What can the optimiser actually discover?

Which rules exist?	How can they fit together?	What can it afford to test?

The rules decide what PetaBricks can choose

Sorting transform

- Quicksort
- Merge sort
- Insertion sort

The tuner can't pick radix sort.

Could pruning early cost us a better result?

Approach	Small task 1,000 items	Larger parallel run 100,000 items
Quicksort-based	**0.8 ms**	18 ms
Merge-sort-based	1.0 ms	**12 ms**

Best on a small task may not be best in a larger one.

Do more choices mean a better result?

Smaller search space

Best possible	Choices tested
10 ms	100 / 100

10 ms

Larger search space

Best possible	Choices tested
5 ms	100 / 1M

13 ms

How does PetaBricks keep the search manageable?

01 Start small

02 Measure and keep

03 Build and repeat

What does tuning cost?

2.14×

faster on 8 cores after tuning for 8 cores rather than 1 core

2.35×

slowdown on Xeon when using a sort setup tuned for Niagara

How long did the tuner take?