

AI Cloud · Pricing

AI Cloud pricing

One seat price covers the workspace, two open-weight models and the security controls. Add-ons price support levels, compliance evidence and extra capacity.

Annual plan

\$20 per user a month
Billed annually, by card or invoice.

Monthly plan

\$25 per user a month
Billed monthly, by card.

5-seat minimum

No setup fee

No platform fee

10% off seat licenses at 1,000 seats or more

What the seat includes

Security and control

- ✓ Single sign-on through SAML or OIDC, SCIM provisioning, roles and audit logs, all kept in your account
- ✓ Controls documentation aligned to SOC 2
- ✓ Zero data retention at Clover Tech, and conversation retention set by your admins
- ✓ Spend controls: per-user metering, monthly and daily caps and a per-agent kill switch **In development**

Models

- ✓ Two open-weight models, plus Claude or GPT on your own vendor key with no Clover Tech fee
- ✓ Automatic routing of routine work to open-weight models and harder tasks to frontier models **In development**

Workspace

- ✓ The workspace, with projects that switch between models per chat
- ✓ Agents with instructions, file search and code execution
- ✓ Three RAG connections, each to one internal data source

Support and reporting

- ✓ Standard support, covering installation, production issues and model updates
- ✓ A usage dashboard that reports by team, with a minimum group size **In development**

Add-ons

Add-on	What it covers	Price
Products		
Clover Code	VS Code in the browser, connected to an open-weight coding model in your account	+\$5 per developer a month, coding model included
MCP server	Lets an AI tool deploy and run Clover Tech Marketplace apps in your account	\$1,000 a year; included with agent payments
Agent payments In development	USDC on Base and Solana, Bitcoin over Lightning, a dedicated BTCPay node	\$1 per 1,000 stablecoin payments, plus \$6 a month per BTCPay node
Support		
Standard	Email, 8 AM to 8 PM ET (5 AM to 5 PM PT), Monday to Friday; production issues answered the next business day	Included
Priority	Production issues answered within 4 business hours	+\$2 per seat a month, \$2,000 a year minimum
24 × 7	Production issues answered within 1 hour, around the clock	+\$3 per seat a month, \$10,000 a year minimum
Compliance evidence		
Security report	Quarterly deployment report and annual SOC 2 evidence pack, built in your account from audit logs	\$1,200 a year
Compliance mapping	Controls documentation for one framework beyond SOC 2, starting with GDPR	\$15,000 a year per framework
Model evaluation report In development	Cost per session for the same users before and after a model change	\$5,000 a year
Capacity		
RAG connection	Each internal data source beyond the three included	\$1,000 a year
Open-weight model	Each model beyond the two included. A new release of an included model runs beside it on trial at no charge.	\$1,000 a year
People		
Integration and workflow expert	Integrations and workflow design, which support does not cover	\$12,000 a year

Cloud compute

Compute is on your own cloud bill

Your cloud provider bills GPUs, storage and network to your account, and frontier model usage goes on your own vendor or cloud bill. One NVIDIA L40S on AWS serves both included models for up to about 50 users, an estimated \$12,300 to \$16,300 a year, so for teams under 50 users compute costs as much as the seats or more. Page 3 estimates compute for each example.

Per deployment

RAG connections, extra models, compliance mapping, the security report and the model evaluation report apply per deployment, so a second deployment, such as one in an EU region, repeats them. Add-ons bill on the same cycle as seats.

Every support level

Covers installation, production issues, model updates and replay of error logs, plus log forwarding to your MSSP. Your cloud provider runs the compute and network, so Clover Tech commits to response times and offers no uptime SLA. End-user support stays with your staff.

Build or buy

**\$1,500–
\$2,500**

a month of operations work

Running the same open-source stack yourself adds an estimated \$1,500 to \$2,500 a month of operations work (Navya AI, Cost to host a private LLM). That is the work the maintained install takes on: patches, tested upgrades and controls documentation.

Example annual bills

Five configurations on annual billing, with the compute each would need on AWS.

Buyer and seats	Add-ons	Clover Tech per year	Estimated compute per year
Family office, 12 users Seats \$2,880	Security report \$1,200	\$4,080	\$12,300–\$16,300 1 L40S
InfoSec team, 20 users Seats \$4,800	24 × 7 support \$10,000 (minimum), 3 extra RAG connections \$3,000, 1 extra open-weight model \$1,000, security report \$1,200	\$20,000	\$24,700–\$32,600 2 L40S
Trading team, 25 users Seats \$6,000	Agent payments: 240,000 payments \$240, one BTCPay node \$72; MCP server included	\$6,312	\$12,300–\$16,300 1 L40S
DevOps team, 40 users Seats \$9,600	Clover Code for 40 developers \$2,400, Priority support \$2,000 (minimum), 2 extra RAG connections \$2,000	\$16,000	\$24,700–\$32,600 2 L40S
Public company, 3,000 users Seats \$648,000 after 10% off	24 × 7 support \$108,000, integration and workflow expert \$12,000, compliance mapping for GDPR \$15,000, security report \$1,200, model evaluation report \$5,000, 5 extra RAG connections \$5,000, Clover Code for 300 developers \$18,000	\$812,200	About \$482,000 8 H100

Counts and add-ons are examples, and seats are at the annual rate. Compute estimates use AWS us-east-1 GPU prices, from a one-year savings plan to on-demand, sized for two Qwen3.6 models with 1 in 10 to 1 in 20 users active at once: one NVIDIA L40S for up to about 50 users, two for coding or three-model teams, eight H100s for 3,000 users. A load test on your own usage sets the final size. Frontier model usage and blockchain network fees are billed separately.

Next step

Request a quote for your firm

Send your seat count, your cloud provider and the add-ons you are weighing. We reply by email with a quote.

[Request a quote](#)

hello@clover-tech.ai