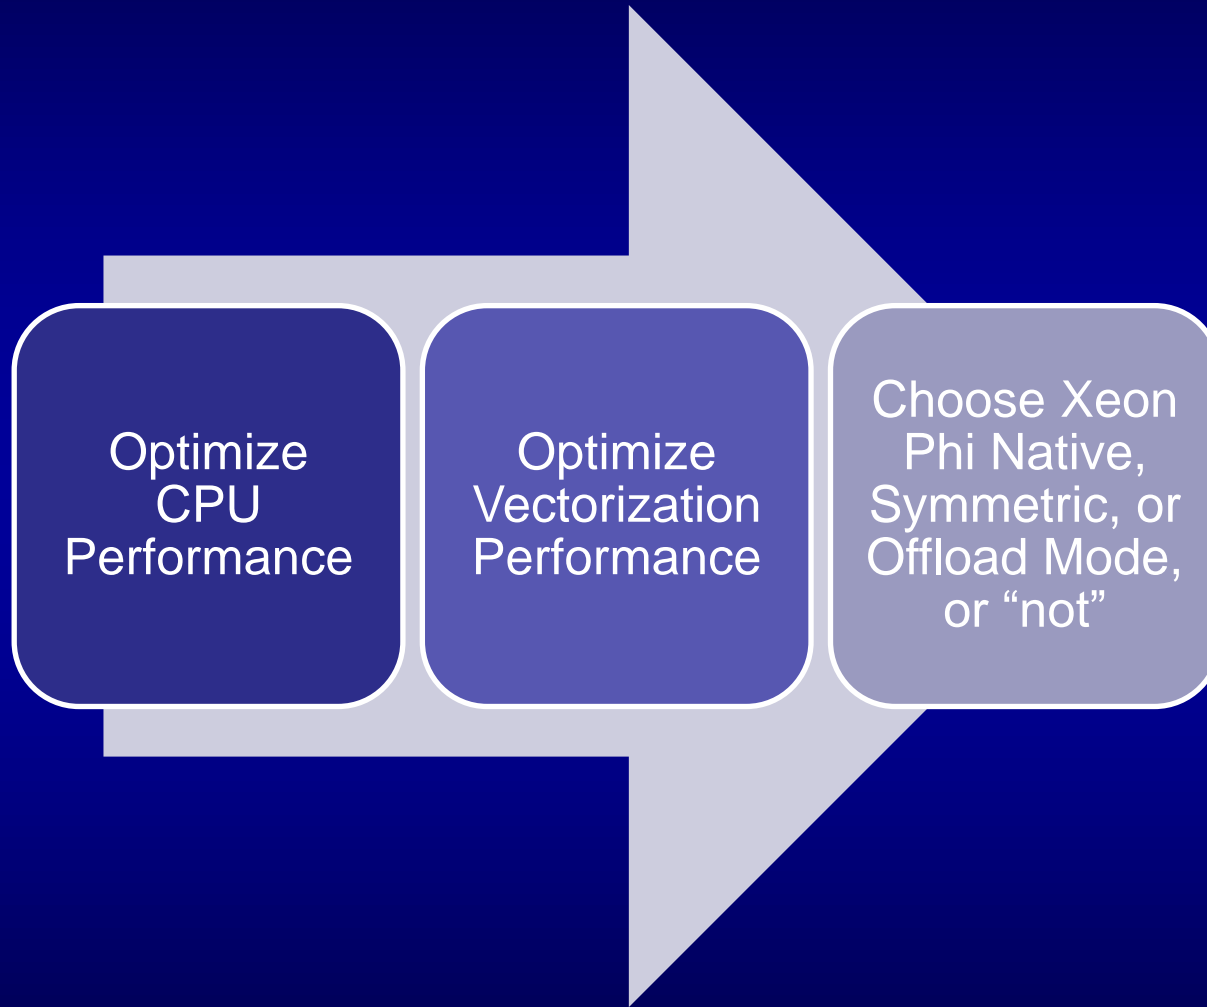
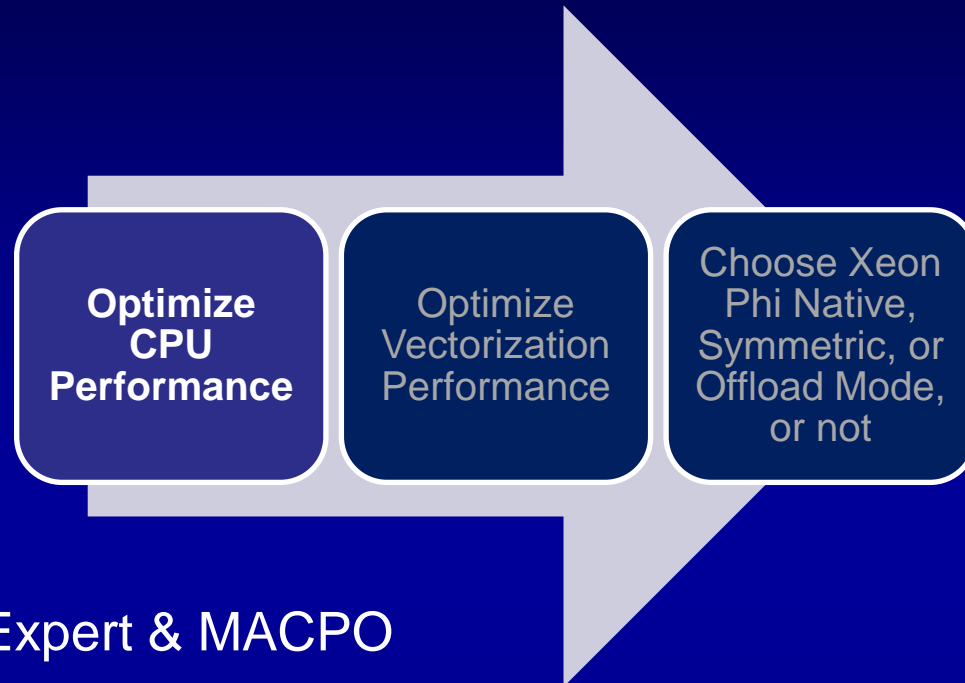


Xeon Phi Recommendor Process - 1

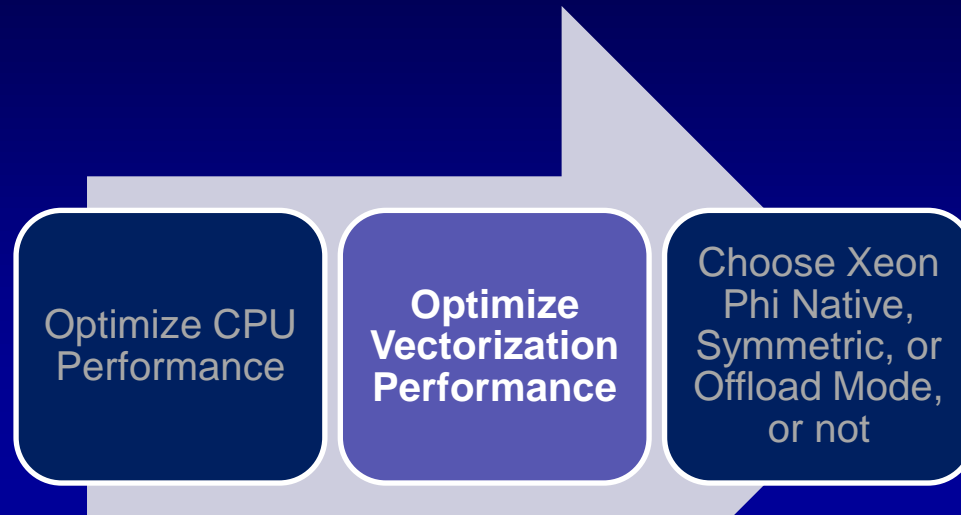


Xeon Phi Recommendor Process - 2



- Use PerfExpert & MACPO
- Focus on vectorization, scaling, & finding optimal task/thread count per node
- Determine execution time of each function at optimal task/thread count

Xeon Phi Recommendor Process - 3



- Assess Degree of Vectorization
 - Use Intel compiler vectorization and optimization reports (<http://software.intel.com/en-us/articles/vecanalysis-python-script-for-annotating-intel-compiler-vectorization-report>)
- Optimize Vector Processing

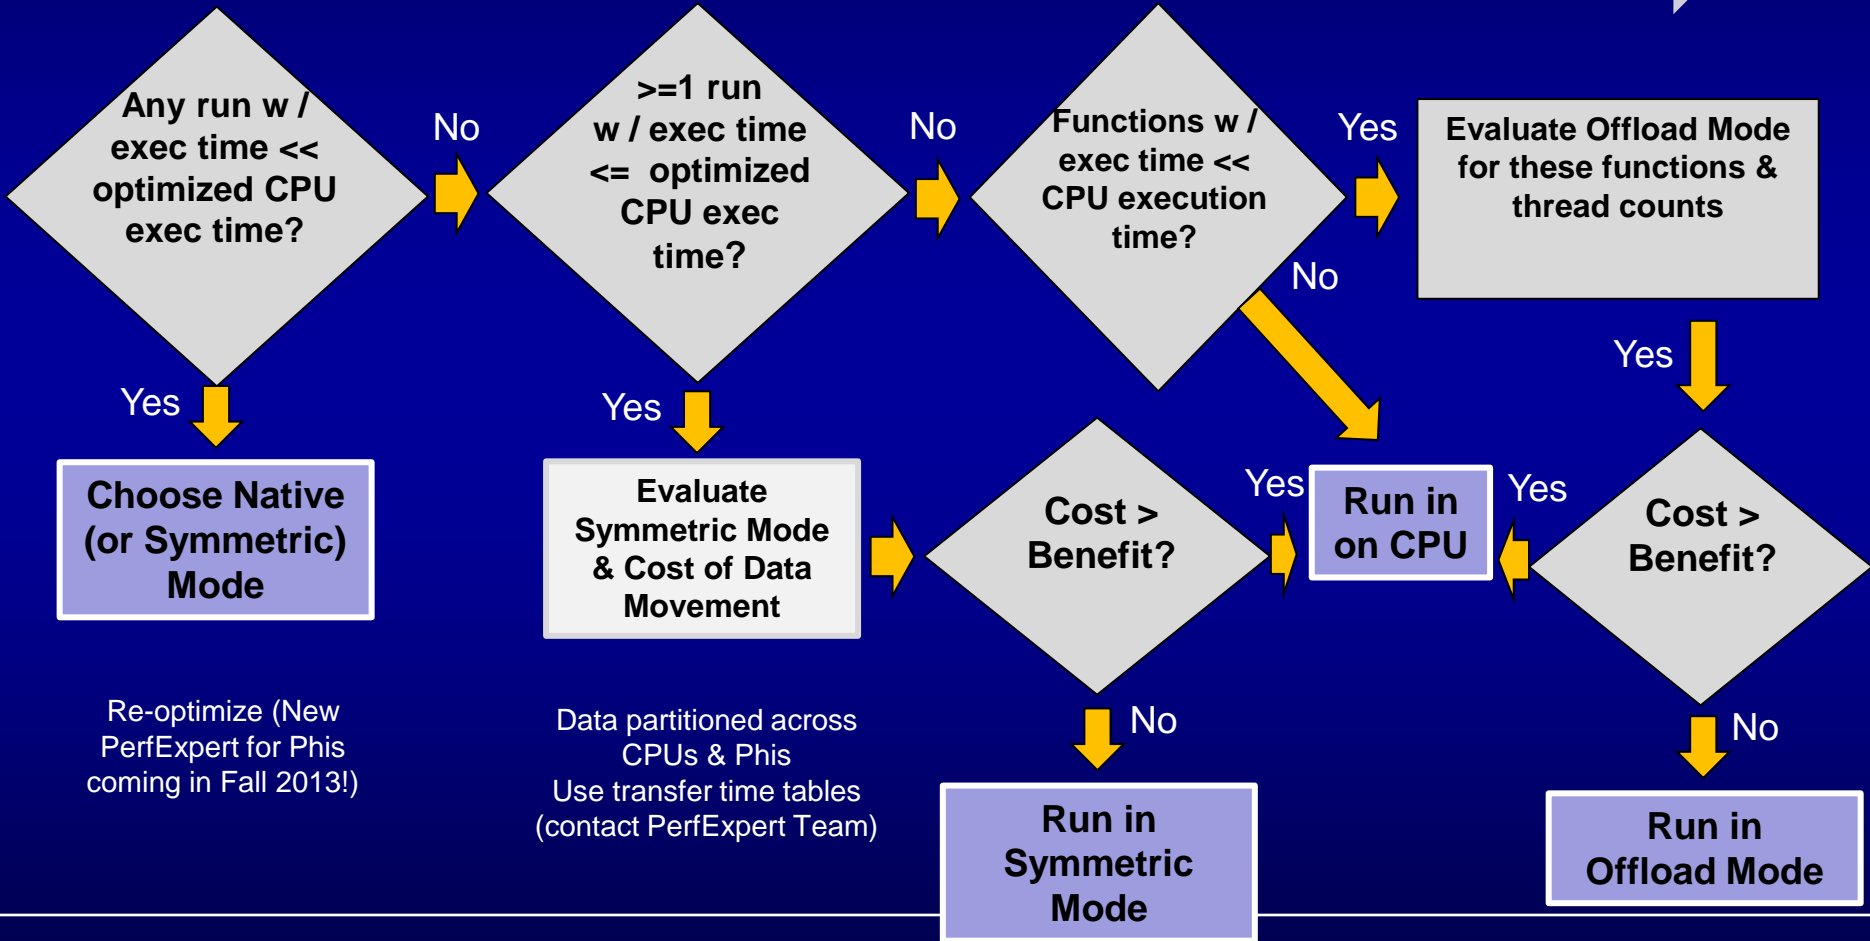
Xeon Phi Recommendor Process - 4

Study Native Mode Scaling

Run optimized app in Native Mode (Simple!) using different core/thread counts - try 30, 60, 90, 120 & 240



Use profiling data



Re-optimize (New PerfExpert for Phis coming in Fall 2013!)

Data partitioned across CPUs & Phis
Use transfer time tables (contact PerfExpert Team)