

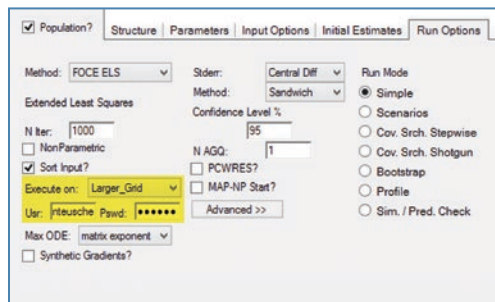
Grid and Cloud Computing with Phoenix® NLME™ Frequently Asked Questions



One of the most challenging aspects of PK/PD modeling is lack of computing power required to solve complex models in a reasonable time frame to support rapid drug development decisions. The explosion of cloud computing resources has provided access to significant compute power to solve these complex models; however, access to these cloud computing systems can be complex and confusing, and use of these systems generally requires knowledge of command-line tools. Phoenix 7.0 gives today's modeler a modern tool to easily access the power of these computing environments from the comfort of their desktop. A user can initiate a remotely executed NLME model from the Phoenix 7.0 application with the click of a button. Easy!

How do I run an NLME job on a cloud compute platform?

To send an NLME job to a remote/parallel computing platform, a user can select the desired remote/parallel computing platform using the dropdown box from the Execute option in the Run Option tab. If a username and password is required for the grid, it can be entered in the available boxes below the selected remote/parallel computing platform. The username and password will be saved by Phoenix for all future models. If a username and password is not needed (eg, MPI cluster), then the username and password boxes will be greyed out.



If a remote/parallel computing platform is selected in the Execute option, the user can click the "Execute" button or press F7 to initiate the NLME job to be executed on the selected remote/parallel computing platform. The NLME job can be executed on the local computer by selecting the Execute option as "Local" (uses 1 core) or "Local_MPI_n" (uses n cores on the local computer).

Phoenix 7.0 provides first-in-class integrated support for execution of NLME jobs on remote/parallel compute platforms including Message Passing Interface (MPI) clusters, remote Linux computers, and remote Linux grid (on your network or in the cloud including Amazon cloud), without restricting the number of cores a user can leverage.

What happens when I run an NLME job on a remote/parallel computing platform?

The use of an **MPI cluster** works in similar fashion to using MPI on local computer processors. An iteration progress status window will be opened and progress can be monitored just as if the NLME job was using the local processors. The difference is that all of the nodes on the MPI cluster will be used. The user must remain connected to the MPI cluster until the job is completed. In addition, the Phoenix application will remain locked until the NLME job is complete.

If a **Linux** grid is used, Phoenix will copy the required files to the remote grid and initiate the NLME job on the grid. An iteration progress status window will not be opened; however, the progress bar at the bottom of the Phoenix window will show the “by sample” progress of the job. For example, if 200 bootstraps are selected, the progress bar will show “10 of 200 completed” when 10 models have finished. After the NLME job has started, the user can disconnect from the Linux grid by saving the Phoenix project and then closing the Phoenix application. When the Phoenix application is restarted and the project is loaded, Phoenix 7.0 will automatically re-connect to the Linux grid and update the progress bar. If the NLME job has completed, the results will be copied from the Linux grid to the Phoenix project. All files are removed from the Linux grid when the NLME job finishes.

How are NLME jobs parallelized on cloud computing platforms?

Shotgun Covariate Search and Bootstrap involve execution of many (often hundreds) of independent models. Execution of Shotgun Covariate Searches and Bootstraps on all remote/parallel computing platforms will result in “by sample” parallel processing. A “sample” consists of one independent model of the set of models required by the method. For example, consider a 100 node Linux grid and a bootstrap NLME with 300 samples. In this scenario, each of the 100 nodes would execute 3 NLME models ($100 * 3 = 300$). Next consider a shotgun covariate search that contains 25 models. When executed on a 100 node Linux grid, only 25 nodes would be used with one model per node.

For MPI clusters, Shotgun Covariate Search and Bootstrap will be handled in the same way as for Linux grids. All other job types are executed by dividing each sample across multiple cores by subject. For example, consider a 100 core MPI cluster and a simple estimation job with 1000 subjects. Approximately 10 subjects would be assigned to each core during calculation steps until the model has completed.

How much does grid and cloud computing cost?

All support for remote/parallel computing platform is included with your NLME license at no additional cost. You can run NLME jobs on as many cores as you would like with no extra fees. Competitors charge you by the number of cores, but with Phoenix 7.0, you can use unlimited cores to allow you to meet your drug development needs without additional costs.

What types of cloud computing platforms are supported?

Phoenix NLME jobs can be executed on (1) MPI clusters, (2) Linux computer, and (3) Linux grids. The MPI Clusters, Linux computers or Linux grids can be user-maintained systems or cloud systems such as an Amazon Cloud computing platform. The number of cores/nodes utilized for each type of NLME job for each of these 3 remote/parallel computing platforms is shown in the following table:

NLME Job Type	MPI Cluster	Linux Computer	Linux Grid
Simple Estimation	All nodes	1 core	1 node
Scenarios	All nodes	1 core	1 node
Stepwise Covariate Search	All nodes	1 core	1 node
Shotgun Covariate Search	All nodes	All cores	All nodes
Bootstrap	All nodes	All cores	All nodes
Likelihood Profiling	All nodes	1 core	1 node
Simulation/Predictive Check	All nodes	1 core	1 node

Supported Cloud Computing Platforms

- MPI Cluster
- Linux Computer
- Linux Grid

What needs to be installed to use cloud computing?

A user's existing Phoenix installation is all that is needed. Phoenix NLME does *not* need to be installed on the remote/parallel computing platform. All required execution files are transferred to the remote/parallel computing platforms by Phoenix 7.0 and then the results are transferred back to the local computer and stored in the Phoenix project.

How do I set up my remote/parallel computing platform with Phoenix?

Phoenix 7.0 is out-of-the-box ready-to-use remote/parallel computing platforms. To start using the power of cloud computing, update a single informational text file (see details in Appendix A of the Phoenix NLME User Guide). Many users set up multiple grids and an MPI cluster to maximize the available remote/parallel computing power for modeling. There is no limit to the number of remote/parallel computing platforms that can be set up for each user. Now you have the power of the cloud in your hands with Phoenix NLME.

Some free software packages are required on the remote/parallel computing platforms including:

- MPI Cluster: R version 3.1.1 and the following R libraries (Rmpi, XML, Certara.NLME)
- Linux Computer and Linux Grid: R version 3.2 and the following R libraries (R BatchJob, XML, Certara.NLME)

Note that the R package Rmpi is supported on R 3.1.1, but is not supported on R 3.2 or higher, and Rmpi is no longer updated by the authors. Thus, users must have R 3.1.1 installed to utilize a remote MPI cluster.



About Certara

Certara is a leading provider of decision support technology and consulting services for optimizing drug development and improving health outcomes. Certara's solutions, which span the drug development and patient care lifecycle, help increase the probability of regulatory and commercial success by using the most scientifically advanced modeling and simulation technologies and regulatory strategies. Its clients include hundreds of global biopharmaceutical companies, leading academic institutions and key regulatory agencies.

For more information visit www.certara.com or email sales@certara.com.