
Nexus[®] User Guide

© 2009 Halliburton

HALLIBURTON | Landmark Software
& Services

December 2009

**© 2009 Halliburton
All Rights Reserved**

This publication has been provided pursuant to an agreement containing restrictions on its use. The publication is also protected by Federal copyright law. No part of this publication may be copied or distributed, transmitted, transcribed, stored in a retrieval system, or translated into any human or computer language, in any form or by any means, electronic, magnetic, manual, or otherwise, or disclosed to third parties without the express written permission of:

Halliburton | Landmark Software & Services

2107 CityWest Blvd, Building 2, Houston, Texas 77042-3051, USA

P.O. Box 42806, Houston, Texas 77242, USA

Phone: 713-839-2000, FAX: 713-839-2015

Internet: www.halliburton.com/landmark

Trademarks

3D Drill View, 3D Drill View KM, 3D Surveillance, 3DFS, 3DView, Active Field Surveillance, Active Reservoir Surveillance, Adaptive Mesh Refining, ADC, Advanced Data Transfer, Analysis Model Layering, ARIES, ARIES DecisionSuite, Asset Data Mining, Asset Decision Solutions, Asset Development Center, Asset Development Centre, Asset Journal, Asset Performance, AssetConnect, AssetConnect Enterprise, AssetConnect Enterprise Express, AssetConnect Expert, AssetDirector, AssetJournal, AssetLink, AssetLink Advisor, AssetLink Director, AssetLink Observer, AssetObserver, AssetObserver Advisor, AssetOptimizer, AssetPlanner, AssetPredictor, AssetSolver, AssetSolver Online, AssetView, AssetView 2D, AssetView 3D, BLITZPAK, CasingLife, CasingSeat, CDS Connect, Channel Trim, COMPASS, Contract Generation, Corporate Data Archiver, Corporate Data Store, Data Analyzer, DataManager, DataStar, DBPlot, Decision Management System, DecisionSpace, DecisionSpace 3D Drill View KM, DecisionSpace AssetLink, DecisionSpace AssetPlanner, DecisionSpace AssetSolver, DecisionSpace Atomic Meshing, DecisionSpace Nexus, DecisionSpace Reservoir, DecisionSuite, Deeper Knowledge. Broader Understanding., Depth Team, Depth Team Explorer, Depth Team Express, Depth Team Extreme, Depth Team Interpreter, DepthTeam, DepthTeam Explorer, DepthTeam Express, DepthTeam Extreme, DepthTeam Interpreter, Desktop Navigator, DESKTOP-PVT, DESKTOP-VIP, DEX, DIMS, Discovery, Discovery Asset, Discovery Framebuilder, Discovery PowerStation, DMS, Drillability Suite, Drilling Desktop, DrillModel, Drill-to-the-Earth-Model, Drillworks, Drillworks ConnectML,DSS, Dynamic Reservoir Management, Dynamic Surveillance System, EarthCube, EDM, EDM AutoSyn, EDT, eLandmark, Engineer's Data Model, Engineer's Desktop, Engineer's Link, ESP, Event Similarity Prediction, ezFault, ezModel, ezSurface, ezTracker, ezTracker2D, FastTrack, Field Scenario Planner, FieldPlan, For Production, FZAP!, GeoAtlas, GeoDataLoad, GeoGraphix, GeoGraphix Exploration System, GeoLink, Geometric Kernel, GeoProbe, GeoProbe GF DataServer, GeoSmith, GES, GES97, GESXplorer, GMaplus, GMI Imager, Grid3D, GRIDGENR, H. Clean, Handheld Field Operator, HHFO, High Science Simplified, Horizon Generation, i WellFile, I2 Enterprise, iDIMS, Infrastructure, Iso Core, IsoMap, iWellFile, KnowledgeSource, Landmark (as a service), Landmark (as software), Landmark Decision Center, Landmark Logo and Design, Landscape, Large Model, Lattix, LeaseMap, LogEdit, LogM, LogPrep, Magic Earth, Make Great Decisions, MathPack, MDS Connect, MicroTopology, MIMIC, MIMIC+, Model Builder, Nexus, Nexus View, Object MP, OpenBooks, OpenJournal, OpenSGM, OpenVision, OpenWells, OpenWire, OpenWire Client, OpenWire Direct, OpenWire Server, OpenWorks, OpenWorks Development Kit, OpenWorks Production, OpenWorks Well File, PAL, Parallel-VIP, Parametric Modeling, PetroBank, PetroBank Explorer, PetroBank Master Data Store, PetroStor, PetroWorks, PetroWorks Asset, PetroWorks Pro, PetroWorks ULTRA, PlotView, Point Gridding Plus, Pointing Dispatcher, PostStack, PostStack ESP, PostStack Family, Power Interpretation, PowerCalculator, PowerExplorer, PowerExplorer Connect, PowerGrid, PowerHub, PowerModel, PowerView, PrecisionTarget, Presgraf, Pressworks, PRIZM, Production, Production Asset Manager, PROFILE, Project Administrator, ProMAGIC, ProMAGIC Connect, ProMAGIC Server, ProMAX, ProMAX 2D, ProMax 3D, ProMAX 3DPSDM, ProMAX 4D, ProMAX Family, ProMAX MVA, ProMAX VSP, pSTAx, Query Builder, Quick, Quick+, QUICKDIF, Quickwell, Quickwell+, QUIKRAY, QUIKSHOT, QUIKVSP, RAVE, RAYMAP, RAYMAP+, Real Freedom, Real Time Asset Management Center, Real Time Decision Center, Real Time Operations Center, Real Time Production Surveillance, Real Time Surveillance, Real-time View, Reference Data Manager, Reservoir, Reservoir Framework Builder, RESev, ResMap, RTOC, SCAN, SeisCube, SeisMap, SeisModel, SeisSpace, SeisVision, SeisWell, SeisWorks, SeisWorks 2D, SeisWorks 3D, SeisWorks PowerCalculator, SeisWorks PowerJournal, SeisWorks PowerSection, SeisWorks PowerView, SeisXchange, Semblance Computation and Analysis, Sierra Family, SigmaView, SimConnect, SimConvert, SimDataStudio, SimResults, SimResults+, SimResults+3D, SIVA+, SLAM, SmartFlow, smartSECTION, Spatializer, SpecDecomp, StrataAmp, StrataMap, StrataModel, StrataSim, StratWorks, StratWorks 3D, StreamCalc, StressCheck, STRUCT, Structure Cube, Surf & Connect, SynTool, System Start for Servers, SystemStart, SystemStart for Clients, SystemStart for Servers, SystemStart for Storage, Tanks & Tubes, TDQ, Team Workspace, TERAS, T-Grid, The Engineer's DeskTop, Total Drilling Performance, TOW/cs, TOW/cs Revenue Interface, TracPlanner, TracPlanner Xpress, Trend Form Gridding, Trimmed Grid, Turbo Synthetics, VESPA, VESPA+, VIP, VIP-COMP, VIP-CORE, VIPDataStudio, VIP-DUAL, VIP-ENCORE, VIP-EXECUTIVE, VIP-Local Grid Refinement, VIP-THERM, WavX, Web Editor, Well Cost, Well H. Clean, Well Seismic Fusion, Wellbase, Wellbore Planner, Wellbore Planner Connect, WELLCAT, WELLPLAN, WellSolver, WellXchange, WOW, Xsection, You're in Control. Experience the difference, ZAP!, and Z-MAP Plus

are trademarks registered trademarks or service marks of Halliburton.

All other trademarks, service marks and product or service names are the trademarks or names of their respective owners.

Note

The information contained in this document is subject to change without notice and should not be construed as a commitment by Halliburton. Halliburton assumes no responsibility for any error that may appear in this manual. Some states or jurisdictions do not allow disclaimer of expressed or implied warranties in certain transactions; therefore, this statement may not apply to you.

Halliburton acknowledges that certain third party code has been bundled with, or embedded in, its software. The licensors of this third party code, and the terms and conditions of their respective licenses, may be found at the following location:

..\Nexus-Nexus-VIP5000.0.2\help\com\lgc\dsp\ThirdParty.pdf

Table of Contents

Chapter 1: Introduction	9
About This Manual	10
Additional Nexus® Manuals	11
Chapter 2: Nexus® and VIP® Comparison	13
VIP® Functionality Not Available in Nexus®	14
Nexus® Functionality Not Available in VIP®	15
PVT Differences	16
Typical Conversion issues	17
Chapter 3: Workflows and Data Conversions	19
Related Documentation.....	20
Creating a Single Reservoir Model.....	21
Convert Existing Simulation Data to Nexus® Input Data	21
Combine 3D Grid and Well Data with Nexus® Input Data	30
Multireservoir Modeling.....	46
Structure of Case File	46
Network Data	48
PVT Data	48
Map Data	49
Creating a Multireservoir Nexus® Model.....	51
Connect Well Networks and Merge Surface Pipeline Networks	54
Merge Reservoir Models in a Single Multireservoir Model	57
Running Simulation	62
Using the Job Submit Option	62
Using the Command Line	83
Writing Restart Data	97
Specifying Restart Data in the Runcontrol File	97
Specifying Restart Data with a Procedure	97
Specifying Restart Data by an Interactive Request	98
Specifying Restart Data Automatically	98
Generating a Nexus® Model with a Structured Grid	99
Generate Structured Grid Data in SimDataStudio™	99
Chapter 4: File Descriptions	105
Related Documentation.....	106
Standard Input Data Processing.....	107
Structured Grid Data Processing.....	109
Input Files	111
Output Files.....	112

Files Output by the Keyword Parser	112
Files Output by the Simulator	114
Intermediate Files	118
Map and Plot File Formats.....	119
Map File	119
Initialization Map File	121
Simulation Recurrent Map File	123
Plot File Format	125
Chapter 5: Region Summaries	127
Related Documentation.....	128
Using Regions.....	129
The REGDATA Table.....	131
Differences in Implementation between VIP® and Nexus®	132
Controlling Surface Volume Calculations within a Region	133
Chapter 6: Equilibrium	135
Related Documentation.....	136
Differences between Nexus® and VIP®	137
Useful Debugging Output	138
Equilibration Depth Table Report	139
Equilibration Overread Data Report	141
Debug Report	143
Reconciliation Report	145
Chapter 7: Parallel Processing	147
Related Documentation.....	148
Decomposition of Data	149
Unstructured Grid Data	150
Structured Grid Data	150
Assignment of Grids to Processes	151
Limitations of Current Parallel Implementation.....	154
Chapter 8: Well Data	157
Related Documentation.....	158
Well Data Tables and Files.....	159
Requirements for Unstructured Grids.....	160
Requirements for Structured Grids	161
Wellbore Geometry.....	162
WELLSPEC Data Used by Procedures	164
WELLSPEC Data Used by Drawdown/Buildup Constraint.....	165
WELLSPEC Properties That Can Be Modified by Procedures.....	166
Modifying WELLSPEC Data Using WELLMOD	167

Chapter 9: Production and Injection Controls	169
Related Documentation	170
How Production and Injection Are Modeled	171
Constraints for Individual Connections	172
Flow Control Devices	173
Surface Rate Constraint Calculations	174
Targets for Group Connections	175
Target Rate Calculation	177
Control of the Target Rate	179
Chapter 10: Run Control	183
Related Documentation	184
Time Step Control	185
Primary Timestep Controls (DT and DCMAX)	186
Timestep Cuts	188
Timestep Size Statistics	188
Newton or Non-linear Iteration Controls (TOLS)	190
Control of Maximum Changes During Implicit Newton Iterations	192
Well and Network Controls on Newton Iterations	193
Newton Iteration Statistics	195
IMPES Stability Controls (IMPSTAB)	196
Multiple Grids and Reservoirs	197
Simulation Statistics	198
Tuning	202
Chapter 11: Formulation	203
Related Documentation	204
The Coupled System of Equations	205
Reservoir Equations	206
Global Matrix Equations	209
Network Coupling and Solution	211
Chapter 12: Restarts	213
Related Documentation	214
Restarting a Simulation Run from Restart Data	215
Redefining Simulation Grids on Restart	217
Structure of the Restart Data	218
Appendix A: VIP® Functionality not Available in Nexus®	219
Appendix B: Nexus® Functionality Not Available in VIP®	221
Appendix C: Common Issues in Converting VIP® Data to Nexus® Data	223

Known ECLIPSE™-to-Nexus® Conversion Issues	227
Appendix D: Nexus® Desktop Features	229
Tools Menu	230
Session Tree	232
Viewer Panel	233
Item Identity	233
Nexus View Units	234
Case Contents	235
Job Parameters	236
Parallel Panel	238
Job Progress	246
Text File Contents	247
Icons	249
Tools Icons	249
Study and Case Management Icons	251
Job Submitter Panel	253
Job Control Panel	253
Job Settings Panel	254
Network File Chooser	255
Timestep Table and Plot	261
Queue Chooser	264
Version Chooser	266
Copy Example Dataset	267
Miscellaneous Tasks	269
Remove Job	269
Edit Files	270
Close Multiple Studies	271
Close Multiple Cases	271
Archive Data	271
Index	273

Introduction

Nexus® is Landmark's next-generation reservoir simulation package. Designed for the DecisionSpace environment, Nexus delivers a significant speed increase over older simulators while enabling true integrated asset management by providing the ability to model subsurface reservoir systems simultaneously with surface facilities in a tightly coupled solution environment.

About This Manual

This manual contains information to help the reservoir engineer understand how to work with Nexus® data and workflows. The manual does not go into deep technical detail and does not present detailed equations. It is designed to explain data and process requirements rather than to provide a mathematical understanding of Nexus. It contains chapters on the following topics:

- Nexus and VIP® comparison
- common workflows
- file descriptions
- region summaries
- equilibrium calculations
- parallel processing techniques
- well data structure
- production and injection controls
- run control logic
- formulation
- restarts
- conversion issues

Additional Nexus® Manuals

Nexus® documentation includes two additional manuals--the *Nexus Technical Reference Guide* and the *Nexus Keyword Document*.

- The *Nexus Technical Reference Guide* contains more detail on various technical aspects of Nexus, including solver technology, PVT calculations, and the Nexus network implementation. It is designed for simulation specialists, but all users of Nexus can use the manual to gain a better understanding of how Nexus performs certain tasks.
- The *Nexus Keyword Document* contains detailed descriptions of the input data required for Nexus simulation models. Nexus input data is typically constructed by using the pre-processing application SimDataStudio. The *Nexus Keyword Document* provides the detailed information required for modifying the data generated by SimDataStudio or creating a Nexus model manually. The manual is largely organized by data groups. These groups consist of grid and array data, PVT data, relative permeability and rock data, equilibrium data, aquifer data and recurrent data. The recurrent data includes simulation control data, well and surface network data, plus any additional data associated with wells or the surface network.

Nexus® and VIP® Comparison

Nexus® and VIP® have many functional similarities. However, Nexus is not derived from VIP. It has a completely separate code base and architecture. This being the case, not all Nexus functions map directly to VIP functions. Where they do, the functionality in Nexus may be implemented differently than in VIP.

This chapter will clarify the primary differences between Nexus and VIP. It will explain how to avoid potential problems when moving from VIP to Nexus. It will also describe some issues that may be encountered when moving from ECLIPSE™ to Nexus.

VIP® Functionality Not Available in Nexus®

The Nexus® simulator processes transmissibilities and pore volumes that have been calculated externally, either by other applications or by a Nexus preprocessor. Currently, the Nexus preprocessor does not handle some of the types of grid data that can be processed by VIP® -- for example, radial grids and nine-point difference schemes. However, if transmissibilities and pore volume data are generated from radial grids or nine-point differences externally (by VIP or a third party application), Nexus can use the data.

The inputs to Nexus are basic reservoir properties such as pore volume, depth and integer array assignments, and the transmissibilities. These inputs are treated as cell-to-cell connections. Since there is no pre-existing grid geometry requirement, Nexus is not limited to radial or rectilinear grid geometry. It can solve highly unstructured grids as easily as it can solve structured grids, provided that the data is entered in proper formats.

Other than grid geometries, most commonly used VIP functionality has been implemented in Nexus. However, some VIP features have not been implemented yet. These will be implemented based on client demand and effort required. For example, non-Darcy gas treatment will probably be added fairly soon. Other features, such as thermal simulation, require considerably more resources. These features will be added over a longer period of time.

Please see "[VIP® Functionality not Available in Nexus®](#)" in this manual for a list of VIP features that are not available in Nexus.

Nexus® Functionality Not Available in VIP®

Nexus® provides significant functionality that does not exist in VIP®, for example, the ability to model surface facilities that are tightly coupled with the subsurface reservoir model

Please see "[Nexus® Functionality Not Available in VIP®](#)" in this manual for descriptions of the functionality that is unique to Nexus.

Note: The list in Appendix B is dynamic. New features are being added to Nexus in response to our users' input. Questions about the latest Nexus functionality should be directed to Nexus support.

Note: Nexus also has functionality that exists in VIP, but is difficult to use. For example, Nexus has the capability of solving complex grid geometries much more easily and efficiently than VIP.

PVT Differences

The primary difference between Nexus® and VIP® for black oil models lies in the conversion of VIP BOTAB data.

VIP BOETAB and related data can be converted directly to Nexus. The only requirement is reformatting of the data, which is performed automatically by SimDataStudio™ when you export a SimDataStudio case containing the data to a set of Nexus files.

VIP BOTAB data must be adjusted prior to being used by Nexus. This data adjustment is currently performed by SimDataStudio and results in modified PVT tables, plus black-oil separator input.

For ease of use, you should analyze any VIP models that use BOTAB data to see if you can substitute BOETAB or related data without impacting results.

Typical Conversion issues

Converting data from a VIP® data format or another simulator data format to the Nexus® format can involve several potential pitfalls. Some of these are obvious -- for example, data for a specific VIP feature may not map to a parallel data format in Nexus because the feature is not yet available in Nexus. Others are more subtle and may only be apparent in the divergence of results.


A summary of known VIP-to-Nexus and Eclipse-to-Nexus conversion issues is shown in "[Common Issues in Converting VIP® Data to Nexus® Data](#)". This appendix will be updated frequently. If questions arise, please contact Nexus support for the most up-to-date summary.

Workflows and Data Conversions

While Nexus® data can be created manually, the preferred workflow is to use the SimDataStudio™ software to consolidate, edit, and store data as well as to generate the required Nexus files.

The SimDataStudio software enables you to specify the inputs to be used for the Nexus simulator. Input data can be imported into the SimDataStudio software from an existing Nexus case (.fcs), a VIP® project or an ECLIPSE™ project, or it can be defined from scratch.

Once data has been loaded to the SimDataStudio software and required parameters have been defined, the SimDataStudio case data can be exported as Nexus data.



Related Documentation

Procedures for creating a SimDataStudio™ case and editing the parameters of an existing case can be found in the *Nexus SimDataStudio™ User's Guide*.

Details on the E2V conversion utility and a list of unsupported ECLIPSE™ keywords can be found in the *E2V User's Guide*.

Instructions on how to use the Nexus® View, PowerGrid™, and DecisionSpace® Well Planning software can be found in these products' online help systems, available from their respective help menus.

Detailed descriptions of the keywords that appear in the Nexus .fcs file, and its associated files, can be found in the *Nexus® Keyword Document*.

Creating a Single Reservoir Model

A Nexus® single reservoir model contains the following types of input data:

- grid (unstructured or structured)
- reservoir parameters
- PVT and rock properties tables
- aquifer data (optional)
- well and perforation data
- surface network data
- hydraulics tables
- well management data (production constraints to meet targets, etc.)
- array data (reservoir property multipliers, etc.)
- schedule data (start date, end date, output types, etc.)

The simplest procedure for creating Nexus® input data is to import the input data from an existing reservoir simulation model into the SimDataStudio™ software, or to import existing grid and well data from an active Nexus® Desktop working set into the SimDataStudio software.

Alternatively, you can enter all of the input data manually from within the SimDataStudio software.

Convert Existing Simulation Data to Nexus® Input Data

Nexus provides several methods for converting input data from different sources into a Nexus model, using a SimDataStudio .vds file as the intermediary data format. These methods include:

- Importing Nexus® data from an existing Nexus case into a new Nexus model.
- Converting VIP® data to Nexus data
- Converting ECLIPSE™ data to Nexus data
- Converting existing 3D grids and well planning projects to Nexus data

In each case, you import data to a SimDataStudio case (.vds) and then use the SimDataStudio software to generate a Nexus model (.fcs).


When you import only 3D grids and well plans, you must specify the input simulation data manually in the SimDataStudio software.

Note: Each of the conversion methods are mutually exclusive. You cannot import data from multiple simulation case files into the same SimDataStudio case.

Note: SimDataStudio case files (.vds) are not explicitly tied to a specified Nexus study or case. The SimDataStudio software can generate multiple sets of Nexus data and corresponding .vdb structures via interactive modification of data in the case file. As this can quickly get confusing, a good policy is to use a single SimDataStudio case file for each single Nexus study and case. You should also assign the SimDataStudio case file the same name as the Nexus .vdb case file.

VIP® to SimDataStudio™ to Nexus®

You can convert an existing VIP model to a Nexus model by importing it into the SimDataStudio software and then generating a Nexus model.



To derive a Nexus case from a VIP dataset, perform the following steps.

1. In Nexus® Desktop, click on the **Open or Create New Simulation Study** icon.




2. In the Network File Chooser dialog box, perform the following steps.

- a. Navigate to the location where you want your new Nexus study to reside.
 - b. Specify a study name with a .vdb extension.
 - c. Click on **Open**.

A Create Case dialog box displays the name of a Nexus case that will be created automatically.

 - d. Rename the case if you want it to have a different name than the Nexus study.
 - e. Click on **OK**.
3. In the Nexus® Desktop tree, select the case name under the new study.



4. Click the SimDataStudio icon on the Nexus toolbar.




The SimDataStudio software opens with the New SimDataStudio Case dialog box open. The dialog box lists a case name for the new .vds file, file names, and file locations.

The VDS file name does not have to use the same suffix as the Nexus study or case names. However, using the same suffix will simplify case management.

5. Toggle the **Nexus Case** checkbox ON.
6. Toggle the **Create a new case by parsing an existing VIP data set** option ON.

7. Select the i.dat and r.dat files to be used in the conversion.



8. Click on **OK**.

The SimDataStudio software parses the VIP files and stores the data as a binary SimDataStudio case file (.vds). As the files load, you see the file names along with a progress indicator at the bottom of the window.

Unrecognized data: If you receive a prompt indicating that the SimDataStudio software encountered unrecognized data that can be written to an include file, click the **Yes** button to close the prompt.

Most VIP data that is applicable to Nexus in this release should be parsed by the SimDataStudio software, but review any created INCLUDE files to make sure. The INCLUDE files created by the SimDataStudio software will be located in the study folder with the filename *{date}.inc. Note that a SimDataStudio option can be set to allow INCLUDE files for grid-related data (large files) to be maintained as INCLUDE files. If this option is selected, the SimDataStudio software will only load the data when Nexus data is to be generated.

9. Examine the SimDataStudio data tree to determine whether any of the nodes are colored red. Look at both the CORE tree and the EXEC tree. If you find red nodes, enter the relevant input parameters manually.
10. In the SimDataStudio software, add any input parameters that have not been defined by the imported VIP data.
11. In the SimDataStudio software, select **Simulation Data -> Generate Simulator Data File -> Nexus Files** to generate Nexus files.

The Nexus Data Set Generation dialog box appears. It will list the types of input data that will be written to the Nexus case file (.fcs). If any “error detected” messages appear, correct the errors in the SimDataStudio software.

12. Navigate to the directory where you want to save the Nexus case file, specify a case file name, and click on **Generate**.

The Nexus case data is generated. A log file is displayed in the the SimDataStudio software lower panel. This file posts information about the data transformation.

For a more detailed description of how to generate a Nexus model from within the SimDataStudio software, see [Export SimDataStudio™ Case Data as Nexus® Data](#).

13. Select **File -> Save** case to save the SimDataStudio case (.vds).


You must save the SimDataStudio case so that it will be available for subsequent edits. It is stored separately from the Nexus case.

ECLIPSE™ to E2V to VIP® to SimDataStudio™ to Nexus®

You can convert an existing ECLIPSE model to a Nexus model as follows:

- Convert the ECLIPSE model to a VIP model.
- Import the VIP model into the SimDataStudio software.
- Generate a set of Nexus project files from the SimDataStudio case.

Use this procedure if you want to create an intermediary VIP model or if you want to have access to the full set of E2V conversion options. If you do not need the VIP files, use the [ECLIPSE™ to SimDataStudio™ to Nexus®](#) procedure.



Note: Some ECLIPSE keywords may not be converted, so you should inspect the results of this conversion carefully. You may have to enter additional parameters in the SimDataStudio software.

To derive a VIP model and a Nexus model from an ECLIPSE dataset, perform the following steps.

1. In Nexus® Desktop, create a new study and a new case.

The new study and case are added to the Nexus® Desktop session tree.

2. Click on the case in the session tree to highlight it.
3. Click on the **E2V** icon on the Nexus® Desktop toolbar.



The E2V Converter dialog box displays.

4. Specify the ECLIPSE file to be converted by clicking on the **ECLIPSE File to be Converted** Browse button, browsing to the directory where the file resides, and selecting the ECLIPSE file.
5. Specify the VIP .dat file where the VIP data will be stored by clicking on the **VIP File to be Created** Browse button, browsing to the directory where VIP files should reside.

In the file name text box, replace the asterisk with the name you are using for the case.

6. On the E2V Converter, click **Convert to VIP**.

The Progress bar indicates the progress of the conversion. The data files are listed in the bottom panel of the E2V Converter as the data is converted. When the conversion is complete, a popup notifies you that the data has been processed.

7. Click **OK** to close the popup.


This process will convert ECLIPSE files into VIP format. Once you have the VIP model, you can follow the [VIP® to SimDataStudio™ to Nexus®](#) procedure to convert the VIP data to Nexus data and to generate the Nexus case files

Additional E2V Options

- Use the E2V Converter **Tools -> Properties** options to set preferences for converted models.
- If your ECLIPSE model has many gridblocks that number on the order of millions, you can use the **Large Model Optimization** option to reduce the memory usage during the processing of ECLIPSE files (**Tools -> Properties -> Advanced**).

ECLIPSE™ to SimDataStudio™ to Nexus®

You can convert an existing ECLIPSE model to a Nexus model by importing the ECLIPSE model directly into the SimDataStudio software and generating a set of Nexus project files.



This procedure converts ECLIPSE input data to Nexus input data by running E2V in the background. It skips the intermediate creation of a VIP dataset and directly stores the applicable ECLIPSE data in the SimDataStudio case file.

Note: This procedure runs E2V in the background. As a result, you do not have the ability to set preferences for the converted model. If your ECLIPSE model has many gridblocks (on the order of millions), you may want to use the [ECLIPSE™ to E2V to VIP® to SimDataStudio™ to Nexus®](#) procedure so that you can enable the **Large Model Optimization** option to reduce the memory usage during processing ECLIPSE files.

To derive a Nexus case from an ECLIPSE model without generating a VIP model, perform the following steps.

1. In the Nexus® Desktop, create a new study and a new case.

The new study and case are added to the session tree in the Nexus® Desktop.

2. Click on the case in the session tree to highlight it.

3. With the case name still highlighted, click on the **SimDataStudio** icon.



The SimDataStudio software opens, with the New SimDataStudio Case dialog box open. The case name (for the SimDataStudio case), file names, and file locations are filled in by default using the selected Nexus case.

4. Toggle the **Create an empty Data Studio case** and **Nexus case** checkboxes ON.
5. Click **OK** to close the dialog box and create an empty VDS case.

Note: Use the SimDataStudio **View -> User Options -> ECLIPSE Import** tab to set importing preferences.

6. Click on the **Load ECLIPSE Data** icon in the SimDataStudio toolbar to open the Select ECLIPSE Data File dialog box.



7. Navigate to the directory where the ECLIPSE file resides, select the data file, and click on **Open**.

A message box warns that, when parsed, the ECLIPSE data will be written to the SimDataStudio case, overwriting any existing data. Click on **OK**.

E2V parses the ECLIPSE file and stores the data as a binary SimDataStudio case file. As the files load, you see the file names along with a progress indicator at the bottom of the window.

Unrecognized data: If you receive a prompt indicating that the SimDataStudio software encountered issues during the ECLIPSE translation, click the **Yes** button to continue the import.

Most ECLIPSE data that is applicable to Nexus in this release should be parsed by the SimDataStudio software, but some ECLIPSE keywords may not be supported. If data is not recognized, it will be written to an include file in the study folder. The include file will have the the filename `*{date}.inc`.

8. Examine the SimDataStudio data tree to determine whether any of the nodes are colored red. If you find red nodes, enter the relevant input parameters manually.
9. In the SimDataStudio software, add any input parameters that have not been defined by the imported ECLIPSE data.

10. In the SimDataStudio software, select **Generate Simulator Data File -> Nexus Files** to generate Nexus files.

The Nexus Data Set Generation dialog box appears. It will list the types of input data that will be written to the Nexus case file (.fcs). If any “error detected” messages appear in red, correct the errors in the SimDataStudio software.


11. Click on the **Nexus case data file** folder icon and navigate to the directory where you want to save the Nexus case file. Specify a case file name and click on **Generate**.

For a more detailed description of how to generate a Nexus model from within the SimDataStudio software, see [Export SimDataStudio™ Case Data as Nexus® Data](#).

Combine 3D Grid and Well Data with Nexus® Input Data

You can create a Nexus model by entering the simulation parameters manually in the SimDataStudio™ software or importing them from ASCII files, and then importing a 3D grid, well plans, and a surface network from the “working data set” of an Nexus® Desktop session. The “working data set” is the set of data that is currently available to all Nexus applications.

Nexus provides the functionality to perform a complete simulation workflow while taking advantage of time-saving features offered by the integrated environment. This procedure includes the option of upscaling an existing fine-scale geological model (Rescue, ECLIPSE, or vdb) with PowerGrid™. It uses Well Planning's AssetPlanner™ to create new platforms and wells, based on specified target and well parameters, and NetworkPlanner to create a surface network. The upscaled 3D grid, well plans, and surface network plan are imported into the SimDataStudio software. Additional reservoir data is added in the SimDataStudio software.



In this workflow, you can use Nexus[®] View to load and view your data. You can also interactively view the wells, targets, and surface pipeline network generated in DecisionSpace[®] Well Planning.

Launch Nexus[®] Desktop

1. Select **Start -> Landmark -> Nexus[®] Desktop**.

Load a 3D Grid in Nexus[®] View


For the current workflow, the 3D grid must be in the working data set.

1. In Nexus[®] Desktop software, select the case that contains the grid that you want to view.
2. Click on the Nexus[®] View icon.




Nexus® View displays the first property listed in the grid's Property list.

The Data Selector also displays. The grid and its contents are displayed in the Model tree of the Active Project panel.



3. Select the property that you want to display in the Property column.



4. Click on **Apply**.

The selected property displays in Nexus® View.

Upscale the 3D Grid

Upscaling is optional. It provides a method for reducing the number of cells in the grid, while maintaining the general characteristics of the reservoir.

Check the Numbering Sequence Before Specifying Layer Numbers

For user-defined coarsening, you need to specify layer numbers, so you need to be aware of whether to count from the top or bottom.

All VDBs created by Landmark software are numbered from the top down. However, some software applications export Rescue models with their layers numbered from the bottom up. Bottom-up layering will affect the layer lumping when you use constant coarsening or user-defined grouping. With constant coarsening, if the layering is not evenly divisible into the original grid, the leftover thin layer will be on the top of the grid instead of the bottom.

To check the order of layer lumping, display the 3D grid in the Nexus® View 3D view and then use the Grid Inquiry tool (accessible from an icon on the toolbar) to determine the layer number for layers on the top and bottom. If using automatic layer lumping, you will see the layer numbering in the first column of the table on the Layer Lumping panel.

1. Select **Launch -> PowerGrid** in the Nexus® View toolbar.

The **Upscaler Workflow** icon appears on the Nexus® View toolbar.



2. Click on the icon.

The Upscaler Workflow wizard launches.

3. Select the fine grid to be used for the upscaling.
 - a. Select the fine grid that you want to upscale by clicking on the 3D Grid Model drop-down list and clicking on the name of the 3D grid model.

The Grid Information panel displays information about the selected grid. The three tabs contain tables with information on any Properties, Faults, and Units in the grid. Units can be created using the Reservoir Unit Manager. The first 3D grid loaded in the session is selected by default.

- b. If you have a template, then load a template containing upscaling parameters that was saved in a previous upscaling session. [Optional]

The upscaling parameters in the template include any property mapping that was done. By default, these parameter files are saved to an UpscalerSessionPersistence.xml file in the user's HOME directory as specified in your user environment variables.

- c. Select a region to exclude from the coarsening from the drop-down list of all regions in the selected grid. [Optional]
 - d. Click **Next** to move to the next step in the workflow.
4. Map existing properties in the input grid to valid properties in the simulation model and specify an upscale method for each of the property mappings.

- a. Locate the properties that you want to map to the output grid.

You can use the **Look for fields** option to search or filter the table by specifying a text string and selecting a column.

If you click **Search**, the rows containing the specified text are highlighted.

If you click **Filter**, the table is filtered to show only the rows containing the specified text.

Click **Reset** to restore the original table.

You can use the icons above the table to perform an advanced search/filter or to select/deselect all rows.

- b. Map each 3D grid property to be converted to the appropriate Nexus keyword (target property).

Left click in the Target Property cell to access the drop-down list of simulation properties (i.e., Nexus keywords) assigned to your grid.

Double-click on a simulation property to select it.

Alternatively, you can create a user-defined property that has no equivalent simulation keyword by right clicking in

the cell to activate edit, left clicking to activate the insert mode, and then typing in a name.

The original property units are maintained. If you map a property to KX but map nothing to KY and/or KZ, the upscaler will use KX to calculate and output KX, KY, and KZ for the coarsened model.

You can also select **Auto Map** to map source properties automatically to any existing simulation properties.

- c. Select or change the upscale method for each pair of mapped properties.

The "upscale method" is the upscaling algorithm used to upscale the grid. When you select a target property from the drop-down list of simulation properties, the application automatically selects an appropriate upscale method.

Click in the Upscale Method cell to see if there are additional upscaling methods available for the selected target property. Click on a different method in the drop-down list to select it if desired.

- d. Click **Map** to map the source properties selected (highlighted) in the table to the specified target property.

You must specify a target property for the selected rows. Use **Unmap** to remove the property mapping for the source properties selected (highlighted) in the table. When not mapped, the Target Property column will indicate "None." Toggle on one of the following export options:

Export Well PI Permeabilities and Flow-based transmissibility multipliers (High resolution upscaling)

- Exports well PI permeability to the coarsened model.

(See the PowerGrid™. help system for information on the well PI permeability calculations.)

The following properties will be included in the coarsened model: three well PI permeabilities (KX, KY, KZ) and six half-transmissibility multipliers (TMXMT, TMXPT, TMYMT, TMYPT, TMZMT, TMZPT).

Export Flow based permeabilities (Low resolution upscaling) - Exports only the three flow base permeabilities (KX, KY, and KZ). This is the default.

The direct solution of pressure equation by simulation layer upscaling algorithm is not available if this method is selected.

- e. Click **Next** to move to the next step in the workflow.


5. Define the coarsening parameters.

If the direct solution of pressure equation by simulation layer upscaling algorithm is selected, six additional half-transmissibility multipliers are output (TMXMS, TMXPS, TMYMS, TMYPS, TMZMS, TMXPS).


- a. Select the coarsening method to be used for the columns (I), rows (J), and layers (K) from the drop-down list for each and enter the desired values in the text fields to the right of each direction. Different methods can be used in each direction.

The available coarsening methods are:

Constant Coarsening Factor - Lets you specify the number of refined cells you want in each coarsened cell. In the example shown below, a factor of three is specified. The software adjusts the odd-numbered cells as necessary to best fit the constant value.




Number of Coarse Cells - Lets you specify the total number of coarsened cells you want in the upscaled grid. In the example shown below, 5 coarsened cells are specified. The software merges the odd-numbered cells as necessary to ensure a total of 5.



User Defined Grouping - Lets you specify the grouping for each coarse cell. The format for this entry is a list of cell indices separated by commas. In the example below,

the values 1, 4, 10 were given. As a result, the first coarse cell includes blocks 1 through 3; the second coarse cell includes blocks 4 to 9; and the last coarse cells includes all the remaining cells from 10.



Automated Lumping - Allows you to quickly and accurately upscale the refined grid into a simulation grid using a residual optimization technique. (This option is available only for layers.)

If this option is selected, the layer lumping step in the workflow is activated. When this option is selected, the value turns red to indicate that it is not being used. When you click **Next**, the automatic layer lumping is performed. If you then click **Previous** on the Layer Lumping panel or click on **Coarsening Parameters** in the workflow tree, the first layer in each lumped group is shown in the values field.

This option is not available if no target properties are mapped to the source properties since the application then uses the input names.

- b. If the 3D grid that you are using has more than one unit, select from the following options for Layers (K):

Work on entire grid - Uses the selected coarsening method and values for all layers in the grid.

Work per unit - Allows you to select a different coarsening method and/or values for each layer.

- c. Click **Next** to move to the next step in the workflow.

The number of coarsened cells that will be created using the current parameters is shown at the top of the panel.

6. Specify the vdb and case names and a location to save vdb.

All upscaled grids are saved in VDB format.

7. Click **Upscale** to start upscaling process.

The VDB is added to the Working Data Set. When the upscaling is completed, information on the number of root cells, inactive cells, cut faces, LGRs, and LGR cells is shown at the bottom of the panel.

8. Click **Finish** to close the Upscaler Workflow wizard.

Define the Well Plans and Surface Networks


The next step in the workflow is to define targets for new wells and planning wells through those targets. You use the DecisionSpace® Well Planning application to perform this task.

1. In Nexus® View, click on the **Launch -> Well Planning** icon to activate the Well Planning icon bar.
2. Click on the **Well Planning** icon to launch the Well Planning application.



Well Planning launches.

In the Modules pulldown menu on the left panel, select **AssetPlanner**.



The AssetPlanner™ software opens the Targeting tab.

3. Create targets as follows:
 - a. Toggle on **3D Grid** and then browse to and select the upscaled .vdb.

The case and class definitions will be automatically entered in the Grid specification string.
 - b. Specify a maximum number of targets in the **Max # of targets** field.
 - c. Specify an objective property in the **Objective Property** field.

- d. Specify filters for refining the targeting using appropriate property ranges.

- e. Specify other parameters as appropriate.

See the AssetPlanner help for information on how to set up target parameters.

- f. Click on **Create Targets**.

The AssetPlanner software creates the targets and they display in the session tree on the left side of the window. You can combine these targets into well plans.

4. Create well plans as follows:

- a. Click on the Planning tab.
- b. Select the type of wells to be used.
- c. Specify the type of platforms to be used.
- d. Click on the Plans sub-tab.

- e. Enter a user prefix for the well plans.

- f. Specify engineering parameters as appropriate.

See the AssetPlanner help for information on how to set up well plan parameters.

- g. Click on **Create Plans**.

The AssetPlanner software combines the targets into well plans and connects the well plans to platforms. The resulting well plans are listed in the session tree on the left.

- h. Click on **Delete Unused Targets** to remove targets that did not get incorporated into the well plans.

5. Create perforations as follows:

- a. Click on the Perforations tab.
- b. Toggle on **Grid** and, if necessary, browse to the .vdb that you selected on the Targeting tab.
- c. Set the parameters as desired.

See the AssetPlanner help for information on how to set up filters.


d. Click **Create Perforations**.

The AssetPlanner software perforates the wells using the specified parameters and displays the results in a Results table.

6. Use the NetworkPlanner module in Well Planning to create a surface pipeline network that connects the platforms and wells.

Network Planner provides a graphical interface to create a simple surface network. You must save all the nodes, connections, and elevations profiles the Nexus .vdb for use within the SimDataStudio software. To create the surface pipeline network, perform the following steps.

a. In the Modules pulldown menu on the left panel, select **NetworkPlanner**.



b. On the Well Planning menu bar, select **Tools -> Automatic Network** to open the Automatic Pipeline Network Generation window.

Using this option, you can automatically connect wells and nodes to each other based on the specified maximum pipeline segment length. The software will find optimum locations for connection points, when connecting two or more nodes or wells.

c. Enter the proper parameters and click **Apply** to create a pipeline.

Important: Be sure to assign a unique name in the **Node Type** field. If you create multiple surface networks using the default, the multiple networks will contain nodes with the same node names. Later, if you try to combine the networks in a multireservoir model, the duplicate names will cause the surface network data to become scrambled.

As you create these pipelines, NetworkPlanner will automatically connect them in a network. You can repeat this step as many times as necessary. Click **Cancel** to close Automatic Pipeline Network Generation.

- d. In the Well Planning main tool bar, click **View > Plan View** to open a plan view and examine the resulting surface network.

You can create additional nodes and connections in the plan view as desired, or edit existing nodes and connections. You can also add additional equipment facilities like separators, chokes and valves.

- e. Click on **File -> Save Project As**.

The Save As dialog box displays. Click on the **VDB File** tab and save the well plans and surface network plan *in the same .vdb file as the upscaled VDB grid*.

Important: Do not create a separate case. If you do, you will only be able to access either the grid or the surface network plan in the SimDataStudio software. You want a case that contains both.

At this point, you may want to examine the surface pipeline network in Nexus® View to determine if you want to make any further changes. If not, you are ready to import the contents of the .vdb into the SimDataStudio software and create the case file for simulation.

- f. Select **File -> Close** to exit Well Planning.

Import the Well Plans and Surface Network into SimDataStudio™

When you save your network plan in the VDB that contains the upscaled grid, the surface network data gets written to an .xml file that is located in an SPN folder. The SimDataStudio software can use this data to create simulation perforation data (with the SimDataStudio Perforation Wizard) and to generate a Nexus® surface network.

Note: Communication is one-way from NetworkPlanner to the SimDataStudio software. Modifications made to the network in the SimDataStudio software will not be exported to the .xml file stored in the .vdb. If you change the elevation profile in the SimDataStudio software, your changes will not show up when you subsequently view the surface network in Network Planner.

Perform the following steps to import the well plans.

1. In the Nexus® Desktop session tree, highlight the case where you stored the upscaled model and surface pipeline network data and then click on the **SimDataStudio** icon.



The New SimDataStudio Case dialog box appears with **Fill case with VDB grid data** and **Fill case with VDB well/SPN data** checked on at the bottom.

Click **OK**.

The data from your study .vdb, including grid and SPN data, are loaded into the SimDataStudio software.

2. Click on the **EXEC Data** tab of the SimDataStudio tree.
3. Select the **Well Names and Locations** node in the EXEC Data tree.

Note that the Well Names and Locations table is populated with well names, but that the I Location and J Location columns are empty.

4. Click on the **Well Perforations** node in the EXEC Data tree.
Note that the Well Perforations table is empty.
5. Click on the **Perforation Wizard** icon.



Follow the instructions in the wizard to import the perforations into the SimDataStudio case.

When the import is complete, note that the Well Perforations table is now populated with the perforation data that you created in Well Planning. A pulldown menu at the bottom of the table allows you to view the perforations for each of the wells.

6. Click on the **Well Constraints** node in the EXEC Data tree.
Set the appropriate constraints for each of the wells. At a minimum, you must set at least one constraint per well.
7. Click on the **Surface Network Nodes** node in the EXEC Data tree.

Specify the node pressure limits and node temperatures for the nodes in the surface network. At a minimum, you must set these parameters for the root node of the network.

Populate the SimDataStudio™ Case with the Reservoir and Recurrent Input Data Necessary for Simulation

At this point in the workflow, the SimDataStudio case study has been populated *only* with well data, perforation data, and surface network data. You need to provide the reservoir properties, PVT table data, rock property data, recurrent data, and any other data required for the type of simulation that you intend to perform.

1. Enter the reservoir initialization parameters from within the SimDataStudio software.

See the *SimDataStudio User Guide* for instructions on how to enter the required parameters. If you have existing reservoir initialization data in ASCII VIP format (for example, PVT data or relative permeability tables), you can import it into the SimDataStudio software.

2. Enter the recurrent data parameters from within the SimDataStudio software.

See the *SimDataStudio User Guide* for instructions on how to enter the required parameters.

Export SimDataStudio™ Case Data as Nexus® Data

Perform the following steps in the SimDataStudio software to generate the Nexus case file (<case>.fcs) and its associated include files. This procedure assumes that the SimDataStudio case has all required grid, reservoir and well properties defined.


1. Click the **Generate Nexus Files** icon on the SimDataStudio toolbar to open the Nexus Data Set Generation dialog box.



It is recommended not to change the path or the vdb or case name for Nexus data.

The Nexus Data Set Generation dialog box appears. This dialog box lists the various data types to be generated as outputs.

A Model Summary lists the types of data that are available to be exported.




Before you generate the files, you can click in the file directory tree to select or deselect specific files. For example, if you do not need new map files, you can save time and disk space by deselecting the map file option.

2. Review and correct the data. Deselect the data types that you do not want to generate.

Use the plus (+) and minus (-) buttons to open and close branches of the data tree. Review the corresponding Status field, which will display warnings of missing or incomplete data. Correct the data as indicated.

The SimDataStudio software will check for data completeness and consistency, but will not perform error checking during the file generation. Nexus will perform error checking when a Nexus job is submitted with the Job Submitter.

If data is incomplete, the SimDataStudio software will flag it.



- At the bottom of the dialog box, select the comments mode to be used:

Standard - Include only those comments generated by the SimDataStudio software (usually short descriptions of data cards).

User - Include user-generated comments found while parsing. (Comments must adhere to comment protocol.)

None - Do not include any comments.

- Toggle on **Structured grid format** if you want Nexus to generate a Nexus input model with the structured grid format.

For more information on the structured grid format, see [Generating a Nexus® Model with a Structured Grid](#).

- Click the **Generate** button to generate the Nexus files.

If the selected .fcs file exists, you will be prompted to confirm that you want to overwrite the existing file. A message will display if external include files were generated while parsing the SimDataStudio data. Make sure the include files are checked for any data required by Nexus. A progress monitor displays while the conversion is in progress.

Multireservoir Modeling

A multireservoir case allows the input data of several independent Nexus® models to be combined into a single coupled model without modifying the individual models. The models can then be connected to a common surface facility model, allowing all reservoirs to be controlled by a common set of production and injection constraints.

If you have multiple Nexus models, each with a surface network, you can create the multireservoir model using the NetworkPlanner and the SimDataStudio multireservoir wizard. See [Creating a Multireservoir Nexus® Model](#) for a discussion of this workflow.

The following sections describe how a multireservoir model differs from a standard single-reservoir model.

Structure of Case File

The Nexus® multireservoir case (.fcs) file is somewhat different from the single-reservoir case file. It contains a limited set of data definitions which pertain principally to the surface network that combines the reservoirs.

The case file defines the following:

- standard case utility data (input units, run units, etc.)
- reservoirs in the multireservoir model
- reservoir properties such as reservoir temperature and reservoir pressure
- surface network data, such as nodes, connections, constraints, etc.
- run control parameters such as arrays to be written out at particular times or intervals

The multireservoir case does NOT include the following:

- grid data
- rock data
- initialization parameters

- array data
- aquifer data

The keywords that can appear in the multireservoir case file are shown in the table below.

Keyword	Data File Type	Description
RESERVOIR		Identifies the individual reservoirs that belong to the multireservoir. A RESERVOIR keyword appears for each of the individual reservoirs, followed by the reservoir name and the reservoir case file name..
OPTIONS	PROPERTY_FILES	Introduces a file that defines parameters such as reservoir pressure; reservoir temperature; region names, separator, and PVT methods; etc. See the “Options Data” chapter of the <i>Nexus Keyword Document</i> for details.
PVT	PVT_FILES	Introduces a file that specifies PVT parameters for the surface network only.
SURFACE	RECURRENT_FILES	Introduces a file that specifies a surface network.
RUNCONTROL	RECURRENT_FILES	Introduces a file that specifies run control parameters for the multireservoir. Individual reservoir run control settings will be ignored.
HYD, VALVE, GASLIFT, PUMP	NET_METHOD_FILES	Introduce files that specify equipment methods for the multireservoir surface network.

For additional information on the multireservoir case file options, see the “Multiple Reservoir Case File” section of the “Case File” chapter of the *Nexus Keyword Document*.)

Network Data

The surface network model can be as simple or complex as required to model the controlling features of the reservoir performance adequately.

The surface model is connected to the individual reservoirs by connections (NODECON) between nodes in the surface model and nodes in the individual models. The nodes in the individual model may be wellhead nodes or any type of gathering nodes where one or more well streams come together. If an individual model has surface network data, all data downstream of the connection point will be ignored. The connections between the multireservoir surface model and the individual models are typically connections with zero pressure drop, but this is not a requirement.

Replacing the well models of the individual reservoirs with a generalized network model of the wells and facilities allows a more accurate representation of the complex wellbore geometry and downhole equipment required for a multireservoir. This technique avoids inefficiencies inherent in solving the various network models separately for each reservoir.

PVT Data

If no PVT methods are entered in the multireservoir case, the default PVT method for the surface connections will be the first PVT method from the first individual model specified using the RESERVOIR keyword. Typically, the PVT data in the various reservoirs is not similar enough to be modeled with a single PVT method, so the PVT data in the multireservoir case must incorporate one or more PVT methods containing API interpolation methods. This approach allows the mixing of different fluid streams in the network to be correctly modeled and volumetric consistency to be maintained. For example, if Reservoir 1 has a PVT method with oil density of API=28 and Reservoir 2 has a PVT method with oil density of API=30, a multireservoir PVT method could be created using the PVT data from the two individual methods. The multireservoir method would then contain two sets of PVT data with different oil densities. When the production streams from the two reservoirs mix in the surface network, this new PVT method could be used to model the mixing of the two fluids. Depending on the complexity of the network, several multireservoir PVT methods might have to be created from different combinations of individual methods.

Map Data

If the results of multireservoir simulation are to be viewable in a 3D viewer, the Nexus® data set must include an initial map (i.map) file. This file must combine the grid properties of the individual reservoirs.

Due to the complex structure of this combined map file, you should create the initial map file using the SimDataStudio software. When you run the SimDataStudio multireservoir builder wizard (described in [Creating a Multireservoir Nexus® Model](#) workflow), a combined i.map and/or .vdb file will be created as an output. The combined i.map and/or .vdb for a multireservoir will be written to the same directory as the master .fcs file file.

The map file is based on the local grid refinement (LGR) concept. A ROOT grid has one cell for each reservoir model in the case.

The ROOT grid cells have cornerpoint data (CORP) based on the geometry of the individual models. This cornerpoint data was created when the individual model map files were created. It is stored in MAP_ORIGIN files. When the ROOT grid is generated in the multireservoir map file, each component reservoir's initial map file is copied to the multireservoir map file. The GRID name for each reservoir model consists of the first eight characters of the RESERVOIR name in the multireservoir case file (.fcs). This naming convention requires that the first eight characters of each reservoir name be unique.

Note: The order of the data in the multireservoir map file must match the order in the multireservoir case file. If the map grid data is in a different order, the map results from the multireservoir simulation run will not load and display.

Individual reservoir map files may use different origins for their map data. If they do, you must determine a common origin and apply OFFSETS to the individual map data in order to define correct relative locations. For example, if a case consists of three reservoirs with two of them using UTM coordinates for map data and the other one using a local origin of (0,0), an offset will be required in order for the maps to appear in the correct relative location in the 3D viewer. The MAP_ORIGIN files for the individual models can be helpful in determining the actual map origin for each reservoir and any necessary offsets can be calculated from that data.

Map files can be generated for individual reservoir models using the structured grid format. However, they cannot be generated for multireservoir models using this method. The only way to create a multireservoir map file is to set up a multireservoir case in the

SimDataStudio software or to use a custom or third party external application.

Once the map file is created, the multireservoir grid system must remain unchanged. If the order of the reservoirs in the case file changes, a new map file must be generated. Likewise, if any of the individual models are modified such that the grid properties change, a new map file must be generated.

Creating a Multireservoir Nexus® Model

A Nexus multireservoir model (.fcs) allows several independent Nexus models to be combined into a single coupled model without modifying the individual component models. Simulation jobs can be run using a common set of production and injection constraints.


The multireservoir model specifies the following types of input data:

- PVT tables for network connections
- well and perforation data
- surface network data
- well management data (production constraints to meet targets, etc.)
- schedule data (start date, end date, output types, etc.)

Other data, such as grid data, array data, rock properties, etc., resides in the individual Nexus models. (See [Creating a Single Reservoir Model](#).)

In this workflow, you combine multiple standalone Nexus models into a single Nexus multireservoir model with common surface facilities and a common pipeline network.


We assume that you already have your Nexus standalone models. Those models have their own surface facilities and pipeline networks.



Three Separate Surface Networks

The first step in this workflow is to connect the surface networks manually into a single surface network system. You perform this step in the Plan View after merging the various networks in the same network

planner project. If necessary, you can add additional facilities (like separators, pumps, chokes, etc.) to the common surface network.



Three Surface Networks Connected Manually in Plan View

The next step is to merge your subsurface models into a single model in the SimDataStudio software and link it to the common surface network model. You perform this step with the SimDataStudio™ multireservoir wizard.

The final step is to generate a single Nexus multireservoir model.

Connect Well Networks and Merge Surface Pipeline Networks

This procedure assumes that you have already created well networks for the reservoirs that you want to combine, and you have stored the well networks in the .vdb's of the different reservoirs to which the networks belong. Please see the Well Planning online help system for instructions on how to create the surface networks.

Rules for Creating Surface Network Plans

When creating the surface network plans, observe the following rules:

- All well names, wellhead names, and node names must be unique.
 - Surface network names should share the same prefix. Sharing the same prefix will simplify the process of creating the multireservoir project in Nexus®.
-

Perform the following steps to merge multiple surface pipeline networks into a single network that spans multiple reservoirs.


1. Open a Nexus session.
2. Specify the session units mode that is appropriate for the reservoirs that you intend to merge.
3. From the Nexus® View main menu, activate the Well Planning icon bar (**Launch -> Well Planning**).

The Well Planning icon bar appears.

4. Click on the Well Planning icon.



5. In Well Planning, select **Network Planner** from **Modules** pulldown menu.



6. In Well Planning, select **File -> Open Project From -> VDB** file. Select the .vdb containing the first surface network plan that you want to merge.

An Open VDB Case dialog box appears with the case names that reside in the .vdb.


7. Select the appropriate case name and click **OK**.

The wells and surface network load to Well Planning and display in the nodes, wells, and curves tables. They are also displayed in Nexus® View.

8. Repeat steps 6 and 7 to browse to the .vdb containing the second surface network that you want to merge.

A dialog box appears with the case names that reside in the .vdb.

9. Select the appropriate case name. This time click on **Merge with existing models**.



The contents of the second case are merged with the contents from the first case.


The surface network plans should not have nodes or wells with the same names. If they do, the merge will scramble the data. For example, if there are two nodes with the same names, one of them will get dropped.

Note: The merging may take some time as the networks and wells get displayed in Nexus® View. Do not attempt to move the dialog boxes or the Nexus® View window while the merge is being performed.

10. Repeat the previous steps until you have merged all of the surface pipeline networks that correspond to the reservoirs that you want to combine in the multireservoir model.

When you have merged all of the surface network plans, the nodes, wells, and well sections will all be listed in a common hierarchy in the Well Planning tree view.

11. View the merged surface pipeline networks in a single Nexus® View scene.



Note: At this point, the different networks will not actually be physically connected. However, their data will be combined in the same well planning project. You must connect the networks manually.

12. In the Well Planning Plan View (**View -> Plan View**), connect the various surface network plans into a single network.

Use the **Node** option to create the connections between the networks.

Do not use separators or valves as the primary connections.

13. Once the interconnection network has been created, save it to a new .vdb file.
 - a. In Well Planning, select **File > Save Project As**.
 - b. Select .vdb file tab and save the merged surface pipeline network in its own .vdb.
14. Close Well Planning.


Merge Reservoir Models in a Single Multireservoir Model

This procedure assumes the following:

- You already have full field simulation standalone models (.fcs), each with its own surface network.
- You have created interconnection surface pipeline networks for the Nexus® multireservoir model and merged the pipeline networks.
- You have stored each Nexus model in a separate folder and stored those folders under multireservoir model folder. (Recommended.)

Perform the following steps to combine multiple Nexus models into a single multireservoir model.

1. In Nexus Desktop, create a new multireservoir study for the multireservoir model in the folder where you keep your folders with standalone Nexus models (**File -> Create New Study**).
2. Select the case in the Nexus tree.



3. Click on the **SimDataStudio** icon on the toolbar.



The SimDataStudio software opens with the New SimDataStudio Case dialog box open on top.


4. Perform the following operations in the New SimDataStudio dialog box.
 - a. Toggle the **Use Assistant to Create a Nexus Multireservoir Case** option ON.
 - b. Click on **OK**.

The Multireservoir Builder Wizard displays.

5. Perform the following operations in the Multireservoir Builder Wizard dialog box.
 - a. Click on **New** (insert) button at the bottom right corner of the Reservoir Data panel.



This will open the Multiple Case File Selection dialog box. If your standalone reservoirs are on the same hierarchy tree as your .vdb, you should see all the standalone Nexus models as being available for merging.



- b. Toggle ON the checkboxes for those reservoirs that you would like to combine into multireservoir model and click **OK**.
 - c. Select the timestep size control method to use. The available options are:

Merge the timestep size controls - This option applies the most restrictive criteria across the reservoirs.

Keep timestep size controls from the first reservoir - The timestep controls of the first reservoir in the list. If necessary, use the arrow controls to position the desired reservoir as first in the list.

- d. Select the simulation output option to use. The available options are:

Merge the output options - Uses output options from the different models and applies them at the highest specified frequency.

Keep output options from the first reservoir - Applies the output controls of the first reservoir in the list and ignores the output options specified in the other models. If necessary, use the arrow controls to position the desired reservoir as first in the list.

6. Click on **Next**.


The Setup Interconnect Network and Node Pressure Limits panel displays. You can specify interconnection surface pipeline network components manually in this panel, or you can import the interconnection surface pipeline network that you created in Well Planning.

7. In the Setup Interconnect Network and Node Pressure Limits panel, perform the following steps.

- a. Click on the **Import Network** option.

The Select File Containing Interconnect Network Definition dialog box appears.

- b. Select "VIP database file (*.vdb)" as the file type to be used as a filter.




- c. Select the .vdb of the study where you stored the surface pipeline network for multireservoir model that you created in Well Planning.

The Select Case with Simulation Grid dialog box displays. This dialog box lists the case files in the study.

- d. Select the case where you stored the merged surface pipeline network.
- e. Click on **OK**.


If the name of the case file name in Nexus model (.fcs) does not match the name of case in the .vdb for each individual reservoirs, then you will get message similar below one.



Click **Yes** and map your reservoir names with case names in the provided panel and click **OK**.

- f. Set a pressure limit for the top node in the merged network (required) and for subsidiary nodes.
 - g. Set other parameters as appropriate.
8. Click the **Next** button.

The "Connect individual reservoirs to interconnect network" panel displays how your reservoirs are connected in the multireservoir model.




9. Click the **Next** button.

The "Define interconnect network properties and constraints" displays. In the "Type" and "Method" columns, you can specify different methods for pressure drop calculation for certain connections. If you select one of the correlations, you should input a diameter for that connection. Also, you should input the PVT method to use for pressure calculations. By default, PVT Method 1 will be used for all connections.

10. Click the **Next** button.

A summary panel displays listing the outputs that will be generated and the files to which they will be written. If the panel does not post error messages in red, the wizard is ready to generate the multireservoir model.



11. Click **Finish** and the Nexus model will be generated.

You can now save your project in a .vds file. Once you have saved it, you can edit it, adding or removing reservoirs and generating new multireservoir models as needed.

Running Simulation

You can run simulation by invoking the **Job Submit** option from Nexus Desktop. Alternatively, you can submit jobs from the command line.


Using the Job Submit Option

You can submit local and remote jobs from the Nexus Desktop Job Submitter panel.

Local Serial Job

Once you have created a Nexus® model, you can run simulation from Nexus® Desktop.

1. In the Nexus® Desktop session tree, select the case that you want to submit to the simulator.




The Study panel displays the case file information, including the location of the study to which it belongs and the types of content that it contains.

2. Click on the icon for the Nexus Job Submitter.



A new job node appears in the session tree. The job node is identified by a string that concatenates the case name with “Nexus Simulation.”

The session tree also expands to list the case's .fcs file and the .dat files that it references.



Information about the job appears in the Study panel.

Study:	Case:	Job:
C:/GRIDS/apivar.vdb	apivar	Nexus Simulation
Job Location	[localhost] C:/GRIDS	
☐ Copy VDB		
FCS File	[localhost] C:/GRIDS/apivar.fcs	
☐ Case Initializer	[localhost] C:/GRIDS/nexus_data/apivar_ne...	
☐ Parallel Mode		
Queue Name	NONE	
Version	5000_0_1	

3. Verify the default job location. If you want to use a different location, enter it instead.

The job location is the path to a host and disk directory in which the job results will be written. By default, the job location is the directory that contains the .vdb for the study and case that you selected in Step 1.

4. Confirm that the appropriate case file (.fcs) is displayed in the **FCS File** field.

By default, the field posts the case file for the case that you selected in Step 1.

5. If you want your simulation results to be available for display in the Nexus View or 3DView viewer, toggle on the **Case Initializer** option and select the i.map file that was output when you generated the Nexus model in the SimDataStudio software.


The i.map file will be located in the same folder as the case .fcs file.

6. Select the simulator version that you want to use.

The **Version** pulldown menu lists the simulator versions that are available in the nexussimulators directory of your Nexus installation. If you copy the R2003.19.3 simulator to this directory, it will show up as being available.

Important: Do not select a version that is not compatible with your dataset, or the job may not run.

7. Click on **Start**.



The Job Submitter submits the job. The right panel of the Nexus® Desktop displays progress messages.

When the job is complete, the session tree lists the results files that are available for viewing. A timestep table report provides timestep outputs. A standalone report file provides details about the actions performed during the job. A Nexus report file (the .out file) provides comprehensive details of the data that has been generated.

Note: If the job failed, look for warnings and errors in the standalone file to determine what went wrong.


Nexus restart data can be written under a variety of conditions as described in the [Writing Restart Data](#) discussion.

Local Parallel Job

For any parallel job submission, your Nexus data must include decomposition settings.

The following instructions assume that you are setting up a two-process job. You can modify the steps to set up more than two processes by keying in the desired number.

1. In the Job Settings panel, click the **Parallel Mode** option to enable it.



The Assign Grids to Processes dialog box displays. If a message notifies you that a grid to process file could not be found, make certain that your model data is set up for a parallel job with a decomposed simulation grid.

Note: If a red X appears beside the **Parallel Mode** option when you toggle it ON, the X indicates that a parameter for the parallel option has not been set or has not been appropriately set.


2. Click on the **Open Parallel Mode** button (to the far right of the Parallel Mode check button).

The Assign Grids to Processes dialog box displays.

Note: If a message notifies you that a grid to process file could not be found, make certain that your model data is set up for a parallel job with a decomposed simulation grid.


3. On the left of the parallel settings panel, select the tab labeled "Defined."

A list of defined remote hosts will be displayed.




4. Click on "localhost" in the defined remote hosts list to select it.
5. With "localhost" selected, type the number "2."

After one second, the name "localhost" will be listed twice in the selected nodes list in the middle of the parallel panel.



6. On the right of the panel, expand the parallel grid tree to view all the simulation grids.




7. Click on a parallel grid in the tree to select it.

Note: You can select more than one grid by clicking and dragging the mouse or by clicking on multiple grids while pressing the 'Ctrl' or 'Shift' key.

8. Type the number '1' or '2'.

The number typed will appear next to the grid name, and the host name "localhost" will appear on the same line. The color of the background behind the typed number will match the color of that number in the selected nodes list.

9. Repeat Steps 7 and 8 until all the grids have been assigned a node for processing.
10. Click **OK**.
11. Click on **Start**.



The Job Submitter submits the job. The right panel of the Nexus® Desktop displays progress messages.

When the job is complete, the session tree lists the results files that are available for viewing. A timestep table report provides timestep outputs. A standalone report file provides details about the actions performed during the job. A Nexus report file (the .out file) provides comprehensive details of the data that has been generated.


Note: If the job failed, look for warnings and errors in the standalone file to determine what went wrong.

Remote Parallel Job

For any parallel job submission, your Nexus data must include decomposition settings. For remote parallel jobs, the NEXUS_SCRIPT and NEXUS_SIM environment variables must be set on the remote system. See the Nexus-VIP 5000.0.2 release notes for information on how to configure a remote system for parallel jobs.

The following instructions assume that you are setting up a two-process remote parallel job. You can modify the steps to set up more than two processes.


1. In the Job Settings panel, click on the **Job Location** navigation button.



The Select a Job Location dialog box appears.

2. Click on **Edit**.

The Network Host Chooser dialog box displays.



3. Specify a network host as follows:
 - a. Specify the host name in the **Host** field.
 - b. Specify the Linux user name in the **User** field.

- c. Specify the Linux user password in the **Pass** field.
- d. Click on **Save**.

The host that you have specified will be listed in the Defined Remote Hosts list of the Node to Grid Assignment dialog box when you open it.

4. In the Select a Job Location dialog box, click on the **Open** button at the top of the dialog box (next to the Hosts field).

The dialog box opens the specified Linux host to the home directory of the Linux user that you provided.

5. Navigate to the directory where you want the job results to be written. If necessary, use the **Create New Folder** option to create a folder for the remote job on the remote host.

The location that you have specified will be the job location. The Nexus .vdb will be written here if you toggle the **Copy VDB** option on.

6. Click the **Open** button at the bottom of the Select a Job Location dialog box (next to the File Name field).

The location in the Job Location field is changed to the remote host and folder path that you defined. The **Copy VDB** option becomes active and is toggled ON if you have a local .vdb file, and it has not already been copied to the remote folder.

Note: The **Copy VDB** option specifies that the Nexus .vdb will be copied to the remote job location when the job is run. The option will become active only if the .vdb is NOT accessible by a network path (i.e., they are not in a shared folder that can be accessed via the network).

7. Toggle the **Case Initializer** check button ON and select the appropriate i.map file if you want the Nexus case to be initialized with map data using an i.map file.
8. Toggle the **Parallel Mode** check button ON.

Job Location: [schubert] /home/filer2home/snewland/apivar_runs

☐ Copy VDB

FCS File: [schubert] /home/filer2home/snewland/apivar_runs/apivar.fcs

☒ Case Initializer: [schubert] /home/filer2home/snewland/apivar_runs/apivar_nex...

☐ Parallel Mode

Queue Name: NONE

Version: 5000_0_2

Note: If a red X appears beside the **Parallel Mode** option when you toggle it ON, the X indicates that a parameter for the parallel option has not been set or has not been appropriately set.

9. Click the Open Parallel Mode button for the parallel mode option.

Job Location: [schubert] /home/filer2home/snewland/apivar_runs

☐ Copy VDB

FCS File: [schubert] /home/filer2home/snewland/apivar_runs/apivar.fcs

☒ Case Initializer: [schubert] /home/filer2home/snewland/apivar_runs/apivar_nex...

☒ Parallel Mode

Queue Name: NONE


Version: 5000_0_2

The Assign Grids to Processes panel will appear.


Note: If a message notifies you that a grid to process file could not be found, make certain that your model data is set up for a parallel job with a decomposed simulation grid.

10. On the left of the parallel settings panel, select the tab labeled "Defined."


11. Select the name of the remote location host and type the number “2.”



12. After one second, the name of the remote location host will be listed twice in the selected nodes list in the middle of the parallel panel.



13. On the right of the panel, expand the parallel grid tree to view all the simulation grids.



14. Click on a parallel grid in the tree to select it.

Note: You may select more than one grid by clicking and dragging the mouse or by clicking on multiple grids while pressing the ‘Ctrl’ or ‘Shift’ key.

15. Type the number '1' or '2'.

The number typed will appear next to the grid name, and the remote location host name will appear on the same line. The color

of the background behind the typed number will match the color of that number in the selected nodes list.

Grid Name	Process Id	Host Name
Parallel Grid		
Reservoir 1		
Grid 1	2	schubert
Grid 2	2	schubert
Grid 3		
Grid 4		

16. Repeat Steps 13 and 14 until all the grids have been assigned a node for processing.
17. Click **OK**.
18. Click on **Start**.



The Job Submitter submits the job. The right panel of the Nexus® Desktop displays progress messages.

When the job is complete, the session tree lists the results files that are available for viewing. A timestep table report provides timestep outputs. A standalone report file provides details about the actions performed during the job. A Nexus report file (the .out file) provides comprehensive details of the data that has been generated.


Note: If the job failed, look for warnings and errors in the standalone file to determine what went wrong.

Remote Parallel Job with a Cluster (Linux Only)

For any parallel job submission, your Nexus data must include decomposition settings. For remote parallel jobs, the NEXUS_SCRIPT and NEXUS_SIM environment variables must be set on the remote system. See the Nexus-VIP 5000.0.2 release notes for information on how to configure a remote system for parallel jobs and how to configure a remote cluster.

The following instructions assume that you have access to a remote host that has a cluster of nodes configured on it.

1. In the Job Settings panel, click on the Job Location navigation button.



The Select a Job Location dialog box appears.


2. From the Hosts pulldown menu, select a remote host that has a cluster configured for parallel job submission.

See [Remote Parallel Job](#) for a discussion of how to add hosts to the Hosts list.

Note: The **Copy VDB** option specifies that the Nexus .vdb will be copied to the remote job location when the job is run. Toggle this option ON if you have a local copy of the Nexus .vdb that you want to be written to the remote job location.


3. If you want the Nexus case to be initialized with map data using an i.map file, toggle the **Case Initializer** check button ON and select the appropriate i.map file.

4. In the Job Settings panel, click the **Parallel Mode** option to enable it.



Note: If a red X appears beside the **Parallel Mode** option when you toggle it ON, the X indicates that a parameter for the parallel option has not been set or has not been appropriately set.

5. Click the **Open Parallel Mode** button for the parallel mode option.




The Assign Grids to Processes panel will appear.


Note: If a message notifies you that a grid to process file could not be found, make certain that your model data is set up for a parallel job with a decomposed simulation grid.

6. On the left of the parallel settings panel, select the tab labeled "Defined."


7. Select the name of the remote location host and type the number of processes on the remote location host.



After one second the remote job location host will appear in the selected nodes list in the middle of the parallel panel copied by the number that you typed.




8. Select the tab labeled "Named" on the left of the parallel panel.
9. Double-click the empty line of the named parallel hosts list to add a host name.
10. Type the name of a parallel host on the cluster and hit **Enter**.
You must enter a valid name of an existing host on the cluster.
11. The parallel host typed in the previous step should be selected. If not, select it.



12. Type the number of processes upon the parallel host that you wish to use. After one second the host name will be added to the selected nodes list.

13. Repeat Steps 9 through 12 until all desired cluster hosts have been added to the selected nodes list.




14. On the right of the panel, expand the parallel grid tree to view all the simulation grids.
15. Click on a parallel grid in the tree to select it.

Note: You may select more than one grid by clicking and dragging the mouse or by clicking on multiple grids while pressing the 'Ctrl' or 'Shift' key.

16. Type the appropriate number on your selected nodes list.


The number typed will appear next to the grid name, and the remote location host name will appear on the same line. The color of the background behind the typed number will match the color of that number in the selected nodes list.

17. Repeat Steps 15 and 16 until all the grids have been assigned a node for processing.




18. Click **OK**.

The Assign Grids to Processes dialog box closes. Note that in the Job Submit panel, the Queue Name is posted as “NONE” because you are not using a queuing system to submit the job.



19. Click on **Start**.



The Job Submitter submits the job. The right panel of the Nexus® Desktop displays progress messages.

When the job is complete, the session tree lists the results files that are available for viewing. A timestep table report provides timestep outputs. A standalone report file provides details about the actions performed during the job. A Nexus report file (the .out file) provides comprehensive details of the data that has been generated.


Note: If the job failed, look for warnings and errors in the standalone file to determine what went wrong.

Remote Parallel Job with Queuing System (Linux Only)

The following instructions assume that you have access to a remote host with a queuing system installed on it. See the Nexus-VIP 5000.0.2 release notes for instructions on how to set up a queuing system.

The instructions assume that you are setting up a two-process job. You can modify the steps to set up more than two processes by keying in the desired number of processes that you want to use.

1. In the Job Settings panel, click on the Job Location navigation button.



The Select a Job Location dialog box appears.


2. From the Hosts pulldown menu, select a remote host that has a cluster configured for parallel job submission.

See [Remote Parallel Job](#) for a discussion of how to add hosts to the Hosts list.

Note: The **Copy VDB** option specifies that the Nexus .vdb will be copied to the remote job location when the job is run. Toggle this option ON if you have a local copy of the Nexus .vdb that you want to be written to the remote job location.


3. If you want the Nexus case to be initialized with map data using an i.map file, toggle the **Case Initializer** check button ON and select the appropriate i.map file.

4. In the Job Settings panel, click the **Parallel Mode** option to enable it.



Note: If a red X appears beside the **Parallel Mode** option when you toggle it ON, the X indicates that a parameter for the parallel option has not been set or has not been appropriately set.

5. Click the Open Parallel Mode button for the parallel mode option.




The Assign Grids to Processes panel will appear.

Note: If a message notifies you that a grid to process file could not be found, make certain that your model data is set up for a parallel job with a decomposed simulation grid.


6. On the left of the parallel settings panel, select the tab labeled "Queue."

Selectors will be displayed for setting the number of hosts and processes.



7. In the Hosts selector, click the up button to select two hosts.
8. In the Processes selector, click the down button to select one process.

Note: You should now see the names "QueueHost1" and "QueueHost2" in the selected nodes list.



9. On the right of the panel, expand the parallel grid tree to view all the simulation grids.
10. Click on a parallel grid in the tree to select it.

Note: You may select more than one grid by clicking and dragging the mouse or by clicking on multiple grids while pressing the 'Ctrl' or 'Shift' key.

11. Type the number '1' or '2'.

The number typed will appear next to the grid name, and the remote location host name will appear on the same line. The color of the background behind the typed number will match the color of that number in the selected nodes list.

12. Repeat Steps 10 and 11 until all the grids have been assigned a node for processing.
13. Click **OK**.

14. In the Job Submit panel, click on the **Queue Name** button.

Job Location: [cnxb201] /home/filer2home/snewland/apivar_runs

☐ Copy VDB

FCS File: [schubert] /home/filer2home/snewland/apivar_runs/apivar.fcs

☒ Case Initializer: [schubert] /home/filer2home/snewland/apivar_runs/apivar_nexu...

☒ Parallel Mode

Queue Name: NONE

Version: 5000_0_2

The Please Select a Queue dialog box displays.

15. Select a queue from the pulldown list and click **OK**.

Job Location: [cnxb201] /home/filer2home/snewland/apivar_runs

☐ Copy VDB

FCS File: [schubert] /home/filer2home/snewland/apivar_runs/apivar.fcs

☒ Case Initializer: [schubert] /home/filer2home/snewland/apivar_runs/apivar_nexusi.map

☒ Parallel Mode

Queue Name: Local_Queue

Version: 5000_0_2

Note: The “Local_Queue” is relative to the job location, not to the system where the job is being defined. If the job location is remote (as it is in this case), the local queue is also remote.

16. Click on **Start**.



The Job Submitter submits the job. The right panel of the Nexus® Desktop displays progress messages.

When the job is complete, the session tree lists the results files that are available for viewing. A timestep table report provides timestep outputs. A standalone report file provides details about the actions performed during the job. A Nexus report file (the .out

file) provides comprehensive details of the data that has been generated.

Note: If the job failed, look for warnings and errors in the standalone file to determine what went wrong.

Job Submission to a Windows HPC Server 2008 Cluster

This is the initial release of Nexus on Windows HPC Server 2008 clusters. Best practices for using this workflow are still evolving. Please contact Landmark Support for assistance if you wish to run Nexus on a Windows HPC Server 2008 cluster.

Using the Command Line

You can run the Nexus® simulator from a Linux shell or Windows command prompt, rather than using Nexus Desktop. You will need to set certain environment variables in order to make this easy to do. These are listed in the table below. The proper values for these variables are discussed in the *Nexus and VIP Release Notes*.

You can run the commands given in the following sections exactly as presented if the environment variables have been set properly and you replace the italicized items with their appropriate values.

The following environment variables must be set in order to the job submit commands from the command line.


Environment Variable	Description
STAND_EXE	Must be set to the appropriate Standalone executable. Setting must include the complete path.
DBMAP_EXE	Must be set to the appropriate DBMAP executable. Setting must include the complete path.
NEXUS_EXE	Must be set to the appropriate Nexus executable. Setting must include the complete path.
MPI_ROOT	Must be set to the root directory of the MPI installation. Typically this variable is set automatically when you install Nexus.
PATH	Must include settings for MPI bin directory. Typically this variable is set automatically when you install Nexus.

Detailed instructions for setting the environment variables appear in the *Nexus and VIP Release Notes*.

Parallel Environments

Nexus can run simulations in either single-process (serial) or parallel mode, in which the reservoir model is decomposed into a number of sub-grids. The computations for these sub-grids are run on one member of a set of processor cores or cluster nodes defined by the user. The *Nexus Keyword Document* discusses the assignment of grids to processors with the keyword GRIDTOPROC.

You can run Nexus in parallel mode in two system environments: a single computer with multiple processor cores, and a cluster with multiple computer systems (each of which might have multiple processor cores).



Parallel System Environments for Nexus

Using the Command Line under Windows XP

On Windows XP, you can run serial jobs and parallel single host jobs from the command line.

Running a Serial Job on Windows

1. Log into the system from which you will submit the job, open a Command Prompt window (**Start -> All Programs -> Accessories -> Command Prompt**), and change the working directory to the appropriate location.

The .vdb and .fcs files normally reside in this directory.

2. Enter the following command to run Standalone

```
%STAND_EXE% options fcs_file_name
```

Standalone is the initialization step of the workflow, which verifies the input data and transforms it to the binary unstructured-grid files that the Nexus® simulator requires. You must run Standalone every time you change input data, whether the change is in the .fcs file, in one of the files included by the .fcs file, or in a file included within a structured-grid file.

Use the following command-line options:

Option	Effect
-b	Forces recompilation of all data files into binary form. Ordinarily only files with changes need to be recompiled.
-s <i>arg</i>	Sets the name of the study (.vdb directory) for structured grid cases. This command line argument is required.
-c <i>arg</i>	Sets the output case name. This command line argument is required. The case name is used as the root name of the .rpt and .err files output by standalone, and, for structured grid cases, as the name of a subdirectory in the study directory containing the INIT class map data for the run.

<i>-o arg</i>	Forces all output files from Standalone to be located in the directory given as the argument. This will cause problems if two or more files have the same name but are in different directories.
---------------	--

3. (Optional) Enter the following command to load the i.map file into the .vdb.

```
%DBMAP_EXE% < mapin
```

where *mapin* is a file which contains the following data:

```
MAP study-name.vdb case-name
```

```
FILE casename_nexusi.map
```

```
EOF
```

```
STOP
```

If the .vdb directory does not exist at the time dbmap is run, it will be created.

Note: This step is optional. However, if you want your results to be visible in Nexus View, you will need to load the i.map file into the VDB. Standalone will create this file if you are running a structured-grid model. If you used SimDataStudio to create an unstructured-grid model, SimDataStudio should have created the file. In both cases, this file normally is located in the same directory as the case .fcs file.

4. Enter the following command to run the Nexus simulator.

```
%NEXUS_EXE% options fbc_file_name 2> errfile
```

Use the following command-line options:

Option	Effect
<i>-s arg</i>	Sets the name of the study (.vdb directory). This command line argument is required.

<i>-c arg</i>	Sets the output case name. This command line argument is required. The case name is used as the root name of all output files, such as the .out, .sum, etc., except for the study directory. Within the study directory, a subdirectory with the case name is used to contain the results of the run.
<i>-nocpu</i>	Eliminates the reporting of CPU time and memory information in the .sum and .max files. This is useful when trying to compare the results from two different cases.

A timestep summary will be printed to the screen for each successful timestep. You only need to include the “2> *errfile*” portion of the command if you want to save any system error messages to a file. If you do not include this, system error messages will be printed in your Command Prompt window.

Running a Single-Host Parallel Job

1. Log into the system from which you will submit the job, open a Command Prompt window (**Start -> All Programs -> Accessories -> Command Prompt**), and change the working directory to the appropriate location.

The .vdb and .fcs files normally reside in this directory.

2. Enter the following command to run Standalone

```
%STAND_EXE% options fcs_file_name
```

Standalone is the initialization step of the workflow, which verifies the input data and transforms it to the binary unstructured-grid files that the Nexus simulator requires. You must run Standalone every time you change input data, whether the change is in the .fcs file, in one of the files included by the .fcs file, or in a file included within a structured-grid file.

Use the following command-line options:

Option	Effect
--------	--------

-b	Forces recompilation of all data files into binary form. Ordinarily only files with changes need to be recompiled.
-s <i>arg</i>	Sets the name of the study (.vdb directory) for structured grid cases. This command line argument is required.
-c <i>arg</i>	Sets the output case name to <i>arg</i> . This command line argument is required. The case name is used as the root name of the .rpt and .err files output by standalone, and, for structured grid cases, as the name of a subdirectory in the study directory containing the INIT class map data for the run.
-o <i>arg</i>	Forces all output files from Standalone to be located in the directory given as the argument. This will cause problems if two or more files have the same name but are in different directories.

3. (Optional) Enter the following command to load the i.map file into the .vdb.

```
%DBMAP_EXE% < mapin
```

where *mapin* is a file which contains the following data:

```
MAP study-name.vdb case-name
FILE casename_nexusi.map
EOF
STOP
```

If the .vdb directory does not exist at the time dbmap is run, it will be created.

Note: This step is optional. However, if you want your results to be visible in Nexus View, you will need to load the i.map file into the VDB. Standalone will create this file if you are running a structured-grid model. If you used SimDataStudio to create an unstructured-grid model, SimDataStudio should have created the file. In both cases, this file normally is located in the same directory as the case .fcs file.

4. Run the following command to run the Nexus simulator.

```
mpirun -np N %NEXUS_EXE% -mpi options
```

```
fbc_file_name 2> errfile
```

where “-np” designates that the following number will specify the number of processors, and “N” is the actual number. (For example, you would specify “-np 2” to run a job on two processors.)

Use the following command-line options:

Option	Effect
-s <i>arg</i>	Sets the name of the study (.vdb folder). This command line argument is required.
-c <i>arg</i>	Sets the output case name. This command line argument is required. The case name is used as the root name of all output files, such as the .out, .sum, etc., except for the study directory. Within the study directory, a subdirectory with the case name is used to contain the results of the run.
-nocpu	Eliminates the reporting of CPU time and memory information in the .sum and .max files. This is useful when trying to compare the results from two different cases.

A timestep summary will be printed to the screen for each successful timestep. You only need to include the “2> *errfile*” portion of the command if you want to save any system error messages to a file. If you do not include this, system error messages will be printed in your Command Prompt window.

Using the Command Line under Linux

The commands in this section are appropriate for use with the C shell and shells compatible with it (csh and tcsh). To convert these to commands appropriate for Bourne-compatible shells (sh, bash, ksh, and zsh), the output re-direction should be changed from

```
>& screen_file
```

to

```
> screen_file 2>&1
```

Running a Serial Job

1. Log into the system from which you will submit the job, open a shell or terminal window, and change the working directory to the appropriate location.

The .vdb and .fcs files normally reside in this directory.

2. Enter the following command to run Standalone

```
$STAND_EXE options fcs_file_name
```

Standalone is the initialization step of the workflow, which verifies the input data and transforms it to the binary unstructured-grid files that the Nexus simulator requires. You must run Standalone every time you change input data, whether the change is in the .fcs file, in one of the files included by the .fcs file, or in a file included within a structured-grid file.

Use the following command-line options:

Option	Effect
-b	Forces recompilation of all data files into binary form. Ordinarily only files with changes need to be recompiled.
-s <i>arg</i>	Sets the name of the study (.vdb directory). This command line argument is required.
-c <i>arg</i>	Sets the output case name. This command line argument is required. The case name is used as the root name of the .rpt and .err files output by standalone, and, for structured grid cases, as the name of a subdirectory in the study directory containing the INIT class map data for the run.
-o <i>arg</i>	Forces all output files from Standalone to be located in the directory given as the argument. This will cause problems if two or more files have the same name but are in different directories.

3. (Optional) Enter the following command to load the i.map file into the .vdb.

```
$DBMAP_EXE < mapin
```

where *mapin* is a file which contains the following data:

```
MAP study-name.vdb case-name
FILE casename_nexusi.map
EOF
STOP
```

If the .vdb directory does not exist at the time dbmap is run, it will be created.

Note: This step is optional. However, if you want your results to be visible in Nexus View, you will need to load the i.map file into the VDB. Standalone will create this file if you are running a structured-grid model. If you used SimDataStudio to create an unstructured-grid model, SimDataStudio should have created the file. In both cases, this file normally is located in the same directory as the case .fcs file.

4. Enter the following command to run the Nexus simulator.

```
$NEXUS_EXE options fbc_file_name >& screen_file &
```

Use the following command-line options:

Option	Effect
-s <i>arg</i>	Sets the name of the study (.vdb directory). This command line argument is required.
-c <i>arg</i>	Sets the output case name. This command line argument is required. The case name is used as the root name of all output files, such as the .out, .sum, etc., except for the study directory. Within the study directory, a subdirectory with the case name is used to contain the results of the run.
-nocpu	Eliminates the reporting of CPU time and memory information in the .sum and .max files. This is useful when trying to compare the results from two different cases.

The >& *screen_file* portion of the command redirects all the screen output to a file. The output includes both the one-line summary of each timestep and any system error messages. The trailing & causes the job to be run in the background, disconnected

from the shell window. You can see the screen output by running the following command from the shell window:

```
tail -f screen_file
```

To exit from the tail command, press Ctrl-C.

Running a Single-Host Parallel Job

1. Log into the system from which you will submit the job, open a shell or terminal window, and change the working directory to the appropriate location.

The .vdb and .fcs files normally reside in this directory.

2. Enter the following command to run Standalone

```
$STAND_EXE options fcs_file_name
```

Standalone is the initialization step of the workflow, which verifies the input data and transforms it to the binary unstructured-grid files that the Nexus simulator requires. You must run Standalone every time you change input data, whether the change is in the .fcs file, in one of the files included by the .fcs file, or in a file included within a structured-grid file.

Use the following command-line options:

Option	Effect
-b	Forces recompilation of all data files into binary form. Ordinarily only files with changes need to be recompiled.
-s arg	Sets the name of the study (.vdb directory). This command line argument is required.
-c arg	Sets the output case name. This command line argument is required. The case name is used as the root name of the .rpt and .err files output by standalone, and, for structured grid cases, as the name of a subdirectory in the study directory containing the INIT class map data for the run.

<i>-o arg</i>	Forces all output files from Standalone to be located in the directory given as the argument. This will cause problems if two or more files have the same name but are in different directories.
---------------	--

3. (Optional) Enter the following command to load the i.map file into the .vdb.

```
$DBMAP_EXE < mapin
```

where *mapin* is a file which contains the following data:

```
MAP study-name.vdb case-name
FILE casename_nexusi.map
EOF
STOP
```

If the .vdb directory does not exist at the time dbmap is run, it will be created.

Note: This step is optional. However, if you want your results to be visible in Nexus View, you will need to load the i.map file into the VDB. Standalone will create this file if you are running a structured-grid model. If you used SimDataStudio to create an unstructured-grid model, SimDataStudio should have created the file. In both cases, this file normally is located in the same directory as the case .fcs file.

4. Enter the following command to run the Nexus simulator.

```
mpirun -np N $NEXUS_EXE -mpi options fbc_file_name
>& screen_file &
```

Use the following command-line options:

Option	Effect
<i>-s arg</i>	Sets the name of the study (.vdb directory) for structured grid cases. This command line argument is required.

<i>-c arg</i>	Sets the output case name. This command line argument is required. The case name is used as the root name of all output files, such as the .out, .sum, etc., except for the study directory. Within the study directory, a subdirectory with the case name is used to contain the results of the run.
<i>-nocpu</i>	Sets the output case name. This command line argument is required. The case name is used as the root name of all output files, such as the .out, .sum, etc., except for the study directory. Within the study directory, a subdirectory with the case name is used to contain the results of the run.

The `>& screen_file` portion of the command redirects all the screen output to a file. The output includes both the one-line summary of each timestep and any system error messages. The trailing `&` causes the job to be run in the background, disconnected from the shell window. You can see the screen output by running the following command from the shell window:

```
tail -f screen_file
```

To exit from the tail command, press Ctrl-C.

Running a Cluster Parallel Job (Linux)

1. Log into the system from which you will submit the job, open a shell or terminal window, and change the working directory to the appropriate location.

The .vdb and .fcs files normally reside in this directory.

2. Enter the following command to run Standalone

```
$STAND_EXE options fcs_file_name
```

Standalone is the initialization step of the workflow, which verifies the input data and transforms it to the binary unstructured-grid files that the Nexus simulator requires. You must run Standalone every time you change input data, whether the change is in the .fcs file, in one of the files included by the .fcs file, or in a file included within a structured-grid file.

Use the following command-line options:

Option	Effect
-b	Forces recompilation of all data files into binary form. Ordinarily only files with changes need to be recompiled.
-s <i>arg</i>	Sets the name of the study (.vdb directory) for structured grid cases. This command line argument is required.
-c <i>arg</i>	Sets the output case name. This command line argument is required. The case name is used as the root name of the .rpt and .err files output by standalone, and, for structured grid cases, as the name of a subdirectory in the study directory containing the INIT class map data for the run.
-o <i>arg</i>	Forces all output files from Standalone to be located in the directory given as the argument. This will cause problems if two or more files have the same name but are in different directories.

3. (Optional) Enter the following command to load the i.map file into the .vdb.

```
$DBMAP_EXE < mapin
```

where *mapin* is a file which contains the following data:

```
MAP study-name.vdb case-name
FILE casename_nexusi.map
EOF
STOP
```

If the .vdb directory does not exist at the time dbmap is run, it will be created.

Note: This step is optional. However, if you want your results to be visible in Nexus View, you will need to load the i.map file into the VDB. Standalone will create this file if you are running a structured-grid model. If you used SimDataStudio to create an unstructured-grid model, SimDataStudio should have created the file. In both cases, this file normally is located in the same directory as the case .fcs file.

4. Create a file *hostfile* containing a list of the names of the cluster nodes on which you want to run the job. These should be listed one per line.
5. Enter the following command to run the Nexus simulator.

```
mpirun -np N -hostfile hostfile -udapl $NEXUS_EXE
-mpi options fbc_file_name >& screen_file &
```

Use the following command-line options:

Option	Effect
<i>-s arg</i>	Sets the name of the study (.vdb directory). This command line argument is required.
<i>-c arg</i>	Sets the output case name. This command line argument is required. The case name is used as the root name of all output files, such as the .out, .sum, etc., except for the study directory. Within the study directory, a subdirectory with the case name is used to contain the results of the run..
<i>-nocpu</i>	Eliminates the reporting of CPU time and memory information in the .sum and .max files. This is useful when trying to compare the results from two different cases.

The `>& screen_file` portion of the command redirects all the screen output to a file. The output includes both the one-line summary of each timestep and any system error messages. The trailing `&` causes the job to be run in the background, disconnected from the shell window. You can see the screen output by running the following command from the shell window:

```
tail -f screen_file
```

To exit from the tail command, press Ctrl-C.

Writing Restart Data

“Restart data” is information written to disk during a job so that a subsequent job can be started at some designated point later than the original starting point. In Nexus®, restart data consists of a series of directories and files contained within a main restart directory -- “casename.rst.” A subsequent run that continues from a previous run is called a “restart run.”

See [Restarts](#) for a discussion of the restarts files and how they are structured.

Specifying Restart Data in the Runcontrol File

The user can direct Nexus® to write restart data by means of the OUTPUT RESTART keywords in the RUNCONTROL file.

Example RUNCONTROL data:

```
.
.
OUTPUT
  RESTART TIMES
ENDOUTPUT
.
.
TIME 01/01/2008 ! 2557 days
TIME 01/01/2009 ! 2922 days
TIME 01/01/2010 ! 3288 days
TIME 01/01/2011 ! 3653 days
STOP
```

In this example, restart data will be written at 2557, 2922, 3288, and 3653 days. The user can then restart a subsequent run by picking up and continuing from any of these times.

Specifying Restart Data with a Procedure

The user can use the procedures function “abandon (true)” in the SURFACE file to write restart data and stop the run.

Example SURFACE data:

```
.
.
PROCS
INTEGER I = 0, J = 0
REAL X
X = Q("OS",FIELD)
```

```
IF(X < 200.0) THEN

! When surface oil rate at connection FIELD drops
! below 200 STB/D,
! write restart data and stop the run.

ABANDON(TRUE)

ENDIF

ENDPROCS
.
.
```

Specifying Restart Data by an Interactive Request

A user can stop a simulation run at any point in time while a job is running and be able to restart from that time. To stop the run, you simply create a file with name “SIGNAL2” in the same directory as the casename.out file for the run. You create this file after launching the simulation run.

The SIGNAL2 file can be empty. Do not rename a useful file to SIGNAL2 since the file will be deleted automatically by the run. The simulator checks for the existence of file SIGNAL2 at the completion of each time step. If the simulator finds the SIGNAL2 file, it writes restart data and stops.

In the event that multiple simultaneous runs are writing output to the same directory, you can stop a specific run by creating the file “casename.SIGNAL2.” Simply creating a SIGNAL2 file will stop the first run that detects the file’s presence.

Specifying Restart Data Automatically

Restart data is automatically written when the simulation cannot converge. The run will progressively cut the time step size to help convergence, but if convergence is still not achieved with a time step size of $0.1 * DTMIN$, the simulator writes restart data and stops. This gives the user the opportunity to restart the run with changed data or simulator controls.

Generating a Nexus® Model with a Structured Grid

Nexus is normally a “gridless” simulator. It processes unstructured input data that includes pore volumes, connection transmissibilities, and well properties that have been precalculated by the SimDataStudio™ software or a conversion utility. Grid data and initialization properties that were used to calculate the pore volumes and transmissibilities are not included in the input data since they are no longer required.

However, you can use a SimDataStudio option to generate a Nexus input model that includes structured grid data and initialization parameters. These data are output in ASCII format so that they can be used for sensitivity analysis or as input to various decision management applications.

A structured Nexus model contains the following data which is not contained in an unstructured Nexus model:

- grid files
- grid physical parameters (porosity, permeability, compressibility, etc.)

A structured Nexus model does NOT contain the following data, which are contained in an unstructured Nexus model:

- property files that contain pore volume, connections, and x, y, and z indices
- input data already preprocessed in the SimDataStudio software

Note: Nexus case files that use the structured grid format incorporate some arrays that are adopted from VIP®.

Related Documentation

For an overview of the structured grid case file, see the “ASCII Structured Grid Case File” section of the *Nexus® Keyword Document*.

For a description of the arrays that are supported in the structured grid case file, see “Appendix A: Structured Grid Data” in the *Nexus Keyword Document*.

Generate Structured Grid Data in SimDataStudio™

The following steps assume that you have a SimDataStudio case that contains all of the grid data and reservoir property definitions needed to generate a standard set of Nexus® project files.

1. In the Nexus main window, select the case for which you want to generated structured data.

(Selecting the case will cause the SimDataStudio icon to become active.)

2. Click the SimDataStudio icon on the Nexus toolbar as shown below.




3. In the SimDataStudio software, perform the following steps.
 - a. Select **View -> User Options**.
 - b. At the top of the Default Options dialog box, scroll sideways and select the Nexus Case Generation-Initialization tab.
 - c. In the Grid Arrays subpanel, toggle the **Generate grid dependent data using structured grid input format** option ON.

This option causes the **Structured grid input format** option to be displayed in the Nexus Data Set Generation dialog box where you select the Nexus data to be generated. This option will be toggled on by default.

- d. Click **OK** to save the setting.
4. Select the **Simulation Data -> Generate Simulator Data File -> Nexus Files** option in the SimDataStudio main menu.

The Nexus Data Set Generation panel appears. It should look similar to the example shown below.



Note that the Property Data section typical of unstructured data output has been replaced with a Grid Data section. Property array files will not be created. Instead, a structured grid file is created. This file contains the grid and array data, along with the utility data that controls the array property calculations.

Note also that a Structured grid input format appears in the lower left corner of the dialog box. You enabled this option when you toggled on the **Generate grid dependent data using structured grid input format** option in Step 3.

5. Select the data that you want to export and then click on **Generate**.

The structured data that is created supports all of the VIP array input formats as well as the array modifiers. Although not apparent from the panel, the well data will also be modified such that it will appear in structured IJK format.

A sample structured grid case (.fcs) file is shown below.

```
DESC                                SPE COMPARISON PROBLEM #9
DESC                                3-D HETEROGENEOUS EXAMPLE
DESC                                APRIL 25, 1994
RUN_UNITS ENGLISH
DATEFORMAT MM/DD/YYYY

GRID_FILES

    STRUCTURED_GRID nexus_data/run_spe9-grid.dat
    OPTIONS nexus_data/run_spe9-prop_options.dat

ROCK_FILES

    ROCK Method 1 nexus_data/run_spe9-
rock_rock.dat

    RELPM Method 1 nexus_data/run_spe9-
prop_relm.dat

PVT_FILES

    PVT Method 1 nexus_data/run_spe9-pvt_pvt.dat

    SEPARATOR Method 1 nexus_data/run_spe9-
pvt_separator_dflt.dat

    WATER Method 1 nexus_data/run_spe9-
pvt_water.dat

INITIALIZATION FILES

    EQUIL Method 1 nexus_data/run_spe9-
init_equil.dat

RECURRENT_FILES

    RUNCONTROL nexus_data/run_spe9_runcontrol.dat

    WELLS Set 1 nexus_data/run_spe9_wells.dat

    SURFACE Network 1 nexus_data/
run_spe9_surface.dat
```

When this case file is processed by the simulator, Nexus will recognize the structured input data and will process it accordingly.

File Descriptions

The Nexus® simulator reads input information from a main input file -- the case file (.fcs). This is a master file and provides a link to all input information. In the Nexus model, each set of properties necessary to run simulation is saved in a separate file in Nexus keyword format. The Nexus case file .fcs provides links to those files. The case file has the following sections:

- Header with general information (i.e., RUN_UNITS, DATEFORMAT, etc.)
- PROPERTY_FILES or GRID_FILES
- INITIALIZATION_FILES
- ROCK_ARRAY_FILES
- ROCK_FILES
- PVT_FILES
- AQUIFER_FILES
- RECURRENT_FILES
- NET_METHOD_FILES

If the user requests structured grid data when generating the Nexus case file, the case file will have a different format than a standard case file. A case file that references structured grid data has a Grid_Files section instead of a Property_Files section. This alternative case file is used to create ASCII files for sensitivity analysis or as input to various decision management applications. For a detailed description of the structured and unstructured case files, please see the *Nexus Keyword Document*.

The Nexus case file is generated by the SimDataStudio™ software. Once the SimDataStudio software populates the .fcs with data necessary for generating a Nexus model, the same information can also be saved in a SimDataStudio .vds file. This file could be used later for editing parameters and regenerating the Nexus model.

The case file name extension must be “.fcs.” All other data files can have arbitrary names, extensions, and path locations.

Case files for different simulation runs can share common data files. Typically the data files referenced from within the case file are formatted ASCII files; however, they can also be FORTRAN binary files.

When you generate a Nexus model from the SimDataStudio software, the SimDataStudio software creates one Nexus case file (.fcs) and one

folder named “nexus_data.” It places all property files under that folder. If you have requested unstructured data, the SimDataStudio software also generates an i.map file and/or .vdb with INIT class data. The i.map and/or .vdb are located in the same folder as the .fcs file. This map data has information about simulation grids and properties. The map data also has links between I J K (structured) data and NB (unstructured) data. These allow you to display simulation results in the 3D Viewer application.

When submitting a job that uses structured data, the i.map file is generated in batch mode during Standalone.exe run.

Related Documentation

Details of the data structure of the Nexus® case file (.fcs) can be found in the “Case File” chapter of the *Nexus Keyword Document*.

Standard Input Data Processing


Nexus processes the standard (unstructured) input data files in two steps.

- On the first step, Nexus (standalone.exe) reads the data files, checks for data consistency and errors, and writes corresponding FORTRAN binary files (.fbc and .gbl files from the .fcs and .fbin files from all other ASCII files).

The FORTRAN binary files are placed in bin directories that are automatically created in a nexus_data folder at the path location of the ASCII data files. (If the data is generated using the structured format option, the files are placed in an sbin directory.) No binary files (.fbin) are generated for any referenced files that are already FORTRAN binaries.

- On the second step, Nexus (simulator.exe) reads and processes the FORTRAN binary files generated from the first step.

The following diagram shows the details of the input data processing.



Standard Input Data Processing

Structured Grid Data Processing

Structured grid data is a special type of Nexus data that includes grid data and reservoir initialization data in ASCII format. See "[Generating a Nexus® Model with a Structured Grid](#)" for a description of this type of data and how to generate it. A Nexus case file that includes structured data has a STRUCTURED_GRID entry in its case file (.fcs).

Nexus processes the structured input data files in two steps.


- On the first step, Nexus (standalone.exe) reads the data files, checks for data consistency and errors, and writes corresponding FORTRAN binary files (.fbc and .gbl files from the .fcs and .fbn files from all other ASCII files).

The FORTRAN binary files are placed in sbin directories that are automatically created in a nexus_data folder at the path location of the ASCII data files. No binary files (.fbn) are generated for any referenced files that are already FORTRAN binaries.

A .vdb folder, or an i.map file, or both are generated in the case directory, depending on what map formats are specified in the grid.dat file.

- On the second step, Nexus (simulator.exe) reads and processes the FORTRAN binary files generated from the first step.

The specified grid data file is processed as follows:



Structured Grid Data Processing

Input Files

The main Nexus input file is required to use the .fcs extension. All other input files have no name restrictions.

.fcs

This is the case file. It is the main file that defines a simulation run.

Output Files

The output files produced by Nexus have names that combine the case name and a data-specific suffix -- for example, “casesim.out.” The case name is specified by the “- c case_name” executable parameter. The Nexus output files and their content are described below.

Files Output by the Keyword Parser

The following files are generated by the standalone keyword parser.

.vdb

Database directory containing INIT class data. Map data are output in this format unless a different format is specified in the structured_grid.dat file. The selection of which map arrays are written to the .vdb for structured grid data is controlled by the MAPOUT keyword in the structured_grid.dat file. See “Appendix A: Structured Grid Data” of the *Nexus Keyword Document* for these keyword descriptions.

.err

This is an output file produced when the keywords are parsed and the data checked for consistency. If the data generates large numbers of warning messages, the first few warnings are output to the .rpt file, and the remainder to the .err file.

.rpt

This is the output file produced when the keywords are parsed and the data checked for consistency. It contains a report of the data read, error, and warning messages.

_nexusi.map

This is the map file containing INIT class data produced when structured grid data are processed. It is located at the path location of the .fcs file.

.srpt

This is the output file produced when structured grid data are processed. It is located in the sbin directory at the path location

of the STRUCTURED_GRID file. The contents of this file are also included in the .rpt file.

Files Output by the Simulator

_alias_names.xref

This file contains a cross reference of well names between the short and long forms. Short, i.e., 8 character well names are used for some output fields to be compatible with other application.

(.out.#, .wrn.#, .dbg.#, .log.#, where # is a process number)

These are debug files from parallel runs. They are created when the keywords DEBUG PARALLEL are specified in the RUNCONTROL file. See the DEBUG keyword in the Run Control section of the Nexus Keyword Document.

.alloc

This file can be present when a simulation run fails due to memory allocation error. It contains information about the size of memory requested.

.dbg

This file contains debug output. It is created when the keywords DEBUG LINEARIZATION HYDRAULICS are specified in the RUNCONTROL file.

.hyd

This file contains hydraulics output if requested. See the OUTPUT and HYDRAULICS keywords in the Run Control section of the Nexus Keyword Document.

.lcklog

This file contains a record of files locked and released by a simulation run. Nexus uses a file lock mechanism to prevent data from changing while the data is being read. File locking is a mechanism designed to prevent interference from simultaneous runs that share the same files.

nexus.map

This ASCII or FORTRAN binary file contains map data. The file is created if the map format is specified as MAPFORM or

MAPBINARY. The map data written is controlled by the MAPOUT and OUTPUT keywords in the RUNCONTROL file. (Note: Where map data gets written depends on whether the user specifies MAPFORM, MAPBINARY or MAPVDB and whether the vdb with INIT class data already exists. If the user specifies MAPVDB, but a .vdb with INIT class does not already exist, the map data will be written to a binary .map file.)

.max

This file contains a summary of maximum changes and damping for each Newton iteration. For an explanation of the entries in this report, refer to section “Description of Max File Entries” in the “Output Files” chapter of the Nexus Keyword Document.

.maxijk

This file contains a summary of maximum changes and damping for each Newton iteration, with i, j, k locations of grid blocks. This file is only present if an OPTIONS file containing STRUCTURED_MAPS data is specified in the .fcs file.

.netsum

This file reports the field production/injection summary (one line per time step), obtained by summing the production/injection at the network sinks/source. This file is only present if NETSUM is specified in the RUNCONTROL file. The volumes reported in the .netsum file may be different from those reported in the .sum file for various reasons, which include mixing of fluids in the network, separation or gas plant nodes in the network and different pvt/separator methods used in different parts of the network.

.osc

This file contains debug oscillation output. It is created when the keywords DEBUG OSCILLATION are specified in the RUNCONTROL file.

.out

This is the main output file from the simulator. It contains field summary data, region reports, recurrent data input reports, well, perforation, network, etc., reports as requested by the OUTPUT

keyword in the RUNCONTROL file. The field summary output is also written to the .sum file.

.plt

This ASCII or FORTRAN binary file contains plot data and is present if the plot format is specified as PLOTFORM or PLOTBINARY in the RUNCONTROL file. The plot data written is controlled by the OUTPUT, FIELDPLOT, NETPLOT, PERFLOT, REGIONPLOT, TARGETPLOT and WELLPLOT keywords in the RUNCONTROL file.

.rft

This file contains any Repeat Formation Tester (RFT) output that has been requested. The RFT outputs pressures at each perforation of the well. The frequency that data is written to the .rft file is controlled by the OUTOUT and RFTFILE keywords in the RUNCONTROL file. The RFTWELLS keyword in the Surface Network file specifies which wells to included in the report.

.rst

This directory contains restart data that can be read by a new simulation run desiring to pick up the exact simulation state at a specified time. The frequency of restart data in the .rst file is controlled by the OUTPUT and RESTART keywords in the RUNCONTROL file.

.runlog

This file contains a log of events that have occurred in the network, such as wells shutting in or reopening, actions taken as a result of procedures, etc.

.ss_field, .ss_network, .ss_regions, .ss_targets, .ss_wells

These files contain output from the simulator that is in a format suitable to be imported into a spreadsheet. Output type and frequency are controlled by the SPREADSHEET keywords in the RUNCONTROL file.

.sum

This file reports the field production/injection summary (one line per time step). The field production/injection is obtained by summing the production/injection from all the perforations in each well, and flashing them through the separator assigned to each well. The rates and cumulatives calculated in this manner should agree with the sum of the well rates and cumulatives, but may be different than what is ultimately produced/injected at the network sinks/sources

.tssum

This file reports information on how the time step size was determined by the simulator.

.vdb

Database directory containing map and plot data. Map and plot data are output in this format unless a different format is specified in the RUNCONTROL file. The i.map file description above also has details about when simulation map data does not get written to the .vdb database. The frequency of plot and map data written is controlled by the OUTPUT keyword. The selection of arrays mapped is controlled by the MAPOUT keyword. See the “Run Control” section of the *Nexus Keyword Document* for these keyword descriptions. The .vdb name is specified by the “-s study_name” executable parameter.

.wrn

This file contains warning messages that may be output by the simulator. It also contains statistics on memory usage.

Intermediate Files

The Nexus keyword parser creates FORTRAN binary files which are then read by the simulator. These files are placed in a bin or sbin directory at the path location of the keyword files. On subsequent runs, a FORTRAN binary file is recreated only if the timestamp of the keyword file is newer than that of the binary file.

.fbin

Files ending in .fbin are FORTRAN binary files created by the Nexus keyword parser. These files can also be created by the simulator for scratch output and will be automatically deleted when the run ends.

.fbc

This is the FORTRAN binary file equivalent of the .fcs file. It is the main file read by the simulator.

.gbl

This FORTRAN binary file contains global data read by the simulator.

.lck

Files ending in .lck are binary lock files that are used to prevent interference from simultaneous runs that use common data files.

Map and Plot File Formats

The MAP and PLOT files written out by Nexus can be post-processed by a variety of programs. To facilitate user access to these files, the content and formats are described below. Each line corresponds to a different record. The "Type" column specifies the variable type for each field in the record. The "Format" column specifies the format used to write the record to a formatted file.

The meanings of the symbols in the "Type" column are:

C*4	CHARACTER*4
C*6	CHARACTER*6
C*8	CHARACTER*8
C*20	CHARACTER*20
C*80	CHARACTER*80
I	INTEGER
R	REAL
R*8	DOUBLE

Map File

The MAP file contains the grid data and property data. INIT map file is the initialization data prepared for the simulation and RECUR map file is the time-dependent property data of the simulation. Each grid may have local grid refinement (LGR) and keywords PARENT, PAREA, and CHILDREN in the GRID section and are used to indicate the relationship between LGR and its parent grid. Keyword PARENT indicates the LGR grid's parent grid number; PAREA defines which portion of the coarse grid is to be refined; CHILDREN indicates the number of LGRs the current grid has. DAD and KID array data are used to indicate if a gridblock has LGR. DAD array defines the parent grid number for each gridblock and KID array defines the child grid number for each gridblock.

Grid numbering starts from 0. Multiple ROOT level grids are allowed, and these have PARENT = -1.

The MAP file may have some inactive cells where the CORP and property data are null values. DEADCELL array is used to indicate the inactive cells (0=active, 1=inactive).

Formatted MAP arrays are written using 6 numbers per line. The numbers are not always blank separated and may run together. The number field lengths can vary between arrays.

The possible simulator names in the map file are:
CORE, EXEC, NEXUS, SDS

Initialization Map File

Type	Format	Record
C*6	(A,1X)	MAP <FORM BIN> <map file version> <simulator><simulator version> [DUAL]
C*8	A8	CASE
C*256	A256	<case name>
C*8	A8	TITLE
I	I5	<# titles>
		Loop for each title
C*80	A80	<title i>, i=1,..., number of titles
		End loop for title
C*8	A8	COMPS
I	I5	<# comps>
		Loop for each comp
C*8	A8	<comp i>, i=1,..., number of comps
		End loop for comp
C*8	A8	UNITTYPE
C*8	A8	<units system>
C*8	A8	ORIGIN
R*8	E15.6	origin_x>
R*8	E15.6	<origin_y>
R*8	E15.6	<origin_z>
C*8	A8	HANDED
C*8	A8	<grid handedness>
C*8	A8	CLASS
C*8	A8	INIT
C*8	A8	DATE
I, R	3I52E13.6	<month> <day> <year> <timestep> <time>
		Loop for each grid
C*8	A8	GRID
I	I5	<grid number>
C*8	A8	NAME
C*8	A8	<grid name>
C*8	A8	GRTYPE
C*8	A8	<grid type>
C*8	A8	GRDIMENS
I	I5	<nx>
I	I5	<ny>
I	I5	<nz>
C*8	A8	STATUS
C*8	A8	<grid status>
C*8	A8	PARENT
I	I5	<parent grid number>
C*8	A8	PAREA
I	I5	<lgr imin>
I	I5	<lgr imax>
I	I5	<lgr jmin>
I	I5	<lgr jmax>
I	I5	<lgr kmin>
I	I5	<lgr kmax>
C*8	A8	CHILDREN
I	I5	<# child grids>
		Loop for each child grid
I	I5	<child grid i>, i=1,..., number of child grids
		End loop for child grid

```

                                Loop for each INIT array
C*8   A8       ARRAY
C*8   A8       <array name>
C*8   A8       DESCR
C*80  A80      <array description>
C*8   A8       UNITS
C*20  A20      <array units>
C*8   A8       TYPE
C*8   A8       <array data type>
C*8   A8       SIZE
I     I        <number of array values>
R     6E15.6   <array value i> , i = 1, ..., number of array values
      6E17.8e3 for CORP array
                                End loop for INIT array

                                End loop for grid

C*8   A8       STOP

```

Example

```

MAP    FORM    3.0.0  SDS    5000.0.3
CASE   pvlgrover
TITLE  1
Nexus Arrays for VDS Conversion of Case pvlgrover
COMPS  2
LIGHT  HEAVY
UNITTYPE ENGLISH
ORIGIN 0.0000000000000000.0000000000000000.0000000000000000
HANDED LEFT
CLASS  INIT
DATE   1      1      1980  0.0000 0.0000
GRID   0
NAME   ROOT
GRTYPE CARTSIAN
GRDIMENS10  10      3
STATUS ON
PARENT 0
CHILDREN1
1
ARRAY  UNPACK
DESCR  (PACKEDARRAY INDICATOR)
UNITS  (NONE)
TYPE   INTEGER
SIZE   300
0.100000E+010.200000E+010.300000E+010.400000E+010.500000E+010.600000E+01
0.700000E+010.800000E+010.900000E+010.100000E+020.110000E+020.120000E+02
0.130000E+020.140000E+020.150000E+020.160000E+020.170000E+020.180000E+02
0.190000E+020.200000E+020.210000E+020.220000E+020.230000E+020.240000E+02
0.250000E+020.260000E+020.270000E+020.280000E+020.290000E+020.300000E+02
.....
.....
STOP

```

Simulation Recurrent Map File

<u>Type</u>	<u>Format</u>	<u>Record</u>
C*6	(A,1X)	MAP<FORM BIN> <map file version> <simulator> <simulator version> [DUAL]
C*8	A8	CASE
C*256	A256	<case name>
[C*8	A8	CASPAR
C*256	A256	<parent name>]
C*8	A8	TITLE
I	I5	<# titles>
		Loop for each title
C*80	A80	<title i>, i=1,..., number of titles End loop for title
C*8	A8	COMPS
I	I5	<# comps>
		Loop for each comp
C*8	A8	<comp i>, i=1,..., number of comps End loop for comp
C*8	A	UNITTYPE
C*8	A8	<units system>
C*8	A8	ORIGIN
R*8	E15.6	<origin_x>
R*8	E15.6	<origin_y>
R*8	E15.6	<origin_z>
C*8	A8	HANDED
C*8	A8	<grid handedness>
C*8	A	CLASS
C*8	A8	RECURR
		Loop for each timestep
C*8	A8	DATE
I, R	3I52E13.6	<month> <day> <year> <timestep> <time>
		Loop for each grid
C*8	A	GRID
I	I5	<grid number>
C*8	A8	NAME
C*8	A8	<grid name>
C*8	A8	GRTYPE
C*8	A8	<grid type>
C*8	A8	GRDIMENS
I	I5	<nx>
I	I5	<ny>
I	I5	<nz>
C*8	A8	STATUS
C*8	A8	<grid status>
C*8	A8	PARENT
I	I5	<parent grid number>
C*8	A8	PAREA
I	I5	<lgr imin>
I	I5	<lgr imax>
I	I5	<lgr jmin>
I	I5	<lgr jmax>
I	I5	<lgr kmin>
I	I5	<lgr kmax>
C*8	A8	CHILDREN

```

      I      I5      <# child grids>
                                Loop for each child grid
      I      I5      <child grid i>, i=1,..., number of child grids
                                End loop for child grid

                                Loop for each RECURRE array
      C*8    A8      ARRAY
      C*8    A8      <array name>
      C*8    A8      DESCR
      C*80   A80     <array description>
      C*8    A8      UNITS
      C*20   A20     <array units>
      C*8    A8      TYPE
      C*8    A8      <array data type>
      C*8    A8      SIZE
      I      I      <number of array values>
      R      6E13.6 <array value i> , i = 1, ..., number of array values
                                End loop for RECURRE array

                                Loop for each well perforation
      C*8    A8      WELL
      I      I      <well number>
      C*8    A8      BLOCK
      I      I      <gridblock number>
      C*8    A8      PERFORD
      I      I      <perforation order>
      C*8    A8      WLNAME
      C*8    A8      <well name>
      C*8    A8      WLSTATUS
      I      I      <well status>

                                End loop for well perf block

                                End loop for grid

                                End loop for timestep

      C*8    A8      STOP

```

Plot File Format

The PLOT file contains the time-dependent injection/production data of wells, well perforations and surface facilities of the simulation.

The possible plot class names are:

WELL WLLYR GATHER FLOSTA AREA REGION FIELD
SURFACE NODE CONN CONNLIST TARGET SLPERF SLAM

The possible units system names are:

FIELD METRIC METBAR METKG LAB

Type	Format	Record
6C*6,C*8	(A,1X)	PLOT <FORM BIN><plot file version><simulator> <simulator version>[DUAL] <units system>
C*256	(A,1X)	<case name><case parent>
C*80	A80	<title 1>
C*80	A80	<title 2>
C*80	A80	<title 3>
I	(1X,I)	<# plot classes><day><month><year><nx><ny><nz><# comps>
C*8	(A,1X)	<class #1 name><class #2 name>...
I	(1X,I)	<# variables in class #1><# variables in class #2>... Loop for each class
C*4	18A4	TIME<variable #1 name><variable #2 name>... End loop for class Loop for each timestep Loop for each class
C*8	A8	Class name
R	5G15.5	<timestep #><time><# items><max # items><max # perfs> Loop for each item in class
C*8	8A8	<item name><item #><1st perf i><1st perf j><gc #><fs #><area #>
R	1P6E13.6	<variable #1 value><variable #2 value>... End loop for item End loop for class End loop for timestep
C*8	A8	STOP

Example

```
PLOT FORM 5.0.0 NEXUS 01.00 FIELD
Small impes_sat
Small Impes Black Oil Dataset
```

```
8      18      3      1994  1      1      1      2
```

```

WELL      WLLYR  NODE   CONN   REGION FIELD  CONNLIST      TARGET
33        29      2      15      28      28      11      25
TIME COP CGP CWP CGI  CWI  QOP QGP QWP QGI  QWI  BHP WPH WKH THP COWPQOWPGOR
WCUT WOR  QGLGCGLGQ1P Q1I  C1P C1I X1P Y1P Q2P Q2I  C2P C2I X2P Y2P
TIME QOP QGP QWP COP CGP CWP CGI  CWI  PRESHEADRW  SKINWI  WBC PWB QGI  QWI
COWPQOWPGOR WCUTQ1P Q1I  C1P C1I Q2P Q2I  C2P C2I
TIMEPNODTNOD
TIMEQGASQOILQWTRCGASCOILCWTRCBFGCBFOQBFWP_INPOUTT_INTOUTQ1  Q2
TIMECOP CGP CWP COI  CGI  CWI  PAVTPAVHOIP  GIP WIP  QOP QGP QWP  QOI  QGI  QWI
OIN GIN WIN SO   SG   SW  ORECGRECCIP Z2  Z2
TIMECOP CGP CWP CGI  CWI  QOP QGP QWP QGI  QWI  COWPQOWPGOR  ORECGRECPAVTPAVH
QGLGCGLGWCUTQ1P Q1I  C1P C1I Q2P Q2I  C2P C2I
TIMEQGASQOILQWTRCGASCOILCWTRCBFGCBFOQBFWQ1  Q2
TIMESQO SQG SQW SQL SQA SQH RQO RQG RQW RQL RQA RQH TSQOTSQGTSQWTSQQLTSQA
TSQHTRQOTRQGTRQWTRQLTRQATRQHP
WELL
572.      2800.0 9.0000 9.0000 1.0000
P-1      1      0      1      1      1      1      1
2.253355E+068.021942E+053.862873E+050.000000E+000.000000E+001.843999E+02
6.564635E+016.191530E+020.000000E+000.000000E+002.751833E+035.159651E+03
5.159650E+03 2.000000E+02 2.639642E+06 8.035528E+02 3.560000E-01 7.705193E-01
3.357665E+00 0.000000E+00 0.000000E+00 5.487237E+04 0.000000E+00 6.705369E+08
0.000000E+00 1.000000E+00 0.000000E+00 4.324529E+03 0.000000E+00 5.284547E+07
0.000000E+00 0.000000E+00 1.000000E+00
.....
.....
STOP

```

Region Summaries

The Nexus[®] software calculates reservoir volumes, pressures, and fluids in place by summing these properties in “regions” of the reservoir. The user can assign particular calculation methods to each region. These calculations can include the following:

- fluids in place
- average pressures
- surface volumes of phases
- compositions of phases

The simulator outputs region summaries at periodic user-requested times. The region summary lists data for each region about in-situ moles/mass in place, the reservoir volumes of phases, and the surface volume of phases, the compositions of the phases, and average pressures.

Region summaries appear in the .out output file

Related Documentation


Descriptions of the various keywords and parameters which pertain to regions can be found in the “Options Data” chapter of the *Nexus® Keyword Document*.

Procedures for defining regions in a SimDataStudio™ case can be found in the “Initialization Data” and “Defining Grid Array Data” chapters of the *Nexus SimDataStudio User’s Guide*.

Using Regions

The Nexus® software distinguishes between regions and region groups. Regions are individual groups of gridblocks. Region groups are sets of regions. A region group includes ALL of the gridblocks in the reservoir. However, all gridblocks do not need to be assigned to regions.

Region groups provide the ability to model multiple regions using different sets of properties and different calculations.



Region groups and regions are constrained by the following rules:

- At least one region group always exists for a project. This is the default region group, designated in the region group table as “FIELD.”
- Each region group may contain as many regions as there are gridblocks.
- There are no restrictions on the number of region groups.
- Every region belongs to a region group.
- Gridblocks may be assigned to any regions or none of them (by assigning the gridblock to region 0).
- A gridblock can only be assigned to a single region that exists within a region group.
- Within each region group, all values are totalled. Since some gridblocks may be excluded from a region group, the totals for a region group do not have to equal the reservoir totals.

The REGDATA Table

A region group corresponds to an individual IREGION data file defined in the .fcs file. All the regions defined within the IREGION data file belong to that region group. You can define the region group in the SimDataStudio™ software using the IREGION array settings in the Region Arrays panel.

Water, oil, and separator PVT methods are assigned to the various regions of the region group in the REGDATA table. This table resides in an OPTIONS input file that is identified in the PROPERTY_FILES section of the Nexus case file.

A typical REGDATA table is shown below:

```
REGDATA
NUMBER   NAME      IBAT
      1   Reg1      2
      2   Reg2      2
ENDREGDATA
```

Methods are assigned to regions as follows:

NUMBER entries correspond to the region numbers in the IREGION array.

NAME entries correspond to names that have been assigned to regions.

IBAT entries correspond to the numbers specified for separator methods in the PVT_FILES .dat file specified in the case file. The IBAT specification only affects compositional simulation.

Other possible entries in this table are IPVT for the numbers assigned to PVT methods and IWAT for the numbers assigned to water PVT methods. IPVT only affects black oil simulation. IWAT affects both black oil and compositional simulation.

Note: You assign the region numbers and names in the **SimDataStudio Utility Data -> Output Data** panel.

Differences in Implementation between VIP® and Nexus®

The VIP implementation of regions is different from the Nexus implementation in the following ways:

- A VIP project can have only one region group (i.e., one IREGION.dat file).
- In the VIP software, the sum of the individual region in-situ pore volumes and fluid volumes will equal the reservoir in-situ volumes. This may not be the case in the Nexus software, since the user may exclude some grid cells from the calculation by excluding them from the region group where the summing occurs.
- In the VIP software, extra regions can be defined with the XREG keyword.
- In the VIP software, a gridblock can be assigned to as many as three extra regions, and there may be as many as 99 regions and 99 extra regions.

The Nexus software equivalent to a VIP region would be a single region group with the same number of regions. A VIP “extra” region would be represented in the Nexus software by a region group containing a single region.

Controlling Surface Volume Calculations within a Region

Two factors control how surface volumes are calculated.

- The first factor is the specification of the PVT method number and the battery number. For black oil models without separators, the surface volume conversion information is totally included in the PVT method.
- The other factor is how the moles (or mass) of all the gridblocks in a region group calculated. There are two options: (1) The moles can first be summed and then flashed; or (2) the individual gridblock moles can be flashed and then summed. The second approach is much more costly, but is likely to be more accurate.

By default, the moles are summed and then flashed. To invoke flashing of individual gridblocks, the region group name must be included in the REGBLK table of the OPTIONS section.

Equilibrium

This chapter discusses the methodology employed by the Nexus[®] software for model initialization. The major topics covered are:

- Differences in modeling initialization between the Nexus software and the VIP[®] software
- Useful debugging output

Related Documentation

Descriptions of the various keywords and parameters which pertain to equilibrium can be found in the “Equilibration Data” chapter of the *Nexus® Keyword Document*.

Descriptions of the calculations performed when generating equilibrium data can be found in the “Equilibrium” chapter of the *Nexus Technical Reference Guide*.

Procedures for defining equilibrium tables in a SimDataStudio™ case (which can then be exported as a Nexus case file) can be found in the *Nexus SimDataStudio User's Guide*.

Differences between Nexus® and VIP®

The Nexus initialization algorithm differs significantly from the VIP initialization algorithm. In some circumstances, the different algorithms can produce moderately different in-place fluids and pressures. These differences are listed below:

- The VIP software uses evenly-spaced, reconstructed relative permeability and capillary pressure tables, whereas the Nexus software honors the input tabular data. The VIP reconstructed tables can produce different results for initialization of the transition zone, particularly when the capillary pressure is large and varies rapidly. This artifact can be virtually eliminated by decreasing the reconstructed table spacing. See the *VIP-CORE Reference Manual* for the appropriate dimension settings.
- For three-phase, gas-water-contact situations, the VIP software determines saturations by an inverse table lookup using the water-oil capillary pressure curves and the gas and water hydrostatic pressure difference. The Nexus software creates a synthetic gas-water capillary curve and performs the lookup based on the gas-water hydrostatic pressure difference. (The Nexus calculations are discussed in the “Mainstream Path” section of the “Equilibrium” chapter in the *Nexus Technical Reference Guide*.)
- The VIP software provides no special treatment for situations where a thin oil zone exists and the water transition penetrates the gas-oil contact. The Nexus software detects this situation and uses the synthetic gas-water capillary pressure for the inverse lookup.
- During initialization, the VIP software assigns pressure based on saturations determined prior to overreads. The Nexus software assigns pressures based on saturations determined after saturation overreads.
- In transition zone gas-water-contact situations, the VIP software assigns pressure based on the gas pressure and the gas-oil capillary pressure. The Nexus software assigns pressure based on the water pressure and the water-oil capillary pressure.

Useful Debugging Output

There are many output reports available for confirmation of results and debugging purposes. Entering the OUTSTART and DEBUG options in the run control file will generate these reports as Nexus outputs. An example of how to specify these options in the run control file is shown below.

```
OUTSTART
EQUILREGION      DEPTH TAB
                  GRIDBLOCK
                  OVERREAD
                  ALL
ENDOUTSTART

DEBUG
EQUILREGION
ENDDEBUG
```

At a minimum, the depth table report is very useful in confirming data entry and even PVT data, since the pressure profiles are dependent on the calculated density. You specify the depth table report as an output with the DEPTH TAB keyword. The resulting table is written to the .out file.

R5000.0.2	Equilibrium: Useful Debugging Output	139
Profrac Holding Corporation	PGR2023-00023	Ex. 1048, p. 139 of 276
U.S. Well Services, LLC		

196	8315.0	4013.4	4014.2	3993.1	0.9	20.3	38.8546	13.2931	60.4883	4014.2*	4014.7
GOC-PINIT->											
201	8320.0	4014.7	4014.7	3995.2	0.0	19.5	38.8542	13.2946	60.4884	4014.7	4014.7
206	8325.0	4016.0	4015.2	3997.3	-0.9	18.8	38.8549	13.2991	60.4887	4014.7	4014.7
211	8330.0	4017.4	4015.6	3999.4	-1.8	18.0	38.8556	13.3036	60.4889	4014.7	4014.7
216	8335.0	4018.7	4016.1	4001.5	-2.7	17.3	38.8564	13.3080	60.4892	4014.7	4014.7
221	8340.0	4020.1	4016.5	4003.6	-3.5	16.5	38.8571	13.3125	60.4894	4014.7	4014.7
226	8345.0	4021.4	4017.0	4005.7	-4.4	15.8	38.8578	13.3170	60.4897	4014.7	4014.7
231	8350.0	4022.8	4017.5	4007.8	-5.3	15.0	38.8586	13.3215	60.4899	4014.7	4014.7
236	8355.0	4024.1	4017.9	4009.9	-6.2	14.3	38.8593	13.3260	60.4901	4014.7	4014.7
241	8360.0	4025.5	4018.4	4012.0	-7.1	13.5	38.8600	13.3304	60.4904	4014.7	4014.7
246	8365.0	4026.8	4018.9	4014.1	-8.0	12.8	38.8607	13.3349	60.4906	4014.7	4014.7
251	8370.0	4028.2	4019.3	4016.2	-8.9	12.0	38.8615	13.3394	60.4909	4014.7	4014.7
256	8375.0	4029.5	4019.8	4018.3	-9.8	11.3	38.8622	13.3439	60.4911	4014.7	4014.7
261	8380.0	4030.9	4020.3	4020.4	-10.6	10.5	38.8629	13.3483	60.4914	4014.7	4014.7
266	8385.0	4032.2	4020.7	4022.5	-11.5	9.8	38.8636	13.3528	60.4916	4014.7	4014.7
271	8390.0	4033.6	4021.2	4024.6	-12.4	9.0	38.8644	13.3573	60.4919	4014.7	4014.7
276	8395.0	4034.9	4021.6	4026.7	-13.3	8.3	38.8651	13.3618	60.4921	4014.7	4014.7
281	8400.0	4036.3	4022.1	4028.8	-14.2	7.5	38.8658	13.3663	60.4924	4014.7	4014.7
286	8405.0	4037.6	4022.6	4030.9	-15.1	6.8	38.8666	13.3707	60.4926	4014.7	4014.7
291	8410.0	4039.0	4023.0	4033.0	-16.0	6.0	38.8673	13.3752	60.4928	4014.7	4014.7
296	8415.0	4040.3	4023.5	4035.1	-16.8	5.3	38.8680	13.3797	60.4931	4014.7	4014.7
301	8420.0	4041.7	4024.0	4037.2	-17.7	4.5	38.8687	13.3842	60.4933	4014.7	4014.7
306	8425.0	4043.0	4024.4	4039.3	-18.6	3.8	38.8695	13.3886	60.4936	4014.7	4014.7
311	8430.0	4044.4	4024.9	4041.4	-19.5	3.0	38.8702	13.3931	60.4938	4014.7	4014.7
316	8435.0	4045.7	4025.4	4043.5	-20.4	2.3	38.8709	13.3976	60.4941	4014.7	4014.7
321	8440.0	4047.1	4025.8	4045.6	-21.3	1.5	38.8717	13.4021	60.4943	4014.7	4014.7
326	8445.0	4048.4	4026.3	4047.7	-22.1	0.8	38.8724	13.4066	60.4946	4014.7	4014.7
WOC -->											
331	8450.0	4049.8	4026.8	4049.8	-23.0	0.0	38.8731	13.4110	60.4948	4014.7	4014.7
336	8455.0	4051.1	4027.2	4051.9	-23.9	-0.8	38.8738	13.4155	60.4951	4014.7	4014.7
341	8460.0	4052.5	4027.7	4054.0	-24.8	-1.5	38.8746	13.4200	60.4953	4014.7	4014.7
346	8465.0	4053.8	4028.2	4056.1	-25.7	-2.3	38.8753	13.4245	60.4955	4014.7	4014.7
351	8470.0	4055.2	4028.6	4058.2	-26.6	-3.0	38.8760	13.4290	60.4958	4014.7	4014.7
356	8475.0	4056.5	4029.1	4060.3	-27.4	-3.8	38.8767	13.4334	60.4960	4014.7	4014.7
361	8480.0	4057.9	4029.6	4062.4	-28.3	-4.5	38.8775	13.4379	60.4963	4014.7	4014.7
366	8485.0	4059.2	4030.0	4064.5	-29.2	-5.3	38.8782	13.4424	60.4965	4014.7	4014.7
371	8490.0	4060.6	4030.5	4066.6	-30.1	-6.0	38.8789	13.4469	60.4968	4014.7	4014.7
376	8495.0	4061.9	4031.0	4068.7	-31.0	-6.8	38.8797	13.4514	60.4970	4014.7	4014.7
381	8500.0	4063.3	4031.4	4070.8	-31.9	-7.5	38.8804	13.4558	60.4973	4014.7	4014.7
386	8505.0	4064.6	4031.9	4072.9	-32.7	-8.3	38.8811	13.4603	60.4975	4014.7	4014.7
391	8510.0	4066.0	4032.4	4075.0	-33.6	-9.0	38.8818	13.4648	60.4978	4014.7	4014.7
396	8515.0	4067.3	4032.8	4077.1	-34.5	-9.8	38.8826	13.4693	60.4980	4014.7	4014.7
401	8520.0	4068.7	4033.3	4079.2	-35.4	-10.5	38.8833	13.4738	60.4982	4014.7	4014.7
406	8525.0	4070.0	4033.8	4081.3	-36.3	-11.3	38.8840	13.4782	60.4985	4014.7	4014.7
411	8530.0	4071.4	4034.2	4083.4	-37.2	-12.0	38.8848	13.4827	60.4987	4014.7	4014.7
416	8535.0	4072.7	4034.7	4085.5	-38.0	-12.8	38.8855	13.4872	60.4990	4014.7	4014.7
421	8540.0	4074.1	4035.2	4087.6	-38.9	-13.5	38.8862	13.4917	60.4992	4014.7	4014.7
426	8545.0	4075.4	4035.6	4089.7	-39.8	-14.3	38.8869	13.4962	60.4995	4014.7	4014.7
431	8550.0	4076.8	4036.1	4091.8	-40.7	-15.0	38.8877	13.5006	60.4997	4014.7	4014.7
436	8555.0	4078.1	4036.6	4093.9	-41.6	-15.8	38.8884	13.5051	60.5000	4014.7	4014.7
441	8560.0	4079.5	4037.0	4096.0	-42.4	-16.5	38.8891	13.5096	60.5002	4014.7	4014.7
446	8565.0	4080.8	4037.5	4098.1	-43.3	-17.3	38.8898	13.5141	60.5005	4014.7	4014.7
451	8570.0	4082.2	4038.0	4100.2	-44.2	-18.0	38.8906	13.5186	60.5007	4014.7	4014.7
456	8575.0	4083.5	4038.4	4102.3	-45.1	-18.8	38.8913	13.5230	60.5010	4014.7	4014.7
461	8580.0	4084.9	4038.9	4104.4	-46.0	-19.5	38.8920	13.5275	60.5012	4014.7	4014.7
466	8585.0	4086.2	4039.4	4106.5	-46.9	-20.3	38.8928	13.5320	60.5014	4014.7	4014.7
471	8590.0	4087.6	4039.9	4108.6	-47.7	-21.0	38.8935	13.5365	60.5017	4014.7	4014.7
476	8595.0	4088.9	4040.3	4110.7	-48.6	-21.8	38.8942	13.5410	60.5019	4014.7	4014.7
481	8600.0	4090.3	4040.8	4112.8	-49.5	-22.5	38.8949	13.5455	60.5022	4014.7	4014.7
486	8605.0	4091.6	4041.3	4114.9	-50.4	-23.3	38.8957	13.5499	60.5024	4014.7	4014.7

```

491 8610.0 4093.0 4041.7 4117.0 -51.3 -24.0 38.8964 13.5544 60.5027 4014.7 4014.7
496 8615.0 4094.3 4042.2 4119.1 -52.1 -24.8 38.8971 13.5589 60.5029 4014.7 4014.7
501 8620.0 4095.7 4042.7 4121.2 -53.0 -25.5 38.8979 13.5634 60.5032 4014.7 4014.7

```

(Actual Grid Block Depth Range -- Min = 8130.0 [ft]; Max = 8620.0 [ft])

Equilibration Overread Data Report

The overread report can also be helpful. It details the capillary adjustment calculations. You specify the overread report as an output with the OVERREAD keyword. The report is written to the .out file.

```

-----
|
|                               Overread/Intsat Report ( Grid Block Values )
|                               Subgrid ( 1)   Equilibrium Region ( 1)
|                               PVT Type      ( 1)
|                               Water Type    ( 1)
|
|
|-----
|
|                               -Read- -Read- |                               -Read- |                               -Off -|...| | | | | |
|                               |           |           |           |           |           |           |           |           |
|                               | Po         Po         | Po         Po         | Po         Po         | Po         Po         |
| GLOBAL MDEPTH | So   So   |...| Sw   Sw   |[psia] [psia] |[psia] [psia] |[psia] [psia] |[psia] [psia] |...|P|SUBGRID
| NODE   [ft]   |(GBC ) (OVER) |...| (GBC ) (OVER) | (CALC) (OVER) | (GBC) (OVER) (ADJ) (USED) |...|M|NODE
|-----
3619 10062.2| 0.1185 0.1185|...|0.8815 0.8815|4003.4 4003.4| -2.75 -2.75 0.00 -2.75 |...|4| 1819
3620 10115.1| 0.1185 0.1185|...|0.8815 0.8815|4026.5 4026.5| -2.75 -2.75 0.00 -2.75 |...|4| 1820
3621 10168.0| 0.1185 0.1185|...|0.8815 0.8815|4049.6 4049.6| -2.75 -2.75 0.00 -2.75 |...|4| 1821
3622 10220.9| 0.1185 0.1185|...|0.8815 0.8815|4072.7 4072.7| -2.75 -2.75 0.00 -2.75 |...|4| 1822
3623 10273.8| 0.1185 0.1185|...|0.8815 0.8815|4095.8 4095.8| -2.75 -2.75 0.00 -2.75 |...|4| 1823
3624 10326.7| 0.1185 0.1185|...|0.8815 0.8815|4118.8 4118.8| -2.75 -2.75 0.00 -2.75 |...|4| 1824
3625 9110.0| 0.8326 0.8260|...|0.1674~0.1740|3629.3 3629.3|38.49 32.12 0.00 32.12 |...|1| 1825
3626 9162.9| 0.8301 0.8247|...|0.1699~0.1753|3649.9 3649.9|36.07 30.85 0.00 30.85 |...|1| 1826
3627 9215.8| 0.8276 0.8234|...|0.1724~0.1766|3670.6 3670.6|33.64 29.59 0.00 29.59 |...|1| 1827
3628 9268.7| 0.8251 0.8221|...|0.1749~0.1779|3691.2 3691.2|31.22 28.32 0.00 28.32 |...|1| 1828
3629 9321.6| 0.8225 0.8207|...|0.1775~0.1793|3711.9 3711.9|28.80 27.66 0.00 27.66 |...|1| 1829
3630 9374.5| 0.8165 0.8165|...|0.1835 0.1835|3732.5 3732.5|26.37 26.37 0.00 26.37 |...|1| 1830
3631 9427.4| 0.8084 0.8084|...|0.1916 0.1916|3753.2 3753.2|23.95 23.95 0.00 23.95 |...|2| 1831
3632 9480.3| 0.8004 0.8004|...|0.1996 0.1996|3773.8 3773.8|21.52 21.53 0.00 21.53 |...|2| 1832
3633 9533.2| 0.7833 0.7824|...|0.2167~0.2176|3794.5 3794.5|19.10 19.02 0.00 19.02 |...|2| 1833
3634 9586.1| 0.7583 0.7579|...|0.2417~0.2421|3815.1 3815.1|16.68 16.64 0.00 16.64 |...|2| 1834
3635 9639.0| 0.7145 0.7145|...|0.2855 0.2855|3835.8 3835.8|14.25 14.25 0.00 14.25 |...|2| 1835
3636 9691.9| 0.6543 0.6596|...|0.3457~0.3404|3856.4 3856.4|11.83 12.04 0.00 12.04 |...|2| 1836
3637 9744.8| 0.6485 0.6554|...|0.3515~0.3446|3877.1 3877.1| 9.41 11.87 0.00 11.87 |...|3| 1837
3638 9797.7| 0.6467~0.6512|...|0.3533~0.3488|3897.8 3897.8| 6.98 11.71 0.00 11.71 |...|3| 1838
3639 9850.6| 0.6449~0.6471|...|0.3551~0.3529|3918.4 3918.4| 4.56 7.42 0.00 7.42 |...|3| 1839
3640 9903.5| 0.6429 0.6429|...|0.3571 0.3571|3939.1 3939.1| 2.13 2.11 0.00 2.11 |...|3| 1840
3641 9956.4| 0.6366~0.6325|...|0.3634~0.3675|3959.7 3959.7| -0.29 -0.95 0.00 -0.95 |...|3| 1841
3642 10009.3| 0.1342 0.1343|...|0.8658 0.8657|3980.4 3980.4| -2.72 -2.72 0.00 -2.72 |...|3| 1842
3643 10062.2| 0.1185 0.1185|...|0.8815 0.8815|4003.4 4003.4| -2.75 -2.75 0.00 -2.75 |...|4| 1843
3644 10115.1| 0.1185 0.1185|...|0.8815 0.8815|4026.5 4026.5| -2.75 -2.75 0.00 -2.75 |...|4| 1844
3645 10168.0| 0.1185 0.1185|...|0.8815 0.8815|4049.6 4049.6| -2.75 -2.75 0.00 -2.75 |...|4| 1845
3646 10220.9| 0.1185 0.1185|...|0.8815 0.8815|4072.7 4072.7| -2.75 -2.75 0.00 -2.75 |...|4| 1846
3647 10273.8| 0.1185 0.1185|...|0.8815 0.8815|4095.8 4095.8| -2.75 -2.75 0.00 -2.75 |...|4| 1847
4201 9118.0| 0.8322~0.8258|...|0.1678~0.1742|3632.4 3632.4|38.12 1.93 0.00 31.93 |...|1| 2101
4202 9170.9| 0.8297~0.8245|...|0.1703~0.1755|3653.0 3653.1|35.70 30.66 0.00 30.66 |...|1| 2102
4203 9223.8| 0.8272~0.8232|...|0.1728~0.1768|3673.7 3673.7|33.28 29.39 0.00 29.39 |...|1| 2103
4204 9276.7| 0.8247~0.8219|...|0.1753~0.1781|3694.3 3694.4|30.85 28.13 0.00 28.13 |...|1| 2104
4205 9329.6| 0.8222~0.8205|...|0.1778~0.1795|3715.0 3715.0|28.43 27.60 0.00 27.60 |...|1| 2105
4206 9382.5| 0.8153 0.8153|...|0.1847 0.1847|3735.7 3735.7|26.01 26.01 0.00 26.01 |...|1| 2106
4207 9435.4| 0.8072 0.8072|...|0.1928 0.1928|3756.3 3756.3|23.58 23.58 0.00 23.58 |...|2| 2107

```

```

4208 9488.3| 0.7992~0.7981|...|0.2008~0.2019|3777.0 3777.0| 21.16 20.83 0.00 20.83 |...|2| 2108
4209 9541.2| 0.7795 0.7795|...|0.2205 0.2205|3797.6 3797.6| 18.73 18.73 0.00 18.73 |...|2| 2109
4210 9594.1| 0.7545~0.7532|...|0.2455~0.2468|3818.3 3818.3| 16.31 16.18 0.00 16.18 |...|2| 2110
4211 9647.0| 0.7054 0.7054|...|0.2946 0.2946|3838.9 3838.9| 13.89 13.89 0.00 13.89 |...|2| 2111
4212 9699.9| 0.6499~0.6590|...|0.3501~0.3410|3859.6 3859.6| 11.46 12.02 0.00 12.02 |...|2| 2112
4213 9752.8| 0.6482~0.6548|...|0.3518~0.3452|3880.2 3880.2| 9.04 11.85 0.00 11.85 |...|3| 2113
4214 9805.7| 0.6465~0.6506|...|0.3535~0.3494|3900.9 3900.9| 6.61 11.68 0.00 11.68 |...|3| 2114
4215 9858.6| 0.6447~0.6464|...|0.3553~0.3536|3921.5 3921.5| 4.19 6.54 0.00 6.54 |...|3| 2115
4216 9911.5| 0.6426~0.6422|...|0.3574~0.3578|3942.2 3942.2| 1.77 1.43 0.00 1.43 |...|3| 2116
4217 9964.4| 0.6345~0.6288|...|0.3655~0.3712|3962.8 3962.8| -0.66 -1.23 0.00 -1.23 |...|3| 2117
4218 10017.3| 0.1185 0.1185|...|0.8815 0.8815|3983.8 3983.8| -2.75 -2.75 0.00 -2.75 |...|3| 2118
4219 10070.2| 0.1185 0.1185|...|0.8815 0.8815|4006.9 4006.9| -2.75 -2.75 0.00 -2.75 |...|4| 2119
4220 10123.1| 0.1185 0.1185|...|0.8815 0.8815|4030.0 4030.0| -2.75 -2.75 0.00 -2.75 |...|4| 2120
4221 10176.0| 0.1185 0.1185|...|0.8815 0.8815|4053.1 4053.1| -2.75 -2.75 0.00 -2.75 |...|4| 2121

```

Notes: (a) '(GBC)' - determined by gravity-capillary equilibrium with grid-block-center as basis
 (b) '(OVER) or (INTS)' - result after processing overreads or INTSAT option
 (c) '~' indicates a differences exists between the two columns --
 > 0.10000E-03 for saturations and > 0.10000E-01 for pressure
 (d) '?' indicates that Pc adjustment does not restore original Pc by >0.10000E-03
 This may be desirable if, for example, overread/intsat saturations are mobile and.
 GBC based saturations were not.

Debug Report

The debug report provides extra information that can be useful for identifying potential equilibrium issues. For example, with compositional runs, you get the overall composition with depth report shown below.

You request the debug report by embedding the DEBUG keyword in the run control file. The debug report will be written to the directory where your projects case file (.fcs) resides. It will have the extension .dbg.

	Zi (Overall Composition)									
	Vs.									
	DEPTH									
	Equilibrium Region (1)									
	PVT Type (1)									
	Water Type (1)									

Table	DEPTH	Temp.	PSAT	PSAT						
Index	[ft]	[F]	(Input	(Calc.)	P1	P2	.	P4	P5	WATER
			[psia]	[psia]						
1	5030.0	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
2	5030.2	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
3	5030.3	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
4	5030.5	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
5	5030.6	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
6	5030.8	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
7	5030.9	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
8	5031.1	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
9	5031.2	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
10	5031.4	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
11	5031.5	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
12	5031.7	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
13	5031.8	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
14	5032.0	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
15	5032.1	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
16	5032.3	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
17	5032.4	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
18	5032.6	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
19	5032.7	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
20	5032.9	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
21	5033.0	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
22	5033.2	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
23	5033.3	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
24	5033.5	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
25	5033.6	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
26	5033.8	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
27	5033.9	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
28	5034.1	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
29	5034.2	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
30	5034.4	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
31	5034.5	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000

32	5034.7	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
33	5034.8	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
34	5035.0	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
35	5035.2	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
36	5035.3	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
37	5035.5	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
38	5035.6	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
39	5035.8	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
40	5035.9	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
41	5036.1	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
42	5036.2	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
43	5036.4	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
44	5036.5	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
45	5036.7	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
46	5036.8	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
47	5037.0	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
48	5037.1	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
49	5037.3	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
50	5037.4	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
51	5037.6	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
52	5037.7	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
53	5037.9	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
54	5038.0	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
55	5038.2	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
56	5038.3	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
57	5038.5	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
58	5038.6	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
59	5038.8	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000
60	5038.9	276.0	4389.46	4389.46	0.5627000	0.2429000	.	0.0842320	0.0134960	0.0000000

-----\

Reconciliation Report

Another detailed report is the reconciliation report shown below.

DETAILED Post-Reconciliation Report			Equilibrium Region (	1)
PVT	Type		(	1)
Water	Type		(	1)

NODE	DEPTH [ft]	TEMP [F]	Po [psia]	-Pre- PSAT [psia]	-Post- PSAT [psia]	So	Sg	Sw
1	5050.0	220.0	5000.0	4389.5	4389.5	0.8000	0.0000	0.2000
(Pre-Reconciliation)								
----- MASS/MOLE FRACTIONS -----								
COMPONENT	OVERALL (a)	Oil	Gas	Wat	Kv	Kw		
P1	0.56270	0.56270	0.58155	0.00000	1.0335	0.0000		
P2	0.24290	0.24290	0.24048	0.00000	0.99005	0.0000		
P3	0.09667	0.09667	0.09329	0.00000	0.96504	0.0000		
P4	0.08423	0.08423	0.07505	0.00000	0.89101	0.0000		
P5	0.01350	0.01350	0.00962	0.00000	0.71297	0.0000		
WATER	0.00000	0.00000	0.00000	0.00000	0.0000	0.0000		
(Post-Reconciliation)								
----- MASS/MOLE FRACTIONS -----								
COMPONENT	OVERALL (b)	Oil	Gas	Wat	Kv	Kw		
P1	0.56270	0.56270	0.58155	0.00000	1.0335	0.0000		
P2	0.24290	0.24290	0.24048	0.00000	0.99005	0.0000		
P3	0.09667	0.09667	0.09329	0.00000	0.96504	0.0000		
P4	0.08423	0.08423	0.07505	0.00000	0.89101	0.0000		
P5	0.01350	0.01350	0.00962	0.00000	0.71297	0.0000		
WATER	0.00000	0.00000	0.00000	0.00000	0.0000	0.0000		
(a) Interpolated Overall and Interpolated Phase Compositions.								
(b) Interpolated Overall and Equilibrium Phase Comps (From PSAT at interp. overall)								
-- Note reconciliation step for Supercritical case only involves PSAT calc.---								

Parallel Processing

Running jobs in the parallel mode enables faster simulation and often allows finer resolution models than running jobs in the non-parallel mode. The Nexus® software provides the ability to set up and run parallel simulation jobs with relatively few steps.

The Nexus software uses a domain-decomposition approach to parallelism. The entire Nexus model is subdivided or decomposed into collections of cells called “grids.” When the Nexus software is run in parallel, a specified number of Nexus processes is started. Each of the grids is assigned to one of these processes. The assignment can be specified by the user or determined by the Nexus software. Each process performs the model computations upon the grids which were assigned to it.

Some computations require data from grids assigned to other processes (for example, when a cell has one or more of its immediately adjacent neighbors in a different grid). In these cases, data from cells in a grid assigned to one process must at times be communicated to another process. The Nexus software uses MPI, a standard message passing interface, to transfer data between processes of a parallel Nexus job. The passing of data from one process to another never changes the assignment of a cell to a particular grid or the assignment of a grid to a particular process; only a copy of the needed data is transferred.

Related Documentation

A description of the debug option that pertains to parallel simulation can be found in the “Run Control” chapter of the *Nexus® Keyword Document*.

A description of the keywords used to assign grids to processes manually can be found in the “Options Data” chapter of the *Nexus Keyword Document*.

Descriptions of how to configure systems to run parallel jobs on Windows and on Linux can be found in Appendix A and Appendix B of the *Nexus and VIP® Release Notes*.

A discussion of the “grid coloring” algorithm, used for distributing multiple grids automatically to parallel processes, can be found in the *Nexus Technical Reference Guide*.

Decomposition of Data

The Nexus® software is designed as an "unstructured grid" simulator, in which the physical extent and location of the cells are not used by the Nexus software in the timestep computations. The relevant physical data required to run the simulation consists of the cell-to-cell connectivity graph, the pore volume of each cell, and the transmissibilities between connected cells. These properties are usually computed outside the simulator in the SimDataStudio™ software. However, they can be computed by the simulator itself if the user elects to export data from the SimDataStudio software in a structured grid format. The manner in which the data is decomposed for parallel simulation runs is different in these two cases.

Unstructured Grid Data

The standard Nexus® case file does not reference structured grid data. Its cell information has already been processed into a simple list, with no data indicating the location of the cell in the reference coordinate system. Accordingly, the grid assignment file consists simply of a list of grid numbers, whose length is the number of active cells. This file can be easily decomposed for parallel job submission.

Note: Unlike the VIP® software, any Nexus decomposition that is created automatically is done only along one coordinate direction. It is possible to create a grid decomposition file by hand, since the grid file consists only of a list of integers. However, this is not advised because in order to get good parallel performance, most cell-to-cell connections should be within a single grid. Almost all users will want to use grid assignment files.

Structured Grid Data

If the Nexus case file has been generated with the structured grid data option enabled in the SimDataStudio™ software, it references a text file which contains information about the simulation grid. The grid must be Cartesian, and logically a rectangular prism of cells, so that the grid can be described by the number of planes in each coordinate direction: NX, NY, and NZ. The underlying geometry does not have to be orthogonal in nature, but the connectivity of the cells must be that of a rectangular prism: each cell, with the exception of those on the outer boundary, must be connected to exactly six nearest neighbor cells, with two nearest-neighbor connections along each coordinate direction. Additionally, groups of these cells can be further subdivided by local grid refinements.

Once the grid is fully specified, it can be decomposed using the DECOMPOSE keyword. The decomposition can be viewed as a coarse representation of the grid, which must follow the guidelines given above for the grid itself. Unless instructed otherwise by an X or Y keyword inside the DECOMPOSE data section, the Nexus software will decompose each grid named in the DECOMPOSE block into a set of subgrids, attempting to equally distribute active cells among the subgrids.

Note: Unlike the VIP software, at this time the Nexus software decomposes the structured grid along only one coordinate direction only, even if X and Y keywords are given.

Assignment of Grids to Processes

Once each cell in the problem has been assigned to a grid, the grids must be assigned to individual parallel processes. As in the case of the grid decomposition, the Nexus software is capable of automatically generating a grid-to-process assignment. It is also straightforward for the user to assign grids to processors manually through the GRIDTOPROC keyword. Except for special cases noted below, it is permissible to assign any grid to any process. While there are many different permissible assignments, the performance of the parallel model is dependent upon the choice of assignments, so care must be taken if grids are assigned to processes manually.

Automatic Assignment

If multiple grids have been created, a parallel run has been initiated, and no GRIDTOPROC keyword is present in the OPTIONS data, the Nexus software will distribute the grids internally to its parallel processes. The algorithm which distributes the grids uses a concept named “grid coloring.” The goal of grid coloring is to allow an uncoupled set of grids to be distributed to the processes, which reduces the number of messages to be passed among the processes. However, the nature of the computed grid coloring can lead to cases where the load is not balanced as well as might be expected. If this occurs, it is advisable to switch to manual grid assignment.

For a technical discussion of “grid coloring,” see the *Nexus Technical Reference Guide*.

Manual Assignment

The GRIDTOPROC keyword and associated data is used to assign individual grids to processes for parallel runs. In its most basic form, the GRIDTOPROC data consists of a two-column list; one column lists the grids present in the model, the other lists the number of the process to which that grid is to be assigned. A third column is required for a multifield run, since the grids in each reservoir are numbered beginning at 1.

In all cases, if manual assignment is used, it must be used for all grids and reservoirs in the run. If manual grid assignment is desired for a multifield case, the GRIDTOPROC data must be entered at the level of the global OPTIONS file, not in the per-reservoir OPTIONS data. Any per-reservoir GRIDTOPROC data is ignored.

The assignment of grids to processes most strongly affects the performance of the linear solver, and most specifically the preconditioning chosen. Most of the preconditioner choices require a minimum of two grids per process. The exceptions to this rule are the algebraic multigrid preconditioners for the pressure solve (PRECON_AMG and PRECON_AMG_RS), and the ILU preconditioner for the full system solve (IMPLICIT_PRECON_ILU). These preconditioners will work properly with one grid per process.

Performance Implications of Grid Assignments

The performance of the Nexus® software when run in parallel is affected significantly by the decomposition of the model cells into grids, and the manner in which the grids are assigned to processes. Because the reservoir grids are coupled through the pressure equation, data from adjacent grids affects the solution within a grid. Several times within every timestep, information needs to be passed from each grid to all grids physically connected to it. The assignments of cells to grids, and grids to processes, must account for the physical connections between cells in the model. This is one reason why a typical decomposition often looks like a collection of building blocks: the inter-grid communications are minimized. Some considerations when designing a strategy for decomposition:

- The individual grids should be *logically convex*: i.e., in order to move from any cell in the grid to any other cell in the grid, the path with the fewest number of cells should be contained fully within the grid.
- If more than one grid is assigned to a process, grids assigned to the same process should be connected to one another. The Nexus software will recognize connected grids assigned to the same process and will not pass messages, since the data is already locally accessible.
- If more than one grid is assigned to a process, inasmuch as it is possible, the collections of grids assigned to each process should be logically convex. In addition to being logically convex, the collections of grids should minimize the number of off-process connections.
- Grids should be assigned to processes such that the total workload for each process is approximately equal. This can be difficult to do in models where some regions have single-phase flow and other regions have multi-phase flow. It is more straightforward to

balance the number of active cells per process, which often is a reasonable enough approximation.

Better data decomposition and grid-to-process assignments can greatly increase parallel efficiency, in some cases, doubling it for 20-process runs.

Limitations of Current Parallel Implementation

The interaction between the surface network and the reservoir has not lent itself to easy parallelization, so at this time the network computations are repeated on every process in the parallel job. This causes the parallel efficiency to be less than optimal, especially in cases with a very large number of processes or in cases where a large fraction of the time is spent in the network solution. For this reason, given any model, there is some number of processes beyond which the performance gain is negligible. A reasonable estimate of this maximum number of useful processes can be determined from Amdahl's law¹, which states that the speedup of a parallel run is the sum of the serialized portion of the runs and the time that the remaining parallel portion takes:

$$t_N = t_1 \left(s + \frac{1-s}{N} \right)$$

The parallel efficiency is then the ratio of the serial time to the product of the number of processes used and the parallel time:

$$E = \frac{t_1}{Nt_N}$$

It can be shown that if the efficiency E is $\frac{2}{3}$ or less, doubling the number of processors reduces the efficiency below $\frac{1}{2}$.

The number of processors which would result in an efficiency of $\frac{2}{3}$ is $1 + \frac{1}{2s}$, where s is approximated by doubling the fraction of combined time spent in the TOTAL NETWORK, FACILITY SOLVE, FACILITY FACTOR, and FACILITY SETUP computations as reported by the Nexus software in a serial run. As an example, if a serial model spends 5% of its time in these code sections, $s = 2 \times 0.05 = 0.1$, and the maximum number of processors estimated to keep the efficiency above two-thirds is 6. If a serial model spends more than 25% of its time in

1. Amdahl, G.M. Validity of single-processor approach to achieving large-scale computing capability, Proceedings of AFIPS Conference, Reston, VA. 1967. pp. 483-485.

the FACILITY and NETWORK computations, the benefit of running in parallel is minimal at best, as N will be at most 2. It must be noted that this estimate needs further evaluation, as it is based upon a limited number of cases.

The timestepping method used in the model also has a significant effect on the scalability of the model. Models run using IMPES (implicit pressure explicit saturation) timestepping tend not to scale as well as models run using IMPLICIT timestepping. For a given input dataset, the linear solver step in an IMPES run requires much less work per processor than does the same step in an IMPLICIT run, due to most of the unknowns being computed explicitly (outside the linear solver) in an IMPES run. The same number of messages is passed per iteration of the linear solver in the two cases, assuming the same preconditioner is used; because less work is being done per iteration, the communication cannot be overlapped with computation as effectively. This limitation is noted mostly in cases where there are fewer than 20,000 active cells per processor.

An additional limitation on parallel performance is imposed by the architecture of the typical microprocessors used for running the Nexus software. These multicore CPUs, and their predecessors, all share some part of the path between the CPU's caches and the system main memory. The Nexus software can use more than half of the memory bandwidth available between the CPU and the memory when only one core is used, so the scalability or efficiency of the code suffers when multiple cores are used in a single CPU. For best performance, only one core per CPU socket should be used, as this will minimize contention between CPUs and local memory. The newest quad-core CPUs and associated memory subsystems (available in 2008) can usually supply sufficient memory bandwidth to allow use of two of the four cores.

Well Data

The Nexus® software treats wells as part of the surface network and represents each well with nodes, connections between nodes, and perforations which connect nodes to gridblocks. The network representation for each well is automatically generated from the data input in the WELLSPEC table, which appears in the wells.dat file.

An example WELLSPEC table is shown below:

WELLSPEC well1									
CELL	KH	SKIN	RADW	RADB	DEPTH	FLOWSECT	PARENT	MD	MDCON
18	432.1	0.0	0.5	3.71	#	1	NA	5150	NA
231	364.2	0.0	0.5	3.71	#	1	NA	5250	NA
468	295.4	0.0	0.5	3.71	#	1	NA	5350	NA
92	489.9	0.0	0.5	3.71	#	2	1	5175	5010
359	395.5	0.0	0.5	3.71	#	2	NA	5275	NA
592	934.5	0.0	0.5	3.71	#	2	NA	5375	NA

This chapter describes some of the specialized uses of the WELLSPEC inputs.

Related Documentation

Detailed descriptions of the keywords used to specify wells can be found in the “Well Specification Data” chapter of the *Nexus® Keyword Document*.

Descriptions of the calculations performed when generating well data can be found in the “Well Data” chapter of the *Nexus Technical Reference Guide*.

Procedures for defining well tables in a SimDataStudio™ case can be found in the *Nexus SimDataStudio User's Guide*.

Well Data Tables and Files

Well data is input in a WELLSPEC table in the WELLS file. The table contains the following types of data:

- cell indices for each completion
- KH values (permeability * wellbore length) for each completion
- grid block radii for each completion
- wellbore radii for each completion
- skin factors for each completion
- well indices for each completion
- partial perforation factors for each completion
- relative permeability methods used for each completion
- sections to which the completions flow
- workover group assignments for completions
- zone assignments for completions
- angles open to flow for completions
- depths for sections of the wellbore
- x coordinates for sections of the wellbore
- y coordinates for sections of the wellbore
- temperatures for sections of the wellbore
- flow sections for completions
- measured depths for completions
- parent flow sections for completions
- measured depths where flow sections connect to parents
- lengths of perforated intervals

During simulation, WELLSPEC data can be input at multiple times for the same well in the WELLS file. Wells can be deepened, side-tracked, partially plugged, etc., at any time. For example, the WELLSPEC data might change dynamically due to well recompletion, where new perforations are added or existing perforations are plugged.

If a new WELLSPEC table is input for a well that already has well data, any previous data is completely replaced. If the new WELLSPEC table does not input an optional column that appeared in the previous table (for example, a SKIN column), the old values are not retained. They are replaced by the default value (0 in the case of SKIN).

Note: Even though a WELLSPEC table may have been input for a well, the well data will not be activated until the well is used in a WELLS table in the surface network file. The WELLSPEC data for a well referred to in a WELLS table must be specified at or before the time that the WELLS table is input.

Requirements for Unstructured Grids

For unstructured grids, the WELLSPEC table must contain the following:

- a cell ID, corresponding to the GLOBALCELL index for the perforated gridblock
- the effective gridblock radius (RADB)
- either the KH of the perforation (permeability*wellbore length), or the effective permeability (K) and the wellbore length (L)

Requirements for Structured Grids

For structured grids, the WELLSPEC table must contain the following:

- the i, j, k gridblock indices (IW, JW, L)
- the grid name if the model uses local refined grids (LGR)

With this information, the program will calculate KH and the effective gridblock radius radw.

Wellbore Geometry

The default geometry for the network wellbore representation is determined by the input for the SECT, FLOWSECT, DEPTH, X, Y, MD, PARENT and MDCON columns.

SECT - Column that specifies the section of the wellbore to which this perforation flows. Each section is modeled with a single wellbore node in the network, so SECT input can be used to control the number of wellbore nodes manually.

FLOWSECT - Column that specifies contiguous sections (flow sections) of the wellbore with common flow properties (such as tubing diameter).

MD - Column that specifies the measured depth (usually length measured following the well from a known surface reference point like Kelly Bushing). Measured depth should increase continuously in each flow section. The combination of FLOWSECT and MD uniquely identifies a point in the well, and can be used in a WELLNODE table in the surface network file to specify the locations of wellbore nodes. These nodes can be connected to model complex wells, such as smart wells.

DEPTH - Column that specifies the depth of the perforation. This defaults to the gridblock center.

X and Y - Columns that specify the x and y coordinate location of the perforation in the grid coordinate system.

PARENT and MDCON - The default network wellbore representation connects the perforations sequentially, starting with the first perforation specified in the WELLSPEC table. The PARENT and MDCON columns can be used to create more complex well configurations, such as multilateral wells, automatically. The PARENT column specifies the parent flow section that the perforation will connect to, and the MDCON column specifies the measured depth at which the perforation will connect. However, even if PARENT and MDCON are input, specification of any control equipment such as a valve still requires well nodes to be identified in a WELLNODE table, so that valve connections connecting well nodes can be input.

An example of the use of PARENT and MDCON appears below.

WELLSPEC well1									
CELL	KH	SKIN	RADW	RADB	DEPTH	FLOWSECT	PARENT	MD	MDCON
18	432.1	0.0	0.5	3.71	#	1	NA	5150	NA
231	364.2	0.0	0.5	3.71	#	1	NA	5250	NA
468	295.4	0.0	0.5	3.71	#	1	NA	5350	NA
92	489.9	0.0	0.5	3.71	#	2	1	5175	5010
359	395.5	0.0	0.5	3.71	#	2	NA	5275	NA
592	934.5	0.0	0.5	3.71	#	2	NA	5375	NA

If an elevation profile is not provided in a WELLBORE table in the surface network file, the X, Y, and DEPTH input is used to calculate the depth of each wellbore node, and the elevation change and length of wellbore connections in the network representation of the well.

If X and Y are not input, the length is the absolute value of the elevation change, which is appropriate if the well is vertical or a hydrostatic gradient is used for determining pressures in the wellbore (which is usually the case).

If an elevation profile is provided in a WELLBORE table, the length and elevation change is calculated from the elevation profile. In this case, the program does not check for consistency between the X, Y, DEPTH data and the elevation profile.

WELLSPEC Data Used by Procedures

The LENGTH and GROUP columns of the WELLSPEC table can be used by Nexus® procedures.

The LENGTH column is used to specify the length of the perforated interval in the gridblock. This can be used in conjunction with the permeability (K) to calculate the KH for the perforation. Note that this column is not used to determine connection lengths for the network representation of the wellbore, because this length may not span the entire gridblock, and grid blocks may be separated by nonproducing intervals. The perforation length is used by the wplug procedure function, which plugs a user-specified length of perforations depending on user-defined criteria.

The GROUP column is used to specify groupings of perforations, which can then be used by procedure functions such as GLR to retrieve the flow properties of the group of perforations. In this case, the function GLR(“GROUP”,well_name) would return as the GLR of each perforation, the result obtained by summing the production of all perforations with the same GROUP index.

WELLSPEC Data Used by Drawdown/Buildup Constraint

The ZONE column can be used to identify the section of the wellbore to which a drawdown or buildup constraint is applied. The drawdown/buildup constraint (DPBHMX keyword) restricts the production/injection of a well, so that the maximum drawdown or buildup at any perforation is no more than the specified amount. The allowable drawdown may vary with the conditions in the wellbore.

It may be desirable to exclude perforations with little or no productivity from the drawdown/buildup constraint. In this case, ZONE input can be used to identify which perforations should be included in the drawdown/buildup calculation.

WELLSPEC Properties That Can Be Modified by Procedures

The workover procedure function provides the capability to modify the PPERF, KH, SKIN, WI and RADW values for any perforation, based on logic that the user can supply in a procedure.

Modifying WELLSPEC Data Using WELLMOD

Well data can be input at multiple times by respecifying the entire WELLSPEC table, in which case all previous data for the well is cleared and must be input again (including data that can be defaulted). Alternatively, for some completion properties, WELLMOD input can be used to change only the specified property, without affecting the other properties that were previously input.

Production and Injection Controls

This chapter describes the production and injection constraints options in detail.

The Nexus® software imposes controls on production and injection in two ways:

- constraints defined for individual network connections
- constraints for groups of connections defined by network nodes

Related Documentation

Detailed descriptions of the keywords used to specify injection and production constraints can be found in the “Surface Network” chapter of the *Nexus® Keyword Document*.

Descriptions of the calculations performed when imposing constraints can be found in the “Network Calculations” chapter of the *Nexus Technical Reference Guide*.


Procedures for defining injection and production constraints in a SimDataStudio™ case can be found in the *Nexus SimDataStudio User's Guide*.

How Production and Injection Are Modeled

The Nexus[®] software models the well and surface facility system as a set of nodes with the following types of connections:


- connections from the nodes to reservoir grid cells
- connections from the nodes to sinks
- connections from the sources to nodes

Production is modeled as connections from nodes to sinks.



Production

Injection is modeled as connections from sources to nodes.



Injection

Constraints for Individual Connections

Production and injection can be constrained by specifying maximum flow rate constraints for individual network connections. These constraints are defined in the CONSTRAINTS table, which appears in the surface.dat input file.

The following constraints can be defined for individual connections:

- maximum surface oil rate
- maximum surface gas rate
- maximum surface water rate
- maximum surface liquid rate
- in-situ maximum oil rate
- in-situ maximum gas rate
- in-situ maximum water rate
- in-situ maximum liquid rate
- in-situ maximum total rate of all components
- in-situ maximum total rate of all hydrocarbons
- maximum reverse surface oil rate
- maximum reverse surface gas rate
- maximum reverse surface water rate
- maximum reverse surface liquid rate
- in-situ reverse maximum oil rate
- in-situ reverse maximum gas rate
- in-situ reverse maximum water rate
- in-situ reverse maximum liquid rate
- in-situ reverse maximum total rate of all components
- in-situ reverse maximum total rate of all hydrocarbons
- total reservoir barrel rate for all fluid phases combined

See the “Network Rate Constraints” section of the “Surface Network” chapter of the *Nexus® Keyword Document* for detailed description of the CONSTRAINTS table and the keywords that apply to network connections.

Flow Control Devices

When a constraint is input for a particular network connection, and the connection is capable of producing more than that rate, a flow control device that restricts flow is implemented. This flow control device imposes a pressure drop sufficient to reduce the rate to the specified constraint. In the Nexus® software, the specification of a constraint implies that a control device is present.

By default, the Nexus software will assume that the control device is located in one of the following places:

- the connection for which a rate constraint is specified
- the connection immediately downstream of the node for which a minimum pressure is specified
- the connection immediately upstream of the node for which a maximum pressure is specified.

The Nexus® software does not need to know what this device is and does not need to model it in detail. It needs only to compute the pressure drop necessary to reduce the rate to the specified constraint, or to reduce or increase the node pressure to the specified value.

Surface Rate Constraint Calculations

Surface rate constraints are evaluated by feeding the mass flow rates for the connection to the separator assigned to the connection (or by using the stock tank properties in a black oil table if no separator is assigned). This would be the rate that would be obtained by connecting the connection to a test separator.

The separator data is defined in the SEPNODES table. In many cases, the production from many connections is commingled before being separated, so the sum of the surface flow rates of the connections feeding a separator battery may not sum to the offtake from the separator battery. The Nexus® software provides a COMMINGLE option for the SEPNODE table. This option instructs the test separation calculation to use the partial mass/molar volumes from the common separator node to which multiple connections flow. If you use the COMMINGLE option, the sum of the surface flow rates of connections feeding a separator battery will be equal to the sum of the offtake.

In situ rate constraints are evaluated by flashing the mass flow rates for the connection at the average temperature and pressure of the inlet and outlet nodes.

The default treatment of constraints is to control them locally -- i.e, in the same connection for which a rate constraint is specified (or in the upstream or downstream connection for a pressure constraint). However, the Nexus software does provide for constraints to be controlled remotely. For example, the bottom hole pressure can be controlled by a control device in the offtake connection from the well head node. The only limitations on this capability are that the location of the control device does permit control of the specified constraint.

Targets for Group Connections

Production and injection for group connections can be constrained by using “targets.” You assign maximum or minimum pressure constraints to the targets and use methods to assign percentages to the individual connections to which the targets refer. The constraints are defined in the TARGET table, which appears in the surface.dat input file.

See the “Production and Injection Targets” section of the “Surface Network” chapter of the *Nexus® Keyword Document* for detailed description of this table and its keywords.

Targets permit you to assign constraints that are governed by conditions that you impose with methods. For example, a gathering center might have a maximum water handling capacity of 100,000 STBD, which requires the production from the wells feeding into the gathering center be controlled so that this limit is not exceeded. This is illustrated below.




Illustration of Targeting

Since the Nexus software allows constraints to be defined at any point in the network, you could assign a water rate constraint of 100,000 STBD to the outlet connection from the gathering center, *gc_offtake*. The constraint would implement a control device located at the connection *gc_offtake*, which would reduce the production rate from all five wells by applying a back pressure. However, this type of

constraint would probably be too simplistic, since it would arbitrarily reduce the production of all the wells, when alternative scenarios are possible. For example, four wells could produce the 100,000 STBD, with one well shut in. Because they can not model complex conditions, constraints should be used only used when they control a single well, i.e., only when they control a linear series of connections flowing to or from a single well.

Instead of using a constraint at `gc_offtake`, you could use a target. For example, the water production target could be apportioned between each of the wells and controlled by adjusting control devices at the wellhead connection for each well (connections `wh1`, `wh2`, `wh3`, `wh4` and `wh5`). This strategy would result in the program having five degrees of freedom with which to satisfy a single condition. In addition to specifying the target rate and the controlling connections, you would have to specify a method that uniquely determines what fraction of the target rate to allocate to each well. For example, to control water production rate, you might specify that the highest water cut wells be shut in until the target is not exceeded.

The Nexus software provides a flexible way of specifying targets. The target inputs specify how to calculate the target rate and how the target rate is to be controlled. Specifying the target rate can be as simple as specifying a specific flow rate.

Target Rate Calculation

The target rate is calculated using the results of the first network solution (if there are no procedures) or the network solution following execution of the procedures. This network solution occurs at fixed reservoir pressure and perforation fluid mobilities. All recurrent data for the time step has been processed, so any changes to the network and wells are incorporated, but the previous Newton iterations target allocations are still applied. (On the first Newton of the first time step, there are no target allocations). This is an explicit calculation. For example, if the target rate is calculated from the sum of production of several wells, it will be calculated at the start of the Newton iteration, and will differ to some degree from the sum of production of those wells at the end of the Newton iteration.

The methods available to calculate the target rate are listed below. These methods are defined in the CALCMETHOD column of the TARGET table. The keywords are column heading keywords in a TARGET table.

- A specified value. (Assigns a particular value as the target rate.)
- The sum of the flow rates of a specified phase in the list of CALCCONS connections, evaluated at the conditions specified by CALCCOND.

Note that flow rates are positive values when flow is in the direction from the inlet node to the outlet node, so the sum of a list of connections that contains both injectors and producers will be additive. The injection rates will be added to the production rates, assuming that all connections flow in the forward direction.

- The sum of the potential flow rates of a specified phase in the list of CALCCONS connections, evaluated at the conditions specified by CALCCOND.
- The sum of the net production or injection flow rates of a specified phase in the list of CALCCONS connections, evaluated at the conditions specified by CALCCOND.

For net production, the flow rates in connections which are part of the injection network are subtracted from flow rates in connections which are part of the production network. For net injection, the flow rates in connections which are part of the production network are subtracted from flow rates in connections which are part of the injection network.

- The total injection or production of a specified phase, into a specified region, evaluated at the conditions specified by CALCCOND (which could be the average region pressure and temperature).
- The injection (or production) rate required to maintain a specified pressure in a specified region. The specified pressure can be a pore volume weighted average, or a hydrocarbon pore volume weighted average. Details of the algorithm used are given below.
- The injection (or production) rate required to ensure that net voidage is replaced in a specified region. Details of the algorithm are given below.

Control of the Target Rate

In order to allocate the target rate among the controlling connections, the maximum flow rate that each connection is capable of making must be known. To determine the maximum flow rates, the Nexus® software solves the network with all targets removed, but with all constraints still active.

This potential solution gives the maximum mass flow rates in every connection in the network. The potential solution is then summed over all connections in the CTRLCONS connection list for the phase specified by CTRL and at the conditions specified by CTRLCOND to get the potential control rate.

If the potential control rate is less than the target rate, no control is needed. If the potential control rate is greater than the target rate, the connections in the CTRLCONS connection list are constrained, so that the sum of the constrained production equals the target rate.

These allocated constraints are fixed values, which are honored implicitly, so the target rate will be exactly honored as long as all the connections remain capable of producing the rate that they are allocated when the full reservoir or network solution is computed.

If there is more than one target, the reductions in flow rates necessary to meet a target may reduce the potential flow rate for subsequent targets and, therefore, the allocations that the subsequent target calculates. As a result, the order in which the targets are processed can affect the result. This order can be set using PRIORITY settings in the TARGET table.

The Nexus software provides ten methods of allocating the target rate, plus a gas lift optimization method that is discussed separately.

Method	Action
SCALE	Scales back connection rates in proportion to their productivity or injectivity.
SCALEQP	Scales back connection rates in proportion to their productivity or injectivity
AVG	All connections which are not limited by a pressure constraint or a maximum rate constraint flow at the same average rate.
GUIDERATE	All connections which are not limited by a pressure constraint or a maximum rate constraint are scaled back in proportion to their specified guide rate.
WCUT	Reduces the production rate of the highest watercut wells to the user-specified minimum rate until the target is reached.
GOR	Reduces the production rate of the highest GOR wells to the user-specified minimum rate until the target is reached.
WRATE	Reduces the production rate of the highest water rate wells to the user-specified minimum rate until the target is reached.
GRATE	Reduces the production rate of the highest gas rate wells to the user-specified minimum rate, until the target is reached.
AVGGOR	Resets the rates of the highest GOR wells to an average rate.
AVGWCUT	Resets the rates of the highest water cut wells to an average rate.

The GUIDERATE, AVG, SCALE, SCALEQP, AVGGOR, and AVGWCUT all use the same basic algorithm to allocate the target rate, where the fraction of the target rate allocated to each of the controlling connections is proportional to a guide rate.

- For the GUIDERATE method, the guide rate is input by the user or calculated by a simple formula.
- For the AVG method, the guide rate is 1.
- For the remaining methods, the potential flow solution is used, but adjusted slightly to allow for production decline.
- For the SCALE method, the guide rate is set equal to the adjusted potential flow rate, which has been reduced due to the action of any previously applied targets.
- For the SCALEQP method, the guide rate is the adjusted potential flow.
- For the AVGGOR method, the guide rate is the potential oil/gas ratio, and for the AVGWCUT method, the guide rate is the potential oil/(water+oil) ratio.

The WCUT, GOR, WRATE and GRATE methods also all use the same basic algorithm to allocate the target rate, differing only in the criteria used to determine which wells to shut in, or scale back. In each case, the criteria is calculated using the adjusted potential network solution, without any reduction due to the effect of other targets. The criteria are:

- the water cut of each connection for the WCUT method
- the gas/oil ratio for the GOR method
- the water rate for the WRATE method
- the gas rate for the GRATE method.

Using this criteria, the connections with the highest values are shut in, until the target rate is not exceeded, then the last connection shut in is reopened at a reduced rate so that the target rate is exactly met.

The minimum allowed flow rate for a control connection can also be specified, in which case the flow in these connections will not be reduced below the specified minimum. However, if as a result, the target rate would be exceeded, then some connections will be shut in.

Run Control

During simulation, the Nexus[®] software advances the reservoir solution in time while honoring the timing of changes to data input and obtaining output at selected points. As long as the correct answer is produced, the scheme employed should minimize work performed (as measured by the CPU time consumed). There are trade-offs between speed and solution quality and, while the default settings will give good performance for most problems, you may need to adjust them in some cases.

Three levels of control exist:

- timestep control
- non-linear or Newton iteration parameters
- iterative solver parameters

The monitoring and adjustment of performance at the first two levels is described in this chapter.

Related Documentation

Detailed descriptions of the keywords used to specify run control parameters can be found in the “Run Control” chapter of the *Nexus® Keyword Document*.

Procedures for defining run control parameters in a SimDataStudio™ case can be found in the “Simulation Control” chapter of the *Nexus SimDataStudio User’s Guide*.

Time Step Control

The simulator moves the reservoir model through a succession of time points. The interval of time between points is called the “timestep.” Within each timestep, the simulator performs Newton iterations to resolve the non-linearity in the pressure- and saturation-dependent properties while determining the change in pressure and component masses, which occur over the timestep. For each Newton-iteration, an approximate linear system of equations is constructed, which can be solved by either direct or iterative matrix solvers. Due to their high efficiency, iterative solvers are used almost exclusively.

The simulator can select its own timesteps. They are constrained only by the maximum component mass/mole changes scaled by the total cell mass/mole, specified on the DCMAX keyword, or by cell throughput limitations when the IMPES formulation is used. (See the IMPSTAB keyword description in the *Nexus® Keyword Document*.) Under automatic timestep control, the maximum mass/mole fraction constraint may be exceeded to save the work required to repeat the timestep.

Timesteps are controlled by the run control parameters set in the runcontrol.dat input file. They are defined by the following Nexus keywords: DT, DCMAX, DCRPT, VOLRPT, MAXNEWTONS, MAXBADNETS, and IMPSTAB.

These keywords are discussed in the following sections of this chapter and described in detail in the *Nexus Keyword Document*.

Note: If you do not explicitly set these parameters when you build a case in SimDataStudio™, they will not appear in a Nexus run control file that you output from the SimDataStudio. In this case, the Nexus software will simply use the defaults for these keywords, where appropriate. See the *Nexus Keyword Document* for the default settings.

Primary Timestep Controls (DT and DCMAX)

The DT keyword sets up primary timestep controls and timestep size limits. Constant timestepping, which is seldom used, is selected with the CON keyword, which also defines the step size. Automatic timestepping, which is the recommended mode, is selected with the AUTO keyword, which also specifies the initial timestep size. Subsequent timestep sizes are calculated by linear interpolation or extrapolation of the maximum change in cell component mass/mole normalized by the total cell mass/mole (dcmx), over the previous timestep, so that it will not exceed the value specified by the DCMAX keyword.

Timestep size is altered automatically to coincide exactly to the times or dates specified on the TIME keyword, even when fixed timesteps are used. The number of steps over which the adjustment is performed is user controlled by the ADJUSTTOTIME keyword. This eliminates very small timesteps that would otherwise be needed to exactly hit the specified time.

Under automatic timestep control, the computed maximum component mass/mole fraction, dcmx, is allowed to exceed the constraint value specified for DCMAX, to save the work required to repeat the timestep. By default, dcmx is allowed to exceed DCMAX by a factor of two, but the factor may be changed with the DCRPT keyword. When dcmx exceeds $DCMAX * DCRPT$, the timestep size will be cut and repeated. However, after two timestep cuts, the DCMAX constraint will not be honored and a timestep will be accepted. Additionally, when running with the implicit formulation, the DCMAX constraint violation will only trigger a timestep size cut and repeat for timesteps that converge on the first Newton iteration.

The DCMAX and DCRPT values can be set different by grid or reservoir. The DCMAX value can also be set differently by IMPES or IMPLICIT grids.

Automatic timestep sizes are bounded by the minimum and maximum values specified with the MIN and MAX keywords. By default, timestep size increases are limited to a factor of 2; however, the increase factor can be set by the MAXINCREASE keyword. Values of MAXINCREASE should be set greater than 1. A value of 1 will timestep the entire run with the initial size value, and values less than 1 will reduce the timestep until the MIN value is reached.

Timesteps will automatically reset down to specified values each time a network connection QMAX constraint changes. Different values can be specified for injection and production connections with the keywords QMAXINJ and QMAXPROD, respectively.

Timesteps are automatically set to specified values each time a WAG injection well switches cycles. The keyword WCYCLE is used to specify the initial timestep size when a well switches to the water cycle. The keyword GCYCLE is used to specify the initial timestep size when a well switches to the gas cycle.

The timestep control parameters used by a run are printed in the casesim.out file in the 'Run Control Data' table, an example of which follows.

Run Control Data

```
-----
MO/DAY/YR:  01/01/2008                0.00000  DAYS
-----
```

Time step sizes :

```
Initial ..... 1.000E-02
Minimum ..... 1.000E-05
Maximum ..... 3.100E+01
Qmaxprod ..... 1.000E+30
Qmaxinj ..... 1.000E+00
Max Increase .... 2.000E+00
Adjust to Time ..      5
```

Timestep Cuts

Timestep cuts are used to compensate for errors. Two distinct conditions can trigger a timestep cut:

1. violation of the maximum allowable mass/mole change
DCMAX*DCRPT *and* convergence failure of the non-linear iterations
2. maximum volume error VOLRPT for grids with non-EOS fluids.

After each failure, the timestep size is cut by a factor and the timestep is retried. The cut factor is 2 for the first cut and n squared for any subsequent n th cut.

The maximum number of timestep cuts allowed is five. In general, the final timestep cut in a sequence results in the minimum timestep size possible, assuming that DT MIN has been set. A convergence failure on this "last" timestep causes an attempt to save the run by using one tenth of this size ($0.1 * \text{MIN}$). If the timestep is still not acceptable, the run is terminated and restart data written at the end of the last successful timestep.

Timestep Size Statistics

The Nexus® software outputs a report of the timestep size statistics to a file casesim.tssum. This file contains information on how each timestep size was determined, the values of the limiting parameters, and the related keywords that control the limits. At the end, the file contains a summary report of the number of occurrence for each time size limit type. This is followed with the values for the maximum, minimum, and average timestep sizes for the run.

Example timestep size report -- file casesim.tssum

```
*****
*
*                               *
*          NEXUS      TIMESTEPSIZESUMMARY      REPORT          *
*                               *
*                               *
*                               *
*                               *
*****

-----LIMITING-----
TIMESTEP  TIMESTEP  SIZEFROM      TO      REASON FOR      NCUTS  GLOBAL  RELATED
NO.      (DAYS)    (DAYS)    (DAYS)    TIMESTEP SIZE      GRID  IJK   VALUE  PARAMETER
-----
1         1.0000    0.0000    1.0000    SPECIFIED/DEFAULT  0      -    -    -    (dtime)
2         2.0000    1.0000    3.0000    MAX INCREASE      0      -    -    -    -
3         4.0000    3.0000    7.0000    MAX INCREASE      0      -    -    -    -
4         8.0000    7.0000   15.0000    MAX INCREASE      0      -    -    -    -
```

5	16.0000	15.0000	31.0000	MAX INCREASE	0	-	-	-	-
6	8.0771	31.0000	39.0771	DCMAX LIMIT	0	1	100	-0.594 (dmax)	
7	16.154	39.077	55.2314	MAX INCREASE	0	-	-	-	
8	7.0464	55.2314	62.2778	DCMAX LIMIT	0	1	99	-0.688 (dmax)	
9	14.0927	62.2778	76.3705	MAX INCREASE	0	-	-	-	
10	28.1854	76.3705	104.5559	MAX INCREASE	0	-	-	-	

*
*
*
*
*
*
*
*
*
*
*

NEXUS TIMESTEP SIZE SUMMARY STATISTICS REPORT

REASON FOR NUMBER OF TIMESTEP SIZE	% OF TOTAL OCCURRENCES	
DCMAX LIMIT	37	19.3717
DCMAX CUT	0	0.0000
IMPES STABILITY LIMIT (S	0	0.0000
IMPES STABILITY CUT (S	0	0.0000
IMPES STABILITY LIMIT (M	0	0.0000
IMPES STABILITY CUT (M	0	0.0000
IMPES STABILITY LIMIT (d	0	0.0000
IMPES STABILITY CUT (d	0	0.0000
DTQMAX(Production Wells)	0	0.0000
DTQMAX(Injection Wells)	0	0.0000
DTMIN	4	2.0942
DTMAX	101	52.8796
TIME CARD	10	5.2356
MAX INCREASE	38	19.8953
SPECIFIED/DEFAULT	1	0.5236
PMAX OR PMIN VIOLATED	0	0.0000
CONVERGENCE FAILURE	0	0.0000
DAMP FACTOR TOO SMALL	0	0.0000
NEGATIVE MASS	0	0.0000
RESTART VALUE	0	0.0000
ATTEMPT TO SAVE RUN 0.1*	0	0.0000
DTWCycle	0	0.0000
DTGCycle	0	0.0000
*(Influenced by TIMECARD)	40	20.9424
Totals (excluding *)	191	100.0000

Maximum Time Step Taken [day]	30.0000
Minimum Time Step Taken [day]	1.0000
Average Time Step Taken [day]	19.1257

Newton or Non-linear Iteration Controls (TOLS)

The Nexus[®] software uses a numerical formulation similar to that described by Shiralkar, et al.^{1, 2} and uses a generalized compositional approach^{3, 4}. The primary reservoir cell equations are the mass conservation equations for each component, and a volume constraint equation. The primary Newton variables for both the IMPES and IMPLICIT cases are the cell pressures and component masses. The formulation is such that it conserves mass exactly. However, fluid volumes and component mobility are non-linear functions of cell pressure and component masses. The non-linearity is handled with a Newton iteration procedure.

The convergence criteria for both IMPES and IMPLICIT is the difference between the cell pore volume and the volume of the occupying fluids (normalized by cell pore volume), also called “cell volume error.” By default, a timestep converges when the maximum cell volume error is less than 0.005. The volume error tolerance is specified with the TOLS VOLCON keyword.

With the IMPLICIT formulation, mass balance errors can be introduced because the iterative linear solver does not solve exactly, but only solves within a specified tolerance. By default, mass balance errors are not used to control convergence; however, they can be used when the TOLS MASS data is specified. A timestep will not be considered converged when the material balance residual is more than the specified tolerance.

The maximum number of Newton iterations that are allowed during a timestep is controlled by the keyword MAXNEWTONS. Exceeding

1. Shiralkar, G.S., Stephenson, R.E., Joubert, W., Lubeck, O. and van Bloemen Waanders, B.: "Falcon: A Production Quality Distributed Memory Reservoir Simulator", SPE 51969, SPE Reservoir Evaluation & Engineering, October 1998.

2. Shiralkar, G.S., Fleming, G.C., Watts, J.W., Wong, T.W., Coats, B.K., Mossbarger, R., Robbana, E. and Batten, A.H.: "Development and Field Application of a High Performance, Unstructured Simulator with Parallel Capability", SPE 93080, SPE Reservoir Simulation Symposium, (31 January-2 February 2005)

3. Young, L.C., and Stephenson, R.E.: "A Generalized Compositional Approach for Reservoir Simulation", SPEJ (October 1983) 727; Trans., AIME, 275.

4. Watts, J.W.: "A Compositional Formulation of the Pressure and Saturation Equations", SPERE (May 1986) 43; Trans., AIME, 281

the maximum number of iterations will trigger a timestep cut. (See above.)

Control of Maximum Changes During Implicit Newton Iterations

To improve convergence for implicit timesteps, limits may be applied to the changes of various variables during iterations. If greater changes occur during a Newton-iteration for any of these variables, the solution for that iteration will be damped.

The following types of limits can be set:

- fluid volume change for each newton (DVOLLIM)
- maximum phase mass/mole fraction change for each newton (DZLIM)
- maximum saturation change for each newton (DSLIM)
- maximum pressure change for each newton (DPLIM)
- maximum fractional change in phase mobility for each newton (DMOBLIM)
- maximum emerging gas saturation for each newton (DSGLIM)
- maximum flow out of a cell with zero mobility (NEGFLOWLIM)
- relative permeability damping (on by default) (KRDAMP)
- saturation tolerance for grid blocks whose gas saturation crosses S_{gc} during a newton (SGCTOL)
- saturation tolerance for perforated grid blocks whose gas saturation crosses S_{gc} during a newton (SGCPERFTOL)
- line search algorithm for improving implicit convergence (off by default) (LINE_SEARCH)

These variables are described in detail in the *Nexus® Keyword Document*.

Well and Network Controls on Newton Iterations

The Nexus® software tests for counter current component flow at well perforations and will not converge if the condition exists. This is a non-physical solution (no perforation capillary pressures are considered). However, if allowed, the Nexus software can improve performance with little affect on well production and injection. The PERFREV ALLOW keyword allows for convergence even when counter current component flow exists. Note that this may cause inconsistencies in the reported perforation rates, like producing water while injecting gas.

The Nexus software also provides controls for when inactive network constraints or inactive hydraulics relationships (i.e., hydraulics tables, pressure drop correlations, or static gradients) are violated within a Newton-iteration. By default, the Nexus software will not converge when inactive constraints or inactive hydraulics relationships are violated. The keywords OVERSHOOT ALLOW will allow convergence as long as the violation is not more than a specified fraction, as defined by the OVERSHOOT MAX and OVERSHOOT HYDMAX keywords. If the violation is beyond the MAX or HYDMAX values, the solution is damped and another Newton performed.

An example of an inactive constraint is a producing well which is operating on a rate constraint (the active constraint), but which also has a minimum bottomhole pressure constraint (the inactive constraint). The simulator will exactly honor the active constraint, which may result in the inactive bottomhole pressure constraint being violated after the pressures are updated at the end of a Newton-iteration. If overshoot is allowed, and the fraction by which the bottomhole pressure constraint is violated is less than MAX, the simulator will proceed to the next timestep (assuming all other convergence criteria are satisfied). The next timestep will then impose the minimum bottomhole pressure constraint as the active constraint, and the rate constraint will become inactive. Because the constraint is only violated for one timestep, allowing the overshoot usually will not make a material difference to the simulation results, and it saves the CPU expense of an extra Newton iteration.

Similarly, when a rate constraint is active, the hydraulics equation for that connection is not active. The rate constraint will be exactly honored, which implies that a valve has restricted flow in the connection, introducing a pressure drop over and above that which would be calculated from the hydraulics relationship for that connection. If, after the pressures are updated at the end of a Newton-iteration, the pressure drop due to this implied valve actually turns out to be a pressure increase, the hydraulics equation has been violated, and

the hydraulics equation should have been active instead of the rate constraint.

The keyword MAXBADNETS controls the maximum number of Newton iterations taken because of network reasons that are allowed during a timestep. Network reasons include damping due to perforation flow reversals (Perf_Rev), maximum perforation cell pressure change (Cell_Dp), network pressure constraint violation (Net_Pmax), bad network solutions (Bad Net) and any inactive constraint overshoot. The terms in parentheses are printed in the casesim.max and casesim.maxijk files. Exceeding the maximum number of MAXBADNETS iterations will trigger a timestep cut.

Newton Iteration Statistics

The Nexus® software outputs reports of the Newton iteration statistics to the files casesim.max and casesim.maxijk. These files contain a summary of maximum changes and damping for each Newton iteration. For an explanation of the entries in these reports, refer to section "Description of Max File Entries" in the "Output Files" chapter of the *Nexus Keyword Document*. The file casesim.maxijk contains the cell i, j, k locations for the maximum changes. It is created only if an OPTIONS file containing STRUCTURED_MAPS data has been specified in the .fcs file.

IMPES Stability Controls (IMPSTAB)

The IMPSTAB keyword is used to control the simulator treatment of IMPES stability restrictions. By default, the Nexus® software does not perform IMPES stability tests to control timestep sizes.

However, two stability algorithms are available. The first is based on the work of K. Coats^{5, 6} and the second is similar to what is used in VIP®⁷.

Using an IMPSTAB option often involves a trade-off of performance versus enhanced stability. When no stability issues are suspected, the option should not be used because its algorithms will usually select more conservative timestep sizes than would otherwise be used, resulting in slower performance. However, when stability issues are a problem, one of the options should be tried. Of the two, the COATS option is more conservative in timestep size selection and may be more effective in some cases.

When these algorithms are active, the IMPSTAB calculated timestep size competes with the other limiting timestep quantities described in the sections above to determine the new timestep size. The timestep summary report (example above) will show the controlling entity.

5. Coats, K.H.: "IMPES Stability: The CFL Limit", SPE 66345, SPE Reservoir Simulation Symposium, Houston, Texas, (Feb. 11-14, 2005)

6. Coats, K.H.: "IMPES Stability: The Stable Step", SPE 69225, SPE Reservoir Simulation Symposium, Houston, Texas, (Feb. 11-14, 2005)

7. *VIP Technical Reference Manual.*

Multiple Grids and Reservoirs

Some of the control parameters mentioned in this chapter can be defined by grid (or by reservoir for a multi-field run). Encapsulating the corresponding input data with the keywords GRID and ENDGRID will apply the data to the specified grids or reservoir. For details, see the *Nexus® Keyword Document* on GRID and ENDGRID.

For runs with multiple grids, the Newton convergence criteria are applied to each individual grid and a timestep is converged when all grids converge. Similarly, the automatic timestep size is computed from the most limiting grid.

Multiple grid (or multireservoir) cases can be complicated because each grid may present a different degree of numerical difficulty. Since each grid can potentially represent a very different problem, the largest permissible timestep for the different grids can vary significantly, and the timestep taken will be the smallest of the maximum permissible among the individual grids. For this reason, you should model all grids with the same level of implicitness; otherwise, a small timestep needed by an IMPES grid will drive the fully implicit grids, and result in computation and CPU inefficiency. A secondary effect is that the grid with the most non-linear behavior (i.e., requiring the most number of Newton iterations) will determine the number of Newton iterations for all the grids.

The high performance of the simulator enables shared timestepping with the simultaneous simulation of the large number of cells associated with multiple coupled reservoirs, thereby maintaining the tight coupling and consistency across the entire system. With multiple grids, no difference in approach occurs between simulation of a single grid and simulation of a number of coupled grids, other than the potentially larger number of cells.

Simulation Statistics

Various run statistics are printed at the end of the output file casesim.out. The summaries are grouped in three tables--Performance Statistics, CPU Time Distribution, and Process CPU Summary. The Performance Statistics table contains various timestep, iteration, and damping counts, and the cumulative mass balance errors by component. The CPU Time Distribution table shows CPU and elapsed time usage by program sections and subsections. The Process CPU summary lists time usage, memory usage, cell and well distribution among processes. Examples are shown below.

Performance Statistics

MO/DAY/YR: 04/09/1961 1194.74805 DAYS

	THIS RUN	TOTAL RUN
TOTAL TIME STEPS	1408	1408
SUCCESSFUL TIME STEPS	825	825
TIME STEP REPEATS	583	583
Repeats due to Linear solver failed	583	583
TOTAL GLOBAL NEWTON-RAPHSON ITERATIONS	1409	1409
TOTAL GLOBAL LINEAR SOLVER ITERATIONS	57560	57560
TOTAL PRESSURE LINEAR SOLVER ITERATIONS	191566	191566
TOTAL NETWORK NEWTON-RAPHSON ITERATIONS	0	0
TOTAL NON-PHYSICAL NETWORK ITERATIONS	0	0
TOTAL WELL NETWORK SOLVES	2427	2427
TOTAL WELL NETWORK ITERATIONS	1448	1448
NEWTONS SUBJECT TO DAMPING	0	0
DAMPS LIMITED BY GRID(S)	0	0
DAMPS LIMITED BY NETWORK(S)	0	0
NEWTONS RESCALED BY LINE SEARCH	0	0
NEWTONS WITH VOL ERR NOT CONVERGED	0	0
NEWTONS WITH MASS BAL NOT CONVERGED	0	0
NEWTONS WITH NON-PHYSICAL NETWORK SOLUTION	0	0
CUMULATIVE MASS BALANCE ERROR:		
OIL1	0.0E+00	
OIL2	2.0E-15	
GAS	-1.7E-15	
WATER	-6.0E-16	

CPU Time Distribution

	CPU SECONDS -----	% TOTAL RUN -----	ELAPSED SECONDS -----	% TOTAL RUN -----
TIMESTEPS:				
PRE NEWTON	486.150	1.15	488.590	1.15
NEWTONS	41901.070	98.80	41953.630	98.75
WELL UPDATES	0.250	0.00	0.180	0.00

OUTPUTS	17.180	0.04	25.210	0.06
MISC.	2.970	0.01	6.690	0.02
TOTAL TIMESTEPS:	42407.620	99.99	42474.300	99.98
PRE NEWTONS:				
NUMERICAL DKRDS	0.000	0.00	0.000	0.00
REL PERM	431.920	1.02	432.900	1.02
EOS PHASE TRANSITION	0.000	0.00	0.000	0.00
VISCOSITY	0.000	0.00	0.000	0.00
IMPES COEFFICIENTS	45.160	0.11	45.050	0.11
NETWORK/WELL SETUP	0.040	0.00	0.580	0.00
INPUT	0.000	0.00	0.320	0.00
MISC.	9.030	0.02	9.740	0.02
TOTAL PRE NEWTONS:	486.150	1.15	488.590	1.15
NEWTONS:				
COEFFICIENTS	2097.220	4.94	2097.660	4.94
FLASH	727.670	1.72	728.800	1.72
REL. PERM	254.690	0.60	254.270	0.60
PERF EQUATIONS	0.190	0.00	0.210	0.00
MOBILITY DERIVS	318.970	0.75	319.030	0.75
SURFACE NETWORK	3.830	0.01	3.920	0.01
WELL MANAGEMENT	0.000	0.00	0.000	0.00
TARGETS	0.000	0.00	0.000	0.00
LINEAR SOLVER	38399.800	90.54	38407.410	90.40
CONVERGENCE CHECK	1.620	0.00	1.600	0.00
IMPSTAB CALCS	0.000	0.00	0.000	0.00
OUTPUT	39.670	0.09	82.080	0.19
MISC.	57.410	0.14	58.650	0.14
TOTAL NEWTONS:	41901.070	98.80	41953.630	98.75
IMPSTAB CALCS:				
Common Calc.	0.000	0.00	0.000	0.00
Coats/Peaceman Calc	0.000	0.00	0.000	0.00
Misc.	0.000	0.00	0.000	0.00
TOTAL IMPSTAB CALC:	0.000	0.00	0.000	0.00
SURFACE NETWORK:				
INITIALIZATION	0.530	0.00	0.790	0.00
COEFFICIENTS	0.120	0.00	0.130	0.00
CONSTRAINT CHECKING	0.020	0.00	0.020	0.00
Mobilities	0.000	0.00	0.010	0.00
Zero rate	0.050	0.00	0.050	0.00
Rate and pressure	0.040	0.00	0.030	0.00
elimination	0.020	0.00	0.020	0.00
FLASH	0.020	0.00	0.060	0.00
SOLVER	0.000	0.00	0.000	0.00
VARIABLE UPDATE	0.150	0.00	0.190	0.00
HYDRAULICS	0.040	0.00	0.000	0.00
-PDCorr Flash	0.000	0.00	0.000	0.00
-Output	0.000	0.00	0.000	0.00
GRID INTERFACE	0.540	0.00	0.380	0.00
SOLVE WELLS	1.940	0.00	1.910	0.00
-SOLVE WELLS SOLVER	0.560	0.00	0.550	0.00
-WELLS FLASH	0.080	0.00	0.070	0.00
-WELLS HYDRAULICS	0.000	0.00	0.000	0.00
-NET WELLS TRANSFER	0.300	0.00	0.260	0.00

-WELLS ALLOCATE	0.000	0.00	0.000	0.00
MISC.	0.470	0.00	0.440	0.00
	-----	-----	-----	-----
TOTAL NETWORK:	3.830	0.01	3.920	0.01
	-----	-----	-----	-----
LINEAR SOLVER:				
STRUCTURE SETUP	208.670	0.49	209.610	0.49
COMPUTATIONS	38186.240	90.04	38192.760	89.90
MISC.	4.890	0.01	5.040	0.01
	-----	-----	-----	-----
TOTAL SOLVER:	38399.800	90.54	38407.410	90.40
	-----	-----	-----	-----
STRUCTURE SETUP:				
CONST. RED. SYSTEMS	94.240	0.22	94.430	0.22
CONST. LINKED LIST	8.420	0.02	8.430	0.02
APPROX. TEMPLATE 1	95.080	0.22	95.580	0.22
APPROX. TEMPLATE 2	9.950	0.02	10.200	0.02
MISC.	0.980	0.00	0.970	0.00
	-----	-----	-----	-----
TOTAL:	208.670	0.49	209.610	0.49
	-----	-----	-----	-----
SOLVER COMPUTATIONS:				
CONST. RED. SYSTEM	1753.630	4.13	1755.760	4.13
FACILITY SETUP	14.870	0.04	14.880	0.04
FACILITY FACTOR	0.910	0.00	1.030	0.00
REDUCED SYSTEM SOLVE	36335.770	85.68	36340.170	85.54
BACK SOLVE	80.820	0.19	80.740	0.19
MISC.	0.240	0.00	0.180	0.00
	-----	-----	-----	-----
TOTAL:	38186.240	90.04	38192.760	89.90
	-----	-----	-----	-----
REDUCED SYSTEM SOLVE:				
EQUATION SETUP	266.900	0.63	266.820	0.63
LOAD LDU COEFS	39.950	0.09	39.820	0.09
LDU FACTORIZATION	141.160	0.33	141.310	0.33
PRESSURE SOLVE	433.110	1.02	433.480	1.02
PRESSURE PRECOND	3602.100	8.49	3596.540	8.47
IMPLICIT PRECOND	18461.700	43.53	18466.170	43.47
FACILITY SOLVE	251.780	0.59	252.400	0.59
RESIDUAL UPDATES	8151.540	19.22	8154.300	19.19
ORTHOMIN	4960.910	11.70	4962.820	11.68
NORM CALCULATIONS	2.000	0.00	2.230	0.01
MISC.	24.620	0.06	24.280	0.06
	-----	-----	-----	-----
TOTAL:	36335.770	85.68	36340.170	85.54
	-----	-----	-----	-----
TOTALS:				
INITIALIZATION	2.630	0.01	7.010	0.02
INITIAL INPUT	0.830	0.00	2.620	0.01
TOTAL TIMESTEPS	42407.620	99.99	42474.300	99.98
	-----	-----	-----	-----
TOTAL RUN:	42411.080	100.00	42483.930	100.00
	-----	-----	-----	-----

Process CPU Summary

CPU TIME ELAPSED		Active		Active	MilliSeconds/	PEAK MEMORY	CURRENT MEMORY	PROCESS	HOSTNAME
Seconds	Seconds	Cells	Wells	Perfs	Cells*Timesteps	MegaBytes	MegaBytes		
42411.08042483	9301291132			68	0.233281	873.83544	35.29560	1	nexus.lgc.com

Tuning

Orthomin is an algorithm used to accelerate the convergence of iterative methods for solving matrix equations. It works by combining the latest preconditioning with some number of earlier solutions in such a way as to minimize the L2 norm (or square root of the sum of squares) of the residuals. In other words, it finds the best solution in an L2 sense that can be obtained from a combination of these earlier solutions and the current preconditioning.

Orthomin uses two parameters.

- **CYCLE_LENGTH** - This parameter specifies how many prior solutions are to be retained and used in the recombination process described above. At the peak, a total of CYCLE_LENGTH vectors (CYCLE_LENGTH - 1 prior solutions plus the latest preconditioning) are combined.

The solution proceeds as follows. On the first iteration, there is no previous solution, so only the current preconditioning can be used. It is scaled so as to minimize the residual L2 norm, and the result is the first solution. On the second iteration, the second preconditioning and the first solution are combined in a way that minimizes the residual L2 norm. On the third iteration, the third preconditioning and first and second solutions are combined in a way that minimizes the residual L2 norm. This procedure continues through iteration number CYCLE_LENGTH. On iteration CYCLE_LENGTH + 1, the first CYCLE_LENGTH solutions are discarded and the process restarts as if it is again the first iteration, with the difference being that it is starting from the current solution estimate.

- **MAXCYCLES** - This parameter determines the maximum number of iterations allowed, which is MAXCYCLES times CYCLE_LENGTH.

As long as the iteration converges quickly, you do not need to modify these parameters. Sometimes poor convergence can be improved by increasing CYCLE_LENGTH, but this is expensive, since the orthomin computational requirement is proportional to the square of the number of prior solutions being used.

Formulation

The Nexus[®] software is a simulator that solves a fully coupled reservoir and surface network model^{1, 2}. The simulator can be run in black oil mode (and variations of black oil) and in fully compositional mode using an equation of state (EOS). The conventional well model is replaced with a generalized network model of the wells and facilities. The detailed discussion of the surface network model is in [Production and Injection Controls](#).

This chapter describes the overall global system of equations. It also describes the reservoir equations in detail.

1. Coats et al. "A Generalized Wellbore and Surface Facility Model, Fully Coupled to a Reservoir Simulator" SPE 79704 presented at the SPE Reservoir Simulation Symposium (Feb., 2003) Houston, TX.

2. Shiralkar et al. "Development and Field Application of a High Performance, Unstructured Simulator with Parallel Capability" SPE 93080 presented at the SPE Reservoir Simulation Symposium (Feb., 2005) Houston, TX.

Related Documentation

Detailed descriptions of the run control keywords used to modify formulation can be found in the “Solver Parameters” section of the “Run Control” chapter of the *Nexus® Keyword Document*.

Procedures for defining run control parameters in a SimDataStudio™ case can be found in the “Simulation Control” chapter of the *Nexus SimDataStudio User’s Guide*.

The Coupled System of Equations

The global system of equations can be divided into two separate domains -- a network domain and a reservoir domain.

The network domain consists of a collection of nodes linked to each other by connections. Molar flow rates (or mass flow rates) for each component are connection variables, while pressures are node variables.

The network equations consist of the following:

- perforation rate equations
- component mass balance equations at network nodes
- hydraulics equations
- network constraint equations
- composition specification equations

Network constraint equations may include constraints on rates (surface or in-situ conditions) applied to particular connections, or constraints on pressures, applied to particular nodes. Wells are also included within the network domain.

During each Newton iteration of a timestep, the network domain is solved to convergence before the entire system of equations is solved simultaneously. The system of network equations is always solved implicitly. During the solution, the reservoir variables are held constant.

The global system of equations is then assembled and sent to the matrix solver. By default, the timestep convergence is based only on the reservoir conservation and constraint equations. The network equations will be converged within the accuracy of their linearizations, assuming continuity of active constraints. Convergence of the global system is not allowed if any inactive rate or pressure constraint is violated (within a user-specified tolerance) as a result of the final solution.

The reservoir domain may itself be broken up into separate grids. The grids do not necessarily represent a collection of gridblocks that are contiguous in space. Rather, they represent a group of blocks that share a common fluid description (for example, black oil or equation of state) and use a common solution method (for example, implicit or IMPES).

Within each grid, the reservoir equations include molar conservation equations for EOS fluids, or mass conservation equations for black oil fluids. Flow may be between gridblocks within the same grid or gridblocks that belong to connected grids. Additionally, a volume balance equation is specified for each gridblock.

Reservoir Equations

The reservoir equation formulation is a volume balance formulation which is similar to the one proposed by Watts³.

For each component, a molar (or mass) conservation equation can be written for each gridblock of the form

$$R_{M,i} = \sum_{c=1}^{nconn} F_{c,i}^k + Q_i^{n+1} - \frac{(M_i^{n+1} - M_i^n)}{\Delta t}$$

where

$nconn$ is the number of connections to the gridblock

M_i^{n+1} is the moles (or mass) of component i in the gridblock at timestep $n+1$

M_i^n is the moles (or mass) of component i in the gridblock at timestep n

Δt is the timestep size

Q_i is the well source/sink term for component i

$F_{c,i}$ is the molar (or mass) flux rate for component i and connection c .

The $n+1$ superscript on Q indicates that the source/sink terms are always treated implicitly.

The source/sink terms are perforation equations which provide a link between the network and the reservoir.

The k superscript on F indicates that the flux term may be treated implicitly or explicitly.

The equation is written for $nc+1$ components, where the first nc components are hydrocarbons, and the last one is water.

3. Watts, J.W.: "A Compositional Formulation of the Pressure and Saturation Equations", SPERE (May, 1986) 43; Trans., AIME, 281.

For a single porosity system, the discretized form of the molar flux for component i written for connection c is equal to

$$F_{c,i} = \left[T \left(\sum_{j=1}^{nphases} \left(\frac{k_{rj}}{\mu_j} \rho_j x_{j,i} \right)^k \right) (\Delta p_o^{n+1} + \Delta P_{c,oj}^k - (\rho_j M_j g)^n \Delta D) \right]_c$$

$$i = 1, \dots, nc + 1$$

where

T is the transmissibility

$nphases$ is the number of phases

k_{rj} is the relative permeability of phase j

ρ_j is the molar density of phase j

μ_j is the viscosity of phase j

$x_{j,i}$ is the mole fraction of component i in phase j

$\Delta p_{o,c}$ is the oil phase pressure difference for connection c

$\Delta P_{c,oj}$ is the capillary pressure (the offset between the oil phase pressure and the pressure for phase j) difference for connection c

M_j is the molecular weight of phase j

g is the gravitational acceleration

ΔD_c is the depth difference for connection c .

The superscripts indicate the time level at which variables are evaluated.

The variables pre-multiplying the depth difference are always evaluated explicitly, while the pressure difference is always treated implicitly.

Capillary pressure and the phase-specific coefficients multiplying the potential difference can be treated implicitly or explicitly.

A similar equation for F can be written for a mass conservation equation.

All the variables in the above connection exist at connection c . Usually, single point upstream rules are used to determine how variables at a connection are evaluated from variables that exist at gridblock centers.

In addition, a volume balance equation is written for each gridblock. The volume balance is

$$R_V = V_p - \sum_{j=1}^{npnases} V_j$$

In other words, for each gridblock the pore volume (V_p) is constrained to be equal to the fluid volume which is equal to the sum of the individual phase volumes (V_j).

Global Matrix Equations

At each Newton iteration, it is necessary to solve a matrix equation. Depending on the number of cells used to represent the reservoir, the reservoir part of this equation must be solved iteratively. For the simulation to be successful, this iteration must converge in a reasonable number of iterations.

In general, the overall convergence rate will be limited by convergence of a particular part of the overall process. Potentially limiting the convergence rate are the following:

- Convergence problems related to the coupling between the facility network and the reservoir.
- In implicit computations, convergence of the CPR computation.
- Convergence of the reservoir pressure solution.
- In multiple subgrid runs, which typically use multiple processors, convergence problems related to coupling between neighboring subgrids.

In selecting solver parameters, the Nexus® software attempts to resolve issues related to the potential problem areas listed above. The following is intended to help in doing so.

The global matrix equation takes the following form.

$$\begin{bmatrix} A_{nn} & A_{nr} \\ A_{rn} & A_{rr} \end{bmatrix} \begin{bmatrix} x_n \\ x_r \end{bmatrix} = \begin{bmatrix} r_n^0 \\ r_r^0 \end{bmatrix}$$

A_{nn} and A_{nr} contain the network equation coefficients that multiply network and reservoir unknowns, respectively. A_{rn} and A_{rr} contain the reservoir equation coefficients that multiply these same unknowns. x_n and x_r are the network and reservoir unknown vectors, respectively. The network unknowns are pressures and compositions at network nodes and total component mass rates on network connections, including perforations. The reservoir unknowns are cell pressures if IMPES is being used and cell pressures and component masses if the calculations are fully implicit. r_n^0 and r_r^0 are the initial residual vectors.

The Nexus software normally solves the global matrix equation iteratively. Before beginning the iteration, the Nexus software performs a network solve and updates the reservoir residuals. The network is solved exactly, within roundoff error, so the Nexus software assumes that $r_n^0 = 0$. That solution leads to the state of the global matrix equation. The Nexus software begins each iteration with a reservoir solve, which generally is only approximate.

$$A_{rr} \delta x_r^{k+1} = r_r^k$$

This solve can be performed in several ways, which are discussed below. As the reservoir solve proceeds, the Nexus software continually updates the reservoir residuals. At the end of the solve, they are current with the reservoir solution, which means that in effect the following computation has been performed.

$$r_r^{k+1/2} = r_r^k - A_{rr} \delta x_r^{k+1}$$

The Nexus software then computes the network residuals resulting from the change in reservoir solution, recalling that the network residuals were initially zero.

$$r_n^{k+1/2} = -A_{nr} \delta x_r^{k+1}$$

Then the Nexus software solves the network equation.

$$A_{nn} \delta x_n^{k+1} = r_n^{k+1/2}$$

Following this, the network residuals are again zero. The Nexus software updates the reservoir residuals.

$$r_r^{k+1} = r_r^{k+1/2} - A_{rr} \delta x_n^{k+1}$$

Network Coupling and Solution

To perform reservoir simulation, the Nexus® software uses what mathematicians call a multiplicative Schwartz procedure. The Nexus software solves for the reservoir unknowns, updates the network equations based on the new reservoir solution, solves the network equations, updates the reservoir equations based on the new network solution, solves the reservoir equations, etc. The Nexus software continues this process until convergence criteria are achieved.

Generally, this procedure is adequate. Better performance and convergence may sometimes be achieved by incorporating a more complex system of equations for the solution of the network equations. The user may choose to extend the network solution to incorporate connected regions of the reservoir by invoking the SOLVER FACILITIES EXTENDED option in the RUN CONTROL data file or in the SimDataStudio™ software. A number between 1 and 5 may be placed after the “SOLVER FACILITIES EXTENDED” keywords indicating the number of grid-blocks that the facility is extended into the reservoir from the perforations. (The default is 1 if no number is entered.) The use of an extended network solve implies that the reservoir equations and variables in the extended grid-blocks are added to the equations and variables of the network. The result is an overlapping Schwartz procedure, where the overlap region consists of the reservoir cells included in the network set. When the extended network is used, the user must also specify SOLVER IMPLICIT_COUPLING NONE for implicit grids or SOLVER IMPES_COUPLING NONE for IMPES grids. Using the extended network solve means that each network iteration is more computationally expensive, but the convergence of the global solution may be improved.

The meaning of n is as follows: If $n = 1$, the network set includes the normal network plus the cells containing the perforations. If $n = 2$, the network set includes the $n = 1$ set plus the perforated cells' neighboring cells. The $n = 3$ set includes the $n = 2$ set plus the perforated cells' neighbors' neighbors, and so on for $n = 4$ and 5.

In practice, it best not to use n larger than 2 because, as n increases, the cost of the network solve increases and can become quite expensive. It may be possible to reduce this expense by performing the network solve iteratively by specifying FACILITIES ITERATIVE.

Another choice relating to the network-reservoir coupling is the IMPLICIT_COUPLING or IMPES_COUPLING. The default for these

is ROWSUM. Other possibilities are COLSUM and NONE. Sometimes using COLSUM reduces cpu time. As noted above, NONE should be used if the facilities are extended.

Restarts

At any time during a simulation run, information can be written to disk so that a subsequent run can be started by picking up data that was generated by the first run. The information saved this way is called “restart data.” In the Nexus® run process, restart data consists of a series of directories and files contained within a main restart directory - - “casename.rst.” A subsequent run that continues from a previous run is called a “restart run.”

Related Documentation

Detailed descriptions of the restart keywords can be found in the “Case File” chapter of the *Nexus® Keyword Document*.

Procedures for defining restart dates in a SimDataStudio™ case can be found in the “Defining Output Options” chapter of the *Nexus SimDataStudio User’s Guide*.

Restarting a Simulation Run from Restart Data

To launch a restart run, the Nexus® software needs some of the input parameters used by the initial run and the restart data generated by the initial run. Parts of the input data that are of dynamic nature (for example, the data in the RECURRENT_FILES section of the case file that apply to times preceding the restart time) are not necessary. However, that data can be included.

The RESTART keyword in the case file is used to specify the data to restart from. The recommended procedure is to create a copy of the initial run's case file, insert the RESTART keyword, and make other necessary data changes. For a detailed description of the RESTART keyword and related options, see the *Nexus Keyword Document*.

Example of Restart Case File (.fcs):

```
.
.
RESTART   TIME 01/01/2011   initial_test.rst

PROPERTY_FILES
  GLOBALCELL Set 1 nexus_data/uid.dat
  PV Set 1 nexus_data/pv.dat
  CONNECTIONS Set 1 nexus_data/connections.dat
  DEPTH Set 1 nexus_data/depth.dat
  OPTIONS nexus_data/options.dat

ROCK_FILES
  ROCK Method 1 nexus_data/rock.dat
  RELPM Method 1 nexus_data/relpm.dat

PVT_FILES
  PVT Method 1 nexus_data/pvt.dat
  SEPARATOR Method 1 nexus_data/separator.dat
  WATER Method 1 nexus_data/water.dat
INITIALIZATION_FILES
  EQUIL Method 1 nexus_data/init_equil.dat

RECURRENT_FILES
  RUNCONTROL run_control.dat
.
.
```

In the example above, the simulator will read restart data for TIME 1/1/2011 from directory initial_test.rst. The minimum amount of data required in the RUNCONTROL file is TIME data that extends past 1/1/2011.

Additional Considerations

- Any time-dependent data that precedes the restart time and that is included in the input data will have no effect on the restart run. Time-dependent data is stored and read from the restart .rst data. Therefore, changing the CONNECTIONS data will have no effect on a restart run. CONNECTIONS data are time dependent and can be changed with the OVERRIDE file in the RECURRENT_FILES section of the Case File.
- Static data (i.e., data that cannot be changed with time in a simulation run) is not saved as restart data. Therefore, if static data is changed on a restart run, it will have an effect on the run. Although changing static data is possible and allowed, the user needs to exercise caution when choosing to do so.

Redefining Simulation Grids on Restart

The Nexus® restart logic is designed to allow simulation grids to be redefined from the initial run. The user can redefine the IGRID data used for the restart run in any arbitrary manner, including changing the number of simulation grids. Note that changing the number of simulation grids will also require changing the GRIDTOPROC data of the OPTIONS file.

Structure of the Restart Data

The restart data within the .rst directory consists of FORTRAN binary files that are organized in subdirectories by time. An ASCII file with the name 'directory' contains a list of corresponding time step numbers and time in days and date formats for which restart data have been written. The first two columns of data in the directory file are used to identify the subdirectory that contains the data from a given time or date format. The third and fourth columns are matched against the user-supplied data on the RESTART keyword. The 'directory' file can sometimes be useful to determine the format of data to use on the RESTART keyword.

Example contents of the .rst 'directory' file:

TS	Dir Time	Run Time	Run Date
0	0.00000	0.00000	1/01/2001(0: 0: 0)
149	2557.00000	2557.000000000000	1/01/2008(0: 0: 0)
162	2922.00000	2922.000000000000	1/01/2009(0: 0: 0)
175	3288.00000	3288.000000000000	1/01/2010(0: 0: 0)
188	3653.00000	3653.000000000000	1/01/2011(0: 0: 0)

Appendix A: VIP® Functionality not Available in Nexus®

The following is a list of commonly used VIP features which have not yet been implemented in the Nexus software. Less used VIP options are not covered. This list will change over time. If there are questions about the most current functionality, please contact Nexus support.

- Thermal simulation
- Polymer injection
- Todd-Longstaff miscible displacement
- Hydraulic fracture option
- Dual porosity and dual permeability (in limited beta test for R5000.0.2)
- Vertical equilibrium
- Creep option (time-dependent compressibility)
- Two-point upstream weighting
- Integrated saturation initialization (VAITS)
- Refined grid activation and deactivation
- Predictive well management (some functionality may be implemented through the use of Nexus procedures)
- Automatic recompletion units (can be implemented with use of Nexus procedures)
- Pattern balancing
- Gaslift performance curve option (the Nexus software has optimal GLR table method and benefit function optimization method)

- Gas field operations
- Production well optimization in SPN
- Automatic tuning of WI and tubing string
- Radial root grid (radial LGR can be used)
- Diffusion (molecular, pressure, thermal)

Appendix B: Nexus® Functionality Not Available in VIP®

The following is a subset of the functionality available in the Nexus software that either is not available in the VIP software or is difficult to model. This is not an exhaustive list. The Nexus software's ability to model complex networks and well configurations, coupled with its ability to consume array properties for very complex or unstructured grid geometries, enables it to solve a much broader range of subsurface and surface facility configurations than the VIP software.

- Surface facilities model tightly coupled with sub-surface model.
- Smart wells simulation.
- Gaslift network.
- Smart wells can be modeled in detail using wellbore nodes and inline downhole devices such as chokes and valves.
- Full modeling of gaslift distribution networks.
- Wellbore can be modeled in fine detail, essentially modeling a gridded wellbore.
- Option to run the Nexus software with different formulation (IMPES/IMPLICIT) for different defined regions, without using the local grid refinement (LGR) option.
- Pump model available in network.
- Compressor model available in network.
- Benefit function available for gas lift optimization.
- Surface network not limited to tree structure – i.e., loops, crossovers, network branches can be modeled.
- Network PVT treatment can be a superset of the reservoir PVT

- Individual models can be easily coupled into a single full-field model with no modifications.
- Models with different fluid types (gas-water, water-oil, black-oil and API) can be coupled with a common surface network.
- Well management levels are virtually unlimited.
- Pressure maintenance can use any defined region.
- Flexible production and injection target capability.
- Ability to model different hysteresis treatment in different sections of the reservoir.
- Ability to model different compaction treatment in different sections of the reservoir.
- Macro language provides capability to customize well and network treatment.
- Relative permeability scaling input for completions.
- Multiple output region arrays.
- Ability to switch between IMPES and IMPLICIT during runtime.

Appendix C: Common Issues in Converting VIP® Data to Nexus® Data

The following are known VIP®-to-Nexus conversion issues. This section is intended to offer guidance in areas where VIP and Nexus function differently. Please see the “Known Issues” section of the Nexus-VIP Release Notes for additional information on specific items which may not yet be converted correctly by SDS or have known problems in the Nexus implementation.

- Converting RFLOW data may result in significantly different well performance if the drainage radius is not entered. If no drainage radius is entered, the VIP software will use the geometry and permeabilities of the top cell in the well column to calculate the Well Index for all perforations of a given well. The Nexus software will always use the actual cell properties for this calculation.
- In VIP jobs, if no PI or WI data is input or calculated, the VIP software will dynamically calculate the value required for a well to produce or inject at the input maximum rate. The Nexus software will calculate WI values from cell geometry if no WI or PI value is input.
- Saturation table reconstruction – The Nexus software uses the saturation table (relative permeability) data exactly as input. The VIP software will reconstruct input SWT and SGT data into tables with equally spaced saturation entries. By default the number of entries in the reconstructed table is 51. If saturation spacing is closer than 2% in some sections of the table the data integrity can be lost. In this situation, the VIP software should be rerun with DIM entries for NSWDIM and NSGDIM. Dimensions of 1001 (0.1% saturation spacing) is typically sufficient to maintain data integrity.
- The VIP software models using SPN data will not necessarily calculate the same BHP from the reservoir well model as from the SPN model. To enable convergence, the VIP software will switch wells from network control to BHP control after a user specified number of iterations. The tightly coupled surface/sub-surface model in the Nexus software assures convergence between the surface and sub-surface.

- The VIP SPN model uses a default of 60 degrees F in network calculations if no network temperature data is input. Additionally, if temperature-dependent PVT data is not entered, the VIP software will adjust fluid densities based on the ratio of network absolute temperature to reservoir absolute temperature. This capability has been added to the Nexus software (DEN_TEMP and NETTEMP data), but by default the Nexus software will perform no fluid adjustments.
- Using VIP timestep control data can severely restrict the Nexus software performance. It is always best to allow the Nexus software to run with default control data and convergence tolerances prior to adjusting these parameters.
- The Nexus software does not have a VAITS initialization option. INTSAT is usually sufficient.
- The VIP software defaults to well crossflow turned off while the Nexus software defaults to crossflow turned on.
- The Nexus software by default always applies relative permeability endpoints to both cells and perforations. The VIP software has special methodology for injector wells, otherwise the behavior is the same as in Nexus. Nexus has a VIP emulation option that will mimic the VIP methodology. See keyword VIP_INJ_PERF_KR_SCALING in the *Nexus Keyword Document* for more information.
- Beginning a VIP TITLE line with the word 'INCLUDE' will cause an error when converting to Nexus data.
- SDS will not attempt to parse data within a comment block beginning with a comment line including the words 'Begin – External Include File'. This is typically not a problem except for older VIP models built when SDS had much less functionality than currently exists.
- By default, the Nexus software will not perform a linear extrapolation on hydraulic table data. To force linear extrapolation, LIMITS data must be input in each table to specify the maximum values of table input to which the Nexus software will perform linear extrapolation. By default, the Nexus software uses an asymptotic extrapolation algorithm for each table parameter.
- Through the SimDataStudio™ conversion process, ISAT/ISATI and ICMT/IWIRC arrays are combined into Nexus IRELPM and

IROCK arrays respectively. The corresponding table data is combined into Nexus methods as necessary to assure coverage of all required combinations.

Only combinations currently being used by the model being converted are created. Modifying IRELPM and IROCK arrays should be done carefully and the modifications cross-referenced with existing table data to maintain consistency.

- The Nexus software uses the node depth and ddepth (NODECON) to calculate a bottom hole depth. If the bottom hole node depth is different from the depth of the first perf, the bottom hole connection will have a non-zero ddepth, and, by default, use a hydrostatic gradient to compute the pressure difference from the first perf to the bottom hole node. This calculation should be essentially the same as the VIP calculation if the bottom hole node is above the first perf. If the bottom hole node is below the first perf, the Nexus software will not truncate the well connection. The VIP software, on the other hand, will truncate the tubing string at the top of the first perf. In addition, if the WELCON data in the VIP input file includes the parameter WDZ (depth to the first perf) and the tubing string is below the perfs, the VIP software will truncate it at WDZ.

Beginning with version R5000.0.1, an option was implemented in the Nexus software to truncate the well connection at the first perf. See keyword 'ALLOW_BHMOVE' in the “Elevation Profiles” section of the *Nexus Keyword Document* for more information on this option.

- The default PMIN in a Nexus input file is defined at sinks and is 50 psi. If a tubing head does not connect to a sink, it is not assigned a default PMIN. In the VIP implementation, the minimum THP is set to 14.7 psia, by default.
- The capillary pressure adjustment in VIP is limited to the maximum capillary pressure entered in the saturation table. There is no upper limit set for this adjustment in the Nexus software.
- It is strongly recommended that VIP XREG data be manually converted to named IREGION data when using the structured grid workflow. The named IREGION option allows greater flexibility and assists in generating region output that is easier to understand. The Nexus software will internally assign region group names to XREG data which can occasionally be difficult understand.

- When the SimDataStudio software converts VIP models with BOTAB data, it will apply the same adjustments to the data that the VIP software applies internally. These adjustments will result in Nexus PVT data that will be somewhat different from the original VIP BOTAB data plus separator data. The modified input data will allow the Nexus software to more closely model VIP's PVT calculation and generate consistent volumetrics at both reservoir and standard conditions.
- There may be some confusion between Nexus and VIP concerning the component tracking and tracer options. The Nexus tracer option is based on the VIP water and hydrocarbon component tracking. The VIP tracer option, which is actually based on particle tracking, is not available in the Nexus software.
- If any perforation is defined in an inactive grid in the VIP data model, the entire well will be shut in with a zero rate. The Nexus software will ignore these perforations and try to satisfy the input well targets through any active perforations.

Known ECLIPSE™-to-Nexus® Conversion Issues

- Well potentials may be calculated differently. If production or injection well constraints are input in ECLIPSE, they are ignored when ECLIPSE calculates potentials. The Nexus software will honor the well rate constraints in calculating well potentials. Target allocations to individual wells often rely on guide rates calculated from well potentials. This reliance can cause significant differences in the production profile between wells. The ECLIPSE methodology can generally be modeled in Nexus with some simple PROC code.
- In the Extended Network model, ECLIPSE uses the network results to assign THP values to wellheads. Well potentials are then calculated based on well THP limits, ignoring the downstream network connections. The Nexus software honors the network data when calculating well potentials.

Appendix D: Nexus® Desktop Features

Nexus Desktop provides a common environment for the tools that you can use to build Nexus models, submit jobs, and analyze the job results. It provides the following:

- a Session tree, where you can access existing Nexus studies, cases, simulation jobs, and files
- a Viewer panel, where you can view the contents of the various cases in a study
- a Tools menu and toolbar, where you can launch other Nexus applications
- a Job Submitter panel, where you can create and submit jobs. Jobs can be run on the local machine, on a remote machine, and in serial or parallel mode
- a Study and Case management toolbar that allows you to create, manipulate, and remove Nexus studies, cases, and jobs.

Tools Menu

The Tools menu and toolbar provide access to additional software applications, which you can use to create and manipulate Nexus input data and analyze Nexus results.


- **SimDataStudio™ software** - A Nexus application which automates the creation of the initialization and recurrent data files that are submitted to the reservoir simulator.
- **GRIDGENR™ software** - A Nexus application that allows you to map the three-dimensional structure and properties of a hydrocarbon reservoir, then compile the data into a format that can be used to drive reservoir simulation models.
- **Array Calculator** - A Nexus application developed for use with GRIDGENR or with Landmark's DESKTOP-VIP product. Array Calculator calculates arrays of reservoir properties. The array values are assigned to the cells in a reservoir grid.
- **Grid Calculator** - A Nexus application that helps you describe the three-dimensional structure and properties of a hydrocarbon reservoir, then compile the data into a format that can be used to drive reservoir simulation models.
- **Region Calculator** - A VIP post-processing application which lets you generate region reports for areas not initially defined as IREGIONS.
- **E2V** - A Windows application to assist in the conversion of Eclipse data files into an equivalent VIP format. The program works by looking at the entire Eclipse data file and mapping the Eclipse keywords and associated arguments on to the VIP equivalents where possible.
- **SimConvert™** - An application that allows you to filter and move data to and from VIP.
- **DESKTOP-PVT™** - An application designed to simulate the behavior of hydrocarbon fluid mixture and generate PVT properties which can be used in a compositional reservoir simulator to analyze oil and gas production characteristics.
- **SimResults™** - An application that you can use to view, graph, plot, and analyze data from simulation runs.

- **PlotView™** - A 2D view application for viewing plots of simulation results.
- **3DView™** - A 3D view application for viewing 3D volumes of simulation results.
- **Nexus® View** - An integrated Nexus 3D view application that allows you to view simulation models (.vdb's, RESCUE, and ECLIPSE).
- **User Editor** - An application that the user selects as the default text editor to be used by Nexus View for editing text files.

You can also access the applications listed in the **Tools** menu from the Tool Application toolbar.

Session Tree

The Session tree lists the studies that you have loaded to Nexus® Desktop and the cases that reside in those studies. You can load as many studies to the session tree as you want.




Multiple sessions can be created for better case management. To create a new session, select the **Open/Create Session** option from the File menu on the desktop. Type a name into the Session dialog box to create an empty session. All session changes will automatically be saved by the desktop. To open an existing session, select the name from the dropdown box on the Session dialog.

When you create a job for a case, the job and its primary input and output files will be displayed in the Session tree.

As a general rule, actions that you perform using Nexus Desktop tools pertain to the study, case, job, or file currently selected in the Session tree. There are some exceptions. For example, some utility applications may be launched without any selected study or case.


Viewer Panel

When you select an item in the Session tree, the Viewer panel on the right side of Nexus Desktop posts specific information about the item.



Item Identity

For any item selected in the Session tree, four fields appear at the top of the Viewer panel.




These provide information about the case location and identity.

- **Study** - The path to the study for the selected item.
- **Case** - The name of the case for the selected item.
- **Job** - The name of any job for the selected item.
- [...] - The path and description of the selected item.

Nexus View Units

When the Session node is selected in the Session tree, the Viewer panel displays the **Measurement Units to Be Used for Nexus View** selection tool. This tool specifies the units that will be used by Nexus View if it is launched.




The selection tool lists the following options:

- **Field Units. X,Y: Feet Z: Feet**
- **Mixed Units. X,Y: Feet Z: Meters**
- **Mixed Units. X,Y: Meters Z: Feet**
- **Mixed Units. X,Y: Meters Z: Meters**

The selected units will be applied only to the display of subsequently launched instances of Nexus View. The selected units do not modify data in any way. Changing the selected units on this panel will not change the units being used by a currently opened instance of Nexus View.

Case Contents

When you select a case in the Session tree, the Viewer panel displays the types of data currently stored in the study database (VDB) for the case. It also displays the user-defined comment for the case.



Boxes that contain check marks indicate that the corresponding data type has been created and resides in the study database for the selected case. The various data types are:

- **Calculated Data** - Geologic data generated by GridGenr or PowerGrid.
- **Geologic Data** - Geologic data generated by GridGenr or PowerGrid.
- **Initialization Data** - Reservoir properties generated by SimDataStudio, VIP, or a third party product.
- **Recurrent Data** - Simulation production data generated by a VIP or Nexus simulator.
- **Map Data** - 3D map array data for viewing the simulation results in a 3D viewer such as Nexus View, 3DView, or SimResults.
- **Plot Data** - Plot data for viewing in PlotView or SimResults.
- **Observed Data** - Observed real world data for existing wells. Generally loaded in ASCII format.

- **Surface Pipeline Network Data** - Surface pipeline network data generated with Network Planner, a DecisionSpace Well Planner utility.

Note: If the content of a case is modified by a Nexus simulation job, the case information should be refreshed automatically by Nexus Desktop shortly after the job finishes. There is no dynamic update if the content of a case is modified with an external Nexus tool application. In this case, you must click the **Reload Selected Simulation Study** icon in the Study and Case Management toolbar to refresh the displayed case information.

Job Parameters

When you create a job for a case, the job is automatically set up using default values based upon the selected case. If the job was previously created, the previously defined values for that job will be loaded. This job creation requires time to perform disk and possibly network interrogation.

For a Nexus job, the input files referenced in the .fcs file and their include statements must be parsed. Files are checked for existence, network hosts are tested, etc. Some jobs will require more creation time than others. After a job has been created, the job parameters are displayed in the Viewer panel.

Study:	Case:	Job:
C:/GRIDS/apivar.vdb	apivar	Nexus Simulation
Job Location	[localhost] C:/GRIDS	
☐ Copy VDB		
FCS File	[localhost] C:/GRIDS/apivar.fcs	
☐ Case Initializer	[localhost] C:/GRIDS/nexus_data/apivar_ne...	
☐ Parallel Mode		
Queue Name	NONE	
Version	5000_0_1	


The job parameter information includes the following:

- **Job Location** - The network host upon which the job will be run and the directory on that host where the job results will be written.
- **Copy VDB** - Toggle that determines whether or not the .vdb will be copied to a remote location that you have selected as the job location. This option becomes available when you select a remote job location. If the job location is on the local host, this option is disabled. This option will also be disabled if the selected study exists on the remote host. If the option is enabled, it may be selected to copy the VDB for a simulation or left unselected if a VDB is not required for a simulation. Nexus and VIP simulations can be set up so that they don't require a VDB.
- **FCS file** - The Nexus case file to be used as the input for the job. The .fcs file and all of the files that it references must exist. If any file does not exist, a red X will appear beside this option and the job's START button will be disabled. The missing files will appear in red in the Session tree.
- **Case Initializer** - A check button that determines whether the Nexus case will be initialized from an i.map file. If you toggle this check button ON, the information in the i.map file will be used to initialize the job's case in the VDB with map data. You can select an existing i.map file, or allow the job submitter to use an i.map file created by Standalone.
- **Parallel Mode** - A toggle that elects the running of a parallel job on multiple processes instead of a serial job on a single process. Parallel setup using the job submitter is discussed in the [Parallel Panel](#) section.
- **Queue Name** - The name of the queue to be used for parallel jobs. Two listings appear by default -- **NONE** and **Local Queue**. Other listings will be present if you have third party queuing software configured on the job location's host. If you select **NONE**, no queue will be used for the job. If you select a queue name, the job will be submitted to that queue instead of being directly run on the job location host. Nexus Desktop has a simple first-in, first-out single-host job queuing system called "local_queue." However, you may have an installation of SGE, LSF, PBS, or another 3rd party Queue System. You can easily modify the Nexus installation to use these third party queue systems instead of the simplistic default local_queue. More information is provided in the [Queue Chooser](#) section.

- **Version** - The version of the Nexus software that will be used for the job. Versions will be listed for simulators that exist on the selected job location host. By default, only the version for the current installation will be available. Any simulation version that exists in the nexussimulators directory can be used -- past, present, and future. However, the Nexus-VIP 5000.0.2 installer installs only the 5000.0.2 version. You must manually place any other version in the nexussimulators directory of the Nexus-VIP 5000.0.2 installation. More information is provided in the [Version Chooser](#) section.


Parallel Panel

When you toggle the Parallel Mode parameter ON in the Job Parameters panel, the **Open** button becomes enabled. When you click the **Open** button, the Parallel panel displays. You specify the parallel simulation options in this panel. The primary use of the panel is the assignment of parallel nodes to decomposed simulation grids.



General Parallel Options

At the top of the panel is an area containing two components that are used for general parallel setup.




- **Host List** - A field that allows you to select a file in which to store the list of selected nodes upon which the parallel simulation will run. The host list is required to submit a parallel job. By default, the host list file will be named after the selected case and placed in the currently selected job location. If the host list file already exists, it is loaded when the parallel panel appears. When you close the parallel panel, the host list file will be saved to disk. If you constantly use the same hosts over and over (for instance, in a cluster), then selecting a previously saved host list file may be faster than reselecting hosts for a new job.
- **Parallel Debug** - A toggle that enables the running of a parallel job with debug output. The debug output created depends upon the simulator that is used. Please reference the simulator documentation for more information.

Node to Grid Assignment

The center of the parallel panel contains three areas that facilitate the assignment of parallel nodes to decomposed simulation grids. These areas have no labels on the parallel panel but are described below in order from left to right.

Available Hosts

The subpanel on the left contains three tabs that allow the user to select one or more available hosts upon which to run a parallel job.



The three tabs contain different types of available hosts that may be used for a parallel job.

- **Defined Remote Hosts** - A list of the names of all hosts that have been fully defined using the Network File Chooser. This list is automatically populated and can not be edited.

Several approaches may be taken to select defined hosts for a parallel job: 1) Select a defined host and type the number of

processes that should be used on this host. Once you begin typing a number, you have one full second to complete typing. After the second is over, the selected host will appear in the Selected Nodes list in the center of the Parallel panel. 2) Select a defined host and choose the number of processes in the lower-left number selector box. Click on the **Add** button to add this host to the Selected Nodes list in the center of the Parallel panel. 3) Click on the desired Defined host and drag it to the Selected Nodes list in the center of the Parallel panel. Defined hosts can be used alone or with Named Parallel Hosts (next tab).

- **Named Parallel Hosts** - A list of the names of parallel hosts that have been entered manually. This list is useful when hosts are unreachable on a remote cluster or when you simply do not want to fully define all of the hosts. This list is automatically populated when the Parallel panel opens and automatically saved when it closes. Unlike the Defined Remote Hosts, you can edit this list. To edit a host name, double-click the box containing the name. To add a new host name, double-click the empty box at the bottom of the list. To remove one or more host names, select the host names and press the Delete or Backspace key.


Several approaches may be taken to select Named Hosts for a parallel job: 1) Select a Named Parallel host and type the number of processes that should be used on this host. Once you begin typing a number, you have one full second to complete typing. After the second is over, the selected host will appear in the Selected Nodes list in the center of the Parallel panel. 2) Select a Named Host and choose the number of processes in the lower-left number selector box. Click on the **Add** button to add this host to the Selected Nodes list in the center of the Parallel panel. 3) Click on the desired Named Host and drag it to the Selected Nodes list in the center of the Parallel panel. Named parallel hosts can be used alone or with defined remote hosts (previous tab).

- **Queue** - Contains two number selectors that allow you to select the number of queue system hosts and the number of processes upon those hosts to be used for a parallel simulation job submitted to a queue system. The allowable number of selectable queue hosts and processes per host ranges from a minimum of 1 to a maximum of 128. To change the number of hosts or processes, click the up or down arrows next to the displayed value or click on that value and type a new number. Once the value of queue hosts or the value of processes has changed, the Selected Nodes list in the center of the Parallel panel will be repopulated with hosts named "QueueHost####" where #### is numbered from 1 to the

selected value in the Hosts field. Multiple hosts will appear in accordance with the value selected in the Processes field. The generic host names will be replaced by real host name appropriated by a queue system during a run on the queue system. Queue hosts can not be used in conjunction with user-defined remote hosts or with user-named parallel hosts.

Selected Nodes

The subpanel in the center provides a list of selected nodes that will be used during a parallel simulation run.




Each selected node has an index number with a color associated with it. The node indices and colors are used with the parallel simulation grid. If a host list file exists, the Selected Nodes list will be automatically populated with the nodes contained in that file when the Parallel panel is opened. The content of the Selected Nodes list is also based on your selections in the available hosts area on the left side of the Parallel panel. To remove one or more host names from the Selected Nodes list, select the host names and press the Delete or Backspace key.

Upon closing the Parallel panel the host names in the Selected Nodes list will be saved to the selected host list file. The number of hosts in the Selected Nodes list must match the number of hosts assigned to the parallel simulation grid. If a mismatch exists between the number of selected nodes and the number assigned to parallel simulation grids, the parallel run will fail.

Parallel Simulation Grid

The subpanel on the right of the Parallel panel displays the parallel simulation grids and the selected nodes to which they are assigned for processing.



The columns that display simulation grid information differ in number, name, and content between the Nexus and VIP simulators, but three columns are the same for each simulator type.

- **Grid Name** - The name of decomposed simulation grid that should be assigned to a node for processing.
- **Process ID** - The index of the node assigned to the simulation grid for processing.
- **Host Name** - The name of the host assigned to the simulation grid for processing.

To assign a node to one or more parallel simulation grids, select the row containing the grids and type the number of the index corresponding to the desired node in the Selected Node list. Once you begin typing a number, you have one full second to complete typing. The host name

in the Selected Nodes list for that index will appear in the Host Name column for the selected grids.

Note: You may select more than one grid by clicking and dragging the mouse or by clicking on multiple grids while pressing the `Ctrl` or `Shift` key.


There are two buttons below the parallel simulation grid that can be used to facilitate rapid grid to node assignment.

- **Distribute** - Automatically distributes selected nodes to parallel simulation grids that do not currently have assigned selected nodes. The distribution algorithm will attempt to assign nodes to grids as best is possible, but the automatic distribution may not be optimized for the grids and nodes being used.
- **Clear** - Clears all assignments of nodes to grids.

Existing grid-to-node assignments will be loaded and displayed when the Parallel panel is opened. When the Parallel panel is closed the grid-to-node assignments will be saved to disk.

Job Progress

When a job is running, the Viewer panel displays the job progress information in real time. The exact contents of the job progress panel will vary from job to job and from run to run. This information displayed in the job progress panel is written to a file named [casename].progress in the job location directory when the job is complete.



```

=====
Sat Nov  8 08:51:00 CST 2008  apivar Main
=====
Initialization
=====
Running Nexus Initialization
=====
EXE:      C:/Landmark/Nexus-VIP5000.0.1/n
FCS:      C:/GRIDS/apivar.fcs
CASE:     apivar
PWD:      C:/GRIDS
CMD:      C:/Landmark/Nexus-VIP5000.0.1/n
          -c C:/GRIDS/apivar apivar.fcs

Errors      0 Warnings      1
Return code 0

Erasing:   C:/GRIDS/apivar.pid


Initialization Pre-processing Complete

=====
VDB Case Setup
=====
Deleting Case Information
=====
  
```

Text File Contents

The Viewer panel displays the content of a text file that you have selected in the Session tree. The file may be one of the following:

- the Nexus case file (.fcs)
- a .dat file that the case file references
- an output file created by a job



The file viewer displays 1000 lines of a text file at a time. Four navigation buttons allow rapid browsing of large text files. These four buttons will be enabled and disabled based upon the total length of the file and the currently viewed portion of the file.

- **First** - Display the first 1000 lines of text of a file. If the file has fewer than 1000 lines of text then the complete text of the file is displayed.
- **Back** - Display the 1000 lines of file text previous to the first line of the currently displayed text.
- **Next** - Display the 1000 lines of file text after the last line of the currently displayed text.
- **Last** - Display the final 1000 lines of file text or the largest segment possible.

A field to the right of the navigation buttons posts the line range and total number of lines. The field displays with the format: "[first displayed line index] - [last displayed line index] / [total line count]"

To the far right of the navigation buttons is a **Reset** icon that can be used to refresh the file viewer.



This option is useful if the file size has changed since the last attempt to view the file. One instance in which resizing the file may be experienced is when a file is viewed during a job run that writes to the currently displayed file. Since the file grows as the job run progresses, a refresh may be necessary to display the most recent contents of the file.

Icons

Nexus® Desktop uses icons extensively to launch applications, perform actions, and adjust displays.

Tools Icons

The Tools menu and toolbar contain the following icons:

Icon	Description
	Launches the SimDataStudio™ software. You can use the SimDataStudio software to view and edit existing Nexus models, convert ECLIPSE™ and VIP® cases to Nexus cases, and to build entirely new Nexus cases.
	Launches the GRIDGENR™ software. You can use the GRIDGENR software to create grids and populate them with reservoir properties.
	Launches the Array Calculator software. You can use the Array Calculator software to populate grids with arrays.
	Launches the Grid Calculator software. You can use the Grid Calculator software to perform specialized calculations on the properties in a grid.
	Launches the Region Calculator software. You can use the Region Calculator software to generate specialized reports on regions in the project that are not defined as IREGIONS.
	Launches the E2V utility. You can use the E2V software to convert ECLIPSE data to VIP data.
	Launches the SimConvert™ software. You can use the SimConvert software to data to and from a VDB.

	Launches the DESKTOP-PVT™ software. You can use the DESKTOP-PVT software to generate PVT properties or develop a mathematical model which can be used in a compositional reservoir simulator such as VIP-COMP to analyze oil and gas production characteristics.
	Launches the SimResults+3D™ software. You can use the SimResults+3D software to view and analyze the results of a simulation job. The SimResults+3D software combines a 3D viewer with a variety of different types of plots.
	Launches the PLOTVIEW™ software. You can use the PLOTVIEW software to view plots of simulation data.
	Launches the 3DVIEW™ software. You can use the 3DVIEW software to view grid data in 3D.
	Launches the Nexus View software. You can use the Nexus View software to view grid data in 3D and to run the DecisionSpace® Well Planning and the PowerGrid™ software.
	Launches the User Editor tool. This editor is whatever editor application you have configured as the default editor for plain text files or files with specific name extensions (such as .fcs).

Note: When the SimDataStudio software or the user editor is launched using a selected Nexus .fcs file or a VIP primary .dat file, a "file watcher" is started on that file, checking it every 30 seconds or so for modifications. If the file is changed or deleted, Nexus Desktop will attempt to reload that file, modifying the Session tree and Viewer panel as necessary. If Nexus Desktop does not appear to update the information after 30 seconds, you must select a different file in the Job Settings panel and then reselect the edited file to see any modifications. Data files referenced by the .fcs or VIP .dat files will not have watchers placed upon them, so any change in the file hierarchy made within one of those files will require a manual reload of the .fcs or .dat file to update the Session tree and Viewer panel.

Study and Case Management Icons

The study and case management toolbar contains icons that allow the user to create, manipulate, and remove Nexus studies, cases, and jobs. These icons also appear in the **File** menu.

Icon	Description
	Launches the Network File Chooser dialog box, where you can select an existing study or create a new one.
	Reloads the selected study. You can use this tool to refresh the study when you have generated new data that does not display in the main Study panel.
	Closes the selected study. Closing the study will not delete it from the disk.
	Creates a new simulation case in the currently selected study. Note: The case is not permanently written to disk until it is populated with data using a Nexus tool application or by running a simulation job.
	Renames the selected case with a new name that you provide. This tool also allows editing of the comment for the selected case.
	Removes the selected case. <i>Note:</i> This option will remove the case permanently from disk.
	Creates a Nexus job for the selected case. Only one job per case is allowed.
	Creates a VIP job for the selected case. Only one job per case is allowed.
	Creates a STREAMCALC™ job for the selected case. Only one job per case is allowed. To run a StreamCalc job after a Nexus job, remove the Nexus job and create a StreamCalc job.

	Removes the selected job. No job files are removed from disk.
---	---

Job Submitter Panel

The Job Submitter panel provides information about the current job and controls for running it

Job Control Panel

The Job Control panel appears at the top of the Viewer panel when a job or some file belonging to a job is selected in the Session tree.



The panel provides the following controls:

- **Start Button** - Starts the currently selected job. If the **Start** button is grey, then a job option has an invalid value that can not be used to successfully complete a job.

Note: Shortly after the start button has been clicked, the Viewer panel will automatically display the job progress for the job.


- **Stop Button** - Stops a currently running job.

Note: The job may not immediately stop running, as the current process of the job may need to complete and some job cleanup may be required.

- **Status Field** - Displays the current or last status of a job, which may be running, complete, stopped, failed, or empty (no action taken).

Job Settings Panel

The Job Settings panel provides various job options and settings.



- **Job Options** - Used to change the manner in which a job will be run. Placing the mouse arrow over an option name will display a tool tip box containing more information about the option.

Some job options have toggle boxes next to them. These are used to turn the options on or off. If a toggle box contains a checkbox, the corresponding job option will be used for the job when it is run. If a toggle box is not present next to an option, use of the option is mandatory. If the box is grey, the option usage can not be changed.

Note: If a red X appears beside the name of a job option, the current value for the option is invalid and a valid value must be selected to start the job.

- **Value Fields** - Display the current values for the job options. The value field is just a display and the job value can not be changed in the value field.


Note: If an option has an empty value field, the option has only an on/off or use/don't use value such as the **Copy VDB** option. If the value field has a horizontal line such as that in the **Parallel Mode** option, the option has a complex value parameter that can not be displayed with one line of text.

- **Select Buttons** - Used to select values for the options. Clicking a select button normally opens a new window containing possible alternative values for a job option. Selection of a new value will

replace the previously displayed value in the value field with the newly selected value.

Network File Chooser


The Network File Chooser allows you to select a file or directory on the local machine or on a remote host on the network.



To select a file or directory on a remote host, that host must have a current installation of Nexus.

Host Controls

The host controls allow you to select or edit a network host.




The host controls include the following:

- **Host Selector** - Displays the name of the currently selected or selectable host. The host selector may be clicked to show a selectable list of all defined network hosts. If a host name in the host selector is displayed in red text then that host is not available.

Note: The host named "localhost" is your local host, is always available for selection, and can not be deleted.

- **Open** - Opens the host whose name is displayed in the host selector to browse files and directories on that host. When the host name displayed in the host selector is that of the currently open host, the **Open** button does nothing. When the host name is that of a different host, the other host will be opened for browsing.

Clicking the **Open** button of the host controls opens a window in which parameters required to communicate with a remote host may be defined.



Host Name - Used to type the name a remote host that you wish to use for file browsing.

Note: Immediately after you enter a host name, Nexus Desktop will attempt to find that host on the network. During this time the host name may appear in grey text. If the host can not be resolved on the network, the host name will be displayed in red text. If this color change to red occurs, please check your host name. The full host address may be required (for example, [hostname].[networkname].com), or even the IP address (for example, 134.132.01.01).

User Name - Used to type your username for the remote host.

Note: Your user setup on the remote host must have access to a Nexus installation, and your user setup must have the environment variables \$NEXUS_SCRIPT and \$NEXUS_SIM defined and correctly pointing to the nexusscript/ and nexussimulators/ (consecutively) directories of that Nexus installation.

Password - Used to type your password on the remote host for your username.


Note: If the remote host is set up to accept remote connections using your username without a password, this field may be left blank.

The contact parameters for new host can be defined by simply typing a new host name in the hostname field. The previously defined and displayed host will still be saved, and the new host definition will be created.

- **Edit** - Opens a new window to change the contact parameters for the host currently displayed in the host selector or to create a new selectable host. Host contact parameters are those needed to communicate with a remote host. They include host name, user name on that host, and user password on that host.
- **Delete Host** - Removes the name of the host currently displayed in the host selector from the list of selectable hosts.

File Controls


The file controls allow the selection of a file or directory on the currently selected host.



The file controls of the network file chooser are similar to the file chooser controls of any standard application.

Progress Bar


The progress bar displays the current status of the network file chooser.



If the chooser is busy opening a directory, reading a file, or writing a file, the progress bar will display text describing the process currently being performed and a moving graphical bar.

Refresh Button


The **Refresh** button reloads the currently displayed directory.



For speed purposes (especially for remote hosts), directory information is stored in memory when that directory is first accessed. If the directory contents have changed since the last access, you should click the **Refresh** button to reload and display the current contents of the directory.

Timestep Table and Plot


A Nexus or VIP simulation will create a file containing basic status information for the field as it runs. When available, this status file can be displayed in the Viewer panel by selecting "Timestep Table" in the Session tree.



Selecting the timestep table is possible when "Timestep Table" appears in black text. If the timestep table is selected during a simulation run, the status information will be displayed and updated in real time.

Updates of the timestep table for a local run will be much faster than updates for a remote run simply because of the network access required for such updates. Contents of the timestep table are not kept in memory. Therefore, if you select another item in the Session tree and then later reselect the timestep table, the entire table must be reloaded, which may take some time, depending upon the amount of timestep information and the network speed for a remote run.

When you select “Timestep Table” Session tree, the timestep table and plot are displayed in the Viewer panel.



The table and plot incorporate the following components:

- **Color Tools** - Allow you to change the background color for the data names and timestep plot or the color palette for the data names text and the timestep plot lines.
- **Data Name** - The name of the type of data displayed in the underlying timestep table column. Clicking on the data name will add a line displaying that column's data to the timestep plot. Clicking again will remove that column's data from the timestep plot. Double-clicking will cause the data for that column to be used as the X-axis data in the timestep plot. The text of data name is displayed using the same color as the line for that data in the timestep plot. The background color for all data names is the background color used in the timestep plot.
- **Timestep Table** - Displays the data of a simulation for each timestep.
- **Timestep Plot** - Displays timestep data in a graphical format.

Y-Axis labels appear on the left and right of the plot. Axis labels are displayed with text in the same color as the plot line to which


they apply. Moving the mouse arrow over a Y-axis label will cause the corresponding plot line to appear in bold and the plot will display horizontal reference lines for each number on the axis display.

Plot lines are displayed in the same color as their corresponding axis text and to their corresponding timestep table data name. Moving the mouse over a plot line will display a data name for that line. Right-clicking on a plot line allows you to change the color of that line or the type of scale for the plot line's axis (linear or logarithmic).

Queue Chooser

When you create or open a job, a **Queue Name** setting is listed among the job parameters in the Nexus Desktop Viewer panel.

The **Queue Name** setting allows you to submit the job to a queue on a queue system instead of running it immediately in the selected job location. Clicking on the **Open Queue Name** selection button opens the Queue Chooser dialog box, in which you can choose a desired job queue.



- **Queue Selector** - Allows you to select the desired queue to which the job should be submitted. Selecting "NONE" causes the job to run normally instead of being submitted to a queue system.

Note: If there is no proprietary queue system on your network or you have not configured Nexus to use that queue system, you may select the default queue named "Local_Queue." Local_Queue will cause the job to be submitted to a first in, first out sequential queue on the job location host.

- **List Jobs Button** - Allows you to view the names of jobs on the currently selected queue.

Note: Retrieving the job list from the queue system may take some time depending upon the queue system setup and network speed.


- **Job List** - Lists the jobs on the selected queue after the list jobs button is clicked. The label at the top of the job list displays the name of the queue to which the job list applies.

- **Progress Bar** - Shows the progress of the list jobs process.

Version Chooser

When you create or open a job, a **Version** setting is listed among the job parameters in the Nexus Desktop Viewer panel.

The **Version** option allows you to specify a different simulator version than the current Nexus installation version. It opens The Available Selections dialog box in which you can choose a desired simulator version.




Versions listed will be those available on the currently selected job location host in the directory pointed to by the \$NEXUS_SIM environment variable. If that host does not have additional versions of simulators in that directory, no versions other than the current Nexus version will be available.

Note: To make available a version other than the current Nexus version, copy the executable and libraries of the desired simulator version to the directory pointed to by the \$NEXUS_SIM environment variable. If the executable is not present in that directory, the executable version will not be available for selection. If the executable is present but its required libraries are not, the version will be selectable but the job run will most likely fail

Copy Example Dataset

The Nexus Desktop **Help** menu lists an item named “**Examples...**”. Selecting this menu item will open a window that allows you copy an example dataset from the Nexus installation to another location for training purposes.



↑
Example Dataset Tree

↑
Example Dataset Viewer

↑
Copy Dataset Button

- **Example Dataset Tree** - Displays the example datasets available in the Nexus installation. Select the desired dataset for copying in this tree.
- **Example Dataset Viewer** - Displays information about the current selection in the example dataset tree. If the name of an example dataset type is selected, a description of that dataset is displayed. If a dataset text file (.fcs, .dat) is selected, the contents of that text file are displayed.
- **Copy Dataset Button** - Copies the currently selected example dataset from the Nexus installation to another location. Clicking the **Copy Dataset** button opens a network file chooser window, in which a destination directory for the example dataset copy may be selected. When the **Open** button is clicked on the network file


chooser, the example dataset will immediately be copied to the destination location.

Miscellaneous Tasks

Remove Job

To remove an existing job, perform the following steps.

1. Select the job in the Session tree.



2. Click on the **Remove Selected Simulation Job** icon.




The Remove Job message box prompts you to confirm the deletion.

3. Click **OK** in the Remove Job message box.

Edit Files

To edit a file related to a Nexus job, perform the following steps.

1. Select the file to be edited in the Session tree.



2. Select **Tools -> User Editor** from the Nexus Desktop Tools menu.

The user's default editor for text documents will launch with the file loaded.

Note: You can also launch the editor by right clicking on the file and selecting **User Editor**.

3. Make the desired edits and save the file.

Note: When the user editor is launched using a selected Nexus .fcs file or a VIP primary .dat file, a "file watcher" is started on that file, checking it every 30 seconds or so for modifications. If the file is changed or deleted, Nexus Desktop will attempt to reload that file, modifying the Session tree and Viewer panel as necessary. If Nexus Desktop does not appear to update the information after 30 seconds, you must select a different file in the Job Settings panel and then reselect the edited file to see any modifications. Data files referenced by the .fcs or VIP .dat files will not have watchers placed upon them, so any change in the file hierarchy made within one of those files will require a manual reload of the .fcs or .dat file to update the Session tree and Viewer panel.

Close Multiple Studies

To close multiple studies from a session at the same time, perform the following steps.

1. Click on multiple studies in the session tree while pressing the 'Ctrl' or 'Shift' key.
2. Select the **Close Study** option from the File Menu.

Close Multiple Cases

To remove multiple cases from a single study, perform the following steps.

1. Click on multiple cases in the same vdb while pressing the 'Ctrl' or 'Shift' key.
2. Select the **Remove Case** option from the File Menu.

Archive Data

To archive data associated with a parent input file (.fcs or .dat), perform the following steps:

1. Select the **Zip Files** option from the File Menu on the Desktop.
2. Using the File controls, navigate to the appropriate directory and select an .fcs or .dat file for archiving.

All data associated with the parent file will be added to a tar archive and compressed.

The time required for this process is dependent upon the size of the dataset. Large datasets may take several minutes.

Index

C

case file (.fcs) 105
 command line
 running simulation from 83
 convergence
 for runs with multiple grids 197
 improving for implicit timesteps 192
 network convergence 205
 timestep convergence 205
 convergence criteria 190

D

data flows
 export SimDataStudio case data as Nexus Data 44
 import 3D data to SimDataStudio 30
 import data from Eclipse to SimDataStudio 26, 28
 import data from VIP to SimDataStudio 22
 import well data to SimDataStudio 30
 debug 143, 148
 decomposition
 for structured grid data 150
 for unstructured grid data 150

E

environment variables
 for command line job submission 83
 equations
 global matrix 209
 mass conservation 190
 reservoir volume balance formulation 206
 volume constraint 190

F

file extensions
 .alloc 114
 .dbg 114
 .err 112

.fbc 118
 .fbn 118
 .fcs 105, 111
 .gbl 118
 .hyd 114
 .lck 118
 .lcklog 114
 .map 114
 .max 115
 .maxijk 115
 .netsum 115
 .osc 115
 .out 115
 .plt 116
 .rft 116
 .rpt 112
 .rst 116
 .runlog 116
 .srpt 112
 .ss_field 116
 .ss_network 116
 .ss_regions 116
 .ss_targets 116
 .ss_wells 116
 .sum 116
 .tssum 117
 .vdb 112, 117
 .wrn 117
 _alias_names.xref 114
 files
 input 111
 intermediate 118
 map 119
 INIT 119
 example 122
 format 121
 RECUR 119
 format 123
 output 112
 plot 125
 example 125
 format 125

flow control devices 173

I

IMPES

- timestepping efficiency 197
- timestepping performance 155

IMPES stability algorithms 196

IMPLICIT

- timestepping performance 155

injection constraints

- for group connections 175
- for individual connections 172

K

keywords

- ADJUSTTOTIME 186
- ALLOW 193
- AUTO 186
- CALCCOND 177
- CALCCONS 177
- CALCMETHOD 177
- COLSUM 212
- COMMINGLE 174
- CON 186
- CONNECTIONS 216
- CONSTRAINTS 172
- CYCLE_LENGTH 202
- DCMAX 186
- DCRPT 186
- DEBUG 143
- DECOMPOSE 150
- DEPTH 162
- DEPTHTAB 139
- DMOBLIM 192
- DPBHMx 165
- DPLIM 192
- DSGLIM 192
- DSLIM 192
- DT 186
- DVOLLIM 192
- DZLIM 192
- ENDGRID 197
- EXTENDED 211

FLOWSECT 162

GASLIFT 47

GCYCLE 187

GLOBALCELL 160

GRID 197

GRIDTOPROC 151, 217

GROUP 164

HYD 47

HYDMAX 193

IBAT 131

ILU 152

IMPES 186, 190, 196

IMPES_COUPLING 211

IMPLICIT 186, 190

IMPLICIT_COUPLING 211

IMPLICIT_PRECON 152

IMPSTAB 196

ITERATIVE 211

IW 161

JW 161

K 160

KH 160, 161, 166

KRDAMP 192

L 160, 161

LENGTH 164

LGR 161

LINE_SEARCH 192

MASS 190

MAX 193

MAXBADNETS 194

MAXCYCLES 202

MAXINCREASE 186

MAXNEWTONS 190

MD 162

MDCON 162

NAME 131

NEGFLOWLIM 192

NODECON 48

NONE 212

NUMBER 131

OPTIONS 47, 151

OVERREAD 141

OVERRIDE 216

OVERSHOOT 193

PARENT 162

PERFREV ALLOW 193
PPERF 166
PRECON_AMG 152
PRECON_AMG_RS 152
PRIORITY 179
PUMP 47
PVT 47
QMAX 187
QMAXINJ 187
QMAXPROD 187
RADB 160
RADW 166
REGDATA 131
RESERVOIR 47, 48
RESTART 215
ROWSUM 212
RUNCONTROL 47
SECT 162
SEPNODES 174
SGCERFTOL 192
SGCTOL 192
SKIN 159, 166
SOLVER 211
SURFACE 47
TARGET 175, 177
TIME 186
TOLS 190
VALVE 47
VOLCON 190
WCYCLE 187
WELLBORE 163
WELLMOD 167
WELLNODE 163
WELLSPEC 157
WI 166
X 162
Y 162
ZONE 165

M

MPI 147
multireservoir modeling 46
 case file structure 46
 map data 48

network data 48
PVT data 48

O

orthomin (algorithm) 202
overview 46

P

parallel simulation 147
 limitations of 154
procedures 164
processing standard data 108
processing standard input data 109
production constraints
 for group connections 175
 for individual connections 172

R

region group 129
region summaries 127
reports
 timestep size statistics 188
restart data 97
 specifying in run control file 97
 specifying interactively 98
 specifying with a procedure 97

S

separators 174
simulation
 running 62
structured grid data
 parallel job submission 150
 processing 109
surface rate constraints 174

T

targets
 for controlling constraints for group connections 175
timestep 185

timesteps 185

U

unstructured grid data
described 150
parallel job submission 150

W

workflows
create multireservoir model 51
connect well networks 54
merge reservoir models 57
create Nexus model with structured grid 99
create single reservoir model 21
define well data and surface network data
38
upscale a 3D grid 33