



AI Looks Like Web3 in 2017–18: L1s Today, Apps Next

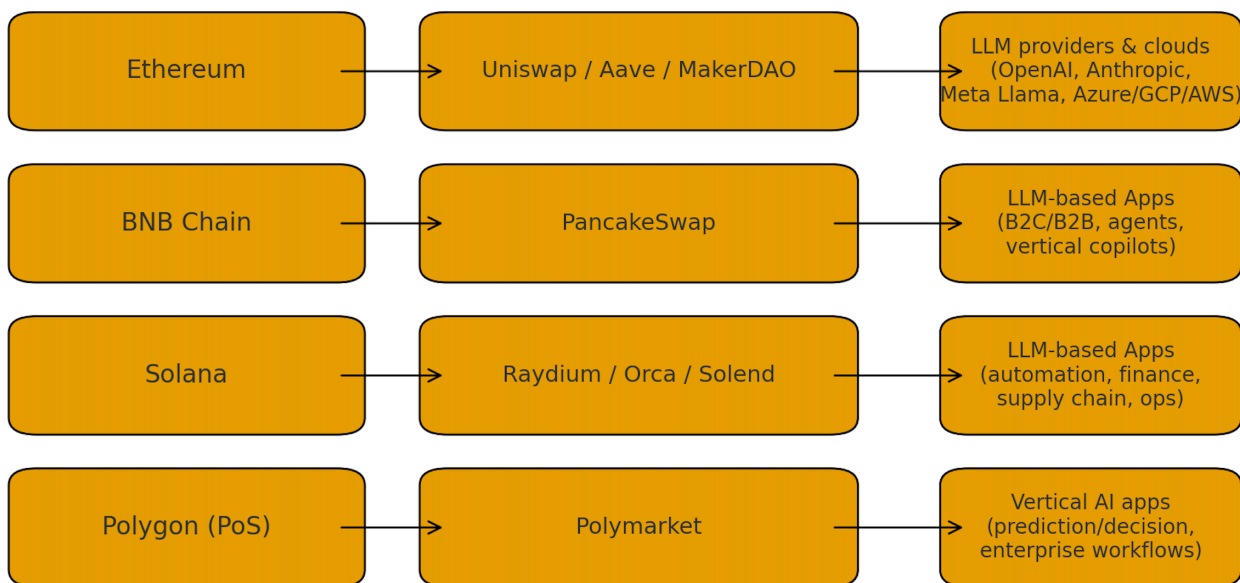
Executive summary

AI’s “foundation-model build-out” mirrors crypto’s L1 boom (2017–18): hyperscalers and frontier-model labs are spending unprecedented capex to lay rails, while API prices fall and credits proliferate. As in DeFi Summer (2020), the next outperformance window shifts to **applications**—vertical B2B/B2C products on top of standardized “L1-style” LLM platforms.

1) The analogy in one chart

In crypto, value first pooled in **L1 networks** (Ethereum, BNB Chain, Solana, Polygon). The breakout came from **apps** built on those rails: AMMs, lending, and prediction markets. In AI, **LLM providers + clouds** play the L1 role; the next leg belongs to app-layer companies (agents, copilots, vertical SaaS).

See Figure 1 — L1→Apps→AI map



Analogy: L1 buildout → App wave (DeFi → AI Apps)

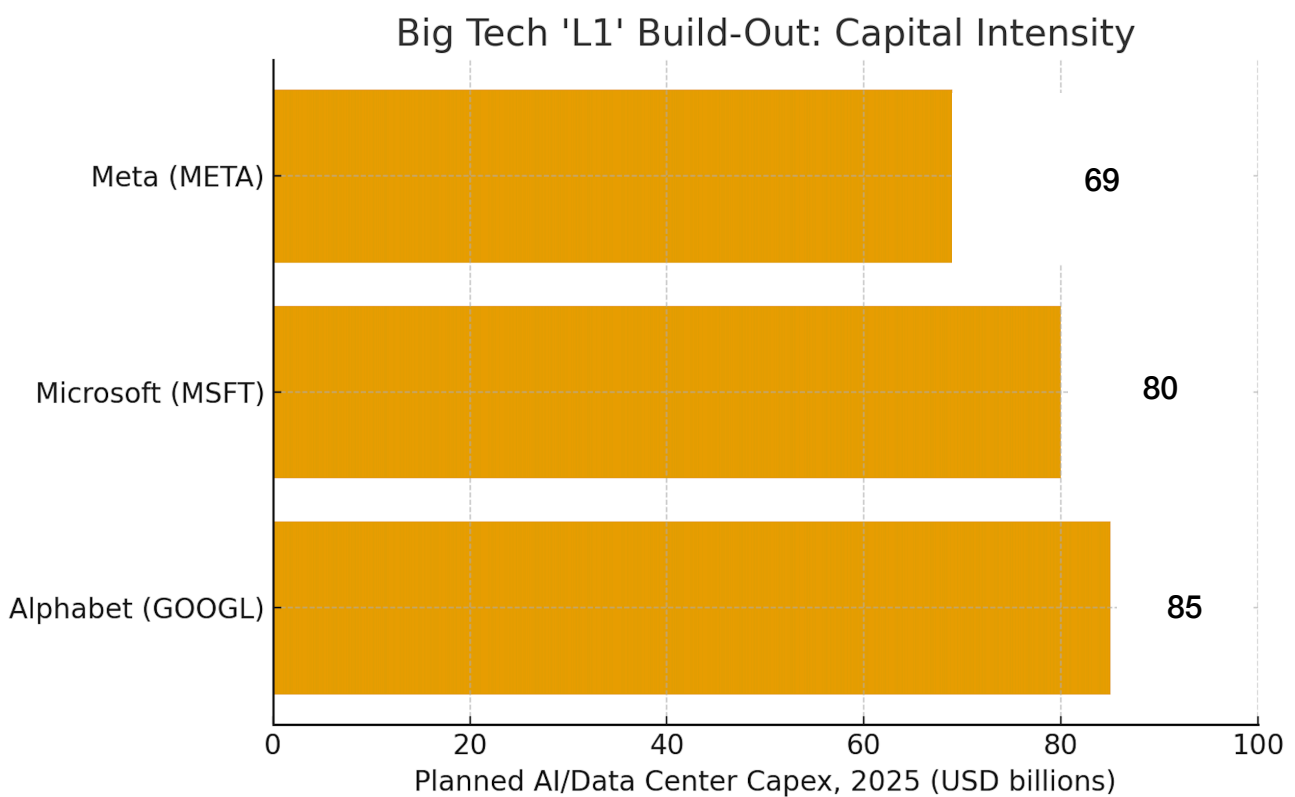
- Ethereum → Uniswap, Aave, MakerDAO.
- BNB Chain → PancakeSwap (leading DEX on BNB Smart Chain).
- Solana → Raydium (AMM), Orca (DEX/CLMM), Solend (lending).
- Polygon → Polymarket (prediction markets) on Polygon PoS.



2) Evidence we're in the AI "L1 build-out" phase

- **Capex shock.** Alphabet raised 2025 capex target to **~\$85B**; Microsoft's fiscal-2025 AI/datacenter spend is **~\$80B**; Meta's 2025 capex outlook is **\$66–72B**. This is the rails phase.
- **Monetization scale at the platform layer.** OpenAI reached **~\$12B annualized revenue** by July 2025, indicative of mature L1-like demand aggregation.
- **Open distribution baselines.** Meta's Llama 3/3.1 models are widely available under a community license (source-available, not OSI-open); availability still standardizes the "L1" for app developers.


See Figure 2 — 2025 AI capex (Alphabet, Microsoft, Meta).



3) The historical parallel: from L1 to killer apps (2017–22)

- **ICO/L1 boom.** EOS's 2018 ICO raised **~\$4.1B**—a symbol of capital concentration at base layers.
- **DeFi Summer (2020).** DeFi TVL jumped from **~\$1B** → **~\$15B** in months; Uniswap's **\$15.4B** September 2020 volume **surpassed Coinbase's \$13.6B** monthly volume. Apps captured the cycle.

See Figure 3 — timeline: ICO → DeFi → AI-L1 build-out.



4) Why the AI “app wave” is next (2025–28)

- **Costs are falling.** a16z documents “LLMflation”: **~1000x** drop in inference \$/token vs 2021; premium-tier costs down **~62x** since GPT-4 launched. Official pricing pages for OpenAI/Anthropic corroborate a competitive, declining trend with caching.
- **Standardized distribution.** Major LLMs (GPT-4.x, Claude, Llama) are packaged via hyperscalers and APIs—equivalent to L1 composability.
- **Resource pull for founders.** Clouds/Labs are subsidizing the app layer: AWS **\$230M** in AI credits; Google for Startups Cloud up to **\$350k** credits; Anthropic/LLM vendor startup programs provide API credits + higher rate limits.
- **Demand readiness.** Enterprise adoption is moving from pilots to workflow rewiring, per McKinsey 2025. App-layer companies can capture this with vertical integration.

5) What to build (pattern catalog)

B2B (vertical SaaS & agents):

- Finance ops, supply-chain planning, legal ops, healthcare RCM—where LLMs act (not just chat), integrate with systems, and produce measurable ROI; think DeFi’s “composables” but for enterprise workflows. Use retrieval + action models + governance.

B2C (assistants with actions):

- Booking, creative tooling, consumer finance—“agentic” surfaces with real transactions, logging, and preference memory. a16z/Sequoia lists show where usage clusters today.

Analogy anchors:

- **AMMs ↔ automation primitives:** routing, “quote → action” pipelines (Uniswap/Raydium → task/agent routers).
- **Lending ↔ policy-guarded workflows:** risk gating, limits, audit trails (Aave/Solend → enterprise policy engines).



- **Prediction markets ↔ decision-intelligence:** Polymarket’s market-based beliefs ↔ enterprise “decision scores.”

6) Risks & how this cycle differs from crypto

- **No token GTM by default.** Distribution favors product-market fit and integrations (SSO, SOC2, data-governance) over token incentives.
- **Platform dependency.** Manage API/price/limit risk with multi-provider support; keep a migration path to open weights (e.g., Llama family) and batch/offline modes.
- **Compliance & IP.** Respect license scopes (Llama is community-licensed, not OSI-open), data provenance, and sector regulations.

7) Founder playbook (90 days)

1. **Pick a vertical** with measurable unit economics (e.g., net time saved per seat).
2. **Design for action** (not chat): tie LLMs to systems of record; log actions; add guardrails.
3. **Benchmark costs** across providers; exploit prompt caching and small models for most calls.
4. **Stack the subsidies:** apply to AWS/Google/Anthropic/OpenAI startup programs; pre-negotiate rate limits.
5. **Ship a narrow, undeniable win;** expand via adjacent workflows; publish ROI.
6. **Keep optionality:** instrument a lightweight path to open-weights for latency/cost spikes.

Conclusion

Therefore, NextSpark Ventures is focused on backing exceptional teams building B2B and B2C applications on top of LLM “infrastructure.” We seek product-led founders who turn models into outcomes—agentic workflows, vertical copilots with system-of-record integrations, auditable retrieval, and real actions, not chat. Our playbook: multi-model architecture to reduce platform risk, aggressive cost control (caching, smaller models), and ROI proven in weeks. With inference costs falling and credits plentiful, distribution will reward fast executors; we partner to accelerate GTM, compliance, and enterprise security—capturing the application-layer alpha.