

DevDay 2026 Recap

PRIMARY

Publisher	OpenAI
Published	2026-09-29
Retrieved	2026-09-30
Cite key	openai2026devday
URL	https://openai.com/index/devday-2026-recap/

Assessment. OpenAI's own consolidated announcement list. Primary for product names, availability tiers and the quoted performance and pricing language.

Claims drawn from this source. CLAIM-2026-09-30-01, CLAIM-2026-09-30-02, CLAIM-2026-09-30-03, CLAIM-2026-09-30-04

OpenAI DevDay 2026, published 29 September 2026. Consolidated announcement list.

dots. "Remarkably capable, always-on agents built to handle everything. They're a whole new way to work with AI - one that gets to know what matters to you." Availability: Pro and Business Premium in eligible markets; Enterprise, Edu and Healthcare in beta, off by default.

GPT-6.1 Sol. A major upgrade to GPT-6 Sol with "exceptionally strong performance on agentic coding, along with computer use, and professional work". "Delivers near-Astra intelligence to everyone at a fifth of its standard input and output token prices." Available to all API, Plus, Pro, Business, Enterprise and Edu users.

Ultrafast. A premium speed tier offering "up to 8x faster token generation (300 tokens per second) in Codex and up to 6x in the API". GPT-6 Astra Ultrafast in the API and ChatGPT Work/Codex on Pro 500 and Enterprise; GPT-6.1 Sol Ultrafast coming soon.

Agents API with Computer Use. Now "supports computer use so developers can build agents that interact with software to complete tasks"; includes multi-agent capabilities and tool calling. API; Codex and ChatGPT Work on Pro 500 and Enterprise.

Private Intelligence. Includes "Zero Data Retention with Private Safety Processing" enabling "automated safety reviews without giving OpenAI personnel access to the underlying content". Private Inference preview "coming this fall, combines confidential computing with strict, verifiable controls".

Codex Security Cloud. "Scan entire GitHub repositories on demand or on a schedule, with ongoing checks of new commits." Pro, Business, Enterprise and Edu.

Codex in the Cloud. Developers can "run Codex wherever they need it: on a computer, remotely from a phone, or in the cloud from any device".

MCP Events for Plugin Automations. "Support for the proposed MCP Events specification, so plugins can start automations when something happens in a connected app." All plans.

Decisions API. "Enables real-time decision-making by focusing Luna's intelligence on a specific set of user-defined questions with finite pre-defined answers." Limited preview.

Bedrock Managed Agents. Partnership with Amazon allowing "use of Bedrock Managed Agents to build OpenAI agents that run entirely in AWS".

Sign in with ChatGPT. Plan allowance usable "across 16 partners including Cognition's Devin, Notion, Vercel, T3, OpenClaw, and Dactyl".

Pro 500. "Offers our highest usage allowance at 25 times the ChatGPT Plus allowance and includes access to Ultrafast."