



Spot and on-demand GPU hours land builders on a Kubernetes-native AI cloud; reserved capacity (up to 60% off) locks in scale; multi-year take-or-pay contracts with AI labs, enterprises, and hyperscalers drive backlog; storage, networking, and Mission Control attach around the same clusters.

KEY PARTNERS

- **NVIDIA** — GPUs, Exemplar Cloud (GB200 NVL72), collaboration / investment news (FACT — IR + company news).
- **Data-center and power providers for MW-scale AI factories (FACT — IR capacity).**
- **AI labs & enterprises: e.g. Cognition, Cursor, Midjourney, Runway, CrowdStrike (FACT — IR).**
- **Hyperscaler cloud customers (expanded relationships — FACT IR; Microsoft/OpenAI historically disclosed).**
- **Capital / asset-level financing secured by take-or-pay contracts (SECONDARY — investor materials).**
- **Platform M&A: OpenPipe, Marimo, Monolith (FACT — FY25 IR highlights).**

KEY ACTIVITIES

- Deploy and operate NVIDIA GPU / CPU clusters at MW scale.
- Run Kubernetes-native cloud + Mission Control observability.
- Price and fulfill spot, on-demand, and reserved capacity.
- Negotiate multi-year take-or-pay enterprise / hyperscaler deals.
- Attach storage, networking, managed software, Federal.
- Expand active and contracted power (IR capacity milestones).

KEY RESOURCES

- NVIDIA GPU fleets (H100, H200, B200, GB200 NVL72, etc.).
- **>850 MW active power · ~3.1 GW contracted (YE2025 FACT — IR).**
- Mission Control + Kubernetes-native platform software.
- AI Object Storage, networking, managed services IP.
- **Revenue backlog \$66.8B (FACT — IR YE2025).**
- Public company capital access (Nasdaq CRVV).

VALUE PROPOSITIONS

- Purpose-built AI cloud on latest NVIDIA GPUs.
- Kubernetes-native DX + Mission Control reliability.
- Spot / on-demand land → reserved expand.
- Free egress + AI Object Storage for data gravity relief.
- Take-or-pay scale for labs and hyperscalers.
- Federal / managed software for regulated workloads.

CUSTOMER RELATIONSHIPS

- Self-serve Login + published pricing for on-demand/spot.
- Contact sales for reserved (up to 60% off) and custom.
- Solutions / field engineering for cluster design.
- 24/7 support claims for production AI workloads (homepage).
- Named pioneer accounts as reference relationships (IR).
- **Expand path — hours → reserved → take-or-pay renewals.**

CHANNELS

- **Primary — coreweave.com + /pricing rate card.**
- Login / Contact us conversion paths.
- Direct enterprise & hyperscaler sales.
- NVIDIA / ISV partner co-sell and announcements.
- IR + newsroom for procurement credibility.
- Docs, status, Mission Control for operators.

CUSTOMER SEGMENTS

- AI labs / model companies (training + inference at scale).
- **AI-native startups & platforms (Cursor, Midjourney, Runway — IR).**
- **Enterprises (e.g. CrowdStrike partnership — IR).**
- Hyperscalers needing supplemental GPU capacity (IR).
- VFX / rendering and Physical AI workloads (site solutions).
- Government / CoreWeave Federal buyers.

COST STRUCTURE

- GPU / server capex and PP&E depreciation (IR).
- Power, colocation, and DC leases.
- Technology & infrastructure opex (\$2.929B FY2025 FACT).
- Cost of revenue / delivery (\$1.453B FY2025 FACT).
- Interest on asset-level and parent financing.
- S&M + G&A (smaller vs infra; FACT IR lines).

REVENUE STREAMS

- **Spot GPU hours —** discounted, interruptible (FACT rates).
- **On-demand PAYG —** e.g. H100 8× \$49.24/hr; B200 \$68.80 (FACT).
- **Inference single-GPU lane —** e.g. H100 \$6.16/hr (FACT).
- **Reserved / committed —** up to 60% off on-demand (FACT).
- **Multi-year take-or-pay —** backlog \$66.8B YE2025 (FACT — IR).
- Storage (Hot \$0.06/GB-mo...) + networking + managed attach (FACT).
- **FY2025 revenue \$5.131B (FACT — IR).**

Sources: investors.coreweave.com/news/news-details · coreweave.com/news/nvidia-and-coreweave-strengthen-collaboration-to-accelerate-buildout-of-ai-factories · coreweave.com/news/coreweave-announces-agreement-with-openai-to-deliver-ai-infrastructure · coreweave.com · coreweave.com/pricing
Not affiliated with CoreWeave.