



# Distributed Optimization

Use multiple machines for maximum performance.

Distributed optimization lets you leverage multiple machines to dramatically reduce solve times. Some optimization models solve 15 times faster with 32 machines, and speed-ups of 2-3x are common with eight machines.

Gurobi offers three distributed algorithms:

- Distributed MIP --- where multiple machines work together to solve a single MIP model
- Distributed concurrent --- where multiple machines use different algorithmic strategies in a race to solve an LP or MIP model
- Distributed tuning --- where multiple machines do experimental solves to find parameter settings that improve performance

All three distributed algorithms are easy to use. Once your machines are set up, simply specify the algorithm and the number of machines to run it on. The Gurobi Optimizer handles all the work of dividing the computation among the machines.

The speed-up achieved by distributed optimization varies depending on the model. Read on to understand when to use distributed optimization.

## Distributed MIP

The distributed MIP solver divides the work of solving a single MIP model among multiple machines. A manager machine sends different portions of the MIP search tree to each worker machine to solve, and it periodically rebalances the remaining work across the workers.

### *When to use distributed MIP*

Distributed MIP works best on difficult models with large search trees. These models have enough work to ensure all the workers stay busy. Easier models that solve at the root won't benefit from the extra workers.

## Distributed Concurrent

The distributed concurrent solver uses a simple approach to take advantage of multiple machines. It starts an independent solve using a different strategy on each machine. The machines then race to see which can solve the model first. The solve is done when the first machine crosses the finish line. By trying different strategies at the same time, the concurrent optimizer can often find a solution faster than if it had to choose a single strategy.

### *When to use distributed concurrent*

Distributed concurrent is particularly effective on models whose solve times vary substantially depending on the data used to define the model or the algorithm used to solve it. Trying different strategies on different machines increases robustness and can smooth out variations in solve times.

## Distributed Tuning

The Gurobi Optimizer has a wide variety of parameters that control the algorithmic strategies used during a solve. It can be challenging to find the combination of parameter settings that yields the best performance on your specific model. The distributed tuning tool automates this search by performing a set of experimental solves on multiple machines.

### *When to use distributed tuning*

Distributed tuning is useful when you need to find settings that maximize the performance of a single machine solve.

## Machine Requirements

The distributed algorithms require a set of machines to serve as distributed workers. These algorithms work best on identical machines, but it is fine if the machines vary slightly.

Distributed MIP works best on 8-32 machines, distributed concurrent is most effective on 2-10 machines, and distributed tuning can take advantage of all available machines.

## Licensing

You can use any of the distributed algorithms by adding the distributed capability to an existing named user, single machine, or compute server license. A license is only required for the manager machine that launches and coordinates workers. Worker machines do not need separate licenses; they only need Gurobi Remote Services installed.

## Know Before You Buy

Do you want to know if your models can benefit from distributed optimization? Simply send us your models as MPS or LP files. We will run them on a cluster of machines and tell you if:

- ✓ distributed MIP solves your models faster
- ✓ distributed concurrent decreases variations in solve times
- ✓ distributed tuning finds parameter settings that speed up single machine solves