NPDOLLARS PROJECT

P-Time.
NP-dollars.

A Bedrock-smooth universal CNF endpoint. Every NP-complete problem, billable by tokens, decidable in wall-clock — or returned as a strictly easier formula with a proof trail.

MATHIEU HÉLIN
CEO @ Monce · Building the future of industrial commerce

linkedin.com/in/mathieu-helin · INSEAD

CHARLES DANA
AI/ML @ Monce · Making AI trustworthy · CPU-based AI

linkedin.com/in/charles-dana · npdollars lead

DASHA ZUYEVA
Building Monce · AI-native industrial commerce

linkedin.com/in/dashazuyeva · FH Oberösterreich

THE PAIN

SAT is the bottleneck.

Routing. Scheduling. Verification. Cryptanalysis. Glass-cutting optimization. All reduce to SAT. None are easy to run.

What “using a SAT solver” looks like today

- **Install, tune, babysit.** Every solver has its own flags, quirks, silent failure modes.
- **Guess a timeout.** Too short → no answer. Too long → stuck. Solvers return nothing on timeout.
- **Scale blindly.** No way to trade wall-clock time for dollars. One machine, hope it finishes.
- **Integrate from scratch.** DIMACS in, undocumented stdout out.

> **If your hard problem times out at hour 4, you walk away with nothing.**

NPDOLLARS V8

Wall-clock stays polynomial.
Cost scales in dollars.

ONE ENDPOINT

POST /cnf/solve

Any CNF, 10 bytes to 10 MB.

STABLE CONTRACT

4 outcomes

SAT · UNSAT · REDUCED · TIMEOUT

BAR-PASSING

budget + 3s

Every call returns inside its deadline.

Tier routing by CNF token count

TIER	ENGINE	TYPICAL WALL	LAMBDA COST
small (< 10 KB)	Kissat on EC2	**< 250 ms**	**$0**
medium (10 KB – 500 KB)	Kissat || v6 race	**100 ms – 5 s**	**~$0.000002**
large (500 KB – 50 MB)	v7 swarm + S3 offload	**10 – 40 s**	**~$0.003**
push: laptop	300 workers × 240 s	**up to 4 min**	**≤ $12 (capped)**
push: extreme	1000 workers × 70 s	**up to 70 s**	**≤ $12 (capped)**

Push mode enforces $12/call + $50/month ceilings before any Lambda fires. Ceilings persisted in S3.

THE CORE INNOVATION

Timeouts that hand back something useful.

Most SAT solvers return TIMEOUT with zero information. We return REDUCED — a strictly easier formula, in DIMACS, with a proof trail.

Example — BMC IBM #6 (industrial verification CNF)

	INPUT	AFTER 76 S, $0.03
Size	**6.9 MB**	**–55.2% clauses**
Backbones proven	**0**	**26 426**
Response body	**(nothing)**	**DIMACS + audit header**

> **$0.02 buys 1.28 M learned clauses.**
> Re-POST at lower budget — usually solves.

ENVELOPE: result · engine · tier · tokens_in/out · total_ms · cost_usd · assignment · partial_backbones · learned_2sat · cnf · extras

/CNF/SOLVE RETURNS THIS SHAPE EVERY TIME

SELF-DEMONSTRATING PRODUCT

Crowdsource SATLIB.

npdollars.aws.monce.ai/satlib serves 2 107 instances from 23 canonical categories. Anyone picks one, it solves, leaderboard updates. The page is the demo.

Live progress

1965 solved 97 reduced 25 remaining

INSTANCES TOTAL
2 107
across 23 categories

TOTAL TOKENS_IN
118 MB
of DIMACS

SPENT SO FAR
$0.085
for 98.8% done

COST TO FINISH THE ARCHIVE

$0.15 total

Push-mode on stubborn residuals: ≤ $12 per call. The game is self-funding.

LIVE AT /SATLIB

FOUR IDEAS, ONE PRODUCT

Three discoveries, one shipping API.

YEAR	WORK	WHAT IT GAVE NPDOLLARS
2023	**AUMA** A O(2^(n/2)) Universal Maximization Algorithm — MSc thesis, X-HEC / Ecole Polytechnique	Bridge: every function f : {0,1}^n -> R is a weighted MAXSAT instance.
2024	**Dana Theorem**	Any indicator function over a finite domain is encodable as SAT in polynomial time.
2026	**LogicSpace** logicspace.aws.monce.ai	DualTree + PolySAT: polynomial-time knowledge compilation for SAT.
2026	**npdollars v8** (this)	Tier router + bar-passing watchdog + REDUCED contract + SATLIB game.

STACK

FastAPI on EC2 t3.medium · Kissat 4.0.4 · AWS Lambda 10 GB (20 000 concurrent) · S3 state + daily archives · Route53 · systemd watchdog with cgroup memory caps. $37/month fixed infra.

WHERE IT GOES

From SATLIB demo to industrial reasoning API.

NEAR TERM

Finish SATLIB under $0.15. Publish the reduced-CNF archive as a public dataset. Package the monceai SDK on PyPI.

MEDIUM TERM

Karp-21 encoders at /encode/<problem>: graph coloring, TSP, scheduling -> CNF. Token-aware B2B pricing tier.

LONG TERM

Plug into Monce's industrial commerce stack. Commercial push-mode pool with SLAs.

ASK

Design partners for the B2B tier. Send us a CNF from your domain - we'll show you the REDUCED envelope.

TEAM · CONTACT

Mathieu Hélin
CEO @ Monce
linkedin.com/in/mathieu-helin

Charles Dana
AI/ML @ Monce · npdollars lead
linkedin.com/in/charles-dana

Dasha Zuyeva
Building Monce
linkedin.com/in/dashazuyeva