Barracuda Virtual Reactor User Manual

Still have a question? Email us anytime (support@cpfd-software.com), or call us at +1.713.429.1252 from 9am-5pm Central Time.

2. Installation

2.1. Overview

This chapter documents the procedures for installing *Virtual Reactor* and its license server, the Reprise License Manager (RLM). The following steps lead you through the process of getting started with *Virtual Reactor*:

1. Create a CPFD Support site account in order to access downloads
2. Acquire hardware that meets recommended System Requirements
3. Plan the Client-Server Configuration
4. Download Virtual Reactor
5. Install the RLM Server
6. Install the RLM License
7. Install Virtual Reactor
8. Point Client Machines to the RLM Server
9. Install NVIDIA Drivers

In order to use *Virtual Reactor*, you must have a valid RLM software license. *Virtual Reactor* uses a client-server model for licensing, meaning that your license(s) can be placed on one machine, which acts as a license server, and other machines on the same network can point to and request licenses from that server. The server may also use the license(s) that it hosts. *Virtual Reactor* also supports RLMCloud, which is a cloud-based solution for serving RLM licenses.

The licensing policy for *Virtual Reactor* is as follows:

- Each solver license allows the user to run one instance of the *Virtual Reactor* solver. Optionally, add-ons to each solver license may be purchased to include Chemistry capabilities and/or Parallel capabilities.

- The number of *Virtual Reactor* solvers that can be run simultaneously is limited to the number of solver licenses that have been purchased.

- With any license, an unlimited number of *Virtual Reactor* Graphical User Interfaces (GUIs) can be open at the same time. This allows users to set up simulations, review previously set up models, and use all functions of the GUI.

- Running a single time step does not count against the number of *Virtual Reactor* solver licenses in use. This allows users to review a project setup at any time.

- An unlimited number of Tecplot for Barracuda post-processing instances can be open at the same time.

- If Parallel add-on licenses have been purchased, *Virtual Reactor* can utilize NVIDIA GPU cards to increase calculation speed. The number of simulations that can use GPU acceleration simultaneously is limited to the number of Parallel add-on licenses purchased.

- If Chemistry add-on licenses have been purchased, *Virtual Reactor* can include user-specified chemical reactions in the simulation. The number of simulations that can solve chemistry simultaneously is limited to the number of chemistry add-on licenses that have been purchased.

- If Parallel and Chemistry add-on licenses have been purchased, *Virtual Reactor* can use multiple CPU cores for parallelization of volume-average chemistry to increase calculation speed. The number of simulations that can use CPU parallel acceleration simultaneously is limited to the number of Parallel add-on licenses purchased.

- The RLM license server must be run on a physical machine. A virtual machine (VM) cannot be used as the RLM license server host.

- The RLMCloud service is only available to customers with a current lease or maintenance contract for *Virtual Reactor*.

2.2. System Requirements

Virtual Reactor can be used on either Linux or Windows. The following table lists minimum and recommended system requirements.

	Minimum	Recommended
Operating System	Recent 64-bit Linux 64-bit Windows 7, 8, or 10	64-bit CentOS 6 (RHEL 6) or higher Windows 10 Pro 64-bit
CPU	Any 64-bit Intel compatible from the last 5 years	Intel Core i7-9800X (4.4 GHz, 8 cores, 16.5 MB cache) or better. Higher clock speed and newer Intel architecture are better.
Memory (RAM)	8 GB	Twice as much as your GPU memory. Faster is better.
Hard drive space	500 GB	Boot Drive: M.2 NVMe SSD Data Drives: Two 8 TB 7200 RPM HDDs
GPU *	NVIDIA GPU required	NVIDIA Titan RTX NVIDIA Quadro RTX 8000 for larger simulations where 24 GB is not enough.
CUDA Compute Compatibility *	2.0	3.5 or higher
GPU RAM *	4 GB	24 GB or more (see above)

* - GPU requirements only apply if *Virtual Reactor* will be running in GPU parallel mode. The required GPU RAM is also dependent on the size of the simulation being run. A larger simulation will require more GPU RAM. We recommend running a single Barracuda calculation per GPU, so if you have more than 1 license of Barracuda, you should buy a GPU for each license.

Though *Virtual Reactor* simulations can be run on laptops, or lower-performance desktop machines, doing so is generally not recommended. Investing in an up-to-date calculation machine, with the fastest hardware currently available, will provide much faster calculation speeds. Additionally, since computer hardware advances in capacity and speed at such a fast pace, it is recommended to purchase updated hardware every 2 to 3 years to obtain the fastest performance.

Virtual Reactor can be installed on compute nodes of a cluster. However, it will not take advantage of the multi-node parallel computing capabilities of the cluster. Each *Virtual Reactor* simulation utilizes the computing resources of a single machine. Parallelization is available through the use of an NVIDIA GPU, or through CPU-parallel solution of volume-average chemistry. Since each individual node of a cluster is not usually optimized for the fastest possible single-machine CPU performance, it is often the case that running *Virtual Reactor* on a cluster node will not give the best possible calculation speed. Instead, it is generally better to purchase a very fast single-CPU standalone calculation machine on which to run *Virtual Reactor*. This standalone machine will outperform a cluster node in the majority of cases, for the purpose of running *Virtual Reactor*.

2.3. Plan the Client-Server Configuration

It is possible to configure RLM in several ways, depending on your needs and preferences. *Virtual Reactor* can use an on-premise RLM server, RLMCloud's hosted cloud server, or a combination of the two. If you only have a single machine that will be running *Virtual Reactor* simulations, it may be most convenient to install the RLM server directly on that machine, so that it acts as both RLM server and RLM client:

Fig. 2.1 Single machine acting as RLM server and client

It is also possible for such a machine to act as the RLM server for additional client machines. In the example shown below, two additional large calculation machines are clients, and one laptop is also a client. In general, laptops are not powerful enough to be useful for running large simulations, but it is often convenient to set up simulations on a laptop and then transfer them to a larger calculation machine.
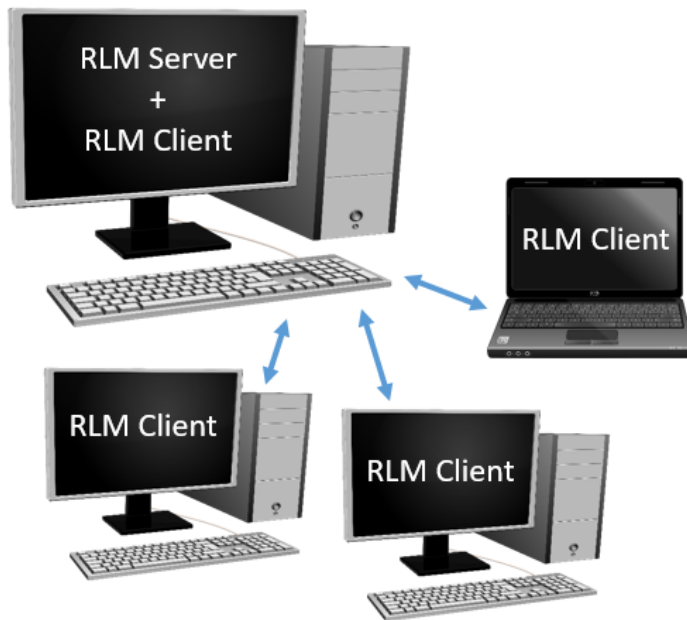

Fig. 2.2 Single machine acting as RLM server and client, with multiple clients

Another option is to have a standalone RLM license server which is not intended to run simulations. This is common when a centralized license server already exists, or when an always-on file server is able to act as the RLM server as well.
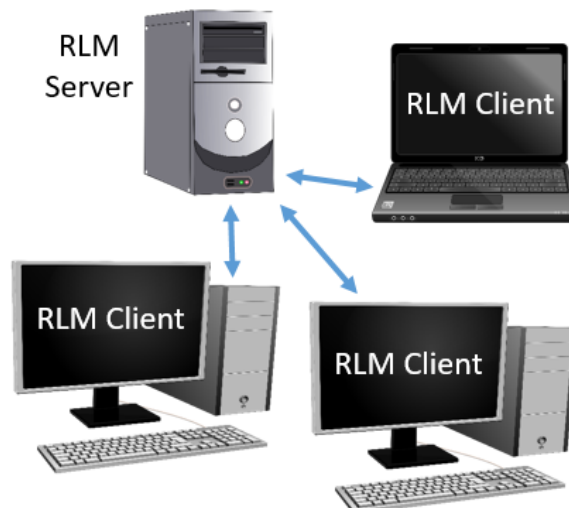

Fig. 2.3 RLM server with multiple clients

For any of the above on-premise RLM server configurations, it is important to note that the RLM server machine must remain on and accessible to the RLM client machines during simulations. If the RLM server is

powered off, or loses network connectivity to a client machine that is running a simulation, the *Virtual Reactor* solver will pause and wait for its license to become available again.

> **Note**
>
> 1. When using an on-premise RLM server, the RLM server and all RLM client machines must be on the same network, or the network must be configured to allow all machines to see each other, in order for the *Virtual Reactor* licenses to work.
> 2. An on-premise RLM server must be a physical machine. A virtual machine (VM) cannot be used as the RLM license server host.

Virtual Reactor also supports RLMCloud, an easy-to-use, cloud-based solution for serving RLM licenses. RLMCloud is hosted by [Reprise Software](). Advantages of using RLMCloud include:

1. Users and system administrators do not need to configure and maintain an on-premise RLM server. Instead, a user-level RLMCloud license file is specified in the *Virtual Reactor* GUI license server manager, and *Virtual Reactor* is ready to run.
2. RLMCloud is compatible with cloud-based computing platforms such as [Amazon EC2](), [Microsoft Azure](), and [Google Cloud]().
3. With RLMCloud, increasing your *Virtual Reactor* license count or adding short-term licenses is a seamless and no-hassle process. CPFD, or your local in-country distributor, can quickly adjust licenses as needed, with no license administration burden for users.
4. RLMCloud enables future licensing options such as on-demand overflow licensing capabilities. Contact [licensing@cpfd-software.com]() to learn more.
5. RLMCloud includes a customer portal so that *Virtual Reactor* users have visibility into their license servers. This portal allows you to view current license server status as well as users of licenses, view and download server log files, and edit license server options.
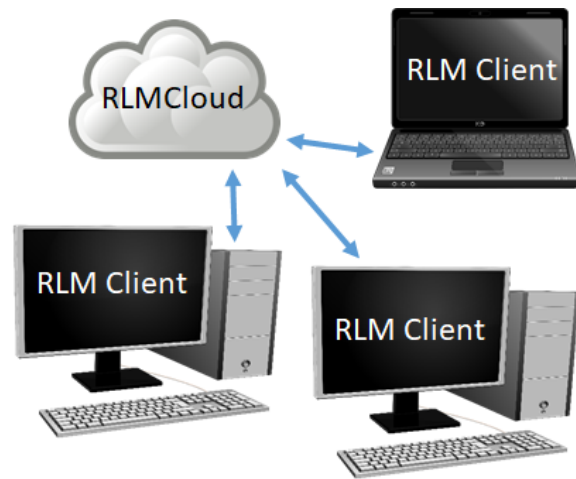

Fig. 2.4 RLMCloud with multiple clients

> **Note**
>
> The RLMCloud service is only available to customers with a current lease or maintenance contract for *Virtual Reactor*.

2.4. Download *Virtual Reactor*

The current version of *Virtual Reactor* is available for download on the CPFD support site:

https://cpfd-software.com/customer-support/knowledge-base/category/release-versions

When downloading, be sure to choose the appropriate distribution for your operating system (Linux or Windows).

The RLM license server installation program is a separate download from the *Virtual Reactor* distribution. If you will be running an on-premise RLM license server, be sure to download the RLM installation package for the appropriate operating system of the RLM license server (Linux or Windows).

Once downloaded, the installation files must be extracted from the compressed download files by using a decompression utility. This will generate the installation files in the same directory as the compressed file. You may now Install the RLM Server (if necessary), or Install Virtual Reactor.

Physical installation media can be sent upon request via email to licensing@cpfd-software.com (please specify Windows or Linux and USB stick or CD).

2.5. Install the RLM Server

Note

1. If you will be using RLMCloud, you do not need to perform this step. No RLM server installation is needed. Instead, please proceed to Install Virtual Reactor and Point Client Machines to the RLM Server.

2. You do not need to install or update the RLM server for each new version of *Virtual Reactor*. Once the RLM server is set up the first time, it will work with all versions of *Virtual Reactor*.

3. Administrator (on Windows) or `root` (on Linux) privileges are required in order to perform the RLM server installation.

4. An on-premise RLM server must be a physical machine. A virtual machine (VM) cannot be used as the RLM license server host.

The Reprise License Manager is installed by running **reprise-installer.run** (on Linux) or **reprise-installer.exe** (on Windows), extracted from the RLM server download. At the end of the Reprise License Manager installation, be sure to leave the box checked for *Show system information*.
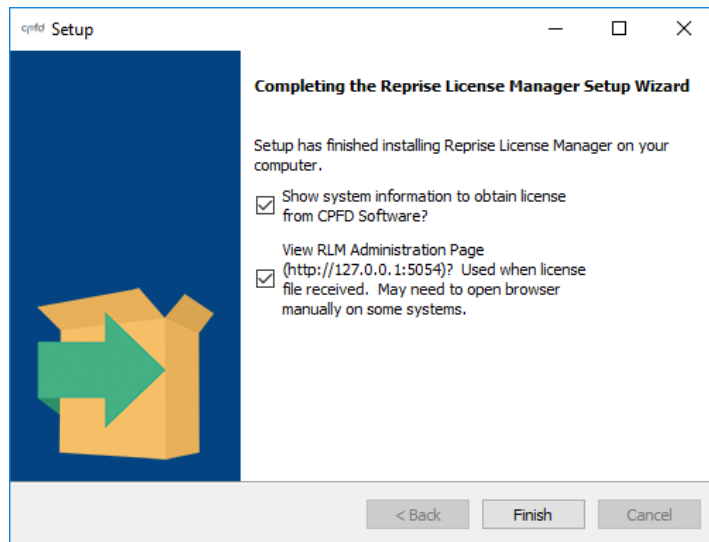

Fig. 2.5 Final screen of RLM installer

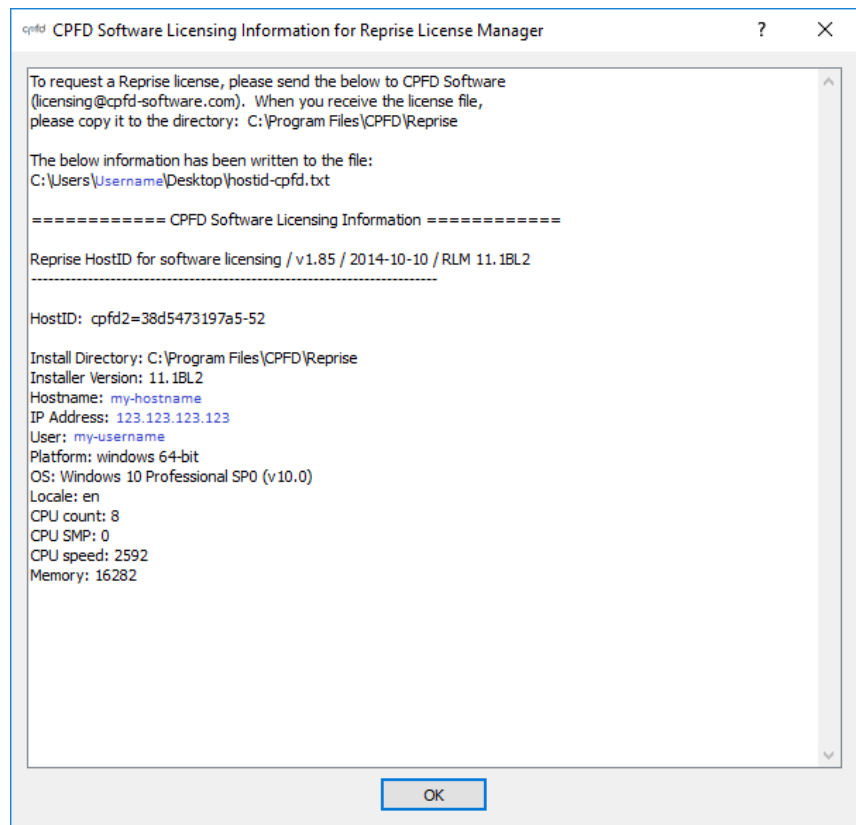A window will pop up showing your system information.

Fig. 2.6 RLM installer host information

Additionally, a text file containing the same information will be saved to your `Desktop` with a file name of `hostid-cpfd.txt`. Email this file to licensing@cpfd-software.com, and CPFD Software will generate a license file and email it to you.

An alternative method of finding the license server's `hostid` is to run the command **getcpfdid** in a terminal or command prompt.

On Linux, the default installation location of this command is:

`/usr/local/bin/reprise/getcpfdid`

On Windows, the default installation location of this command is:

`C:\Program Files\CPFD\Reprise\getcpfdid.exe`

2.5.1. Firewall Configuration

In order for RLM client machines to communicate with the RLM server, the server's firewall must have ports 5054, 27015, and 27016 open.

- If the OS of the RLM server is **CentOS 7 Linux**, the firewall ports can be opened by typing the following commands in a terminal as `root`:

```
firewall-cmd --permanent --zone=public --add-port=5054/tcp
```

```
firewall-cmd --permanent --zone=public --add-port=27015/tcp
```

```
firewall-cmd --permanent --zone=public --add-port=27016/tcp
```

```
firewall-cmd --reload
```

- If the OS of the RLM server is **Windows**, the firewall ports can be opened by using an [Administrator Command Prompt](). In the prompt, type the following commands:

```
netsh advfirewall firewall add rule name="RLM Web Server" dir=in action=allow
protocol=TCP localport=5054
```

```
netsh advfirewall firewall add rule name="RLM License Server" dir=in action=allow
protocol=TCP localport=27015
```

```
netsh advfirewall firewall add rule name="RLM ISV Server" dir=in action=allow
protocol=TCP localport=27016
```

2.6. Install the RLM License

> **Note**
>
> 1. If you will be using RLMCloud, you do not need to perform this step. Instead, please proceed to Install Virtual Reactor and Point Client Machines to the RLM Server.
>
> 2. The license file only needs to be installed on the RLM server. It does not need to be installed on client machines.

Once you have requested and received your RLM license file from CPFD Software, it is necessary to install it on the RLM server. To install the RLM license file, follow these steps:

1. Save the RLM license file to your `Desktop`.

2. Copy the license file to the following location depending on the OS of your license server:

   ### Linux

   `/usr/local/bin/reprise/`

   You will need root permission to do this. In a terminal opened to the same directory as the license file, type:

   ```
   su
   <Enter root password>
   cp <license file>  /usr/local/bin/reprise/
   ```

   where `<license file>` is the RLM license file name.

   ### Windows

   `C:\Program Files\CPFD\Reprise`

   Note that Windows Explorer will ask for administrator permission to copy a file into any directory inside `C:\Program Files`.

3. Open a web browser to: http://127.0.0.1:5054

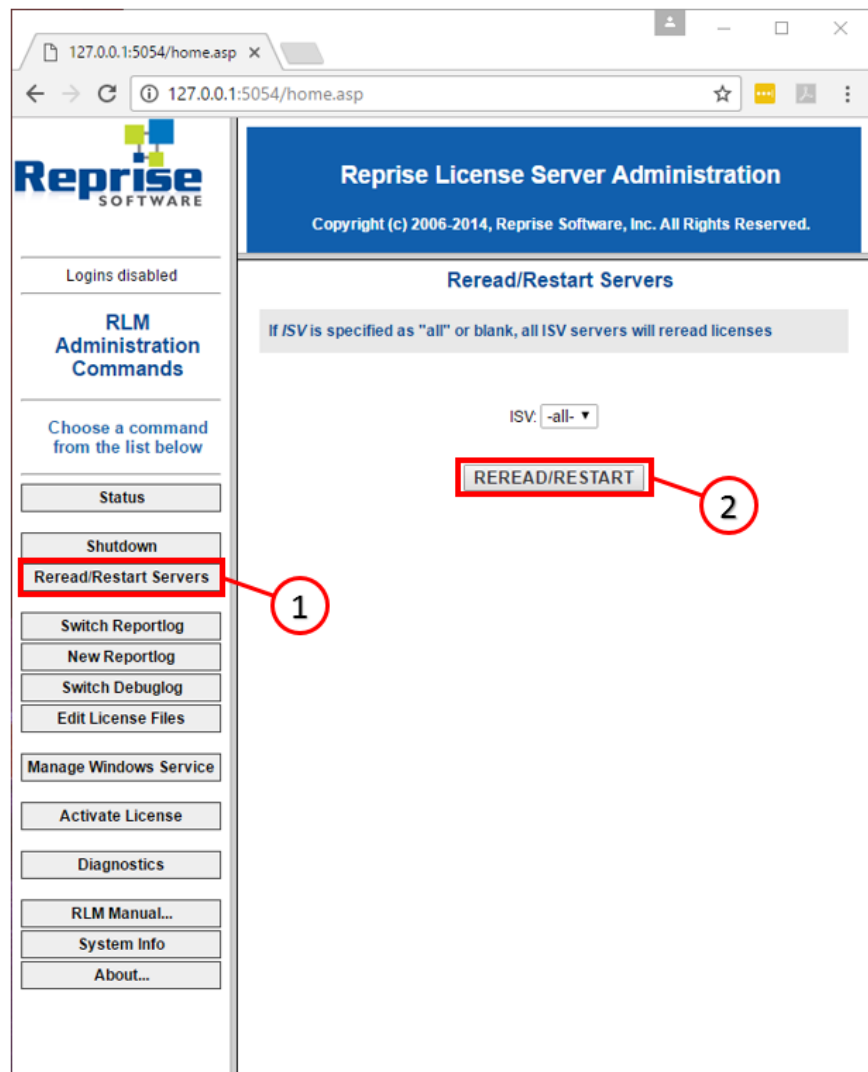4. Click *Reread/Restart Servers*, then *REREAD/RESTART*:

Fig. 2.7 RLM Reread/Restart Servers

It is also possible to request that the RLM server reread its license file(s) via the command-line, in case you do not have access to the web interface of the RLM server. In the `reprise` folder, running the script named `rlmreread` will trigger a reread of the license file(s) for the RLM server. You must have administrator privileges to run this script.

5. Click *Status*, and then click the first *cpfd* button under the *Server Status* column. This lists the current status of your Reprise server licenses and shows licenses available.
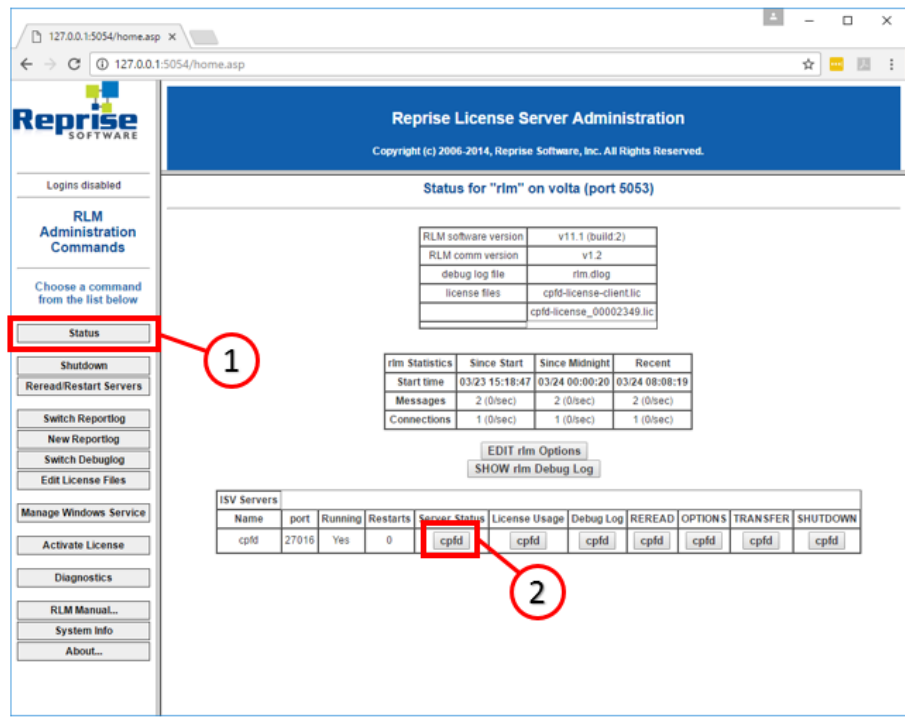
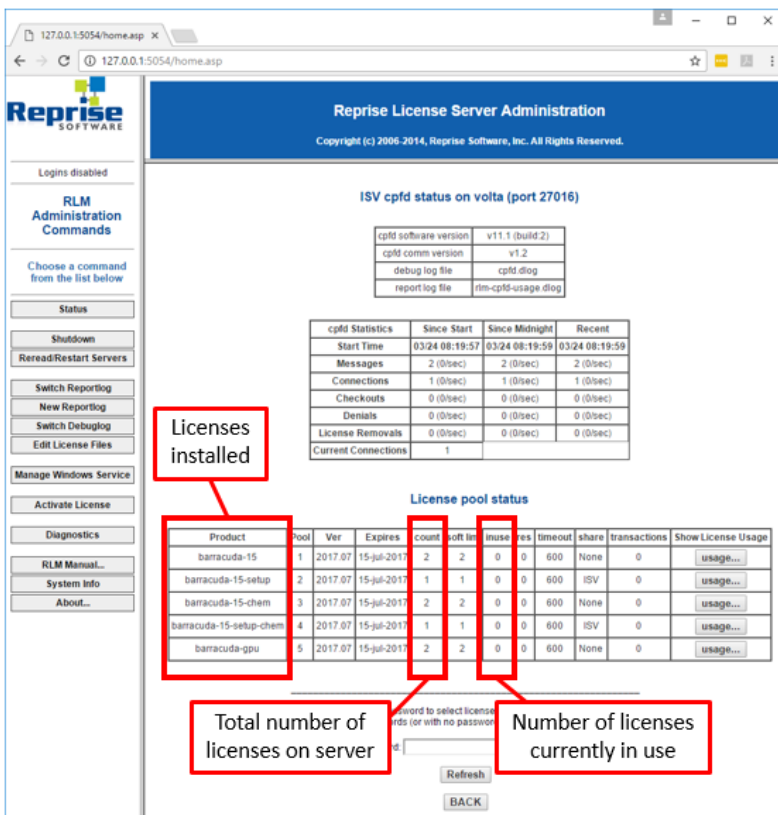*Fig. 2.8* Checking RLM license status

Fig. 2.9 RLM license status table

2.7. Install *Virtual Reactor*

Virtual Reactor must be installed on every client machine. Multiple versions of *Virtual Reactor* can be installed on a single machine, and each version operates independently. It is not necessary to uninstall old versions of *Virtual Reactor* when installing a new version. Download Virtual Reactor and then follow the steps below to install the software:

1. Navigate to the *Virtual Reactor* installation folder that was extracted.

2. Double-click on `barracuda_setup.run` (on **Linux**) or `barracuda_setup.exe` (on **Windows**) to start the *Virtual Reactor* Setup Wizard. Click *Next*.
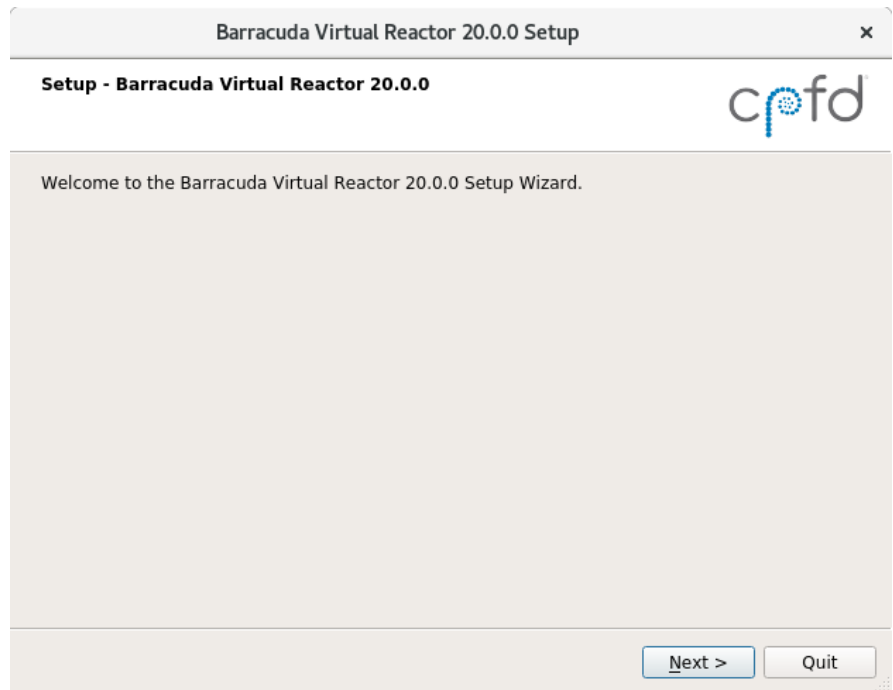

Fig. 2.10 *Virtual Reactor* Setup Wizard

3. This window shows the default directory in which the *Virtual Reactor* version will be installed. It is recommended to use the default directory. Click *Next*.
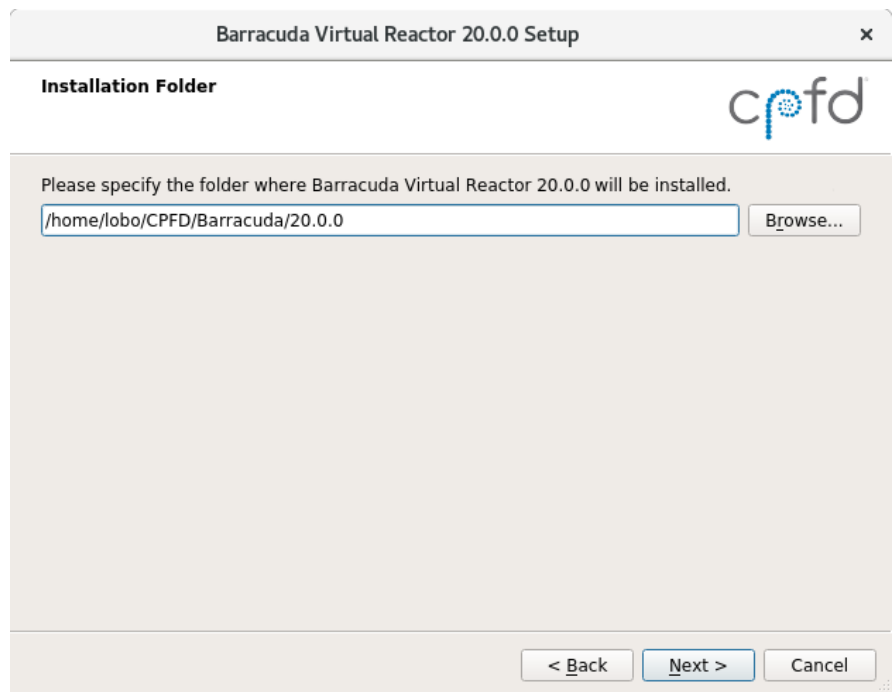

Fig. 2.11 *Virtual Reactor* Installation Directory

4. This window allows you to choose the components to be installed. It is recommended that you leave all components checked and click *Next*.
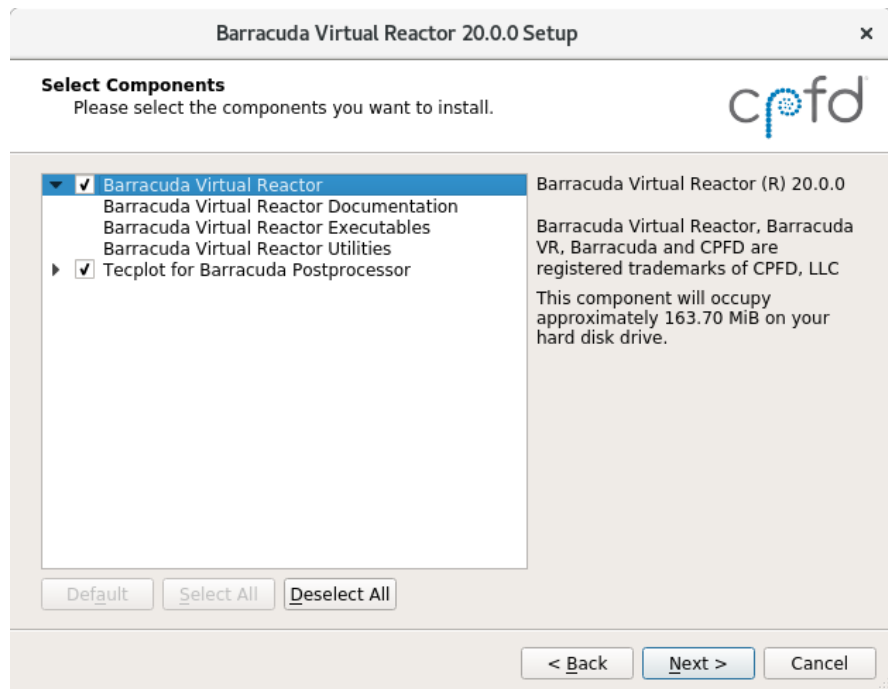
Fig. 2.12 Virtual Reactor Setup Select Components

5. Read and accept the License Agreement and click *Next*.
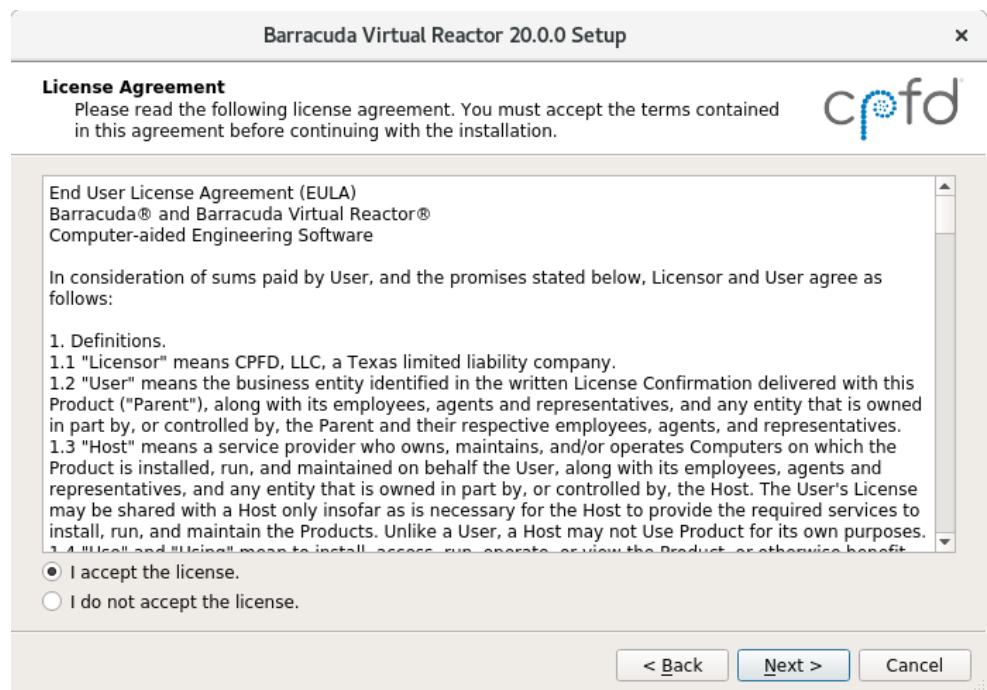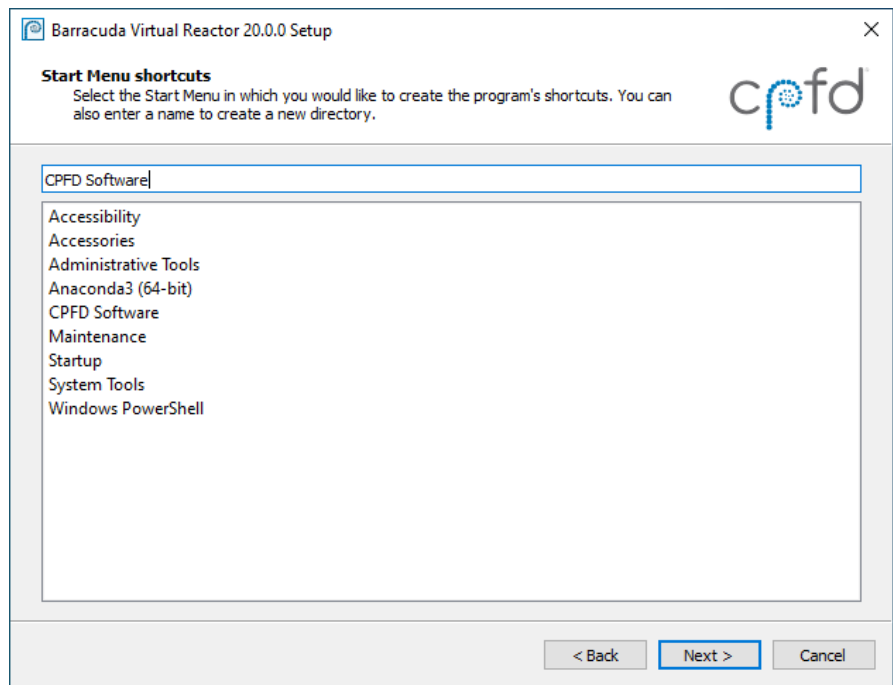

Fig. 2.13 Virtual Reactor License Agreement

6. On Windows, create a Start Menu shortcut and click *Next*.

7. Click *Next* in order to begin the installation.

The installer will display a progress bar as it installs *Virtual Reactor*.
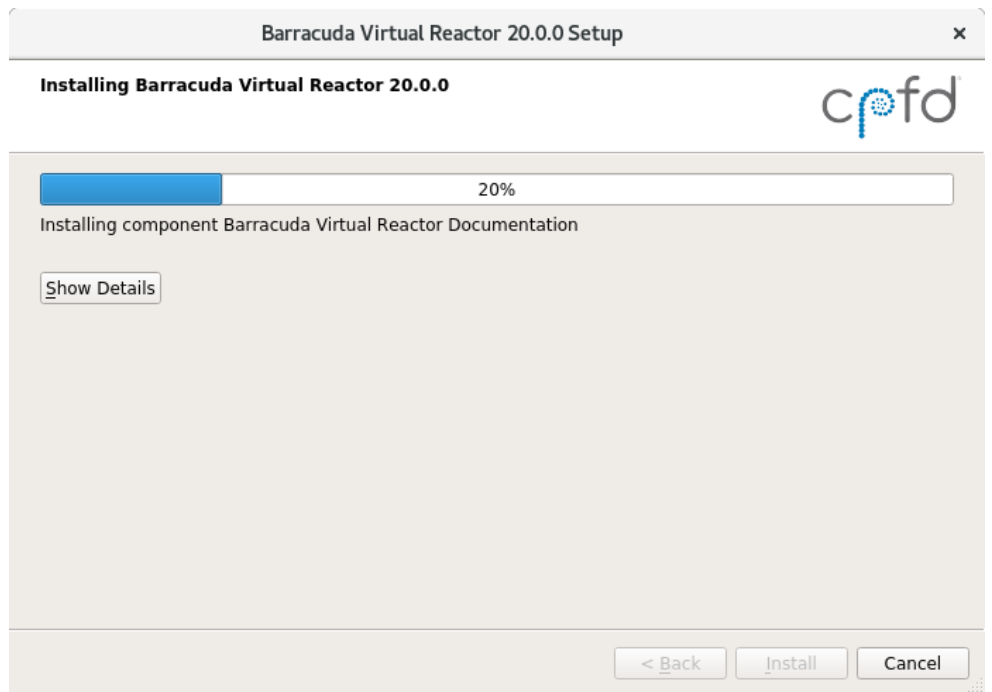
Fig. 2.16 Virtual Reactor Setup Installing

8. Once the installation is complete, the final window of the *Virtual Reactor* Setup Wizard will be displayed.
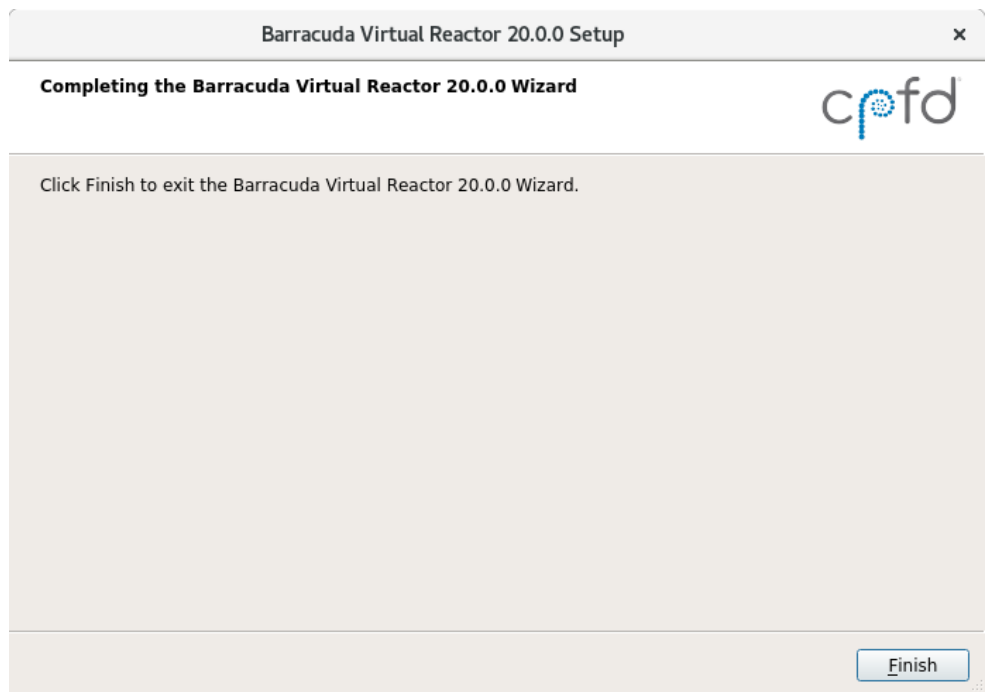

Fig. 2.17 Virtual Reactor Setup Final Page

> **Note**
>
> After installing *Virtual Reactor* on a Linux machine, it may be necessary to log out and log back in to ensure that *Virtual Reactor* is placed correctly in the system menu.

2.8. Point Client Machines to the RLM Server

Every RLM client machine must point to an RLM server in order to access the *Virtual Reactor* license(s). Clients can point to any combination of one or more RLMCloud license servers and one or more on-premise RLM license servers. The steps for pointing clients to the RLM license server(s) must be performed on the following occasions:

- The first time you set up the RLM server and its clients
- Any time the RLM server changes its IP address or hostname; in this case, all client machines need to be updated to point to the new IP address or hostname
- When a new client machine is being configured

The recommended method for pointing clients to the RLM license server(s) is by using the *Barracuda Virtual Reactor License Manager* dialog, which is accessed by choosing *Manage License Servers* in the *Virtual Reactor* GUI *Help* menu.
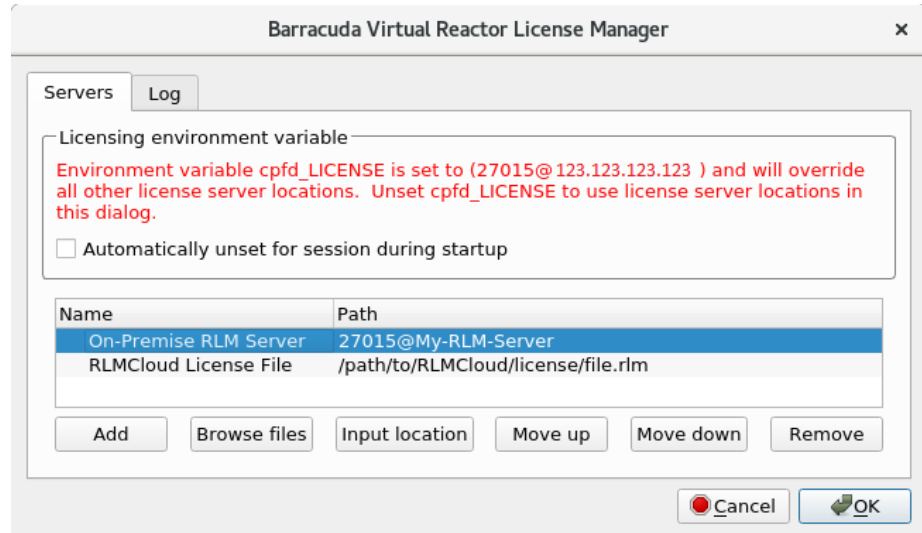

Fig. 2.18 Servers tab of *Barracuda Virtual Reactor License Manager*

The *Servers* tab, shown in Fig. 2.18, is used to define one or more RLM license servers to be used by the client. The order in which license servers are listed in this dialog is the order the client will use to request licenses. The first available license from the configured servers will be checked out.

The *Licensing environment variable* section will appear if a legacy cpfd_LICENSE environment variable has been set. This environment variable was used in pre-20.0.0 versions of *Virtual Reactor*, but is not the recommended method for pointing to the RLM license server in versions 20.0.0 and later. It is recommended that the option to *Automatically unset for session during startup* be checked so that Tecplot for Barracuda is able to check out its license successfully.

In the *Servers* table, there are two columns:

Name This column is used to give each defined RLM license server a name for easy identification. There are no restrictions on the *Name* specified.

Path If an RLMCloud license server is being defined, this column holds the path to the RLMCloud license file (see Pointing Clients to an RLMCloud License Server). If an on-premise RLM license server is being defined, this column holds the port number and host name (or IP address) of the RLM server (see Pointing Clients to an On-Premise RLM Server).

The following buttons are present below the list of defined RLM license servers:

Add This button adds a new entry to the list of defined RLM license servers. After clicking *Add*, the entry can be given a *Name*.

Browse files This button opens a file browser to select an RLMCloud license file. When a file is selected, the full path to the file will be shown in the *Path* column.

Input location This button allows text to be directly typed into the *Path* column for the currently selected entry. Double-clicking in the *Path* column for the current entry has the same effect. Direct text entry is most commonly used for specifying the port and host name (or IP address) of an on-premise RLM license server.

Move up This moves the currently selected entry up in the list of defined RLM license servers.

Move down This moves the currently selected entry down in the list of defined RLM license servers.

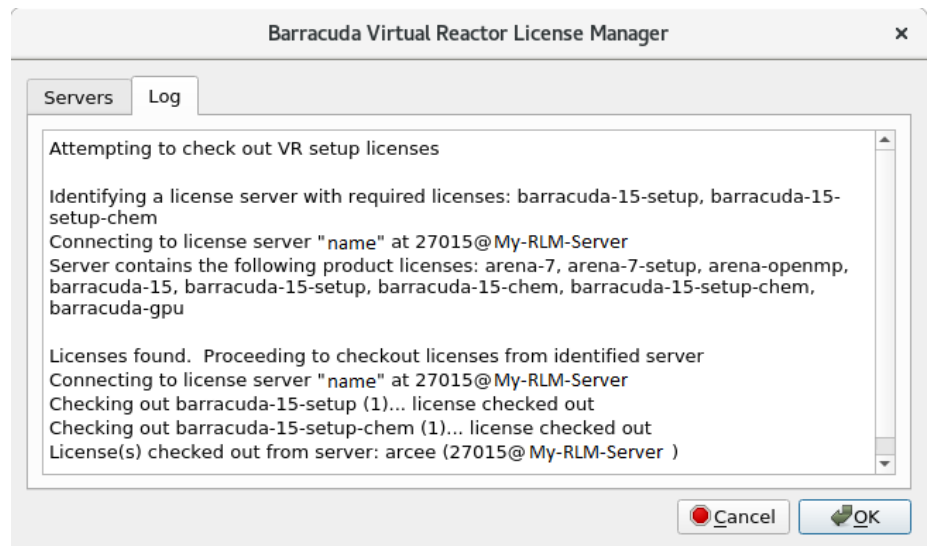**Remove** This removes the currently selected item from the list of defined RLM license servers.

Fig. 2.19 *Log* tab of *Barracuda Virtual Reactor License Manager*

The *Log* tab, shown in Fig. 2.19, is used to view messages related to RLM license checkout. This is often useful for diagnosing and troubleshooting license issues.

2.8.1. Pointing Clients to an RLMCloud License Server

If an RLMCloud license server is being defined, the *Path* column holds the full absolute path to the RLMCloud license file. The RLMCloud license file can be located anywhere on the client's file system, as long as it is readable by the user(s) running *Virtual Reactor*. To open a file browser and select an RLMCloud license file, use the *Browse files* button as shown in Fig. 2.18.

2.8.2. Pointing Clients to an On-Premise RLM Server

Client machines are configured to point to an on-premise license server by specifying the port and host name (or IP address) of the RLM license server in the *Path* column of the *Servers* table in the *Barracuda Virtual Reactor License Manager*. To enter the port and host name (or IP address) of an on-premise RLM license server, use the *Input location* button, as shown in Fig. 2.18, or double-click in the *Path* column.

Port number The default port for the RLM license server is `27015`. If you have intentionally changed it, use your custom port number instead.

Host name If the client machine can identify the RLM license server by its host name, then the *Path* can be specified as `27015@<hostname>`, where you need to replace `<hostname>` with the license server's host name.

IP address Instead of the host name, the IP address of the RLM license server may be used. The form of the *Path* entry in this case would be `27015@<IP address>`, where you need to replace `<IP address>` with the license server's IP address.

- If the OS of the RLM server is **Linux**, you can find its IP address by opening a terminal on the server and typing:

```
ifconfig
```

  The IP address is listed as `inet addr` in the printed output.

  If you want to use the `<hostname>` of the RLM server, you can find it by opening a terminal on the server and typing:

```
hostname
```

  This will print the `<hostname>` of the server machine.

- If the OS of the RLM server is **Windows**, you can find its IP address by opening a command prompt on the server and typing:

```
ipconfig
```

  The IP address is listed as `IPv4 Address` in the printed output.

If you want to use the `<hostname>` of the RLM server, you can find it by opening a terminal on the server and typing:

```
hostname
```

This will print the `<hostname>` of the server machine.

Legacy Environment Variable Method

A legacy method of pointing clients to an on-premise RLM license server is by defining an environment variable named `cpfd_LICENSE`. Using the IP address or `<hostname>` of the RLM server, the `cpfd_LICENSE` environment variable can be defined for the client machine(s) as follows:

- If the OS of the RLM client is **Linux**, you can define the environment variable based on the RLM server's `<IP address or hostname>` by opening a terminal on the client and typing:

```
echo "export cpfd_LICENSE=27015@<IP address or hostname>">>~/.bashrc
```

Logout and login in order for settings to update.

- If the OS of the RLM client is **Windows**, you can define the environment variable based on the RLM server's `<IP address or hostname>` by opening a command prompt on the client and typing:

```
setx cpfd_LICENSE 27015@<IP address or hostname>
```

2.9. Install NVIDIA Drivers

If you will be running *Virtual Reactor* in GPU parallel mode, you must install the NVIDIA drivers for your graphics cards.

- If your calculation machine's OS is **CentOS 6 Linux**, the NVIDIA drivers can be installed from the ELRepo repository by running the following commands in a terminal as the `root` user:

```
rpm --import https://www.elrepo.org/RPM-GPG-KEY-elrepo.org
```

```
yum install https://www.elrepo.org/elrepo-release-6.el6.elrepo.noarch.rpm
```

```
yum install kmod-nvidia
```

Once installation of the NVIDIA drivers is complete, reboot the system so that they are loaded by the Linux kernel.

- If your calculation machine's OS is **CentOS 7 Linux**, the NVIDIA drivers can be installed by running the following commands in a terminal as the `root` user:

```
rpm --import https://www.elrepo.org/RPM-GPG-KEY-elrepo.org
```

```
yum install https://www.elrepo.org/elrepo-release-7.el7.elrepo.noarch.rpm
```

```
yum install kmod-nvidia
```

Once installation of the NVIDIA drivers is complete, reboot the system so that they are loaded by the Linux kernel.

- If your calculation machine's OS is **Windows**, the NVIDIA drivers can be installed by following the instructions on NVIDIA's webpage for driver downloads:

http://www.nvidia.com/Download/index.aspx

2.10. Troubleshooting

We hope you never need to read this section, but in case things aren't working as expected after following the steps outlined above, please review the following troubleshooting tips.

2.10.1. License not found

If the *Virtual Reactor* GUI is unable to find a license file, it will not open. Instead, an error dialog box such as the one shown below will open:
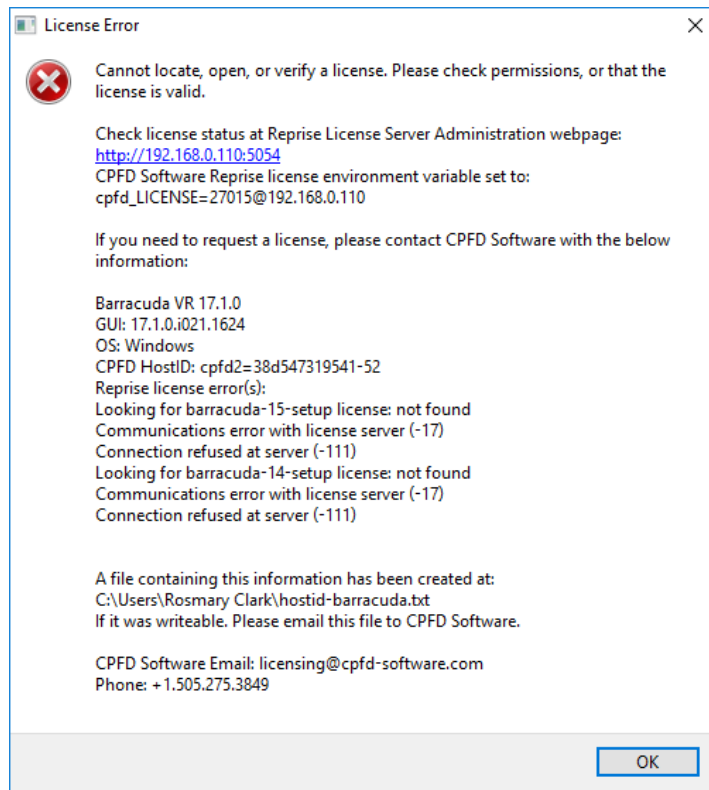

Fig. 2.20 Error dialog shown when the *Virtual Reactor* GUI cannot find a license

If this dialog appears, check the following things:

1. Is the license expired? The license file (.lic) emailed to you will have the license expiration date as will the RLM license status table (see Fig. 2.9).

2. Is the IP address or host name of the RLM server correct? Here's how to check.

3. Is the port number of the license server correct? The default port is 27015, so double-check that the `cpfd_LICENSE` environment variable is using that port. Here's how to check:

   On **Linux**, type in a terminal:

   ```
   echo $cpfd_LICENSE
   ```

   On **Windows**, type in a CMD command prompt:

   ```
   echo %cpfd_LICENSE%
   ```

   If the port number is not correct, go through the steps in Point Client Machines to the RLM Server.

4. Can the RLM web-page be opened successfully using the hyperlink in the error dialog? If not, then the RLM server is either not running on the server machine, or the client machine is not able to communicate with the RLM server. How to check:

   - Ensure that the RLM server is installed.
   - Can the client machine **ping** the RLM server? If not, then the client machine is not able to communicate with the server machine. The client and server machines must be able to communicate in order for RLM licenses to be used.

     Note

     The `ping` service is not allowed to pass through the Windows firewall by default. It is necessary to enable ping responses in order to test communication to a Windows RLM server machine.

     In order to **ping** a machine, open a terminal (on Linux) or a CMD command prompt (on Windows) and enter the command:

     ```
     ping <hostname or IP address of RLM server>
     ```

5. Is there a firewall running on the server machine, preventing the client machine from obtaining an RLM license? See Firewall Configuration for details on how to open the necessary ports for the RLM server.

6. Is the computer's date and time set correctly? The license manager will not function properly if the date and time is not correct.

2.10.2. RLM Server Issues

Issues with the RLM server can most often be diagnosed by using two log files:

1. The *debug log file* is written to continuously by the RLM server as it is running. By default, the file name is `rlm.dlog`. The file name can be verified by using the *Status* page of the RLM web interface.

Fig. 2.21 Checking RLM debug log file name

If the RLM server is running on Linux, the default location of this file is:

`/usr/local/bin/reprise/rlm.dlog`

If the RLM server is running on Windows, the default location of this file is:

`C:\Program Files\CPFD\Reprise\rlm.dlog`

2. The *Diagnostics* file is written only upon request. To create this file, use the *Diagnostics* page of the RLM web interface. The *Output File* can be specified, which refers to a location on the RLM server's hard drive. Click the *Run Diagnostics* button to run the RLM diagnostics routine and create the *Output File*.
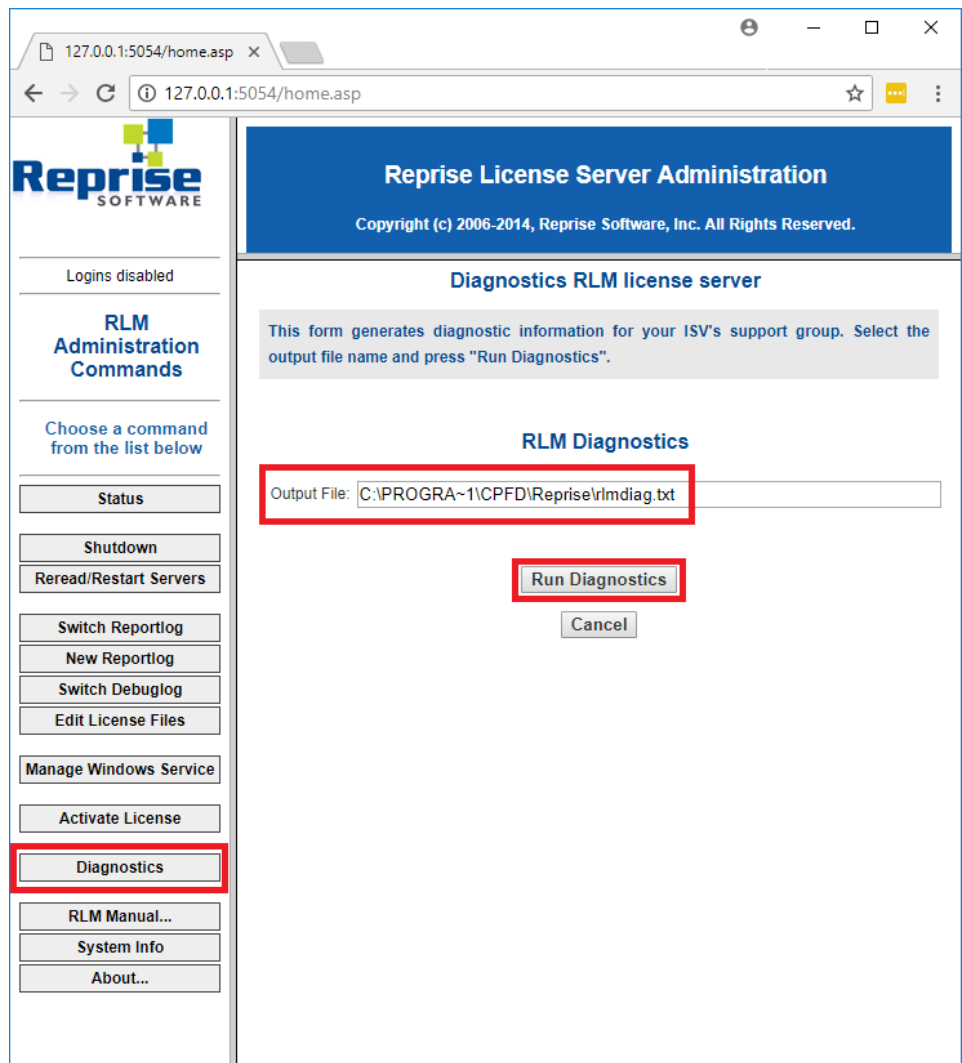
Fig. 2.22 RLM Diagnostics file

2.10.3. NVIDIA Driver Issues

1. After the installation of a new GPU card, some Linux systems may have trouble booting up because the kernel refuses to load the driver module. This can happen because of the secure boot option in the system BIOS. To fix this, either disable secure boot, or add the new GPU card's information to the system's UEFI database.