealed artifact but ordinary source that a human engineer can read, change, and own.

That control is formalized in a governance layer baked into the OS rather than bolted on around it. Specifications are approved before they are built. Deployments to staging and to production pass through explicit approval gates. Every release produces a control-mapped evidence trail. These are not workflows the enterprise has to assemble from a GRC platform, a CI/CD tool, and a ticketing system stitched together — they are native properties of the substrate, applied uniformly to every application it produces. Sovereignty is not only about where data lives; it is about who must say yes before software moves, and the OS makes that an explicit, auditable act rather than a tribal convention that varies by team.

What a build actually produces

An executive can take "transparent and governed" on faith; the reader this paper is written for cannot, and should not. So here, concretely, is what a single build emits — not a screenshot of a demo, but the shape of the artifacts every application leaves behind. A spec a human approved and signed. The schema generated from it, as ordinary source you can read and change. An append-only audit trail with every gate, scan, and approval dated and attributed. And the policy that scopes the agent permitted to act on it.

① SPEC — APPROVED

```
app:      recertification-tracker
owner:    Compliance
entities: Credential, Provider, Review
rule:     flag Review 90d before
          Credential.expiresAt
access:    Compliance(write),
          Provider(read-own)
approved: v3 · j.reyes · 2026-06-01
```

② GENERATED SCHEMA · PRISMA

```
model Credential {
  id          String  @id
  providerId  String
  type        String
  expiresAt   DateTime
  status       Status  @default(ACTIVE)
  reviews     Review[]
}
```

③ AUDIT TRAIL — APPEND-ONLY

```
09:14 spec.approved    v3 j.reyes
09:31 build.passed     sast+dast clean
09:33 deploy.gate      staging ✓ m.chen
02:00 secops.scan      0 criticals
—      every event signed + dated
```

④ AGENT POLICY — RECERT-COUNSELOR

```
population: Compliance
reads:      Credential, Provider,
            Review (knowledge graph,
scoped)
writes:     Review.status (gated)
autonomy:   advisory → approval-gated
guard:      logged · minimum-necessary
```

Illustrative — the artifacts a single build leaves behind: the approved spec, the schema generated from it, the append-only audit trail, and the policy scoping an agent. Representative of every build's output, not a screenshot of one.

None of these is exotic. The schema is Prisma a human can edit; the audit trail is the dated, signed evidence the compliance stage produces; the policy is the same per-population, minimum-necessary boundary that governs the knowledge graph and the agents alike. The point is not that any one of them is novel — it is that they exist, by default, for every application the OS produces, and that you can read them. "Verification is the product" is only a slogan until you can see the receipts. These are the receipts.

A new security posture

OSaaS does not merely match the security model of the software it replaces. It changes the model, in five ways that matter most to a security officer.

First, the default unit of software is a signed binary, not another web app exposed to the open internet. Most of what an organization builds is for its own people, and those applications ship as signed native binaries that run where the organization chooses — no public endpoint to probe, no sprawling client-side surface to exploit. A public web surface is the deliberate exception, not the default: when an app is genuinely meant for the world — the corporate website, a customer portal — it is provisioned to the public on purpose, and that one surface is hardened and watched as a public surface, instead of every internal tool carrying that exposure for no reason. The model inverts. Today every app is a web app and the binary is the rarity; in OSaaS the binary is the default and the public web surface is a choice made app by app, with eyes open.

Second, the organization controls where development and execution happen. Autonomous builds can run on an employee's laptop, inside a local container, or in the cloud — the enterprise decides where the code is written and where it runs. For regulated industries this is decisive: data residency and execution location become configuration, not a vendor's architecture you are forced to accept.

Third, ownership is immediate and mechanical. The enterprise connects its own Git provider, and the source lands in its own repository from the first commit. There is no escrow arrangement, no "request your code," no dependence on the vendor's continued existence to keep what you built. You own the source the moment it exists, in infrastructure you already control.

Fourth, isolation and unity coexist. Each application gets its own dedicated CRUD database, so a fault or a breach in one app is contained to one app — the blast radius is a single database, not the whole estate. Those isolated stores are unified only at the knowledge layer — the knowledge graph this paper turns to next. This is the architectural resolution of a tension that sits at the heart of the whole design: how an organization can know everything about itself without making everything reachable from everywhere. The data is *physically separated* for containment and *logically unified* for intelligence. The organization gets the blast-radius discipline of strict isolation and the perfect-information benefit of a single knowledge graph at the same time, because the two live at different layers.

Fifth — and this is the question a security officer asks within the first five minutes — availability is designed, not assumed. The honest worry about one operating system is that it is also one outage: lose the substrate and you lose the front line, the back office, and the public site at once. The answer has three parts, and none of them is faith. The substrate runs multi-zone on managed, replicated infrastructure with point-in-time recovery — the posture a serious enterprise already demands of its own cloud. Applications are independently deployable units rather than one monolith, so a fault in one app no more takes down the others than its breach could reach their data; the failure domains that data isolation draws are availability boundaries too. And the knowledge graph is a derived, rebuildable projection, not a system of record — if it falls behind, the apps keep serving from their own stores and it catches up. For the surfaces an organization truly cannot lose, the default-to-binary posture is itself the continuity story: a signed native client keeps working on a desk when the network does not. One OS is not one fragile thing. It is one well-understood thing, engineered for the failure modes that a sprawl of independent vendors never coordinated in the first place.

The knowledge graph: the organization that knows itself


Underneath the OS sits the piece that makes everything above it true rather than aspirational: a knowledge graph that absorbs the entire organization. Every chat, every document, every file is parsed into the knowledge graph. And — this is the part with no precedent — the database of every application scaffolded onto the OS is *part of the knowledge graph too*.

This is the place to answer the objection that buried every master-data program in history: a single coherent model over an entire organization's data is the dream that always died on contact with reality. It died because those programs tried to impose one ontology onto dozens of systems they did not build and could not change — reconciling one vendor's notion of "customer" with another's, forever. OSaaS is the one setting where the dream is tractable, because the OS authored the schemas in the first place. The knowledge graph does not reverse-engineer meaning out of someone else's databases; it maps stores whose shape it defined, on a vocabulary it controls, extended deliberately as each new app is

born. It is still real, ongoing work — an ontology is grown and curated, not declared into being — but it is work done with the grain of the data instead of against it, which is exactly why it can succeed here and could not there.

Recall the deepest cost of sprawl, the one beneath the integration bills: every application is a place where the truth about the organization goes to hide. Each silo makes the enterprise dumber about itself. Buy a new tool, and you add a new room the organization cannot see into. The knowledge graph inverts this completely. When every app's data is part of one substrate, **adding an application makes the organization smarter, not dumber**. Apps stop being silos and become sense organs — each one a new stream of perception feeding a single nervous system. No point-solution vendor can ever make that claim, because their business model *is* owning their slice of your data. In the old world every new app fragments what the organization knows. In OSaaS every new app deepens it.

A fair objection arrives here, and it is the sharpest one the architecture faces: if every app keeps its own isolated database and the knowledge graph unifies them, is the knowledge graph not simply iPaaS by another name — the very data-ferrying the whole thesis condemned, smuggled back inside the boundary? The answer turns on what integration actually costs. The expense and the seam in iPaaS were never *movement*; they were *reconciliation* — the endless work of forcing two writable systems, each convinced it owned the customer's address, to stop disagreeing. That cost exists only where the same fact has more than one home that can overwrite it. Inside the OS, every fact has exactly one authoritative home: the database of the one app responsible for it. The knowledge graph does not mint a second writable copy to reconcile against the first; it is a unified read-projection over stores that never compete, because none of them is a duplicate of another. Movement still happens, but it is one-directional projection into a knowledge layer, not bidirectional treaty-keeping between rival owners. iPaaS exists to settle disputes between systems that each claim the truth. The knowledge graph has no disputes to settle, because under sovereignty the truth has one address. That is the categorical difference, and it is why integration genuinely retreats to the perimeter rather than merely changing its uniform.



Physical isolation, logical unity: every application keeps its own store; the knowledge graph reads upward into one projection, and data never moves sideways between apps — so there is nothing to reconcile.

The knowledge graph is also what turns "self-improving OS" from a slogan into a mechanism. How does the system know which applications to harden and which to let die? The knowledge graph sees all usage. How does the agent know what enhancement to build next? It has near-complete context. Self-improvement was magic when it was a vibe; it is *mechanical* when it is perfect information feeding a build pipeline. Knowledge graph plus pipeline is a closed loop — the organization perceives its own friction and revises itself in response, without a roadmap and without a vendor.

The ambition is an agent with as close to perfect information about the organization as can be achieved. And here the architecture must confront its own most dangerous sentence, because the property that thrills a CEO terrifies a CISO — and in regulated industries it is not merely frightening, it is a legal problem. Healthcare runs on minimum-necessary access and purpose limitation; these are not preferences but obligations. A single knowledge graph that ingests every conversation, every file, and every application's records, queryable by one all-seeing agent, is the largest blast radius a security officer has ever been asked to approve.

Sovereignty is the answer, and it converts the scariest property of OSaaS into its strongest. The knowledge graph is not *raw* omniscience. It is **governed omniscience**. The same policy boundary that makes applications born-compliant governs what the agent may see — per user, per purpose, per context. The agent is not omniscient to everyone; the *organization* is omniscient about itself, and every actual query is scoped, logged, and minimum-necessary by construction. This is a genuinely novel claim, and it is the one that wins regulated markets instead of losing them: **OSaaS is the first architecture in which perfect information and perfect access control are the same mechanism**. In a sprawled estate, every silo is simultaneously a knowledge gap and an ungoverned access risk. Sovereignty closes both at once. The knowledge graph knows everything, and because it is one governed substrate, every grain of that knowledge is policy-bound. And as the security model makes

concrete, the knowledge is unified while the underlying data stays physically isolated per application — so omniscience never requires a single store large enough to lose all at once.

Agents, provisioned

The knowledge graph gives the organization ambient intelligence about itself. Agents are how that intelligence is aimed.

Treat an agent the way OSaaS treats everything else worth governing: as a first-class entity an administrator creates, configures, and provisions to a population — exactly as they provision an application. An agent is not a single omniscient oracle bolted to the top of the company. It is a thing you make. You give it a name and a purpose. You scope the knowledge it may draw on. You set how much autonomy it is permitted — from advisory, where it can only suggest, through approval-gated action, to fully autonomous operation inside a bounded domain. And then you provision it to the people it is meant to serve. A clinician's agent, a caseworker's agent, and a finance team's agent are different beings with different knowledge, different powers, and different reach, because they were configured to be, by someone accountable for the configuration.

What makes these agents categorically different from the chatbots stapled onto point solutions is the substrate beneath them. Every application scaffolded on the OS has its own dedicated database, and those databases are not sealed behind a vendor's API built to keep you out. They are addressable. An agent provisioned on the OS is given a *raw, governed connection to the databases of the applications in its scope* — through MCP tools that let it read and act against those stores directly, not through a brittle integration negotiated after the fact, but through a native channel that exists because the OS built both the app and the agent in the same world. The agent does not "integrate with" the recertification app the way an outside bot would. It reaches into the recertification app's database, because inside this sovereignty there is no border between them to cross.

This is the moment "agent" stops meaning *a thing that answers questions* and starts meaning *a thing that operates the organization*. An agent with a direct line to the databases of the apps it serves, scoped knowledge from the knowledge graph, and a configured degree of autonomy is not a search box. It is a provisioned operator — a member of staff made of software, who knows what it is permitted to know and may do what it is permitted to do.

And here the same sentence that thrilled the CEO returns to terrify the CISO, sharper than before: an agent with raw access to production databases is, on its face, the most dangerous thing in the building. Sovereignty is again the answer, and again it converts the fear into the strongest part of the design. Raw access is not ungoverned access. The connection an agent holds is bound by the same policy boundary that makes applications born-compliant and the knowledge graph governed rather than omniscient — per population, per purpose, per autonomy setting, every action logged and minimum-necessary by construction. The agent reaches directly into the database not because the door stands open, but because the OS holds the key and decides, on every reach, whether to turn it. Direct is not the opposite of governed. In OSaaS, direct *is* governed — the same mechanism is both.

One structural question follows from all this isolation, and it deserves a straight answer: if every fact has exactly one home and the knowledge graph only reads, what happens when a real workflow must write across two apps at once — close the case in one and post to the ledger in another? The OS mediates the

write; a shared database does not. There is no distributed transaction to get wrong because there is no shared store to contend on. The agent or the workflow writes to each app's own home in sequence, inside the same policy boundary, idempotently and fully logged, with the knowledge graph reflecting the result a moment later. Cross-app coordination lives at the OS layer, where it is governed and audited as a single act — not smeared across point-to-point integrations that no one team owns.

Ownership, and the real moat

A natural fear attaches to anything that calls itself an operating system: that the enterprise is trading hundreds of vendors for one super-vendor who now owns its entire operational layer. OSaaS detonates that fear by inverting the ownership model. **The enterprise owns the applications it creates.** Not licenses them. Owns them. You are not the landlord's tenant; you are the holder of the deed. Every other consolidation play asks you to surrender your stack. OSaaS hands it to you. And as the security model details, that ownership is mechanical rather than contractual — the source lives in the organization's own Git repository from the first commit, with no escrow to invoke and no vendor whose survival you depend on.

This is safe to give away precisely because the applications were never the valuable thing. They are the *exhaust*. The value is the pipeline that keeps them alive and the knowledge graph that lets the organization understand itself. So the structure is honest and simple: you own the apps, and you subscribe to the pulse. The moment an enterprise forks its owned applications and walks away, those apps stop metabolizing — no more enhancement-tomorrow, no more automatic re-test when a vulnerability lands, no more compliance staying current as regulation shifts. The organization keeps its software and loses its heartbeat, and the software begins decaying back into the point solutions it replaced.

And to be clear about the one place a sharp reader expects a catch: the knowledge graph is yours too. Not just the applications, and not just the raw records — the accumulated, interconnected self-knowledge the substrate spent years deriving. You can export it, and you can take it. We do not hold an organization's intelligence hostage as the price of having trusted us with it; doing so would make a mockery of every claim of sovereignty in this paper. If you leave, you leave with everything — the apps as source, the data in open formats, and the knowledge graph that sits on top of them.

What you give up by leaving is not access to anything you own. It is the pulse. A knowledge graph carried out of the OS is a snapshot the moment it crosses the border; inside, it is alive — re-derived as every new app ships, every document lands, every agent acts, kept current by the same pipeline that keeps the applications current. The moat was never captivity. It is that everything keeps compounding only while it stays connected to the engine that compounds it. Lock-in you impose is a cage. What OSaaS offers is the opposite: take it all whenever you want — and stay because leaving costs you the living thing, not because you were ever trapped.

The network

Sovereignty is a border, not a wall. The point of drawing one is to choose what crosses it — and the most valuable things that cross a border are not data dumps but relationships. An organization running on OSaaS is not an island; it is a sovereign node on a network of other sovereign nodes, and the same port of entry that admits a bank or a state portal can admit another organization on the platform.

What crosses is more than files. One organization can discover and use another's applications and agents, on terms both sides set — a supplier's ordering app surfaced inside a buyer's OS, a partner's compliance agent invited to answer a narrow question against a scoped slice of the knowledge graph, a referral handed from one institution's workflow straight into another's. People cross too: an individual on the network can interact with an organization's apps and agents without being inside it — the way the public already meets the website, except now the surface is anything the organization chooses to provision outward. The network is sovereignty's positive form: not what you keep out, but what you let in, deliberately.

And there is an economy on it. Because an organization owns the applications it builds — really owns them, as ordinary source in its own repository — it can do the thing no SaaS tenant can: package one and sell it. An app that solved a problem once becomes a product other organizations install, and the internal tool that used to be a cost line becomes a revenue line. The marketplace is where that happens — the place a company turns the software it had to build anyway into something its peers will pay for.

The buyer's protection is the same mechanic as the builder's freedom: the option to fork. Install a marketplace app and run it as shipped, or — because what you bought is real, ownable source — fork the codebase and develop your own implementation, extending it, re-skinning it, bending it to a workflow the original author never imagined. There is no version of this where you are trapped on someone else's roadmap, because the roadmap is a repository you hold a copy of. The lock-in that defines every other software marketplace simply does not exist here: you buy the work, and you keep the right to make it your own.

The model, named

This paper has named its asterisks — the knowledge graph that does not travel, the lock-in that is earned rather than imposed. There is one more, and the reader has been waiting for it since the word "generated" first appeared: the OS does not write applications by magic. It writes them with a frontier model, and that model is the one part of the stack the organization does not itself produce. A paper about sovereignty that declined to mention its own upstream dependency would deserve the suspicion it earned. So, plainly.

Intelligence runs in two places, and they are not the same kind of dependency. At **build time**, a frontier model generates the application — and this is a dependency the way a compiler is a dependency: you need it to produce the artifact, and you do not need it to run the artifact. The output is ordinary owned source in your own repository. It compiles, deploys, and runs with no model in the loop and no call back to anyone; if every model vendor vanished tomorrow, the applications already built would keep running unchanged. The codegen brain is the foundry, not the landlord — and the deed it stamps is yours.

At **run time**, intelligence reappears wherever agents think, and here sovereignty is a configuration rather than a compromise. Execution location is the organization's to set: hosted by the platform for convenience, or pinned inside the organization's own cloud account and boundary for the regulated case — inference running in infrastructure the enterprise already controls, under model terms it has already accepted, with prompts and outputs that never cross the perimeter. The most sensitive institutions can hold the entire loop inside their own walls. The point is not that one model is blessed forever; it is that the dependency is named, bounded, and placed where the organization chooses. It is the same answer OSaaS gives for data — your cloud, your keys, your boundary — applied to cognition.

The objections, met in advance

Isn't generated software inherently risky? Generation is the starting point, not the authority. There is as much human in the loop as the organization wants, the code is fully transparent and editable by hand, and nothing reaches real data without passing spec approval, automated testing and scanning, compliance evidence, and explicit staging and production gates — and once live it is probed continuously by SecOps. The app is interrogated — and, where the organization requires it, authored and reviewed by people — before it is trusted.

Who is accountable when a generated application is wrong? The enterprise that runs it — as with any software it operates — but for the first time with complete provenance and a real commercial counterparty standing behind the tool that built it. Every application's birth, every revision, and every agent query lives in the knowledge graph, logged and auditable, rather than scattered across forty-five vendors' support queues. And because UpliftOS is a product rather than a raw model you prompted alone, the relationship carries what a CIO actually needs when something breaks: named support, defined SLAs, and contractual accountability — liability and indemnification for the platform that generates and verifies the software — negotiated like any other enterprise agreement, not the all-warranties-disclaimed posture of an API key. The generated source is yours and editable; the platform that produced it is a vendor you can hold to terms.

An agent with a direct line to a production database — isn't that the most dangerous thing in the building? It would be, if the access were ungoverned. It is not. An agent's raw connection is scoped by the same policy boundary that governs every app and every knowledge-graph query — per population, per purpose, per configured autonomy level, every action logged and minimum-necessary. The agent reaches directly into the store because the OS holds the key and decides, on each reach, whether to turn it. Direct access and governed access are the same mechanism here, not opposites.

What about systems of record we are legally required to maintain? Sovereignty meets them at the port of entry. The OS does not pretend to replace the bank, the clearinghouse, or the statutory system of record; it speaks the treaty languages to them at the perimeter, in one governed place, on the organization's terms.

Is the omniscient knowledge graph a privacy catastrophe? Only if it is ungoverned, and it is not. Governed omniscience makes perfect information and minimum-necessary access the same mechanism. The knowledge graph is the most controlled data surface the enterprise has ever had, not the least.

The path to OSaaS

None of this requires a forklift, and any CIO who has survived two "end state" pitches a year is right to distrust one that does. OSaaS is not a migration scheduled for a weekend. It is a wedge driven into the part of the estate that hurts most and defends itself least, and it earns the next step by surviving the last one.

It starts at the long tail. Every enterprise runs dozens of internal point solutions and homegrown tools that no one loves, no one will miss, and no incumbent will fight to keep — the recertification tracker, the vendor-risk spreadsheet that quietly became a system, the onboarding workflow stitched from three

SaaS seats. These go first, because they are the cheapest to replace and the fastest to prove, and net-new apps start here too: the thing the organization has wanted and never built because the old pipeline made it a twelve-month project. First value lands in weeks, against software that costs real money and inspires no loyalty.

There is no single on-ramp, because the right first move depends on what the organization needs to prove — so there are four, and you pick the one that fits. The first is a **quick win**: one team, one workflow, a visible result in four to six weeks — replace a single Airtable base and bank the saving, a proof point more than a value driver. The second is a **safe pilot**: ship the net-new app a team has been asking for and could never get built, retiring nothing and risking nothing already in production. The third is **clean economics**: retire a commodity vendor — the form builder, the ticketing tool, the wiki — where the savings are obvious and nothing political protects the incumbent. And the fourth, for the bold, is to go **straight at the fire**: the workflow everyone complains about, the renewal everyone dreads, replaced before it comes due. Quick win, safe pilot, clean economics, or straight at the pain — the point is not which one you choose. It is that every one of them is small enough to start now and prove before it is ever bet on.

And there is a fifth path, for the software a company has already built: bring it with you. Ownership runs in both directions — an organization can connect an existing codebase and bring it under the OS rather than rebuild it from nothing. **HeyMedicaid** is the proof. It is a real web and mobile application, built the old-fashioned way, by hand — born out of deep collaboration with Children's Nebraska, long before any of this — and it was *ported in*, not reconstructed. Once inside, it inherits what the pipeline gives everything else: the verification, the knowledge graph, the agents, the cross-platform projection, the path to ownership. Bring-your-own-repo is not a slide. It is how applications with a history join a platform built for their future. And the arc does not stop at porting: HeyMedicaid will be marketed on UpliftOS itself, offered to other hospitals and state governments across the network — a tool built by hand for one institution becoming a product the marketplace carries to many.

The systems of record wait at the perimeter, and may wait forever. The ERP, the core banking system, the statutory registry — OSaaS does not race to rip these out. It meets them at the port of entry, speaks their treaty languages, and lets them keep doing the one job they do well, while the workflows, portals, and agents that actually touch people move inside the boundary one at a time. Workday does not die on day one; it may not die at all. It simply stops being the center of gravity and becomes one more system the OS reconciles against at the edge.

So the estate and the OS coexist, deliberately, for years — and that is not a weakness in the plan, it is the plan. Sprawl took a decade to accumulate and will not be unwound in a quarter; any vendor who promises otherwise is selling the next near-miss. The honest sequence is unglamorous, which is exactly why it works: start where it is safe, prove it where it is visible, expand where it earns trust, and let the things that must outlive the transition outlive it. The end state is a destination, not a switch — and the first step is small enough to take on Monday.

Not only the enterprise

This paper has spoken in the register of the large institution — hundreds of vendors, ten thousand people, a portfolio that took a decade to sprawl. But nothing in the architecture requires scale, and the most surprising consequence of OSaaS is what it does at the other end of the spectrum.

The floor is one person. The same engine that lets a hospital retire a vendor estate lets an individual or a five-person company build and run software that used to require a department to commission and a budget to defend. A founder can stand up the apps a business needs — the portal, the workflow, the billing, the back office — in the time it once took to write the RFP. And then the multiplier: agents. A small organization does not hire a team for every function; it *provisions* agents to operate them — a finance agent, a support agent, an operations agent, each with scoped knowledge from the knowledge graph and a configured degree of autonomy. The leverage that historically belonged only to the large — an institution's worth of software, run by a handful of people — becomes available to anyone, and it hyperscales not by headcount but by the number of agents provisioned and the apps those agents drive.

This is the quiet inversion underneath the whole thesis. For thirty years, software economics rewarded size: the bigger the organization, the more it could afford to build and the more leverage it could buy. OSaaS breaks the correlation. When building is nearly free and operating is delegated to agents, a small team on a sovereign platform can carry the surface area of a large one — and join the same network, sell into the same marketplace, and own what it makes on the same terms. The end state is not only the enterprise that stops renting. It is the individual who can finally build and run what only an enterprise once could.

The end state

Seen whole, it is one organism: sovereignty the border, the knowledge graph the nervous system, the pipeline the metabolism that turns friction into repair, and the freedom to delete the way a healthy body sheds what it no longer needs. Point-solution sprawl, in that frame, is an organism whose organs have stopped talking to its brain — and which has forgotten how to let any of them go.

The end state of the enterprise is not a better-integrated portfolio. It is the disappearance of the portfolio as a category. Not hundreds of applications wrangled into cooperation, but one sovereign operating system that grows exactly the applications each population needs, provisions them to the people they serve, keeps them alive on the organization's clock, lets them die when their purpose is served, and grows smarter about the organization with every one it produces.

This is the thesis behind UpliftOS: that the right number of vendors is not fewer, but that the right unit of software is no longer the vendor at all. The enterprise stops buying software. It starts owning an operating system that makes its own.

A final musing: the last landlord

Permit the authors a closing thought, offered as musing rather than argument. Follow this paper to its end and one landlord remains — the most intimate of them all. We have moved identity, applications, data, agents, and the pipeline inside the boundary. And yet the signed binary that results still runs on a laptop and a phone whose operating system, whose app store, and whose silicon answer to a consumer-platform vendor. It is the last rented endpoint, and it sits closer to the work than any SaaS tool ever did. The treaty languages reach all the way to the glass in your hand — and stop there.

Having drawn the boundary this far, it is hard not to ask the obvious question: why does it stop at the screen? An organization that owns its operating system, its applications, and its own intelligence has

every reason to want the device those things live on to belong to the same sovereignty — a laptop and a phone that are not consumer products with a mobile-device-management policy bolted on after the fact, but endpoints of the OS itself, where the apps an organization grows, the agents it provisions, and the keys it holds are native to the hardware rather than guests on someone else's.

That is the horizon, and it is as far as this paper will go in describing it. The end state of enterprise technology is not only software a company owns instead of rents. It is a stack that is sovereign all the way down — from the spec a person approves, through the pipeline that makes it real, to the device in their hands. *More on that soon.*

© 2026 UpliftOS · OSaaS: The End State of Enterprise Technology · by Mark McNasby & Michal Oglodek