

Emerging UX Trends

Hi Alex, here's your personalized Deep Research Briefing on **Emerging UX Trends**.

Quick summary: Interfaces are shifting from fixed designs to systems that generate and adapt in real time, while agent UX is consolidating toward persistent context and longer unsupervised runs.

Reviewing releases from August 19 through September 18, 2026, several distinct UX shifts stand out across coding tools, AI agents, and consumer products. Below is a breakdown by theme.

1. Interfaces Are Becoming Generative, Not Just Assisted

The most structurally significant trend is the collapse of the line between "design" and "runtime" in interface generation.

- Runway launched Solaris, [an Interface World Model](#) that generates entire application interfaces frame-by-frame as continuous video in response to user interaction, eliminating the intermediate code layer between visual design and interactive behavior.
- Runway also shipped a [real-time sketch animation app](#) where a live video model generates frames as the user draws, converting static sketches into motion instantly.
- Runway's co-CEO Cristóbal Valenzuela argues this points toward interfaces built on "ultrafast live diffusion models" with software logic shifting to LLM-based state prediction, [trading determinism for adaptability](#). He separately contends that meaningful AI breakthroughs will have no present-day analogue and calls incremental extrapolation from current categories "[steam-powered AI](#)" thinking.
- Vercel's CEO Guillermo Rauch makes a related but distinct claim: that agents now make language choice irrelevant, functioning as a new compilation layer that translates intent directly into working software across Zig, Go, Rust, TypeScript, and Python at [equal iteration speed](#).

This is a bet that "UX" stops being a fixed set of screens and becomes something generated live per interaction. It remains unproven at production scale, and no source here shows this shipping in a mainstream consumer product yet, only Runway's experimental app and research framing.

2. Design-to-Code and Prototyping Tools Are Converging

Multiple companies shipped tools that blur prototyping, design, and implementation into one continuous loop:

Company	Tool	Core UX pattern
Anthropic	Claude Code /design command	Editable UI mockups as artifacts, iterated before implementation
Stripe	Harbor (internal)	AI agent prototypes in-browser, self-inspects work, comments stay anchored through page rewrites
Vercel	design.md system	Markdown design-guidance file plus public stylesheet lets agents generate on-brand pages outside codebases, validated through an evaluation loop
Hex	Generative Chart Controls	Point-and-click chart styling synced back to code, avoiding regeneration loops

The shared pattern: let an agent generate freely, then give humans a lightweight, non-destructive way to correct or steer the output without regenerating from scratch. Anthropic extended this logic further by letting Claude Code generate docs and slides as shareable, editable artifacts alongside code, turning chat [into a draft-share-iterate-build workflow](#).

3. Persistent, Context-Preserving UI Panes

A recurring pattern this month is keeping information visible rather than making users navigate away and lose state:

- Claude Code's `/diff` command now renders as a persistent, scrollable, clickable pane that updates in real time [instead of transient output](#).
- Claude Code Desktop panes can pop out into standalone windows and [dock back on demand](#).

- Figma's agent chat panel can open in its own separate window on desktop, staying visible while switching [design tools or tabs](#).
- Shopify POS added a persistent mobile cart button showing item count and running total across all screens, replacing [a dedicated cart tab](#).
- Shopify redesigned its Payouts page with a list-and-detail layout so merchants review multiple payouts without [losing context](#).

This is less about AI specifically and more a broader move toward "don't make me leave this screen to see the thing I need."

4. Agent Conversation UX Is Maturing Beyond Simple Chat

As agents take on longer, multi-step tasks, several products are rethinking what a chat turn even means:

- Retool's Chat tab now queues up to five messages while an agent works, batching them once the agent responds, [with edit/remove/clear controls](#).
- Linear Agent now reacts with an emoji instead of sending a full reply when no substantive response is needed, [cutting notification noise](#).
- Lovable's `/goal` command lets a user describe an outcome once and have Lovable work uninterrupted for up to 10 hours without asking clarifying questions, at [a higher credit cost](#).
- Lovable also consolidated its Chat actions menu so skills, projects, connectors, designs, and files all live under one "Add context" group, reachable via @, /, [or a + button](#).
- DoorDash rebuilt its Ask DoorDash shopping interface around a grounded product-list artifact rather than a carousel-based chat, distinguishing simple edits (immediate updates) from complex ones (new agent turns); July 2026 sessions showed roughly two UI interactions per list and about a [third converting to cart](#).

The direction is consistent: reduce back-and-forth friction, either by batching, going quiet when unnecessary, or letting the agent run long unsupervised, while giving users structured objects (lists, artifacts) to edit rather than plain text.

5. Voice Interfaces Are Shifting From Turn-Based to Continuous

Voice UX had one of the busiest 30-day windows in the set, with a clear move toward full-duplex, always-listening interaction:

- Google released Gemini 3.8 Live and Gemini 3.8 Live Extended Thinking, supporting near-real-time visual processing, 97-language auto-switching mid-conversation, and background tool execution without interrupting dialogue; the Extended Thinking variant [reasons and speaks simultaneously](#).
- OpenAI's GPT-Live-1, a full-duplex voice model supporting natural pauses and interruptions, was integrated into Yelp Host and Hatch as one [of its first deployments](#).
- HeyGen's Wayne Liang points to a specific failure mode this unlocks: real-time agents often misread a learner's mid-sentence pause as the end of a turn, which breaks language-tutoring interactions; he frames improved pause detection as core to whether a beginner-facing voice agent [feels usable at all](#).
- Twilio published a reference blueprint for handing live calls from an ElevenLabs voice agent back to human agents (via Flex or custom TwiML) [while preserving call context](#), and also brought ElevenLabs TTS voices into its Say product in public beta.
- Meta's Muse expanded outbound calling to US businesses after users [specifically requested phone functionality](#), and Meta Superintelligence Labs shipped Muse Voice Transcribe, a real-time speech model with diarization for 20+ concurrent speakers [across 25+ languages](#).

Full-duplex, interruption-tolerant voice is quickly becoming baseline expectation rather than a differentiator, which raises the bar on handling silence, overlap, and mid-thought pauses gracefully.

6. Trust Signals Are Becoming Explicit UI Elements, Not Just Backend Policy

Several products are surfacing "is this AI, can I escalate, is this verified" directly in the interface rather than leaving it implicit:

- Intercom's 2026 AI Sentiment Report found 49% of end users report positive experiences with AI agents (up 9 points year over year), but trust jumps to 74% only after users watch an agent resolve a real query, versus 49% in the abstract. Top trust drivers were easy escalation to a human and transparency that the user is talking to AI, [not a person](#).
- [Lovable's Insights dashboard](#) and Amplitude's Wave agent both surface adoption and conversion metrics directly to admins, treating onboarding funnel visibility as a

product feature rather than an internal report.

- Approval-before-action patterns are spreading: Grok Bot now drafts messages inline for [user review before sending](#), and OpenAI's Admin plugin for ChatGPT Work routes usage approvals through Slack or Teams while [respecting existing role permissions](#).

The pattern across these: users trust behavior they've watched succeed, not capability claims, so products are building in visible proof points (escalation paths, review-before-send, dashboards) rather than asking for blind trust.

7. Mobile Performance and Native Rebuilds

Several companies are treating mobile as a first-class surface again, not a stripped-down companion:

- Notion is rebuilding its mobile app natively from scratch, replacing the web-wrapper implementation, and is actively soliciting user pain points [during early testing](#).
 - Notion also cut AI chat scroll stutter on iOS by 74%, sped up answer appearance by 9%, and reduced memory use by 26%, alongside adding a double-tap gesture to [copy specific response text](#).
 - Replit's iOS app uses stacked swipe gestures (up for projects, right for tools, further right for a kanban board) to [manage limited screen space](#).
 - Ramp is shipping a mobile app redesign, though details on scope have [not been released yet](#).
 - Gamma released native iOS and Android apps with 40+ language support, push notifications, [and native exports](#).
-

8. Access Control and Credential UX Are Getting Quiet but Real Upgrades

Security-adjacent UX improved in ways aimed at reducing friction without reducing safety:

- Google enabled direct password-and-passkey transfer between Android password manager apps (Bitwarden, 1Password, Dashlane, and more coming), [eliminating CSV export/import](#).
- 1Password redesigned phishing prevention around a 50,000-URL verified list and added native [macOS Credential Provider support](#).

- WorkOS added live, real-time feedback on password requirements as [users type in AuthKit](#).
 - Stripe's Instinct agent can now generate and enter TOTP codes for two-factor authentication itself, removing a common interruption point in automated agent tasks, while still requiring user approval for enrollment and phone-based [step-up checks](#).
-

Where This Is Heading

Two tensions run underneath all of this. First, several leaders are pushing back on agent-as-default-UI: Sentry's David Cramer argues that exposing raw agent tool behavior is often a substitute for actually designing a workflow, and that this leaves users guessing rather than [getting clarity upfront](#). Zed designer Danilo Leal separately flags that even mature products like Linear have features (command palettes) that [go undiscovered](#), suggesting agent-driven and command-driven interaction can worsen discoverability if not paired with deliberate onboarding.

Second, several companies are explicitly designing to keep humans in the loop rather than removing them. Notion's agents can now propose edits for sequential review instead [of applying them directly](#), and Zapier's Wade Foster built adversarial "war-council" subagents specifically to stop agents from defaulting to sycophantic [agreement with the user](#). Gamma's Grant Lee frames the best AI products as making users more capable rather than simply automating more, citing amateur chess players beating grandmasters through superior process and judgment rather than [raw computational power](#).

The releases this month suggest UX is bifurcating: one track toward fully generative, agent-driven interfaces with minimal fixed structure (Runway's Solaris, Vercel's language-agnostic compilation framing), and another toward more deliberate, reviewable, human-gated agent interaction (Notion's suggest-edits, Zapier's adversarial checks, Sentry's workflow-first stance). Which one wins for a given product likely depends on how much users need to trust output without watching it happen first, a gap Intercom's research suggests is still wide.

[Read the full report](#)



Researched and written by AI on the Department of Product Intelligence platform. It cites its sources, and they are worth checking before you act on anything here.

You're receiving this because a teammate set up a recurring research briefing ("UX trends") and added you as a recipient. [Unsubscribe from this briefing](#) · [Manage briefings](#)