

State of Agentic Software Development

Technology press coverage • Compass Lens Report

40 topics measured • Window: Full history • Jun 2025 – Sep 2026 • 22,814 articles • As of 10 Sep 2026

A measured view of how technology-press attention is distributed, moving and connected across 40 related topics — who holds attention, who is gaining or cooling, whether attention is concentrating or spreading out, which topics the press repeatedly connects, and what kind of coverage defines the conversation.

This is a **Lens Report** covering **Agentic Software Development** — a measurement of how the technology press covered the **40 topics** listed below, between **Jun 2025** and **Sep 2026**, across **22,814** articles. Coverage draws on at least **116** distinct outlets; the busiest single outlet accounts for about **5.9%** of the set's measured press items.

How to read this report: Compass reads a curated panel of technology outlets — the same panel, at the same steady pace, month after month — and counts and classifies what the press publishes about the topics in this lens: attention, story types, direction, and what's covered together. Every figure is a measurement of that coverage — not a ranking of the products, and not advice — and each traces to a dated, cited article. The pattern is the product; the conclusions stay yours.

THE READ IN 60 SECONDS

Most-covered topic	GitHub — 17.7% of the lens's coverage
Fastest-gaining	Moonshot AI — +5 pp (0.7% → 5.7% of lens coverage)
Attention concentration	The top 3 topics hold 35% of the coverage — broadly spread
Strongest pairing	Google AI Studio + Antigravity — covered together in 53 articles • distinctiveness 0.58

What not to over-read: this measures press coverage — not adoption, revenue, or quality.

Topics in this lens • 40

Topic share = that topic's measured articles divided by the sum of every lens topic's measured articles. An article covering several lens topics counts once for each of them, so shares measure attention to topics — not exclusive slices of unique articles. Lens entries are coverage topics, not mutually exclusive market participants — a company, its products and its models can each be topics, and their coverage overlaps.

GitHub organization	17.7%	Kiro product	1%
Claude Code product	9%	Block company	0.9%
Model Context Protocol technology	8.2%	Deep Research product	0.8%
Copilot product	7.9%	n8n product	0.8%
AI agents technology	6.1%	GPT-5.6 model	0.8%
Codex product	5.5%	Windsurf product	0.8%
Google DeepMind organization	5%	Replit company	0.7%
Perplexity company	4.8%	Gemini API product	0.7%
agentic AI technology	4.5%	JetBrains company	0.7%
Cursor product	4.4%	A2A standard	0.7%
GitHub Copilot product	2.8%	DataRobot company	0.5%
Fable 5 model	1.8%	Zhipu AI company	0.5%
Moonshot AI company	1.7%	GitHub Copilot CLI product	0.5%
Antigravity product	1.5%	Gemini Spark product	0.5%

Google AI Studio product	1.4%	Agent Builder technology	0.4%
Amazon Bedrock AgentCore product	1.4%	Gemini 3.5 Flash model	0.4%
LangChain organization	1.3%	Muse Spark model	0.4%
MCP servers technology	1.1%	ChatGPT Agent product	0.3%
Canvas product	1%	Agent Mode product	0.3%
Gemini CLI product	1%	Elastic Agent Builder product	0.2%

BOTTOM LINE

Broadly spread attention is masking a rotation in agentic software development coverage: several leaders are cooling as a compact Moonshot AI cluster gains, while announcement-heavy reporting leaves results at the edge.

Key findings

- GitHub is still the coverage leader, but its share fell 4.5 percentage points when the latest three months are compared with the preceding period of equal length, making the rotation visible at the top.
- Moonshot AI recorded the largest gain, rising 5 percentage points from 0.7% to 5.7%, so the strongest recent shift sits outside the leading group.
- Alongside Moonshot AI, Fable 5 gained 2.3 points and Zhipu AI gained 1 point, so all three members of the compact cluster moved upward together.
- Launches and news updates account for 46% and 39% of coverage, respectively, so topical breadth is carried by two dominant story formats.
- The press describes 59% of reported work at announcement, 33% at launch, and 6% at results, which is why outcomes remain peripheral.
- Neutral framing accounts for 86% of coverage, compared with 8% promotional and 5% critical, so shifts in attention are presented mainly as straightforward developments rather than disputes.

What Compass measured

This report describes how the technology press COVERED the 40 topics in this lens over full history — which drew the most coverage, which are gaining or cooling, what's covered together, and how concentrated attention is. It measures coverage patterns, not the products.

Coverage leaders

Topic	Type	Share	Movement
GitHub	organization	17.7%	-4.5 pp
Claude Code	product	9%	-3.1 pp
Model Context Protocol	technology	8.2%	steady
Copilot	product	7.9%	-2.2 pp
AI agents	technology	6.1%	+1.7 pp
Codex	product	5.5%	steady
Google DeepMind	organization	5%	+1.1 pp
Perplexity	company	4.8%	-0.8 pp

Topic	Type	Share	Movement
agentic AI	technology	4.5%	steady
Cursor	product	4.4%	+0.5 pp
GitHub Copilot	product	2.8%	steady
Fable 5	model	1.8%	+2.3 pp
Moonshot AI	company	1.7%	+5 pp
Antigravity	product	1.5%	-0.6 pp
Google AI Studio	product	1.4%	steady
Amazon Bedrock AgentCore	product	1.4%	+0.4 pp

Movement is the share shift defined under Momentum shifts below.

Momentum shifts

Moonshot AI's five-point gain and Fable 5's 2.3-point gain place two members of the compact cluster among the prominent upward moves. GPT-5.6 added 2.6 points and Google DeepMind added 1.1 points, so gaining coverage extends beyond that grouping. GitHub dropped 4.5 points, Claude Code declined 3.1 points, and Copilot lost 2.2 points, so the largest cooling moves are concentrated among major coverage leaders.

Movement is each topic's share of the lens's coverage in the last 3 months versus the 3 months before that, in percentage points (pp). Topics shifting less than 0.2 pp are steady.

GAINING SHARE

- Moonshot AI +5 pp (0.7% → 5.7% of lens coverage)
- GPT-5.6 +2.6 pp (0.5% → 3.1% of lens coverage)
- Fable 5 +2.3 pp (2.9% → 5.2% of lens coverage)
- AI agents +1.7 pp (5.4% → 7.1% of lens coverage)
- Google DeepMind +1.1 pp (4.2% → 5.3% of lens coverage)
- Zhipu AI +1 pp (0.3% → 1.3% of lens coverage)
- JetBrains +0.6 pp (0.3% → 0.9% of lens coverage)
- Cursor +0.5 pp (4.7% → 5.2% of lens coverage)


COOLING

- GitHub -4.5 pp (17.3% → 12.8% of lens coverage)
- Claude Code -3.1 pp (13.2% → 10.1% of lens coverage)
- Copilot -2.2 pp (6.6% → 4.4% of lens coverage)
- Canvas -1.1 pp (1.5% → 0.4% of lens coverage)
- Perplexity -0.8 pp (3% → 2.2% of lens coverage)
- Gemini 3.5 Flash -0.8 pp (1.2% → 0.4% of lens coverage)
- Antigravity -0.6 pp (2.1% → 1.5% of lens coverage)
- MCP servers -0.5 pp (1.3% → 0.8% of lens coverage)

Coverage share across the lens

Top-3 topics ≈ 35% of the lens's coverage, across ≈ 14 effective topics — a measure of how concentrated attention is across the set.


GitHub	-4.5 pp	18%
Claude Code	-3.1 pp	9%
Model Context Protocol	steady	8%
Copilot	-2.2 pp	8%
AI agents	+1.7 pp	6%



Δ = each topic's share shift, as defined under Momentum shifts. The month-by-month share chart is on the interactive report this file was exported from.

What the coverage is made of

Pooled across the whole lens — the kinds of stories the press runs, how far the reported work has progressed (as the press claims), and the framing.



Covered together

Google AI Studio and Antigravity form the most strongly connected pairing, appearing together in 53 articles across the full press record, so launch-and-release treatment anchors part of the larger Claude Code cluster. Across that record, Claude Code appeared with Cursor in 332 articles and with Gemini CLI in 91, giving the eight-subject cluster several recurring editorial links. Moonshot AI appeared with Fable 5 in 45 articles and Zhipu AI in 26, which is why the gaining edge also reads as a compact peer set.

Pairings are ranked by connection strength: distinctiveness weighted by how many confirmed events link the pair and how far apart the two topics sit in the field, with a company's own product-family pairs set aside. So a pair can rank above another that shows a higher raw distinctiveness score or more shared articles — each card carries both numbers. Distinctiveness is normalized co-occurrence (NPML, -1 to 1): how much more often the pair shares stories than the two topics' separate coverage volumes alone would predict. Big topics co-occur often by volume alone; a high score means the pairing itself is the pattern. Each pairing also shows how many outlets its cited stories span: a pairing cited almost entirely by one specialist outlet is that outlet's editorial pattern, not necessarily a field-wide one. A measurement of shared coverage across Compass's full measured record, not only this report's window, and not a partnership or endorsement.

Google AI Studio + Antigravity

actor-actor · launch, other, release

53 shared · 0.58

Cited stories span 7 outlets.

Google Cloud Labs: Accelerate AI with Cloud Run

Cloud Blog · 7 Jul 2026

What's new with Google Cloud

Cloud Blog · 1 Jul 2026

What Google Cloud announced in AI this month

Cloud Blog · 30 Jun 2026

Google's Nano Banana 2 promises Flash speeds with Pro results

The New Stack · 26 Feb 2026

Antigravity + Gemini API

actor-actor · launch, other, release

27 shared · 0.53

Cited stories span 6 outlets.

20 questions for the Agentic Enterprise (and how Agent Platform can help)

Cloud Blog · 7 Jul 2026

What's new with Google Cloud

Cloud Blog · 1 Jul 2026

- What Google Cloud announced in AI this month Cloud Blog · 30 Jun 2026
- With Google's debut, the most important AI agent feature is now the most boring one The New Stack · 27 May 2026

Google AI Studio + Gemini 3.5 Flash actor-actor · launch, other, release

12 shared · 0.48

Cited stories span 5 outlets.

- Google launches Gemini 3.5 Flash to push AI agents deeper into enterprise workflows InfoWorld · 20 May 2026
- Google launches Gemini 3.5 Flash. How to try it for free. Mashable · 19 May 2026
- Google named a Leader in the 2026 IDC MarketScape for Worldwide Foundation Model Software Cloud Blog · 14 Jul 2026
- What Google Cloud announced in AI this month Cloud Blog · 30 Jun 2026

Antigravity + Gemini CLI actor-actor · benchmark, launch, other, release

33 shared · 0.54

Cited stories span 6 outlets.

- I tried Google's Antigravity for a week, and this limitation made me close it for good XDA · 12 May 2026
- Choosing Antigravity or Gemini CLI Cloud Blog · 4 Feb 2026
- Google Antigravity introduces agent-first architecture for asynchronous, verifiable coding workflows VentureBeat · 18 Nov 2025
- 20 questions for the Agentic Enterprise (and how Agent Platform can help) Cloud Blog · 7 Jul 2026

Claude Code + Gemini CLI actor-actor · benchmark, launch, other, partnership

91 shared · 0.51

Cited stories span 7 outlets.

- The DataRobot platform as skills in Claude Code Blog | DataRobot · 15 Jun 2026
- Xiaomi's new open source, agentic AI coding harness MiMo Code beats Claude Code at ultra-long, 200+ step tasks VentureBeat · 11 Jun 2026
- DataRobot for Developers — integrating with the Google Antigravity CLI Blog | DataRobot · 17 Jun 2026
- Valid certificates, stolen accounts: how attackers broke npm's last trust signal VentureBeat · 22 May 2026

Fable 5 + Moonshot AI actor-actor · benchmark, other

45 shared · 0.51

Cited stories span 9 outlets.

- Moonshot's Kimi K3 outperforms Fable 5 in frontend code but lags far behind in complex math The Decoder · 19 Jul 2026
- Kimi-K3 is now #1 on the Frontend Code Arena benchmark, surpassing Claude Fable 5; the model scored 88.3 on Terminal Bench 2.1, only below GPT-5.6 Sol's 88.8 (Michael Nuñez/VentureBeat) Techmeme · 16 Jul 2026
- China's Moonshot AI releases Kimi K3, the largest open-source model ever, rivaling top U.S. systems VentureBeat · 16 Jul 2026
- Anthropic is bringing back Claude Fable 5 globally after US lifts export control order — where can enterprises access it? VentureBeat · 1 Jul 2026

Moonshot AI + Zhipu AI actor-actor · other, price_change, release

26 shared · 0.56

Cited stories span 7 outlets.

- Zhipu AI's GLM-5.1 can rethink its own coding strategy across hundreds of iterations The Decoder · 9 Apr 2026
- Chinese AI developers may shift to 'paid weights' commercial licensing: Goldman Sachs South China Morning Post Technology · 28 Jul 2026
- China's AI models are closing the gap with overseas rivals on a different cost curve TechNode · 27 Jul 2026
- Why Weibo's tiny VibeThinker-3B has the AI world arguing over benchmarks again VentureBeat · 17 Jun 2026

Claude Code + Cursor actor-actor · benchmark, launch, other, partnership

332 shared · 0.58

Cited stories span 7 outlets.

- The Anthropic leader who built Claude Code says he ditched prompting — now he just writes loops. The New Stack · 10 Jun 2026
- How to delegate 40% of tickets to AI The New Stack · 11 Jun 2026
- Observability overload is drowning engineers The New Stack · 10 Jun 2026
- Try the new console experience in Amazon Bedrock, optimized for Anthropic- and OpenAI-compatible APIs AWS News Blog · 5 Jun 2026

Block + Model Context Protocol topic-actor · launch, other, partnership, release

29 shared · 0.35

Cited stories span 8 outlets.

- How Block manages its fleet of AI coding agents from Slack The New Stack · 17 Jun 2026
- Why Block handed Goose to the Linux Foundation The New Stack · 15 May 2026
- OpenAI and Anthropic Donate AGENTS.md and Model Context Protocol to New Agentic AI Foundation InfoQ · 23 Dec 2025
- OpenAI, Anthropic and Block join new Linux Foundation effort to standardize the AI agent era TechCrunch · 9 Dec 2025

Codex + Gemini CLI actor-actor · launch, other, partnership

49 shared · 0.46

Cited stories span 8 outlets.

- OpenAI's Codex desktop app is all about managing agents The New Stack · 2 Feb 2026
- The DataRobot platform as skills in Claude Code Blog | DataRobot · 15 Jun 2026
- Xiaomi's new open source, agentic AI coding harness MiMo Code beats Claude Code at ultra-long, 200+ step tasks VentureBeat · 11 Jun 2026
- JetBrains: AI agents are about to repeat the cloud ROI crisis The New Stack · 31 Mar 2026

Google AI Studio + Gemini CLI actor-actor · launch, other, release Cited stories span 9 outlets. 🔗 Gemini CLI is a free, open source coding agent that brings AI to your terminal Ars Technica · 25 Jun 2025 🔗 What's new with Google Cloud Cloud Blog · 1 Jul 2026 🔗 What Google Cloud announced in AI this month Cloud Blog · 30 Jun 2026 🔗 Google releases Gemini 3.1 Pro with improved reasoning capabilities The Decoder · 19 Feb 2026	25 shared · 0.51
Claude Code + Kiro actor-actor · benchmark, launch, other, price_change Cited stories span 6 outlets. 🔗 Claude Code turned every engineer into three. Now companies need more product thinkers VentureBeat · 27 Jun 2026 🔗 AWS Security Agent adds threat modeling, Kiro power and Claude Code plugin, and more AWS News Blog · 17 Jun 2026 🔗 Beating context rot in Claude Code with GSD The New Stack · 31 Jan 2026 🔗 AWS launches Kiro powers with Stripe, Figma, and Datadog integrations for AI-assisted coding VentureBeat · 4 Dec 2025	57 shared · 0.43
Cursor + Replit actor-actor · benchmark, launch, other, release Cited stories span 6 outlets. 🔗 I tried vibe coding an app as a beginner - here's what Cursor and Replit taught me ZDNet · 19 Jan 2026 🔗 Cursor's \$2.3B Financing Reminds Us: Coding Automation Is Still Ultra-Hot Crunchbase News · 13 Nov 2025 🔗 Why AI coding tools like Cursor and Replit are doomed - and what comes next ZDNet · 31 Oct 2025 🔗 After nine years of grinding, Replit finally found its market. Can it keep it? TechCrunch · 3 Oct 2025	46 shared · 0.51
Cursor + Kiro actor-actor · benchmark, launch, price_change, release Cited stories span 8 outlets. 🔗 Cursor-Opus agent snuffs out startup's production database The Register · 27 Apr 2026 🔗 AWS launches Kiro powers with Stripe, Figma, and Datadog integrations for AI-assisted coding VentureBeat · 4 Dec 2025 🔗 Amazon launches Kiro to streamline AI prototyping The Decoder · 27 Jul 2025 🔗 Claude Code turned every engineer into three. Now companies need more product thinkers VentureBeat · 27 Jun 2026	48 shared · 0.49
Codex + Kiro actor-actor · launch, other, release Cited stories span 7 outlets. 🔗 Intelligence is Free, Now What?
 Data Systems for, of, and by Agents The Berkeley Artificial Intelligence Research Blog · 7 Jul 2026 🔗 Agent Toolkit for AWS includes 20+ agent skills, but your agent might not load them without this one file The New Stack · 25 Jun 2026 🔗 It's safe to close your laptop now: Hosting coding agents on Amazon Bedrock AgentCore AWS Machine Learning Blog · 8 Jun 2026 🔗 Vibe coding or spec-driven development? How to choose InfoWorld · 5 May 2026	30 shared · 0.39
GitHub Copilot CLI + Gemini CLI actor-actor · other Cited stories span 7 outlets. 🔗 AI coding at the command line with Gemini CLI InfoWorld · 22 Jul 2025 🔗 6 Benefits of Sandbox Environments (and How Docker Sandboxes Delivers Them) Docker · 8 Sep 2026 🔗 YOLO Mode: Agent Autonomy Without the Guardrails Docker · 3 Sep 2026 🔗 Cursor launches Origin code hosting platform as GitHub outage exposes opening in AI coding race VentureBeat · 17 Aug 2026	10 shared · 0.47
Claude Code + GitHub Copilot actor-actor · benchmark, launch, other, partnership Cited stories span 5 outlets. 🔗 Microsoft Enlists AWS to Help GitHub Handle Explosive Growth in AI Development DevOps.com · 18 Jun 2026 🔗 Chainguard agent skills matures The New Stack · 17 Jun 2026 🔗 PagerDuty's CAIO says most AI incident tools are missing a critical layer The New Stack · 14 Jun 2026 🔗 Can JetBrains close the IDE skills gap before AI widens it further? The New Stack · 13 Jun 2026	141 shared · 0.46
agentic AI + DataRobot topic-actor · launch, other All cited stories from one outlet. 🔗 Get more from DataRobot with new capabilities for Agentic AI Blog DataRobot · 23 Apr 2026 🔗 What it takes to scale agentic AI in the enterprise Blog DataRobot · 5 Apr 2026 🔗 How to build an agentic AI governance framework that scales Blog DataRobot · 27 Mar 2026 🔗 Agentic AI deployment best practices: 3 core areas Blog DataRobot · 20 Mar 2026	19 shared · 0.40
GitHub Copilot + Model Context Protocol topic-actor · launch, other, partnership, release Cited stories span 5 outlets. 🔗 Best of 2025: GitHub Copilot Evolves: Agent Mode and Multi-Model Support Transform DevOps Workflows DevOps.com · 26 Dec 2025	115 shared · 0.43

- [How the .NET MAUI Team uses GitHub Copilot for Productivity](#) · .NET Blog · 9 Jul 2025
- [Improve Your Productivity with New GitHub Copilot Features for .NET!](#) · .NET Blog · 12 Jun 2025
- [Getting more from each token: How Copilot improves context handling and model routing](#) · The GitHub Blog · 17 Jun 2026

Claude Code + GitHub Copilot CLI · actor-actor · benchmark, launch, other, regulation

30 shared · 0.41

Cited stories span 9 outlets.


['TrustFall' Convention Exposes Claude Code Execution Risk](#) · darkreading · 7 May 2026

[Docker Sandboxes: Run Claude Code and Other Coding Agents Unsupervised \(but Safely\)](#) · Docker · 30 Jan 2026

[From pixels to characters: The engineering behind GitHub Copilot CLI's animated ASCII banner](#) · The GitHub Blog · 28 Jan 2026

[Microsoft MCP server gives AI assistants access to MSBuild logs](#) · InfoWorld · 30 Jun 2026

Coverage clusters



What to watch

The clearest edge thread is the synchronized gain across Moonshot AI, Fable 5, and Zhipu AI, because all three sit inside the same compact coverage cluster. A separate edge forms around GPT-5.6 and Google DeepMind, up 2.6 and 1.1 points, so the rotation is not confined to one press grouping. Across the lens, critical and skeptical framing total 6%, so these edge threads emerge within an overwhelmingly non-confrontational press field.

Representative cited sources (32)

The dated stories behind the read — each leading topic contributes its top stories, each linked.

- [GitHub Workflows Strain Under AI Agent Development](#) · DevOps.com · 9 Sep 2026 · GitHub
- [GitHub Puts Guardrails on Copilot's Sandbox Inside JetBrains IDEs](#) · DevOps.com · 9 Sep 2026 · GitHub
- [I built Claude Code's workflow with a local model, and it's surprisingly close](#) · XDA · 9 Sep 2026 · Claude Code

- [Claude did best on a new benchmark for agents that build agents. It still passed fewer than a quarter of the tests.](#) · The New Stack · 9 Sep 2026 · Claude Code
- [Zayo advances network as a service with agentic networking](#) · ComputerWeekly.com · 9 Sep 2026 · Model Context Protocol

[Automated agent evaluation with Amazon Bedrock AgentCore and GitHub Actions](#) · AWS Machine Learning Blog · 8 Sep 2026 · Model Context Protocol

[Agentic analytics with the Data Agent Kit](#) · Cloud Blog · 8 Sep 2026 · Model Context Protocol

[Why Your AI Agents Need An Employee ID, Not Just An API Key](#) · Forbes - Innovation · 9 Sep 2026 · AI agents

[AI agents can use Quill to talk direct to enterprise SQL databases](#) · Blocks & Files · 8 Sep 2026 · AI agents

[Chamelio Launches 'Always On' AI Agent Hub](#) · Artificial Lawyer · 8 Sep 2026 · AI agents

[OpenAI gave an AI the power to block its own engineers' code](#) · The New Stack · 9 Sep 2026 · Codex

[Alibaba.com Says Accio Ran E-Commerce Tasks at Over 50% Lower Cost](#) · Unite.AI · 9 Sep 2026 · Codex

[Type.com](#) · Product Hunt · 9 Sep 2026 · Codex

[Google is a Leader in the 2026 Gartner® Magic Quadrant™ for Enterprise AI Assistants](#) · Cloud Blog · 9 Sep 2026 · Google DeepMind

[Turn one AI prompt into multiple answers with this \\$55.30 lifetime subscription](#) · PCWorld · 9 Sep 2026 · Perplexity

[Claude Comes to CarPlay as Fifth Major AI Chatbot App](#) · The Mac Observer · 5 Sep 2026 · Perplexity

[NVIDIA Vera: Rebuilding the CPU for Agentic AI](#) · SemiWiki · 8 Sep 2026 · agentic AI

[From RAG to Agentic AI: Building the Next Generation of Intelligent Enterprise Systems](#) · KDnuggets · 8 Sep 2026 · agentic AI

[A16z Doubles Down On AI Coding Agents Months After Cursor Exit](#) · Forbes - Innovation · 9 Sep 2026 · Cursor

[Cognition hits \\$48B valuation, signaling investors believe AI coding is far from a winner-take-all market](#) · TechCrunch · 8 Sep 2026 · Cursor

[GitHub Copilot app for Beginners: Run several agents at once](#) · The GitHub Blog · 3 Sep 2026 · GitHub Copilot

[Harness rebuilt its Git repository for nonstop AI agent traffic](#) · The New Stack · 9 Sep 2026 · GitHub Copilot

[Claude Fable 5.1 vs. Fable 5: On real work, I couldn't tell them apart.](#) · The New Stack · 5 Sep 2026 · Fable 5

[Artificial Analysis Coding Agent Index: GPT-6 Astra scored 67, roughly equal to Claude Opus 5, Fable 5, and Muse Spark 1.3, but trailing leader Fable 5.1's 70 \(Artificial Analysis\)](#) · Techmeme · 4 Sep 2026 · Fable 5

[Six Chinese AI firms accused of aggressively copying US frontier models](#) · Ars Technica · 9 Sep 2026 · Moonshot AI

[US Agencies Warn Chinese AI Firms Are Extracting Advanced AI Models](#) · Security Affairs · 9 Sep 2026 · Moonshot AI

[Customizing Angular Aria Tabs Quickly with Google Antigravity CLI](#) · Angular Blog - Medium · 2 Sep 2026 · Antigravity

[I asked Claude Code, Codex, and Antigravity to build the same game, and one absolutely crushed the others](#) · XDA · 30 Aug 2026 · Antigravity

[Google brings AI music generation directly into the Gemini app with its new Lyria 3.5 model](#) · The Decoder · 6 Sep 2026 · Google AI Studio

[Gemini hands Lyria 3.5 the mic as the music model makes its debut on the app](#) · Android Authority · 4 Sep 2026 · Google AI Studio

[Deploy a multimodal WhatsApp ordering assistant with Amazon Bedrock AgentCore](#) · AWS Machine Learning Blog · 4 Sep 2026 · Amazon Bedrock AgentCore

How Intuit built an agentic disaster recovery assistant with Amazon Bedrock · AWS Machine Learning Blog · 4 Sep 2026 · Amazon Bedrock AgentCore

Methodology

Every figure is drawn from the Enginerds corpus — a curated panel of technology outlets, each read at the same steady pace over time — and is a measurement of coverage WITHIN this lens: article counts, each topic's share of the set's coverage over time, the direction of that coverage, and which topics are covered together. Titles are shown as published; inclusion is not endorsement.

Topic share = that topic's measured articles divided by the sum of every lens topic's measured articles. An article covering several lens topics counts once for each of them, so shares measure attention to topics — not exclusive slices of unique articles. Lens entries are coverage topics, not mutually exclusive market participants — a company, its products and its models can each be topics, and their coverage overlaps.

Movement is each topic's share of the lens's coverage in the last 3 months versus the 3 months before that, in percentage points (pp). Topics shifting less than 0.2 pp are steady.

Pairings are ranked by connection strength: distinctiveness weighted by how many confirmed events link the pair and how far apart the two topics sit in the field, with a company's own product-family pairs set aside. So a pair can rank above another that shows a higher raw distinctiveness score or more shared articles — each card carries both numbers. Distinctiveness is normalized co-occurrence (NPMI, -1 to 1): how much more often the pair shares stories than the two topics' separate coverage volumes alone would predict. Big topics co-occur often by volume alone; a high score means the pairing itself is the pattern.

What this report does not claim. Compass measures coverage patterns in the technology press. It does not predict outcomes, rank technologies by merit or importance, or recommend what to buy, build, or invest in. Coverage is not the same as real-world significance.

Source: Enginerds Compass — measurements over a longitudinal technology corpus. Correlational & coverage-limited; not investment advice. © Enginerds.