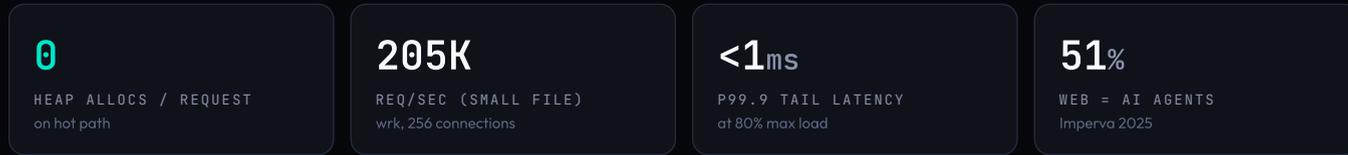


The HTTP Server Built for *the AI Agent Era*

Zero-allocation web infrastructure in Rust. When every microsecond costs tokens, SynapServe delivers.



01 THE PROBLEM

Today's servers see AI agents as attacks.

When an AI agent receives a task, it sends thousands of HTTP requests per second — hitting APIs, scraping data, coordinating with other agents. To nginx, Apache, and modern Rust frameworks, this looks indistinguishable from a DDoS attack. Agent traffic grew 1,300% in Jan–Aug 2025, and 51% of all web traffic is now non-human.



02 THE SOLUTION

HTTP infrastructure redesigned from byte zero.

<p>Zero-Allocation Hot Path</p> <p>Every parsed request lives on the stack. Span-based parsing replaces string copies. io_uring provided buffers mean the kernel manages memory. 0 allocations per request, verified by counting allocator.</p> <p>0 alloc/req</p>	<p>Agent-Native Protocols</p> <p>Built-in IETF Web Bot Auth (Signature-Agent headers), MCP Streamable HTTP, Google A2A discovery, and SSE token streaming. AI agents are the primary design target, not an afterthought.</p> <p>MCP + A2A + SSE</p>	<p>Predictable Tail Latency</p> <p>Thread-per-core architecture with io_uring. No Arc, no Mutex, no Send bounds. Each core owns its connections, buffers, and allocator. p99.9 stays flat under load — because there's nothing to spike it.</p> <p><1ms p99.9</p>
--	---	--

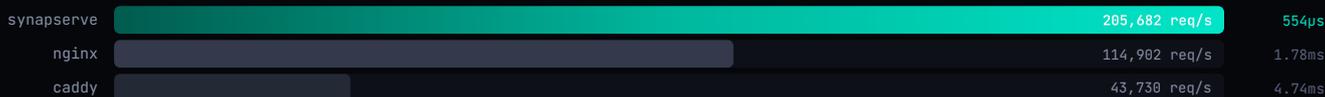
03 TECHNOLOGY

<p>Rust</p> <p>Memory safety without GC</p>	<p>io_uring</p> <p>Linux 6.1+ zero-copy I/O</p>	<p>kTLS</p> <p>Kernel-mode TLS 1.3</p>	<p>Thread-per-core</p> <p>Shared-nothing, no locks</p>	<p>SIMD</p> <p>SSE2/AVX2 header scan</p>
--	--	---	---	---

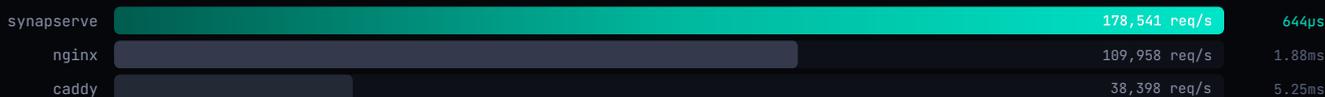
04 BENCHMARKS

STATIC FILE THROUGHPUT — wrk -t4 -c256 -d60s, 8 workers, localhost

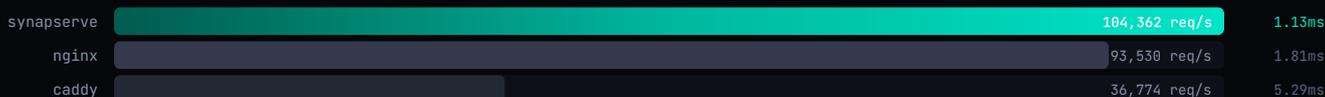
small.json 118 bytes



medium.json 4,005 bytes



large.json 24,203 bytes



Intel i7-8550U, 4C/8T, Linux 6.17.0. ETag & compression disabled. 10s warmup per server per file size. 14.5MB RSS under load.

Three tectonic shifts are converging.

AI agents are the new browsers

- ✓ 51% of web traffic is non-human
- ✓ OpenAI, Anthropic, Google all shipping agents
- ✓ Enterprise deploying autonomous workflows
- ✓ Agent traffic growing 1,300% in 8 months

Protocols are being standardized

- ✓ IETF Web Bot Auth (Signature-Agent)
- ✓ Anthropic MCP (Streamable HTTP)
- ✓ Google A2A (Agent-to-Agent)
- ✓ First-mover advantage is real

io_uring has matured

- ✓ Linux 6.1+ multishot accept
- ✓ Provided buffer rings (zero-copy recv)
- ✓ Zero-copy send (SEND_ZC)
- ✓ The OS is finally ready for this

The AI infrastructure market is **\$200B+ by 2030.**

API Gateway for Agent Traffic

Replace nginx/Envoy at edge. Native agent auth, protocol-aware routing, load balancing.

\$12B

TAM

LLM Inference Serving

Zero-copy SSE streaming for token delivery. Every byte copy costs real money at scale.

\$40B

TAM

Agent-to-Agent Infrastructure

MCP/A2A coordination infrastructure. Machine-to-machine as first-class.

\$150B+

TAM

Open Source Core (MIT / Apache 2.0)

- ✓ HTTP/1.1 + HTTP/2 + HTTP/3 server
- ✓ Zero-alloc parser & io_uring I/O
- ✓ TLS 1.3 (kTLS), static files, reverse proxy
- ✓ Community adoption flywheel

SynapServe Cloud (SaaS)

- ✓ Managed agent gateway (edge deploy)
- ✓ Agent auth + signature verification
- ✓ Advanced analytics & enterprise SLAs
- ✓ Per-request billing; recurring revenue

Playbook: Redis → Redis Labs, Kafka → Confluent, Nginx → Nginx Inc (acquired by F5 for \$670M)

- DONE** ● Zero-allocation HTTP/1.1 parser — Span-based, SIMD scanning, benchmarked at 14.6M req/sec
- DONE** ● io_uring I/O layer — Multishot accept, provided buffer rings, zero-copy send, splice
- DONE** ● HTTP/1.1 server framework — Radix-tree router, virtual hosts, static files, ETag/Range/Brotli
- DONE** ● Reverse proxy & upstream — Keepalive pool, weighted round-robin/least-conn/IP hash, health tracking, retry, backup servers, splice relay, DNS re-resolution
- DONE** ● TLS 1.3 with kTLS — rustls handshake + kernel encryption offload, parallel HTTP/HTTPS listeners, early data support
- Q2 2026** ● HTTP/2 + SSE streaming for LLM token delivery
- Q3 2026** ● Agent protocols: MCP, A2A, IETF Web Bot Auth
- Q4 2026** ● HTTP/3 (QUIC) + SynapServe Cloud managed offering

Seed round to build the HTTP layer for the agent web.

STAGE

Working
Product

HTTP/1.1 server, TLS 1.3 (kTLS),
reverse proxy — built & benchmarked

RAISING

Seed

18 months runway: team,
development, first deployments

USE OF FUNDS

Ship
& Scale

HTTP/2-3, agent protocols,
cloud offering, enterprise

- 55% Engineering
- 20% Cloud infra
- 15% Go-to-market
- 10% Ops & legal

invest@synapserve.io

synapserve.io | We respond to every serious inquiry within 24 hours