



Senior Software Engineer, Short-Message Data Processing

- ✓ Remote
- ✓ Full time
- ✓ Contractor
- ✓ Time zone EST
- ✓ Payment in US Dollar

About the Role

We are hiring a Senior Software Engineer to build a purpose-built processing engine for large-scale short-message and chat data —the kind of communications data that now sits at the center of most litigation and regulatory matters, and that no tool on the market handles well at volume.

A working prototype will be in hand when you start. Your job is to take it from prototype to a production system the practice runs every day: scalable, observably correct, and operable by practitioners rather than by its author.

You will be the senior engineer on a small, tight product team — working directly with a Product Manager and a Designer, with no layers in between. You own the technical decisions: architecture, data model, pipeline design, infrastructure, and the shape of the practitioner-facing application. You will also inherit knowledge from the engineer who built the system this replaces, and turn what currently lives in one person's head into documented, tested, version-controlled software.

About the Team

This role sits within the Platforms at AlixPartners, alongside experienced product managers, designers, and engineers building across the full product development lifecycle. You will not only build products but help establish engineering excellence within a Risk Advisory practice spanning eDiscovery, digital forensics, and beyond — bridging startup-grade engineering craftsmanship with the rigor of a world-class advisory firm.

What You'll Do

Own the build for taking a validated prototype to production — architecture, data model, ingestion design, storage, and integrations.

Build ingestion that survives real-world data: deeply nested export structures (recent single collections have exceeded a million folders), multiple export formats, schema variability, and structural anomalies that must be flagged rather than silently skipped.

Design for volume as a first-class requirement: multi-terabyte collections and tens of millions of records processed without manual splitting, using streaming, chunking, and containerized horizontal scaling.

Model conversation data properly — messages, edits, deletions, replies, reactions, attachments, and time-based participant membership — with stable identifiers and relational links that hold up when a later export rewrites historical records.

Make correctness observable: deterministic quality gates in code, run-level audit trails capturing which build and which settings produced an output, completeness reporting, and explicit reason codes for anything missing. Automated checks are the product, not an afterthought.



Establish engineering foundations: CI/CD, automated testing, observability, coding standards, and security posture on our internal development platform.

Build the self-serve application with the Designer so a processing team — not a single expert — can run jobs, monitor progress, and produce reports in minutes rather than a week.

Work day-to-day with the Product Manager to sequence work, scope complexity, and surface trade-offs, balancing near-term delivery against the platform this becomes. You will be in the room for scoping decisions, not handed a spec.

Absorb and document institutional knowledge from the legacy tool and its author, then retire it via a validated, defensible cutover.

Set the technical bar for engineers who join later as the platform expands to additional data types.

What We're Looking For

Required

- ✓ 7+ years of software engineering experience, with substantial time building data-intensive systems in production.
- ✓ Demonstrated experience designing large-scale data ingestion and processing pipelines — throughput, memory, and
- ✓ parallelism under real constraints; schema variability; deduplication; transformation at scale. You have personally debugged a pipeline that fell over on volume and fixed the design, not just the symptom.
- ✓ Experience in a domain where provenance and reproducibility are the deliverable — regulated, audited, forensic, financial reconciliation, clinical, or compliance data. You treat “we can’t explain why these two runs differ” as a defect of the same severity as a crash.
- ✓ End-to-end ownership of at least one system you built from scratch, including the early architectural decisions that defined it for years.
- ✓ Strong command of modern application and cloud architecture — you can design and implement a scalable service, including infrastructure, APIs, data stores, and the interface practitioners actually use.
- ✓ Comfort working closely with product and design. You can hold your own on scope and sequencing, translate technical constraints into terms a PM can decide on, and partner with a designer on an interface rather than bolting one on at the end.
- ✓ Fluency with modern AI-assisted development. You use these tools daily to move faster, and you have clear judgment about where they belong and where they don’t — our quality gates are deterministic by design.
- ✓ Startup mindset: comfortable with ambiguity, pragmatic about trade-offs, and able to make decisions that balance speed against long-term soundness.

Strongly Preferred

- ✓ Experience with messaging, chat, or communications data — Slack, Teams, mobile message extractions, or similar. Conversation threading, participant attribution, and attachment resolution are harder than they look.
- ✓ Experience replacing a legacy system while it remains in production, including parity validation against the tool being retired.

- ✓ Background in legal technology, eDiscovery, investigations, or compliance tooling — genuinely useful context, but something a strong engineer can learn here. We are more interested in how you think about completeness than in whether you already know our vocabulary.
- ✓ Experience building integrations with third-party platforms via API, particularly where you do not control the target system's data model.
- ✓ Experience introducing modern engineering practices (CI/CD, automated testing, observability) in an organization new to formal software development.
- ✓ Experience working alongside a professional services or consulting delivery team, where your users are practitioners under deadline pressure on high-stakes matters.

Nice to Have

- ✓ Experience with search and query layers over large datasets (Elasticsearch, Splunk, or similar), including saved and reusable query conditions.
- ✓ Experience designing modular, plug-in architectures where new data sources can be added without destabilizing existing ones.
- ✓ Familiarity with entity resolution or identity normalization across inconsistent identifiers — usernames, phone numbers, aliases, email addresses.
- ✓ Exposure to chain-of-custody and data security requirements in litigation or investigations contexts.
- ✓ Experience with incremental and delta processing — recomputing only what changed when a dataset is re-delivered or a request expands.

Why This Role

You will replace a system that is a single point of failure for a practice handling hundreds of matters a year — the impact is immediate and unambiguous.

The hard problems are genuinely hard: volume no commercial tool absorbs, adversarial data structures, and a correctness bar set by regulators and courts rather than by a dashboard.

You will build on a prototype that de-risks the direction, with no legacy architecture constraining how you design the real thing.

Small team, short path: a PM, a designer, and you. Decisions get made in the conversation you're already in.

You will bridge two worlds — the rigor and domain expertise of a top-tier advisory practice and the speed and craft of modern software engineering.