



BEXTIO — 30–90 Day AI Pilot Guide V1.0

Guide

Run a short, governed and measurable AI pilot that produces a real decision.

BEXTIO

Public document • English edition.

Référence / Reference: BXT-GUIDE-PILOT-001-EN • V1.0 • Octobre / October 2026

bextio.ci • direction@bextio.ci • Abidjan, Côte d'Ivoire

1. What a pilot should prove

A pilot is not a demo. It tests a hypothesis in a real, limited and controlled context.

The central question is: "Does this system improve the process enough to justify its cost, risk and the effort required to deploy it?"

- Explicit hypothesis.
- Baseline.
- Real users.
- Metrics.
- Risk controls.
- Decision at the end.

2. Choose 30, 60 or 90 days

Duration depends on complexity and the speed of learning.

- 30 days: limited prototype, few integrations, simple metrics.
- 60 days: multiple user cycles, corrections and first integrations.
- 90 days: more complex process, adoption, governance and scaling decision.

3. Define the problem and hypothesis

The pilot tests a relationship between an intervention and an outcome.

Example: "If advisors receive a source-grounded assistant on the product corpus, meeting preparation time will fall by 30% without a significant increase in correction rate."

- Problem.
- Intervention.
- User.
- Expected outcome.
- Success threshold.

4. Define the scope

Reducing scope improves the quality of evidence.

- One process.
- One user group.
- Defined corpus or dataset.
- Controlled environment.
- Clear exclusions.

5. Establish the baseline

Measure before building.

- Average time.
- Error / correction rate.
- Cost.
- Volume.
- Delay.
- Conversion or satisfaction when relevant.

6. Prepare data and knowledge

Sources determine the quality of the test.

- Authorized data and documents.
- Quality and completeness.
- Versioning.
- Permissions.
- Sensitive data.
- Grounding / citations where possible.

7. Define users and owners

Clear roles prevent an ownerless pilot.

- Sponsor.
- Business owner.
- Pilot owner.
- Test users.
- Data/IT.
- Security/risk.
- Final decision-maker.

8. Design human oversight

Human-in-the-loop must specify where and how people intervene.

- Validation before external action.
- Escalation when uncertain.
- Human override.
- Stop authority.
- Logging of corrections.

9. Build the minimum testable system

Do not industrialize before you have learned.

- Minimal prototype or agent.
- Versioned prompts / instructions.
- Limited corpus.
- Useful logs.
- Metric instrumentation.

10. Test controlled cases

Before real users, test quality on a known set.

- Normal cases.
- Edge cases.
- Ambiguous input.
- Missing data.
- Out-of-scope attempts.
- Appropriate security tests.

11. Launch with real users

Observe real behavior, not only technical performance.

- Actual time saved.
- Usage frequency.
- Correction rate.
- Workarounds.
- Uncovered questions.
- Satisfaction and trust.

12. Track costs and incidents

A pilot should reveal full cost and operational risk.

- Licenses.
- API / compute.
- Integration.
- Support.
- Human time.
- Incidents and near misses.

13. Measure ROI / value

Time saved creates value only when it improves real capacity.

- Financial value.
- Time reallocated.
- Quality.
- Capacity.
- Revenue enablement.
- Risk reduction.
- Full cost.

14. Run the final review

The final review compares baseline, outcomes, adoption, cost and risk.

- What worked.
- What failed.
- Conditions for success.
- Limitations.
- Required controls.
- Transferable lessons.

15. Make a decision

Five decisions avoid the false “deploy or fail” choice.

- STOP.
- REDESIGN.
- EXTEND.
- SCALE.
- INDUSTRIALIZE.

16. Prepare the next stage

A positive decision opens a new stream: scaling, operations and governance.

The next plan should define architecture, ownership, budget, support, monitoring, change management, governance, training and review cadence.

- 3–12 month roadmap.
- Backlog.
- Budget.
- Operating model.
- Risk controls.
- Success metrics.

NEXT STEP

Scope a 30–90 day pilot: </en/businesses/scoping-conversation/>

NOTE

This document is a non-contractual methodological overview. Any engagement must be adapted to the sector, applicable law, data, security requirements and organizational context.