NAVAR

Radical AI Workflow Reinvention Superprompt

Radical AI Workflow Reinvention Superprompt

Step 1: Understand the Process

Instructions for LLM:
1. Ask the user:
- What is the process you want to reinvent? (e.g., "New Product Development", "Candidate Hiring", "Customer Complaint Handling")
- What industry/domain is it in? (e.g., "Healthcare", "Retail", "Banking")
- Upload or describe the current workflow (e.g., step list, flowchart). Optional.

2. If no workflow is provided:
- Based on the process name and industry, generate a **typical current-state workflow** using your knowledge.
- Estimate typical cycle times and human roles.

Step 2: Radically Redesign the Workflow with AI

Instructions for LLM:
- Treat AI as a **real-time, intelligent collaborator** — not just an automation tool.
- Remove all legacy constraints: redesign from scratch.
- Humans should only be required for tasks that either a) would be done better by a human or b) need their governance (eg final judgment, strategic alignment, ethics, and risk).
- AI handles all intermediate tasks, real-time analysis, ideation, optimization, and iteration.

Step 3: Create Two Radical AI-First Versions
Instructions for LLM:
- Design and present two fully AI-powered versions of the reimagined workflow:
- One version based on the Basic Stack capabilities.
- One version based on the Advanced Stack capabilities.

Use the following principles to distinguish the two:
Basic Stack =
"Proven, accessible, lower-risk AI capabilities that can radically speed up workflows while maintaining operational stability and requiring standard human oversight."
Advanced Stack =
"Emergent, less proven but highly promising AI capabilities that offer radical acceleration and transformation possibilities, but require greater care, calibration, and risk management by human overseers."

Notes

- You may reference product names if helpful to pinpoint the functional capability, but explain this in text too, don't rely on users being familiar with specific products.
- Operate with the assumption that AI-native workflows enable real-time or near-real-time execution.
- Only slow down estimates if concrete, unavoidable constraints are identified.
- If extreme compression seems unrealistic for a specific step, provide both a best-case and a conservative-case estimate for clarity.

Guide to Systematic Timeline Compression

When estimating the compressed timeframes for each redesigned workflow step, apply the following structured thinking:

- **Task Complexity:** Assess the complexity of the AI task. Is it simple, medium, or complex?
- (Simpler tasks should be near-instantaneous; complex multimodal tasks may take longer.)
- **Model Inference Speed:** Estimate how quickly the AI model(s) involved can deliver usable outputs once triggered.
- (Consider whether the task is single-output or batch-processing.)
- **Synthetic Generation Time:** If synthetic data, synthetic personas, or simulated users are involved, estimate how long they take to generate.
- (Prioritize seconds to minutes unless complexity demands more.)
- **Pre-Processing Requirements:** Assess whether significant pre-processing, cleaning, or filtering of AI outputs is needed before human review.
- (If yes, estimate the additional time required.)
- **Human Cognitive Review Load:** Estimate how long a competent human would need to review, govern, and approve the curated AI outputs at each checkpoint.
- (This should typically be measured in minutes, not hours, unless the step carries unusually high complexity or risk.)
- **Integration Overhead:** Consider whether multiple AI tools or systems need to coordinate for this step, and estimate any additional latency or integration time.
- (If minimal, treat integration as near-instantaneous.)
- **External Friction Factors:** Identify whether any legal, regulatory, compliance, or external strategic governance elements impose non-compressible delays.
- (If such constraints exist, flag them clearly and estimate the minimum unavoidable delay.)
- **Synthetic Validation Speed Expectation:** Assume that synthetic testing cycles (eg synthetic persona testing, synthetic focus groups, virtual sensory simulations, etc.) can typically be completed in seconds to a few minutes at most, once systems are operational.
- Only if multiple synthetic generations, validation loops, or complex model refinements are required should these steps extend beyond this timeframe — and if so, flag and justify the extension clearly.

- **Bottom-Up Time Estimation Based on AI Capabilities:**
- When estimating timeframes for AI-powered workflow steps, do not simply compress traditional timelines by a fixed factor.
- Instead, estimate time aggressively based on the actual performance characteristics assumed:

- How quickly the AI can generate or process the required outputs.
- How fast human reviewers can engage with curated, high-signal outputs.
- How much integration or system overhead realistically exists.

For Advanced Stack scenarios, if you assume highly autonomous, real-time, or agentic tools, reflect that in the time estimates.
Timings should be based on expected capabilities, not compressed traditional cycles.
Only apply longer durations if the specific AI tool or system complexity genuinely requires it — and flag the reason clearly.

Step 4: Output Format
Before/After Comparison
For each step in the original workflow, provide:
- Original Step
- Original Actor
- Original Timeframe
- AI-Redesigned Step
- AI/Tool Used
- New Human Role
- New Estimated Timeframe

Present these clearly, either as a structured list or formatted text blocks.

Dual Scenario Time Estimation:
For each AI-redesigned workflow step, provide two time estimates:
- **Best-Case Estimate:** Assuming ideal AI performance, minimal human friction, and seamless integration.
- **Conservative Estimate:** Allowing for minor human delays, minor AI re-iterations, or initial system calibration needs.

If the two estimates are very close, note that explicitly. If they differ significantly, flag the source of friction or uncertainty.

Summary of Transformation
- Key bottlenecks eliminated
- Estimated total cycle time compression
- Reduction in human task load (hours, minutes or seconds if feasible)
- Strategic checkpoints (where humans intervene)

Human-AI Calibration Estimates
- For each human checkpoint:
- Estimated time required for review/decision-making
- Assessment of whether this introduces friction compared to AI pace
- Flag any human steps likely to require significantly longer attention.
- Where relevant, recommend pre-filtering and/or pre-prioritisation by AI to reduce load.

Calibration Validation Recommendation
- Advise that the user simulate the new workflow using real (or synthetic) inputs.
- Encourage real human testers to validate time estimates and identify recalibration needs.
- Note that quality must not suffer for speed: **aggressive but realistic calibration is required**.

Stack Used
- Clearly indicate: (Basic Stack or Advanced Stack) for each version.

Final Instruction to the LLM:
This is a radical reimagination of the workflow, not an incremental improvement. AI is treated as a full creative, analytical, and operational partner. Humans govern final judgment, ethics, and strategic integrity only. Build for real-time, high-signal collaboration, maximum compression of timelines, and optimized human-AI workload distribution. Aggressively prioritize speed, but never at the expense of final human judgment, ethical standards, or long-term strategic goals.

IMPORTANT: Once the user has given their initial input, progress seamlessly through each step until the end. Do not ask for, or wait for, any intermediate input or feedback.