

EquipBridge AI Project launch handbook

A practical guide to demonstration, customer discovery and a bounded validation programme.

EquipBridge AI helps Chinese equipment manufacturers organise bilingual after-sales knowledge and service tickets for overseas support. This handbook gives the project team a repeatable way to present the prototype, recruit pilots, measure results and make release decisions.

The next decision

Demonstrate the current workflow honestly, then secure three scoped prospective pilots. Customer outcomes, commercial demand and the business plan financial projections remain unvalidated. The first segment is one construction-equipment family in a Zhejiang exporter cluster.

Use this handbook	Page
Run a three-minute demonstration	2
Understand current capabilities and prepare evidence	3
Assign the next 90 days of work	4
Interview prospective customers	5
Apply pilot acceptance criteria	6
Measure quality, effort, adoption and conversion	7-8
Check release and submission readiness	9
Use the brand and project contact details	10

Live project: equipbridge-ai-bo.bcao0016.chatgpt.site

Basis: the project launch content supplied by the project team and its current implementation scope. This handbook is an operating guide, not an independent code, security or compliance audit. EX-20, EX-35, sample manuals and demonstration tickets are fictional.

Run the demonstration

Use `/demo` for a public sample walkthrough. Its temporary changes reset on refresh. Use `/workspace` for records saved to the signed-in account, reviewed bilingual source input, optional files, approval and closure, and JSON export. Explain which route is on screen.

Prepare before the session

- Use fictional records in a presenter-owned account. Prepare an EX-35 manual-applicability ticket, a reviewed bilingual excerpt and a clear revision and page reference. A file attachment does not imply automatic extraction.
- Prepare one supported and one unsupported question. Check the actual interface and rehearse saving, refreshing and exporting. Explain that one account owner is demonstrating the workflow; separate organisation roles are not implemented.

Time	Action and narration
0:00-0:25	Open the public site and identify the route. Say: "We help manufacturers bring technical knowledge and overseas service tickets together. This is a prototype with fictional equipment and documents."
0:25-0:55	Open the EX-35 ticket. Show the model, available serial information, service question and status. Explain what is known and what requires clarification.
0:55-1:30	Retrieve a maintained bilingual excerpt. Show its revision, page reference and optional attachment. Say: "This retrieves supplied text. It does not call an external language model or automatically read the uploaded PDF."
1:30-2:00	Try the unsupported question. Inspect the actual result and keep the technical response pending. Explain the missing evidence. Do not endorse an unrelated excerpt or imply that every unsafe request is detected.
2:00-2:35	In the owner workspace, review the supported response, record the outcome, approve or close as appropriate, then refresh and confirm the saved state. Show JSON export if time allows. Do not use the public demo to claim persistence.
2:35-3:00	Close with the request: "We seek three prospective pilot manufacturers, equipment-service mentoring and resources for a bounded validation programme. We will test handling effort, answer quality and willingness to pay."

Recovery rule

If a network action fails, state the failure. Use a clearly labelled recording or screenshot from a verified run if needed. Claim a save, approval, closure or export only after the interface confirms it. A recorded walkthrough must be introduced as recorded.

Know the current scope

Area	Available in the prototype	Next requirement
Public demo	Interactive fictional sample workflow; changes reset on refresh.	Keep sample data separate from project evidence.
Saved workspace	User-scoped tickets and knowledge; records persist for the account owner.	Shared manufacturer and distributor organisations, membership and tested role permissions.
Sources and files	Reviewed bilingual text entered into knowledge records; optional attached files.	Automatic PDF extraction, OCR and reviewed document ingestion at scale.
Answers	Retrieval of maintained bilingual excerpts and inspectable source references.	External LLM integration, generated translation and evaluated answer drafting.
Review and export	Human review, approval or closure in the owner workflow; JSON export.	Independent engineer approval responsibilities and validated restore procedures.
Communications	Service information retained in the workspace.	Outbound messaging and notification integration.
Commercial operation	Proposed pilot and subscription offers in the plan.	Billing, payment collection and subscription administration.

Prepare a source record

- Use an authorised document or fictional sample. Record document title, equipment model, revision, source page and applicability information. Record missing serial or version details explicitly.
- A domain reviewer checks the source and manually prepared bilingual excerpt, including identifiers, units, warnings and meaning. Enter the reviewed text and attach a file only if needed and authorised.
- Test a supported question and an insufficient-evidence case. Inspect the retrieved material before approving a response. Source references aid inspection; they do not establish technical correctness.

Operating boundaries

Owner-scoped persistence is not shared organisation collaboration or proof of a complete security model. A JSON export is a portable record; successful restore must be tested separately. No external LLM, automatic PDF extraction, shared organisation roles, outbound messaging, billing or machine control is claimed.

No signed customers, earned revenue, patents, safety certification or completed compliance audit is claimed. Use fictional or expressly authorised data until deployment, access and data-handling reviews are complete.

PROJECT LAUNCH HANDBOOK / 04

Deliver the next 90 days

Day 1 is the agreed start of the next work cycle. These first 90 days sit within the 16-week programme. Pilots starting on Day 71 need measurement through at least Day 98 to complete four weeks of use; aim for the final cycle review by approximately Day 112. Move dates when prerequisites slip.

Period	Deliverable and accountable owner	Decision gate
Days 1-14	Rehearsed demo, accurate feature claims, fictional records, release checklist and persistence tests. Project lead coordinates; Product lead resolves workflow issues; Security lead reviews access paths.	Publish only demonstrated capabilities. Resolve unauthorised record or attachment access before customer-data use.
Days 15-28	15-20 manufacturer and distributor interviews; anonymised workflow notes; at least five qualified firms willing to share sample evidence. Project lead recruits and interviews; Product lead observes workflow.	Require a recurring problem, approved documents, a sponsor and a realistic buying process. Record refusals.
Days 29-56	Seek three scoped pilot agreements; prepare one equipment-family corpus; review excerpts; evaluate 100+ held-out questions. AI lead owns retrieval tests; Product lead workflow; Security lead access checks; Project lead agreements.	Require deployment and access review, a named domain engineer and acceptable technical tests. If shared roles remain unfinished, use sandbox studies and leave collaboration unvalidated.
Days 57-70	Two-week baseline for each ready pilot, with eligible tickets, case mix, timings and measurement owners. Product lead runs measurement; Project lead coordinates sponsors; AI lead checks question categories.	Start comparison only with usable baseline data and agreed definitions. Stagger pilots when capacity is limited.
Days 71-90	Begin measured use, weekly quality review and logs of coverage, errors, support effort and onboarding hours. Present written offers. Product lead owns adoption; AI lead quality; Security lead access incidents; Project lead buyer feedback.	Day 90 is an interim continue, narrow or pause decision. It does not establish four weeks of use for Day 71 starts.

Responsibilities and resources

Assign one accountable owner to each role: Project lead for delivery and customer decisions; Product lead for workflow and adoption; AI lead for retrieval evaluation; Security lead for access reviews. Secure an external equipment-service reviewer; no existing advisor appointment is assumed.

The RMB 100,000 validation cash cap is illustrative and excludes student founder time. The Project lead records spending and contributed hours. The later RMB 8 million envelope is conditional; build a monthly cash forecast before making funding or runway claims.

Interview prospective customers

Start with recent behaviour and records before showing the product. Obtain permission for notes, recordings or sample documents. Use redacted examples and retain only agreed information. Speak with both the operating user and the person who can approve a purchase.

1. Map the work

Walk us through your most recent overseas support request, from the first message to the final outcome.

2. Find repeated effort

Which people and tools handled it? Where did information have to be entered or explained again?

3. Check applicability

How did you identify the model, serial number and correct document revision?

4. Understand language work

Which language steps required an engineer or another employee? Can you show an anonymised example?

5. Find the constraint

What made the response difficult: documents, technical judgement, translation, ownership, parts or something else?

6. Plan measurement

Which service records or timings already exist? What could we measure without adding excessive work?

7. Identify the reviewer

Who approves technical replies today? Which requests must always go to that person?

8. Map access

What may distributors see? Which documents must stay inside the manufacturer?

9. Find the buying process

Who sponsors a trial, approves a purchase and controls the budget? What steps and timing are required?

10. Invite objections

What would prevent the team from using this workspace, even if the demonstration worked?

11. Define value

Which result would justify paying? What evidence would you need at the end of a pilot?

12. Test a real offer

After discussing value, present the written scope and price. Ask whether they would accept it and what would need to change.

Keep an evidence record

Capture interview ID, date, role, current workflow, supporting record, pain frequency, buying authority, objection and agreed next step. Separate observations, direct quotes and researcher interpretation. A positive reaction or letter of interest is not a paid conversion.

Proposed offer: RMB 10,000 for a fixed-scope pilot, credited against the RMB 10,000 onboarding fee on conversion. Do not count it twice. Starter pricing of RMB 49,000 per year is a hypothesis to test, not proven willingness to pay.

Apply pilot acceptance gates

All thresholds are proposed. Passing a finite test set does not establish unrestricted real-world reliability. Record the evidence, reviewer, date, result and unresolved items for every gate.

Gate 1 Demonstration release

- Complete the fictional ticket workflow: inspect source evidence, review, approve or close, refresh the saved workspace and export JSON. Confirm the public demo resets after refresh.
- Test owner isolation using direct record identifiers and attached-file access. A successful interface walkthrough alone does not establish isolation.
- State the scope accurately: excerpt retrieval, file storage and an owner workflow. Leave unsupported questions pending rather than endorsing irrelevant text.

Gate 2 Before customer data pilots

- Target three scoped agreements, each with an accountable customer sponsor and an authorised equipment-service reviewer. Approve data access, deployment, corpus and document applicability.
- Evaluate at least 100 held-out questions covering both language directions, routine cases, absent evidence, version conflicts and safety-sensitive requests. Label cases before tuning and record any test-set reuse.
- Target at least 90% of returned routine answers passing the engineer rubric. Report coverage alongside it. For the current retriever, assess relevance, applicability and sufficiency; this is not a generative-model accuracy result.
- All designated safety cases in the finite test set must route to engineer review. If reliable software routing is unfinished, a supervised demonstration cannot satisfy the release criterion for independent customer use.
- If shared organisations and roles are required for the pilot but unfinished, run a sandbox study and identify the collaboration hypothesis as untested.

Gate 3 End of the validation cycle

- Each evaluated pilot has a two-week baseline and at least four weeks of use, or is reported as incomplete.
- Target at least 20% lower median active handling time for comparable eligible tickets without increased critical technical errors. A critical error pauses affected use for review.
- Target at least two of three pilots converting at the tested commercial price. Count paid conversion only after commercial terms are accepted and required payment is received; report unpaid commitments separately.
- Review onboarding and support costs. More than 40 hours of setup and training per new account triggers investigation of scope, document quality and price.

Decision record

Choose continue, revise or pause. Cite measured results, customer refusals, missing observations and delivery cost. Assign an owner and a recheck date. If a gate fails, adjust the segment, workflow, corpus or price before expanding.

Measure effort and answer quality

These definitions are a measurement plan, not a claim that analytics are automated. Use a manual log until events and timing are implemented. Assign stable ticket IDs, a consistent time zone, eligibility rules and equipment and language categories before collecting a baseline.

Metric	Definition	Reporting rule
Active handling time	Minutes reading, searching, translating, reviewing and responding for an eligible ticket. Exclude waiting.	Report median, interquartile range and ticket count by period and case type. Improvement = (baseline median - pilot median) / baseline median; undefined if baseline is zero.
First meaningful response	Elapsed time from the request to relevant information or a necessary, specific clarification.	Exclude automatic acknowledgements. Choose calendar-hour or business-hour timing before comparison.
Resolution time	Elapsed time from request to a confirmed outcome and closure.	Report separately from handling effort. Parts, travel and customer availability can dominate elapsed time.
Seven-day repeat contact	Cases with further contact about the same unresolved issue within seven days / cases with a complete seven-day follow-up.	Exclude incomplete follow-up windows from the denominator and report their count separately.
Routine-answer pass rate	Returned routine answers meeting the predeclared engineer rubric / returned routine answers reviewed.	Material corrections are failures. Check relevance, model and revision applicability, identifiers and sufficient evidence.
Answer coverage	Eligible routine questions receiving a candidate answer / all eligible routine questions.	Show with pass rate and unsupported-case counts so abstention cannot conceal low coverage.
Reviewer acceptance	Candidate responses accepted without material correction / candidate responses reviewed.	Keep separate from rubric results and end-user satisfaction.
Critical technical errors	Reviewed technical responses with at least one domain-reviewer-defined critical error / all reviewed technical responses.	Show count and denominator. Log errors blocked before release separately from those sent to users. Pause affected use for review.

Comparison discipline

Where feasible, use a phased rollout with a concurrent comparison group. A simple before-and-after change is an association: describe workload, staffing, language and equipment differences. Do not extrapolate three selected pilots to the whole market.

Review a sample with a second domain reviewer and resolve disagreements. Freeze the test rubric before evaluation. Record model or retrieval changes and later test-set reuse so improvements remain interpretable.

Measure adoption and conversion

Metric	Definition	Reporting rule
Escalation rate	Requests needing further engineer intervention for missing, conflicting or safety-sensitive information / all eligible requests.	Routine approval is not escalation. Classify reasons; a lower rate is not automatically better.
Weekly active users	Distinct eligible customer users performing a substantive ticket or knowledge action during seven defined days.	Exclude presenters, test accounts and background actions. Until shared roles exist, report account activity and observed session participants separately.
Onboarding effort	Team hours preparing documents, configuring access, training and setting up an account.	Include founder time and rework. Report paid delivery cost separately from hours.
Paid conversion	Pilot customers with accepted commercial terms and required payment received / pilots reaching the agreed decision date.	Also show the absolute result against the original three-pilot target. List incomplete pilots, refusals and unpaid commitments separately.

Minimum manual observation log

For each eligible ticket, capture: pilot and ticket ID; equipment model and language; question category; request time; first meaningful response time; active handling minutes; candidate and source IDs; reviewer decision; material correction; escalation reason; any critical error; closure time; and repeat contact after seven days.

At account level, record onboarding hours, founder hours, paid delivery costs, active-use observations, commercial offer and decision date, accepted terms, payment received and the evidence reference. Store only authorised information and limit unnecessary personal data.

Weekly review agenda

Review	Action
Quality and safety	Inspect critical errors, unsupported answers, applicability conflicts and correction patterns. Pause affected use when required.
Use and workload	Compare eligible-ticket counts and case mix. Investigate missing logs, low coverage and adoption barriers.
Delivery effort	Review onboarding, support and rework. A 40-hour setup threshold triggers investigation rather than automatic expansion.
Buyer evidence	Record written offers, acceptance, refusal, discount requests and required payment. Avoid substituting enthusiasm for conversion.
Next decision	Assign an owner and date to each action. Continue, narrow or pause the pilot based on documented evidence.

Report value accurately

Recovered engineer hours are capacity value unless actual spending falls. Keep avoided travel, parts mistakes and cash costs separate, with supporting records and no double counting. Financial projections in the business plan remain illustrative model outputs.

Check launch and submission readiness

Use this as a working checklist. Items below are requirements to verify or complete; they are not certifications or a claim that every check has passed. Keep a dated release record with owner, result and evidence link.

Checkpoint	Evidence required before proceeding
Route behaviour	Verify the current public URL and links. Confirm /demo changes reset on refresh; /workspace saves only to the signed-in user; review, close and JSON export behave as described.
Access boundaries	Test two authorised test accounts. Check direct record IDs and file URLs for unauthorised access. Review actual deployment, storage access and sign-in/session behaviour before customer-data use.
Source quality	Use approved or fictional documents. Confirm model, revision and page references. Have a domain reviewer approve bilingual excerpts and technical meaning.
Data handling	Agree allowed data, access, retention, deletion and hosting/processing arrangements with each pilot sponsor. Verify the actual removal process before promising deletion.
Recovery	Keep a labelled backup walkthrough. Verify export contents; test restore separately if it will be offered. Record a support contact and incident triage procedure.
Independent customer use	Complete required role and approval enforcement, safety routing, deployment review and supervised evaluation. Otherwise keep the scope as a sandbox study.
Commercial readiness	Obtain written pilot scope, sponsor and reviewer. Confirm fees, credit treatment, decision date and payment process. Billing automation and messaging remain future work.
Submission consistency	Use matching project contact details, scope and illustrative financial figures across the site, presentation and plan. Keep personal team details in the presentation only. Label fictional data and recorded demos.
Decision evidence	Store test logs, interview notes, pilot status and unresolved issues. Do not claim customers, revenue, patents, safety certification or completed compliance audits without evidence.

Final presenter check

- Show one supported question, one insufficient-evidence case and the human review and closure workflow. Say which account and route are used.
- Keep the request concrete: three prospective pilot manufacturers, equipment-service mentoring and resources for the bounded validation programme.
- Prepare a short answer for each remaining capability: external LLM, extraction, shared roles, messaging and billing. Describe these as planned work, with no completion dates promised before prerequisites are understood.

Use the brand consistently



The logo combines an EB letterform, a bridge and a document motif. The intended meaning is a connection between equipment-service participants, supported by organised technical knowledge. Use the supplied master artwork rather than recreating the lettering.

Element	Value	Use
Cobalt	#2F4DF2	Links, emphasis and key actions
Navy	#10213E	Headings, body text and identity
White	#FFFFFF	Primary background and clear space

Practical usage rules

- Place the full logo on white or a plain light background. Preserve its aspect ratio, colours and legibility. Do not stretch, rotate, crop, add shadows or place it over busy imagery.
- Keep clear space on every side of at least one quarter of the EB symbol height. This is a practical project guideline, not a claim of a registered brand standard. At small sizes, confirm the document and bridge details remain readable.
- Use Arial or a similarly clear sans serif for project materials. Use navy headings, cobalt highlights and generous white space. Maintain readable contrast for labels and links.
- Do not add a registered trademark symbol or imply trademark registration, legal clearance, patents, awards or certification. Those statuses are not established by this project pack.

Public project and contact details

Website

<https://equipbridge-ai-bo.bcao0016.chatgpt.site>

Public sample workflow

<https://equipbridge-ai-bo.bcao0016.chatgpt.site/demo>

Saved owner workspace

<https://equipbridge-ai-bo.bcao0016.chatgpt.site/workspace>

Email caobuo218315@126.com

China phone +86 176 8475 1441

Australia phone +61 401 254 388

Project materials are prepared for coursework and startup validation. Update this handbook when a release changes the implemented workflow, access model or data handling. Keep the evidence status and next decision visible.