

# Fast-Track PoC 48 h – Automated Mock-up Generation Pipeline

---

## 1. Introduction

Leading technology firms are converging on subscription-based, agentic AI delivery models that compress proof-of-concept (PoC) cycles from weeks to days. Meanwhile, advances in multimodal language models, speech-to-text, and text-to-UI generation enable near-real-time conversion of raw stakeholder input into interactive prototypes. This document outlines a 48-hour pipeline that transforms the first virtual meeting with a client into a branded, high-fidelity mock-up, while preserving essential human oversight for quality and feasibility.

## 2. Objective

Deliver a clickable, client-branded mock-up within 48 hours of the initial virtual kickoff meeting, utilising automated AI services for transcription, requirements extraction, and UI generation, with less than four hours of human labour for review and refinement.

## 3. End-to-End Workflow (≤ 48 h)

Elapsed Time	Phase	Automation	Human Checkpoint
0 h	Kickoff Meeting	Platform recording & live transcription	—
+1 h	Requirement Extraction	LLM structured JSON (User Stories, Personas)	Curator validates coverage
+4 h	Wireframe Generation	Text-to-UI generator with brand tokens	Designer 15-min review
+24 h	High-Fidelity Mock-up	Design-system components auto-assembled	Lead Engineer checklist
+30 h	Brand Injection	Script inserts logo & corporate data	—

+42 h	QA & Visual Diff	Automated diff + GPT-Vision heuristics	QA Lead approval
≤ 48 h	Delivery	Figma/InVision link + PDF summary	Account Director “send”

## 4. Technical Architecture

The pipeline is orchestrated by a lightweight agent controller (e.g., LangChain + Celery) that sequences independent services via REST and messaging queues:

- Speech-to-text API captures meeting audio and produces diarised transcripts.
- Extraction agent prompts an LLM for structured requirements in JSON.
- UI-generation agent calls a text-to-design API, seeding colour, typography, and logo tokens.
- Brand layer merges corporate assets stored in an S3-compatible bucket.
- QA bot performs visual regression tests and compliance checks before release.

## 5. Toolchain & Rationale

Domain	Tool / Service	Rationale
Speech-to-Text	Zoom AI Companion or AWS Transcribe	Accurate multilingual diarisation
LLM Extractor	GPT-4o with function calling	Reliable structured output
UI Generation	Uizard Autodesigner API	Editable multi-screen prototypes via text
Brand Tokens	Figma REST + JSON design-system	Pixel-perfect consistency
QA	Playwright + GPT-Vision	Automated visual & semantic checks
Orchestration	LangChain agents + Celery	Retry logic & traceability

## 6. Governance & Roles

- AI Ops Lead – maintains pipelines, monitors token usage, and tunes prompts.
- Product Engineer (Curator) – approves stories and mock-ups within a two-hour SLA.

- UX Designer On-Call – ensures accessibility and edge-case brand alignment.
- Customer Success – gathers feedback during the 48-hour review demo.

## 7. Implementation Roadmap

- Sprint 0 (2 weeks): design-system tokens, minimum viable transcription & extraction flow.  
Pilot (1 month): four simulated PoCs to calibrate KPIs and refine prompts.  
Client Launch (Jul–Aug 2025): subscription-like “Mock-up Pod” delivering one PoC per 48 h.  
Scale-up (Q4 2025): multi-client expansion with token-based billing.

## 8. Value Proposition

- 48-hour time-to-wow equals or surpasses benchmark industry offerings.
- Predictable pricing per PoC token instead of variable effort-based fees.
- Minimal human effort focused on high-value validation tasks.
- Deterministic branding guarantees consistent client identity.
- Extensible pipeline lays groundwork for future automated code generation.

## 9. Next Steps

Host a three-hour PoC-Lab Day: capture a real client case on video, run the 48-hour pipeline, and demo an operational mock-up along with efficiency metrics and financial projections.

## 10. Market Landscape: Rapid PoC Acceleration Players

Multiple organisations already operationalise 48-hour (or shorter) proof-of-concept cycles. Understanding their positioning helps benchmark costs, methodologies, and differentiators for our proposal.

Category	Representative Companies	Key Offering / Relevance
Global Consultancies	Accenture, Deloitte, Infosys, Cognizant, Globant, Capgemini, IBM Consulting, EPAM	Subscription-like AI factories, agentic frameworks, and extensive template libraries that compress prototyping to days.
Cloud Hyperscalers	AWS, Microsoft Azure, Google Cloud	Innovation centres and pre-built GenAI pipelines that translate business

	prompts into deployable prototypes rapidly.
Design-Generative Start-ups    Uizard, Galileo AI	Text-to-UI services yielding editable, branded screens within seconds; ideal engines for automated mock-up stages.

These players illustrate an industry shift toward token-based, outcome-centred delivery. Our pipeline adopts compatible standards (LLM agents, design tokens, cloud-agnostic orchestration) while preserving boutique flexibility and human curatorship.

## References (APA)

Globant. (2025, June 5). Globant reimagines traditional IT services with AI Pods: The first subscription model for AI-powered engineering. Globant News.

<https://www.globant.com/news/globant-introduces-ai-pods-model>

Uizard. (2025). Autodesigner 2.0: AI UI design generator. Uizard.

<https://uizard.io/autodesigner/>

Accenture. (2025, June 11). Accenture launches Distiller agentic AI framework to accelerate scalable industry AI solutions. Accenture Newsroom.

<https://newsroom.accenture.com/news/2025/accenture-launches-distiller-agentic-ai-framework-to-accelerate-scalable-industry-ai-solutions>

Zoom Video Communications. (2025, April 16). Customize Zoom AI Companion to accommodate your organization's needs. Zoom Blog.

<https://www.zoom.com/en/blog/custom-ai-companion-for-your-organization/>

Business Insider. (2025, June 8). A top consulting firm rips up its traditional billing playbook for the AI era. <https://www.businessinsider.com/consulting-firm-rips-up-traditional-billing-playbook-ai-era-globant-2025-6>

Deloitte. (2025, May 30). AI Factory as a Service brings rapid generative AI capability to clients. Deloitte Press Release. <https://www2.deloitte.com/global/en/pages/about-deloitte/articles/ai-factory-service-2025.html>

Amazon Web Services. (2025, March 18). AWS Generative AI Innovation Center helps customers build and scale. AWS News Blog. <https://aws.amazon.com/blogs/aws/aws-generative-ai-innovation-center/>

Microsoft. (2025, February 12). Introducing Azure AI Foundry. Microsoft Azure Blog. <https://azure.microsoft.com/en-us/blog/introducing-azure-ai-foundry/>

Google Cloud. (2025, April 25). Accelerating generative AI prototypes with Vertex AI and Gemini. Google Cloud Blog. <https://cloud.google.com/blog/products/ai-machine-learning/vertex-ai-gemini-generative-ai-prototype>