

Muhammad Asad Rasheed

AI-Powered Tech, Growth & Business Operator

asadrasheed.com asad@asadrasheed.com Lahore, Pakistan Cognilium AI

SUMMARY , THE WHOLE WORK, AND HOW

Muhammad Asad Rasheed is an AI-powered tech, growth and business operator at Cognilium AI (the parent) and its products Paralegent AI and Build Buy Software. He does what a company normally splits across a whole team and ships it end to end: he **builds the tech** (the websites and their design systems, the full SEO/GEO/AEO and News-SEO stack, automation scripts, Python data pipelines, and the Sanity content engine) and **runs the growth** (content, social and animations, the brand, organic reach, and the full outbound and business-development motion , cold email on Instantly, prospecting in Clay, the LinkedIn social-selling loop, B2B targeting, partnerships, and the agency's Upwork pipeline).

How: agentic AI is his execution layer , he deploys and orchestrates it to build and ship, while the judgment stays the job: knowing what to build, catching what the tools get wrong, and moving a business rather than just advising it. The arc ran content and social → growth and strategy → the full tech, SEO and content engine → growth-engineering (a legal-tech directory built to feed a product) → operator across all three brands. Alongside that he built and shipped his own product, **mcpcall** (an MCP server on npm), solo, from his own idea. Self-taught, 19, based in Lahore, Pakistan, shipping on live main , consistency beats credentials.

AI IMPLEMENTATION & AGENTIC WORKFLOWS (THE CORE)

- Claude Code as primary pair-programmer (central, not occasional)
- Agentic AI as the primary execution layer , end to end across the team
- Authored 7 custom Claude Code Agent Skills (runtime contracts: design system, voice/vocabulary, page rhythm, conversion psychology, FAQ builder, visual design + cross-refs)
- Skill-spec authoring (open Agent Skills format , frontmatter spec, glob-triggered activation)
- **CLAUDE.md / AGENTS.md** as an enforcement layer (not documentation)
- Persistent cross-conversation memory architecture (user/feedback/project/reference types)
- Skill-vs-memory separation discipline (PR-time enforcement vs cross-session facts)
- Prompt engineering for production workflows + image generation
- Parallel-version workflow (versioned routes, never edit live in place)
- Multi-agent compatibility (Claude Code, Cursor, Codex CLI, Gemini CLI, Antigravity)
- Multi-project + multi-app context discipline (parent scope rules encoded in writing)
- MCP server engineering

MARKETING & GROWTH

- Content creation + content management for Cognilium AI and Paralegent AI
- Animated content + motion graphics for social / marketing (animated posts, motion design)
- Social media management across platforms
- Organic social reach growth + community engagement (growing the brand's own audience, not just outbound)
- Growth strategy planning and execution
- Engagement frameworks, conversation strategy, customer-journey mapping
- Target-audience research and approach
- Buyer-persona work (GC, CFO, CIO, CPO, Legal Ops)
- UTM tracking strategy per acquisition channel
- 10-platform social banner kit produced as Python code
- Newsletter / audience capture (Resend, double opt-in)

OUTBOUND & BUSINESS DEVELOPMENT

- Cold email infrastructure end to end: sending-domain + mailbox setup, mailbox warming, deliverability management (SPF/DKIM/DMARC-aware, spam-trigger avoidance)
- Cold / warm sequence writing + follow-ups (copy written for replies, not just opens)
- Cold email sending in **Instantly** (campaign setup, sequencing, inbox rotation)
- Prospect-list building + enrichment in **Clay** (targeted ICP lists)
- LinkedIn social-selling loop: target list → post engagement → connection request → conversation → pitch
- Full outbound / SDR motion: open cold → qualify → pitch the product (Paralegent) or the services (Cognilium) → close small or hand off warm
- B2B target strategy: lists of startups, incubators, and VCs opened with the "integrate, AI-powered, end to end" angle
- Partnership outreach: reselling Cognilium services through partner companies
- Agency sales pipeline on **Upwork**: proposal writing + client invite/message handling for the CEO account and the Cognilium agency account (Zenfil job-feed → Telegram alert → audit → proposal → apply)
- Executive / founder brand operations: running the founder's authority profiles, knowledge base, résumé/CV, and remote-job applications across platforms

- JSON-LD @graph schema authoring (incl. the rare **DefinedTermSet**) + validation gates
- Canonical URL management at scale (43-page migration), 301 redirect strategy for rebrands
- robots.txt engineering with 18+ AI-crawler allowlist · llms.txt authoring
- IndexNow across 5 engines · Bing Webmaster API · Google Search Console API · Google Indexing API (service-account auth)
- GA4 Data API + Web Vitals (LCP/INP/CLS/FCP/TTFB)
- News-SEO + Google News carousel (NewsArticle schema, news sitemaps, RSS/Atom/JSON, Publisher Center)
- GEO levers from the Princeton paper · BLUF copywriting for AI extraction
- AEO-friendly FAQ authoring · metadata optimization · internal linking for unindexed pages · anchor-text diversification
- Sitemap automation · image SEO · CTR-to-position diagnostics · "Crawled but not indexed" recovery
- E-E-A-T signal building (experience/expertise/authoritativeness/trust across content + author + entity markup)
- Off-page authority + link building via platform sign-ups (the sameAs / citation loop)
- Entity SEO / personal-brand authority engineering: Person JSON-LD, sameAs networks, namesake disambiguation, knowledge-panel + brand-SERP strategy

CONTENT ENGINE & DISTRIBUTION (COGNILIUM SITE)

- Sanity headless CMS (next-sanity, GROQ, Portable Text, image CDN, tag-based ISR)
- Strapi → Sanity migration + SendGrid → Resend migration
- Sanity webhook receiver (HMAC-SHA256 signed, per-type/per-slug revalidation, IndexNow ping, Google URL_DELETED on deletes)
- **News-ingestion + curation pipeline**: configured sources, a server-side scheduled (cron) fetcher that pulls + relevance-filters, a custom **review dashboard** for human-in-the-loop approval, then publish-to-Sanity
- **Two-engine content system**: tech-news feed (5+ posts/week) + engineering blog (1-2/week), each with its own type/schema/route; manual / auto / skip framework per surface
- **Cross-platform auto-mirroring** (push): orchestrator firing on every publish across 17 registered platforms (Dev.to, Hashnode, Bluesky, LinkedIn Newsletter, Resend, Substack...) tracked in one source-of-truth inventory (live / paused / dormant), with per-post status tracking and an add-a-platform runbook
- **RSS / Atom / JSON feeds** (pull): the pull-distribution layer , auto-discovery for Google News, Feedly/Inoreader, Mastodon readers, and AI crawlers
- **Distribution reliability**: cron safety net (mirror-retry + token-health jobs) so the auto-publish self-heals
- **Blog authorship for authority + traffic**: wrote the engineering-blog posts (JSON-LD-rich, GEO-tuned) that drove measured results
- Topical authority at scale (pillar+spoke clusters, GraphRAG pillar, AI Engineering Glossary) · 56-route site IA

FRONTEND ENGINEERING

- TypeScript 5 (strict, generics, type-safe data layers)
- Tailwind CSS 3.4 / v4 with locked design-token architecture · tailwind-merge + clsx
- Framer Motion (scroll-reveal, mega-menu choreography, hero animations, transitions)
- Responsive / mobile-first (375/768/1024/1440), overflow debugging, layout containment
- Native <video> integration · image optimization (next/image) · accessibility-aware patterns (semantic HTML, ARIA)
- Radix UI + shadcn-style component layer · class-variance-authority · CSS token architecture

DESIGN & BRAND

- Brand identity creation on live production, zero downtime (full color pivot across 40+ pages in one sprint)
- Design-token architecture (color/type/spacing as code) · design-system authoring
- Logo design + recolor, favicon production (multi-size SVG), wordmark
- Mega-menu / header / footer systems · hero pattern design · page-rhythm design (12-section template)
- Multi-brand discipline (Paralegent + BuildBuySoftware, locked tokens, no contamination)
- Brand-asset production-as-code · motion graphics / animated content

PYTHON & DATA ENGINEERING

- Production web scraping at scale (~1,700 polite API requests, identifiable UA, robots.txt respect, retry logic)
- REST API reverse-engineering (51 endpoints documented on dreamlegal.in, schema inference)
- Stack fingerprinting · pagination + enrichment orchestration
- SHA256 evidence preservation + immutable raw layer · idempotent re-runs
- Canonical slug normalization · duplicate/collision detection
- Cross-cutting pattern analysis at scale (companies × 9 dimensions) · JSON normalization at scale

BACKEND / INTEGRATIONS

- Next.js API routes (contact, demo, revalidate, indexnow, google-indexing)
- Resend transactional email · Strapi CMS with ISR · webhook-driven revalidation
- Microsoft Bookings embed · service-account credential management
- Prisma + Postgres data layer + relational schema design (the Build Buy Software directory: 0-100 scoring schema with anti-gaming flags, per-record provenance)

- stdio MCP server exposing 10 tools, bearer-authed + audit-logged
- Scoped npm package (@asad120414/mcpcall, npx-runnable)
- SaaS data model (workspaces/files/versions), version rows only on sha256 change
- Auth + security: tokens hashed (sha256 + timingSafeEqual), prefix-indexed, symlink-defended writes, self-host override
- Conflict-safe sync: sha256-keyed pulls writing .incoming on mismatch, drift-aware cleanup, reversible ops
- Deep git internals: hides files via .git/info/exclude (not .gitignore, which leaks names) , the core insight
- End-to-end solo execution: app + MCP server + data model + migrations in one monorepo

STRATEGY & BUSINESS / RESEARCH

- Competitive-landscape framing from traffic data (SimilarWeb), not narrative · synthetic/bot-traffic detection
- Founder/org profiling, freemium-ladder anatomy, pricing-model decomposition · whitespace validation
- Phased execution planning with load-bearing constraints · gap-analysis report writing
- ICP/persona-driven content strategy · vocabulary locking · conversion psychology · stat-drift management
- Market + positioning research · GEO/AEO mechanics for ranking a PERSON · job-market/platform analysis · identity-firewall design

BUILD TOOLING & AUTOMATION

- TypeScript CLI scripts (tsx) · 22+ SEO automation scripts (audit, submit, report, drift-check, validation)
- Diff-based IndexNow submission · auto-regenerating page-dates.json · canonical freshness pipeline (one file → JSON-LD, OG, sitemap, visible "Last Updated")
- Pre-publish checklist enforcement · GA4 Page Vitals · AI Overview citation tracking

PERFORMANCE, QUALITY & DEVOPS

- Core Web Vitals targeting (FCP <0.4s, LCP <2.5s, CLS 0, INP <200ms) · security headers (HSTS, CSP) · cache strategy
- ESLint + Tailwind plugin · TypeScript --noEmit gates · schema validation gates before push
- Vercel (Edge, ISR) · Cloudflare DNS/CDN
- Domain + email infrastructure: Cloudflare DNS, DNSSEC, DMARC/SPF/DKIM, Email Routing, zone-scoped API tokens (production domain set up + secured end to end)
- Self-hosted Strapi on Linux (SSH, PostgreSQL) · multi-app monorepo
- High-velocity git on live main (346 commits in 20 days) · pre-push hooks · npm-script operational interfaces

AI VIDEO & MEDIA (UPWORK SERVICE)

- Video editing and production: short-form and long-form editing (cutting, captions, sound, color, motion graphics) in CapCut and Premiere Pro
- AI marketing videos + commercials · UGC short-form (TikTok/Reels/Shorts) · voice cloning + AI voiceovers · AI character animation · cinematic product/brand videos · AI image generation · script-to-video
- Stack: Veo-3, Sora, Runway, Kling, Higgsfield, InVideo · ElevenLabs, Suno, Pika, Google Flow · Midjourney, Ideogram, Leonardo · CapCut, Premiere Pro

PRODUCT & FOUNDER-LEVEL EXECUTION

- Solo end-to-end product execution (mcpcall: app + MCP server + data model + migrations)
- Stood up a second production app from zero in parallel with the first
- Two production tracks at once without quality drop · brand-independence as a runtime constraint
- Documentation-as-enforcement · recovered pages from "Crawled but not indexed"

WHAT I'VE SHIPPED (PROOF)

- **mcpcall** , solo AI-native SaaS + stdio MCP server, live in free beta on npm
- **Cognilium AI site** , the SEO/GEO/News-SEO engine, the two-engine content system (blog I write + tech-news with ingestion dashboard), the migrations, the cross-platform auto-mirror, the outbound + BD pipeline
- **Paralegent AI** , go-to-market surface, brand system, marketing-site growth, SEO/GEO, the LinkedIn growth funnel
- **Build Buy Software** , legal-tech directory built as a traffic engine: the data model, the directory → Paralegent funnel, the Python competitive-intelligence scrape (1,610 companies, SHA256 evidence)
- **20-day May 2026 sprint** , two production tracks in parallel, full brand pivot on live main (zero downtime), 22+ SEO automation scripts, ~1,700 polite API calls across 4 platforms