

# Running Highway Assignment Models Faster

## Ohio DOT's experience with Cube Cluster

---

### Background

Running travel demand forecasting models for large metropolitan areas, states, provinces, or nations requires a huge number of calculations and generates massive quantities of data. Cube Cluster, an extension to Cube Base, allows analysts to run models across multiple computer processors and substantially reduce model run times.

The Ohio Department of Transportation (ODOT), a user of Citilabs software for many years, installed Cube Cluster, hoping to improve the run times of their models. ODOT Technical Services staff decided to test Cube Cluster the highway assignment of their Columbus area MPO (metropolitan planning organization) model, which uses Cube Voyager. The population of the Columbus area is approximately 1.5 million. The model has 1877 traffic analyses zones.

### ODOT tests

ODOT tested Cube Cluster on a morning peak-period (AM) scenario. For the base run, the AM scenario generated approximately 1.3 million trips. ODOT developed several test runs to learn how the number of processors and the number of clustered Cube Voyager instances affected run

### How Cube Cluster works

Cube Cluster distributes run processes across multiple computer processors. Cube Cluster uses two forms of distributed processing:

- Intrastep distributed processing — Splits the zone-based processing from a single step into zone groups, and distributes the groups among available processors for separate processing.
- Multistep distributed processing — Distributes blocks of one or more independent modeling steps across available computer processors.

Cube Cluster requires a primary instance and one or more node instances of the processing Cube module, such as Cube Voyager. The primary instance initiates and monitors model runs. Each node instance simply completes assigned processes. To avoid performance degradation, Citilabs recommends using at most one instance per computer processor.

With appropriate configuration, Cube Cluster enables you to run model processes simultaneously, optimizing use of your computer resources and reducing your overall run times.

times. All tests, including the base runs, imposed a 20-iteration limit for model convergence.

ODOT ran their tests on hardware equipped with two dual-core processors. A dual-core processor has two separate computer processors on a single CPU (central processing unit) core. Thus, each ODOT computer had four unique computer processors.

---

**Citilabs, Inc.**  
312 Clay Street, Suite 180  
Oakland, California 94607, USA

**World Wide Web**  
[www.citilabs.com](http://www.citilabs.com)

Copyright © 2006 Citilabs, Inc. All rights reserved.  
Citilabs is a registered trademark of Citilabs, Inc. All other brand names and product names are trademarks, registered trademarks or trade names of their respective holders.

Many factors contribute to the results described. Citilabs does not guarantee results for all customers. Citilabs has carefully reviewed the accuracy of this document, but shall not be held responsible for any omissions or errors that may appear. Information in this document is subject to change without notice.

## Results

For each run, ODOT examined the time per iteration, a statistic they considered the best measure of performance. The results were consistent with expectations: adding clustered instances and new processors reduced run times significantly.

In the first test, adding a single additional clustered Cube Voyager instance nearly reduced the per-iteration run time in half, from just over two minutes to just over one minute. Doubling the number of clustered instances again, to four in the second test to match the number of available processors on the computer, reduced the per-iteration run time by more than half, to just under 40 seconds.

Test	Number PCs	Number Processors Used	Number Voyager instances	Time per iteration (minutes)
Base	1	1	1	2:10
1	1	2	2	1:10
2	1	4	4	0:39
3	1	4	8	0:39
4	2	8	8	0:23

To strain their computer hardware in their third test, ODOT attempted doubling the instances without doubling the processors—a practice not recommended if you wish to avoid performance degradation—and found no change in performance. But adding a second computer and doubling the number of processors and instances to eight in the fourth scenario, reduced the per-iteration run time a further 41 percent, to a mere 23 seconds.

ODOT was pleased with the results. With a very short learning curve, their staff was able to take advantage of Cube Cluster's distributed processing functions and gain significant run-time improvements. They felt with more experience using the software, they would be able to get even more gains in performance.

## More information

To learn more about ODOT, please visit <http://www.dot.state.oh.us/>.

To learn more about Citilabs and Cube Cluster software, please visit <http://www.citilabs.com>.

*Cube Cluster is one of the many innovative and market-leading solutions developed by Citilabs—the leader in transportation planning software solutions.*