



TreScout

DAILY TECHNOLOGY REPORT

Today's Technology Trends

September 11, 2026 (Friday)

A FEW SIPS

In the software world today, the shift of platform preferences back to native languages and the performance of artificial intelligence models in specialized tasks are coming to the fore. Shopify's abandonment of React Native in favor of Swift and Kotlin shows that the search for efficiency in application development architectures is gaining momentum once again.

33 HIGHLIGHTS · ~9 MIN READ

16 GitHub
Trending

5 Hacker
News

5 HuggingFace ·
Daily Models

5 HuggingFace ·
Daily Papers

2 Lobsters

Data captured September 11, 2026 at 11:34 AM.

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

01 ayghri/i-have-adhd

i-have-adhd, developed by Ayghri, is a Python library that transforms complex responses from software development agents into more focused and concise formats. It offers a direct, results-oriented output layout that facilitates access to information for users with Attention Deficit Hyperactivity Disorder (ADHD).

Python · ★ 39,426 · +3,882 today <https://github.com/ayghri/i-have-adhd>

02 bilawalsidhu/gods-eye-view

God's Eye View is a spatial intelligence tool that runs in the browser and displays real-time satellite data on a photorealistic 3D globe. This open-source software provides a satellite simulation that allows for live tracking of global activity.

JavaScript · ★ 25,388 · +1,762 today <https://github.com/bilawalsidhu/gods-eye-view>

03 obra/superpowers

Superpowers offers an agentic skills framework and software development methodology for AI agents. This tool simplifies the development of agent-based systems by standardizing complex workflows.

Shell · ★ 284,958 · +732 today <https://github.com/obra/superpowers>

04 alsk1992/CloddsBot

CloddsBot is an open-source AI trading agent that operates autonomously in over 1000 markets. Based on Claude, this software executes automated trading processes on various crypto exchanges and prediction platforms (Polymarket, Binance) using risk management and machine-to-machine payment protocols.

TypeScript · ★ 1,866 · +277 today <https://github.com/alsk1992/CloddsBot>

05 Tencent/teamai-cli

teamai-cli, developed by Tencent, is a command-line interface that allows software teams to integrate artificial intelligence into their workflows. Written in TypeScript, this project aims to increase team productivity by making development processes AI-assisted.

TypeScript · ★ 4,035 · +841 today <https://github.com/Tencent/teamai-cli>

06 AlexsJones/llmfit

llmfit, developed in Rust, allows you to identify large language models that can run efficiently on your hardware with a single command. It helps you find the most suitable option for your system resources by scanning hundreds of different models and providers.

Rust · ★ 35,928 · +258 today <https://github.com/AlexsJones/llmfit>

07 **liquidslr/system-design-notes**

Notes compiled from the book 'System Design Interview - An Insider's Guide', a preparation guide for system design interviews, summarize fundamental concepts on scalable software architectures. It provides a comprehensive study resource for technical interview processes for those who want to learn how to build complex systems.

★ 19,093 · +900 today <https://github.com/liquidslr/system-design-notes>

08 **cathrynlavery/diagram-design**

Diagram-design offers 38 different editorial diagram templates designed for AI models (Claude Code, Codex, Pi). Avoiding shadow effects to prevent visual clutter, this HTML and Scalable Vector Graphics (SVG) based library provides an alternative visualization method to standard diagramming tools.

HTML · ★ 38,130 · +1,294 today <https://github.com/cathrynlavery/diagram-design>

09 **freestylefly/awesome-gpt-image-2**

Awesome-gpt-image-2 is a comprehensive library containing prompt examples and industrial templates used for AI image generation models. With over 530 examples and reusable skill packages, it helps users standardize their visual production processes.

JavaScript · ★ 31,155 · +962 today <https://github.com/freestylefly/awesome-gpt-image-2>

10 **armory3d/armorpaint**

ArmorPaint is a graphic design tool that performs real-time texture painting on 3D models. With Physically Based Rendering (PBR) support, the software provides digital artists with flexibility in complex texture creation processes.

C · ★ 4,545 · +72 today <https://github.com/armory3d/armorpaint>

11 **diegosouzapw/OmniRoute**

OmniRoute is an open-source AI gateway that consolidates over 350 AI providers into a single endpoint. It integrates with developer tools (such as Cursor and Claude Code) and offers compression methods for token savings and automatic fallback support.

TypeScript · ★ 64,501 · +626 today <https://github.com/diegosouzapw/OmniRoute>

12 **JustVugg/colibri**

Colibri is a C-based engine that enables running large-scale Mixture of Experts (MoE) models on local computers with low hardware requirements. By processing expert layers via disk streaming, it makes it possible to run high-capacity AI models on limited hardware.

C · ★ 27,608 · +98 today <https://github.com/JustVugg/colibri>

13 THU-MAIC/OpenMAIC

OpenMAIC is an open-source platform that provides an interactive learning environment for multi-agent AI systems. It allows users to launch immersive simulations with a single click, facilitating the training and experimentation processes of agent-based systems.

TypeScript · ★ 35,635 · +837 today <https://github.com/THU-MAIC/OpenMAIC>

14 nashsu/llm_wiki

LLM Wiki is a cross-platform desktop application that automatically organizes documents into an interconnected knowledge base. Unlike traditional Retrieval-Augmented Generation (RAG) methods, it creates a persistent wiki from your sources, keeping data updated incrementally.

TypeScript · ★ 18,355 · +142 today https://github.com/nashsu/llm_wiki

15 vercel-labs/skills

Skills, developed by Vercel, is an open-source tool that offers reusable skill packages (agent skills) for AI agents. It allows developers to standardize and extend agent workflows via the command line.

TypeScript · ★ 31,317 · +122 today <https://github.com/vercel-labs/skills>

16 vastsa/PI-Desktop

PI-Desktop is an AI coding assistant developed with a local-first working principle. Using an Electron and Rust infrastructure, this desktop application offers a modular development environment where users can install their own plugins.

TypeScript · ★ 2,513 · +624 today <https://github.com/vastsa/PI-Desktop>

01 Shopify is moving from React Native back to Swift and Kotlin

E-commerce infrastructure provider Shopify is moving away from React Native, the cross-platform technology used in its mobile applications, and returning to native development languages, Swift and Kotlin. With this decision, the company aims to increase application performance and utilize platform-specific features more efficiently.

1,022 points · 704 comments <https://shopify.engineering/back-to-native>

Discussion: news.ycombinator.com/item?id=49643982 (1,022 points · 704 comments)

02 More questions about whether researchers can trust OpenAI with unpublished math

In the world of mathematics, researchers are expressing concerns about whether sharing their unpublished work with OpenAI, the developer of the AI model ChatGPT, poses risks regarding data security and intellectual property. Academics are approaching the possibility of private mathematical findings being used in the model training process with caution.

795 points · 732 comments <https://mathstodon.xyz/@andreasthom/117240535270608201>

Discussion: news.ycombinator.com/item?id=49639408 (795 points · 732 comments)

03 Rust is tier-1 language at Microsoft

Microsoft has included the system programming language Rust among its tier-1 supported languages company-wide. This step indicates that the use of Rust is becoming standardized as an alternative to C and C++ in memory safety-focused projects.

664 points · 400 comments <https://rustfoundation.org/media/guest-post-rust-is-tier-1-language-at-microsoft/>

Discussion: news.ycombinator.com/item?id=49643546 (664 points · 400 comments) lobste.rs/s/eerwba/rust_is_tier_1_language_at_microsoft (67 points · 43 comments)

04 Don't let anyone take away your big box of cables

The habit of keeping old cables is defended because it offers quick solutions when needed and resists the rapid consumption cycle of the digital world. This approach is considered a practical way to keep existing hardware functional even as technological standards change.

490 points · 328 comments <https://blog.jim-nielsen.com/2026/hands-off-my-cables/>

Discussion: news.ycombinator.com/item?id=49645393 (490 points · 328 comments)

05 Cognition launches new SWE-2 model, Rivaling Fable 5.1 and GPT-Astra

SWE-2, developed by Cognition, is a new AI model that automates software development processes. This model competes directly with rival software development tools with similar capabilities, such as Fable 5.1 and GPT-Astra.

407 points · 172 comments <https://cognition.com/blog/swe-2>

Discussion: news.ycombinator.com/item?id=49645443 (407 points · 172 comments)

01 deepseek-ai/DeepSeek-V4.1-Flash

DeepSeek-V4.1-Flash is a multimodal AI model that processes visual and text data to produce text-based outputs. Hosted on the HuggingFace platform, this model is used for visual analysis tasks due to its fast data processing capability.

♥ 1,569 · 6 downloads <https://huggingface.co/deepseek-ai/DeepSeek-V4.1-Flash>

02 openbmb/MiniCPM5-2B

MiniCPM5-2B, developed by OpenBMB, is a small-sized language model focused on text generation. Aiming to offer high performance with low hardware requirements, this model provides an efficient option for on-device AI applications.

♥ 1,143 · 42,289 downloads <https://huggingface.co/openbmb/MiniCPM5-2B>

03 XHToken/Spark-X2.5-4B

Spark-X2.5-4B, published on Hugging Face, is a compact language model focused on text generation. Thanks to its low parameter count, it aims to produce fast responses by using hardware resources efficiently.

♥ 1,076 · 15,930 downloads <https://huggingface.co/XHToken/Spark-X2.5-4B>

04 Qwen/Qwen3.8-27B

Developed by Alibaba, Qwen2.5-27B is a multimodal artificial intelligence model capable of processing both text and visual data. Offering advanced reasoning and coding capabilities, this model stands out as a high-performance option among open-source large language models (LLMs).

♥ 14,683 · 7,322,476 downloads <https://huggingface.co/Qwen/Qwen3.8-27B>

05 ISTA-DASLab/Qwen3.8-27B-GSQ-RCO-GGUF

Developed by ISTA-DASLab, Qwen3.8-27B-GSQ-RCO-GGUF is a multimodal artificial intelligence model that processes visual and text data to produce text output. Thanks to the compressed GGUF format, which provides high memory efficiency, the model performs image analysis and text generation tasks on local hardware.

♥ 809 · 614,850 downloads <https://huggingface.co/ISTA-DASLab/Qwen3.8-27B-GSQ-RCO-GGUF>

01 DRG-MAPPO: Hierarchical Dynamic Role-Graph Multi-Agent Reinforcement Learning for Cooperative Air Combat

DRG-MAPPO improves multi-agent reinforcement learning methods in complex scenarios such as air combat using hierarchical dynamic role graphs. This approach models variable relationships during conflict to increase the tactical coordination capability of autonomous systems.

1 likes · 1 comments <https://huggingface.co/papers/2609.11155>

02 MetroLLM-Bench: Evaluating Language Models as Transit Kiosk Runtimes

MetroLLM-Bench is a 955-question evaluation set that measures the performance of large language models as a policy layer in information kiosks at metro stations. It tests the model's ability to produce accurate results using structured tools for tasks such as route planning, fare calculation, and accessibility.

1 likes · 1 comments <https://huggingface.co/papers/2609.10016>

03 UniH³: Unifying Hierarchical Homogeneity and Heterogeneity for All-in-One Medical Image Restoration

UniH³ combines the diversity between different tasks and common anatomical structures in medical image restoration processes into a single model. This approach aims to increase the training efficiency and generalization ability of universal models developed for different imaging methods and degradation types.

1 likes · 1 comments <https://huggingface.co/papers/2609.11156>

04 TempCloze: Can Video-LLMs Identify the Missing Middle?

TempCloze is a new benchmark test developed to measure the temporal reasoning skills of video-based large language models (Video-LLMs). To reduce erroneous results based on linguistic cues, models are asked to find the missing middle section from four options by looking at the beginning and ending clips of a video.

1 likes · 1 comments <https://huggingface.co/papers/2609.01515>

05 X-AuT: Progressive Audio-Encoder Compression for Speech LLMs with Cross-Scale Distillation

X-AuT reduces inference costs in speech-based large language models (speech LLMs) by gradually compressing audio-encoder layers. The model compensates for data loss during layer pruning by fine-tuning with representation alignment and low-rank adaptation (LoRA).

11 likes · 1 comments <https://huggingface.co/papers/2609.11412>

01 **A rant about phishing: It's not the user's fault (and not DNS either)**

It is argued that instead of placing the responsibility for phishing attacks on users or the Domain Name System (DNS) infrastructure, focus should be placed on solutions that close security vulnerabilities at the system level. It is recommended that software developers and service providers implement technical controls to prevent phishing independently of the user.

64 points · 18 comments <https://maurycyz.com/misc/domains/>

Discussion: lobste.rs/s/9sxn91/rant_about_phishing_it_s_not_user_s_fault (64 points · 18 comments)

02 **It's not the YAML spec's fault, but**

The complex structure of the YAML data serialization format leads to debates among developers regarding the management of configuration files. This complexity, which causes syntax errors and unexpected behaviors, constitutes the reasons why alternative formats (TOML, JSON) are preferred.

33 points · 8 comments <https://slugcat.systems/post/26-09-10-yaml-spec/>

Discussion: lobste.rs/s/bsvc7a/it_s_not_yaml_spec_s_fault (33 points · 8 comments)

ADHD

A neurological condition characterized by difficulty focusing, hyperactivity, and impulsivity.

agentic skills framework

A set of capabilities that enable artificial intelligence systems to make decisions and perform tasks on their own.

trading agent

Software that automatically trades on your behalf in stock or digital markets.

command-line interface

A screen that allows you to control a computer without using a mouse, solely by entering written commands.

large language models

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

SVG

An image file drawn with mathematical formulas that can be scaled indefinitely without losing quality.

prompt

A written instruction or question you give to an artificial intelligence to obtain a result.

skills

The specific capabilities a piece of software or artificial intelligence possesses to perform a particular task.

PBR

A rendering method that makes three-dimensional objects appear according to real-world light and material properties.

AI gateway

A gateway that manages access to different artificial intelligence services and controls security and speed.

Mixture of Experts

A technique for running a large artificial intelligence model by dividing it into smaller parts, each specialized in different subjects.

multi-agent

Multiple artificial intelligence units that communicate with each other to work toward a common goal.

RAG

A method that allows artificial intelligence to provide more accurate answers by retrieving information from current or private documents outside of its own knowledge base.

agent skills

Actions that artificial intelligence software can perform in the outside world, such as browsing the internet or organizing files.

local-first

A software approach that prioritizes storing and processing data directly on your device without the need for an internet connection.

native

A version of software developed to be fully compatible with the operating system of the device it runs on and to operate at peak performance.

TreScout · trescout.com · Every day, on your schedule.