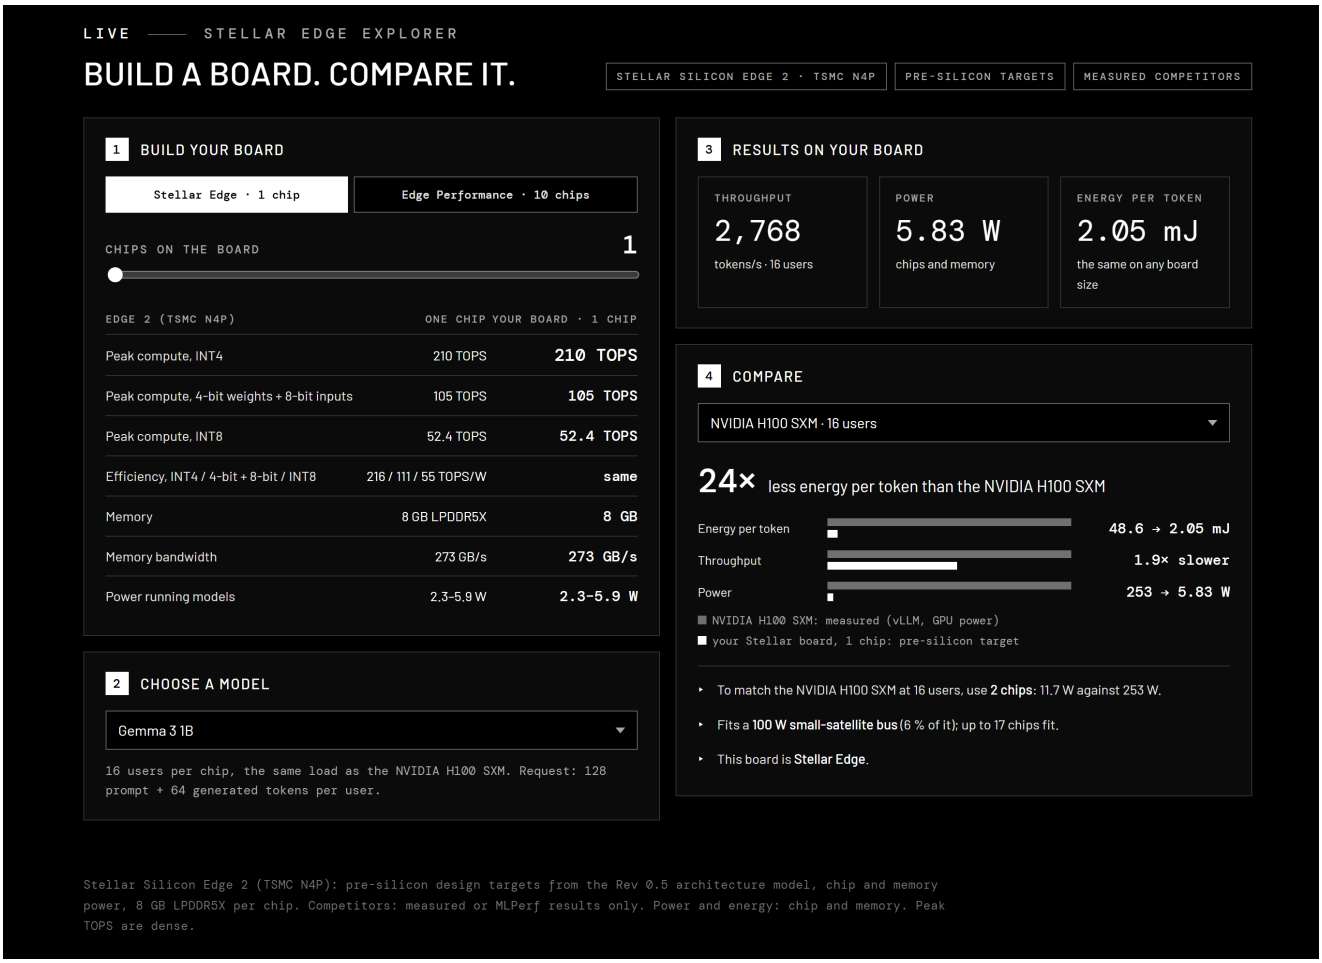
How to use the Stellar Edge explorer

The explorer on **stellarsilicon.space/testbed** lets you build a board from Stellar Silicon Edge 2 chips (TSMC N4P), run a language or vision model on it, and compare the result with measured hardware from NVIDIA, Hailo, Axelera and SiMa.ai. It works in four steps, numbered on the page.



The explorer with one chip, Gemma 3 1B and the NVIDIA H100 SXM at 16 users selected.

1 · Build your board

Pick a preset or drag the slider to choose how many chips are on the board (1 to 16):

- **Stellar Edge**: 1 chip, for satellites and edge devices.
- **Stellar Edge Performance**: 10 chips, for orbital data centres.

The table shows each specification for one chip and for your whole board:

Row	Meaning
Peak compute (INT4, 4-bit weights + 8-bit inputs, INT8)	Dense TOPS at each precision; multiplies with the chip count.
Efficiency	TOPS per watt at each precision; the same on any board size.
Memory and bandwidth	8 GB LPDDR5X at 273 GB/s per chip; adds up across chips.
Power running models	2.3 to 5.9 W per chip, chip and memory, across the models in the explorer.

2 · Choose a model

- **Language models:** Gemma 3 270M, Gemma 3 1B. Each user sends a request of 128 prompt + 64 generated tokens.
- **Vision models:** ResNet-50, YOLO11s (640).

The note under the menu says how the chip is loaded. A language model runs at the same number of users per chip as the competitor you compare with, or at 16 users per chip when there is none. A vision model runs one image at a time (a camera stream) or in its best batch, again matching the competitor.

3 · Results on your board

Number	What it tells you
Throughput	Tokens or images per second for the whole board; doubles when you double the chips.
Power	Watts for the chips and their memory on your board.
Energy per token / image	Millijoules for each token or image. It does not change with board size, so it is the fairest efficiency number.

4 · Compare

Choose a competitor from the menu. Only results that were measured, or published under MLPerf, are listed; there are no vendor claims or estimates. The explorer then shows:

- **The headline:** how many times less energy per token or image your board uses.
- **Three bars:** energy per token or image, throughput and power, competitor (grey) against your board (white).
- **Recommendations:** how many chips match the competitor, whether the board fits a 100 W small-satellite power budget, and which product the board corresponds to.

Worked example

One chip running Gemma 3 1B at 16 users against the NVIDIA H100 SXM at 16 users, the same load on both:

	NVIDIA H100 SXM	Stellar Edge (1 chip)	Result
Energy per token	48.6 mJ	2.05 mJ	24× less energy
Throughput	5,209 tokens/s	2,768 tokens/s	1.9× slower
Power	253 W	5.83 W	

The explorer therefore recommends **2 chips** to match the NVIDIA H100 SXM: 11.7 W against 253 W.

Good to know

- **Stellar numbers are pre-silicon design targets** from the Rev 0.5 architecture model. Competitor numbers are measured hardware or MLPerf results.
- **Power and energy** for Stellar count the chip and its memory; for competitors, what the cited measurement reports (for the H100, GPU power). Peak TOPS are dense (no sparsity).
- Ratios are rounded to whole numbers from 20× upward.

Want to run your own model on the chip? Write to **aditya@stellarsilicon.space** or book a call from the Contact page.