

Partner-Branded AI/Data Readiness Sprint One-Pager

For fractional CTOs, CDOs, and data advisors who need technical evidence behind an executive roadmap.

Why this sprint exists

Executive teams often know AI and data work matters, but they do not know which initiative can safely move first. A roadmap without technical inspection can create budget, hiring, and credibility risk.

The readiness sprint gives the advisor a practical technical layer:

- What systems exist today
- Which data is usable
- Which AI/data initiatives are feasible
- What blocks the first implementation
- What should be funded first

The advisor keeps executive authority. MLDeep supplies the technical reality check.

When to use it

Use this sprint when the client is asking questions like:

- Can current data support AI agents or automation?
- Should the company hire a data leader, data engineer, analytics engineer, or AI engineer first?
- Which roadmap item can ship in the next 30-90 days?
- Is the current warehouse, CRM, BI layer, or spreadsheet process good enough?
- Should the client build internally or bring in specialist delivery?

Inputs needed

The sprint works best with:

- Current roadmap hypothesis
- List of candidate AI/data use cases
- Current data stack and system inventory
- Sample dashboards, reports, or recurring spreadsheet workflows
- Known pain points from leadership, operators, sales, marketing, finance, or CS
- Access to technical owners for short interviews
- Read-only access or screenshots where direct access is not appropriate

Technical inspection areas

1. Data availability

Does the data required for the use case exist, and can it be accessed reliably?

Checks:

- Source systems
- Ownership
- Refresh cadence
- Historical completeness
- API or export path
- Permission constraints

2. Data quality

Is the data consistent enough to drive a model, workflow, dashboard, or automation?

Checks:

- Duplicates
- Missing values
- Stale records
- Definition drift
- Manual overrides
- Unclear lifecycle states

3. Integration path

Can the output land where the team needs to use it?

Checks:

- CRM, BI, Slack, email, app, warehouse, or ticketing destination
- API limitations
- Human approval loops
- Error handling requirements
- Monitoring and ownership

4. Operating readiness

Can the organization maintain the workflow after launch?

Checks:

- Internal owner

- Review cadence
- Runbook needs
- Escalation path
- Security and governance constraints
- Change management burden

Outputs

The sprint should produce:

- System and data flow map
- Readiness score across core dimensions
- Use-case feasibility ranking
- Risks and blockers
- Recommended first-win scope
- 30-90 day implementation path
- Advisor-facing summary for the executive team

Decision rules

Green

Move to a first implementation sprint when the use case has available data, a clear operating owner, manageable integration complexity, and visible business value.

Yellow

Do a foundation repair first when the idea is valuable but blocked by definitions, syncs, ownership, permissions, or missing data model work.

Red

Do not fund buildout yet when the use case is vague, the data does not exist, the operating owner is missing, or the executive team cannot name the decision the system should improve.

NEXT STEP

Bring one client scenario to a partner fit call.

20-min call. Bring the client situation, the technical question, and the decision that needs to happen next. Leave with a written go/no-go — no slide deck.

Book a partner fit call →

or email info@mldeep.io