FOR BUYERS OF AI AND DATA-CENTRE INFRASTRUCTURE

One provider for the AI hardware, cloud and models you need.

Embedica is a provider of AI and data-centre solutions for customers in the UK, US, EU and wider Europe. You tell us what you want to run. We size it, source it from vetted suppliers and give you one all-in quote, with Embedica as your single point of contact.

What we help you buy

Servers, chips and networking
GPU servers, accelerators, switches and racks.

Memory and storage
High-bandwidth memory, NVMe and data platforms.

GPU cloud
Rented GPU capacity from specialist and large providers.

Models, apps and agents
Model access, applications and agent platforms.

How it works

1. **Create a free account** and confirm your email. You can ask questions straight away, and every question and answer is kept in your account history.
2. **Tell us the workload.** Use the Assessment or Sizing Wizard, or simply describe what you want to run, your timeline and your budget range.
3. **We vet the brief.** We check it, screen for export-control and sanctions issues, and only then pass the requirement to suitable suppliers.
4. **You compare quotes.** We put the replies side by side, flag the differences that matter, and you decide. Commercial terms are agreed in writing before any order.

Ballpark prices, November 2026

Item	Typical price (USD)
NVIDIA H100 80GB GPU	$25,000 to $40,000
NVIDIA B200 GPU	$30,000 to $50,000
8-GPU H100 server (DGX class)	$250,000 to $400,000
8-GPU B200 server (DGX class)	$300,000 to $520,000
GB200 NVL72 rack (72 GPUs)	$2 million to $3 million

Figures are rough market ranges from secondary price trackers, not quotes. Configuration, volume, supply and export rules change them. Full tables are in the AI Price Guide on our website.

Cloud and model prices, rough guide

GPU cloud, per GPU-hour	Big cloud providers	GPU cloud specialists
H100	$6.88 to $12.29	$4.39 to $6.16
H200	$7.91 to $10.60	$4.47 to $6.31
B200	About $14.24	$7.69 to $8.60

Model tier	Per million tokens (input / output)
Flagship	$2 to $10 in, $6 to $50 out
Mid-tier	$1 to $2 in, $5 to $10 out
Small and fast	$0.10 to $0.30 in, $0.40 to $1.50 out

Source: public price lists and secondary trackers, checked 18 November 2026. Prices move often, so ask us for a current quote.

Why buyers use Embedica

Not tied to one brand
We work across many manufacturers, clouds and models, so we propose what fits your workload.

Vetted suppliers
Every supplier goes through company, product and reference checks before they see your brief.

Export-control aware
High-end accelerators are controlled. We screen requirements before they are shared.

A record you keep
Your account stores every question, answer and quote so nothing is lost between calls and emails.

Industries we support

Financial services · Healthcare and life sciences · Public sector · Manufacturing · Energy · Telecoms · Retail · Research and education

Good to know

- You receive one all-in quote from Embedica, covering supplier prices, our services and delivery. A Request for Quote carries no obligation.
- Prices in this brief are guides only. Nothing here is an offer or a financial recommendation.
- We keep your details private and share a brief with suppliers only after our review.

Start with a Request for Quote

Create a free customer account at **embedica.ai**, or write to us and we will set it up with you.

Email info@embedica.ai

embedica.ai/customer-sign-up