



FORLAIS

SEE INTELLIGENCE UNFOLD

Rubicon Raise

AVESI

SUMMER 2026

FORLAIS HAS LAUNCHED **THE RUBICON RAISE.**

THE RUBICON RAISE IS FORLAIS' **16 MILLION GBP SERIES SEED** ROUND
THE SEED LEVERAGES FORLAIS' PATENT PENDING PROPRIETARY TECHNOLOGY OF THE CNA ARCHITECTURE.

EVAESI WILL COMMERCIALIZE FORLAIS' RESEARCH OF THE CNA
ARCHITECTURE ALLOWING 40-60% LOWER OPERATING COSTS,
MORE ACCURATE AND RELIABLE AI THROUGH CONSUMER
PRODUCTS AND ENTERPRISE APIS.

WHERE THE RAISE GOES.

£16M PROVIDES TWO YEARS OF OPERATIONAL RUNWAY, SUPPORTING THE GLOBAL LAUNCH OF EVAESI AND CONTINUED DEVELOPMENT OF TALOS.

01 / OPERATIONS

INTERNAL ASSETS

OUR OBJECTIVE IS TO BUILD LONG TERM ENTERPRISE VALUE THROUGH OWNERSHIP OF KEY OPERATIONAL ASSETS RATHER THAN RECURRING SUBSCRIPTION COSTS WHEREVER POSSIBLE.

INFRASTRUCTURE AND OPERATIONAL ASSETS WILL BE ACQUIRED RATHER THAN RENTED, REDUCING LONG TERM OPERATING COSTS AND INCREASING THE COMPANY'S ASSET BASE.

OFFICE SPACE WILL BE ACQUIRED UNDER THIS INVESTMENT IN MULTIPLE STRATEGIC LOCATIONS, THE PLAN COVERS THE MAIN OFFICE IN THE UK AND SMALLER OFFICE BASE WHERE THE INFRASTRUCTURE IS NEEDED.

02 / INFRASTRUCTURE

SELF-SUFFICIENT

OUR ARCHITECTURE ALLOWS US TO SCALE GLOBALLY WITHOUT THE LEVEL OF INFRASTRUCTURE TYPICALLY REQUIRED BY MODERN AI COMPANIES.

FORLAIS PROVIDES THE RESEARCH, EVAESI WILL PROVIDE THE PLATFORM AND DEPLOYMENT. COMPARED TO CURRENT FRONTIER AI PROVIDERS.. EVAESI ONLY NEEDS A FRACTION OF SPACE AND INFRASTRUCTURE TO COMPETE.

SO WE WANT TO TAKE ADVANTAGE OF THIS BY NOT RELYING ON CLOUD COMPANIES AND DATA FARMS AND KEEPING 90% OR MORE INFRASTRUCTURE TO DEPLOYMENT IN HOUSE.

THE INVESTMENT COVERS THE INFRASTRUCTURE FOR LAUNCH ACROSS EUROPE, USA, UK AND CANADA

03 / TEAM

CURATING A TEAM

OUR TECHNOLOGY REQUIRES A WORLD-CLASS TEAM INCLUDING INHOUSE LEGAL, FINANCE, MARKETING AND COMMUNICATIONS TO COMMERCIALIZE AND SCALE. FORLAIS IS READY TO COMPETE WITH THE INDUSTRY LEADERS, SO WE NEED A TEAM THAT CAN SUPPORT AND MAINTAIN THAT.

TWO YEARS OF SALARIES WILL BE ALLOCATED TO A DEDICATED SECURITY ACCOUNT, ENSURING THE TEAM CAN BE RETAINED AND THE COMPANY HAS THE FINANCIAL FLEXIBILITY REQUIRED TO SUCCESSFULLY DELIVER THE GLOBAL LAUNCH OF EVAESI.

BOTH MANAGING DIRECTORS INTEND TO BE COMPENSATED FROM THE INTEREST GENERATED BY THE SECURITY ACCOUNT. PRESERVING AS MUCH CAPITAL AS POSSIBLE TO ACCELERATE THE GROWTH OF FORLAIS UNTIL EVAESI REACHES PROFITABILITY.

FORLAIS HAS MADE MULTIPLE FRONTIER AI
DISCOVERIES.

**DELIVERING 40-60% LOWER OPERATING
COST'S, WITH FASTER, MORE SECURE AND
MORE CAPABLE AI.**

ALL WHILE RETAINING THE QUALITY EXPECTED FROM CURRENT INDUSTRY LEADERS.

WE'RE ABOUT TO DISRUPT A MARKET BUILT ON
DATA CENTRES AND GPU INFRASTRUCTURE.

THE WORLD DOES NOT HAVE
ENOUGH INFRASTRUCTURE TO
MEET DEMAND FOR GPU COMPUTE
AND DATA STORAGE.

EVERY MAJOR AI SYSTEM COMPETES FOR THE SAME
LIMITED GPU CAPACITY AND DATA CENTRE SPACE. FORLAIS
TAKES A DIFFERENT PATH.

CNA FLOATLESS ARCHITECTURE CHANGES THE COST AND ACCESSIBILITY OF AI.

FORLAIS HOLDS THE PATENT-PENDING RIGHTS

40–60% **LOWER AI
OPERATING COST**

At full frontier-scale deployment, Forlais modelling shows EvaEsi delivering performance, speed, quality and knowledge comparable to current industry leaders at 40–60% lower operating cost.

RISING INFRASTRUCTURE COSTS, GPU DEPENDENCE, AND CONCERNS AROUND ACCURACY AND RELIABILITY CONTINUE TO SLOW AI ADOPTION ACROSS CONSUMER AND ENTERPRISE MARKETS.

EVAESI IS BUILT FOR BOTH . IT COMBINES 40-60% LOWER OPERATING COSTS WITH THE RELIABILITY AND ACCURACY AND AUDIBILITY REQUIRED BY HEALTHCARE, FINANCE, AUTOMOTIVE AND OTHER KEY INDUSTRIES.

PHASE 2 : TALOS

TALOS IS A PROCESSOR THAT EXTENDS EVAESI AND ITS ARCHITECTURE BEYOND THE CLOUD, ENABLING FRONTIER EVAESI AI MODELS TO RUN OFFLINE IN MEDICAL EQUIPMENT, VEHICLES, SMARTPHONES AND ENTERPRISE DEVICES WITHOUT TRADITIONAL GPU INFRASTRUCTURE.

THE INDUSTRY SHARES ONE COST BASE. WE DON'T.

GPU-BASED / THE INDUSTRY

\$600BN+

Hyperscaler capital expenditure planned for 2026, with roughly 75% tied to AI infrastructure. Capability is priced by compute, and that cost base is shared by every GPU-bound platform.

No directly comparable CPU-based alternative identified.

Sources: compiled analyst estimates of 2026 hyperscaler capital expenditure; Forlais internal modelling.

THIS OPENS AN UNTOUCHED MARKET TO EVAESI
ONE THAT COULD NOT BE TOUCHED WITHOUT THE RESEARCH FOUND BY FORLAIS.

40–60% CPU-BASED / EVAESI

CNA FLOATLESS ARCHITECTURE RUNS ON CPUS. FORLAIS MODELLING AT FRONTIER SCALE SHOWS 40–60% LOWER OPERATING COST, WITH QUALITY COMPARABLE TO CURRENT INDUSTRY LEADERS.

NO GPU ESTATE

Nothing to buy, rent or queue for. The largest cost line in GPU-bound AI does not exist here.

RUNS ON EXISTING HARDWARE

Standard CPUs in servers, data centres and devices that buyers already own and operate.

IMMUNE TO GPU SCARCITY

Operating cost does not track GPU supply, pricing cycles or allocation queues.

PRIVATE AND LOCAL BY DESIGN

On-premise and offline deployment become economically viable, not a premium exception.



INVESTMENT OPPORTUNITY

THE TECHNOLOGY IS BUILT. THE NEXT STEP IS GLOBAL RELEASE.

The Rubicon Raise is a £16 million seed round to take EvaEsi through final product refinement, global launch and commercial scale. It will also begin Talos prototyping, creating a route for EvaEsi models to run offline inside devices with a further reduction in operating cost.

evaesi



EVAESI IS NOT JUST A MODEL. IT IS A WHOLE PLATFORM.

EvaEsi will launch with three models. A further three are planned shortly after release.

BUILT

CNA + MODEL SYSTEMS

Proprietary CNA technology.

UNDERWAY

EVAESI PLATFORM

Five connected products.

FUNDED NEXT

TEAM + LAUNCH

Commercial scale and delivery.

ELEMENTS IDE


ELEMENTS IS A DOWNLOADABLE EVAESI PLATFORM FOR PEOPLE DESIGNING AND DELIVERING LONG-RUNNING PROJECTS. ITS DYNAMIC CONTEXT MEMORY PRESERVES HOW EACH PROJECT STARTED, EVEN AS THE WORK EVOLVES.



ENTERPRISE API


EVAESI BRINGS GREATER CAPABILITY AND RELIABILITY TO ENTERPRISE WORKFLOWS. ITS ARCHITECTURE IS BUILT FOR INDUSTRIES THAT REQUIRE TRUST, CONSISTENCY AND DEPLOYMENT CONTROL.

SMARTPHONE APP



WEB APPLICATION

THE EVAESI WEB PLATFORM INCLUDES A VERSATILE CODING STUDIO FOR DEVELOPERS AND EVERYDAY USERS. IT ALSO OFFERS A VERIFICATION INDEX FOR RESEARCHERS, ACADEMICS AND ANYONE WHO NEEDS TO VERIFY INFORMATION.



CONSUMER PLANS

The EvaEsi web platform is designed to be available to everyone.

ENTERPRISE API

METERED ACCESS VIA
API.EVAESI.COM, PRICED
PER 1M CHARACTERS.
MARGINS SHOWN ARE
FORLAIS INTERNAL COST
MODELLING ON OUTPUT,
PRE-LAUNCH AND SUBJECT
TO CHANGE.

INSTANT

£0.09 input / 1M characters

£0.30

per 1M output characters

e.g. 70.0% gross margin

REASONING

£0.36 input / 1M characters

£1.20

per 1M output characters

e.g. 67.5% gross margin

ELEMENTS


£0.80 input / 1M characters

£2.70

per 1M output characters

e.g. 67.0% gross margin

 LITE Perfect for getting started with EvaEsi.	 STANDARD More power for everyday creation.
£7 / month	£25 / month
• UP TO 2 MILLION RUNGS Ideal for light use and exploration.	• UP TO 5 MILLION RUNGS Great for regular users.
 e.g. 85.0% GROSS MARGIN	 e.g. 77.2% GROSS MARGIN

 DEVELOPER Built for builders and integrators.	 PREMIUM Maximum access for teams and power users.
£56 / month	£95 / month
• UP TO 10 MILLION RUNGS Advanced access for scale and integration.	• UP TO 20 MILLION RUNGS Full access. Maximum output.
 e.g. 81.5% GROSS MARGIN	 e.g. 79.4% GROSS MARGIN

EVAESI IS FORECAST TO REACH PROFITABILITY WITHIN 18 MONTHS

BASED ON CONSERVATIVE CONSUMER-ONLY ASSUMPTIONS



WE PROJECT OVER 10 MILLION IN CUSTOMER
REVENUE BETWEEN MONTHS 12 AND 18.

A SIMPLE CHOICE:

THE ECONOMIC BENEFITS ALONE WOULD
BE HARD TO IGNORE.

SIMPLY SWITCH APIs OR APPS TO EVAESI AND GET MORE
USAGE FOR LESS, WITH A MORE RELIABLE, ACCURATE,
PRIVATE AND PRACTICAL AI/API PLATFORM.

18-MONTH COMMERCIAL SCALE

Sources: OpenAI, How People Use ChatGPT (2025); Introducing GPT-5.5 (2026).


BENCHMARK / Same workload at GPT-5.5 API list pricing (\$5/1M input; \$30/1M output). 107 messages/user/month derived from OpenAI's 2.5bn daily messages and 700m weekly users; 6k input + 1.25k output tokens/message; £0.75/US\$.

MONTH 18

£19.1M CUSTOMER REVENUE

£5.7m EvaEsi total operating cost

£9.9m GPT-5.5 API-equivalent cost



ENTERPRISE BUYS THE ADVANTAGE. THE API DISTRIBUTES IT.

ACCURATE OUTPUT
LOWER OPERATING COST
CPU-BASED OPERATION
PRIVATE + LOCAL OPTIONS

Bit-accurate return observed in internal tests.
40–60% lower AI operating cost via CNA in Forlais modelling.
Deploy on CPUs without GPU infrastructure.
Private and local deployment options.

APIs ALLOW THE INDUSTRY TO USE OUR MODELS BEYOND THE
TYPICAL USER. THEY CAN BE BUILT INTO SaaS SYSTEMS, OR
DEPLOYED THROUGH SECURE PROTOCOLS FOR INDUSTRIES
SUCH AS HEALTHCARE OR BANKING, WHERE PRIVACY AND
ENCRYPTION ARE ESSENTIAL.

HEALTHCARE
FINANCE
AUTOMOTIVE
RESEARCH

ENTERPRISE IS THE SECOND ENGINE.

Consumer revenue establishes the base. Enterprise converts the platform's cost and accuracy advantages into long-term contracts.

01 / WHO IT SERVES DEFENSIBLE SECTORS

Sectors where output must be defensible: healthcare, banking, legal, sovereign and industrial environments.

02 / HOW IT DEPLOYS API + ENTERPRISE DEPLOYMENT

A metered Enterprise API for integration, with controlled deployment into private and local environments. CPU-based operation removes the GPU estate from the buyer's infrastructure.

THE OPERATING ADVANTAGE

EvaEsi's 40–60% lower AI operating costs allow us to offer better rates and packages, with reviewable outputs and confidence reporting built into the platform.

FORLAIS RESEARCH IMPLEMENTED IN EVAESI EXTENDS FAR BEYOND CNA. THE WORK DOES NOT STOP HERE.

EVAESI IS ROOTED IN FORLAIS RESEARCH.

AS THE RESEARCH ESTATE GROWS, SO

DOES THE COMMERCIAL PLATFORM.

EVAESI WILL ALSO LAUNCH WITH ORIGINAL PLATFORM FEATURES,

INCLUDING A PROPRIETARY REAL-TIME
CONNECTION LAYER AND A MORE SECURE
TRANSFER PROTOCOL. INTERNAL TESTING
INDICATES UP TO 3× FASTER PERFORMANCE.

DYNAMIC CONTEXT MEMORY

AT EVERY STAGE OF A PROJECT OR CONVERSATION, EVAESI
RETAINS THE ORIGINAL POINT. ITS MODELS ORGANISE
SHARED MEMORY ACROSS CONNECTED PLATFORMS AND
WEIGHT RECENT CONTEXT WITHIN EACH OUTPUT.

ACCURATE, RELIABLE AI

WHAT DOES ACCURACY MEAN IN PRACTICE?

EVAESI CAN BE REVIEWED AND AUDITED TRANSPARENTLY
EVERY ANSWER INCLUDES A CONFIDENCE MEASURE. WHEN
CERTAINTY FALLS, THE SYSTEM REPORTS IT HONESTLY.

BUILT FOR ENTERPRISE DEPLOYMENT


FOR ENTERPRISE USERS, INFORMATION CAN MOVE
THROUGH THE SYSTEM WITHOUT LOSSY COMPRESSION OR
QUANTISATION ALTERING THE RESULT.



FORLAIS IS AN ADVANCED AI INSTITUTE OVER THE LAST FEW YEARS WE HAVE DEVELOPED A WIDE RANGE OF DISCOVERIES THAT WE ARE ONLY NOW RELEASING PUBLICLY.

FROM SELF-DRIVEN SYSTEMS TO DIS, OUR RESEARCH PROJECTS ARE CONTINUALLY ADVANCING. EACH PROJECT HOLDS COMMERCIAL POTENTIAL IN ITS OWN RIGHT.

OUR LATEST RESEARCH INDEX HAS BEEN VALUED AT GBP 490 MILLION. BASED ON THE RESEARCH IP AND ANALYSIS OF ITS IMPACT ON CURRENT MARKETS.





EvaEsi Is A Self-Sufficient Platform That Gives Users Access To A New Type Of AI. Its Design Supports Self-Sustaining Economics, With Platform-Wide Operating Savings Of 40-60%


EVAESI BOASTS 3 FLOATLESS MODELS
DESIGNED WITH THE CNA ARCHITECTURE
AND NO GPU DEPENDANCY

THE EVAESI PLATFORM

a faster, smarter platform that will allow users across the globe to access ai at lower cost, while allowing forlais to maintain better margins than current industry leaders.

EvaEsi models run without GPUs. Internal testing indicates this can lower infrastructure cost while retaining comparable performance and output quality.

The product is built and internally verified. We are refining our proprietary data and completing final platform iterations.





**TARGETED FOR DELIVERY IN 2028 OR
EARLIER, DEPENDENT ON INVESTMENT
SCALE AND EXECUTION PACE.**

THE TALOS AI PROCESSOR

TALOS IS A MODIFICATION CPU PROCESSOR ARCHITECTURE THAT IS BUILT UPON OUR CNA ARCHITECTURE

FORLAIS' RESEARCH HAS BROUGHT THE CORE CNA ARCHITECTURE AND COUNTLESS OTHER RESEARCH FINDINGS. THIS WILL ALLOW US TO DEVELOP AND PRODUCE TALOS.

TALOS WILL BE USED ACROSS MULTIPLE INDUSTRIES, ENABLING AI TO RUN INDEPENDENTLY AND OFFLINE WITHIN MACHINERY, MEDICAL EQUIPMENT, VEHICLES, CONSUMER DEVICES AND CONFIDENTIAL ENVIRONMENTS.




Talos will be a purpose-built processor for running large-scale EvaEsi models entirely offline in consumer devices, removing reliance on traditional GPU acceleration while preserving frontier-level performance and model quality.



The core system already exists. It is a modified CPU architecture. Talos extends it into new forms of AI deployment and use.



Forlais already has the designs and underlying technology. Talos provides the most efficient route to deploying it in dedicated hardware.



THE INDUSTRY EXPLORED ONE PATH. WE EXPLORED THE REST.

FOR AI TO KEEP ADVANCING, THE INDUSTRY NEEDS A GENUINE
ALTERNATIVE TO MODEL ARCHITECTURE AND THE HARDWARE
THAT POWERS IT.

FORLAIS PROVIDES THAT ALTERNATIVE.

EXPAND YOUR HORIZONS WITH **FORLAIS** AND SEE INTELLIGENCE UNFOLD

OLIVER COOPER

CHIEF RESEARCH OFFICER

+44 (0) 800 229 4188

+44 (0) 7833 628 631

OCOOPER@FORLAIS.COM

FRASER MACDONALD

CHIEF OPERATIONS OFFICER

+44 (0) 800 229 4188

+44 (0) 7950 746488

FRASER@FORLAIS.COM