

FICTIONAL DEMONSTRATION

A technical talent map with evidence boundaries

Fictional brief: a senior infrastructure engineer for a US-based company operating real-time AI services. This excerpt illustrates how I would segment the search and explain fit. It contains no real candidate identities, verified salaries or availability claims. An actual mapping delivery contains 15–20 researched profiles for one agreed role.

Market segment	Why examine it	What to verify
Real-time communications platforms	Work on low-latency services and operational reliability may transfer.	Actual infrastructure ownership, failure diagnosis and service scale.
Developer platforms	Experience supporting external users and reliable APIs may be relevant.	Depth of distributed systems work and production responsibility.
ML infrastructure teams	Model serving and inference operations may match part of the brief.	Whether the person owns platform reliability or model research only.

Fictional profile	Illustrative evidence	Gap before considering outreach
Profile A	Owned service reliability and a documented incident review.	Confirm actual scale and current role.
Profile B	Built observability for a distributed platform.	Confirm architectural decision ownership.
Profile C	Worked on model-serving cost and latency tradeoffs.	Check platform depth beyond model deployment.
Profile D	Managed API infrastructure used by external developers.	Confirm hands-on technical scope.
Profile E	Led an infrastructure migration with rollback planning.	Verify relevance to the target operating environment.

For an actual profile, record the source URL, date checked, current employer and role, specific evidence for fit, unresolved gaps and a reason to prioritize it. Keep identity and availability separate. Do not infer willingness to move, compensation, eligibility or interest from a profile. Research sources must work without CJ having a personal LinkedIn account.