



AX Assessments

How TheAX gathers the data a diagnostic needs

Asking is easy.

Asking the same way every time is the hard part.

AX Assessments are the interviewing and data-gathering tools within TheAX, the assessment and diagnostic platform. Every assessment is generated from your expertise in your Product Book, so the questions, the evidence requested and the scoring all come from your model rather than being written from scratch. Built in minutes, run by automated bots or by consultants in interview, and comparable with every other assessment your firm has ever run.



FASTER

Built in minutes

Start from a stored template or from scratch with AI drafting the questions against your expertise. No blank page, and no bespoke questionnaire per client.



EASIER

Low effort

Automated bots gather the routine data. Consultants keep the conversations that actually need judgement, instead of the form-filling around them.



MORE CONSISTENT

Repeatable quality

Systemised scales, options and evidence requests mean two engagements ask in the same way with repeatable quality whoever happens to be running them.



COMPARABLE

Answers that add up

Because they were gathered the same way, results combine over time into pathways and across clients into benchmarks. Building your firm's IP as you go.

01 · THE PROBLEM IT SOLVES

ONE BOOK, THREE DEPTHS

LEAD MAGNET

Published, self-serve

Perception only. An instant report for the participant, and a scored reading on your scale for you.

GROUP ASSESSMENT

A named panel

Perception across multiple stakeholders, with evidence where you want it. From an early conversation to a full baseline.

FULL AXAT

Panel, artefacts, symptoms

All three lenses, corroborated and scored against your levels. The defensible baseline.

Same Product Book. Same capabilities, same levels, same evidence rules.

So the results compare

The same instrument at three depths. A lead magnet, a stakeholder questionnaire and a full triangulated AXAT baseline are not three different products — they are the same model asked at different depth, which is why a prospect who filled in a lead magnet already sits on the scale you will assess them against.

Every engagement starts from a blank page

A questionnaire written for this client, by this consultant, this time. Days of senior effort before a single answer arrives.

The answers never add up

Different questions and different scales each time, so nothing aggregates across clients, and a re-read six months on measures something slightly different.

The wrong people do the collecting

Chasing responses and tidying spreadsheets is the least valuable use of the most expensive people in the firm.


02 · CREATING AN ASSESSMENT

From your model. Not from a blank page.

The categories, the levels and the evidence are already systemised in your AX Product Book. An assessment is the instrument that collects against them, and building one is a matter of choosing depth and curating what the AI proposes.

FROM YOUR MODEL

Knowledge & Intelligence	3
Outcome Pathway	3
Execution & Delivery	3
Technology & Platform	3
Commercial Model	3
Client Engagement	3
Leadership & Operating	3
Governance & Assurance	3



QUESTION TYPES

Open ended

Single choice

Multiple choice

Linear scale

Yes / No

Rating

Qualifying questions, skippable questions, and open text the AI reads back as themes.

Nobody writes a questionnaire from nothing. Because the questions are generated against the Consulting Atomic Model that already exists in your Product Book, two consultants building an assessment for two different clients produce instruments that still ask in the same way – which is the only reason the answers can be compared later.



Templates, or bespoke

A finished assessment can be stored in the AX Product Book and reused on any client with two clicks. Where a client needs something different, build it bespoke in minutes against the same categories.



AI drafts the questions

The model already holds the capabilities and what each level looks like, so the AI can propose the questions and the answer options. Your people accept, edit or reject.



Six question types

Linear scale, single and multiple choice, rating, yes/no, and open-ended text the AI analyses rather than leaving for someone to read through.

03 · TEMPLATES AND LEAD MAGNETS

The same build produces two things a practice can publish: an instrument it reuses internally, and a lighter version it puts in front of the market.



Reusable templates

A template is stored in the Product Book and applied to any client. The instrument stops being something the firm rebuilds and becomes something repeatable it owns.



Lead magnets

Publish a light version with a pre-configured report the participant receives immediately. Useful to them, and a demonstration of how the firm thinks rather than a claim about it.




A scored reading

Every submission is measured against your own model, so an early conversation starts from a position on your scale. If that organisation becomes a client, the reading already compares with the full baseline.

04 · COLLECTING THE DATA

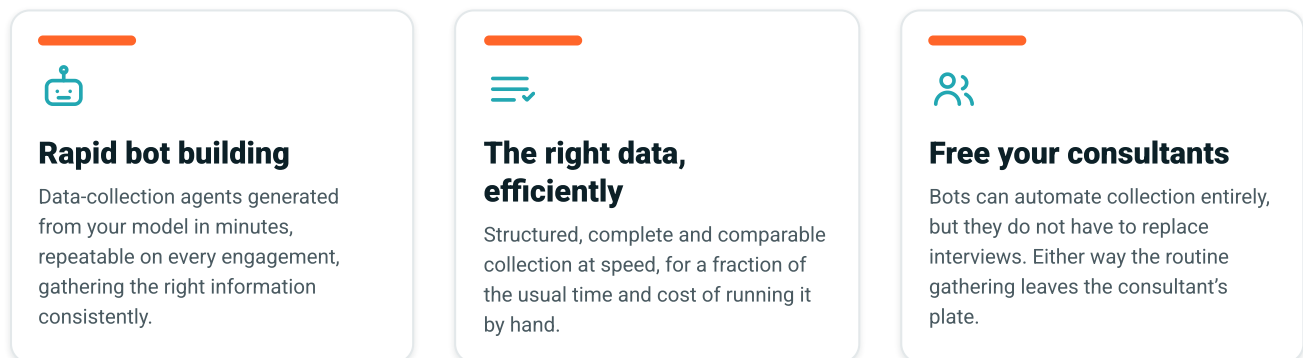
Bots where it is routine. Consultants where it counts.

Collection is where engagements lose their time. TheAX generates automated agents from the same model the assessment was built from, so the routine gathering runs itself — without taking the conversations that genuinely need a consultant in the room.




MIX THE TWO HOWEVER THE ENGAGEMENT NEEDS

Both routes produce the same structured result. A stakeholder answering a bot and a stakeholder answering a consultant are answering the same question, scored the same way, so an engagement can be fully automated, fully human, or any mix of the two without the data becoming inconsistent.



05 · WHO ANSWERS, AND HOW IT RUNS




What it still asks of you. A bot can collect without a consultant in the room, but where evidence is in scope somebody on the client side still has to produce the artefacts. That is the part people underestimate, so we scope it explicitly before anything starts.

06 · WHAT THE DATA DOES NEXT

One assessment is a snapshot. A series is a pathway.

ONE CLIENT, FOUR READINGS



CAPABILITY MOVEMENT AND BUSINESS RESULT, ON THE SAME CHART

This is what makes the work provable. Capability movement and the client's own KPI on one chart — which turns "trust us, it worked" into a number with evidence behind it.



Pathways to the outcome

Readings from successive assessments are combined with KPI readings to join into a pathway, so a client can see the distance travelled against the outcome they actually bought.



Advice that adapts

Each re-read updates the position, and the guidance the model attaches to it updates with it. Advice moves with the client rather than dating from month one.



On track, or not

Targets are set at the baseline. A pulse that falls short shows up as a number while there is still time to act on it.

07 · WHAT COMES OUT

THE REPORT

Word document

.doc

An editable version to adapt, extend or re-brand.

Web page

.html

A self-contained file to keep, send or host anywhere.

THE DATA

Excel data

.csv

Overall position and per-capability scores as a sheet.

Data for AI

.json

Structured figures an LLM can build any new document from.

EVERY REPORT CONTAINS

Position

Capability profile

Four-box placement

Themes

Root cause

Sequencing decision

Risks

Or share it live instead, through a link or an AX Portal, so the client sees the current position rather than the version you exported. The AI file carries the figures, not the calculation detail behind them.

More on TheAX

Every part of the platform has its own fact sheet, and a walkthrough on a completed assessment is the quickest way to judge the method for yourself. Both start at the website.

hello@theax.ai · Marlborough, United Kingdom

FIND EVERYTHING AT

theax.ai