MOCSI

Trusted, local AI, starting in Maine

Draft, November 2026 · Powered by Public AI

AI has enormous potential to help Maine institutions reach their goals while fulfilling their mission. Unfortunately, AI is currently available only through national providers not held to the data privacy, environmental, and other standards Maine institutions require.

Institutions that serve Maine (towns, schools, libraries, nonprofits) need AI that meets privacy standards, that can be trusted and verified, and does not exploit the environment or local communities. National public interest nonprofits have the same need.

THE SOLUTION

Three projects that build trusted, local AI, starting in Maine.

MOCSI (Maine Open Compute Services Initiative) builds AI that institutions can inspect and trust, in three projects that each stand alone and reinforce one another. We are starting in Maine, and building it so other states can follow.

Software and Services

Maine-specific applications built with the institutions that use them, on open-weight models we name and version-pin.

Compute

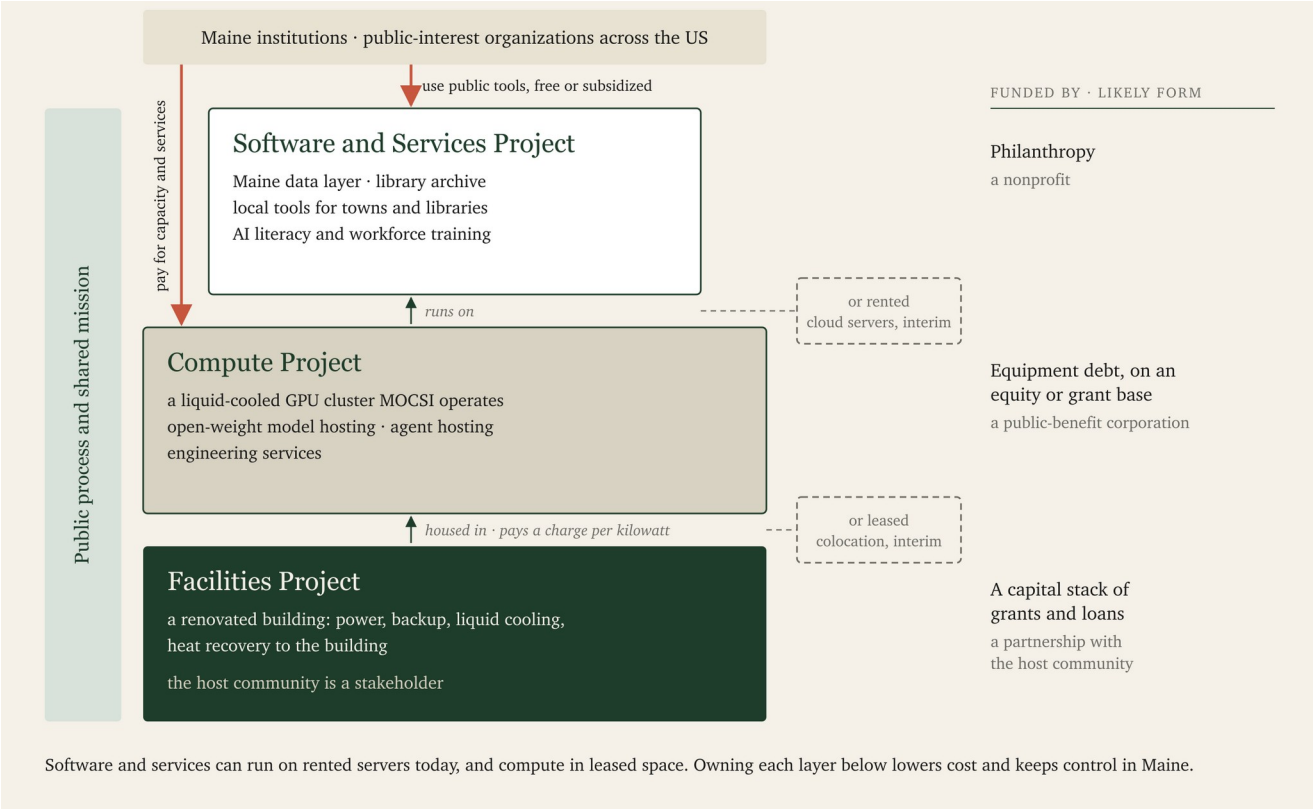
Used NVIDIA GPUs MOCSI operates in Maine, rented as reserved capacity and running the open models Software and Services offers.

Facilities

A retrofitted Maine building, with the host community as a stakeholder. Server heat warms the building, and Compute pays a charge per kW.

HOW IT FITS TOGETHER

Three projects, each independent yet reinforcing each other, make one stack.



Software and services can run on rented servers today, and compute in leased space. Owning each layer below lowers cost and keeps control in Maine.

WHAT WE ARE BUILDING FOR COMPUTE

A liquid-cooled NVIDIA GPU cluster that grows with demand.

64 → 96

GPUs at launch, growing to 96 by year 4

110 kW

peak IT load at full build-out

$1.42M

launch hardware: eight used 8-GPU NVIDIA servers, liquid cooled; each carries about 2,900 assistant users

78%

of server heat captured to liquid and reused to heat the host building

Year 0

8 servers · 64 GPUs

Launch with rental, AI assistant seats and first services projects.

Year 2

+3 servers · 88 GPUs

Add servers, funded from operating cash flow.

Year 4

+1 server · 96 GPUs

Fill the facility's 110 kW design.

PRODUCTS AND SERVICES

One GPU fleet, four ways to serve institutions.

Reserved GPU capacity

90% OF GPUS, YEAR 5

Dedicated NVIDIA GPUs on annual contracts for research and institutions that run their own workloads; spare hours sold on marketplaces.

AI assistant seats

10 GPUS, YEAR 5

Hosted open-weight chat and document assistants for staff, billed per seat per month. 200 seats in year 0, 3,500 by year 5.

Agent hosting

CPU SANDBOX SERVERS

Isolated, secure, audit-logged environments where institutions run multi-step agents safely. Billed per vCPU-hour.

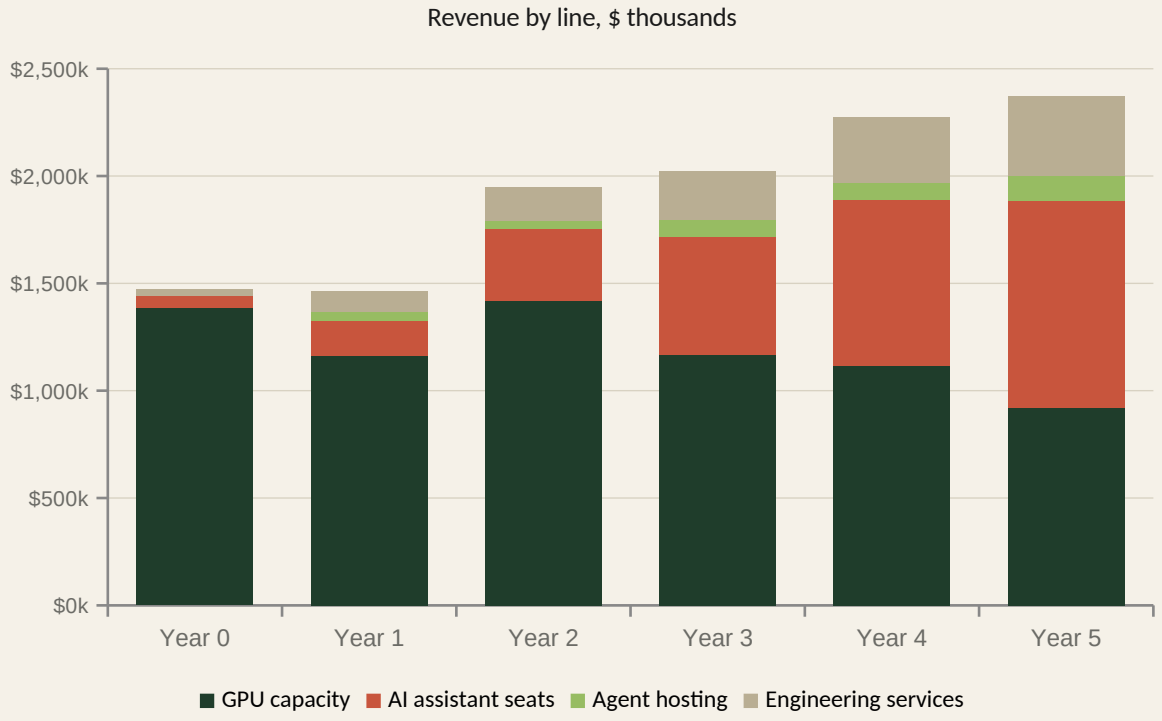
Engineering services

RAG AND INTEGRATION

Retrieval over an institution's own records (a records search, a permit assistant), delivered as fixed-price projects plus support retainers.

REVENUE MODEL

Services grow to two-thirds of revenue by year 5.



$11.55M

six-year revenue across all four lines

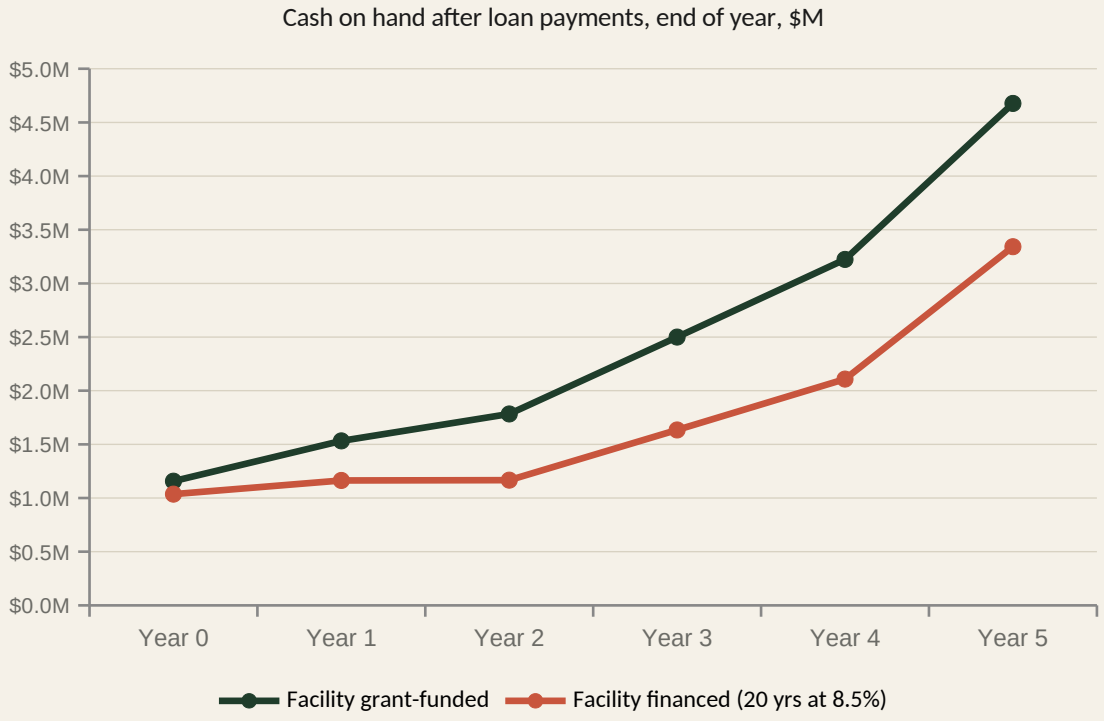
$2.37M

year-5 revenue, up from $1.47M at launch

GPU rental starts strong and eases as contract prices step down; seats, hosting and services build on it.

FINANCIAL OUTLOOK

A financed launch that keeps a cash reserve throughout.



$1.16M

lowest year-end cash after loan payments (facility grant-funded)

$4.68M

cash by year 5, after repaying the loan in full; $3.34M if the facility is financed

$2.15

of revenue for every $1 of operating and hardware cost over six years

THE SITE AND FACILITY

Facility upgrades are a separate project, with the host community as a stakeholder.

Upgrade	Mid estimate
Electrical service and distribution	$695k
Backup power (generator, switch, UPS)	$443k
Liquid cooling plant with heat recovery	$910k
Total (range $1.03M–$3.85M)	**$2.05M**

The upgrades improve the host's building and stay with it. The compute pays a facility charge per kW; how the build is funded sets that charge:

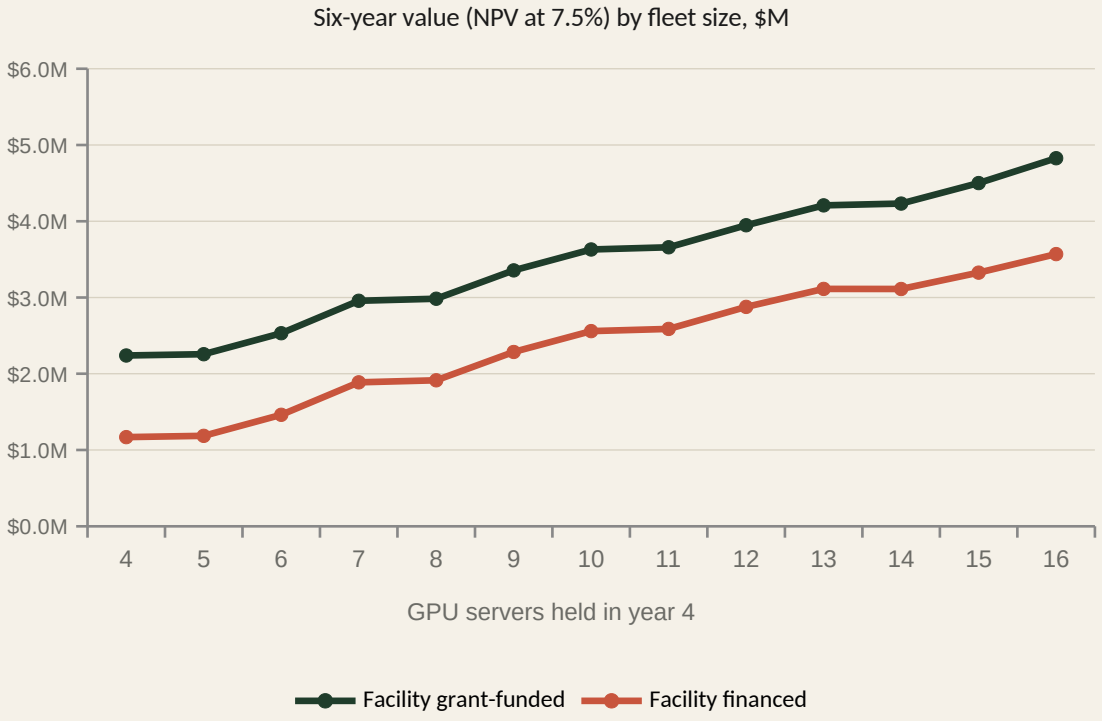
Grant share of facility	Charge / kW-yr	Cash, year 5, loan repaid
100%	$598	$4.68M
50%	$1,570	$4.01M
0% (loan, 20 yrs at 8.5%, guarantee assumed)	$2,541	$3.34M

The host site

A repurposed building in a Maine community that is a stakeholder in the facility. Cold climate cooling, and server heat reused to warm the building. Site to be announced with the host community.

RIGHT-SIZING THE FLEET

Twelve servers fill the facility. Fewer would still make money.



Positive at every size

Each fleet from 4 to 16 servers returns more than it costs, so size can follow demand.

12 servers fits the facility

Peak load 110 kW matches the 110 kW design; hardware $2.06M over six years (11 servers stays under $2M).

Demand sets the pace

Most growth comes from seats and services, which scale with customers more than with GPUs.

WHY IT WILL WORK

The pattern is already working elsewhere.

NEW YORK

Empire AI: Public-private funded AI compute, governed in the public interest.

CALIFORNIA

CalCompute: State framework for public AI infrastructure. Public access, public oversight.

MAINE & BEYOND

Town-owned broadband: Municipalities own the asset. Residents get affordable access. Already established here.

RURAL US

Cooperative ISPs: Member-owned internet infrastructure. Viable in rural and small-state contexts.

Maine has the conditions every other state will eventually face. A pilot here is the template for the next ten.

TRACTION

Site, partners, and funding pathway are taking shape.

SITE	STAKEHOLDERS	BUSINESS PLAN	FUNDING
Host community engaged	**Coalition forming**	**Model built**	**Pathway identified**
A host community with local ownership of its infrastructure is working with us on the first site.	State agencies, legislative leaders, local attorneys, and prospective public and enterprise customers are engaged.	Six-year financial model, facility engineering estimate, colocation comparison and market pricing research completed.	Co-Lab grant supports planning. Community-infrastructure financing and aligned philanthropy in active conversation.

THE TEAM



Matthew Victor

Legal & Partnerships

Tech and infrastructure attorney. Harvard fellow, advancing digital civic infrastructure.



Joshua Tan

Product & Strategy

Public AI. Has led multi-country public AI deployments with libraries, national labs, and public institutions.



Tom Law

Technical Strategy

Oak.AI. Built MOCSI's GPU capacity model and prospectus, and leads its technical analysis.



Fletcher Kittredge

Fractional CEO

Harvard-trained computer scientist with four decades building computing systems. Leads the business plan.

Also contributing: Ben Polsky (Stanford; data center build-outs). Advisors: Kerem Durdag (Indus Fund) and Richard Qualey (Bernstein Shur).

THE ASK

$2.14M

financing for the compute build: the launch fleet plus working capital.

Launch hardware: eight used, liquid-cooled 8-GPU NVIDIA servers (64 GPUs)	$1.42M
Working capital: the launch ramp, before customers pay in full	$0.36M
Working capital: additional cushion	$0.35M
Total financed	**$2.14M**

Repaid from operations

Five years at 8.5%, assuming a loan guarantee: $542k a year. Cash never drops below $1.16M at year-end.

Growth pays for itself

Year-2 and year-4 expansion ($0.61M) and agent-hosting servers come from operating cash flow.

Facility funded separately

Facility upgrades ($2.05M mid), with the host community, by grant, loan or a mix.

MOCSI

contact: matt@partnersindemocracy.us, josh@publicai.co, law@oak.ai, and fletcher@kittredge.com

APPENDIX

Key model assumptions (mid case).

Assumption	Value	Basis
GPU contract rate	about $2.53 per GPU-hour	Model rate; one-year reserved market about $2.30, on-demand median $3.34
Contract share of rental fleet	80%, remainder on marketplaces	Marketplaces at 50% utilization, 12% platform fee
GPUs serving AI assistant seats	2 GPUs at launch, 10 by year 5	Planning choice
Assistant seat price	$20 per seat per month	Hosted-assistant team plans run $15–50
Seats	200 (year 0) rising to 3,500 (year 5)	Estimate; to be firmed with letters of intent
Agent hosting	$0.05 per vCPU-hour, 30% used	Sandbox platforms charge $0.017–0.07
RAG projects	1 rising to 6 a year at $25k; half add a $1,500/month retainer	Small-organization builds run $10–60k
Electricity and cooling	$0.125/kWh; PUE 1.16	Liquid cooling with dry coolers in a cold climate
Hardware	Used NVIDIA datacenter GPUs, model chosen at purchase; 5–7 year life	Price decays 25% a year after introduction