



TreScout

FRESH ONLY · DAILY TECHNOLOGY REPORT

Today's New Trends

September 18, 2026 (Friday)

A FEW SIPS

This report contains trends that have emerged for the first time today, without repeating records included in normal reports from the last 30 days.

24 HIGHLIGHTS · ~6 MIN READ

11 GitHub
Trending

4 Hacker
News

1 HuggingFace ·
Daily Models

5 HuggingFace ·
Daily Papers

3 Lobsters

Data captured September 18, 2026 at 11:47 AM.

Report source counts and item metadata are fixed at capture time. Discovery pages may show newer live source values.

01 cloudflare/security-audit-skill

security-audit-skill, developed by Cloudflare, is a skill package that enables coding agents to perform multi-stage security audits. It automates security standards in software development processes by producing machine-readable and independently verified findings.

JavaScript · ★ 11,472 · +3,607 today <https://github.com/cloudflare/security-audit-skill>

02 Tencent/BrowserSkill

BrowserSkill, developed by Tencent, is a browser automation tool that allows AI agents to manage active browser sessions without interrupting the user's work. Through a command-line interface and browser extension, it allows any AI agent to access existing logged-in browser data.

TypeScript · ★ 4,721 · +1,302 today <https://github.com/Tencent/BrowserSkill>

03 anthropics/knowledge-work-plugins back in the spotlight

knowledge-work-plugins, shared by Anthropic, is an open-source plugin library that allows knowledge workers to create customized workflows on the AI assistant Claude (AI model). This Python-based toolkit enables users to automate productivity processes and standardize complex tasks.

Python · ★ 24,687 · +287 today <https://github.com/anthropics/knowledge-work-plugins>

04 Tencent/WeKnora

WeKnora, developed by Tencent, is an open-source platform that transforms raw documents into a queryable knowledge base (RAG), an autonomous reasoning agent, and a self-updating wiki system. Written in the Go programming language, this tool aims to convert unstructured data into enterprise-level accessible knowledge assets.

Go · ★ 26,667 · +1,125 today <https://github.com/Tencent/WeKnora>

05 abue-ammar/tinycast

Tinycast is a lightweight application launcher, shortcut manager, and clipboard history tool developed for the macOS operating system. Written in Swift, this native software aims to accelerate desktop workflows by using system resources efficiently.

Swift · ★ 6,286 · +739 today <https://github.com/abue-ammar/tinycast>

06 cilium/cilium

Cilium is an infrastructure tool that uses eBPF technology to manage network traffic, provide security, and offer system observability. It is particularly preferred for managing the network layer with high performance in container orchestration platforms like Kubernetes.

Go · ★ 25,329 · +111 today <https://github.com/cilium/cilium>

07 **jamiepine/voicebox** back in the spotlight

Voicebox, developed by Jamiepine, is an open-source AI voice studio where users can perform voice cloning, dictation, and speech generation. Written in TypeScript, this tool consolidates voice processing workflows into a single platform for the user.

TypeScript · ★ 55,005 · +667 today <https://github.com/jamiepine/voicebox>

08 **roboflow/supervision** back in the spotlight

Supervision, developed by Roboflow, is a Python library that provides reusable tools for computer vision projects. It accelerates image processing workflows by standardizing common tasks such as object tracking and visualization.

Python · ★ 50,895 · +329 today <https://github.com/roboflow/supervision>

09 **TencentCloud/Octop**

Octop, developed by TencentCloud, is a multi-user and multi-agent AI assistant that you can host on your own server. This Python-based system allows different users to interact with multiple AI agents simultaneously.

Python · ★ 3,704 · +367 today <https://github.com/TencentCloud/Octop>

10 **cline/cline**

Cline is an autonomous coding agent that automates software development processes. It integrates into developers' projects by operating via an IDE extension, software development kit (SDK), or command-line interface (CLI).

TypeScript · ★ 68,646 · +380 today <https://github.com/cline/cline>

11 **coder/coder**

Coder is a platform that provides secure development environments for software developers and AI agents. Thanks to its cloud-based infrastructure, it allows you to manage development environments centrally by standardizing coding processes.

Go · ★ 14,996 · +145 today <https://github.com/coder/coder>

01 Fujitsu launches made-in-Japan next-generation CPU FUJITSU-MONAKA

Fujitsu has announced its new generation processor (CPU), FUJITSU-MONAKA, which focuses on energy efficiency for data centers. Manufactured in Japan, this processor offers an architecture optimized for high-performance computing (HPC) and AI workloads.

584 points · 225 comments <https://global.fujitsu/en-global/pr/news/2026/09/14-02>

Discussion: news.ycombinator.com/item?id=49715813 (584 points · 225 comments)

02 Astra for Law

Astra for Law, developed by OpenAI, is a specialized AI tool that enables legal professionals to analyze complex legal documents. Designed to increase efficiency in legal processes, this system automates large-scale file review and legal research tasks.

471 points · 500 comments <https://openai.com/index/astra-for-law/>

Discussion: news.ycombinator.com/item?id=49745940 (471 points · 500 comments)

03 Bend – A language that blocks AI mistakes via proof, on CPU and GPU

Bend is a programming language that can run in parallel on both the CPU and GPU, and prevents logical errors through formal verification. It enables software developers to perform complex parallel computations in a safer and more error-free manner.

444 points · 206 comments <https://bend-lang.com/>

Discussion: news.ycombinator.com/item?id=49746163 (444 points · 206 comments)

04 Bonsai 2 27B: Near-Lossless Compression in a 9x Smaller Footprint

Bonsai 2 27B, developed by PrismML, is a compression technology that reduces large language models to nine times their original size while maintaining their original performance. This method lowers the hardware requirements for AI models, allowing them to run on lower-capacity devices.

416 points · 121 comments <https://prismml.com/news/bonsai-2-27b>

Discussion: news.ycombinator.com/item?id=49746618 (416 points · 121 comments)

01 TokenRhythm/NeoHorse-1-4B

NeoHorse-1-4B, developed by TokenRhythm, is an AI model focused on text generation. Thanks to its low parameter count, it offers fast text generation processes by using hardware resources efficiently.

♥ 2,353 · 19,789 downloads <https://huggingface.co/TokenRhythm/NeoHorse-1-4B>

01 FAMOS: Feed-Forward 3D Articulation Modeling from Sparse Observations

FAMOS is a feed-forward model that estimates joint parameters by segmenting moving parts from sparse and irregular point clouds. It aims to model the 3D articulation structure of objects more accurately by processing multiple data points together rather than relying on single observations.

3 likes · 1 comments <https://huggingface.co/papers/2609.20817>

02 When EOS Tokens Disagree: Understanding Length Inflation in On-Policy Distillation

In on-policy distillation processes, the issue of length inflation, which causes student models to generate unnecessarily long responses, is being examined. The research reveals that this behavior is caused by a mismatch regarding end-of-sequence (EOS) tokens between base student models and teacher models that have undergone post-training processes.

26 likes · 1 comments <https://huggingface.co/papers/2609.20511>

03 PACT: Can Enterprise AI Assistants Be Trusted Under Pressure?

PACT is an evaluation framework that measures the extent to which enterprise AI assistants maintain compliance rules under pressure in sensitive processes. It aims to systematically analyze the rule-violation tendencies of large language models used in fields such as recruitment, healthcare, and finance.

2 likes · 1 comments <https://huggingface.co/papers/2609.18605>

04 WeVisDoc: From Coverage to Capability for Robust End-to-End Document Parsing

WeVisDoc addresses the data bias problem encountered in document parsing processes, which convert document images into structured text, with a two-stage framework. The system strengthens end-to-end parsing performance by both increasing content diversity and using a data-centric approach that targets the model's weak points.

8 likes · 1 comments <https://huggingface.co/papers/2609.20423>

05 RetireOPD: Self-Retiring On-Policy Distillation for Agentic Reinforcement Learning

RetireOPD presents a method for AI agents trained with reinforcement learning that regulates the supervision of the teacher model according to stages. This approach optimizes the student model's self-distillation process in situations where the teacher model is not always reliable, thereby increasing training efficiency.

9 likes · 2 comments <https://huggingface.co/papers/2609.20784>

01 I Don't Like LLMs

Software expert Martin Fowler shares his criticisms regarding the integration of large language models into software development processes. He questions the potential negative impacts of these technologies on code quality and developer competence, drawing attention to the risks brought by automation.

72 points · 24 comments <https://martinfowler.com/articles/2026-dont-like-llms.html>

Discussion: lobste.rs/s/pkj7mf/i_don_t_like_llms (72 points · 24 comments)

02 Flock cameras are riddled with security vulnerabilities and hard-coded credentials

Security researchers have discovered hard-coded credentials and various vulnerabilities in Flock security cameras, which provide license plate recognition technology, that facilitate system access. These findings indicate that devices widely used by law enforcement and private properties pose risks regarding data security.

50 points · 3 comments <https://micahflee.com/flock-cameras-are-riddled-with-security-vulnerabilities-and-hard-coded-credentials/>

Discussion: lobste.rs/s/eyg3oe/flock_cameras_are_riddled_with_security (50 points · 3 comments)

03 The Golden Spike, and Resurrecting the Vale(n) Programming Language

The Vale programming language is being redeveloped with Golden Spike, a new compiler architecture that combines memory safety with high performance. This work aims to offer the memory safety features provided by the Vale language on a more efficient infrastructure.

50 points · 8 comments <https://verdagon.dev/blog/golden-spike-reviving-vale-valen>

Discussion: lobste.rs/s/fjnkn4/golden_spike_resurrecting_vale_n (50 points · 8 comments)

LLM

Artificial intelligence models trained on massive amounts of data that can understand human language and generate text.

skill

A specific skill or area of expertise possessed to perform a task.

NSA

The American national intelligence agency that conducts security analysis by monitoring internet traffic and digital communications.

RAG

A method that allows artificial intelligence to retrieve information from up-to-date external sources outside of its own memory when generating answers.

launcher

A home screen or tool that allows you to quickly find and launch applications on computers or phones.

native

An application written in the native language of a device or operating system that runs directly without the need for an extra intermediary.

observability

The ability to understand why a system is not working properly by monitoring the data within a software system.

computer vision

Technology that enables computers to see and interpret digital images or videos like humans.

large language models

Artificial intelligence models trained on massive amounts of data that can understand human language and generate text.

ERP

Software that allows a company to manage all its business processes, such as accounting, production, and human resources, from a single center.
