

Improving Run Times with Cube Cluster

Testing more scenarios more quickly with existing resources

Running a travel forecasting model frequently takes many hours, sometimes days. Citilabs' new product, Cube Cluster, can dramatically cut model run times by distributing run processes across multiple computer processors, those in one or in multiple PCs.

Cube Cluster uses two methods to distribute model processes:

- **Intrastep distributed processing:** Cube Cluster splits the zone-based processing from a single step into zone groups, and distributes groups to available computer processors. For example, when running a mode choice model, Cube Cluster might process origin zones 1–100 on one PC and origin zones 101–200 on a second PC.
- **Multistep distributed processing:** Cube Cluster distributes blocks of one or more independent modeling steps across available processors. For example, Cube Cluster might run the trip generation model on one PC and the network building and skimming on a second PC.

Cube Cluster is an extension to Cube Base. In the initial version, Cube Cluster supports multistep distributed processing with all Cube-initiated

processes and intrastep distributed processing with two Cube Voyager programs:

- Highway — Used for traffic assignments and skimming
- Matrix — Used for demand model and matrix manipulation

Many travel models use these two programs for large and complex processes.

Example time savings

You can experience significant drops in run time with Cube Cluster. For example:

- One model currently runs in about ten hours on one PC. Running the same model across five PCs with Cube Cluster drops the run time to two hours and twenty minutes.
- Another model runs in four hours on one PC. Splitting the model across four PCs with Cube Cluster reduces the run time to one hour and seven minutes.
- Distributing a highway assignment run across two processors reduces each iteration in half, from two minutes to one minute. Distributing the same run across eight

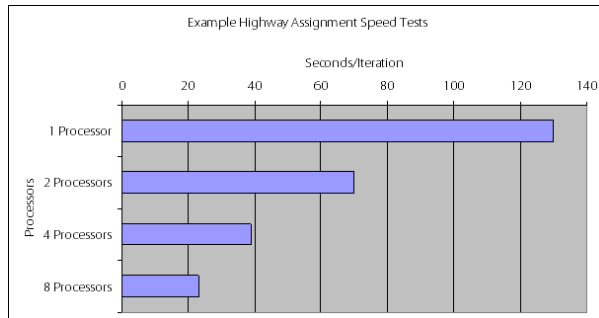
Citilabs, Inc.
312 Clay Street, Suite 180
Oakland, California 94607, USA

World Wide Web
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.

processors reduces each iteration by more than 80 percent, to just more than 20 seconds.



Cube Cluster is truly a huge advance in travel forecasting. Cube Cluster greatly reduces run times, increasing a modeler’s ability to respond quickly to management and client demands and to efficiently run and test many more scenarios. Cube Cluster has no scaling restrictions. You could take a model system and run it across 100 processors, perhaps reducing run times to several minutes!

How does it work?

With Cube Cluster, you configure your Cube Voyager model to run on clustered Cube Voyager instances across the available processors in the PCs in your network. For example, if connected through a standard office Ethernet network, you might configure Cube Cluster to run your model across several of the existing office PCs. Cube Cluster can even run the model across the processors on a single dual-core PC.

The distributed processing approach requires one multiprocessor computer or a network of computers:

- **Primary PC.** You designate one computer as the primary PC and install the model process and input data on this computer. This computer requires a full license of Cube Voyager (standard desktop license) and a license of Cube Cluster.

- **Node PCs.** You configure other computers in your network as node PCs, computers across which Cube Cluster can distribute computing processes. Node PCs do not require Cube Cluster and can operate with either desktop or server licenses of Cube Voyager or with low-cost “node” licenses of Cube Voyager. You cannot use machines with node licenses to start a model run or to run a model process independently. You can only use machines with node licenses in distributed processing, commanded by the primary PC.

Once you start a model run, Cube Cluster distributes the process across the configured processors and executes the model. Cube Cluster stores the resulting data and print files on the primary PC for analysis and mapping.

What do I need to buy?

You must have a full license of Cube Voyager and Cube Base, a license of Cube Cluster, and any required node licenses. You can buy desktop licenses for regular PCs or server licenses for computers with four or more processors.

For example, suppose you currently have one seat of Cube Base and Cube Voyager and you want to run your model across three PCs. You would need to purchase:

- One seat of Cube Cluster
- Two Cube Voyager node licenses

Depending on your model, this configuration can reduce your run times by about two-thirds. If your model runs in six hours today, then with this Cube Cluster configuration, you can reduce the run time to as little as two hours.

Want to make it run faster? Purchase more low-cost Cube Voyager node licenses and distribute the run across additional PCs in your network.

But my mode-choice model is in Fortran (Java, C...). Will it run faster?

Yes and no. For a mode-choice model, Cube Cluster might distribute the computations for different trip purposes across different processors. If you have coded your mode-choice model in a programming language, you must create separate executables for each trip purpose. With separate executables, Cube Cluster can allocate the different trip purposes to different processors or PCs.

Alternatively, you can take advantage of a specially priced offer from Citilabs to convert

models from Fortran, Java, C, etc., to native Cube Voyager scripting. Converting your existing models to Cube Voyager scripting would enable you to take full advantage of the Cube Cluster's complex allocation capabilities and generate large time savings.

Can I try it out?

Definitely. We are currently offering 30-day trial licenses of Cube Cluster to existing clients. Contact your sales representative today at sales@citilabs.com to organize a trial.

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