AutogenAI

autogenai.com

@autogenai

Winners use AutogenAI

Write *more* winning bids.

Harness the power of augmented intelligence to create high-quality bids, proposals and tenders. AutogenAI enables you to create exceptional bids in minutes, not days.

Trusted by 200+ enterprise businesses

serco mitie ooh! EQUANS

Specsavers First arriva Mastek

Discover the *future* of bid and proposal writing with AutogenAI

AutogenAI is the cutting-edge generative AI solution revolutionising the bid and proposals industry globally.

Leveraging the power of Generative Artificial Intelligence, AutogenAI streamlines tasks, improves accuracy, and expedites the entire bidding process. This empowers our clients with a proven competitive advantage.

Our technology is designed to help your business overcome the complexities of bid management, transforming challenges into opportunities for growth and success.

Driving growth and productivity with AI innovation.

AutogenAI leverages advanced AI technology to streamline and automate the complex and time-consuming bid process, ensuring increased efficiency and accuracy.

AutogenAI is the most effective, powerful and secure solution available. We enable large enterprises to submit high-quality, winning bids while minimising resource requirements and time investment.

As leaders in AI innovation, we constantly evolve our solutions to stay ahead of the changing business landscape, ensuring our customers stay competitive, maximise productivity, and grow.

Build vs Buy

Organisations considering the investment in generative AI solutions may face the question of whether to build or buy. Below is a summary of factors to consider when making this decision. From significant upfront investment and long-term maintenance to cost-effectiveness and scalability, these considerations are crucial for organisations looking to make the best choice for their business.

Build

Significant upfront investment: Building software in-house requires significant upfront investment and procurement of expertise in technologies and skills over and above a company's core capabilities.

Time and risks associated with the build, and implementation: Significant time is required to design, build, and test the solution which could lead to a distraction from core business services. Implementation of an in-house system is an exercise that will require careful planning and execution.

Technology is constantly evolving, becoming outdated almost immediately: Updates are crucial to maintain competitiveness, but they can be difficult to prioritise amidst other competing tasks.

Buy *

Cost effective and secure: Buying software is more costeffective than building it in-house, as it eliminates the need for hiring developers, training them, and investing in infrastructure and maintenance. Monthly licence is fixed and all-inclusive.

Available today: Purchasing software allows companies to quickly implement and start using the software, saving time compared to developing it from scratch, which could take years to develop to a similar level. Implementation risks are significantly reduced through the support and expertise of AutogenAI.

Customisation and enhancements: Specific requirements can be built into AutogenAI to satisfy bespoke client needs.

*all is included in an annual subscription

Technical FAQ

Do you use open source LLMs?
No. Where there is a requirement we can provide a bespoke solution using an open source LLM.

Do they use our data to train their models?
No. Your privacy and data are paramount to us. We never train models on your data, and neither do any LLM providers that we use.

This is assured through contractually binding Zero Data Retention agreements which we have reached with each provider and which prevents them from retaining or learning from your content in any way.

What specific models and providers are you using?
We normally agree the choice and location of LLMs and LLM providers as part of onboarding and will always notify and receive approval for any changes in choice of provider during contract delivery.

We may use the following providers:

OpenAI (GPT-3.5, GPT-4)
OpenAI servers are located in the USA. They are SOC 2 compliant, undergo annual penetration testing, and comply with GDPR and CCPA. OpenAI do not retain, train on or learn from our data.

Azure (GPT-3.5, GPT-4 by OpenAI)
Azure OpenAI servers are available in many regions including the USA, EU, UK, and Australia. Azure OpenAI services are FedRAMP authorised. Microsoft and Microsoft Azure possess numerous security accreditations including SOC 2 and ISO 27001. Azure OpenAI does not retain, train on or learn from our data.

AWS (Command by Cohere, Claude 2 and Claude 3 by Anthropic, Large and Mixtral by Mistral)
AWS offers Bedrock and SageMaker LLM services. These services host models created by numerous LLM developers. Both services are governed by standard terms, and we have confirmed that neither service retains, trains on or learns from our data.

Cohere (Command)
Cohere servers are located in the USA. Cohere are SOC 2 compliant and undergo annual penetration testing. Cohere do not retain, train on or learn from our data.

Google (Gemini)
Google Gemini servers are located in the USA. Google Cloud Platform possesses numerous security accreditations including SOC 2 and ISO 27001. Google Gemini does not train on or learn from our data.

Anthropic (Claude 2, Claude 3)
Anthropic servers are located in the USA. Anthropic are SOC 2 compliant, undergo annual penetration testing, and comply with HIPAA. Anthropic does not train on or learn from our data.

Can we plug our models into your engine?
Not by default. However, this can be discussed as it may be possible via custom deployment.

Where is my data stored?
Client data is stored in secure AWS data centres in the client's preferred region. For example, within the UK, USA, Australia, Canada etc. AWS services are provided under a standard agreement.

AWS possess numerous security accreditations including ISO27001, SOC 2, and (for most services) FedRAMP authorisation through JAB. We can host in any AWS region by agreement, including GovCloud.

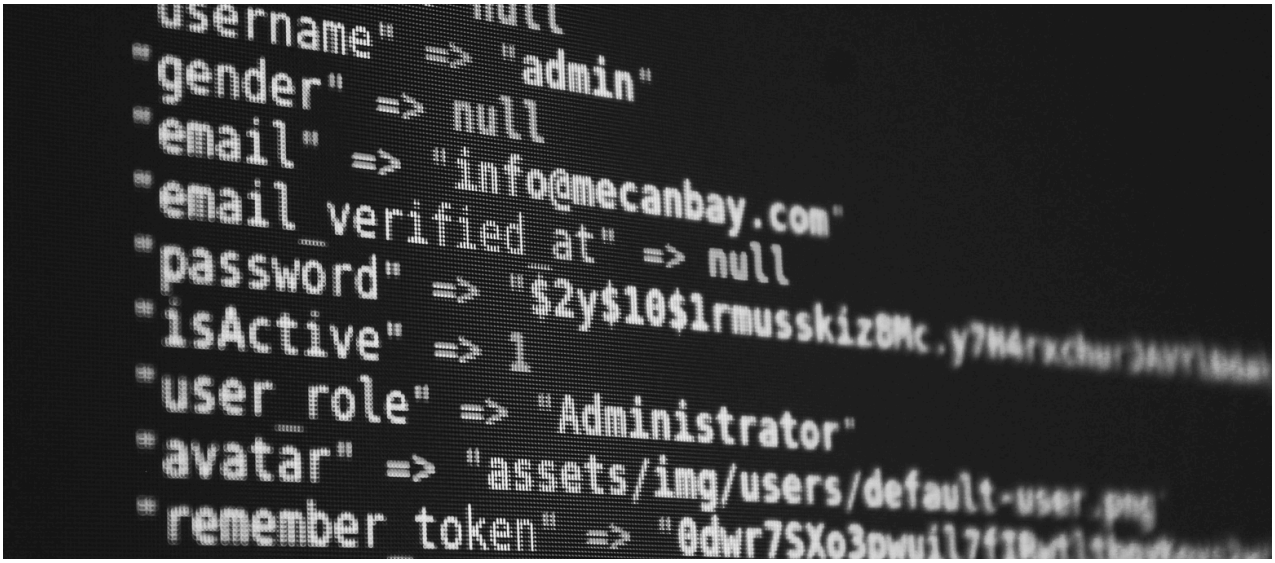

How is my data protected?
We store your data on hardened infrastructure in secure data centres in the geography of your choice. Your data is encrypted at rest and in transit, and it cannot be accessed by anyone without proper authorisation. The only places outside of these data centres that will ever see your data are logged-in, authorised users.

These are LLM providers with whom we have Zero Data Retention agreements. This means that your data is not retained, inspected, or trained on. When using our application's internet search facility, we connect you with an internet search provider that matches your query to webpages with relevant content.

Why should we trust you?
We meet all the major international information security standards including ISO 27001, ISO27017, ISO27018, CSA STAR, and SOC 2, with regular audits by Schellman and BSI Group. We follow the NIST CSF and NCSC CSG standards for cloud security.

We adhere to GDPR and CCPA requirements and undergo regular independent data privacy audits. We commission third-party penetration testing regularly, with additional assessments after major architecture changes. Our online services are under constant active security monitoring, supported by 24/7/365 in-person response to any alerts.

Our dedication to security has earned us the trust of some of the most security-conscious organisations in the world, making us an ideal partner in meeting your security needs.

AutogenAI Improves Efficiency By 85%

With AutogenAI, organisations are unlocking unprecedented levels of efficiency. By integrating AutogenAI into their workflow, Serco was able to streamline its knowledge management processes, reducing the time and effort required to access and organise crucial information 85%.

This saving has allowed Serco to allocate more resources towards other essential areas of their business, resulting in enhanced productivity and ultimately improving its overall performance.

The table below shows the results from an independent trial of AutogenAI in a real-life proposal writing experiment. Writers equipped with AutogenAI were able to write over 4 times faster than writers on their own.

	Quality words per hour	Speed improvement from base
Human Alone	500	0%
Human Plus Chat GPT*	685	+ 37%
Human Plus AutogenAI**	2,209	+ 342%

Customers are already experiencing significant improvements in win rates with AutogenAI, with one customer reported by The Times seeing a 240% increase. On average, win rates increase by 10-20% with AutogenAI's software.

*https://economics.mit.edu/sites/default/files/inline-files/Noy_Zhang_1.pdf
**Independent writing test set to AutogenAI by US-headquartered global technology solutions company

Proven to Increase Win Rates by 241%

Transform The Way You Work

There are no prizes for second place in bidding – it's winner takes all. Find out what AutogenAI can do for you.

85%
Increase in **Productivity**

70%
Increase in **Drafting Speed**

241%
Increase in **Success Rate**

100%
Increase in **Bids Submitted**

Find out *more*

Organisations that adopt AutogenAI experience are unlocking unprecedented levels of efficiency. Serco deployed AutogenAI for knowledge management and collaboration, supporting Serco's work with governments in over twenty countries.

"The first of our 'Impact Pilots' aimed at enabling innovation in public services through our partnership with AutogenAI has already seen us use its functionality more than 6,000 times, generating significant knowledge content.

Mark Irwin Group | Chief Executive at SERCO

Tech Behind the Product

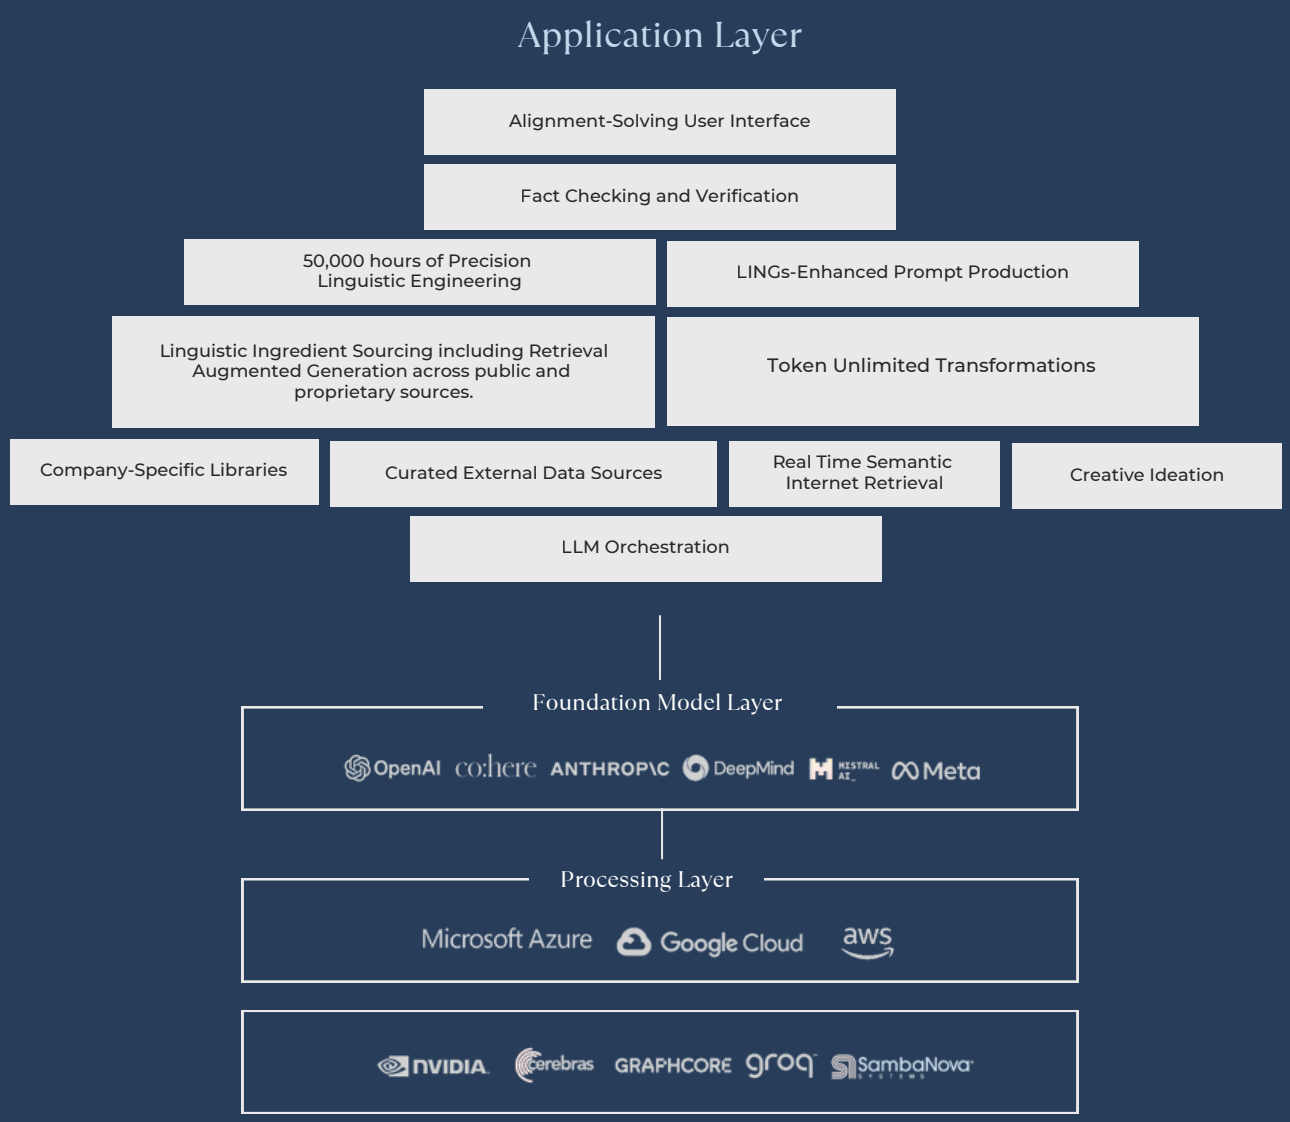

AutogenAI builds Language Engines - these are an application layer technology. They incorporate large language models but do so within a much wider linguistic engineering framework that grounds them with facts.

AutogenAI's language engines have four key components:
LLMs (Language Models), linguistic ingredients (LINGs), LINGs-enhanced independent variable production, and UX/UI design.

More than AI

How AutogenAI Makes the Difference

Faster Proposal Creation: The provider's proposal team now generates high-quality drafts in a fraction of the time, reducing manual effort and enabling faster submissions.

Validated, Cited Sources: Unlike AI models that generate unverifiable content, AutogenAI ensures accuracy by pulling from trusted internal databases, the internet, and proprietary knowledge bases — ensuring fact-based, cited responses.

Major Efficiency Gains: Rolling out AutogenAI delivered measurable efficiency gains of
80%, with tasks that used to take hours now completed in minutes.

Better Quality Submissions: By focusing less on writing from scratch and more on refining winning narratives, their bid team has significantly improved their funding.

Credibility and Capability: The AutogenAI Advantage

This unexpected and unscripted public endorsement highlights an important truth:AutogenAI clients don't just use the platform — they advocate for it because they trust it.

The combination of AI built for proposals and expertise from proposal professionals ensures that organiations don't just submit more proposals — they submit better ones. It isn't just about buying a tool, it's about partnering with the best in the business.

"

Augmented intelligence requires more than just great software - people need to know how to use that software to achieve growth.

Sean Williams | CEO, AutogenAI

The Results (So Far)

Before AutogenAI, our client had spent countless hours assembling proposals manually, often scrambling to meet deadlines. Now, proposal preparation is streamlined, cutting writing time by over 80% and allowing them to focus on crafting winning proposals — not just finishing them on time. And their competitors know all about it, because they shared their success publicly.

Conclusion

Great products don't just solve problems — they inspire loyalty. That's exactly what happened when this community healthcare provider, unprompted, became an advocate for AutogenAI. Their experience spoke volumes — not only about the technology but also about the trust, expertise, and results that come with it.

AutogenAI isn't just another AI tool — it's the most advanced proposal-writing platform available. By combining cutting-edge AI with deep proposal expertise, it ensures teams win more work, faster, and with greater accuracy. No other AI solution delivers this level of precision, credibility, and impact.

Write High-Quality, Winning Bids, Faster than Ever.

Transform your tender writing and management process.

Quality: Determine whether an RFP is right for your company.
Save time and reduce risk in the qualification process with AutogenAI's AI-driven evaluation and questionnaire response tool.

Manage: Work on multiple proposals with ease.
Stay aligned and stay focused with a single, unified platform that brings everything together—deadlines, stakeholders, projects, compliance needs, resources, and context.

Write: Stop starting from scratch.
Turn a blank page into a review-ready draft in minutes.

Refine: Elevate proposals in one-click.
Add flavour and power to your proposals with in one-click. Improve the quality of your proposals at scale.

Review: Leave nothing to chance.
Ensure your bid covers every requirement and stands out, maximising your chances of success.

Submit: Finish with a flourish.
Submit with confidence, knowing you have the most optimized and competitive proposal.

Secure

Built by Experts

Global Support

Find out *more*

AutogenAI

Client Quotes

serco

"The first of our 'Impact Pilots' aimed at enabling innovation in public services through our partnership with AutogenAI has already seen us use its functionality more than 6,000 times, generating significant knowledge content.

Mark Irwin Group | Chief Executive at SERCO

"Since I began using AutogenAI in early August, it has significantly improved the efficiency of my grant writing. It's a huge time-saver and significantly reduces the stress associated with grant application deadlines.

Chrysothemis Brown | Physician Scientist

Twin Employment & training

"AutogenAI, not only elevated the quality of our proposals, but it also increased the speed of responses, especially when faced with challenging questions.

Mark Sargeant | Head of Buisness Development

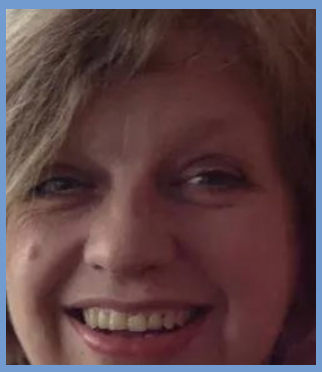

hgs

"Having implemented AutogenAI, it is clear that this is a game changer; significantly increasing the productivity of proposal teams.

Debbie Brockbank | Growth Director

AutogenAI

Case Study

Serco Increases Efficiency by 85% with AutogenAI

The Challenge

Serco, a leading British multinational company, offers a diverse range of services in defence, justice & immigration, transport, health, and citizen services. With a workforce of 50,000 employees and global revenues of $6 billion, Serco faces fierce competition in securing contracts.

Identifying innovation as a key driver for growth, Serco decided to pilot AutogenAI, powered by AWS, with the aim of enhancing their bidding process and knowledge management, resulting in significant improvements and increased efficiency.

Results

By incorporating AutogenAI into their bidding process, Serco achieved an impressive 85% efficiency gain in its bid and proposal writing and knowledge management processes. This enhanced efficiency through advanced AI capabilities allowed Serco to focus on other critical aspects of its operations, leading to improved overall performance.

Get in *touch.*

Let us show you how it works

contactapac@autogenai.com

autogenai

Book a demo now.