AI chips & racks · Product benchmark

Disclosed power distributions: chip / module W and cabinet kW stay separate; each TDP, TGP, TBP and design-power boundary is labeled.

NVIDIA · AMD · Microsoft · Meta · Intel

Chips / accelerators / modules · W

Cabinets · kW (separate boundaries)

Designed rack power

GB200 NVL72
● 120 kW [A3/A8]

GB300 NVL72
● 135 kW [A4/A8]

Cabinet TDP · reference design

Vera Rubin NVL72 (DSX)
◆ 330 kW [A5/A9]

MGX Gen1.1 (DSX)
◆ 198 kW [A9]

GB200 / GB300: Mission Control 2.3.0 designed rack power. MGX / Rubin: DSX reference cabinet TDP, 19 Aug 2026. No common-boundary trend is inferred; neither is measured load or maximum AC input.

Google, AWS and Huawei lack verified absolute power; missing is not zero. Rubin per-GPU, Helios rack and Crescent Island power remain unverified; see notes.

Filled points = disclosed specifications; hollow points = official planned specifications. Meta uses Module TDP, AMD uses OAM TBP and Maia uses SoC TDP. This is a power distribution, not W/cm² heat flux, measured average load or product efficiency.