

AI-native. Systems-first. Design-led. Research-backed.

We don't bolt AI onto apps. We build AI-native systems engineered for the workflow they sit inside, with design conviction and a standing research lab.

AI-native product studio. Software that ships.

A small, technical studio that builds AI-native products for startups and growth-stage companies, and runs an in-house research lab alongside client work, not instead of it. Every engagement is fixed-scope and fixed-price.

We don't bolt AI onto apps. We build AI-native systems engineered for the workflow they sit inside, with design conviction and a standing research lab.

WHO WE WORK WITH

Startups and growth-stage teams worldwide: operators in legal, real estate, healthcare, media, and industrial ops. We've already shipped for clients in the US, UAE, and India. You usually come to us when you need a system in production, not a wrapper.

WHAT MAKES US DIFFERENT

Most AI shops are either a volume agency (spec in, code out) or a thin wrapper over someone else's API. Research-first teams often stop before production. We sit in the middle: research discipline applied to systems that actually ship, with design conviction, at a fixed price and a fixed timeline.

INDUSTRIES

Legal, fintech, travel and hospitality, real estate, media and content, creator operations, healthcare, cybersecurity, quant finance, industrial operations, productivity tooling, and workflow automation.

Three kinds of work. One production standard.

PORTFOLIO

Client work.
Shipped under
constraint.

Client systems across industries. Each one is a scoped engagement for a named team, designed, built, and handed off in production.

PRODUCTS

Built in-house.
Sold in the market.

In-house products we sell and run. Brixloop owns the product, the customers, and the roadmap, this is not client work with a brand on it.

R&D

Research first.
Then it ships.

Applied research we run before it becomes a product. We publish the method, kill what doesn't hold, and pull the rest into production.

01

Tech Consultancy & Systems

Generative AI pipelines, automation, and platform engineering. Positioned as systems, not point solutions.

02

AI SaaS Products

Full-stack AI-native product builds. Zero to one, for startups and technical founders.

03

AI Research & Engineering

Independent research that feeds client systems. Models evaluated, documented, and published honestly.

04

Design

Brand identity and design systems for founders who need visual conviction, not just a logo.

Professionals we've worked with.

The people in the room already operate inside demanding organisations — banks, universities, funds, product companies, and public institutions. These marks are the professional context they bring into a build, not a claim that every institution hired us as an agency of record.

The logo for sarvam, featuring the word "sarvam" in a bold, lowercase, sans-serif font.The logo for kotak, featuring a stylized circular icon with a cross-like shape inside, followed by the word "kotak" in a bold, lowercase, sans-serif font.The logo for RenewBuy.com, featuring the text "RenewBuy.com" in a bold, sans-serif font, with "Smart Tech, Right Advice" in a smaller font below it.The logo for Newton School of Technology, featuring a stylized "N" icon followed by the text "Newton School of Technology" in a bold, sans-serif font.The Microsoft logo, consisting of four colored squares (red, green, blue, yellow) followed by the word "Microsoft" in a bold, sans-serif font.The logo for rishihood university, featuring a stylized "r" icon followed by the text "rishihood university" in a bold, sans-serif font.The logo for rtp global, featuring a circular icon with the letters "rtp" inside, followed by the word "global" in a bold, sans-serif font.The logo for ditch, featuring the word "ditch" in a stylized, gothic-like font.The logo for upwork, featuring the word "upwork" in a bold, lowercase, sans-serif font.The Google logo, featuring the word "Google" in its characteristic multi-colored font.

How work actually runs.

17+

AI SYSTEMS SHIPPED

Designed and built across client engagements and in-house products.

4

RESEARCH-LAB SYSTEMS

Legal NLP, medical imaging, network detection, and quant research. Published honestly.

Fixed-Scope

EVERY ENGAGEMENT

Fixed price and fixed timeline for the build. Optional monthly retainer after launch to keep the system healthy.

Global

WHERE WE WORK

We take on work worldwide. Previous clients include the US, UAE, and India.

PROCESS

I DISCOVERY CALL

A 30-minute call to understand your scope, goals, and constraints. We align on what success looks like before anyone writes code.

30 min · Goals, scope, and timeline alignment

II PROPOSAL + SPEC

Within three days you get a build plan, fixed timeline, and fixed price. All of that before work starts.

3 days · Written spec, milestones, and fixed quote

III BUILD SPRINT

A focused sprint matched to the quoted timeline, with weekly check-ins, a shared Linear board, and steady visibility on progress.

Scoped sprint · Weekly reviews · Shared project board

IV SHIP + HANDOFF

Deployed to production, fully documented, and handed off. Your team owns the codebase from day one. Most clients then add a monthly retainer for model updates and ongoing iteration.

Production deploy · Documentation · Optional monthly retainer

The team

Subham, Shivansh, and Divyanshi co-founded Brixloop. We build AI-native systems, in-house products, brand design, and a standing research lab for startups and growth-stage teams worldwide. Previous client work includes the US, UAE, and India.

Subham Mahapatra

CO-FOUNDER & CEO

Builds AI-native systems end to end: generative pipelines, LLM workflows, and the infra that keeps them running in production. Works directly with founders from scoping through ship.

Shivansh Gupta

CO-FOUNDER & CTO

Leads implementation on client builds: APIs, async job systems, front-end product surfaces, and the glue between models and real user workflows.

Divyanshi Sachan

CO-FOUNDER & TECH LEAD

Ships product surfaces and backend systems on client builds: data models, integrations, and the reliability layer that keeps AI features dependable.

Case studies

20 projects across 3 sections, generated from www.brixloop.com.

01 PORTFOLIO

13

Client work. Shipped under constraint.

- 01 AI Personalised Kids Video Platform · AI-first
- 02 AI Content Production Platform · AI-first
- 03 AI-Native Design Studio · AI-first
- 04 Industrial Intelligence · AI-first
- 05 Equity Research RAG Platform · AI-first
- 06 Finance Document RAG Knowledge Base (Vellum) · AI-first
- 07 Due Diligence Intelligence · AI-first
- 08 Multi-Agent Debate Platform · AI-first
- 09 Fleet Management & Reservation Platform · System-first
- 10 Creator Operations Platform · System-first
- 11 OS-Structured File Explorer · System-first
- 12 Visual Workflow Automation Platform · System-first
- 13 Architecture Studio Portfolio · Design-first

02 PRODUCTS

3

Built in-house. Sold in the market.

- 14 LexVault
- 15 Arth Saathi
- 16 Deskzy

03 R&D

4

Research first. Then it ships.

- 17 Contract Risk Classification Model
- 18 Network Intrusion Detection Model
- 19 Counterfeit Medicine Vision Model
- 20 Regime-Adaptive Equity Model

01

Portfolio

Client systems across industries. Each one is a scoped engagement for a named team, designed, built, and handed off in production.

AI Personalised Kids Video Platform

Personalised animated storytelling for children. Generative video pipelines, parent-safe workflows, and infrastructure tuned for viral traffic and repeat renders.

AI VIDEO

GENERATIVE MEDIA

CONSUMER SCALE

FAMILY TECH

2026 / Global · Live · www.tinytalesvideos.com/

PROBLEM

The client needed a content system that could generate personalised animated story videos for children, with the child's own face and name woven into the narrative, at scale, driven entirely by parent prompts. Off-the-shelf tools either lacked identity preservation or couldn't survive viral traffic spikes without collapsing margin.

WHAT WE BUILT

Tiny Tales targets families who want bespoke animated stories without a production studio budget. The product couples prompt-style personalization with a constrained creative pipeline so outputs stay age-appropriate and on-brand.

Technically, the hardest problems are render throughput and cost control when campaigns spike. Viral curves expose naive queueing fast. We optimized for parallelized generation, caching of reusable assets, and graceful degradation when providers throttle.

On the experience side, parents need clarity on what will be generated, how long it takes, and how to iterate safely. We built flows that separate preview, confirm, and paid render so trust stays high and chargebacks stay low.

OUTCOME

- 01 Prompt-to-finished-video pipeline shipped end-to-end in under 2 weeks
- 02 Async generation queue handles traffic spikes without manual intervention
- 03 Preview → confirm → paid render flow cut chargeback risk to near zero

STACK

AI / ML

Stable Diffusion · IP-Adapter (face insertion) · GPT-4 · ElevenLabs

VIDEO

FFmpeg · Scene assembly pipeline

FRONTEND

Next.js · Tailwind CSS

INFRASTRUCTURE

AWS · Async job queue · Provider fallbacks

[HTTPS://WWW.BRIXLOOP.COM/PORTFOLIO/AI-PERSONALISED-KIDS-VIDEO-PLATFORM](https://www.brixloop.com/portfolio/ai-personalised-kids-video-platform)

AI Content Production Platform

Content and media production platform, bulk Excel-driven workflows for AI-generated images, video, and audio, with pause/resume job control.

CONTENT PLATFORM

AI-NATIVE

MEDIA GENERATION

CLIENT BUILD

2026 · Live

PROBLEM

Teams generating content at volume (dozens of product images, video clips, or voiceovers per batch) typically stitch together multiple AI tools and vendors by hand, with no shared job history, no way to pause and resume a bulk run, and no protection against a provider rate-limiting or the server falling over under concurrent load.

WHAT WE BUILT

A TypeScript monorepo where an Excel upload creates project rows, and generation (image, video, and audio via OpenAI, ElevenLabs, and Fal.ai) is a separate, explicit, pausable and resumable step.

The concurrency problem is real: firing 40 generation jobs at once blows through provider rate limits and can OOM the server. It was solved without adding infrastructure. A custom semaphore bounds how many jobs run per type and per provider, configurable via env vars, with zero Redis or queue service.

Jobs persist as Postgres rows with a JSON payload. Execution can run inline in the API or in an optional dedicated worker. Generated media lands on S3 behind CloudFront.

OUTCOME

- 01 Custom async-semaphore concurrency control replaces Redis/BullMQ/SQS, bulk jobs stay bounded with zero added infrastructure
- 02 Two-layer protection (worker-level job caps + provider-level API rate caps) tuned live via environment variables, no redeploy needed
- 03 Full bulk pipeline from spreadsheet to generated media: Excel upload → project → generation → pause/resume, backed by durable Postgres job state

STACK

AI / ML

OpenAI · ElevenLabs · Fal.ai

BACKEND

Node.js · Express · Drizzle ORM · PostgreSQL · Zod

FRONTEND

React 18 · Vite · TanStack Query · Zustand · Tailwind CSS

INFRASTRUCTURE

AWS S3 · CloudFront · EC2 · ECR · Docker

[HTTPS://WWW.BRIXLOOP.COM/PORTFOLIO/AI-CONTENT-PRODUCTION-PLATFORM](https://www.brixloop.com/portfolio/ai-content-production-platform)

AI-Native Design Studio

AI-native design studio for interior designers: client project management, redesign generation, and team communication in one platform.

REAL ESTATE

DESIGN

AI-NATIVE

PROTOTYPE

2025 / Design · Prototype · paused

PROBLEM

Interior designers run client work across disconnected tools: one place for communication, another for mood boards, another for timelines. Redesign requests happen over email or WhatsApp with no structured history tied back to the room or project.

WHAT WE BUILT

A single workspace where a designer manages client projects, generates and iterates on room redesigns with an AI layer, and keeps project communication and revision history attached to the project itself.

The prototype covers the core two-sided flow (designers and their clients) before a production launch.

OUTCOME

- 01 Unified client-project workspace covering communication, design iteration, and revision history in one platform
- 02 AI-assisted redesign generation built into the core workflow, not bolted on
- 03 Paused pre-launch, architecture and core flows built, not yet in production use

[HTTPS://WWW.BRIXLOOP.COM/PORTFOLIO/AI-NATIVE-DESIGN-STUDIO](https://www.brixloop.com/portfolio/ai-native-design-studio)

Industrial Intelligence

Institutional memory for heavy industry: ingest maintenance documents, build a knowledge graph of assets, people, incidents, and parts, then answer field questions with citations.

INDUSTRIAL

KNOWLEDGE GRAPH

AGENTS

NDA

2025 / Confidential · Prototype · Confidential

PROBLEM

When a senior plant engineer retires or leaves, the institutional knowledge of which machine has which failure history, which procedure applies, and who is the only person who can fix a specific asset leaves with them. That knowledge usually lives in scattered documents and one person's memory, not in a searchable system.

WHAT WE BUILT

A five-agent pipeline (ingest, entity-link, jargon-normalize, graph-traversal, risk-scoring) turns uploaded maintenance documents into a live knowledge graph of Assets, People, Incidents, Documents, Procedures, and Parts.

On top of the graph sits a Risk Radar that flags machines that depend on a single expert, and an Expert Copilot that answers plant-floor questions with citations back to the source document.

OUTCOME

- 01 Answers "if this engineer retires tomorrow, which machines lose their only expert?" deterministically, not as a best-guess LLM answer
- 02 Full document-to-graph pipeline: PDF/DOCX/XLSX ingestion → entity extraction → graph linking, in one flow
- 03 Built and demoed end-to-end in a hackathon timebox on a five-agent tool-use loop

STACK

AI / ML

Anthropic Claude · Agentic tool-use loop

BACKEND

Node.js · Express · TypeScript · MongoDB

FRONTEND

Next.js 15 · React Flow

[HTTPS://WWW.BRIXLOOP.COM/PORTFOLIO/INDUSTRIAL-INTELLIGENCE](https://www.brixloop.com/portfolio/industrial-intelligence)

Equity Research RAG Platform

Enterprise RAG for stock and market research, cited, time-aware answers over SEC filings, news, earnings transcripts, and a user's own notes.

FINTECH

RAG

RESEARCH

NDA

2026 / Confidential · Live · Confidential

PROBLEM

Equity research means reading through dense SEC filings, news, and transcripts, then citing the right source when you make a claim. Off-the-shelf chat tools answer confidently but don't reliably ground claims in the actual filing text, and don't handle a user's own private notes and documents alongside public filings in the same place.

WHAT WE BUILT

A dual-ingestion RAG pipeline: one path automatically pulls SEC 10-K/10-Q filings, news, and transcripts per ticker; a second lets a user paste notes, a URL, or upload a PDF. Both land in the same per-ticker vector store so a query can draw on public filings and private research together.

Every answer is cited back to source. PDF uploads are processed for text and then discarded, not retained, keeping the ingestion path lightweight and auditable.

OUTCOME

- 01 Dual ingest modes (automated filings + manual/BYO documents) into one unified per-ticker knowledge base
- 02 Built-in LLM-as-judge evaluation harness to score answer quality against a held-out eval set
- 03 Query rate-limited and session-authenticated, with a separate admin-key-gated ingestion and eval surface, production access control from the start

STACK

AI / ML	Anthropic Claude · Voyage embeddings · Cohere rerank
BACKEND	Go · River workers
FRONTEND	Next.js
DATA	PostgreSQL · Redis · Qdrant

[HTTPS://WWW.BRIXLOOP.COM/PORTFOLIO/EQUITY-RESEARCH-RAG-PLATFORM](https://www.brixloop.com/portfolio/equity-research-rag-platform)

Finance Document RAG Knowledge Base (Vellum)

End-to-end RAG knowledge base for equity research, ask questions and get cited answers from filings, news, transcripts, PDFs, notes, and URLs.

RAG

FINTECH

VECTOR DATABASE

NLP

2026 / Jul · Live

PROBLEM

Equity research means jumping across SEC filings, news, transcripts, and private notes, then grounding every claim in a source. Generic chat tools answer confidently without reliable citations, and they don't mix public filings with a private document vault in one retrieval path.

WHAT WE BUILT

Built an end-to-end RAG knowledge base for equity research. Users ask questions and get cited answers from filings, news, transcripts, PDFs, notes, and URLs.

Document ingest and chunking feed Postgres metadata storage and Qdrant vectors. Retrieval is hybrid (dense + keyword + RRF) with Cohere rerank and Voyage embeddings. Claude handles routing and grounded answers.

Also shipped a private document vault and a Next.js chat UI, including a Neural Trace graph that shows how a query connects through sources, facts, metrics, and synthesis.

OUTCOME

- 01 End-to-end document retrieval and drafting system for equity research Q&A
- 02 Hybrid retrieval (dense + keyword + RRF) with Cohere rerank and Voyage embeddings
- 03 Reusable patterns for confidential document Q&A, private vault plus grounded Claude answers

STACK

AI / ML

Anthropic Claude · Voyage embeddings · Cohere rerank

RETRIEVAL

Hybrid dense + keyword · RRF fusion · Qdrant

BACKEND

Python · API Development · Postgres

FRONTEND

Next.js · Neural Trace graph UI

[HTTPS://WWW.BRIXLOOP.COM/PORTFOLIO/FINANCE-DOCUMENT-RAG-KNOWLEDGE-BASE-VELLUM](https://www.brixloop.com/portfolio/finance-document-rag-knowledge-base-vellum)

Due Diligence Intelligence

Decision-support for venture firms: upload a startup cap table, screen every owner (including UBOs behind shell layers) against OFAC/EU sanctions lists, and get a plain-English risk report with an ownership graph and AI-generated analyst notes.

FINTECH

COMPLIANCE

AI-NATIVE

NDA

2026 / Confidential · Live · Confidential

PROBLEM

Venture firms doing sanctions/compliance screening on cap tables have to manually trace ownership through shell companies and holding structures, then cross-reference each name against sanctions lists by hand. It's slow, easy to miss a layered UBO, and produces no structured audit trail.

WHAT WE BUILT

The pipeline parses a cap table CSV, builds a directed ownership graph, traverses it to surface ultimate beneficial owners, screens every node against OFAC SDN and EU Consolidated lists, and classifies each as clear, review, or flagged.

Claude writes a plain-English narrative (who, why, confidence, next step) for anything that isn't clean. The tool never renders a verdict; it narrows the review set for a human compliance officer.

OUTCOME

- 01 Full ownership-chain traversal (BFS) surfaces beneficial owners hidden behind shell-company layers, not just direct owners
- 02 Three-tier confidence classification (clear <0.80 / review 0.80–0.95 / flagged ≥0.95) keeps false positives out of the officer's queue
- 03 Production split across Vercel (frontend) and AWS EC2 (backend + sanctions matching engine) for independent scaling

STACK

AI / ML

Anthropic Claude

BACKEND

Node.js · Express · TypeScript · Watchman

FRONTEND

Next.js 15 · React 19 · React Flow

DATA

MongoDB Atlas · Upstash Redis · Cloudflare R2

INFRASTRUCTURE

AWS EC2 · Vercel

[HTTPS://WWW.BRIXLOOP.COM/PORTFOLIO/DUE-DILIGENCE-INTELLIGENCE](https://www.brixloop.com/portfolio/due-diligence-intelligence)

Multi-Agent Debate Platform

Self-hostable open-source platform where AI personas with distinct worldviews debate current topics in structured, fact-checked rounds, bring your own API keys, no central server.

OPEN SOURCE

MULTI-AGENT

LANGGRAPH

BYOK

2026 / OSS · Live · Apache 2.0

PROBLEM

Getting genuinely different AI perspectives on a topic today means running separate chats and stitching the output yourself. There's no structured format for it, and most multi-agent tools assume a central SaaS with your API usage flowing through someone else's server.

WHAT WE BUILT

A LangGraph state machine drives every debate through framing, opening, cross-examination, rebuttals, closing, and synthesis, invoking a Moderator, Debater, Fact-Checker, and Synthesizer at the right point in each round.

Everything runs on the user's own API keys (Anthropic, OpenAI, Google, Groq, or local Ollama), encrypted at rest, with no telemetry back to a central service. Adding a new debate persona is a single markdown file, no code required.

OUTCOME

- 01 Six-phase structured debate with four distinct agent roles, not a freeform chat loop
- 02 Fully self-hostable and BYOK: keys AES-256-GCM encrypted at rest, never leave the instance in plaintext
- 03 Debate state persists to Postgres at every phase, resumable and inspectable, not a one-shot generation

STACK

FRONTEND	Next.js 15 · Tailwind CSS v4 · shadcn/ui
API	Hono · LangGraph.js · Vercel AI SDK
DATA	PostgreSQL 16 · pgvector · Drizzle ORM · Inngest
AUTH	Better Auth · Email magic links · AES-256-GCM key encryption

[HTTPS://WWW.BRIXLOOP.COM/PORTFOLIO/MULTI-AGENT-DEBATE-PLATFORM](https://www.brixloop.com/portfolio/multi-agent-debate-platform)

Fleet Management & Reservation Platform

Luxury and exotic car rental booking for Dubai, search by pickup, dates, and car type, then book directly or via WhatsApp concierge.

TRAVEL

BOOKING

UAE

E-COMMERCE

2026 / UAE • Live • e-srent.vercel.app/

PROBLEM

Dubai's luxury car rental market runs heavily on WhatsApp negotiation and opaque pricing, high security deposits, hidden fees, no real-time visibility into what's actually available. Renters want speed and transparency; operators want a channel that doesn't depend entirely on manual back-and-forth for every booking.

WHAT WE BUILT

The booking site is built around one search flow (pickup location, pickup/return dates, car type, and rental duration (daily / +3 days / weekly)) surfacing the fleet by brand and category.

Positioned against the industry's pain points: minimal-to-zero security deposit, no hidden fees, delivery anywhere in Dubai.

A Quick Book WhatsApp integration sits alongside the standard flow for concierge-style requests on higher-end bookings.

OUTCOME

- 01 Single-flow fleet search (location, dates, car type, duration) replacing the manual WhatsApp-first booking process
- 02 Direct WhatsApp Quick Book integrated into the booking flow itself, not a separate contact channel
- 03 Marketing and product built around a zero-hidden-fee, low-deposit position against incumbent Dubai rental operators

STACK

FRONTEND

Next.js • Tailwind CSS

INFRASTRUCTURE

Vercel

[HTTPS://WWW.BRIXLOOP.COM/PORTFOLIO/FLEET-MANAGEMENT-RESERVATION-PLATFORM](https://www.brixloop.com/portfolio/fleet-management-reservation-platform)

Creator Operations Platform

Internal creator ops product for DITCH LA (submission, approval, and payouts at 100+ creator scale) behind the public brand at ditch.la.

CREATOR OPS

APPAREL

INTERNAL PRODUCT

CLIENT BUILD

2026 / US · Live · ditch.la

PROBLEM

DITCH LA runs UGC and creator content at volume. Onboarding, raw video submission, approval with feedback, earnings, and payment status were scattered across email and DMs. The hard constraint: scale past 100 creators without buying a paid Airtable seat for every creator, or the tool itself becomes the cost problem it was meant to solve.

WHAT WE BUILT

ditch.la is the public storefront. The product we built sits behind it: an internal creator system for the team that actually runs UGC, not another customer-facing shop.

An Airtable backend handles creator records, submissions, approval workflow, and payout tracking, kept strictly admin-only. Creators log into a separate lightweight portal and see only their own submissions, status, feedback, and earnings, with zero Airtable seat cost per creator.

Automated notifications fire on submission and on approval, rejection, or revision request. Approved raw video auto-archives to Dropbox, organized by creator. The approval/payout flow reused the pattern from Arth Saathi rather than designing ops from zero.

OUTCOME

- 01 Full MVP (creator onboarding, submissions, bulk approval/payout workflows, admin dashboard, notifications) delivered on a 7-day build plus a 3–4 day revision cycle
- 02 Architected to scale to 100+ creators with zero additional Airtable seats, backend stays limited to the internal admin team
- 03 Reused the approval/payout architecture from Arth Saathi rather than designing ops from zero

STACK

BACKEND	Airtable · Forms · Interfaces · Automations
CREATOR PORTAL	Softr / Next.js
STORAGE	Dropbox API

[HTTPS://WWW.BRIXLOOP.COM/PORTFOLIO/CREATOR-OPERATIONS-PLATFORM](https://www.brixloop.com/portfolio/creator-operations-platform)

OS-Structured File Explorer

Desktop file explorer and disk-usage visualizer built as a live demonstration of core OS concepts. Every architectural decision maps to a real OS mechanism.

SYSTEMS

DESKTOP

ELECTRON

OS

2025 · Complete

PROBLEM

OS concepts (process isolation, IPC, journaling, scheduling, protection rings) are usually learned from a textbook, not from a real running application. Most student systems projects demonstrate one concept in isolation; few build a complete, usable desktop app where every subsystem maps back to a specific OS mechanism.

WHAT WE BUILT

An Electron desktop app structured like an OS: a main process (kernel-equivalent, owns all file-system access), a sandboxed renderer (user-space, no direct Node/fs access), and a worker thread for disk scanning (kernel thread model), communicating only through a whitelisted IPC layer that mirrors a syscall interface.

File deletion is a journaled two-phase soft-delete, like filesystem journaling. Directory watching uses native kernel event notification (kqueue/inotify). Every operation checks Unix-style permission bits before executing. No AI layer, this is systems engineering.

OUTCOME

- 01
- Full multi-process architecture (main/renderer/worker) with zero shared memory, data crosses process boundaries only via serialized IPC
- 02
- Journaled trash: every delete is recorded to a manifest before the file moves, recoverable on a mid-operation crash
- 03
- Worker-thread disk scanning keeps the UI responsive during large directory scans, verified against main-thread scanning freezing the UI

STACK

RUNTIME	Electron 33 · Node.js
FRONTEND	React 18 · TypeScript · Tailwind CSS · Zustand
VISUALIZATION	@nivo/sunburst
BUILD	Vite · electron-builder

<https://www.brixloop.com/portfolio/os-structured-file-explorer>

Visual Workflow Automation Platform

Full-stack workflow automation, connect Google Drive, Slack, Discord, and Notion through a visual editor and run automations on real-time triggers.

AUTOMATION

WORKFLOWS

OAUTH

STRIPE

2025 · Complete · automation-51k3.vercel.app/

PROBLEM

Off-the-shelf automation tools (Zapier, Make) work, but building one from scratch means solving OAuth flows, webhook-based real-time triggers, and billing/credit metering correctly. That's the plumbing that's usually the hard part, not the drag-and-drop canvas.

WHAT WE BUILT

A ReactFlow visual editor where a user builds a workflow as a trigger node (for example, a new file in Google Drive) connected to one or more action nodes (post to Discord, create a Notion page, send a Slack message). OAuth2 connections are stored securely per user.

When a workflow is published, the app registers a webhook with the source service so actions fire in real time on the actual event. Usage is metered through a Stripe-backed credit system across free, pro, and unlimited tiers.

OUTCOME

- 01 Real-time execution via native webhooks, not polling, for every connected trigger source
- 02 Four native integrations (Google Drive, Discord, Notion, Slack) each with its own OAuth2 flow and action set
- 03 Tiered subscription and credit-metering built in from the start, not bolted on after launch

STACK

BACKEND

Next.js 14 · TypeScript · PostgreSQL · Prisma

FRONTEND

React · ReactFlow · Tailwind CSS · shadcn/ui · Framer Motion

AUTH / PAYMENTS

Clerk · Stripe

STATE / UPLOADS

Zustand · Uploadcare

[HTTPS://WWW.BRIXLOOP.COM/PORTFOLIO/VISUAL-WORKFLOW-AUTOMATION-PLATFORM](https://www.brixloop.com/portfolio/visual-workflow-automation-platform)

Architecture Studio Portfolio

Architecture, interior, and landscape practice with a cinematic portfolio site emphasizing curated materials, passive-house thinking, and premium residential positioning.

CREATIVE PORTFOLIO

ARCHITECTURE

INTERIORS

EDITORIAL UX

2025 / Design · Live · www.inarchdeign.com/

PROBLEM

A multi-disciplinary architecture practice offering residential, interior, and landscape services under one roof lacked a digital presence that reflected their premium positioning. Their previous site felt generic. It didn't communicate the holistic, sustainability-led vision that affluent clients were actually paying for.

WHAT WE BUILT

INARCHDEZIGN's digital presence mirrors how studios win luxury residential work: restrained typography, generous whitespace, and editorial pacing instead of noisy grids.

The narrative blends architecture, interior, and landscape into one holistic pitch. Affluent clients expect one coordinated vision rather than siloed trades.

Passive-house and sustainability cues elevate differentiation in markets where energy performance is both ethics and asset value.

OUTCOME

- 01 Cinematic portfolio experience positioned above commodity architecture agencies
- 02 Blended architecture, interior, and landscape narrative in one coherent pitch
- 03 Sustainability and passive-house framing elevated perceived value for premium clients

STACK

FRONTEND	Next.js · Framer Motion · Tailwind CSS
CMS	Headless CMS for portfolio management
INFRASTRUCTURE	Vercel · Image optimisation pipeline

[HTTPS://WWW.BRIXLOOP.COM/PORTFOLIO/ARCHITECTURE-STUDIO-PORTFOLIO](https://www.brixloop.com/portfolio/architecture-studio-portfolio)

02

Products

In-house products we sell and run. Brixloop owns the product, the customers, and the roadmap, this is not client work with a brand on it.

LexVault

AI-native contract lifecycle platform, draft, review, redline, sign, and collaborate, with every AI answer grounded to an exact, verified line citation.

LEGALTECH

AI-NATIVE

CLM

RAG

2026 / India · Live · invite-only · www.lexvault.in/

PROBLEM

Legal teams manage contracts across email threads, shared drives, and static templates with no structural verification layer. Existing AI contract tools either answer with confident-sounding guesses that aren't traceable to the source document, or bolt AI onto a workflow that doesn't talk to the firm's actual infrastructure, DMS, calendar, e-signature, org hierarchy. Teams need answers that point to a specific line, not a plausible one.

WHAT WE BUILT

LexVault is a full contract lifecycle platform: draft, review, redline, sign, collaborate, wrapped around an AI layer where every claim is server-side validated against exact line ranges in the source document before it reaches the user.

On top of that: DAG-based approval workflows with human-in-the-loop gates, DMS integrations (iManage, NetDocuments), calendar sync, and multi-tenant org management with seat-based billing.

This is an in-house product we sell and operate, not a one-off client build.

OUTCOME

- 01 Every AI citation is validated server-side against exact document line ranges before it's shown, failed citations return a fallback, never a fabricated reference
- 02 20+ database tables span contracts, workflows, vaults, and a 365-day compliance audit trail
- 03 Runs as a 4-app system (product app, marketing site, admin panel, background worker) in production on AWS

STACK

AI / ML

Anthropic Claude · pgvector · Ollama (fallback) · Hybrid RAG

BACKEND

Node.js · Express · Drizzle ORM · PostgreSQL (Neon)

FRONTEND

Next.js 14 · React 18 · TypeScript · Tailwind CSS · BlockNote · Tiptap · Liveblocks

INFRASTRUCTURE

AWS EC2 · Docker · nginx · Cloudflare R2 · Cloudflare Workers

[HTTPS://WWW.BRIXLOOP.COM/PRODUCTS/LEXVAULT](https://www.brixloop.com/products/lexvault)

Arth Saathi

Staff and money OS for Indian small businesses, voice-marked attendance, automated payroll, digital khata ledger, and GST billing in one app, in 14 Indian languages.

FINTECH

SMB OPS

AI-NATIVE

INDIA

2025 / India · Live · arthsaathi.co.in/

PROBLEM

Indian SMB owners (kirana shops, salons, small factories) run attendance on biometrics or memory, payroll on paper or a cousin's spreadsheet, credit tracking (khata) in a physical notebook, and GST invoicing as a separate manual chore. None of it talks to each other, and none of it works in the owner's own language.

WHAT WE BUILT

Arth Saathi starts from the two things an owner does daily (mark the team in, pay them on time) and layers khata, GST billing, and inventory on the same data.

An AI assistant lets an owner mark attendance or approve leave with a spoken sentence ("sabko present kar do") in Hindi, Tamil, Bengali, or eleven other Indian languages, no menus. Salary slips and approvals land on WhatsApp, where the team already reads.

This is an in-house product we take to market.

OUTCOME

- 01 14 Indian languages live across the owner app, employee portal, AI assistant, and WhatsApp notifications, not just English with translated labels
- 02 Monthly payroll close reduced to a single approval tap: attendance + leaves + advances roll up into ready-to-send slips automatically
- 03 One-tap GST invoice generation (HSN codes, CGST/SGST split, IRN) delivered directly on WhatsApp
- 04 Free for the first 10 employees, with tiered pricing (₹500–₹1,000/month) beyond that

STACK

FRONTEND

React · Tailwind CSS · Mobile-responsive design

BACKEND

Node.js · PostgreSQL · Payroll calculation engine

PRODUCT

Voice + GPS attendance · WhatsApp delivery · 14-language AI assistant

[HTTPS://WWW.BRIXLOOP.COM/PRODUCTS/ARTH-SAATHI](https://www.brixloop.com/products/arth-saathi)

Deskzy

Free, no-signup link-sharing and file-tools platform. PDF and image tools run entirely in the browser, files never touch a server.

UTILITY

PRIVACY

DEVELOPER TOOLS

EDGE

2026 / Global · Live · deskzy.xyz

PROBLEM

Most "free" online PDF and image tools require an upload to someone else's server, an account, or a paywall after one use. People sharing a quick link also shouldn't need to sign up just to get a shortened URL with basic analytics.

WHAT WE BUILT

Deskzy is built around a privacy-first constraint: PDF and image operations run client-side in the browser, so files never leave the user's device.

Paired with a link-sharing and shortening tool with Cloudflare-edge analytics for paying users, plus smaller utilities (QR codes, UTM builders, WhatsApp links, bio pages, JSON/text formatters) around the same no-signup philosophy.

This is an in-house product we operate and sell on a freemium model.

OUTCOME

- 01 Full PDF/image suite (compress, merge, split, resize, convert) runs entirely in-browser, files never leave the device
- 02 No-signup access to every free tool, with a \$399/month Pro tier for link analytics
- 03 Global edge presence, live request routing across 10 Cloudflare colos (BOM, IAD, SIN, FRA, GRU, LHR, SJC, NRT, AMS, SYD)

STACK

FRONTEND

Next.js · Client-side PDF/image processing

INFRASTRUCTURE

Cloudflare · Edge analytics · Vercel

[HTTPS://WWW.BRIXLOOP.COM/PRODUCTS/DESKZY](https://www.brixloop.com/products/deskzy)

03

R&D

Applied research we run before it becomes a product. We publish the method, kill what doesn't hold, and pull the rest into production.

Contract Risk Classification Model

Research system for automated contract risk classification: 41 CUAD clause labels plus a contract-level risk rating (Low/Medium/High) with evidence-backed, citation-traceable outputs.

LEGAL AI

RESEARCH

NLP

CUAD

Research / Phase 3 · Complete

PROBLEM

Legal risk assessment on long contracts is either fully manual (slow, inconsistent across reviewers) or handled by black-box AI tools that give a risk score with no traceable reasoning. Neither scales, and neither survives an auditor asking why.

METHOD

A three-phase research pipeline, each phase built to be falsifiable, not just functional. Phase 1 established a classical ML baseline to know what's recoverable from lexical signal alone. Phase 2 fine-tuned Legal-BERT with LoRA adapters for clause detection across long, sometimes-scanned contracts.

Phase 3 converts clause-level signal into a contract-level risk decision. A Bayesian Network collapsed under class imbalance; we diagnosed why and replaced it with a calibrated Random Forest reasoner instead of patching the failure.

FINDINGS

- 01 Macro-F1 progression: 0.719 (classical baseline) $\rightarrow$ 0.759 (Legal-BERT) $\rightarrow$ 0.866 (final hybrid Phase 3 reasoner)
- 02 Failed Bayesian Network (Macro-F1 0.159) documented and replaced, kept as ablation evidence
- 03 Oracle ceiling test (RF fed ground-truth clauses) hits 0.903, remaining headroom is in upstream clause extraction, not the reasoning layer

[HTTPS://WWW.BRIXLOOP.COM/RESEARCH/CONTRACT-RISK-CLASSIFICATION-MODEL](https://www.brixloop.com/research/contract-risk-classification-model)

Contract Risk Classification Model

ARCHITECTURE

Phase 1, Classical ML baseline to measure what lexical signal alone can recover

Phase 2, Legal-BERT with LoRA adapters for 41-label CUAD clause detection on long and scanned contracts

Phase 3, Clause-frequency features into a calibrated Random Forest reasoner after a Bayesian Network collapse under class imbalance

Delivery, Upload UI returns risk class with citation-traceable evidence

METRICS

FINAL MACRO-F1	0.866 · Hybrid Phase 3 reasoner
LEGAL-BERT MACRO-F1	0.759 · Phase 2 clause detection
CLASSICAL BASELINE	0.719 · Lexical signal only
ORACLE CEILING	0.903 · RF on ground-truth clauses
FAILED BN ABLATION	0.159 · Documented, not hidden

STACK

AI / ML	PyTorch · HuggingFace Transformers · LoRA · scikit-learn
BACKEND	FastAPI
FRONTEND	React · Vite
DATA	CUAD v1 · pdfplumber · EasyOCR

[HTTPS://WWW.BRIXLOOP.COM/RESEARCH/CONTRACT-RISK-CLASSIFICATION-MODEL](https://www.brixloop.com/research/contract-risk-classification-model)

Network Intrusion Detection Model

Machine learning intrusion detection system, a full web application that classifies network traffic and flags attacks in real time.

CYBERSECURITY

ML

IDS

RESEARCH

Research • Complete

PROBLEM

Detecting network intrusions and DoS attacks from raw traffic means classifying patterns across a large, imbalanced feature set. Most academic IDS projects stop at a notebook with an accuracy number; few wrap the model in something a security analyst could actually operate.

METHOD

A trained classifier on CSE-CIC-IDS2018-format traffic is served through a Flask API with two prediction paths: batch file upload (CSV/Parquet, with preprocessing for column cleaning, type coercion, and missing-value handling) and manual JSON for ad hoc checks.

A React dashboard shows attack distribution, severity breakdown, and live model status. An operable surface, not a notebook demo.

FINDINGS

- 01 Two independent prediction paths (file upload + JSON) aligned to the model's trained feature order via saved metadata
- 02 Preprocessing handles messy input: column cleaning, numeric coercion, inf/NaN handling, memory downcasting
- 03 Visual analytics dashboard (attack distribution, threat-level summary) on top of the classifier, not just a JSON response

[HTTPS://WWW.BRIXLOOP.COM/RESEARCH/NETWORK-INTRUSION-DETECTION-MODEL](https://www.brixloop.com/research/network-intrusion-detection-model)

Network Intrusion Detection Model

ARCHITECTURE

Jupyter training pipeline on CSE-CIC-IDS2018-format traffic features
joblib-serialized scikit-learn classifier with saved feature-order metadata
Flask REST API with preprocessing for messy real-world inputs
React dashboard for attack distribution, severity, and live model status

METRICS

INFERENCE PATHS	2 · Batch upload + ad hoc JSON
-----------------	--------------------------------

DATASET FORMAT	CSE-CIC · IDS2018-compatible features
----------------	---------------------------------------

DELIVERY	API + UI · Flask classifier + React dashboard
----------	---

STACK

AI / ML	scikit-learn · pandas · NumPy · imbalanced-learn
---------	--

BACKEND	Flask · Flask-CORS
---------	--------------------

FRONTEND	React 18 · Recharts · Axios
----------	-----------------------------

[HTTPS://WWW.BRIXLOOP.COM/RESEARCH/NETWORK-INTRUSION-DETECTION-MODEL](https://www.brixloop.com/research/network-intrusion-detection-model)

Counterfeit Medicine Vision Model

Deep learning system for detecting counterfeit medicine from a photo, transfer-learning image classifier with a full deployment stack around it.

HEALTHCARE

COMPUTER VISION

RESEARCH

ML

Research • Complete

PROBLEM

Verifying medicine authenticity today requires expert inspection or lab testing, neither scalable to a pharmacy counter, a supply-chain checkpoint, or a consumer checking a suspicious package.

METHOD

A ResNet-18 model, transfer-learned to classify medicine images as authentic or counterfeit with a confidence score, deployed as a proper microservices stack rather than a notebook demo.

FastAPI handles inference. Express handles JWT auth and role-based access. React handles real-time upload and verification.

FINDINGS

- 01 Full three-service microservices architecture (React □ Express □ FastAPI) shipped end-to-end, not just a trained model
- 02 JWT-authenticated API with role-based access control around the inference endpoint
- 03 Trained on a Roboflow counterfeit-medicine dataset (1,367 training images) with a documented train/val/test split

[HTTPS://WWW.BRIXLOOP.COM/RESEARCH/COUNTERFEIT-MEDICINE-VISION-MODEL](https://www.brixloop.com/research/counterfeit-medicine-vision-model)

Counterfeit Medicine Vision Model

ARCHITECTURE

ResNet-18 transfer learning for authentic vs counterfeit classification with confidence

FastAPI owns inference

Express owns JWT auth and role-based access around the inference endpoint

React handles real-time upload and verification

Docker Compose packages the full stack

METRICS

TRAINING IMAGES	1,367 · Roboflow Universe, CC BY 4.0
-----------------	--------------------------------------

SERVICES	3 · React · Express · FastAPI
----------	-------------------------------

TASK	Binary · Authentic vs counterfeit
------	-----------------------------------

STACK

AI / ML	PyTorch · ResNet-18 transfer learning
---------	---------------------------------------

BACKEND	FastAPI · Express · MongoDB · JWT
---------	-----------------------------------

FRONTEND	React · Vite
----------	--------------

INFRA	Docker · Docker Compose
-------	-------------------------

[HTTPS://WWW.BRIXLOOP.COM/RESEARCH/COUNTERFEIT-MEDICINE-VISION-MODEL](https://www.brixloop.com/research/counterfeit-medicine-vision-model)

Regime-Adaptive Equity Model

Regime-adaptive equity strategy research on Indian markets (NIFTY 200), as much about disciplined negative-result reporting as the final strategy.

QUANT

FINANCE

RESEARCH

HMM

Research • Complete

PROBLEM

Most retail-adjacent quant research falls for its own backtest: a dashboard shows gross returns, looks great, and the friction, regime risk, and realistic execution constraints get added later, if at all. Building a strategy that survives contact with real costs and real regime shifts requires being willing to kill your own architecture when the evidence says so.

METHOD

Phase 1 built an XGBoost baseline and found that daily-return prediction is mostly noise, the real signal is monthly. Phase 2 built a custom multimodal transformer (regime cross-attention, mixture-of-experts, ranking loss) to exploit that signal directly; a 60-line diagnostic showed it performing worse than plain 21-day momentum with zero ML.

Phase 3 pivoted: production became a 21-day momentum ranking with 3-state HMM regime-based position sizing, sector caps, per-stock stops, and a portfolio drawdown killswitch, benchmarked with 0.22%/rebalance friction. Not a high-frequency system, rebalances every 21 trading days.

FINDINGS

- 01 Production strategy backtest (2024–2026, net of friction): 13.5% CAGR vs 7.7% NIFTY buy-and-hold, Sharpe 0.83 vs 0.65, 64% win rate vs 52%
- 02 Diagnostic script exposed the custom transformer underperforming a zero-ML momentum sort on every metric, published, then rebuilt around the finding
- 03 Regime-based position sizing cut historical 2008-style drawdown by 9.4 percentage points in ablation testing, at the documented cost of capping upside in bull markets

[HTTPS://WWW.BRIXLOOP.COM/RESEARCH/REGIME-ADAPTIVE-EQUITY-MODEL](https://www.brixloop.com/research/regime-adaptive-equity-model)

Regime-Adaptive Equity Model

ARCHITECTURE

Phase 1, XGBoost baseline; daily-return prediction mostly noise, monthly signal holds

Phase 2, Custom multimodal transformer (RAMT); diagnostic showed it underperformed plain 21-day momentum

Phase 3, Production: 21-day momentum ranking + 3-state HMM regime sizing (100%/50%/20%)

Risk controls, sector caps, per-stock stops, portfolio drawdown killswitch, 0.22%/rebalance friction

Validation, 4-window historical ablation (2008, 2010, 2013, 2024)

METRICS

STRATEGY CAGR	13.5% · Net of friction, 2024-2026
BUY-AND-HOLD CAGR	7.7% · NIFTY baseline
SHARPE	0.83 · vs 0.65 buy-and-hold
WIN RATE	64% · vs 52% baseline
2008 DRAWDOWN CUT	9.4 pp · HMM regime sizing ablation

STACK

AI / ML	PyTorch · Chronos-T5 · LightGBM · XGBoost · HMM
DATA	yfinance · pyarrow
DELIVERY	Streamlit · IEEE-format writeup

<https://www.brixloop.com/research/regime-adaptive-equity-model>

Engagement model.

Every build is fixed-scope and fixed-price. No hourly billing during the build phase. You know the cost and timeline before we start.

Once a system ships, most clients move to a monthly maintenance retainer. We built it, we know it best, and AI systems specifically need upkeep most software doesn't: model updates, prompt/eval drift, provider API changes, usage scaling. The retainer isn't a lock-in tactic. It's the honest cost of keeping an AI-native system healthy after launch.

Brand Design

¥50,000+

Brand identity / visual design · 1+ week

AI MVP Build

¥2,00,000+

Core AI system + dashboard · 2+ weeks

Full Product Build

¥3,00,000+

Full-stack AI SaaS product · 3+ weeks

Enterprise Platform

¥5,00,000+

ERP / complex data platform · 4+ weeks

Custom / Other

Quoted

When the brief does not fit a standard path · Scoped

Maintenance

Quoted monthly

Post-launch support, model updates, ongoing features · Monthly

PAYMENT TERMS

50% upfront, 30% at milestone, 20% on delivery.

50%

UPFRONT

Due to start the build

30%

MILESTONE

Due at the agreed mid-build
checkpoint

20%

DELIVERY

Due on ship and handoff

Listed prices are starting points. Custom work is quoted after discovery. Final scope and pricing are confirmed after inquiry review.

Ready to build something real?

We work with a small number of teams at a time. If you're building an AI-native product and need people who can move fast, start an inquiry.

INQUIRY

www.brixloop.com/inquiry

EMAIL

hello@brixloop.com

LEGAL • FINTECH • TRAVEL & HOSPITALITY • REAL ESTATE • MEDIA & CONTENT •
CREATOR OPERATIONS • HEALTHCARE • CYBERSECURITY • QUANT FINANCE •
INDUSTRIAL OPERATIONS • PRODUCTIVITY TOOLING • WORKFLOW AUTOMATION
