RESEARCH BRIEF / TECHNOLOGY BUYING

Before a Nigerian Business Buys an AI Platform

A purchasing brief should test the operating conditions behind the demonstration.

Jennifer Oriwoh

The most useful first question in an AI purchasing meeting is what the business needs to improve. A platform can produce an impressive answer while leaving that question unresolved.

Consider a hypothetical Lagos service business comparing tools for handling customer enquiries. It wants staff to spend less time finding approved information and more time resolving requests. The decision is whether a particular change can improve that work at an acceptable cost and level of risk. Buying AI is one possible answer, alongside better records, ordinary automation and a more organised current process.

The Nigerian setting gives the buyer a concrete place to investigate data handling, payment arrangements, support and connectivity. It does not justify assuming that all Nigerian businesses face the same constraints, or that an overseas product will necessarily fail. A purchasing recommendation should change when the business's circumstances change.

Start with the information the system will receive

For the enquiry use case, assemble an authorised test set that reflects the work: current policies, superseded documents, incomplete questions and requests requiring a person. Record what a correct answer should contain and when the system should decline to answer. Keep private customer records out of an initial demonstration unless their use has been properly authorised and protected.

Then test whether the tool identifies which information is current. An answer that correctly summarises an obsolete policy can still be wrong for today's customer. A source link helps a reviewer investigate; it does not by itself prove that the answer follows the source or that the source applies.

Language and context should be tested through the customers and tasks the business actually serves. If enquiries include multiple languages, informal expressions or sector-specific terms, those belong in the evaluation. Broad claims about African language support cannot replace examples relevant to the intended users.

The World Bank's *Digital Progress and Trends Report 2025* organises AI foundations around connectivity, compute, context and competency. Its discussion of locally relevant data and skills offers a useful starting lens. It is a development-level framework, not evidence that a particular platform will repay this business's investment.

Compare four credible options

Keeping the current process is a legitimate baseline. Measure the time, errors and unresolved requests it produces before claiming that an alternative improves them. The baseline may itself benefit from clearer responsibilities or a better-maintained knowledge base.

Ordinary automation is the second option. If enquiries fall into stable categories with approved responses, better search, templates or routing may address much of the problem. That option deserves a fair trial rather than being dismissed because it sounds less ambitious.

A purchased AI service may offer flexible language handling with less development work. The buyer should establish what it includes, which work remains with staff and how changes to the service will be communicated. A subscription price is only one part of that comparison.

A limited custom component offers another route when a specific integration or control is essential. It also creates maintenance, monitoring and support responsibilities. Building something small does not remove the need for someone to own it after launch.

Calculate the cost of accepted work

For each option, estimate setup, integration, subscription or usage charges, staff review, training, support and likely exit work. Where billing is in a foreign currency, model how a change in the naira cost would affect the business's budget. Use current supplier terms for the calculation rather than assuming an exchange rate or a fixed price here.

The denominator matters. Dividing expenditure by every generated answer rewards volume even when answers need correction. A more useful proposed measure is cost per accepted resolution, with a clear definition of accepted and a separate record of serious errors.

Connectivity needs the same practical treatment. Test the service on the connections and devices staff will use. Interrupt a permitted test request and establish whether work is saved, duplicated or lost. Ask what staff should do while the service is unavailable. These tests address the actual operating environment without turning a national infrastructure assumption into a product verdict.

Establish who controls the data and the decision

Before uploading operational records, ask the supplier to explain where data goes, how long it remains, who can access it and whether it is used to train models. Obtain answers that apply to the service tier and contract under consideration. A generic privacy statement may not settle the specific arrangement.

The business should have its responsible data-protection adviser review the proposed processing against applicable Nigerian requirements before using personal data. This brief does not determine the legal basis for a particular deployment. Commercial convenience is not a substitute for resolving that question.

Operational ownership also needs a name. Someone must maintain the source material, review disputed answers and decide when the tool should be suspended. Ask the supplier to demonstrate export and deletion options using permitted test data. An exit promise is more useful when the buyer understands what can actually be recovered.

Make the pilot capable of saying no

Agree the acceptance conditions before the trial: the quality improvement required, the affordable operating cost, the errors that block rollout and the fallback if the tool is withdrawn. Include ordinary staff in the test rather than relying entirely on a vendor-led demonstration.

A small pilot cannot establish every risk or seasonal demand pattern. It can reveal whether the purchasing case has survived contact with real work. If staff spend more time correcting answers than the tool saves, narrow the use case or stop. If a simpler option performs adequately at lower total cost, buying less technology is a defensible outcome.

The purchase should follow a demonstrated improvement in a defined process. The ability to produce an impressive answer is the beginning of that assessment, not its conclusion.

Source

World Bank, Digital Progress and Trends Report 2025: Strengthening AI Foundations, official report overview.