

RPM Installation Guide

Version 1.5

Revision History

Revision	Description of Change	Date
v1.0	Initial Release	2/2016
v1.1	Updated for OpenCPI Release 1.1	3/2017
v1.2	Updated for OpenCPI Release 1.2	8/2017
v1.3	Updated for OpenCPI Release 1.3	2/2018
v1.3.1	Updated for OpenCPI Release 1.3.1	4/2018
v1.4	Updated for OpenCPI Release 1.4	9/2018
v1.5	Updated for OpenCPI Release 1.5	4/2019

Table of Contents

1	References	3
2	Document Overview	4
3	Installation Process	4
3.1	Background Information	4
3.1.1	Description of Available RPMs	5
3.2	Installation File Host Locations	6
3.3	Acquiring and Installing OpenCPI from Physical Media	6
3.3.1	Install	6
3.3.2	Log Out	6
3.4	Acquiring and Installing OpenCPI from <code>github.io</code>	7
3.4.1	Configure Yum Repository	7
3.4.2	Install External Dependencies	7
3.4.3	List Available Packages	7
3.4.4	Install OpenCPI	7
3.4.5	Install ANGRYVIPER IDE (Optional)	7
3.4.6	Log Out	8
4	Post-Installation Tasks	8
4.1	Installing HDL Simulator(s) and/or Compiler(s)	8
4.2	Using the <code>opencpi</code> Group	8
4.3	Shell Environment Setup	9
5	RPM Uninstallation Process	9
6	Testing the Installation	10

List of Tables

1	References	3
2	RPM Decision Guide	5
3	RPM Descriptions	6

1 References

This document assumes a basic understanding of the Linux command line environment. It does not require a working knowledge of OpenCPI. However, it is recommended that the user read the *Getting Started* document (up to the “Installation of OpenCPI” section) or reference the *Acronyms and Definitions* document for various terms used within.

Table 1: References

Title	Link
OpenCPI Overview	Overview.pdf
Acronyms and Definitions	Acronyms_and_Definitions.pdf
Getting Started	Getting_Started.pdf
Installation Guide ¹	OpenCPI_Installation.pdf
Component Development Guide	OpenCPI_Component_Development.pdf
RCC Development Guide	OpenCPI_RCC_Development.pdf
HDL Development Guide	OpenCPI_HDL_Development.pdf
FPGA Vendor Tools Installation Guide	FPGA_Vendor_Tools_Installation_Guide.pdf
Managing Software with yum (Fedora Project)	https://docs.fedoraproject.org/en-US/Fedora_Core/5/html/Software_Management_Guide/sn-managing-packages.html
RHEL6 Deployment Guide: Useful yum commands (<i>e.g.</i> yum localinstall; Red Hat)	https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/6/html/deployment_guide/ch-yum

¹The RPM installation process is quite different from the process explained in the OpenCPI Installation Guide, but the OpenCPI Installation guide has applicable post-installation information for PCI-based boards, etc.

2 Document Overview

This document describes how to **install OpenCPI at a system level** on a development host for multiple users via RPMs. The host installation allows for local software-based execution of OpenCPI applications and components, cross-building for non-x86 platforms, simulation of HDL, and, when available, hardware testing. **Upon completion of this Guide, the steps described in the *Getting Started Guide* must be followed by *each* OpenCPI user.**

The default host installation platform for OpenCPI development is CentOS 7 Linux x86_64 (64-bit). Other Linux variants and 32-bit systems have been used successfully, but this document expects the OS to be CentOS 7. “Development host” installations can be on physical systems, virtual machines, or containers¹.

This document assumes that CentOS is already installed and proper administrative privileges have been established.

Additional installation options exist for other target processors and technologies such as the Xilinx Zynq SoC (with ARM processor cores and FPGA resources). Preference when targeting non-x86 architectures is given to *cross-building*, rather than self-hosting development. This limits the complexity of installing tools on different development hosts.

3 Installation Process

The ANGRYVIPER Team’s recommended installation method for development is through the use of RPMs. The framework can be built from source for a development host, but is not recommended for beginning users nor Application Developers. The steps herein result in a development system with tooling and runtime software ready to support development and native execution of OpenCPI components and applications.

3.1 Background Information

Understanding OpenCPI RPM naming convention

OpenCPI’s RPM naming follows that of the Red Hat Package Manager recommendations of `<name>-<version>-<release>.<dist>.<architecture>.rpm` where:

1. *name* is the name describing the packaged software
2. *version* is the version of the packaged software
version following the Major.Minor.Sub-minor naming schema
3. *release* is the number of times this version of software has been packaged
this number is independent of the version
4. *dist* is the OS distribution that the package is built for (*e.g.* `.el7`)
5. *architecture* is shorthand name describing the type of hardware the packaged software is to be installed on
6. “*devel*” is sometimes appended to the package’s name to indicate development RPMs which are required for building from source

OpenCPI “hijacks” this term as explained in Table 3.

When to Install

It is recommended that the user install these packages *before* additional tools described in Section 4.1 because the RPMs force the installation of some otherwise-hidden dependencies that allow the other tool installations to be smoother, *e.g.* 32-bit X11 libraries for ModelSim.

¹Commonly provided by Docker

3.1.1 Description of Available RPMs

It is recommended that the user installs all available packages whenever possible. If limited by available disk space, Table 2 can be used to help determine which of the packages should be installed based upon the intended use of the target machine.

Within OpenCPI, there are two types of implementations, called *Workers*, that are used in this framework: Resource-Constrained C Language (RCC) Workers and Hardware Description Language (HDL) Workers. RCC Workers are written using either C or C++ and are designed for either x86 or ARM architecture, while HDL Workers are written in VHDL and are designed for Field Programmable Gate Arrays (FPGAs) or HDL Simulators. For further details regarding RCC and HDL Workers see the *OpenCPI RCC Development Guide* and the *OpenCPI HDL Development Guide* (cf. Table 1).

Table 2: RPM Decision Guide

	Runtime RCC Host	Runtime HDL Host	RCC-Only Development (x86 RCC exclusive)	RCC/HDL Development (x86 RCC, non-SoC ¹ FPGA HDL)	RCC/HDL Development (Targeting non-x86 HW/SW platform)
angryviper-ide...rpm			✓	✓	✓
opencpi-...rpm	✓	✓	✓	✓	✓
opencpi-debuginfo...rpm			✓	✓	✓
opencpi-devel...rpm			✓	✓	✓
opencpi-doc...rpm			✓	✓	✓
opencpi-driver...rpm		✓		✓	✓
opencpi-project-bsp...rpm ²			✓	✓	✓
opencpi-*-platform...rpm					✓

¹“Non-SoC” meaning a standalone FPGA *without* an integrated processor, *e.g.* Xilinx ML605.

²BSP RPMs may not be provided with the standard/basic RPMs, but represent a placeholder for RPMs providing Board Support Package Projects.

The RPMs each have specific usage. Table 3 outlines what each of the RPMs are used for.

Table 3: RPM Descriptions

RPMs	Description
<code>angryviper-ide-*.x86_64.rpm</code>	The ANGRYVIPER IDE (Eclipse with plugins) ^a .
<code>opencpi-*.x86_64.rpm</code>	Base installation RPM includes the runtime portion of the Component Development Kit (CDK) and the source for the <code>ocpi.core</code> and <code>ocpi.assets</code> Projects containing framework essential components, workers, platforms, etc.
<code>opencpi-debuginfo-*.x86_64.rpm</code>	Debug symbols needed to debug the framework.
<code>opencpi-devel-*.x86_64.rpm</code>	Additional header files and scripts for developing new assets as HDL and/or RCC.
<code>opencpi-doc-*.x86_64.rpm</code>	Includes most of the documentation found at github.io . A symlink can be found at <code>/opt/opencpi/documentation.html</code> . If you receive the RPM directly from the AV team, it may include BSP documentation that is not available on GitHub.
<code>opencpi-driver-*.noarch.rpm</code>	OpenCPI driver. Once installed, any subsequent kernel updates will cause the driver to be built automatically on restart.
<code>opencpi-hw-platform-X-Y-*.noarch.rpm</code>	Additional files necessary to build the framework targeting specific hardware platform “X” when running RCC platform “Y” (“Y” <i>can</i> be “no-sw”). This RPM also includes hardware-specific SD Card images when applicable.
<code>opencpi-project-bsp-*.noarch.rpm</code>	A <code>*.bsp.*</code> Project (<i>e.g.</i> <code>ocpi.bsp.e3xx</code>) contains a Board Support Package for a particular physical radio, <i>e.g.</i> RCC/HDL Platform Support, Device Workers, etc. There are certain BSPs which are located in the <code>ocpi.assets</code> Project and therefore do not require their own separate BSP RPMs. As noted in Table 2, these RPMs are <i>only</i> needed for development; the <code>hw-platform</code> RPMs contain all required runtime files to deploy to an SD Card.
<code>opencpi-sw-platform-*.noarch.rpm</code>	Additional files necessary to build the framework targeting specific RCC/software platforms, independent of the final deployed hardware.

^aThis RPM is *not* available on github.io; see Section 3.4.5

3.2 Installation File Host Locations

The ANGRYVIPER Team provides RPMs to their direct customers and other users can find RPMs and other installation files on github.io.

3.3 Acquiring and Installing OpenCPI from Physical Media

If installation from ANGRYVIPER Team-released physical media is desired, follow the steps in this section.

3.3.1 Install

You can choose individual packages to install (cf. Tables 2 and 3). To install all RPMs:

```
$ sudo yum localinstall --nogpgcheck <location of RPMs>/*.rpm
```

3.3.2 Log Out

To enable the various `ocpi*` commands and set other variables, the user must log out and log back in. *Opening a new terminal session is not sufficient.* This step *can* be delayed until after Section 4 is complete.

3.4 Acquiring and Installing OpenCPI from github.io

WARNING: The IDE RPM is NOT hosted on `github.io` due to GitHub Pages' file size limitations.
This section contains additional instructions beyond `yum install` in order to install the IDE RPM.

If RPM installation from `github.io` is desired, installation should be performed with the following commands:

3.4.1 Configure Yum Repository

```
$ sudo yum install yum-utils
$ sudo yum-config-manager --add-repo=https://opencpi.github.io/repo/opencpi-v1.5.0.repo
```

3.4.2 Install External Dependencies

```
$ sudo yum install epel-release
$ sudo yum install libXft.i686 libXext.i6862
```

3.4.3 List Available Packages

To see packages available in the repository (to cross-reference with Tables 2 and 3):

```
$ yum list 'opencpi*'
```

3.4.4 Install OpenCPI

To install *all* non-IDE RPMs:

```
$ sudo yum install 'opencpi-*
```

3.4.5 Install ANGRYVIPER IDE (Optional)

The ANGRYVIPER IDE is implemented using Eclipse's Neon release and a plugin developed by the ANGRYVIPER team. RPM-based users can simply install the single RPM in a standard manner³⁴. The following instructions only apply to public users who were *not* provided an `angryviper-ide-*.rpm` file, but instead have access to the JAR file provided on `github.io`, or users who want to use their own copy of Eclipse for another reason (*e.g.* they want to use Oxygen).

1. Obtain the latest ANGRYVIPER plugin jar file:


```
$ wget https://opencpi.github.io/ide/av.proj.ide.plugin_1.5.jar
```
2. Install prerequisites:


```
$ sudo yum install oxygen-icon-theme jre
```
3. Install either the Neon or Oxygen release of Eclipse:
 - To install the Eclipse Neon release:
 - (a) Download the Eclipse Neon IDE for C/C++ Developers
URL: <https://www.eclipse.org/neon/>
 - (b) Install Eclipse by extracting the archive in the desired location
 - (c) Start Eclipse
Go into the folder where it was installed and click/run `eclipse`
 - (d) Put the `av.proj.ide.plugin_*.jar` file in the `eclipse/dropins` folder
 - (e) Install Sapphire via the Eclipse Marketplace
In Eclipse, navigate to "Help → Eclipse Marketplace". Search for "Sapphire". There should be one search result for Sapphire. Click the "Install" button. Sapphire and its dependencies will be installed.

²Only required on some point releases of CentOS7; others will autodetect

³Possibly as simple as `yum install angryviper-ide`

⁴Once installed, the command `ocpigui` will launch it

- (f) Restart Eclipse when prompted.
- To install the Eclipse Oxygen release:
 - (a) The process to construct the IDE is the same as described above using the Oxygen release for C/C++ Developers.
 - (b) At this time, the ANGRYVIPER Team has not been able to 100% verify using the plugin in Oxygen release. Eclipse Oxygen changed an API that caused problems for Sapphire, and Sapphire 9.1.1 has been released to correct the issue. The unknown part of the process is whether or not the Eclipse Marketplace will have the new version of Sapphire. If it does not, it can be installed manually as follows:
 - i. Add the Sapphire 9.1.1 repository
 - ii. Click the “add” button (to add a new repository site), fill in the popup form:


```
name: Sapphire9.1.1
location: http://download.eclipse.org/sapphire/9.1.1/repository/
```
 - iii. Click “OK” to add it
 - iv. Select the down arrow at the end of the “work with:” input. Select the new Sapphire repository.
 - (c) Select Sapphire. If Samples and Tests appear in the list; deselect them.
 - (d) Install

3.4.6 Log Out

To enable the various `ocpi*` commands and set other variables, the user must log out and log back in. **Opening a new terminal session is not sufficient.** This step *can* be delayed until after Section 4 is complete.

4 Post-Installation Tasks

4.1 Installing HDL Simulator(s) and/or Compiler(s)

For FPGA development and/or HDL simulation, OpenCPI requires vendor-provided tools (*e.g.* Xilinx Vivado, Mentor Graphics ModelSim). Refer to the *FPGA Vendor Tools Installation Guide* from Table 1 for instruction in installing and configuring these tools for use with OpenCPI.

Keep note of where the *license files* are, the *version number* of the tools, and *where the tools are installed*, as this information will be needed to configure the required environment variables.

4.2 Using the `opencpi` Group

At this point, certain users can be added to the `opencpi` group. When a user creates a Project, it is likely that the Project should be **registered**. Registering a Project allows other users and Projects to access its assets. The default Registry on an RPM-configured system is located at `/opt/opencpi/project-registry`. In order for a user to register Projects in this default location, the user will need to be a member of the `opencpi` group. To add a user to the `opencpi` group, run the following command:

```
% sudo usermod -aG opencpi <username>
```

If this command is run as user `<username>`, the user will need to log out and back in to apply this change.

The sharing of projects in this manner has been known to be fragile for various reasons (*e.g.* incorrect permission settings, default “umask” values, etc.) and is *not* recommended for new users.

Users *should* use a personal non-default Project Registry. For more information on this, please visit the *OpenCPI Component Development* document or the *Getting Started Guide* (cf. Table 1).

4.3 Shell Environment Setup

The Framework tries very hard to accept vendor default installation and configuration without additional settings. This section is only required if Section 4.1 and/or the *FPGA Vendor Tools Installation Guide* required a non-standard configuration.

Setting up the environment when installing from RPM requires root privileges. Navigate to `$(OCPI_CDK_DIR)/env.d` and notice the following example scripts:

- `altera.sh.example`
- `modelsim.sh.example`
- `site.sh.example`
- `xilinx.sh.example`

Every time a new `bash`⁵ login shell is opened, all `*.sh` files in `/opt/opencpi/cdk/env.d` are imported (“sourced”), and all `*.sh.example` files in `/opt/opencpi/cdk/env.d` are *ignored*. To enable a script for execution, the name of the script must be changed so that the `.example` suffix is removed. A simple demonstration is below:

```
% sudo cp altera.sh.example altera.sh
```

Now `altera.sh` will execute every time a new shell is opened.

If using the Altera tools, the `altera.sh` will need to be created and the variables `OCPI_ALTERA_DIR`, `OCPI_ALTERA_VERSION`, and `OCPI_ALTERA_LICENSE_FILE` must be defined in `altera.sh`. The `altera.sh` script also calls another script to set up the rest of the variables needed for the Altera tools.

If using the ModelSim tools, the `modelsim.sh` will need to be created and the variables `OCPI_MODELSIM_DIR` and `OCPI_MODELSIM_LICENSE_FILE` must be defined in `modelsim.sh`.

If using the Xilinx tools, the `xilinx.sh` will need to be created and the variable `OCPI_XILINX_LICENSE_FILE` must be defined in `xilinx.sh`. If using an installation of Xilinx Vivado that was *not* installed in the default `/opt` directory, the variable `OCPI_XILINX_VIVADO_DIR` must be defined in `xilinx.sh`. If using a version other than the most recent one installed in that location, the variable `OCPI_XILINX_VIVADO_VERSION` must be defined in `xilinx.sh`. If using an installation of Xilinx ISE that was *not* installed in the default `/opt` directory, the variable `OCPI_XILINX_DIR` must be defined in `xilinx.sh`. If not using the 14.7 version of ISE, the variable `OCPI_XILINX_VERSION` must be defined in `xilinx.sh`. The `xilinx.sh` script also calls another script to set up the rest of the variables needed for the Xilinx tools. See the *FPGA Vendor Tools Installation Guide* (cf. Table 1) for more information on Xilinx license setup.

The script `site.sh.example` has been provided as an example central location where any other variables can be defined globally. *Remember that the names of the scripts do not matter; only the *.sh extension.* More configuration variables can be found in the *Getting Started Guide*.

Once all the desired scripts have been created and edited, log out and back in and check to see that the environment is now set up.

5 RPM Uninstallation Process

In the event that the OpenCPI RPM needs to be uninstalled, or reinstalled, the best way to remove the OpenCPI RPM is to use `yum` to erase the RPMs from Table 3 as seen below:

```
% sudo yum erase <RPM name>
```

⁵Some problems have been reported when the user’s shell is set to `/bin/sh` and not `/bin/bash`.

6 Testing the Installation

To verify the OpenCPI installation, there is a command `ocpitest` that presents various test options. `ocpitest --showtests` will list all available. Some require additional files to be present or Projects to be built, but for a fresh RPM install, you can use:

```
% ocpitest driver os datatype load-drivers container
```

The first test, `driver`, will require `sudo` access. A successful install will output “All tests passed.” at the end of the test.