

Fictional demonstration

Choosing the next hire for an AI team

I'd open the platform lead search first in this example. The immediate constraint is getting the product into dependable pilot use. The CTO can cover the current model work while the team checks whether that assumption holds. I'd keep the applied ML brief ready and revisit the sequence if evaluation shows a different constraint.

The assumed business situation

This fictional seed-stage team has five engineers, a CTO and a founder. It is preparing an enterprise workflow product for pilot use. The founder can approve one engineering hire now. A second hire requires a later budget decision. Every number, observation and plan in this memo is illustrative; none describes an actual client or achieved result.

Assumed evidence	What it changes
Two of five pilot workstreams are waiting for a repeatable deployment process.	Deployment ownership is affecting the next milestone.
Three of twenty staged workflow runs need manual recovery.	Reliability needs investigation before a larger pilot. This is a small test sample.
The CTO has eight hours a week available for applied ML work.	Temporary coverage may be possible, but the founder must protect that time.
Model evaluation has no agreed pilot acceptance threshold.	The team still needs to establish whether model quality is acceptable.

The decision to make

Approve discovery for a platform lead who will own deployment, observability and failure recovery. Do not promise that this hire will solve model quality. The founder and CTO should agree on pilot acceptance criteria before committing the vacancy to a final role brief.

Compare the two hiring choices

Choice	Work it would own	Dependency to resolve
Platform lead first	Reliable deployment, monitoring and recovery; technical ownership across the pilot.	CTO confirms that applied ML work fits the temporary coverage plan.
Applied ML engineer first	Evaluation, error analysis and changes to model behavior.	Existing engineers must own production reliability without delaying the pilot.

I would test the platform recommendation against a short review of the blocked work. If those items are actually caused by poor model behavior, the order could change. A hiring title should follow the work that needs an owner.

Example platform lead brief

Required work	Evidence to discuss
Take a service from manual deployment to a repeatable release process.	A concrete release workflow, rollback decision and explanation of who operated it.
Investigate failures across application and infrastructure boundaries.	A failure the person diagnosed, the observations they used and what remained uncertain.
Make ownership more clear for a small engineering team.	How responsibilities and escalation changed, including work they delegated.

Deep model training expertise would be useful, but it is not a required skill for this brief. The applied ML role would have its own evaluation criteria. Combining both roles into one vacancy would make the tradeoffs harder to assess.

What would change my recommendation

Move the applied ML search first if a reviewed evaluation shows that model behavior prevents the agreed pilot use case and the team already has a workable reliability owner. Keep platform first if failures are chiefly in release, integration or recovery and model work has credible temporary coverage. If both constraints are critical, revisit budget or reduce the pilot scope.

Example hiring sequence and decision record

Review point	Owner	Required output
Before search approval	Founder and CTO	Pilot acceptance criteria, one approved vacancy and a named owner for each blocked item.
Week one	CTO	Review of deployment failures and model evaluation gaps; confirmation of temporary ML coverage.
Week two	Founder	Approve or revise the platform brief based on the review. Record why the sequence still fits.
Weeks three to four	Recruiting owner	Assess relevant profiles against the agreed work; report missing evidence and interview capacity.
Week six or earlier if a constraint changes	Founder and CTO	Decide whether to open applied ML, continue the first search or change the pilot scope.

The six-week review is an example decision point, not a promised hiring timeline. An offer or a second search needs its own approval.

Completed example decision entry

Proposed decision: investigate and open platform first. Decision owner: founder. Basis: deployment work is blocking two pilot workstreams, while the CTO may cover applied ML temporarily. Unresolved: the model acceptance threshold and the true cause of the failed runs. Revisit when either assumption changes.

Advisory scope for this brief

I would help the founder compare the roles, clarify the evidence needed and keep the hiring sequence connected to the milestone. Active sourcing, interview scheduling, implementation and placement work would need a separate agreed scope.

Next step

Bring the current milestone, a list of blocked work and the people who can own it to a calibration session. The useful output is an agreed vacancy and a decision record the team can revisit.