Open eLMS OS
All-in-one Software Solution for Business

The Business Case for Open eLMS OS

Why replacing a stack of SaaS subscriptions with one AI-powered operating system cuts cost, risk and wasted time, and how to work out the numbers for your own organisation.

A briefing for boards, finance and IT leaders

Open eLMS Ltd · Established 2003 · ISO 27001, ISO 9001 & ISO 20000 certified

October 2026

1. Executive summary

Most organisations didn't plan their software estate; it grew. A CRM here, a project tracker there, a marketing tool, an intranet, an LMS, and now an AI add-on for each of them. Every one of those tools came with a sensible business case. Together, they've become one of the fastest-rising costs in the business.

$4,830
average SaaS spend per employee per year, up 21.9% in a year[1]

11.4%
annual SaaS price inflation, against 2.7% general G7 inflation[3]

46%
of SaaS licences go unused on average[2]

~4 hrs
a week per person lost reorienting after switching between apps[5]

Open eLMS OS is an AI-powered operating system for the whole business. It brings the knowledge base, project management, sales, marketing, finance and learning into one platform with one database, one login and one AI layer that understands your organisation.

The case in brief

- **The biggest return: productivity.** Work that took days takes minutes. Marketing materials are produced at a single click, and an RFP that takes a team 33 hours on average[12] is answered in about an hour from the company's own knowledge. Staff learn one interface and one way of working with AI, raising productivity 10X on key tasks. In our worked model, time saved is worth over 85% of the total benefit.
- **Lower, more predictable cost.** One subscription replaces several per-seat licences and their annual uplifts. Organisations save $100Ks a year in software costs, and up to 99% on AI subscriptions.
- **AI that actually reaches across the business.** 95% of IT leaders say integration is a barrier to using AI well.[4] An AI operating system removes the barrier rather than building around it.
- **No wasted AI tokens.** One platform for generating with AI means knowledge and work are reused, not regenerated in silo after silo. One AI account runs and extends the whole system, instead of every employee paying for their own and building tools only their small group can use.
- **Software that fits you.** Open eLMS OS can be customised after implementation, so it follows your processes instead of a vendor's roadmap.
- **Less risk.** One ISO 27001 certified supplier and one data boundary, instead of dozens of vendors to vet and audit.
- **Time back.** Fewer tools to switch between, one place to find answers, and AI agents that carry work from one step to the next.

2. The problem: SaaS sprawl

The bill keeps growing

Zylo's 2025 SaaS Management Index, based on more than 40 million licences and $40 billion of SaaS spend, found average spend of **$4,830 per employee per year**, a 21.9% rise in twelve months and the first increase in both spend and app count since 2021. Zylo puts the rise down to vendor price hikes, more complex licensing and growing AI investment.[1]

Prices themselves are climbing faster than almost anything else on the balance sheet. Vertice's SaaS Inflation Index recorded an **11.4% year-on-year rise** in SaaS pricing in January 2025, compared with 2.7% general inflation across the G7.[3] Renewals rarely go down.

Much of it isn't used

Zylo reports that organisations leave **46% of their licences unused** on average: people change roles, projects end and tools are adopted faster than they're governed.[2] Every separate tool is another set of seats to over-buy.

The estate is too big to join up

MuleSoft's 2025 Connectivity Benchmark, a survey of 1,050 IT leaders, found organisations running an average of **897 applications**, with 45% using 1,000 or more.[4] Each one keeps its own copy of customers, projects and documents, so the same information is entered, synchronised and argued over many times.

People pay the "toggle tax"

Research published in Harvard Business Review followed workers at Fortune 500 companies and found they switched between apps and websites **nearly 1,200 times a day**. Reorienting after each switch cost just under **four hours a week**, roughly 9% of their working time.[5]

Risk is spread across every vendor

IBM's 2025 Cost of a Data Breach Report puts the global average breach at **$4.44 million**. Organisations with high levels of "shadow AI" (AI tools staff adopt outside IT's control) paid around $670,000 more per breach, and 63% of breached organisations had no AI governance policy or were still writing one.[8] Every extra tool, and every extra AI add-on, is another place for data to leak.

3. Why bolting AI onto SaaS doesn't fix it

Agentic AI (AI that doesn't just answer questions but carries out multi-step work) is arriving fast. Gartner predicts that by 2028 **33% of enterprise software applications** will include agentic AI, up from less than 1% in 2024, and that at least **15% of day-to-day work decisions** will be made autonomously by AI agents.[6]

But an agent is only as useful as what it can see and do. A copilot inside a CRM knows about deals; it knows nothing about the support ticket, the overdue project or the policy document that would change its advice. In a SaaS stack, every vendor is adding its own AI, each charging separately and each working from a narrow slice of your data.

The same Gartner research warns that **over 40% of agentic AI projects will be cancelled by the end of 2027** because of escalating costs, unclear business value or inadequate risk controls.[7] The usual culprit is the plumbing: 95% of IT leaders say integration is a hurdle to implementing AI effectively.[4]

> **The operating system approach**
>
> Instead of connecting many tools and hoping their AI can see across them, an AI operating system puts the work and the data in one place first. The AI then works from your whole organisation's knowledge (documents, emails, tickets, projects, customers, courses) and can act across modules, with one set of permissions and one audit trail.

One place to spend AI tokens

In a SaaS estate, AI spend multiplies. Every tool charges for its own assistant, and every team that builds something new starts from scratch: re-explaining the business to the AI, regenerating content another department already has, and paying for the tokens each time. Developer licences multiply the same way, with each department running its own accounts and its own half-finished projects.

Open eLMS OS is a platform for building as well as using. Knowledge, prompts and outputs live in one place, so they're shared and reused rather than regenerated. Development is focused too. In most companies, every employee who builds with AI has their own account, a Claude or ChatGPT seat each, and develops in a silo: the tool or workflow they make helps them and perhaps a few colleagues, and the next person pays to build it again. In Open eLMS OS, a large operation is run and extended from a **single AI account**, and whatever is built on it reaches every user in the company, saving thousands in seats and token use for even the smallest department. Open eLMS built, manages and expands the entire platform this way, on one development licence.

4. What's in Open eLMS OS, and how it compares

The modules

Module	What it replaces or does
AI-powered knowledge base	A secure "company brain" built from your procedures, documents, live data feeds and learning, which every other module and AI agent draws on.
Project management	Kanban boards, dependencies, automatic tickets and AI task breakdowns, connected to the knowledge base.
Sales	A CRM shaped around your own sales process, with contact discovery and outreach tracking from lead to conversion.
Marketing (SMART)	Content discovery, SEO, scripts, thumbnails and multi-platform scheduling across LinkedIn, Reddit and X.
Finance	Payments through Stripe and automated affiliate commission structures.
Learning	The Open eLMS AI-powered LMS and Learning Generator, so staff training is part of the same system.
Enterprise tools	Visual editor, role-based access control, media libraries, API and single sign-on.

SaaS stack vs Open eLMS OS

The same comparison as on our website, with the figures referenced at the end of this report.

	Typical SaaS stack	Open eLMS OS
Productivity	Each tool speeds up only its own step, so a campaign, a tender or a new web page is still days of work handed between people and apps. A single RFP response takes a team of nine 33 hours on average.	The biggest saving of all: work that took days takes minutes. Marketing materials at a single click, an RFP answered in an hour, and every module working from the same company knowledge, raising productivity 10X on key tasks.
Cost	A per-seat licence for every tool, often paying for the same person ten times over. The average organisation now spends $4,830 per employee a year on SaaS.	One platform in place of separate CRM, project management, marketing, knowledge base and learning tools, saving organisations $100Ks in software costs.
Training	Every tool is a different paradigm to learn: separate training for each, retraining whenever a vendor redesigns, and new starters facing the whole stack at once.	One interface and one way of working with AI to learn. The same skills then rebrand a website, create a social media campaign or build a CRM.
Price rises	Renewals arrive with uplifts: SaaS prices rose 11.4% in a year, roughly four times general inflation.	One contract and one renewal to plan for, instead of dozens landing at different times of the year.

	Typical SaaS stack	Open eLMS OS
AI	An AI add-on bolted onto each product, each with its own fee, and each only able to see its own corner of your data.	AI built into every module and grounded in your own company knowledge, cutting AI subscription costs by up to 99%.
Agentic AI	Copilots stop at the edge of their own app, so a person still carries the work from one tool to the next.	AI agents work across the knowledge base, sales, projects, marketing and finance, so a task flows from lead to project to report without the copy and paste.
AI token spend	Every team and every tool pays for its own AI, and tokens are spent again and again rebuilding context and repeating work another silo has already done.	One platform for generating with AI: knowledge, prompts and outputs are shared and reused, so no tokens are wasted on siloed or repeated projects.
Building & expanding	Every employee who builds with AI needs their own AI account, such as a Claude or ChatGPT seat, and what each one builds stays in a silo that helps only them or a small team.	One AI account runs and extends the whole system, and whatever is built on it reaches everyone in the company, saving thousands in tokens and seats for even the smallest department.
Customisation	You get the vendor's roadmap. A change you need waits in a queue shared with every other customer, if it is accepted at all.	Shaped around how you already work, and customisable after implementation: a custom CRM or project board can be built in an hour.
Extendability	A new need means a new tool: another contract, another login, another data silo.	Add a module or build a new tool inside the same system, on the same data, with the same users and permissions.
Integrations	Connectors and middleware to keep tools in sync, each one another cost, and any vendor update can quietly break one.	The modules share one database from the start, and the platform still connects outwards through its API, SSO, HubSpot and Stripe.
Knowledge & data	Information scattered across apps, with no single source of truth and staff chasing data between systems.	One AI-powered knowledge base feeds every module, so answers are consistent and up to date across the business.
Support	Many suppliers, contracts and SLAs, and problems bounced between vendors who each blame the other's integration.	One supplier, one support desk and one SLA for the whole system.
Security & compliance	Your data spread across many vendors, each one to vet, audit and manage access for.	One ISO 27001 certified platform with role-based access control, helping you meet ISO requirements in minutes.

5. The biggest return: productivity

Licence savings are the easiest benefit to count, but they aren't the largest. The largest is what happens to people's work when the whole business runs on one system, with one AI that already knows the organisation. Work that used to take days of drafting, chasing and handing over between tools takes minutes.

Where the time goes today, and what changes

Task	Typical time with a SaaS stack	With Open eLMS OS	Time saved
Respond to an RFP or tender	33 hours per response, with nine people contributing on average. Organisations answer only 55% of the RFPs they receive, and half name bandwidth as a top challenge.[12]	A full draft in about an hour, written from the company knowledge base (policies, past answers, certifications, case studies) and checked against confirmed company facts before review.	~95%
Marketing materials	3 hours 20 minutes to write one blog post,[13] before images, SEO, social posts and scheduling in separate tools.	At a single click, one video becomes a blog post, LinkedIn, X, Facebook, Instagram and TikTok posts, an email newsletter, a carousel and thumbnail hooks, with Shorts cut and everything scheduled and posted automatically.	~90%
Product and explainer videos	Weeks through an agency for each video, and again whenever the product changes.	Walkthrough and explainer videos scripted, voiced by AI and rendered inside the platform, and re-rendered when the product changes.	Weeks to hours
Rebrand a website	Typically two to three months of agency and page-by-page CMS work, with the live site changing piece by piece.	A complete rebranded copy of the site built alongside the live one, compared side by side, then switched over in one step. Open eLMS rebranded its own website this way.	Months to days
Build a CRM or project board	Select a SaaS product, configure it, train staff and then work the way it allows.	Shaped around your own process and built in an hour, on the same data and users.	Weeks to an hour
Find information	Searching for and gathering information takes around 19% of knowledge workers' time.[14]	One AI-powered knowledge base answers from the organisation's own approved content, in every module.	~25% (cautious)
Answer customers	Agents search several systems for each answer.	AI drafts answers from the knowledge base; in a study of 5,179 support agents, AI assistance raised issues resolved per hour by 14%, and by 34% for newer staff.[15]	14–34%

Task	Typical time with a SaaS stack	With Open eLMS OS	Time saved
Press releases and reports	Hours of drafting, finding images and chasing figures from several systems for each one.	A press release drafted from a short brief with images and press contacts matched; weekly work, sales and compliance reports (such as an ISO 27001 risk register and Statement of Applicability) kept up to date automatically.	Hours to minutes
Meetings to actions	Someone writes up the notes, then copies actions into a project tool by hand.	Meeting notes become tasks on each person's list automatically, linked to the right project.	Admin removed

Putting a number on it

For the 250-person organisation in the worked model (section 6), cautiously:

Productivity gain	Assumption	Hours a year	Value a year
Faster information finding	A quarter of the 19% of time spent searching,[14] on a 40-hour week	21,850	$764,750
RFP and tender responses	24 a year, from 33 hours[12] to 2 hours including review	744	$26,040
Articles and marketing materials	One a week, from 3 hours 20 minutes[13] to 30 minutes	147	$5,145
Faster AI-suited work	25% faster[11] on one day a week of AI-suited work	23,000	$805,000
Productivity gain, every year		**45,741**	**$1,600,935**

That is roughly 25 full-time people's worth of time, every year, before counting the tenders that can now be bid for at all. Videos, website work and custom tools are left out because they replace agency spend as much as staff time.

Why it compounds

A field experiment by Harvard Business School and BCG with 758 consultants found that, on tasks suited to AI, people using it completed **12.2% more tasks**, worked **25.1% faster** and produced work rated **40% higher in quality**.[11] Those gains apply to each task where AI is used well; an AI operating system puts AI in reach of nearly every task, with your organisation's knowledge behind it. And because staff learn one interface and one way of working with AI, rather than a new paradigm for every tool, the skill carries across marketing, sales, projects and documents. US organisations alone spent $98 billion on training in 2024,[10] and a share of every training budget goes on teaching people to operate software rather than to do their jobs better.

> **Why this is the largest benefit.** A licence saving happens once a year. A productivity gain happens every working day, for every person, on every task the platform touches, and it grows as staff get more fluent and more of the business moves onto one system.

6. A worked financial model

Every organisation's numbers are different, so here is an illustrative model you can repeat with your own. It assumes a **250-person organisation** spending the average $4,830 per employee on SaaS, and is deliberately cautious: it counts only a quarter of the SaaS bill, only a quarter of the time lost to switching and to searching for information, and the speed gain on just one day a week of AI-suited work.

Assumption	Illustrative value	Basis	Your figure
Employees	250	Example organisation	
SaaS spend per employee	$4,830	Zylo 2025 average[1]	
Total annual SaaS spend	$1,207,500	250 × $4,830	
Share of SaaS spend Open eLMS OS replaces	25%	CRM, projects, marketing, knowledge base, LMS, AI (cautious)	
Price rise avoided on that share	11.4%	Vertice index[3]	
Hours a week regained from fewer tools	1	A quarter of the ~4 hrs lost[5]	
Hours a week of AI-suited work	8	One working day	
Speed gain on that work	25%	HBS / BCG study[11]	
Share of searching time saved	25%	A quarter of the 19% of time spent searching[14]	
RFP responses a year	24	33 hours each,[12] 2 hours with Open eLMS OS	
Articles and marketing pieces a year	52	3 hrs 20 mins each,[13] 30 mins with Open eLMS OS	
Working weeks a year	46	After leave and holidays	
Loaded cost of an hour	$35	Example rate	

Gross benefit	Year 1	Year 2	Year 3	3-year total
Productivity (section 5)	$1,600,935	$1,600,935	$1,600,935	$4,802,805
Licences retired	$301,875	$301,875	$301,875	$905,625
Price rises avoided	–	$34,414	$72,751	$107,165
Time from fewer tools	$402,500	$402,500	$402,500	$1,207,500
Total gross benefit	**$2,305,310**	**$2,339,724**	**$2,378,061**	**$7,023,095**

Net benefit = gross benefit less the Open eLMS OS subscription, implementation and data migration from your quote. The model is illustrative, not a quote or a guarantee; ask us and we'll fill it in with you. Open eLMS made this switch itself, replacing around $100,000 worth of its own software.[9]

Not counted, but worth adding: training on one system, not many; AI subscriptions in one AI layer (up to 99% lower); no tokens spent on siloed or repeated projects, and one AI account for the whole system, not one per employee; integration middleware; unused seats (46% of licences[2]); lower breach exposure;[8] and one contract and support desk.

7. Implementation and managing the risks

A phased move, not a big bang

- **Discover (weeks 1–2).** List your current tools, their cost, renewal dates and who uses them. Pick the two or three with the highest cost and lowest satisfaction.
- **Build the company brain (weeks 2–4).** Load procedures, documents and key data into the AI-powered knowledge base so the AI understands your organisation from day one.
- **Replace in waves (months 2–6).** Move one function at a time (for example projects, then sales, then marketing), timing each wave to a SaaS renewal date so nothing is paid for twice.
- **Customise and extend (ongoing).** Shape modules to your processes and add new tools inside the platform as needs arise.

Risks and how they're handled

Risk	Mitigation
Relying on one supplier	Open eLMS has been trading since 2003. The platform connects outwards through its API, SSO, HubSpot and Stripe, so you can keep specialist tools where they genuinely add value.
Data migration	Move function by function, keeping the old tool read-only until the new module is signed off.
Staff adoption	One login and one interface to learn, instead of a different one for every tool.
Security and compliance	ISO 27001, ISO 9001 and ISO 20000 certified, with role-based access control and one data boundary to audit.
AI projects that stall	AI is built into each module and grounded in your own knowledge, so there's no separate AI integration project to fund or abandon.[7]

Questions to ask any vendor, including us

- Which of our current tools will this replace, and by when?
- Can we change the software after go-live without waiting for your roadmap?
- Can AI agents work across sales, projects, knowledge and finance, or only inside one module?
- How is our data separated, secured and exported if we ever leave?
- What will our total cost be over three years, including price rises?

See it working on your own processes

Book a demo and we'll walk through Open eLMS OS with your current tool list and build this financial model with you: library.openelms.com/book/ai-operating-system

References

1. Zylo (16 January 2025), *2025 SaaS Management Index*. Average SaaS spend of $4,830 per employee, up 21.9% year on year. zylo.com/news/2025-saas-management-index
2. Zylo, *License management*. Organisations leave 46% of licences unused on average. zylo.com/solutions/manage-licenses
3. Vertice, *SaaS Inflation Index*. SaaS pricing up 11.4% year on year in January 2025, against 2.7% average G7 inflation. vertice.one/insights/saas-inflation-rate
4. MuleSoft / Salesforce (2025), *2025 Connectivity Benchmark Report*, with Vanson Bourne and Deloitte Digital, 1,050 IT leaders. Average of 897 applications; 95% report integration as a hurdle to implementing AI. salesforce.com/blog/mulesoft-connectivity-benchmark-2025
5. Murty, R. N., Dadlani, S. & Das, R. B. (August 2022), "How Much Time and Energy Do We Waste Toggling Between Applications?", *Harvard Business Review*. hbr.org/2022/08/how-much-time-and-energy-do-we-waste-toggling-between-applications
6. Gartner (2025), predictions on agentic AI: 33% of enterprise software applications to include agentic AI by 2028, and at least 15% of day-to-day work decisions made autonomously. gartner.com/en/articles/ai-agents-pc1
7. Gartner (25 June 2025), press release: "Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027". gartner.com/en/newsroom
8. IBM & Ponemon Institute (July 2025), *Cost of a Data Breach Report 2025*. Global average breach cost $4.44 million; high shadow AI use adds around $670,000. ibm.com/reports/data-breach
9. Open eLMS, "We Replaced $100,000 Worth of Software With One AI Platform" (video). library.openelms.com/website-video/os-tour-1
10. Training Magazine (2024), *2024 Training Industry Report*. US training expenditure of $98 billion in 2024. trainingmag.com/2024-training-industry-report
11. Dell'Acqua, F. et al. (2023), "Navigating the Jagged Technological Frontier", Harvard Business School working paper with Boston Consulting Group: 758 consultants; on tasks within AI's frontier, 12.2% more tasks completed, 25.1% faster and over 40% higher quality. aiinstitute.hbs.edu/navigating-the-jagged-technological-frontier
12. Loopio (March 2026), *RFP Response Trends & Benchmarks Report*, seventh edition, 1,500+ teams: 33 hours per RFP response, nine contributors, 55% of RFPs received answered. loopio.com/blog/rfp-statistics-win-rates
13. Orbit Media (2026), *Annual Blogger Survey*, 1,042 content marketers: 3 hours 20 minutes on average to write a blog post. orbitmedia.com/blog/blogging-statistics
14. McKinsey Global Institute (July 2012), *The social economy: Unlocking value and productivity through social technologies*: searching for and gathering information takes about 19% of interaction workers' time. mckinsey.com/industries/technology-media-and-telecommunications/our-insights/the-social-economy
15. Brynjolfsson, E., Li, D. & Raymond, L. (2023), "Generative AI at Work", NBER Working Paper 31161: 5,179 customer support agents; 14% more issues resolved per hour, 34% for novice workers. nber.org/papers/w31161

Figures are in US dollars as published by their sources. Open eLMS OS claims (software and AI cost savings, customisation after implementation, ISO certification) are Open eLMS's own. Products and prices change; this report reflects published sources as at October 2026.