

# ECLIPSE

Version 2018.1



INDUSTRY-REFERENCE RESERVOIR SIMULATOR

Installation Guide

**Schlumberger**

## Copyright Notice

**Copyright © 2018 Schlumberger. All rights reserved.**

This work contains the confidential and proprietary trade secrets of Schlumberger and may not be copied or stored in an information retrieval system, transferred, used, distributed, translated or retransmitted in any form or by any means, electronic or mechanical, in whole or in part, without the express written permission of the copyright owner.

## Trademarks & Service Marks

Schlumberger, the Schlumberger logotype, and other words or symbols used to identify the products and services described herein are either trademarks, trade names or service marks of Schlumberger and its licensors, or are the property of their respective owners. These marks may not be copied, imitated or used, in whole or in part, without the express prior written permission of Schlumberger. In addition, covers, page headers, custom graphics, icons, and other design elements may be service marks, trademarks, and/or trade dress of Schlumberger, and may not be copied, imitated, or used, in whole or in part, without the express prior written permission of Schlumberger. Other company, product, and service names are the properties of their respective owners.

ECLIPSE® is a mark of Schlumberger.

An asterisk (\*) is used throughout this document to designate other marks of Schlumberger.

## Security Notice

The software described herein is configured to operate with at least the minimum specifications set out by Schlumberger. You are advised that such minimum specifications are merely recommendations and not intended to be limiting to configurations that may be used to operate the software. Similarly, you are advised that the software should be operated in a secure environment whether such software is operated across a network, on a single system and/or on a plurality of systems. It is up to you to configure and maintain your networks and/or system(s) in a secure manner. If you have further questions as to recommendations regarding recommended specifications or security, please feel free to contact your local Schlumberger representative.

---

# Table of Contents

<b>1</b>	<b>Introduction .....</b>	<b>1</b>
	<b>Simulator computing configurations .....</b>	<b>1</b>
	<b>Supported platforms and software .....</b>	<b>2</b>
<b>2</b>	<b>Install the simulator on a PC .....</b>	<b>5</b>
	<b>Prepare for the installation .....</b>	<b>5</b>
	<b>Install the simulator on a Windows PC .....</b>	<b>6</b>
	<b>Install Intel MPI on a Windows PC .....</b>	<b>7</b>
	<b>Complete the installation .....</b>	<b>7</b>
	<b>Test the installation .....</b>	<b>7</b>
	Test using Simulation Launcher .....	7
	Test using the command line .....	9
<b>3</b>	<b>Install the simulator on Microsoft HPC .....</b>	<b>11</b>
	<b>The HPC installation procedure .....</b>	<b>11</b>
	<b>Check the Infiniband driver firmware .....</b>	<b>12</b>
	<b>Install simulation software on the HPC cluster .....</b>	<b>12</b>
	<b>Configure the HPC cluster .....</b>	<b>13</b>
	Configure cluster information on the head node .....	13
	Configure the HPC web services interface on the head node .....	14
	Configure nodes .....	15
	<i>Install the VCC runtimes on compute nodes .....</i>	<i>16</i>
	<i>Define the file system parameters .....</i>	<i>16</i>
	<i>Set the Windows power configuration .....</i>	<i>16</i>
	<b>Configure PCs to access the cluster .....</b>	<b>16</b>
	<b>Complete the HPC installation .....</b>	<b>17</b>
<b>4</b>	<b>Install the simulator on Linux .....</b>	<b>18</b>
	<b>The Linux installation procedure .....</b>	<b>19</b>
	<b>Configure InfiniBand drivers .....</b>	<b>19</b>
	<b>Configure job schedulers .....</b>	<b>20</b>
	About job schedulers .....	20
	Configure LSF to work with Infiniband .....	20
	Configure LSF to work with Intel MPI .....	21
	Configure LSF HPC .....	21
	Install and configure the integration kits .....	22
	Using the cluster without a scheduler .....	22
	<b>Install the simulation software on a Linux cluster .....</b>	<b>23</b>

Install ECLIPSE in a Linux cluster .....	23
<b>Configure Intel MPI or Platform MPI for Linux .....</b>	<b>25</b>
Check and configure ssh for use with the MPIs .....	25
Intel MPI settings .....	26
Platform MPI settings .....	27
<b>Set up user accounts in the Linux cluster .....</b>	<b>28</b>
<b>Install ECLRUN in the Linux cluster .....</b>	<b>29</b>
Configure the schedulers to work with ECLRUN .....	29
Test the installation .....	29
<b>A Legacy macros and scripts .....</b>	<b>32</b>
Macros: command line flags .....	34
Examples .....	35
<b>B Installation DVD content .....</b>	<b>36</b>

# 1

## Introduction

Using the ECLIPSE or Intersect simulation software, you can create models of your reservoir or field to match historical data and to evaluate different production and recovery scenarios. You can use the simulators with Schlumberger programs such as Petrel or the ECLIPSE pre and post-processor applications (EPP) to model, assemble, visualize and analyze your data.

The simplest simulator configuration is a standalone installation on a Windows computer, but if your organization has a computing cluster, you can use it to run your simulation jobs. This chapter provides an overview of the standalone and cluster installations and lists the supported platforms and software.

### Simulator computing configurations

On a standalone system, you run the simulator software on your workstation, PC or laptop. To do this, you install the software on your computer and set up a link to the license server which controls access to the software. The software is supplied on a DVD and uses an installation wizard similar to those used widely for Windows software installations.



*Figure 1.1. Standalone simulator installation*

In this case, all of the applications and data required for running simulation jobs are on the workstation or computer. The licensing server(s) is usually in the network to provide an access point for all users, but it can be installed on individual Windows computers too.

You can also use network connections to computing clusters to process the simulation jobs. In this case, you might choose to run smaller simulation jobs on your computer and larger, more complex jobs on the cluster. These clusters use either Microsoft Windows or Linux operating systems but share many features. A head node or control node assigns simulation requests/jobs to the available compute nodes. To do this, the cluster requires its own software and data which is kept on local high-speed storage. This includes

scheduling software which controls how and when the simulation requests are processed. The cluster also has the simulation software installed and uses it to process the submitted requests. The cluster and your computer need to share the simulation data so that the cluster can run the simulation and return the results for analysis.

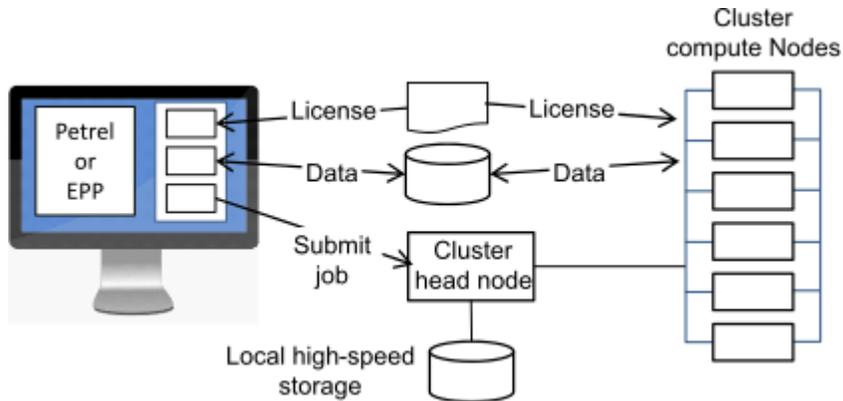


Figure 1.2. Cluster installation

With cluster configurations, there is a great deal of flexibility in where the software components are installed. For example, the license server could be installed on the head node, or be a separate machine in the network. There could be a single installation of the simulation software, with the cluster starting up multiple copies (instances) when needed, or there could be several installations across the cluster. This allows IT professionals to choose a setup that fits with the available cluster equipment.

Installation on a cluster requires a number of steps to configure it, install the software and to set up connections from each Windows computer to the cluster. The installation is normally carried out by IT professionals familiar with network environments and Windows and Linux computing clusters.

## Supported platforms and software

This section describes the supported platforms for a PC/workstation, a Windows HPC cluster, a Linux cluster and for the licensing server.

### Windows PC or workstation

#### Operating systems

Windows 7 Professional 64-bit, Service Pack 1.

Windows 10 64-bit Pro and Enterprise.

No 32-bit systems are supported.

#### Required software

Microsoft .NET 3.5 AND 4+, and Microsoft Installer 3.0

Intel MPI 2018.1. This is used for parallel processing. It is included on the installation DVD and you install it separately from the simulation software.

Carnac 1.2b145 to support the EPP program FloGrid. Installed automatically.

#### Other requirements

Minimum screen resolution of 1024x768

## Windows HPC cluster

For simulator job submission.

### Operating systems

Windows Server 2012 R2 Update 2 x64.

Requires the HPC pack.

The Windows HPC operating systems have their own Message Passing Interface (MPI) and job scheduler.

### Required software

Microsoft C runtimes V9, V10 and V11 which are in the `3rdparty` directory on the DVD. These runtimes are installed automatically on the cluster headnode when you install the simulation software. You must install them manually the other nodes in the Windows HPC cluster.

### Optional software

The License Aware Activation Filter. You can use this to manage license allocation on a Windows HPC cluster. It is supplied as a utility on the installation DVD in the `3rdparty` directory.

## Linux cluster

For simulator job submission. Linux systems must be based on the Intel Xeon x86\_64 chips.

### Operating systems:

Red Hat Enterprise Linux Server 6, Update 6 (x86\_64), or later.

Red Hat Enterprise Linux Server 7, Update 2 (x86\_64), or later.

CentOS Linux Server 7, Update 2 (x86\_64), or later.

### Job schedulers/queueing systems:

LSF 9.1

PBS Pro 12.x

UGE 8.1.3

These are not supplied on the simulator installation DVD and you must purchase and install them separately.

The installation DVD contains an integration kit for each of these schedulers which provides an interface between the license server and the scheduler. Each integration kit enables jobs to be started only when licenses are available and to queue jobs when no licenses are available.

### Message Passing Interfaces (MPI)

Intel MPI 2018.1

Platform MPI 9.1.2

Both of these are installed automatically during Linux installation if you select tools. After installation you select and configure your chosen MPI.

## Licensing platforms

---

**Note:** For the 2015 or later versions of the simulators you must install or upgrade your license server to use Schlumberger Licensing 2015, or later (2017). This is available on the installation DVD for Windows-based licensing and is in the `Tools` directory for Linux licensing.

---

### **Windows licensing platform operating systems**

Windows 7 and Windows 10 64-bit

Windows Server 2012 R2 and 2016

### **Linux licensing platform operating systems**

Red Hat Enterprise Linux Server 6 Update 3 (x86\_64) or later.

### **Licensing software**

2015 license server for Windows and Windows clusters. This must be installed separately from the other software.

2015 license server on Linux. This is installed automatically during Linux installation if tools are selected.

# 2

## Install the simulator on a PC

In a typical standalone installation, you use the simulator software with other applications to model, visualize and characterize data, for example with Petrel which provides full field and reservoir management tools for processing data from seismic survey through to simulation. ECLIPSE is also supplied with a number of legacy pre and post-processor (EPP) applications to help with the construction of simulation datasets. The simulator and associated applications are accessed through the Simulation Launcher. The software components are shown in the following figure.

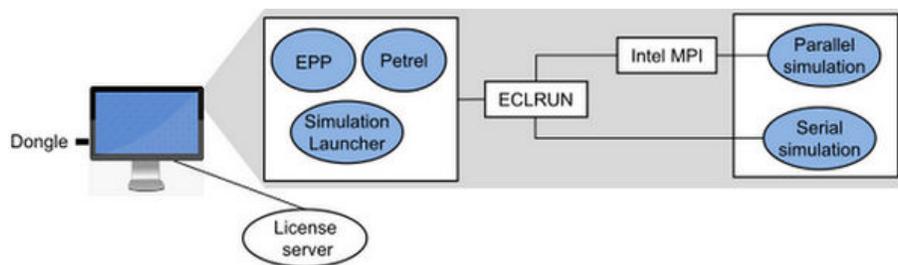


Figure 2.1. Windows installation

The Simulation Launcher and Petrel, if it is installed, send requests to a management and control program called ECLRUN which starts the simulator or EPP application. Intel MPI supports parallel processing and so ECLRUN interacts with it when running simulation jobs which use the multiple processors/parallel processing available on personal computers or workstations.

### Prepare for the installation

Before installing the software, make sure that you are prepared:

- Have administrator privileges on your Windows computer.
- Know the identifiers of your drives, for example C : for the hard drive and D : for the DVD drive.
- Remove previous versions of the simulator, if installed. For example, if you have a Beta version, remove it before starting this installation.
- Temporarily disable your virus checking software while installing the software.
- Read the documentation provided you should have a PDF reader such as Adobe Reader or Foxit.

- Activate .NET framework 3.5 in the `Control Panel` if you have not done so already for another application. Simulation Launcher requires .NET framework 3.5 to function. Windows 10, Windows Server 2012 R2, and Windows Server 2016 all require that the user manually activates .NET framework 3.5. The Microsoft support website contains guidance on activating .NET framework 3.5 for these operating systems.

If you need any advice on these items, consult your IT department.

The installation DVD contains all of the software that you need: the simulator software, associated programs and utilities, the Simulation Launcher, Intel MPI and licensing. You use normal Windows-based installation procedures to install the software and then contact your IT administrator to connect to the license server. Petrel is available separately.

There are several stages to the installation:

1. Install the simulator software and reboot your PC.
2. Install Intel MPI (message passing interface) to allow parallel processing and use with other applications.
3. Complete the installation by setting up the license connection, performing some installation checks, and running some test simulations.

You may need help from your IT department to complete the installation, for example with setting up your connection to the license server.

## Install the simulator on a Windows PC

The installer uses a standard installation wizard. In most cases, clicking **OK**, **Next** or **Continue** moves you through the installation process.

1. Insert the DVD in the DVD drive.

In most cases, this starts the installation automatically. If it doesn't, use Windows Explorer to access the DVD and double click the file `setup.exe`.

This displays the `Install Menu` window.

If you want to access the documentation before proceeding with the installation, click on **Release Notes** or **Installation Guide**.

2. Click **Install Products** and then click the simulator.
3. Click on **Next** or **Yes** in each window to carry on with the installation.
4. Select the location for the installation and then click **Next**.

Either accept the default location of `c:\ecl`, or change to another directory. **Note: directory names cannot contain spaces**, for example `Program Files` is not an acceptable directory.

5. Continue the installation by clicking **Next** or **Yes** in the following windows, accepting the default options and leaving any settings or options unchanged.
6. If prompted  leave this setting and click **Finish** to complete the installation.

A reboot may be required to complete the installation, after which you must install Intel MPI.

## Install Intel MPI on a Windows PC

The MPI installer uses a standard Windows style installation wizard.

1. Insert the DVD in the DVD drive.  
This displays the `Install` Menu, or you can start it by browsing the DVD for the `setup.exe` file.
2. Click **Install Products**, click **Intel MPI**, and then click **Extract** to extract the program files. Do not alter any selections in this window.  
This opens in MPI installation wizard.
3. Click **Next** and step through the installation procedure accepting the license conditions.  
The MPI installation wizard closes when the installation is complete.
4. After the MPI installation is complete, click **Exit** in the `Install` Menu dialog box.

---

**Note:** To complete the installation, you need a connection to the Schlumberger licensing server.

---

## Complete the installation

After installing the simulator and Intel MPI, you need a connection to the license server. Contact your IT department to either:

- Add/update the `SLBSLS_LICENSE_FILE` user variable through the `Control Panel`.
- Install the licensing software and set up your PC as a license server.

During installation, two system variables are added to your PC: `\ecl\HOME` and `\ecl\MACROS` in this order. You can check the system variable settings using the `Control Panel`, or have an IT professional check for you.

## Test the installation

Once you have a license connection, you can test your installation by running a simulation using one of the example datasets in your installation directory. You can run the tests using the `Simulation Launcher` or the command line.

At any stage after installing the simulation software, you can modify or repair it using the installation DVD. In the `Control Panel`, find the software in your installed programs, and select the change option. Follow the prompts to install media and to change or repair selected components.

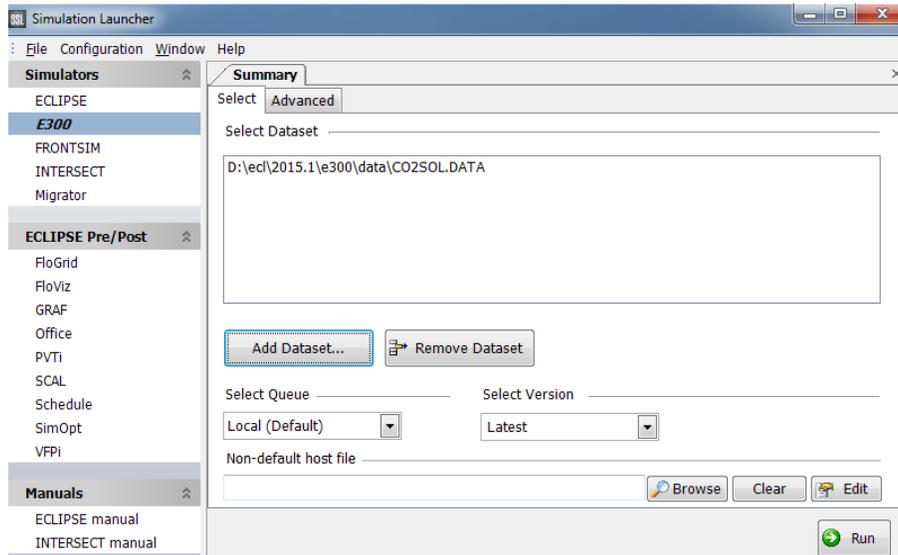
## Test using Simulation Launcher

Once you have a license connection, you can test your installation by running a simulation using one of the example datasets in your installation directory. You can run the tests using the `Simulation Launcher` or the command line.

**Note:** It is recommended that you test the installation with one of the example datasets before running your own models.

## Open the Simulation Launcher

From the Windows Start menu, select **Simulation Launcher**.



## Test the ECLIPSE installation

The ECLIPSE datasets are in your installed directory, for example in `C:\ecl\\eclipse\data` for ECLIPSE 100. To run an ECLIPSE simulation:

1. Select **ECLIPSE**, from the `Simulators` group.
2. Select a dataset from the `eclipse\data` directory and click **Run**.

You can perform similar tests for ECLIPSE 300 and FrontSim if you wish, using datasets in the `e300\data` directory and selecting the E300 simulator or using datasets in `frontsim\data` and selecting the FRONTSIM simulator.

## Test the INTERSECT installation

A sample dataset is provided with the INTERSECT installation in the `C:\ecl\\IX\Examples\Example_2\EX2.DATA` directory. To open the dataset:

1. Select **Migrator** in the `INTERSECT` group.
2. Click **Add Dataset**, select `EX2.DATA` and click **Run**.

This runs the migration. Wait until that has finished before going on to the next step.

3. Select **INTERSECT** in the `INTERSECT` group.
4. Click **Add Dataset**, select the file `EX2.af1` and click **Run**.

The should run and end as shown in the example below:

```

C:\WINDOWS\system32\cmd.exe
REPORT Debug flags utilized:
+-----+-----+
| Flag | Level |
+-----+-----+
| memory_profile_level | 3 |
| nan_debug_level | 10 |
| perf_debug_level | 2 |
| regression_tests | 1 |
| reservoir | 1 |
| use_afi_multi_cmd_syntax | 1 |
| use_afi_restart | 1 |
+-----+-----+
REPORT Message summary
+-----+-----+
| Message Level | Frequency |
+-----+-----+
| INFO | 22 |
| WARNING | 9 |
| ERROR | 0 |
+-----+-----+
INFO Run finished normally.
ix_base checked in
ix_compositional checked in
D:\aim>

```

Alternatively, if you already have an AFI file then you can skip the first two steps and go directly to running INTERSECT.

If you have an ECLIPSE dataset that you want to simulate in INTERSECT, use the Migrator option in the same way as above to convert the ECLIPSE dataset to AFI format.

## Test using the command line

### Create a batch file for testing

As an alternative to running jobs one at a time from the Simulation Launcher you can create a batch file and use it to submit several jobs one after the other. The batch file contains commands for ECLRUN to run the simulation job(s). To do this create a file called RUN .BAT and put in the commands to run each job, for example, for ECLIPSE:

```

cd \data\dir1
eclrun -v 2012.2 eclipse DATASET1
cd \data\dir2
eclrun -v 2013.1 eclipse DATASET2
cd \data\dir3
eclrun -v 2014.1 e300 DATASET3
cd \data\dir4
eclrun -v 2014.2 e300 DATASET4
eclrun eclipse DATASET5

```

To start the runs, double-click on the RUN .BAT in Windows Explorer. The example shows how you can specify particular release versions, or leave out the version number to run the latest version. You can use similar commands but with other simulators, such as `frontsim`, `ix`, and `ecl2ix`.

### Test the Intersect installation

The Migrator can convert an ECLIPSE keyword file into an AFI file, which can then be read by INTERSECT.

1. Use the `cd` command to change to the directory containing the example ECLIPSE dataset.

For example, change to the directory `C:\ecl\<INTERSECT_version>\IX\Examples\Example_2\`.

2. Type `eclrun ecl2ix <basename>` where `<basename>` is the root name of the input dataset, for example `EX2`.

The command generates the AFI file but does not run or simulate it.

3. After the migration is complete, run the simulation by typing `eclrun ix <basename>`.

If the input ECLIPSE dataset is altered, the AFI file has to be regenerated from the dataset first before you can simulate it using INTERSECT.

An AFI file consists of `{basename}.afi` and its associated files `basename_ECL2IX.gsg`, `basename_ECL2IX_IX.ixf`, `basename_ECL2IX_FM.ixf`, `basename_fm_edits.ixf` and `basename_reservoir_edits.ixf`.

## 3

## Install the simulator on Microsoft HPC

This section describes how to set up and configure a Microsoft Windows Server 2012 HPC cluster, to allow simulation jobs to be run on it. This involves some cluster-side and PC-side configuration. An example network is shown in the following figure.

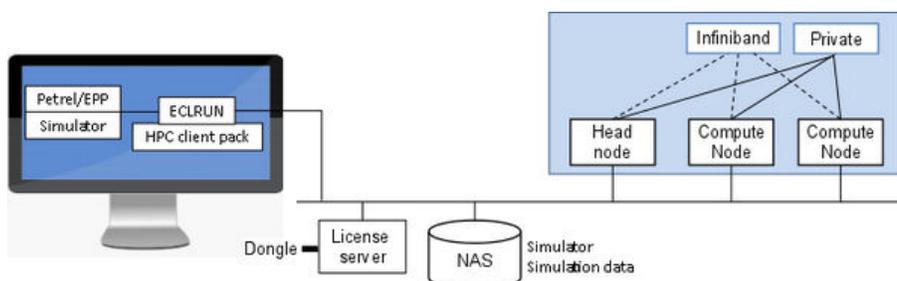


Figure 3.1. Windows HPC installation

The installed software on the personal computer or workstation is the same as for the standalone Windows installation and you can use Petrel or EPP to run the simulation job on your PC. However, ECLRUN is configured with extra connection options to access the Windows HPC network to run simulation jobs. The simulator is typically installed in a file system on Network Attached Storage (NAS) but can be installed on the cluster head node. The simulator must be accessible from the PC and the cluster. To run a simulation job, ECLRUN sends job requests to the HPC Client Pack. The simulation dataset is on a shared disk visible to both the workstation and the Windows HPC cluster. The cluster head node allocates the job to the compute nodes which process the job in parallel. On completion, the simulation results are sent to the shared disk and the HPC Client Pack sends ECLRUN information about the job's status.

On the cluster, the Infiniband network provides high throughput communications links between the cluster nodes. The private network contains data for the operation of the cluster.

The workstation, cluster, license server and shared disks are on a network, typically ethernet, which they use for communication and to share data.

### The HPC installation procedure

Before installing the simulation software you must have met the following prerequisites:

- Installed and configured the HPC cluster according to Microsoft's instructions.
- Set up the user accounts and cluster access.
- Run the cluster diagnostics and corrected any issues.
- Ensure that Microsoft MPI version 5 or later is installed.

The instructions in this guide describe how to:

1. Check the Infiniband interconnect in the cluster.
2. Configure the cluster nodes.
3. Install the simulation software on the cluster head node.

You must install Schlumberger licensing if this is not already in your network.

4. Configure client PCs to access the cluster.

To complete the installation:

- (Optional) Install and set up the license-aware activation filter to manage job submission to the cluster. The software and documentation for this is supplied as a third party PC resource on the installation DVD.
- Set up ECLRUN to submit jobs to the cluster.

You can use ECLRUN's command line options to debug and test the setup during installation and configuration.

## Check the Infiniband driver firmware

The current Infiniband (IB) drivers released from Mellanox require the firmware on the cards to be at the latest level in order to work with HPC. In most cases you can download the firmware from the Mellanox website <http://www.mellanox.com> but you may need to determine what type of card you have. To do this:

1. Download the Mellanox Firmware Tools (MFT) from [http://www.mellanox.com/content/pages.php?pg=management\\_tools=34](http://www.mellanox.com/content/pages.php?pg=management_tools=34)
2. Follow the instructions to get your card type.
3. Then download the correct firmware for your card.

---

**Note:** It is important that you download the correct firmware or you can risk damaging your card. The firmware needs to be installed before the IB drivers. The current generation of IB drivers (2.2.1) are available from <http://www.openfabrics.org> for the generic drivers. There may be specialist drivers for certain modern IB cards that you obtain directly from the Host Channel Adapter (HCA) manufacturers. You need to check on this with your card manufacturer.

---

Once the drivers are installed check that you can see the IB HCAs in the device manager. If they are not in the device manager, the HPC installation will not be able to use them.

## Install simulation software on the HPC cluster

Be aware of the following when installing the simulation software on the cluster from the installation DVD:

- Only install simulation software. If you are installing ECLIPSE, clear the EPP components option, when prompted for the components to install.
- Do not install Intel MPI as Windows HPC has its own MPI.

## Set up the license server

---

**Note:** For the 2015 or later versions of the simulators you must install or upgrade your license server to use Schlumberger Licensing 2015, or later (2017). This is available on the installation DVD for Windows-based licensing and is in the `Tools` directory for Linux licensing.

---

For a new installation, you must set up and configure a license server to provide your users with access to the software (if you are updating an existing installation the license server will already be configured). Typically, you use Network Attached Storage (NAS) for the license server and choose a location that is accessible by the cluster and your users' PCs.

The licensing software is on the installation DVD (insert the DVD, select **Install Products**, then **Schlumberger Licensing** and follow the installation instructions).

After installing the software, open the "*Schlumberger Licensing User Guide*" which is located in the directory where you installed the licensing software (for example `C:\Program Files (x86)\Schlumberger\Schlumberger Licensing\<version>`). Use the guide to configure the licensing server and to provide your users with access to it.

## Configure the HPC cluster

HPC requires specific settings in the cluster and on PCs connecting to the cluster to run the simulation software. This section describes these requirements.

### Configure cluster information on the head node

You must set some environment variables in order to use the simulation software with the cluster. Do this with administrator access rights on the cluster head node.

---

**Note:** The use of Universal Naming Convention (UNC) paths is vital to the current implementation of HPC as the simulator is installed on a file system that must be visible to the head node and the compute nodes. All of the required DLLs are installed not only in the `lib` directory as previously but also in the `bin` directory. This means you do not have to install the simulator software on the compute nodes.

---

For Windows Server 2012, set the following:

```
cluscfg setenvs SLBSLS_LICENSE_FILE=7321@licenseserver
cluscfg setenvs ECLPATH=\\139.10.10.1\ec1
cluscfg setenvs MPIEXEC_AFFINITY_AUTO=1
cluscfg setenvs MPIEXEC_AFFINITY=spr:p:L
cluscfg setparams AffinityType=NoJobs
cluscfg setenvs MSMPI_CONNECTIVITY_TABLE=1
cluscfg setenvs MSMPI_ND_EAGER_LIMIT=2000000000
cluscfg setenvs MSMPI_SHM_EAGER_LIMIT=2000000000
```

The first line sets the location of the license server for the simulator and the optional License Aware Activation Filter. The second line identifies the location of the simulator installation which must be shared

and accessible by all users and machines that submit simulations to the cluster. The remaining configuration parameters provide information for the Message Passing Interface (MPI).

You may have to set other environment variables depending on how you provide access to your cluster:

- If your users are going to submit Multiple Realization (MR) jobs to the cluster, you need to set the MR scheduling environment variable:

```
cluscfg setenvs ECL_MR_SCHEDULING=true
```

You must also add the following to the `eclrun.config` file on each client machine:

```
<LSFLicenses><True></LSFLicenses>
```

- The `DefaultHoldDuration` parameter specifies the time between license checks. By default, if no licenses are available, the scheduler waits 900 seconds before checking again. A setting of 90 is recommended:

```
cluscfg setparams DefaultHoldDuration=90
```

- To check the environment variables settings, use

```
cluscfg listenvs
```

This sets the variable for all users and all compute nodes. If these are not set, ECLRUN will return an error. These settings require HPC Service Pack 3 to work.

- If you want to alter any of the system parameters you can find those that are available using the command:

```
mpiexec -help, mpiexec -help2 or mpiexec -help3
```

Be careful when changing settings as they can cause unforeseen problems.

- Set the license server by IP address rather than name, although both methods should work.
- Wherever simulation data is stored it must be visible to the compute nodes for both read and write operations. The ECLRUN program checks this before it submits a job.

For HPC the `eclrun` command takes the form:

```
eclrun -s localhost -q <cluster name> -u <user on cluster> eclipse  
<datafile>
```

You can also use the debug option when you are testing the connection:

```
eclrun -s localhost -q <cluster name> -u <user on cluster> --  
debug=both eclipse <datafile>
```

---

**Note:** The server for submitting to HPC is always localhost. This option can only be used on systems where the HPC client pack has been installed.

---

## Configure the HPC web services interface on the head node

This is an optional procedure to allow your users to use the Web services or REST interface to monitor their jobs through a remote interface, as well as monitoring the cluster load through the basic Job Manager interface. You can also set up a job submission web page for the simulation software similar to that provided by EnginFrame. To install the web interface:

1. Install Web services components from

Program files\Microsoft HPC Pack 2012\Bin\HpcWebComponents

2. Run a Powershell Window as administrator.

---

**Note:** Powershell scripts and arguments are case sensitive.

---

3. Set unrestricted policy executions so that you can run scripts:
  - a. To check the current execution policy settings, run the command  

```
Get-ExecutionPolicy -List
```
  - b. If the execution policy is restricted, run the following command to set it to unrestricted:  

```
Set-ExecutionPolicy -ExecutionPolicy unrestricted
```
4. Change to the \bin directory:  

```
cd program files\Microsoft HPC Pack 2012\bin
```
5. Enable the service portal.  
If you want to use the default port number of 443, type:  

```
.\Set-HpcWebComponents.ps1 -Service Portal -enable
```

  
If you want to use a different port, type:  

```
.\Set-HpcWebComponents.ps1 -Service Portal -enable -Port  
<web_port_number>
```
6. When prompted for the certificate, enter 0 (zero) for a self-signed certificate.
7. Then enable the REST interface making sure that you use the same port number as for the *Enable the service portal* step:  

```
.\Set-HpcWebComponents.ps1 -Service REST -enable
```

  
or  

```
.\Set-HpcWebComponents.ps1 -Service REST -enable -Port  
<web_port_number>
```
8. To restart the scheduler, type:  

```
Net stop hpcscheduler then  
Net start hpcscheduler
```
9. To test access to the web interface, go to another machine and type  

```
https://<machine_name>/HpcPortal
```

  
If you cannot access the web service, repeat the setup procedure.
10. After setting up web access you should restrict execution policy again by running the command:  

```
Set-ExecutionPolicy -ExecutionPolicy restricted
```

## Configure nodes

To set up the nodes to run the simulation software, you need to:

1. Install the VCC runtimes.
2. Define the file system parameters (although this is not always necessary).
3. Set up the power configuration.

## Install the VCC runtimes on compute nodes

You need to make sure that all the required VCC runtimes are installed on all compute nodes. These can be found on the installation DVD in the following directories:

```
3rdparty\PC\resource\vcredist_vs9  
3rdparty\PC\resource\vcredist_vs10  
3rdparty\PC\resource\vcredist_vs11
```

## Define the file system parameters

Windows Server 2012 (and Windows 8) use the new SMB3 file system. However, the SMB2 file system is still widely used. If you are using this file system, and want to continue using it, run PowerShell as an administrator and run this command (a single command line):

```
Set-ItemProperty -Path "HKLM:\SYSTEM\CurrentControlSet\Services  
\LanmanWorkstation\Parameters" RequireSecureNegotiate -Value 0 -Force
```

You must run this command on all compute nodes and the head node.

## Set the Windows power configuration

On Windows HPC you must set the power configuration to high performance. An incorrect setting can have a 15% performance impact on simulation runs.

To do this, you must have administrator rights on the head node and all of the cluster nodes.

To see what configuration you have, type `powercfg -l` in a DOS prompt on the head node. It should return:

```
Existing Power Schemes (* Active)
```

```
-----
```

```
Power Scheme GUID: 381b4222-f694-41f0-9685-ff5bb260df2e (Balanced)
```

```
Power Scheme GUID: 8c5e7fda-e8bf-4a96-9a85-a6e23a8c635c (High performance) *
```

```
Power Scheme GUID: a1841308-3541-4fab-bc81-f71556f20b4a (Power saver)
```

If you do not have high performance set, you will not see the entry marked with an asterisk (\*). To add this setting, use the `clusrun` command in node management to run the command on all cluster nodes:

```
clusrun /all powercfg -s 8c5e7fda-e8bf-4a96-9a85-a6e23a8c635c
```

## Configure PCs to access the cluster

After installing the HPC client pack on each PC/workstation, you must add an environment variable so that Petrel and the Simulation Launcher can work. On each client PC, add:

```
CCP_SCHEDULER=<name of head node>
```

If you do not set the CCP\_SCHEDULER environment variable, you see the error:

```
D:\>cluscfg listensvs
No connection could be made because the target machine actively refused it
127.0.0.1:5800
```

If you do not set the scheduler, ECLRUN and remote job submission does not work.

## Complete the HPC installation

After installing the simulation software, you can optionally install and set up the license-aware activation filter to manage job submission to the cluster. The filter software and documentation are in a zipped file on the installation DVD in the `3rdparty/PC/resource/HPCUtilities` directory. Use the documentation to set up and configure the filter.

You must set up ECLRUN on all of the PCs that can submit simulation jobs to the HPC cluster. The *"ECLRUN User Guide"* is part of the installed documentation set. To get the guide, start `Simulation Launcher`, select one of the simulator manual sets and open the PDF documentation. The guide is available from the documentation bookshelf. Use the configuration information to provide your users with access to the cluster.

## 4

## Install the simulator on Linux

This section describes how to set up and configure a Linux cluster to run simulation jobs. This involves some cluster-side and PC-side configuration. An example network is shown in the following figure.

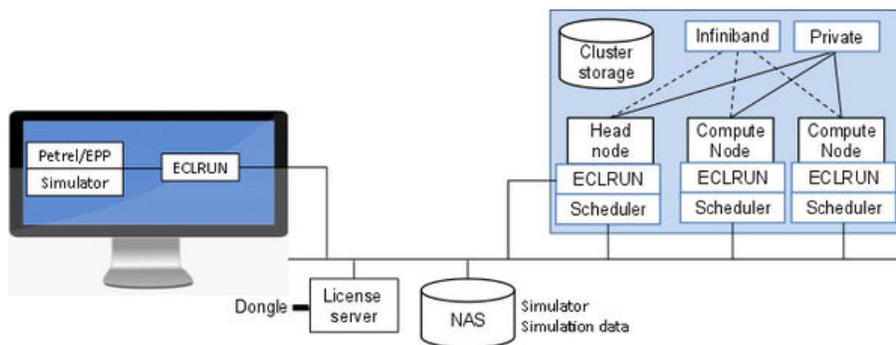


Figure 4.1. Linux installation

The installed software on the personal computer or workstation is the same as for the standalone Windows installation and you can use Petrel or EPP to run the simulation job. However, ECLRUN is configured to access the Linux network to run simulation jobs. ECLRUN is installed on both the workstation and in the Linux cluster. It works with various scheduling systems to manage job submission and completions. LSF, PBS and UGE. ECLRUN and a scheduler are installed on each cluster node. The cluster head node allocates the job to the compute nodes which process the simulation the job in parallel. On completion, the simulation results are sent to the shared disk and information about the job status is sent to the workstation. The simulator is typically installed in a file system on Network Attached Storage (NAS) but can be installed on the cluster head node for example.

On the cluster, the Infiniband network provides high throughput communications links between the cluster nodes. The private network contains cluster data.

The workstation, cluster, license server and shared disks use the ethernet network to communicate and share data.

For the shared disk configuration to work, the shared disk must be accessible by both the Linux and Windows file systems. That is, the disk must support both the Windows and Linux file naming conventions. The user id and permission mapping must work on both Windows and Linux. Where this is not possible, there are separate data disks with the PC/workstation transferring data to a Windows disk and the Linux cluster using a Linux disk.

## The Linux installation procedure

Before you install the simulator software you must already have:

- Installed and configured the Linux cluster according to the manufacturers' instructions.
- Set up storage, user accounts, and cluster access.
- Run the cluster diagnostics and corrected any issues.

To set up the cluster to support the simulators follow the steps below:

1. Configure the Infiniband interconnect in the cluster.
2. Configure the job scheduler and install the job scheduler integration kits.
3. Install the simulator.

You must install Schlumberger licensing if this is not already in your network.

4. Configure MPI settings.
5. Set up ECLRUN to allow your users to submit jobs to the cluster.

You can use ECLRUN's command line options to debug the setup.

## Configure InfiniBand drivers

You must install the OFED drivers if you want to use InfiniBand. See <http://www.openfabrics.org> for more information.

---

**Note:** You can also install the drivers from the Redhat installation disk. Depending on the type of hardware in your system, you may need to download and install a later driver from the OpenFabrics web site.

---

## General InfiniBand notes

To use InfiniBand (IB) cards, you must assign an IP address to the card. This allows you to "ping" the other computers via the IB interface. If the ping succeeds then you know all the hardware is physically connected properly.

## InfiniBand drivers

Follow the installation instructions from the InfiniBand supplier. You may find that IP over IB is working properly but the simulator is not. If you are using Red Hat 5 or later, you can correct this by editing the file `/etc/security/limits.conf` and adding the following two lines:

```
* soft memlock 10240000000
* hard memlock 10240000000
```

The value represents the number of kilobytes that may be locked by a process. The file above contains further documentation.

**Note:** The steps described above will allow any user in the system to lock as much as the whole memory set in the configuration files.

---

## Infinipath drivers

To obtain drivers and instructions, visit this site: <http://www.qlogic.com>.

## Configure job schedulers

This section describes how to set up and configure the job schedulers on the Linux cluster. You can use LSF, PBS or UGE. This section:

- Gives an overview of schedulers.
- Describes how to configure LSF to work with Infiniband and Intel MPI.
- Recommends what to do if you have LSF HPC set up.
- Describes how to install the integration kits.
- Describes how to use a cluster without a scheduler.

## About job schedulers

It is common practice to run simulation jobs using a commercial job scheduler (also called a queueing system). If a single user is using a machine, it may be adequate. However, the main advantage of using parallel processing is speed. This is lost if a machine is overloaded, as all parallel tasks slow down. This can often create run times longer than for serial simulation runs. Job schedulers allow greater throughput of jobs by using the available resources more efficiently.

Using the `eclrun` macros with a job scheduler allows you to make sure that the appropriate control parameters are used when your users submit simulation jobs to the Linux cluster for processing.

## Configure LSF to work with Infiniband

If you are using LSF, you need to edit the `lsf` startup file `/etc/init.d/lsf` to add an entry for `ulimit` so that both Intel MPI and Platform MPI can run when under LSF control. To do this, open the file and add an entry for `ulimit`, for example as shown here in the in the two lines starting `#line` added so simulator:

```
#!/bin/sh
# $Id: startup.svr4,v 1.10 2008/04/08 06:13:09 xltang Exp $
#
# Start and stop LSF daemons, System V / OSF version
# Make sure we're running a shell that understands functions
#
# The following is for the Linux chkconfig utility
# chkconfig: 35 99 01
# description: Load Sharing Facility
#
# The following is for the Linux insserv utility
### BEGIN INIT INFO
# Provides: lsf
# Required-Start: $remote_fs
# Required-Stop: $remote_fs
# Default-Start: 3 5
```

```
... # Default-Stop: 0 1 2 6
... # Description: Start LSF daemons
... ### END INIT INFO
...
... #line added so simulator can run over Infiniband when under LSF control.
... ulimit -l 102400000000
...
... if test "$SH5" != "/bin/sh5" -a -r /bin/sh5
... then
...     SH5=/bin/sh5
...     export SH5
...     exec /bin/sh5 "$0" "$@"
... fi
...
... check_env () {
... if [ x$LSF_ENVDIR = x ]; then
...     # Using default path of lsf.conf...
...     LSF_CONF=/lsftop/lsf/conf/lsf.conf
... fi
... }
```

---

**Note:** LSF by default saves temporary files in a hidden directory called `.lsbatch` which is inside the user's home directory. This can cause problems if the home directories don't have much free space, or quotas are enabled. This can be avoided by adding the following setting in `lsf.conf`:

```
LSB_STDOUT_DIRECT=Y
```

You must restart LSF for this change to take effect.

---

## Configure LSF to work with Intel MPI

The introduction of the HYDRA loader from OpenMPI, caused some changes to the way jobs launch, and this can cause issues.

To go back to the old method of using `ssh` you need to set the environment variable `I_MPI_HYDRA_BOOTSTRAP` to `ssh`. This can be done in one of two ways:

under `csh/tcsh` using `setenv I_MPI_HYDRA_BOOTSTRAP ssh` in your `.cshrc` file, or

under `bash` using `export I_MPI_HYDRA_BOOTSTRAP=ssh` in the `.profile` file

---

## Configure LSF HPC

**Note:** We highly recommend that you do not use any LSF extensions for LSF HPC. These can cause problems if the versions of LSF HPC and Intel MPI are not compatible.

---

### Recommended disabling of LSF HPC

The use of LSF HPC extensions is controlled by an environment variable `ECL_LSFHPC`. If the variable exists, the extensions are invoked. It is recommended that you removed the `ECL_LSFHPC` variable from any settings. This causes simulation jobs to launch using the standard and more flexible method.

### Configuration changes if you require LSF HPC

If you require the extensions to support other cluster operations, you must edit the MPI location in the wrapper script, for example wrapping the Intel MPI in the `intelmpi_wrapper` script. The following example assumes that the simulator has been installed to `/ec1` and `lsf` has been installed to `/lsf`. It also assumes that you are using Intel MPI. If you have installed elsewhere, please use the appropriate path.

1. In the LSF installation directory (/lsf in this example), edit the file  
`/lsf/9.0/linux2.6-glibc2.3-x86_64/bin/intelmpi_wrapper`
2. Search for the line `MPI_TOPDIR=" . . . . . "`
3. Replace the line with the correct location of the Intel MPI. If the default settings have been used this line should look like this  
`MPI_TOPDIR="/ecl/tools/linux_x86_64/intel/mpi/5.0/"`
4. Find all occurrences of "`$MPI_TOPDIR/bin`" and replace them with "`$MPI_TOPDIR/bin64`"
5. If you wish to use SSH to start the MPI daemons:
  - a. Search for the line  
`MPDBOOT_CMD="$MPI_TOPDIR/bin64/mpdboot "`
  - b. Change it to  
`MPDBOOT_CMD="$MPI_TOPDIR/bin64/mpdboot -r /usr/bin/ssh"`

## Install and configure the integration kits

The installation DVD contains an integration kit for each of the schedulers for: PBS Pro, LSF, and UGE. Each integration kit enables jobs to be started only when licenses are available and to queue jobs when no licenses are available. The kits are in the `3rdparty/PBSPro/resouce`, `3rdparty/LSF/resouce` and `3rdparty/UGE/resouce` directories on the installation DVD. Follow the instructions to set up your chosen kit.

## Using the cluster without a scheduler

If you decide not to use a job scheduler/queueing system, you can use a host file to determine how requests are made to computing nodes to run simulation jobs. If you do not specify a name for the host file, a host file containing the name of the current host is auto-generated. This file can then be updated using the relevant formats:

- The standard hostfile format is of the form  
`Node01`  
`Node02`  
`Node03`  
`Node04`  
This runs a process on each node in turn.
- In the case of nodes with twin CPUs, to specify that a process is to be run on each CPU, use a hostfile of the form  
`Node01`  
`Node01`  
`Node02`  
`Node02`

## Install the simulation software on a Linux cluster

The software can be installed independently of previous major or incremental releases. It can be installed over existing installations. For the installation please ensure the following prerequisites are met:

- Only use static mounts for any directories holding home directories, simulator data, or simulator executables. Problems have been noted with auto-mounted file systems.
- Reservoir simulation software should always be installed in version order. Do not install the current version followed by an older version.
- If you are installing both ECLIPSE and Intersect, install ECLIPSE first. For example, install ECLIPSE 20XY.Z and then Intersect 20XY.Z.
- If you do install an earlier version of the software, do not install the macros. To do this, clear the macros option during installation.

---

**Note:** It is assumed in this chapter that the software is installed in the default location `/ecl`.

---

After installing the simulator software, you must install the licensing software if you do not already have a licensing server in your network. You then configure the MPI settings to support the simulators.

## Install ECLIPSE in a Linux cluster

1. Insert the installation DVD into the disk drive.

The DVD should mount automatically and you can continue with the installation.

2. Run the install C-shell script which is supplied on the disk:

```
/media/cdrom/ECLIPSE/UNIX/install/cdinst.csh
```

If this is successful, you are prompted to continue the installation.

```
The following versions were found for installation
0) [ ] Exit without installing
1) [ ] documentation_and_data 2015.1
2) [x] linux_x86_64 2015.1
Choose a version to install (1 - 2) [default : Linux x86_64] :
```

If this is not successful, it means that the DVD has not mounted automatically and you have to mount it manually. To do this:

- a. Unmount the installation disk using the command: `umount -k /dev/cdrom`
- b. Mount the installation disk use the command `mount /dev/cdrom /tmp/a`.

Please note that `/tmp/a` should be an existing directory.

- c. To start the installation, type `/media/cdrom/ECLIPSE/UNIX/install/cdinst.csh`

You are prompted to continue the installation.

3. Press enter to select the default Linux installation.

This displays a list of options and prompts you for the installation directory.

```
Options for linux_x86_64
1) E300
2) ECLIPSE
3) Expand
4) Extract
5) FrontSim
6) ParallelE300
7) ParallelECLIPSE
8) Rename
9) Tools

Select A - Install all of the above
      N - Install none of these
      S - Select a subset

Please enter choice :
Enter proposed location for installation (default /ecl) :
```

4. For the choice, type a and for the installation directory, type /ecl. Then press enter.

**Note:** Always install ECLIPSE to the same location as any previous versions. There is an internal directory structure to prevent any old versions being overwritten.

ECLIPSE for different architectures should also be installed into the same location

5. Select the install directory.

The software then installs:

```
Enter proposed location for installation (default /ecl) : /ecl
Installing E300 (size 15877259 bytes)
Installing Eclipse (size 16229034 bytes)
Installing Expand (size 356035 bytes)
Installing Extract (size 572208 bytes)
Installing FrontSim (size 25773064 bytes)
Installing ParallelE300 (size 33849984 bytes)
Installing ParallelEclipse (size 30258247 bytes)
Installing Rename (size 336972 bytes)
Installing Tools (size 155826405 bytes)
Installing Resource
(size 20530476)
(size 67967635)
Resource installed (total size 88498111)
Do you want to install the macros [default n]?
```

6. At the prompt Do you want to install the macros [default n]? type y and press the return key.

Always answer y (yes) to the prompt unless you are installing an earlier version, in which case, do not install the macros. The current macros directory, if any, will be backed up with a suffix consisting of the date and time of install, for example macros.backup.13:24:52.230209.

7. If you want to install the ECLIPSE documentation and data, rerun the cdinst.csh script and select the option when prompted.

## Setup the license server

---

**Note:** For the 2015 or later versions of the simulators you must install or upgrade your license server to use Schlumberger Licensing 2015, or later (2017). This is available on the installation DVD for Windows-based licensing and is in the `Tools` directory for Linux licensing.

---

For a new installation, you must set up and configure a license server to provide your users with access to the software (if you are updating an existing installation the license server will already be configured). Typically, you use Network Attached Storage (NAS) for the license server and choose a location that is accessible by the cluster and your users' PCs.

The licensing software is installed along the Linux tools. You can get the software and the "*Schlumberger Licensing User Guide*" from the tools directory `/ecl/tools/linux_x86_64/flexlm1114/`.

## Configure Intel MPI or Platform MPI for Linux

Intel MPI and Platform MPI are installed with the simulator installation. No separate installation is required. Please note the following:

- Both of these MPIs use ssh to set up the communication between nodes, so please make sure that password-less ssh is set up between any computers that may be used to run a simulation job in serial or parallel.
- Add some tuning parameters to the Intel MPI configuration if you are using this MPI.
- You must set Platform MPI interconnect settings if you are using this MPI.

## Check and configure ssh for use with the MPIs

To set up the ssh connections, first check to see if the ssh connections are working. If they are not, generate the public and private keys for each user and make sure that the `authorized_keys` file exists and has the correct keys inside. You then update each user's ssh configuration files with the correct settings and access permissions. Using this ssh setup means that connections to new machines are automatically accepted (if you don't do this it means that you have to log on to every machine in the cluster before you could guarantee that a MPI job will work correctly).

---

**CAUTION:** If you support other applications which rely on ssh to work, you may need to re-evaluate these instructions as they could affect the other applications.

---

## Check the ssh connections

Check that the connections are working by trying to connect to other machines in the cluster. For example, try to connect from `comp001` to `comp002` using ssh with no password:

```
comp001:/home/user>ssh comp002
Last login: Thu Apr 9 10:37:03 2009 from comp001.geoquest
comp002:/home/user>
```

If the connection is successful, follow the instructions in [Update the user configuration information](#) to set up each user's configuration files. If not, follow the instructions in the flowchart to generate the authentication keys.

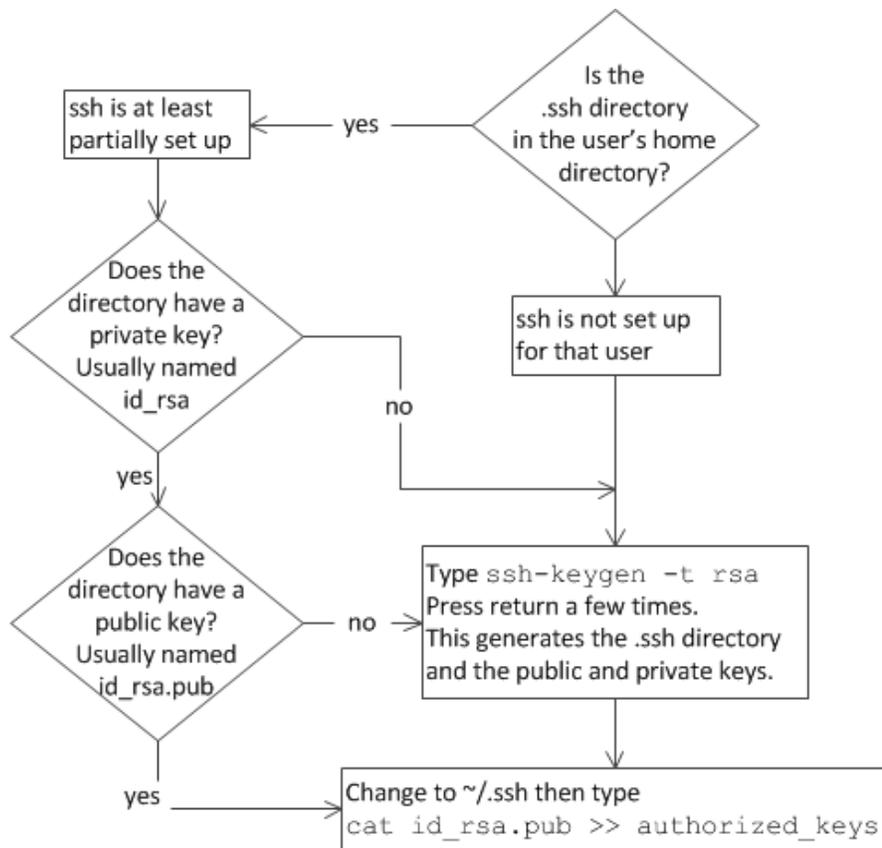


Figure 4.2. ssh setup

## Update the user configuration information

To create or update each user's configuration files and set the correct access permissions:

1. Either create the `config` file in the user's `~/.ssh` directory, or add the line if it already exists.

```

comp002:/home/user>cd ~/.ssh
comp002:/home/user/.ssh>cat config
StrictHostKeyChecking=no
  
```

2. Change the file permissions for the `config` file to 400 if necessary.
3. Change the file permissions for the `authorized_keys` to 600 if necessary.
4. Set the permissions of the `$HOME` and `$HOME/.ssh` directories so that other users do not have write access to them.

## Intel MPI settings

There are some tuning parameters which you can add to the `mpiexec.conf` file which can improve performance. You can find this file in the directory

`ecl/tools/linux_x86_64/intel/mpi/<mpi_version>/etc`

and `ecl/tools/linux_x86_64/intel/mpi/<mpi_version>/etc64`

or in the local work directory. For the <mpi\_version>, see [Linux cluster](#). Add the following tuning parameters to the file:

```
-genv I_MPI_PIN_PROCESSOR_LIST allcores:map=scatter
-genv I_MPI_PIN_CELL core
```

The Intel MPI detects and uses the correct interface. Setting an environment variable will force an interconnect to be chosen. The interconnect settings are shown in the following table.

Device Type	Description	Setting
OFED	OFA interface is used by 10G iWarp, InfiniBand and Infinipath cards.	ofa
Standard Ethernet	Standard 100Mbps, 1Gbs and 10Gbs network cards.	tcp
Infinipath	Infinipath cards from QLogic.	tmi

Table 4.1: Intel MPI interconnect settings

Set the I\_MPI\_FABRICS\_LIST environment variable as follows:

```
setenv I_MPI_FABRICS_LIST ofa,tcp
```

This sets Intel MPI to try the chosen devices in order.

I\_MPI\_FALLBACK and I\_MPI\_DEVICE are no longer required and should be removed if set.

To print out the connection or interface type used set the following environment variable:

```
setenv I_MPI_DEBUG 2
```

For more information, see <http://software.intel.com/en-us/intel-mpi-library/>

**Note:** The latest Intel MPI HYDRA launch mechanism is a 32-bit application that starts the simulator processes. This means 32-bit libraries must be installed for this to work. On Redhat 6, the required rpm is called `glibc-2.12-1.47.el6.i686.rpm`. Follow the on-screen instructions and prompts to complete the installation.

## Platform MPI settings

The Platform MPI should automatically determine and use the best interface. If you wish to set the interface, then it is possible by setting an environment variable. The search order for the interconnect is determined by the environment variable MPI\_IC\_ORDER (which is a colon-separated list of interconnect names). The most common options are shown in the following table

Device Type	Description	Setting
InfiniBand (IB)	InfiniBand can be SDR, DDR or QDR.	IBV
Infinipath	Infinipath made by QLogic	PSM
uDAPL (10g iWarp Ethernet and InfiniBand)	uDAPL interface is used by 10G iWarp network cards. Additionally it can also be used by most types of InfiniBand and Infinipath.	UDAPL
Standard Ethernet	Standard 100Mbps, 1Gbs and 10Gbs network cards.	TCP

Device Type	Description	Setting
Myrinet	GM based Myrinet cards.	GM

Table 4.2: Platform MPI interconnect settings

## Example

To set the MPI to use the DAPL device, add the following line to the user's `.cshrc` file.

```
setenv MPI_IC_ORDER UDAPL
```

or

```
setenv MPI_IC_ORDER udapl:tcp
```

Using upper case for UDAPL sets the instruction to use that device or exit. Using lower case sets the instruction to try the device and, if it doesn't work, try another. Shared memory is always implemented if possible. The standard ECLIPSE scripts will add a few extra lines into the top of the output files and these will display the interconnects used.

The following extract from the `FILENAME.OUT` file uses TCP to communicate between nodes.

```
host | 0 1
=====|=====
0 : SHM TCP
1 : TCP SHM
```

Alternatively you can override it by using command line arguments to the `eclrun` script. The following arguments force the Platform MPI to use the IBV setting, regardless of the environment variable setting.

```
eclrun -c plmpi --mpi-args="-IBV" eclipse DATASETNAME
```

## Set up user accounts in the Linux cluster

After installation, each user's environment needs to be set up to reflect the simulator installation. A sample `.cshrc` file is shown below. Edit the paths if required, and copy to each user's home directory.

```
#!/bin/tcsh
#Edit the line below to reflect the location of the license server.
#LM_LICENSE_FILE was used by pre 2007.2
setenv SLBSLS_LICENSE_FILE 1234@licenseserver.com

#Edit the line below to set the install location.
setenv ECLPATH /ecl

#Set Endian so Linux files line up with other OS's
if ( `uname` == "Linux" ) then
    setenv F_UFMTENDIAN big
endif

#Source eclrunsetup if it exists, so eclrun works
if ( -r $ECLPATH/macros/@eclrunsetup.csh ) then
    source $ECLPATH/macros/@eclrunsetup.csh
endif

#Source LSF environment script
if ( -r /lsf/conf/cshrc.lsf ) then
    source /lsf/conf/cshrc.lsf
endif
```

## Install ECLRUN in the Linux cluster

This section provides some information on how to set up and configure ECLRUN to allow your users to run simulation jobs on the HPC cluster. For a full description of the configuration options and settings, see the *"ECLRUN User Guide"*.

### ECLRUN passwords

The *"ECLRUN User Guide"* contains a list of restricted characters which cannot be used in passwords as ECLRUN cannot pass them to the cluster.

### Migrate ECLRUN settings

Migrate any custom settings from the old configuration files into the configuration files in the new macros directory. Settings may be stored in any of the following files:

- CONFIG.ECL
- eclrun.config
- @eclrc

Do not just copy the old files into the new macros directory. From 2009.1 onwards, `eclrun.config` uses an XML format which differs from that required in previous releases.

### Configure the schedulers to work with ECLRUN

License scheduling is available with the simulator and PBSPro, LSF or UGE. To install and configure the PBSPro, LSF or UGE integration kits, see the instructions on the installation disk in the folders `3rdparty/PBSPro`, `3rdparty/LSF` or `3rdparty/UGE`.

Regardless of your choice of scheduler, you must update the ECLRUN configuration settings:

1. In `eclrun.config` find the setting:  
`<LSFLicenses>False</LSFLicenses>`
2. Change the setting to:  
`<LSFLicenses>TRUE</LSFLicenses>`

The `LSFLicenses` setting is used by all of the supported schedulers.

---

**Note:** Only ECLRUN is supported. The old scripts starting with the character @ are not able to schedule licenses.

---

## Test the installation

The following information may be useful for a simple tests

## Test ECLIPSE with an example data file

To test an ECLIPSE 100 dataset:

1. Create a temporary directory, and copy an example .DATA file from /ecl/<version>/eclipse/data/ into this folder.
2. Change to the temporary directory and run the example using the command `eclrun eclipse <filename>`.

You can perform similar tests on ECLIPSE 300 and FrontSim datasets.

## Run INTERSECT from the command line

The Migrator can convert an ECLIPSE keyword file into an AFI file, which can then be read by INTERSECT.

- The following ECLRUN command line generates the AFI file but does not run or simulate it:  
`eclrun ecl2ix basename` where **basename** is the root name of the input dataset.
- To run or perform the actual simulation using the AFI file the following command is used: `eclrun ix basename`. INTERSECT will read the AFI file and perform the simulation.

If the input ECLIPSE dataset is altered, the AFI file has to be regenerated from the dataset first before being given to INTERSECT to be simulated.

An AFI consists of {**basename**}.afi, and its associated files **basename\_ECL2IX.gsg**, **basename\_ECL2IX\_IX.ixf**, **basename\_ECL2IX\_FM.ixf**, **basename\_fm\_edits.ixf** and **basename\_reservoir\_edits.ixf**.

## Test INTERSECT with an example data file

1. Create a temporary directory, and copy the example EX2.DATA from /ecl/2013.1/ix/Example/Example\_2/ into this folder.
2. Change to the temporary directory and run the example using the command `eclrun ecl2ix EX2..`

You should see a simulation run and exit normally, as below:

```

...
REPORT   Keyword migration summary:
+-----+
| Fully migrated      | 49 |
| Partially migrated  |  5 |
| Not migrated        |  6 |
| Ignored              |  8 |
| Total                | 68 |
+-----+

INFO EX2.afi is created.
ix_tools checked in
INFO Successfully checked in license feature 'ix_tools'
INFO Migration complete. Elapsed time 0hrs:0mins:10secs (10secs)
REPORT   Message summary
+-----+
| Message Level | Frequency |
+-----+
| INFO          |         16 |
| WARNING       |         17 |
| ERROR         |          0 |
+-----+

```

```
+-----+  
INFO Run finished normally.
```



## Legacy macros and scripts

---

---

**Note:** We recommend you use ECLRUN to run most of the simulators and other programs. Please refer to the *"ECLRUN User Guide"* for more details. The legacy macros are still released but may have limited capabilities compared to ECLRUN.

---

For Linux systems, the software can be run using scripts supplied in the macros directory in place of ECLRUN. The following scripts are available to run the License Manager:

@flexstart Start up the License Manager  
@lmdown Shut down the License Manager  
@lmhostid Get FLEXlm hostid for the machine  
@lmreread Reinitialize License Manager  
@lmstat Check license status

The default location for the following files for use with the License Manager is the `ecl/macros` directory:

license.dat License file  
flex.log Activity log for License Manager

The following macros are available to run the principal simulator software programs:

@e300 Compositional simulator  
@eclipse Black oil simulator  
@frontsim Streamline simulator  
@graf Post processor  
@grid Gridding and mapping  
@mpieclipse Run ECLIPSE 100 simulator in parallel mode

@mpie300      Run ECLIPSE 300 simulator in parallel mode

The following macros are available to run the utility programs (those marked \* are no longer released):

- Batch operations:
  - @fill      Corner point geometry generation
- File format conversions:
  - @convert      Converts formatted / unformatted ECLIPSE output
  - @convertgrid      Converts between .GRID and .EGRID formats
  - @dos2unix      Converts ASCII from MS-DOS to UNIX format
  - @rescueconvert      Converts from older versions to RESCUE Version 19 format
  - @unix2dos      Convert ASCII from UNIX to MS-DOS format
- File manipulation:
  - @expand      Merge INCLUDE files into master file
  - @extract      Extract subset of data from ECLIPSE output
  - @glkread      Checks the contents of a Stratamodel binary .glk file
- File name changes:
  - @change\_prefix      Changes prefix of file name
  - @change\_suffix      Changes suffix of file name
  - @lower      Lowercase file name
  - @rename      Change ECLIPSE filename case and extensions
  - @upper      Uppercase file name
- Viewing:
  - @manuals      Launches Manuals bookshelf
  - @pdf      Launches Adobe Acrobat

---

**Note:** Not all programs are supplied with the 2013.2 release; some older applications have been retired. The scripts have been left so they can be used to run older versions.

---

- The following environmental variables can be set to control LSF functionality within the macros:
  - ECL\_LSF      set true if LSF is installed
  - ECL\_LSF\_BYDEFAULT      If this and ECL\_LSF are set to TRUE then jobs are submitted using LSF without the need for the `-lsf` flag. For compatibility with ECLIPSE

Office the `-nolsh` flag has been added as a new option to stop nested `bsub` commands.

<code>ECL_LSFHPC</code>	This switches on the LSF HPC extensions, but it is recommended that you do not use this extension and remove it from all settings. However, if you must use LSF HPC make sure that you follow the configuration steps in <a href="#">"Configuring changes if you require LSF HPC"</a> .
<code>ECL_LSF_PARALLELQ</code>	defines the default queues for parallel jobs
<code>ECL_LSF_SERIALQ</code>	defines the default queues for serial jobs
<code>ECL_LSF_LICCHECK</code>	set if LSF is to control FLEXlm licensing and the LSF-SIS integration kit has been installed (contact Platform Computing <a href="http://www.platform.com">www.platform.com</a> for details).
<code>ECL_MR_SCHEDULING</code>	set this variable to correctly queue multiple realization jobs. Follow the instructions in the LSF-SIS integration kit.

## Macros: command line flags

The Linux macros are written using c-shell scripts. It is essential that the macros directory is placed on the user's path; this is not necessary, however, to ensure their successful execution. The macros will automatically detect all versions of software installed, and will invoke questions as necessary to launch the program required.

The macros require little or no configuration to launch interactive products. However, it is possible to configure the macros using command line flags. This is especially useful for running the ECLIPSE simulators. The main command line flags for running the ECLIPSE 100 and ECLIPSE 300 simulators are listed below:

<code>-data /home/user/workarea</code>	to set default working directory, or other directory.
<code>-file &lt;filename&gt;</code>	to define a dataset file name.
<code>-help</code>	to obtain a list of available flags with brief description.
<code>-second &lt;Title&gt;</code>	to define a secondary run title. <code>&lt;Title&gt;</code> should be enclosed in quotes.
<code>-fore</code> or <code>-foreground</code> or <code>-fg</code>	defines a foreground run (the default)
<code>-back</code> or <code>-batch</code>	defines a background (queued) run
<code>-priority</code>	defines priority of a background run
<code>-ver</code> or <code>-vers</code> or <code>-version</code>	defines version of software to run, for example: <code>-ver 2005a _1 -file TESTCASE</code>
<code>-proc</code> or <code>-procs</code> or <code>-processors</code>	defines number of processors for a Parallel run.
<code>-local</code>	uses the local configuration file <code>ECL.CFG</code> .

For parallel ECLIPSE, namely @mpieclipse and @mpie300 the following special options are available:

-hostfile hosts.txt Configuration file for Parallel run (same as -machinefile)

-machinefile hosts.txt Configuration file for Parallel run (same as -hostfile)

-lsf For runs using LSF.

-ilmpi Option for processing parallel jobs using Intel MPI

-plmpi Option for processing parallel jobs using Platform MPI

-ibmmpi Option for processing parallel jobs using IBM POE MPI

If -procs is not set, the macro will attempt to read the PARALLEL keyword in the dataset to get the number of processors.

-nolocal Do not run on front end when a parallel job is run.

## Examples

```
@eclipse -version 2018.1 -local -lsf TESTCASE
@mpie300 -lsf -lsfqueue normal -version 2018.1 -local TESTCASE
```

# B

## Installation DVD content

---

The software is distributed on a disk which contains:

- The simulators for all x86 64-bit architectures.
- Intel MPI Runtime installation.
- Documentation in PDF (Portable Document Format).
- License server and dongle drivers.
- Utility resources which contain:
  - Benchmark and datasets (ECLIPSE)
  - PC resources
  - Scheduler integration kit.