

ASHAY KUBAL (HE/HIM/HIS)

(201) 554-9947 | ashay.kubal@gmail.com | [LinkedIn](#) | [Website](#) | [Github](#)

Professional Summary

Product builder with 20+ years in Engineering, including 10+ years leading Product. I'm a hands-on, AI-native builder who lives in Claude Code and Cowork, prototyping and shipping from the CLI. I build Skills and MCP connectors, plus the plugins that wire them into real systems and application prototypes that can be iterated upon at speed. I've shipped AI features: LLM-powered document summarization for Equity Research, and automated interaction capture from Bloomberg chat, where a deterministic pipeline classifies and an LLM surfaces trading signals before logging to CRM. An agentic search assistant is in discovery & design. I prototype my own product hypotheses and back the ones the data supports, defining the evaluation, guardrails and auditability that AI needs in a regulated, high-stakes domain. I grew daily platform adoption from roughly 35% to 80-90% by sitting with the people who used it, while analyzing usage to refine the roadmap.

AI/ML Leadership: Currently part of the AI leadership working group at Jefferies, that enables engineers and builders to set up effective, governable & observable workflows in Claude Code. As part of this group, I've launched plugins on the org managed plugin marketplace that augment the Claude Code harness. Also building an internal application for Product that connects knowledge and delivery systems, with AI-powered features for user research, PRD creation, automated user story generation & delivery reporting.

Skills

- **Product Management & Strategy:** Product Strategy, Product Sense, 0-to-1 Product Lifecycle Management, Market Research, Competitive Analysis, User Research, Stakeholder Management, Roadmap Prioritization, Technical Documentation
- **AI/ML & Agent Engineering:** Agentic RAG, Multi-agent architecture (orchestration, memory, tool usage), Agent skill design, Sub-agent patterns, Context engineering & memory management, LLM-as-a-judge & HITL evals, Prompt engineering. Familiar with: LLM classification & fine-tuning, Reinforcement learning (RL), Guardrails & safety, Natural Language Processing (NLP), Machine Learning (ML) concepts
- **GenAI Platforms & Tooling:** Claude Code, Cursor, OpenAI Codex, CoWork & Design, Figma & Figma Make, multi-agent patterns, Workflow orchestration, Plugin & skill development, MCP server & connector development, Design systems & functional prototyping with Claude, familiar with Claude Agent SDK & API
- **Infrastructure & Delivery:** Familiar with API design, performance & latency optimization, cloud architecture & infrastructure (AWS, microservices & event-driven architecture)
- **Search & Data Platforms:** Federated & semantic search, Relevancy ranking, Real-time indexing, Data pipeline architecture, Search quality evaluation, Usage & feedback analysis
- **Enterprise Platforms:** Deep experience in CRM & CMS systems. Familiar with Data modeling, Data warehousing, ETL, ELT, Data governance.
- **Tech stacks:** Salesforce, MongoDB, AWS OpenSearch, Kafka, Redshift

Open-Source Portfolio

- **[SDLC Enforcement & Governance Framework](#):** Claude Code plugin providing multi-stage orchestrator skills and sub-agents that guide SDLC from ideation, research, brainstorm & planning to code review & test audit, with governance and logging. Code and test governance enforced deterministically using programmatic hooks.
- **[Code Intelligence & Knowledge Management System](#):** a Claude Code plugin that builds an interlinked web of markdown knowledge as code is written and keeps it current as files change. When a file is touched, its relevant bite-sized knowledge loads automatically into context; the knowledge base is pruned and updated programmatically through deprecation and supersession, with on-demand additions, and is organized into primitives: lessons learned, technical decisions, business rules, architecture patterns, people, and more. Filesystem-based, with no external MCP or subscription.
- **[Agent & Skill Toolkit for Structured Development](#):** Toolkit of meta-skills and sub-agents that enable structured development with observable and measurable output.

Advisory Experience

Human Presence Inc. – New York, USA | **Advisor - Product Strategy & Applied AI**

May 2026 – Present

- Building production-grade custom Claude plugins and Skills for RIAs, plus the MCP connectors that wire them into CRM and portfolio systems on top of AWS Bedrock.
- Advising product strategy for Advisor OS, a governed AI workspace for wealth-management firms, translating advisor and paraplanner jobs-to-be-done into per-tenant plugin configurations.

- Led competitive and architecture research across the wealth-tech AI stack, weighing build-versus-partner options and matching solutions to each firm's pain points.
- Advised buy-to-stand-up with a build roadmap for the AI context layer: scanned the market and recommended partnering on a semantic-layer platform (Clarista) and graph database (ArangoDB) to stand up fast, with a path to build in-house as the firm scales.
- Authored the enterprise MCP and plugin platform proposal: custom MCPs with auditing gated behind AWS Agentcore gateway and a relationship-aware entitlement model with per-persona tool scoping.

Professional Work Experience

Jefferies LLC – New York, USA | Vice President | Product Lead – CRM, AI, Data & Search

Jul 2018 – Jun 2026

LLM Assisted automated interaction capture from Bloomberg Terminal (Jan 2026 – Present)

- Conceptualized and implemented a rules engine for signal validation, participant identification, and structured conversation extraction from raw Bloomberg Terminal XML
- Developed multi-stage data ingestion & funnel pipeline using Python (spaCy, NLTK) to transform raw Bloomberg Terminal chat data into clean, structured conversations
- Participated in fine-tuning & evaluations of LLM classification of Bloomberg chat signals with guardrails for filtering small-talk, trade signals, exclusions, and sensitive commentary, including feedback loops for reinforcement learning (RL)
- Architected the enrichment of raw interactions with CRM reference data
- Prototyped human-in-the-loop review pipeline with model RL feedback cycle, enabling rapid evaluation and iterative model output improvement

Enterprise AI Search Assistant (Jan 2026 – Present)

- Defined product vision for an AI-powered CRM search assistant using advanced RAG and multi-agent patterns
- Assessed modern LLM search architectures: single-agent with advanced RAG vs. multi-agent orchestration patterns
- Finalized manager-worker pattern after detailed architectural assessment, selecting the pattern with best tradeoffs for query complexity and maintainability
- Led product discovery, strategy, and UX design for the AI-powered search solution
- Used Claude Code and Figma Make to rapidly prototype, capture early user feedback, and iterate on search patterns

Research LLM Summaries (Jul 2025 – Dec 2025)

- Launched an AI-powered LLM solution for rapid summarization of research documents, accelerating content distribution pipeline
- Led discovery and prioritization of multi-document summarization, portal integration, and enterprise search integration
- Reduced time-to-distribute from roughly 15–20 minutes (manual) to under a minute via pre-formatted LLM templates, delivering 85–95% time savings for content operators
- Introduced source citation and hyperlinked references for transparency and user trust
- Conducted in-depth human-in-the-loop evaluation sessions with analysts to refine summarization quality, tailor output to document types, and improve user trust in LLM output
- Implemented feedback mechanisms within the summary interface for continuous model refinement via RL

Enterprise Global Search (Feb 2025 – Jul 2025)

- Launched enterprise federated search aggregating CRM data from 20+ datasets, replacing universally disliked legacy system
- Authored detailed API contracts for typeahead & advanced search services, defining request/response payloads, error handling, latency targets (<100ms p50) and uptime SLAs, demonstrating developer-centric API design
- Applied semantic search by indexing user-specific datasets with differential weighting, driving 3x shift in search traffic to the CRM application; 75% of all queries originated from CRM despite homepage remaining the default browser URL
- Built a real-time keyword-based typeahead, reducing response time from 5s to <100ms measured via Playwright performance harness
- Designed advanced search with filtering, delivering results in 200–500ms at p50 measured via production instrumentation
- Increased user productivity with saved filters and full keyboard navigation with hotkeys for power users
- Architected real-time indexing between MongoDB on AWS and Amazon OpenSearch
- Managed UAT with diverse user groups, refining search algorithm for relevancy based on direct feedback
- Achieved universal satisfaction (n=50 surveyed users), with search consistently cited as the first tool users used - up from widespread dissatisfaction with prior system

Contact Recommendation Engine (Jul 2024 – Dec 2024)

- Launched an ML based contact recommendation application for marketing and analyst-sales touchpoints, delivering measurable business impact
- Led discovery & design of a customizable ML scoring engine using interaction, readership, and coverage data, achieving 70% adoption across roughly 700 research analysts and sales users within the first quarter

- Shipped pre-configured scoring templates for common use cases, reducing setup friction for non-technical users

CRM, Data & Intelligence Platform for Capital Markets (Jan 2021 – Present)

- Led a cross-functional team (25+ engineers, designers, product owners, analysts) delivering a 0-to-1 CRM, Data & Intelligence platform serving multiple business groups
- Led product strategy and launch of the CRM platform using event-driven data pipelines & microservices on AWS; unified data from Salesforce and 10+ enterprise sources
- Architected core CRM workflows such as Client Onboarding, KYC, Contact & Coverage Management, Meeting Management II Campaigns, Opportunity & Deal Management, and Revenue Attribution
- Implemented Master Data Management (MDM) for contact records by consolidating contact & client records with golden IDs
- Grew daily active usage from roughly 35% to 80–90% of in-scope users (measured via normalized daily logins), replacing a universally disliked SaaS product and offline Excel workflows
- Implemented strict schema validation checks within the custom CRM, achieving and maintaining 95% data accuracy across datasets (validated via weekly automated reconciliation)
- Maintained full compliance with ethical wall requirements between Research and Banking
- Led discovery, user research, product design, and implementation of core CRM features: Client Onboarding & Management, Contact Management, Client & Contact Coverage, Interaction Management, Lists Management
- Owned the end-to-end data lifecycle, defining ingestion, enrichment, and modeling requirements for core CRM entities
- Launched Reporting Dashboards enabling senior management transparency and intra-team collaboration
- Architected data cleanup strategy achieving 98% data accuracy prior to platform launch
- Implemented hierarchical taxonomy of regions and products for granular regional reporting down to city level
- Led daily scrums, sprint planning, product backlog maintenance, prioritization and all scrum ceremonies including retrospectives

Capgemini America Inc – New York, USA | **Sr. Consultant – Financial Services & Fintech**

Nov 2010 – Jul 2018

Research CMS & Distribution (Product Owner for Morgan Stanley)

- Revamped firm's Research platform, transitioning from legacy Word/Excel add-ins and PDF distribution to bespoke CMS (authoring, publishing, distribution, analytics)
- Developed bespoke research authoring platform with rich-text editing and embedded media support (SVG, PNG), enabling analyst-created printer-friendly, HTML-based documents with integrated PDF downloads
- Designed automated publishing workflow replacing manual multi-step review process with digital pipeline
- Built a custom Research Portal that became the primary access point for the firm's research content, with detailed usage logging by region, channel, and document type
- Implemented usage logging and tracking for research content across dimensions (region, channel, document type)

Previous Work History

Citiustech Solutions – Mumbai, India | **Sr. Quality Assurance Engineer**

Jan2010 – Nov 2010

Mindtree Ltd. – Pune, India | **Sr. Software Engineer**

Jul 2006 – Jan 2010

Education

Mumbai University – India | **Bachelor of Engineering (Computer Science)**