datacentrix

Cutting through the complexity for smart enterprise AI adoption

When it comes to selecting artificial intelligence (AI) software, enterprise leaders are finding themselves overwhelmed by choice, the complexity of pricing models, and a lack of clear value pathways.

The initial excitement surrounding AI tools like ChatGPT has rapidly evolved into a multi-layered marketplace, with vendors racing to offer increasingly sophisticated models and solutions. However, as the landscape expands, so too do the challenges for buyers. And hence, enterprises are no longer asking “what can AI do?”, but rather “how do I make this work for my organisation – without losing control of costs or value?”

This is according to Phil Anderson, National Sales Manager: Digital Business Solutions at Datacentrix, who maintains that the key to AI adoption lies in cutting through the noise and zeroing in on enterprise-grade use cases.

Personal productivity tools, he argues, may improve individual efficiency. However, the challenge of tracing the benefits back to metrics that prove substantial business impact means that justifying large-scale investment is rarely straightforward.

“Enterprise use cases, integrating directly into business process, that require the use of a large language model (LLM), are where the highest return can be seen in the market. It's here that the real business impact can be found, through the standard return on investment (ROI) and time to value (TTV) calculations with which every organisation is familiar.”

> “Enterprises are no longer asking “what can AI do?”, but rather “how do I make this work for my organisation – without losing control of costs or value?”

Avoiding the token trap

A core issue, Anderson says, is the industry's current fixation on token-based pricing models.

“Tokens are the unit of measurement used by LLMs to calculate consumption… but this is where the simplicity ends. Token cost comparisons between LLMs are misleading,” he warns. “Different models define and use tokens in inconsistent ways; some count spaces, others don't, some cache full words, others cache characters. The result is that cost per token becomes an unreliable metric.”

> Enterprises must begin by defining what they need their AI to do, whether that's summarisation, classification, extraction, translation, sentiment analysis or more. From there, model capabilities, size, tuning potential and transparency of data training practices become critical selection factors.

Instead of making the mistake of cost-per-token comparisons, Anderson proposes three evaluation strategies:

1. Don't try to map every token nuance – the pace of change renders this effort quickly outdated.
2. Testing the use case across multiple models would be ideal, but this approach is often expensive and resource intensive.
3. Focus instead on model fit and applying the smallest possible model to the problem statement. This is the most practical approach, aligning use case requirements with model capabilities and using smaller models which cost less to tune and deploy.

From foundation to function

Enterprises must begin by defining what they need their AI to do, whether that's summarisation, classification, extraction, translation, sentiment analysis or more. From there, model capabilities, size, tuning potential and transparency of data training practices become critical selection factors.

Anderson also stresses the growing importance of AI governance, particularly with many LLMs operating via cloud-based APIs. Trust in the vendor's handling of data, ethical AI training standards and compliance with data sovereignty regulations are all crucial to mitigating risk.

“The LLM is really just the engine in a much bigger vehicle. To truly scale AI across an enterprise, you need the whole car: integration, orchestration, data governance and model management.”

Prime components to be considered should include the following:

Integration – in most use cases, there will be a requirement to interact with systems already in use within the enterprise, transporting data to and from models. This must be managed through an integration bus using common API standards. If there is no platform in place, this must be considered in any proposed solutions. However, those with existing integration platforms and associated standards in place should design according to those.

Orchestration – this is a significant consideration. How will all of the data pipelines, execution of various models, workflows and data sources be managed in association with the wider objective? This is at the core of the 'agentic' movement currently under wide discussion, which combines AI and different automation disciplines, such as

Phil Anderson, National Sales Manager for Digital Business Solutions at Datacentrix

robotic process automation (RPA), to co-ordinate agent activity across the enterprise in a single control plane. If an organisation is planning AI deployment at scale over time, this is an absolutely foundational component.

Data governance and management - this has been said many times by many people about AI, but it matters. If there is no proven understanding of data being provided to the models, then dark times lie ahead. Don't underestimate the architecture and method required here and make it a focus.

Model management - what is the home for all these AI exploits? Using an LLM as standalone for an enterprise is possible but it brings risk, especially in cases where tuning is required. What is the AI management platform that will be used to manage the models and data sources, manage model drift over time and make tuning a more straightforward task? A home is required for all of this innovation.

Beware the hidden costs

Despite vendors claiming transparent pricing, the real question enterprises must ask is: “What can go wrong – and how much could it cost?” From failed proofs of concept to repeated tuning cycles, the 'school fees' for AI deployment can be steep.

“A hidden cost in AI is the iteration required to tune models to your specific needs,” explains Anderson. “The more specialised the business area, the more tuning – and possible trial and error – required. A vague business requirement here is much more costly than in a traditional project with more deterministic outcomes.

To mitigate this, he encourages clients to spend more time than ever before on the definition of 'done'. Deep understanding and collaboration with your business to define how the anticpated benefits will be delivered becomes fundamental to success.

Secondly, be sure to negotiate intellectual property (IP) terms when collaborating with vendors. By co-developing re-usable IP and securing IP terms that are mutually beneficial, organisations can reduce costs while giving vendors commercial incentives to innovate with them.

Achieving cost certainty through AI as a service

One of the most effective ways to reduce risk and increase cost certainty is to embrace AI as a service (AIaaS). Rather than building in-house infrastructure or managing unpredictable consumption costs, enterprises can shift to outcome-based buying.

“With AI as a service, you're buying the result, not the technology stack,” says Anderson. “We take the risk of delivery and performance, while enabling customers to focus on business value.”

This approach is particularly beneficial in the early stages of enterprise AI adoption, as it reduces the need for capital investment, simplifies operational complexity and provides a clear path to return on investment (ROI).

Even for organisations pursuing a hybrid strategy, AIaaS offers a low risk starting point that enables learning and value delivery in parallel.

Key considerations for enterprise leaders

“AI represents a foundational shift,” concludes Anderson. “Success does not come from jumping on the latest model, but from building the right architecture, applying governance and focusing relentlessly on value.”

For enterprise leaders grappling with the complexities of AI, a reliable partner can offer support by helping to define strategy, reduce cost and accelerate transformation, responsibly and at scale. AIaaS enables early-stage AI deployments to be de-risked, enabling organisations to focus on meaningful business outcomes.

Datacentrix's AIaaS offering encapsulates all the essential components, from model access, and orchestration platforms, to data governance and managed delivery, into a cohesive, scalable solution.

Access the full white paper here:
www.datacentrix.co.za/uploads/8/3/1/1/83111140/buying_ai_for_the_enterprise_white_paper_final_062025.pdf