THE LINUX FOUNDATION
OPEN SOURCE SUMMIT
NORTH AMERICA

Embedded Linux Conference

Welcome

Jim Zemlin, CEO, The Linux Foundation

Thank You to Our 2026 Program Committee and Program Chairs!

Thank You to Our Diamond Sponsors

Thank You to Our Diamond Sponsors

Thank You to Our Platinum Sponsors

Solutions Showcase

Open Daily
After Keynote
Sessions

- Visit Sponsor Booths
- Lounge Areas
- Activities
- Job Board
- Coffee Breaks & Lunch

Women and Non-Binary Lunch

Better Together Lunch

- **Time:** 12:00 – 1:30 PM
- **Location:** Located in Seasons (Level Two)
- Open to all attendees who identify as women or non-binary to connect and network over lunch

Join Us Tonight!

Attendee Reception

- Mill City Museum
- 6:30 – 9:30 PM
- An evening of connection, conversation, and Minneapolis character. Enjoy locally inspired food and drinks.

Cap off the evening with a special drone show presented by Uvify!

Stick Around for the Raffle!

Today's Raffle Prizes

- Retro Gaming Console
- Raspberry Pi Starter Kit
- Keychron K2 Keyboard
- You must be present to win!

Code of Conduct

- Be respectful
- Be Professional
- Help us create a welcoming environment for everyone
- For any concerns, visit the registration or information desk

THE LINUX FOUNDATION
OPEN SOURCE SUMMIT
NORTH AMERICA

Embedded Linux Conference

New at the Linux Foundation in 2026

Intent to Form Agent Name Service (ANS)

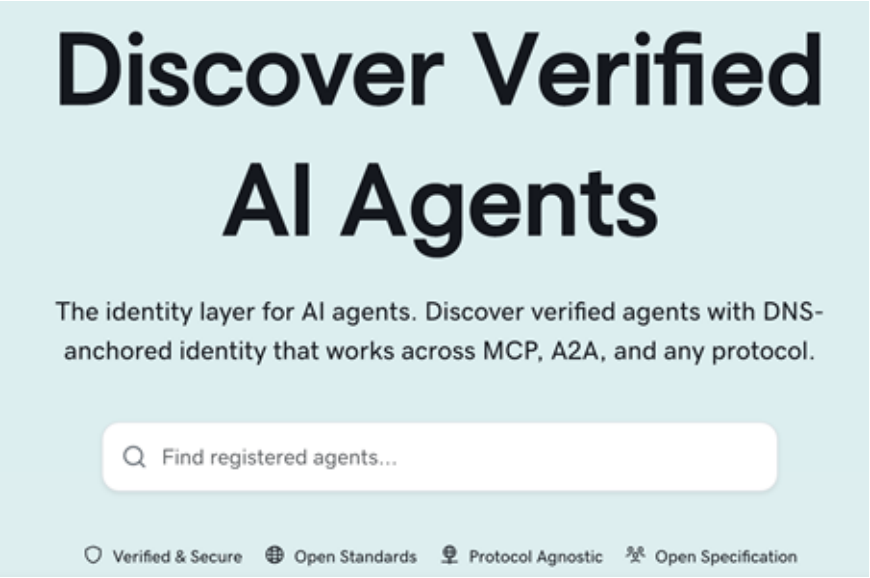

THE LINUX FOUNDATION
OPEN SOURCE SUMMIT
NORTH AMERICA

Embedded Linux Conference

THE LINUX FOUNDATION

THE LINUX FOUNDATION
OPEN SOURCE SUMMIT
NORTH AMERICA

Sometimes you eat the bear. Sometimes *the bear* eats you.

Open source and AI — what we built, what it's building, and the security reckoning we can't outrun.

Jim Zemlin · Executive Director, The Linux Foundation

Open Source Summit North America · June 2026

The bear we ate.

Open source is no longer a layer of the AI stack. It IS the stack.

Every era of computing has been built on open infrastructure

OS / INTERNET

1991–2014

100%

of public cloud

Linux · Apache · Python · MySQL

CLOUD NATIVE

2015–2021

230+

CNCF projects

Kubernetes · CNCF · containerd

AI NATIVE

2022–2025

80%

of training on PyTorch

PyTorch · vLLM · Ray

AGENTIC

2026 →

120M

MCP installs/month

MCP · A2A · AAIF · x402

Three decades — almost four — of neutrally governed open source as the substrate for whatever comes next.

THE LINUX FOUNDATION | OPEN SOURCE SUMMIT NORTH AMERICA

Sources: Linux kernel adoption — CNCF Annual Survey 2024 · Top500.org Nov 2025 · CNCF project landscape (Apr 2026) · Anthropic MCP SDK downloads dashboard

The infrastructure being built right now is the largest in tech history

$700B

2026 hyperscaler capex (Big Five)

$1.15T

Cumulative 2025–2027 (Goldman Sachs)

60%

YoY hyperscaler capex increase

1.3Q

Tokens/month Google now serves in inference

> *"Demand exceeds capacity. Power and data-center capacity are the single biggest constraint we have."*
>
> — Andy Jassy, CEO, Amazon (Q1 2026 earnings call)

Sources: Dell'Oro Group AI Infrastructure Forecast (Apr 2026) · Goldman Sachs Global Investment Research (Mar 2026) · Amazon Q1 2026 10-Q · Alphabet Q1 2026 10-Q · CNBC, May 2026

Open source runs at every layer above the hardware stack

AGENTS	MCP · A2A · x402 · AGENTS.md · Goose	**120M / month** MCP installs
MODELS	Llama · Qwen · DeepSeek · GLM · Kimi · Mistral	**~30%** of global usage (up from 1%)
INFERENCE	vLLM · SGLang · Ray · llm-d · AIBrix	**10× / yr** cost decline
TRAINING	PyTorch · DeepSpeed · Ray · Kubeflow · TorchTitan	**80%** Hugging Face share
INFRASTRUCTURE	Linux · Kubernetes · containerd · Volcano	**90%+** of cloud

THE LINUX FOUNDATION | OPEN SOURCE SUMMIT NORTH AMERICA

Sources: Anthropic MCP SDK metrics (May 2026) · Artificial Analysis open-vs-closed share Q1 2026 · a16z State of AI 2025 (inference cost) · Hugging Face Open Models Report 2026 (PyTorch share) · CNCF Annual Survey 2024 (Linux cloud share)

Layer 1 — Infrastructure: Linux + Kubernetes own the substrate

LINUX

Kernel released 1991. The substrate of everything.

100% of public cloud runs Linux

100% of the Top 500 supercomputers

80%+ of mobile devices (Android = Linux)

*Every **AI training cluster on Earth***

KUBERNETES + CNCF

Donated to CNCF in 2015. The orchestration plane.

230+ CNCF projects

300K+ contributors

Volcano · Kueue · KServe · Ray · Kubeflow

*The control plane of **every serious AI cluster***

If you're running AI infrastructure, you are running open source software.

THE LINUX FOUNDATION | OPEN SOURCE SUMMIT NORTH AMERICA

Sources: Top500.org Nov 2025 list · CNCF Annual Survey 2024 · Linux Foundation 2025 Kernel Development Report · Android Open Source Project metrics

Layer 2 — Training: PyTorch is the path to market for every chip

80%

of training jobs on Hugging Face run on PyTorch

LF-governed. Vendor-neutral. The default.

ACCELERATORS THAT HAVE A PATH THROUGH PYTORCH

NVIDIA H100 / Blackwell	AMD Instinct MI300
Google TPU v5p / Trillium	AWS Trainium 2
Intel Gaudi 3	Huawei Ascend 910C
Apple Silicon (M4 / M5)	Cerebras WSE-3

"Vendor neutrality is the feature. PyTorch is governed by the LF, not by any chip vendor — that's exactly why every accelerator has a path through it."

THE LINUX FOUNDATION | OPEN SOURCE SUMMIT NORTH AMERICA

Sources: Hugging Face Open Models Report 2026 · PyTorch Foundation 2025 ecosystem report · Vendor accelerator backend disclosures (NVIDIA, AMD, Google, AWS, Intel, Huawei, Apple)

Layer 3 — Inference & scaling: the open stack crashed the cost curve

280×

cost-per-token reduction at the frontier in 2 years

67%

of all AI compute spend is now inference

The cost curve is open source's curve.

THE OPEN INFERENCE & SCALING STACK · TANGIBLE 2026 GAINS

vLLM	**Up to 24× throughput**	PagedAttention + continuous batching. Default open serving engine. *Kwon et al., SOSP '23 · LF AI & Data*
Ray	**88% ↓ latency · 11.1× throughput**	Distributed scaling. Powers OpenAI's training + serving. *Anyscale + Google Kubernetes Engine, Mar 2026*
llm-d	**Up to 70% more tokens/sec**	K8s-native disaggregated prefill/decode; 13.9× with hierarchical KV at 250 users. *CNCF sandbox · Red Hat · Google · IBM · CoreWeave · NVIDIA*
AIBrix	**+50% throughput · −70% latency**	Distributed KV cache, cluster-scale autoscaling for LLM serving. *ByteDance, arXiv:2504.03648*

Every dollar saved on inference at hyperscale, every cost-per-token chart bending down — that's the open stack at work.

THE LINUX FOUNDATION | OPEN SOURCE SUMMIT NORTH AMERICA

Sources: a16z State of AI 2025 (cost decline) · Kwon et al. SOSP '23 (vLLM) · Particula/Spheron H100 benchmarks 2026 (SGLang) · Anyscale + GKE blog Mar 2026 (Ray Serve) · CNCF / Red Hat llm-d benchmarks 2026 · ByteDance arXiv:2504.03648 (AIBrix)

Layer 4 — Models: parity has compressed from years to months

3 – 6 MONTHS

from closed-model release to open-model parity

In 2022 the gap was years. In 2024 it was a year. In 2026 it's measured in weeks.

FRONTIER OPEN MODELS, 2026

- **Llama 4**
- **Qwen 3**
- **DeepSeek R2**
- **GLM-5 (745B, Ascend-trained)**
- **Mistral · Mixtral**

THE OPEN-MODEL ADVANTAGE

~6×
cheaper than closed alternatives

~90%
of frontier performance

30%
of global usage (from 1% in 2024)

$24.8B
in unrealized enterprise value/yr (MIT/LF)

THE LINUX FOUNDATION | OPEN SOURCE SUMMIT NORTH AMERICA

Sources: Frank Nagle (MIT Sloan / LF Chief Economist) — "The Latent Role of Open Models in the AI Economy" 2026 · Artificial Analysis open-vs-closed parity tracker (Q1 2026) · Red Hat State of Open Source AI 2026 · LMSYS Chatbot Arena leaderboard

Layer 5 — Agents: open from day one

MCP

120M

installs / month

The protocol Anthropic open-sourced 18 months ago is now the lingua franca of agentic computing. Adopted by ChatGPT, Claude, Gemini, Copilot, Cursor, Replit, AWS, Cloudflare, Google Cloud, Azure. React took 3 years to hit 100M downloads. MCP did it in 16 months.

AAIF

184+

members in <4 months

The Agentic AI Foundation — the fastest-growing foundation in LF history. Platinum: AWS, Anthropic, Block, Bloomberg, Cloudflare, Google, Microsoft, OpenAI. Seven active working groups defining the agentic stack.

x402

75.41M

transactions in the last 30 days

An open agent-native payment protocol, donated by Coinbase and joined by Stripe in April. The plumbing for agentic commerce — already a foundation, already open. Supporters include Amazon Web Services (AWS), Cloudflare, Google, Microsoft, Anthropic, Vercel, Adyen, American Express, Ant International, Circle, Fiserv, Mastercard, PPRO, Visa, Base, Coinbase, Polygon Labs, Solana Foundation, Thirdweb, Shopify and KakaoPay.

In Agentic AI the OPEN layer arrives BEFORE the closed one.

THE LINUX FOUNDATION | OPEN SOURCE SUMMIT NORTH AMERICA

Sources: Anthropic MCP SDK metrics (May 2026) · Agentic AI Foundation member roster (linuxfoundation.org/agentic-ai) · x402 / Coinbase announcement Apr 2026 · Stripe Sessions 2026

Inside Agentic AI Foundation — what it takes to deploy a production agent

AGENTIC APPLICATIONS	Vertical apps built by domain operators
AGENT RUNTIMES	Goose · AGENTS.md · OpenClaw · CrewAI · LangGraph
COMMERCE & PAYMENTS	x402 · agentic-native payment protocols
INTEROP PROTOCOLS	MCP (tools/context) · A2A (agent-to-agent)
CONTROL PLANE	Gateways · routing · workflow orchestration · scheduling
TRUST FABRIC	Agent identity · attestation · observability · security · evals
AI FOUNDATION · open models · inference · K8s · Linux · already won (Part 1)	

Every layer of the agentic stack needs a neutrally-governed open standard or implementation. AAIF will be the home.

Hosted projects

MCP · Goose · AGENTS.md

Active working groups (7)

Identity & Trust
Security & Privacy
Observability & Traceability
Accuracy & Reliability
Agentic Commerce
Governance, Risk & Regulatory
Workflows & Process Integration

Next: control-plane and identity projects expected to land at AAIF over the next 12 months.

AAIF is becoming to agents what CNCF became to cloud — the neutral, open home for the whole stack.

THE LINUX FOUNDATION | OPEN SOURCE SUMMIT NORTH AMERICA

Sources: AAIF Technical Charter v1.0 (Feb 2026) · AAIF Working Group launch announcements (Feb–May 2026) · Anthropic MCP spec · Google A2A spec · x402 Foundation announcement (Apr 2026)

But one layer of the stack is still NOT open — and it's the most important one

THE AI STACK — STATUS CHECK

AGENTS	OPEN
MODELS	OPEN
INFERENCE	OPEN
TRAINING	OPEN
INFRASTRUCTURE	OPEN
DATA	CLOSED

DATA — THE NEXT OPEN FRONTIER

Models will be commodities. Code is open. Data is the moat.

- The frontier labs aren't just winning because of better code
- They're winning because of proprietary training data and exclusive licensing deals
- Open weights without open data is half a movement
- **If the data layer stays closed, the value at the model layer stays closed too**

Every layer of compute is open. The last unfought layer is data — and the window to open it is closing fast.

Three walls are closing in on AI training data — fast

LEGAL CLOUDS

Copyright lawfare is reshaping what can be trained on

- NYT v. OpenAI
- Getty v. Stability still unresolved
- EU AI Act training-data disclosure rules
- Anthropic books settlement $1.5B

INFRASTRUCTURE WALLS

Scraping the open web is no longer the open web

- Cloudflare default-deny for AI crawlers (Jul 2025)
- Reddit, X, Quora behind paid API walls
- robots.txt + signed-agent enforcement spreading
- Common Crawl quality + coverage degrading
- Publishers pulling back on Internet Archive / Wayback Machine

TECHNO-NATIONALISM

Data is now a strategic asset

- EU data sovereignty + GDPR carve-outs
- US-China decoupling on training
- India, Japan, Korea pursuing sovereign AI stacks

Each wall, on its own, is fixable. All three at once is an real challenge for open AI. The window is now.

Sources: NYT v. OpenAI complaint (S.D.N.Y. 2023) · Bartz v. Anthropic settlement docs (Sep 2025) · Cloudflare blog "Default-deny for AI crawlers" (Jul 2025) · Reddit Inc. S-1 · EU AI Act final text (Mar 2024) · China DSL/PIPL official text

Every closed wave triggers an open one — and data is next

THE PATTERN — FOR FORTY YEARS, CLOSED HAS TRIGGERED OPEN

1990s

Closed UNIX → Linux

2000s

Closed databases (Oracle, DB2) → PostgreSQL · MySQL

2010s

Closed cloud stacks → Kubernetes · CNCF

Data is next. The action has already started — and so has the reaction.

THE ACTION — THE CLOSED LICENSING ARMS RACE

$1.5B	Anthropic settlement w/ authors (Bartz v. Anthropic, Sep 2025)
$250M	OpenAI + News Corp, 5-yr licensing deal
$200M+	Reddit data — combined deals with Google + OpenAI
~$39M	OpenAI + Axel Springer, 3-yr (~$13M/yr)

THE REACTION — ALREADY UNDERWAY AT THE LINUX FOUNDATION

CDLA (Community Data License Agreement)

Purpose-built license for sharing data for AI/ML training. Permissive 2.0 specifically addresses model training rights.

Overture Maps Foundation

Meta + Microsoft + AWS + TomTom pooling proprietary map data into open datasets. Hosted by LF — the template for cross-competitor data sharing.

Data is more widely held than code or models — every company owns some. The bottleneck is coordination, not supply.

THE LINUX FOUNDATION | OPEN SOURCE SUMMIT NORTH AMERICA

Sources: Bartz v. Anthropic court filings (Sep 2025, NPR/Axios) · News Corp + OpenAI joint announcement (May 2024) · Reuters "Reddit signs $60M data deal with Google" (Feb 2024) + OpenAI Reddit announcement · TechCrunch "OpenAI inks deal with Axel Springer" (Dec 2023) · linuxfoundation.org/press/CDLA + Overture Maps Foundation announcement (Dec 2022)

The way through — consortium for open data

> *“Foundation models will be infrastructure — they should be built by international consortia, not controlled by a handful of private companies.”*
>
> — the Yann LeCun thesis: open AI needs a CERN-style consortium for data and models

$700B

US data-center capex in 2026 alone

~$1–3B

to license the entire major-publisher corpus

5

publishers control >60% of English-language books

3+

frontier labs already paying for data 1-on-1 — duplicating cost

If we spend $470B on the metal, we can spend $5B on the substance. This is a money problem, not a technology problem.

THE LINUX FOUNDATION | OPEN SOURCE SUMMIT NORTH AMERICA

Sources: Dell'Oro AI Infrastructure Forecast (Apr 2026, US share) · Author's Guild industry analysis of publishing concentration · Yann LeCun, public remarks at NeurIPS 2024 + Meta AI blog (paraphrased) · Bartz v. Anthropic, OpenAI/News Corp + Reddit/Axel Springer deal disclosures

The cheat code goes both ways.

Open source is the cheat code for AI. AI is the cheat code for open source.

AI is the cheat code for open source — the dam just broke

> *"AI enables people who are not professional developers to contribute to open source for the first time."*
>
> — Rich Aberman, Open Core Ventures (April 2026)

36M

new GitHub developers in 12 months

fastest absolute growth in platform history

2×

monthly active OSS contributors

84K → 175K in one year

<1%

of the world used to build all software

that ceiling is now lifting

Source: GitHub Octoverse 2025 · GitHub blog · Open Core Ventures (Apr 2026)

GitHub projects have increased nearly twenty five-fold in just the past six months

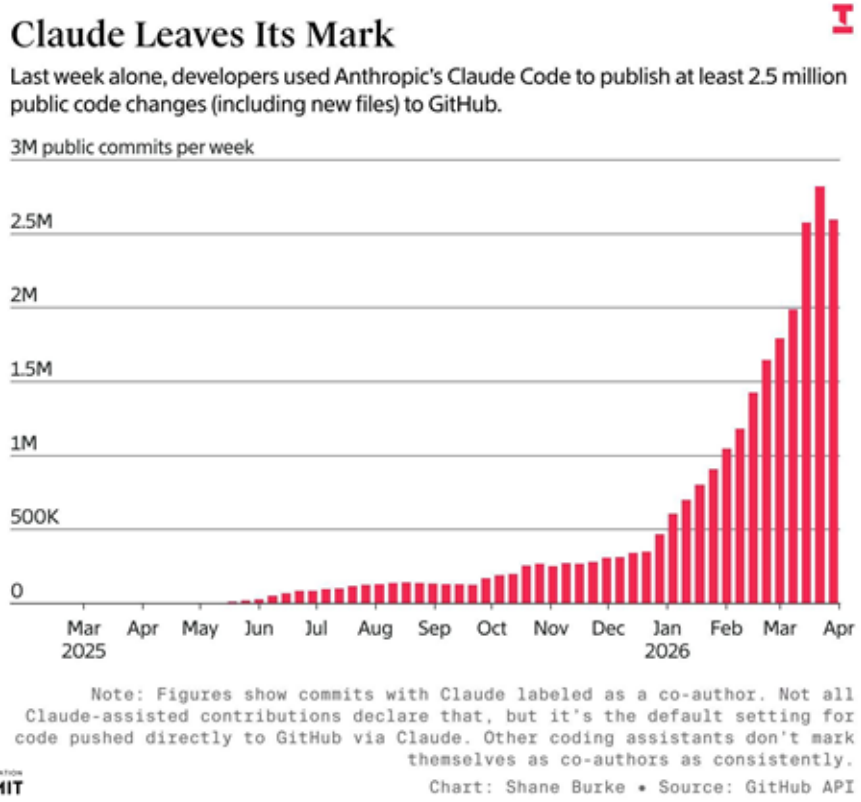

AI is already insanely useful — but the context still lives in people

FORGET THE AGI DEBATE

AI is transformative now — whether or not AGI ever shows up.

- **25%** of Google's new code is AI-generated
- **79%** faster time-to-market at Rakuten with Claude Code
- **470K** employees at Deloitte deploying Claude Enterprise
- **>50%** of Claude Code users at Epic are non-developers

WHAT AI STILL CAN'T DO

The context lives in people, not the model.

- Know which customer call shaped your roadmap
- Remember the regulator's unwritten preferences
- Sense the edge case that broke production in 2021
- Read the room in a partnership negotiation
- **Decide what's worth building in the first place**

Software isn't dead. The barrier to entry just dropped. The moat moved from code to context, judgment, and speed.

AI is a force multiplier on human context — not a replacement for it. The companies that win will pair both.

AI isn’t going to kill tech jobs - it is going to redistribute them - we are already seeing it

- Net technical hiring projected to **increase 31%** in 2026
- Entry-level IT roles expected to **grow by 8%**
- 97% of organizations are **actively investing in AI** implementation
- 57% report a **major gap in AI security** and risk management capacity
- Companies are **prioritizing upskilling over external hiring**
  - 7.9x advantage in business context
  - 5x advantage in total cost vs. hiring externally

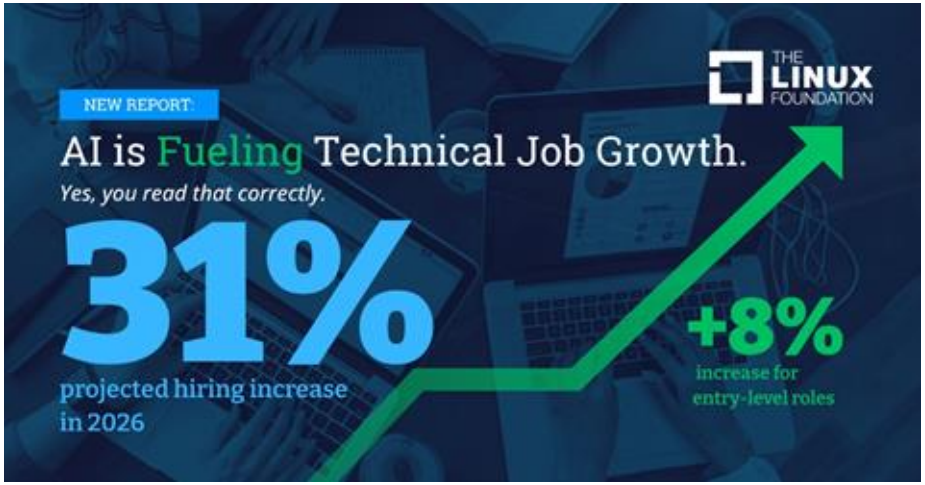

THE LINUX FOUNDATION | OPEN SOURCE SUMMIT NORTH AMERICA

AI coding is the next abstraction layer — and open source is the substrate

> *Every abstraction in computing lifted who could build. AI is just the next one — and it runs on us.*

1950s	→ Machine code	*A few hundred specialists*
1970s	→ C / Unix	*Tens of thousands of engineers*
2000s	→ Python · Ruby · JS	*Millions of web developers*
2010s	→ Stack Overflow + GitHub	*Anyone who could copy-paste*
2026	→ **AI agents conducting OSS**	**Anyone who can describe what they want**

The new abstraction doesn't replace open source — it makes OSS the standard library it uses.

The boring revolution — every “unsexy” vertical is now buildable

The next million apps are not new chatbots.

They are the boring, valuable, vertical workflows that were too expensive to custom-build — until now.

Cost of writing custom software → 0. Bottleneck moves from “can we build it?” to “what should we build?” — and only domain experts can answer that.

Healthcare

OpenEMR, FHIR, HAPI

Practices customize EHR workflows themselves

Legal

OpenAJ, Document automation

Firms compose their own matter-management

Construction

OpenProject, IFC, BIM

Site managers ship custom safety + RFP tools

Finance ops

Ledger CLI, Beancount

Controllers build their own close + audit

Education

Moodle, H5P, Open edX

Districts assemble bespoke LMS workflows

Public sector

OpenStreetMap, CKAN, GovStack

Agencies replace decade-old RFPs in weeks

Open-source primitives are the operating system. AI is the compiler. The application layer is finally unbundling.

Value accrues to the application layer — same pattern, again

WHERE VALUE LASTED

ERA	INFRASTRUCTURE (won early, commoditized)	APPLICATION LAYER (where the durable value landed)
INTERNET ERA	*Cisco · UUNET · Sun · Oracle*	**Google · Amazon · Uber · Airbnb · Netflix**
MOBILE ERA	*Qualcomm · Apple silicon · ARM*	**Instagram · WhatsApp · TikTok · Stripe**
AI ERA (2026 →???)	*NVIDIA · frontier model labs*	**the next decade — built on OSS + AI orchestration**

FRONTIER MODELS = ECONOMIES OF SCALE WITH SHRINKING MARGINS

Capex-heavy. Cost per token down 280× in 2 years. Open models hit parity in 3–6 months. Pattern matches commodity infrastructure — not durable software value.

Lasting value will sit at the application layer — open-source primitives, orchestrated by AI, owned by operators closest to the workflow.

Caveat: We are not in a bubble - the hardware layer will boom for quite some time as agents break the demand curve for hardware for years to come

Sometimes the bear eats you.

The cybersecurity ground just moved under our feet. Most of you don't yet know how badly.

Mean Time to Exploit, 2026: negative seven days

−7 DAYS

Exploitation now happens, on average, BEFORE a patch is even released. Source: Google M-Trends 2026.

MEAN TIME TO EXPLOIT, OVER TIME

2018	2022	2024	2026
63 days	32 days	5 days	−7 days

The traditional discover → disclose → patch → deploy cycle was designed for a slower adversary. That adversary no longer exists.

THE LINUX FOUNDATION | OPEN SOURCE SUMMIT NORTH AMERICA

Sources: Google Mandiant M-Trends 2026 (mean-time-to-exploit) · CERT-EU "AI is changing the economics of vulnerability discovery" (Apr 2026) · Verizon DBIR 2025

What one frontier model found in a few weeks — a generational jump

CLAUDE MYTHOS PREVIEW — APRIL 2026

Autonomously discovered thousands of high- and critical-severity zero-days across every major OS, browser, and core library — many of them undiscovered for decades.

1000s	zero-days found across major OSs and browsers	**27 yrs**	old: RCE in OpenBSD that survived three decades of audits
181 ×	working shell exploits vs. Firefox 147 JS engine (predecessor: 2)	**12 of 12**	CVEs in the Jan 2026 OpenSSL release found by autonomous agent (AISLE)

"I've found more bugs in the last couple of weeks than the rest of my life combined." — Nicholas Carlini, on using Mythos

THE LINUX FOUNDATION

Sources: Anthropic red.anthropic.com/2026/mythos-preview (Apr 2026) · Anthropic Project Glasswing launch · AISLE blog "12 of 12 OpenSSL CVEs" Jan 2026 · Nicholas Carlini (Google DeepMind) — public Mythos remarks Apr 2026

The bill is coming due for decades of underinvestment

> ***"Software was never designed for perfect security. That choice is now catching up with us."***
>
> — Thomas Dullien (Halvar Flake), on the Mythos disclosures, FAZ April 2026

THE DEBT

- Decades of "ship faster" trumping "ship safer"
- Latent vulns externalized onto users
- Auditing was too expensive — until now
- **AI just made auditing nearly free**

THE SILVER LINING

- Every vuln found and fixed is a closed door for that attack IF THE FIXES ARE ADOPTED
- Long-term: a fundamentally cleaner ecosystem
- Short-term: emergency patching at unprecedented tempo
- **We pay the bill faster than we'd like — but we pay it**

The signals that the bear is winning — for now

Attackers need ONE working exploit. Defenders must triage every finding and remediate it correctly.

The asymmetry within the asymmetry — and right now, the attackers are winning it.

JANUARY 2026

cURL shuts down its bug bounty

Daniel Stenberg: AI-generated reports are overwhelming the project. The cost of triaging hallucinated vulns exceeds the value of finding real ones.

APRIL 2026

HackerOne pauses Internet Bug Bounty

AI-generated submissions, mixed legitimate and low-quality, overwhelm both triage capacity and remediation across the open-source ecosystem.

ROLLING

FunkSec, DDoS-for-hire, low-skill actors

Cybercriminal groups with minimal technical proficiency briefly became the most prolific ransomware actors in the world — using AI-generated tooling.

The defenders' choice.

If AI can find vulnerabilities, it can fix them. The only thing scarce now is collective will.

Defensive AI is already working — at scale, at speed, at low cost

54 vulns

in 4 hours across 54M LOC

DARPA AIxCC, August 2025 — seven finalist teams, average $152 per finding. 18 real vulnerabilities responsibly disclosed.

72 fixes

into open source projects

Google DeepMind CodeMender has already autonomously identified and patched complex vulnerabilities in critical OSS.

92% recall

on seeded vulnerability benchmarks

OpenAI Aardvark / Codex Security — autonomous SecOps agent now reviewing production codebases at major enterprises.

FFmpeg+

patches authored by AI

Raptor — Halvar Flake's own framework. Same person who warned us is shipping patches with the same tools.

Sources: DARPA AIxCC Final results, darpa.mil Aug 2025 · Google DeepMind CodeMender blog (Mar 2026) · OpenAI Aardvark / Codex Security launch (openai.com) · Dark Reading "Raptor framework" (Apr 2026) · CERT-EU vulnerability discovery brief (Apr 2026)

The defender stack — and where the Linux Foundation stands

OpenSSF

The foundation for OSS security: Alpha-Omega, Scorecard, Best Practices Badge, GUAC. Funded by every major tech company. The center of gravity for community-led defense.

Sigstore

Cryptographic signing for software artifacts. Cosign + Rekor + Fulcio. Adopted by Kubernetes, npm, PyPI, and GitHub Actions. Makes "is this the code I think it is" a one-line check.

SLSA

Supply-chain Levels for Software Artifacts. The framework hardening build provenance end-to-end. v1.1 nearing graduation under OpenSSF.

SPDX + SBOM

ISO/IEC 5962 software bill of materials. The audit trail that lets defenders know what's in their software — and what to patch when news breaks.

OpenChain

ISO/IEC 5230 for license + supply chain compliance. Adopted across automotive, finance, government.

AI / model security

LF AI & Data + AAIF Security & Privacy working group — extending the same playbook to models, weights, and agent toolchains.

These exist. They work. The job is to deploy them universally — and to wire AI into all of them.

Coalitions of the willing — Project Glasswing and what it signals

Project Glasswing — launched April 7, 2026

50+ organizations — including the Linux Foundation — receiving exclusive defensive access to Mythos-class AI capabilities. $100M in usage credits, $4M to OSS security orgs. Defenders organized around the same technology the attackers are using.

THE COALITION

Anthropic	**Apple**	**AWS**	**Broadcom**
Cisco	**CrowdStrike**	**Google**	**JPMorgan**
Linux Foundation	**Microsoft**	**NVIDIA**	**Palo Alto Networks**

The signal that matters: industry is willing to organize defensively. The job is to widen this circle to maintainers, governments, and the long tail of critical OSS.

The only scarce resource is collective will

Whose side has more will?

Right now, the attackers do. That's the part we have to fix.

Attackers are organized, well-funded, and using AI today. Defenders are larger but fragmented. Fragmentation is the bug we can fix in this room.

WHAT ATTACKERS HAVE

- Frontier AI — directly or via open-weight
- No quality-control burden — one working exploit is enough
- Permission to act fast and break things

WHAT DEFENDERS HAVE — IF WE USE IT

- The same or better AI — Glasswing, Trusted Access, OSS models + new cyber-focused harnesses like MDASH
- Visibility into every system we run
- 100x the headcount — if we coordinate
- OpenSSF, Sigstore, SLSA, SBOM, Alpha Omega— already built

What I'm asking you to do — five commitments

01	**Adopt the defender stack — now**	Sigstore signing. Scorecard. SBOMs via SPDX. OpenSSF Best Practices Badge. If your project doesn't have these in the next 90 days, it's not because the tools aren't there.
02	**Wire AI into your secure SDLC**	AI-augmented SAST, DAST, dependency scanning, and patch generation in your CI. Use the same tooling the attackers are using — on the side of light. Treat it as table stakes by year-end.
03	**Fund the maintainers carrying the load**	OSS maintainers are drowning in AI-generated reports. If your company depends on a critical project, you owe that project funding, FTEs, or both. Alpha-Omega is one channel; there are many.
04	**Join AAIF, x402, OpenSSF, your CNCF SIG**	The standards being written in the next 12 months — for agent identity, agent payment, agent provenance — will shape the next decade. Be in the room or live with what gets written without you.
05	**Coordinate. Disclose. Share defensively.**	Glasswing-style coalitions. Cross-vendor disclosure. Threat intel sharing. Stop defending alone. We've always beaten closed systems by being more coordinated than they are.

So — did we eat the bear?

Yes. The entire AI stack runs on open source code.

Is the bear eating us?

Right now, yes. But we can recover.

> ***Which story wins isn't decided by the AI. It's decided by us - the open source community - over the next 18 months.***
>
> **Let's go eat the bear again.**

Thank you.

The Big Lebowski was written and directed by Joel Coen and Ethan Coen, the **Minnesota-born** Coen brothers, whose distinctive cinematic voice helped define a generation of American independent film.